LOOKALIKE TARGETING ON OTHERS' JOURNEYS: BRAND VERSUS PERFORMANCE MARKETING

By

K. Sudhir, Seung Yoon Lee, and Subroto Roy

September 2021

COWLES FOUNDATION DISCUSSION PAPER NO. 2302

COWLES FOUNDATION FOR RESEARCH IN ECONOMICS
YALE UNIVERSITY
Box 208281
New Haven, Connecticut 06520-8281

Lookalike Targeting on Others' Journeys: Brand Versus Performance Marketing

K. Sudhir, Seung Yoon Lee
Yale School of Management, k.sudhir@yale.edu, seungyoon.lee@yale.edu

Subroto Roy*
University of New Haven, sroy@newhaven.edu

Lookalike Targeting is a widely used model-based ad targeting approach that uses a seed database of individuals to identify matching "lookalikes" for targeted customer acquisition. An advertiser has to make two key choices: (1) who to seed on and (2) seed-match rank range. First, we assess if and how seeding by others' journey stages impact clickthrough (upstream behavior desirable for brand marketing) and donation (downstream behavior desirable in performance marketing). Overall, we find that lookalike targeting using other's journeys can be effective—third parties can indeed identify factors unobserved to the advertiser merely from others' journey stage to improve targeting. Further, while it is sufficient to seed on upstream journey stages for *brand marketing*, seeding on more downstream stages improves *performance marketing* outcomes. Second, we assess the effectiveness of expanding the target audience with lower match ranks between seed and lookalikes. The drop in effectiveness with lower match rank range is much greater for performance marketing (donation) than for brand marketing (click-through). However, performance marketers can alleviate the reduction in ad effectiveness for low match ranks by making targeting more salient; but increasing salience has little impact for high match rank. Overall, by increasing salience, performance marketers can make acquisition cost comparable for high and low match ranks.

Key words: digital advertising, targeting, algorithmic targeting, lookalike targeting, nonprofit marketing

---

* The authors thank Vineet Kumar, Soheil Ghili, Aniko Oery, Jiwoong Shin, Kosuke Uetake and participants at the Yale SOM Lunch Workshop and the IIM Bangalore Behavioral Sciences Lab Webinar, IIM Jammu for comments and helpful discussions. We also thank Anja Lambrecht, Stephan Seiler and Song Yao for constructive feedback on an earlier draft.

1. Introduction

Lookalike advertising is a targeting approach in digital advertising where the advertiser uses a "seed database" of customers with desired behaviors to algorithmically identify "matching lookalikes" in a much larger third party database for targeting. The technique, originally introduced by Facebook in 2013, helps advertisers target and acquire customers using the depth and breadth of third party data available to Facebook, by finding similarities with customers that exhibit "desirable" behaviors (e.g., browse, visit, purchase, donate). The assumption behind the targeting technique is that greater similarity (correlation) in behaviors and descriptors (e.g., demographics/psychographics) with a focal firm's desirable seed individuals will lead to greater responsiveness to the advertising and more efficient acquisition. Several large data-driven advertising platforms such as Google, Twitter, LinkedIn, Outbrain and Taboola now offer Lookalike Ad targeting services.[1]

Lookalike targeting is conceptually different from other targeting methods facilitated by digital technologies. Most new digital technologies use information on the focal individual's position along the "customer journey" or the "purchase funnel" to improve targeting. For example, keyword search (e.g., Rutz and Bucklin 2011) uses information on a customer's needs captured in search terms; retargeting (e.g., Lambrecht and Tucker 2013) uses knowledge of products browsed, and contextual targeting (e.g., Ghose et al. 2019) uses the journey context of the focal individual to target. More broadly, past purchasing/browsing behaviors (e.g., Rossi et al. 1996, Pancras and Sudhir 2007, Rafieian and Yoganarasimhan 2020) have also been used to target the same focal individual.

In contrast to these studies in which advertisers use information on an individual's current journey or past behaviors to target the *same individual*, lookalike targeting targets an individual using the behaviors of *another individual* if there are similarities in behaviors on other categories/brands, interests and demographics. While it is reasonable to expect improvements in targeting based on one's own journey moving down the purchase funnel, whether others' journey stage can serve as effective seeds for targeting is not a priori obvious and is an empirical question. On one hand, it is

[1] Google calls its service "Similar Audiences."

possible that there are unobserved (to the advertiser) characteristics that move individuals along the journey, and if such factors are observable to the third party platforms (given its rich tracking of online behaviors) and can be correlated with lookalikes, a person's position and movement along the journey stage may indeed be effective for targeting a lookalike. But, to the extent that movement along the journey to purchase is based more on contextual and "transient" needs unrelated to permanent (and observed to third party) characteristics of the seed that are correlated with the lookalike, lookalike targeting is unlikely to be effective.

From an advertiser's new customer acquisition perspective, Lookalike targeting not only helps to expand reach efficiently outside of people who are already engaging with the firm, but is also easier to implement. With conventional third party targeting, advertisers specify their desired target profiles (e.g., Neumann et al. 2019) in terms of demographic/psychographic and behavioral characteristics. In contrast, Lookalike Targeting does not need pre-specification of the customer profiles, but requires that the advertiser only specify "seeds"—individuals identifiable by the lookalike targeting platform with advertiser desirable behaviors along the customer journey (e.g., browse, visit, social media engagement, purchase, donate, repeat purchase).

Overall, the possibility of "shifting responsibility" to third parties for specifying characteristics on which to generate similarity can be very valuable to advertisers given that the number of customers and scope of variables tracked about these customers by third party platforms is substantially greater than for any single advertiser. For platforms, being able to offer such improved targeting efficiency to advertisers generates a potentially valuable approach to monetize their data. Lookalike Targeting thus offers the potential for win-win complementarities to both advertisers and third party platforms by identifying "desired customers" based on those who engage with the advertiser and then using the depth and breadth of third party data to generate large numbers of potentially effective targets for efficient new customer acquisition. On large platforms like Facebook, Twitter and Google with hundreds of millions of customers being tracked on hundreds of thousands of behaviors across many categories and brands, the potential for both reaching new customers at greater scale and improved targeting accuracy is indeed very large.[2]

Despite Lookalike targeting's promise and widespread use across most major platforms, it has received little academic attention and little is known either about its effectiveness. In this paper, we address two sets of research questions related to lookalike targeting effectiveness. First, we consider questions related to choice of seeds by journey stage in lookalike targeting: can seeding based on the journey stage of *another individual* improve targeting effectiveness? And if effective, what journey stage should an advertiser use? For an advertiser seeking an upstream behavior (click-through to a brand site for brand building), would there be an incremental value in seeding on behaviors further down the journey (e.g., purchase, donation), or is seeding on the upstream clickthrough behavior sufficient? Similarly, for an advertiser seeking downstream performance outcomes, would there be an incremental value in seeding further down the journey on loyalty (e.g, repeat purchase, lifetime value, WOM), versus seeding on single past purchase event? On the one hand, due to selection, using seeds that have moved further down the customer journey can improve targeting, regardless of brand building or performance marketing objectives if the selection is correlated with both upstream and downstream behavior. However, if the selection is correlated primarily with downstream behaviors, then finer filtering is wasteful for upstream performance; in that it may eliminate potentially viable audience targets. Thus choice of journey stage seeds as a function of marketing objectives is an empirical question.

Our second set of research questions is on the interaction effect between journey stage of seeds and lookalike match rank. First, should an advertiser choose match rank based on the journey stage of the seeds, i.e., upstream journey stage versus downstream journey stage seeding? Second, do the differential effects of seeding and match rank vary based on the desired advertising outcomes, i.e., for brand marketers seeking upstream outcomes such as clickthrough versus performance marketers

[2] Google claims Lookalike Targeting improves performance relative to standard digital ad networks based on firm-specified targeting by over 41%.

seeking downstream outcomes such as donations? The choice of journey stage seed and match rank range involves an exploration-exploitation tradeoff. If the lookalike targeting algorithm's match score (and rank) with the seed is highly predictive of the lookalike's desired behavior (e.g., click-through, donations), then an exploitation strategy focusing on high match rank is likely preferable. However, as the predictive accuracy of the match score with the seed declines, more exploration over lower ranks may be fruitful. Due to selection as one moves along the journey, downstream seeds will have more information content as it relates to interest in the relevant category, firm or brand than upstream seeds. Further upstream behaviors are more widely prevalent and less discriminating for targeting than downstream behaviors. We therefore conjecture that performance marketers will find it more effective to use more *downstream journey stage seeds* with *higher match rank* as selection make the downstream seeds more informative. But brand marketers seeking click-through and using more upstream journey seeds may find it more effective to expand reach with reduced match rank. Whether this conjecture is true is an empirical question based on whether there is an increase in information content available in third party data with downstream seeds relative to upstream seeds.

To the extent that lower match rank leads to reduced targeting effectiveness because they are less relevant, we consider whether increasing match salience of the ad can compensate for the lower relevance. We examine the possibility of increasing effectiveness by making targeting more salient so that its perceived relevance (Shin and Yu 2019, Anand and Shachar 2009) or persuasiveness (e.g., Summers et al. 2016) can be increased. Specifically, we make targeting more salient by disclosing that they are being shown the ad as others like them have been interested in the nonprofit. However, there are concerns in the literature that a reference to similarities with others can lead to potential reactance as privacy concerns can become salient (e.g., White et al. 2008, Tucker 2014). We conjecture that the effect of salience on relevance/persuasiveness and reactance can be heterogeneous by match rank. Specifically, when match rank is high, the ability to increase relevance/persuasiveness might be limited as the fit is naturally very high. By the same argument, the gains from salience can be greater when match rank is low. On the flip side, privacy concerns due to increased salience may be higher with higher match rank because it would be seen as more intrusive/obtrusive (e.g., Van Doorn and Hoekstra 2013, Goldfarb and Tucker 2011). Therefore the net gain from targeting salience may be greater with lower match rank.

To study these questions, we conduct a set of field experiments around Lookalike targeting for new donor acquisition at a nonprofit on the Facebook Lookalike Audiences platform. Facebook is an ideal platform to study lookalike targeting for multiple reasons. First, Facebook was the pioneer and remains a leader in custom audience creation based on lookalike modeling.[3] Second, from a matching perspective, Facebook has the most extensive and richest data on individual behavior (Statista 2020) due to the size and scale of its social platform of over 2.6 billion active monthly users worldwide. Third, in contrast to the cookie-based tracking used in nonsocial ad platforms such as Google, where tracking may be broken when users clear cookies, Facebook's user-based tracking is more stable and can better identify individuals by their profile identifiers. User-based tracking also allows more effective cross-device tracking and thus an even richer basis for matching seeds and finding lookalikes. Fourth, as a practical matter, it provides a straightforward and seamless third party interface for customer acquisition that allows uploads of existing first party data as seeds and also has good options to seed on various points along the journey– and especially so in the context of our nonprofit. Finally, for our empirical context, Facebook is ideal as the focal nonprofit believes that its potential donor base is active on Facebook.

Our key findings are as follows: Overall, lookalike targeting using other's journey stages can be effective—third parties such as Facebook can indeed identify factors unobserved to the advertiser merely from others' journey stage to improve targeting. However the gains from moving seeds along the journey differs by advertiser objective. For a performance marketer seeking downstream journey outcomes such as donors, donation rates increase as one moves further down the journey stage in seeding. In contrast, for a brand marketer seeking more upstream outcomes (e.g., clickthrough), going further down the journey stage does not improve outcomes.

[3] `https://www.adweek.com/digital/lookalike-audiences/`

Second, we find an interaction effect between upstream/downstream stage seeding and Lookalike match rank. For downstream stage seeding (has donated), ad effectiveness decreased significantly when match rank is reduced from top 1% to 1%-2%. However, when seeded in the upstream stage of the journey (Visited website), there is little performance difference between the top 1% and the top 1-2% match rank ranges. Therefore we conclude that performance marketers seeking downstream outcomes should use an "exploitation" strategy by targeting their advertising on lookalikes that have the highest match with seeds (Top 1%), while brand marketers can benefit from a more "exploratory" strategy by expanding their targeting to lower ranked matches (1%-2%).

Finally, we find that performance marketers can alleviate the sharp drop in targeting effectiveness with low match ranks by making targeting more salient. When Lookalike targeting is made salient through the disclosure that the ad is being shown due to their similarity to other donors to the non-profit, ad effectiveness increased significantly, and particularly so for those with lower match rank. This means that by making targeting more salient, performance marketers can make donor acquisition cost comparable for high and low match ranks. As differences in match rank brings in different lookalike customers, the segmented message salience strategy is entirely additive for performance marketers across match rank ranges in terms of incremental customer acquisition.

The rest of the paper is organized as follows. §2 discusses the related literature. §3 describes the lookalike targeting problem and the experimental setting. §4 and §5 describes the experiments and results associated with journey stage seeding and the moderating role of match rank respectively. §6 examines the role of targeting salience in improving ad effectiveness as moderated by match rank range. §7 concludes.

2. Related Literature

This paper is connected primarily to the literature on targeting in digital advertising. Table 1 provides an overview of how Lookalike targeting differs from targeting methods that have been studied in the literature. Digital channels have facilitated many novel ways of targeting; as digital traces left by consumers reveal contextual, real time information about preferences, immediate intent and their stage along the customer journey. We classify targeting strategies into two broad groups: (1) contextual advertising based on contemporaneous information and (ii) those that are based on user history. Examples of contextual advertising include keyword search as in (Agarwal et al. 2011, Rutz and Bucklin 2011), mobile targeting that leverages location and time information (as in Luo et al. 2014, Fong et al. 2015). Targeting based on user history include those based on very short run history along a customer purchase journey (retargeting as in Lambrecht and Tucker 2013) or those based on a longer history of consumer behavior, interests, and preferences as in Trusov et al. (2016).[4] In contrast to all of this literature which focuses on improving targeting effectiveness based on behaviors of the same targeted individual, lookalike targeting uses information on a seed group of *other individuals* to target new prospects based on their similarity with the seeds.

Further, the above individual level targeting—contextual and user profile based—literature is built around conversion of customers who have already embarked on the customer purchase journey. As they are closely tied to "when" the customer is likely to be interested in the product, the conversion rates from such targeting tend to be quite good. However, these techniques are not particularly amenable to new customer acquisition, where customer needs and interest has to be initially stirred by advertisers and "others" even if they share similarities are unlikely to be interested "at the moment." The most common approach for customer acquisition tends to be traditional display or banner advertising—with targeting around desired demographics, interests and behaviors proxied by the content sites that they visit (Manchanda et al. 2006, Goldfarb and Tucker 2011, Neumann et al. 2019). With third party providers having considerable access to user demographics, interests and behaviors, advertisers can now request desired user profiles for new customer acquisition, but there are concerns about their accuracy and effectiveness (e.g., Neumann et al. 2019). Lookalike targeting is particularly useful for customer acquisition because it provides advertisers an opportunity to leverage on massive amounts of third party data to seek out new

[4] Another relevant paper related to the effect of advertising along the purchase funnel in an offline supermarket setting is Seiler and Yao (2017). They show the differential impact of advertising based on the customer's position along the conversion funnel, but their focus is not on targeting.

Table 1: Relationship to Previous Literature on Targeting

Targeting strategies	Example studies	Target based on history of	Firm Objective	Target Criteria Prespecified?
Keyword search	Agarwal et al. (2011)	Focal	Convert during journey	Yes
	Rutz and Bucklin (2011)			
Contextual based on journey	Luo et al. (2014)	Focal	Convert during journey	Yes
	Fong et al. (2015)			
	Li et al. (2017)			
	Ghose et al. (2019)			
Past purchase/browsing	Rossi et al. (1996)	Focal	Convert based on history	Yes
	Pancras and Sudhir (2007)			
	Rafieian and Yoganarasimhan (2020)			
Retargeting on recent browsing	Lambrecht and Tucker (2013)	Focal	Convert during journey	Yes
	Sahni et al. (2019)			
	Jiang et al. (2021)			
Lookalike targeting	**Our paper**	**Others**	**Acquire customer without journey/purchase history**	**No**

customers that are not currently in a firm's database, nor have recently take actions that indicate interest in the brand or category. To the best of our knowledge, we are not aware of research that has investigated how advertisers can improve targeting efficiency by using information in the journey stage of *others*.

Moreover, targeting approaches typically require marketers to enumerate attributes and conditions of the target audience for acquisition. Lookalike targeting, on the other hand, does not require advertisers to pre-specify target behavioral/descriptive profiles; advertisers can simply provide a list of people with desired behaviors, and then let the third party match customers within their database and identify those with similar characteristics. This makes targeted advertising easier to implement in practice, but also can be used even by firms seeking to acquire new customers, but don't have a large existing customer database.

We also note that while there has been some prior research focused on engineering issues in designing lookalike profiling and targeting algorithms to maximize ad performance (Liu et al. 2016, Popov and Iakovleva 2018, Cotta et al. 2019), there has been little work on studying the effectiveness of lookalike targeting and managerial choices in "implementing" lookalike targeting as has been done with various other digital advertising methods. To the best of our knowledge, ours is the first paper that focuses on assessing the effectiveness of Lookalike targeting and giving guidance to managers on various managerial choices that an advertiser implementing Lookalike targeting makes.

Finally, we note similarities and contrasts between lookalike targeting based on seeds and the distinct literature on seeding strategies within networks. In the network literature, seeding strategies are focused on selecting the optimal set of individuals in the social network who, given their position in the network, are most likely to exert peer influence in spreading word of mouth (e.g., Aral and Walker 2012, Domingos and Richardson 2001, Hinz et al. 2011, Kumar and Sudhir 2019). Examples include targeting opinion leaders who are highly connected and located in the central hub (e.g., Goldenberg et al. 2009), or located in dense regions of the network (e.g., Kitsak et al.

2010). Others underscore the importance of accounting for homophily—correlated behaviors or similarities among neighboring individuals in the social network (McPherson et al. 2001)—in measuring the effectiveness of seeding strategies (e.g., Aral et al. 2013, Nejad et al. 2015). On the other hand, seeding in Lookalike targeting is based on the idea that similar individuals profiled from a seed database are likely to behave in a similar manner desired by the firm. The key difference is that network based targeting relies on the observed choices of individuals to be close to each other on some chosen network dimension (geography, social connection etc). However, with Lookalike targeting there is no such observable choice in terms of relationships—we are merely relying on a correlation, which may be causally rooted in an underlying set of latent variables–and through unobservable selection in who moves through stages of the journey. Thus the effectiveness of Lookalike targeting has to be empirically assessed based on whether there are effective latent variables proxied in third party data that causally drive similarities between seeds and lookalikes.

3. Background

In this section, we provide relevant background on: (1) the lookalike targeting problem; (2) Facebook's lookalike targeting platform on which we conduct our empirical work; and (3) the empirical context in which we conduct our field experiments.

3.1. The Lookalike Targeting Problem

Advertisers have to make two key choices that determine Lookalike targeting success: 1) a seed set based on the assumption that user similarity correlates with the likelihood of same (desired) behavior, and 2) a match set from which the lookalike targeting algorithm will determine ad targets through adaptive learning. Given the large pool of potential Lookalike candidates in the ad platform, which typically can reach hundreds of millions of users, performance can greatly depend on declaring the right match set from which adaptive learning has to pick the right targets for advertising.

We formalize the problem as follows: The search space for Lookalike targeting is determined by two variables: the initial seed set $s \in S$ and match set $m \in M$, where M is the universal set of possible choices for the match set provided by the third party provider to the advertiser. The space of seed sets S that an advertiser can choose includes all possible seed sets available to the advertiser, both from its own CRM database and the seed options provided by the third party platform (e.g., Facebook Page engagement). Targeting performance depends on the strength of unobserved (to advertiser) correlation in desired behavior between seeds and matching lookalikes in the match set. The possible seeding strategy is infinitely large, given the large degrees of freedom coming from choosing the source and size of the seed sample. The third party ad platform typically provides a (finite) set of options for the choice set M from which m can be specified by the advertiser. Facebook, for example, allows up to 10% lookalike match rank range (in one percentile increments) in terms of its internal similarity score between the seeds and the lookalikes.

It is important to understand that the pair (s, m) jointly determines the quality of the Lookalike Audience set, $L_{(s,m)}$, the candidates for ad targeting. In other words, advertisers determine $L_{(s,m)}$, the search space for the third party Lookalike targeting algorithm to perform adaptive learning and show the ad to the targeted set of $l \in L_{(s,m)}$ individuals who the algorithm selects to explore and exploit. Hence, the interaction between the unobserved correlations in behavior and potential noise present in the seeds s and match set m influences the exploration-exploitation tradeoff in effective ad targeting.

The advertiser's challenge is to identify the best seed and match set π, from among the infinitely large set of possibilities for $L_{(s,m)}$. While the size of seed set $|s|$ is usually in the range of thousands, the size of Lookalike Audience $|L_{(s,m)}|$ is usually in the millions. For example, Facebook requires a minimum of 100 seeds and recommends between 1,000 to 50,000 seed individuals as a basis for generating a Lookalike Audience.[5] For choosing the match set, advertisers can choose between top 1-10% of the total users on the platforms from a given targeting region/country, with top 1% being the top 1% of users that best match the seeds in terms of user similarity constructed from a high dimensional feature space—based on the large number of variables available to the third

[5] See `https://www.facebook.com/business/help/164749007013531?id=401668390442328`

party. In India, which is our empirical setting, given the large number of Facebook users, even the most selective top 1% Lookalike match rank range, comprises of 3.7-4.1 million users. After the advertiser declares the match set (by specifying the match accuracy range), the third party algorithm uses adaptive learning to find the best Lookalike individuals in its ad platform who are most likely to succeed on the campaign objective, denoted by y. Examples of campaign objectives include maximizing conversion (i.e., purchase/donation), and ad engagement (clickthrough).

Given the advertiser's choice of $\pi_{(s,m)}$, the third party's adaptive learning algorithm, denoted by a, determines the Lookalike Audience targets $l \in L$.[6] We define $l_{optimal}$ as the set of best (i.e., in terms of the likelihood of maximizing the campaign objective) target individuals that reside in the third party platform. The dispersion and characteristics of $l_{optimal}$ individuals in the L sample space are unknown to the advertiser. The advertiser's choices boils down to specifying the search space by selecting the match rank range (m) that maximizes the likelihood of containing $l_{optimal} \in L$, given a (which is exogenous to the advertiser). Hence, the optimal seed-match policy that maximizes the campaign objective can be written as follows:

$$\pi := \underset{s \in S, m \in M}{\arg\max} P(l_{optimal} \in L_{(s,m)} | y, a).$$

Our primary research questions focus on generating insights around how an advertiser should choose s and m. Further, we recognize that the effectiveness of lookalike targeting may be moderated by the content of the ads. While this moderation can kickstart a rich research agenda, in this paper, we explore a particular issue as it relates to our primary research questions: how increasing targeting salience in the advertising content by highlighting similarity with others moderates the effect of s and m.

[6] For example, ad campaigns run on Facebook first goes enters the learning phase in which the algorithm learns about the best audience set for ad delivery. For details, see `https://www.facebook.com/business/help/112167992830700?id=561906377587030`

3.2. Lookalike Targeting Platform

As discussed earlier, many firms like Facebook, Google, Twitter and LinkedIn offer Lookalike Targeting services. While the general approach to lookalike targeting is similar across these platforms as described in the last sub-section, we provide some specific background on Facebook Lookalike Audiences and the specific lookalike targeting problem on which we conduct our experiments.

Advertisers that seek to use Lookalike targeting can either use proprietary seed data or seeds available through Facebook. The advertiser can either upload proprietary seed data or specify journey stage of seeds (e.g., those who have visited, liked on Facebook, visited purchase page etc.). Facebook then uses these individuals as seeds to create an audience of individuals who are similar to, or "lookalike" to the seeds. If the advertiser want to use proprietary seed data, it can upload the information about the seeds through the Facebook platform. Typically, an advertiser can include individual information about the seeds such as name, date of birth, gender, city, country, and email/phone numbers as personal identifiers. Facebook will use the proprietary donor seed data to identify the matching profiles in Facebook through a secure hashing process. After the Lookalike Audience generation process, Facebook deletes all uploaded information.[7]

Facebook generally recommends a seed size between 1,000 to 50,000 to ensure good lookalike matching. Then, advertisers decide the match rank and size a Lookalike Audience from an interface shown in Figure 1. In our context, for example, requesting for top 0-1% Lookalike Audience in India comprises of around 3.7M individuals on Facebook platform.[8] Similarly, requesting for the next match rank range of top 1-2% gives the subsequent 3.7M individuals ranked in Lookalike similarity. Creating a coarser audience diversifies the potential reach for new customer acquisition but reduces the algorithmically computed level of similarity between the Lookalike audience and the seed set.

The Facebook platform allows advertisers to conduct 'A/B/n' audience split test experiments based on different seed set and various types of match rank range with respect to the seeds to

[7] https://www.facebook.com/business/help/112061095610075?id=2469097953376494

[8] Actual audience size would vary from 3.4-3.7M, depending on the active user condition at the time.

Figure 1 Lookalike Audience Generation

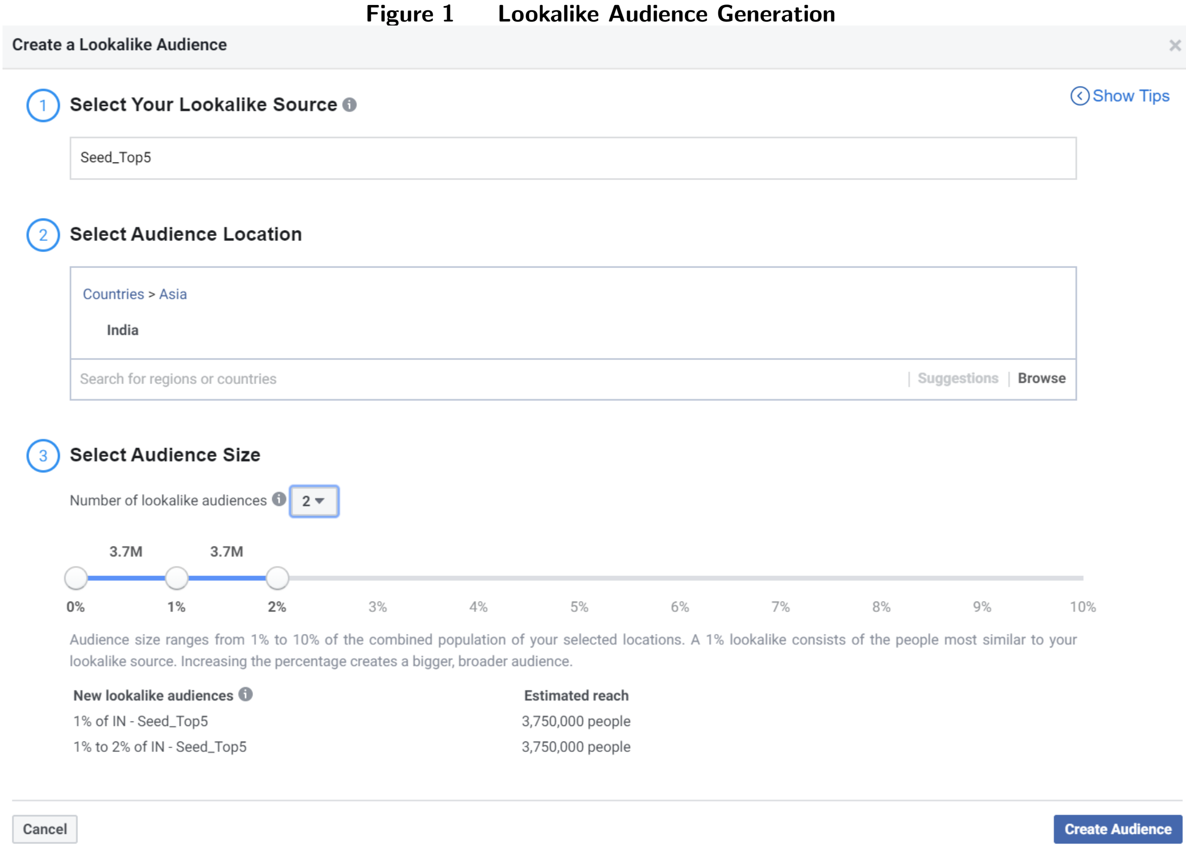

search for Lookalike individuals on Facebook. Facebook allows up to 5 treatments in any test, i.e., $n \leq 5$. Facebook's A/B test feature enables advertisers to compare performance between different audience selection strategies by ensuring that each treatment condition has equal chance of winning the bid using the lowest cost bid strategy, given the same ad budget.[9] For further details on the experiment setup and timeline of the ad campaigns, refer to Appendix B.

3.3. The Advertiser and Empirical Context

We conduct the research in partnership with an advertising partner, HelpAge India, a leading Indian non-profit organization that provides charitable support for the elderly in India. For seeding along the journey stage, it allowed Facebook to use Facebook Pixel technology to identify seeds on who had visited its website and its donation page. In addition, Facebook can use as seeds individuals who had engaged with HelpAge India's Facebook page in terms of either likes, or sharing of content on the page.

[9] https://www.facebook.com/business/help/1159714227408868

In addition to seeding using third party data, which is based on stages of the journey (up to website donation), the nonprofit advertiser also has a first party donor history database that can be used for seeding. For this lookalike targeting campaign, the organization used a monthly history of individual donor giving from April 2016 to February 2020 to segment consumers. The data included individual demographics such as name, date of birth, gender, city, country, and email that were uploaded on Facebook to be used as personal identifiers for seeding.

The CRM system at the nonprofit classified the firm's existing donors using RFM—based on prior donation behavior in the order of recency, frequency, and monetary value quintiles. To assess how donor quality differentiation impacted seeding, we used the RFM metric, as the organization felt that insights based on such donor differentiation can be valuable to them as these metrics will be available on an ongoing basis for their targeting campaigns. In unreported analysis, we indeed found that recency had the highest predictive power for predicting probability of donation, followed by frequency, and monetary value.[10] Table 2 details how the RFM descriptive statistics differ between Top 5% and Top 10% of donors. Note that selecting a larger seed base increases the potential sample size for Facebook profile identification, but reduces seed data quality. Doubling the seed size from the top 5% to top 10% led to a 79% increase in the standard deviation of RFM scores.

4. Seeding by Journey Stage

In this section, we address questions related to the advertiser's problem of determining the seed set based on past observed behaviors along the journey stage. We also consider how this seed choice should be moderated by desired ad outcomes; specifically, an upstream outcome such as clickthrough to the advertiser's website or a downstream outcome such as purchase/donation.

To fix ideas, Figure 2 displays a schematic of a donor journey with respect to the focal firm as a funnel. The funnel shape indicates that only a fraction of individuals move from one journey stage to the next and there is potential selection along the journey. We label specifically

[10] Although many alternative techniques have been proposed to estimate customer value (e.g., Fader et al. 2005, Zhang et al. 2014), RFM remains widely used in industry for its ease of use and minimal data requirements.

Table 2 Seed Data Statistics

	RFM Score	Recency (Month)	Total Donation Frequency	Total Donation Amount ($)	N
Top 5%	554.95	6.23	5.23	611.99	1829
	(0.22)	(4.41)	(5.98)	(2166.73)	
Top 10%	543.32	6.78	3.55	387.10	3659
	(17.65)	(4.58)	(4.67)	(1689.56)	
Total	356.96	24.19	1.33	89.34	36595
	(140.16)	(13.21)	(1.77)	(565.59)	

Note: Standard deviation in parenthesis

Figure 2 Seeding by Customer Journey Stage

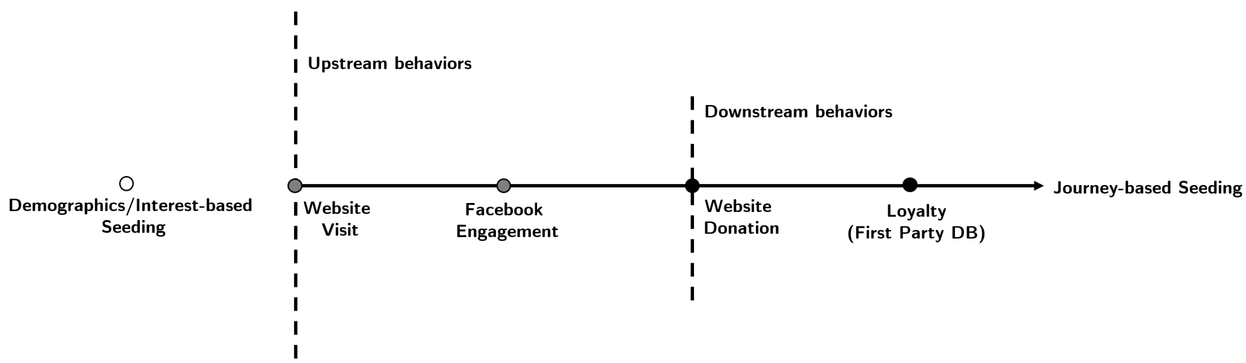

the journey stages for which seeding data is available. The advertiser can seed based on various journey stage behaviors: (i) visited the charity website, (ii) engaged with charity social media (Facebook) site and (iii) made a donation through charity website. Further, the advertiser has data on repeat donations through the charity's CRM system, which can be used to segment further and seed on levels of loyalty. Finally, as a benchmark for comparison, we consider targeting based on demographic/interest-based lookalikes without using journey data. We place this data on the left of the journey in the figure to indicate such demographic/interest-based targeting is available before seeds have embarked on journeys relevant to the target firm.

We label the initial visit and social media engagement behaviors as upstream behaviors, and the donation and loyalty behaviors (based on donation history) as downstream behaviors. This maps to the conventional nomenclature where desired behaviors for brand building (awareness, liking) are considered upstream behaviors, while sales/fundraising performance related behaviors (e.g., purchase or donation) are considered downstream behaviors.

4.1. Research Questions

Our first research question is to assess if seeding based on the journey stage of *another individual* is effective in targeted new customer acquisition. As discussed, it is an empirical question as to whether greater matching in terms of variables available in the third party database between the seeds and targets will lead to similar advertiser desired behaviors seen among the seeds for the lookalikes; i.e., if a seed has clicked through or donated, it does not logically follow that a highly matched lookalike is also likely to engage in the same behavior. The answer depends on whether there are some variables captured by the third-party on seeds and lookalikes that causally lead to the desired advertiser behavior. Given that typically third party databases and lookalike matching algorithms are blackboxes to advertisers (and researchers), advertisers and researchers can only answer whether lookalike targeting by journey stage is an effective targeting tool by comparing targeting performance relative to other relevant benchmarks.

An obvious and relevant benchmark is demographic targeting used by the non-profit; does lookalike targeting based on journey stage seeding perform better relative to demographic targeting? Further, to the extent that movement along the customer journey involves a funneling process, where only a subset of individuals move downstream along the customer journey, it would be useful to consider whether more downstream seeds are more effective for targeted new customer acquisition. For this to happen, there should be selection in who moves along the customer journey, and the third party's database should possess variables that proxy for the selection and use them in the lookalike targeting algorithm. Further, we note that if reaching further down the journey stage is based mostly on contextual or transient factors, then even if these factors are recorded by the third party, it is less likely that similarities between seeds and lookalikes alone would be a useful targeting predictor, because lookalikes also need to be in similar contextual and transient factors. As such it is necessary to compare targeting performance across journey stages and demographics to empirically assess whether an advertiser can use movement along the journey stage as effective seeds for lookalike targeting.

If indeed, there is evidence that Lookalike targeting effectiveness increases as one moves downstream along the journey, we then assess a second set of research questions. These questions are related to the managerial problem of whether the choice of seeding along the journey stage is moderated by the advertiser's advertising objectives—brand versus performance marketing. The former needs to optimize on upstream behaviors (proxied by clicks), while the latter needs to optimize on downstream behaviors (proxied by purchase/donation). For an advertiser seeking an upstream outcome such as click-through for an ad to facilitate brand building, would it be sufficient to seed on individuals who had previously shown interest in the firm by visiting its website? This is because there is little benefit to the advertiser from the selection effects that occur as consumers move down the journey towards purchases and donations, as the desired targeting outcome is an upstream behavior. However, selection imposes a cost to the advertiser in that it narrows the size of the potential target audience with whom the advertiser can build the brand. By the same token, if an advertiser seeks downstream outcomes (e.g., purchase, donation), would it be sufficient to only seed on individuals who have donated once in the past? Or would there be incremental value in seeding based on further downstream behaviors along the journey (e.g., loyalty level, lifetime value, WOM)?

4.2. Experiment Design

For our experiment, we compare the ad performance (in terms of clickthrough and donations) of five sets of audiences using Facebook's audience A/B test feature.[11] Given our primary research focus around journey based seeding, we consider four types of seeds that reflect different stages down the customer journey: upstream journey seeding on (i) Website Visits and (ii) Facebook Page Engagements and downstream journey seeding on (iii) Website Donations and (iv) Customer Value using First Party Purchase History Data. For seeding based on customer value, we use the

[11] Note that Facebook restricts testing feature to five treatments.

Table 3 Seeding by Journey Stage: Experiment Details

Journey	Seeding/Targeting Criterion	Lookalike Match	Seed Data Size	Target/Lookalike Audience Size
Upstream	Website Visit	Top 1%	75,000	3,300,000
	Facebook Page Engagement	Top 1%	32,000	3,900,000
Downstream	Website Donation	Top 1%	2,000	3,600,000
	Top 5% RFM	Top 1%	2,164	3,700,000
Baseline	Demographics	-	-	660,000

Note: 1) *Website Visit*: individuals who visited the firm's website in the past 120 days tracked by Facebook pixel.
2) *Facebook Page Engagement*: individuals who engaged in the firm's Facebook page in the past 120 days.
3) *Website Donation*: individuals who donated online in the past 120 days tracked by Facebook pixel.
4) *Demographics*: Details of demographic and interest-based targeting presented in Appendix A.1.

top 5% of customers based on RFM from the donor database of HelpAge. Table 3 presents the audience split testing experiment design with the five treatments. For all treatments, we keep the match rank range identical— top 1%, and allocated the same ad budget. However, we should note that the size of the seed database is larger for more upstream journey stages due to the funneling nature of the selection that occurs as customers move over journey stages. For the fifth treatment, we consider demographic-based targeting as a benchmark. This treatment does not involve seeds, but advertisers can specify "desired demographics." In our context, for a managerially relevant benchmark, we use the demographic targeting criterion that the non-profit uses for demographic-based targeting to acquire a desired list of prospects.[12] Refer to Appendix B for further details on the experiment setup.

4.3. Results

The outcomes for different Lookalike targeting treatments and the demographics benchmark are presented in Figure 3. First, we find that seeding based on journey stage of *others* improves lookalike targeting performance, relative to the demographic/interest-based benchmark. Indeed, the latent attributes embedded in the journey stage information transfer to correlated behaviors of the

[12] Appendix A.1 provides details on the firm's desired criteria for demographic targeting.

lookalikes. Hence, advertisers can expect similar behaviors on ad outcomes to the journey stage actions of the initial seed set of individuals and customize the seeding strategy with respect to the marketing objectives along the journey.

Furthermore, we find that the incremental benefit of additional movement down the journey only exists to the extent to the threshold behavioral correlation between the journey stage information and observed behavior. Figure 3a shows that indeed, the differences in clickthrough rates based on the various upstream (website visits and Facebook page engagement) and downstream journey seeds (website donations, top 5% RFM) are not statistically significant. Thus, our findings show that past the interest stage journey proxied by visits, the journey stage correlations that move customers down the journey does not add incremental value in improving the likelihood of upstream behaviors. Hence, advertisers with the goal of brand marketing need not use downstream journey data to seed lookalike targeting.

We also find that in terms of downstream advertising goal of acquisition (donation), lookalike targeting effectiveness increases down the journey stage seeding up to seeding on website donations. Figure 3b illustrates this monotonic improvement in donation rates for lookalikes as seeds move down the journey stage. Specifically, the difference in donation rates from using website visits (upstream stage) to website donations (downstream stage) as seeds is statistically different ($p<0.05$). However, we did not find statistically significant differences between website donations and Top 5% RFM ($p=.49$). This finding is consistent with our conjecture that seed stage does not need to go further than the marketer desired outcome (donation).[13]

Overall, seeding on journey stage of *others* matters, and journey stage has differential impact on upstream versus downstream behavioral outcomes of the targeted lookalikes. For upstream marketing, in which the marketer's objective is to increase interest (proxied by clicks), our results suggest that upstream journey stage seeds are sufficient, and investing on later journey seeds does not yield incremental value. For downstream marketing, in which the marketer's primary goal

[13] Additional unreported experiments with seeding based on different ranges of RFM confirmed this conclusion.

Figure 3 Seeding by Journey Stage

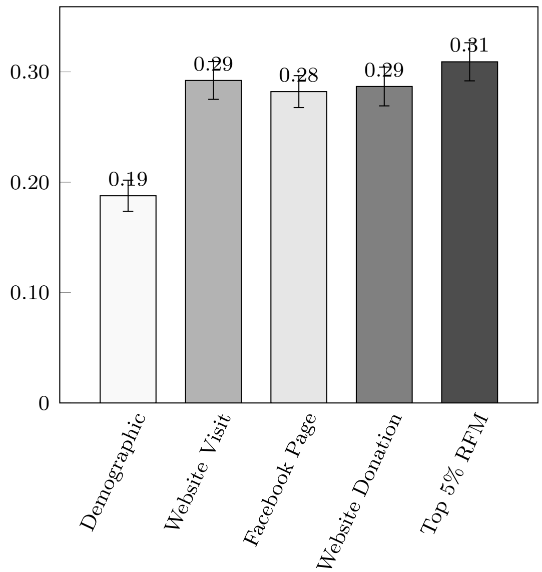

(a) Click-through Rate (%)

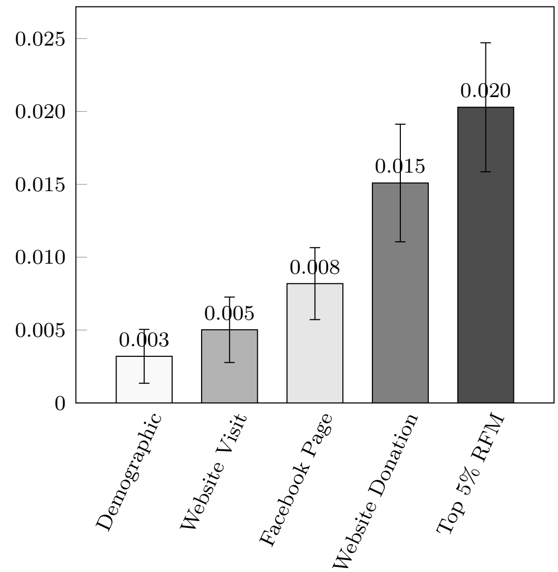

(b) Donation Rate (%)

Note: Error bars represent standard error.

is to drive conversions/donations, there is a monotonic increase in donation rates as one moves further down the journey, but beyond the purchase journey stage, fine-tuning of seed quality using customer history and loyalty information adds limited incremental value.

We note that seeding data using (i)-(iii) up to website donation is available through the third party, i.e., from Facebook through its Facebook Pixel (cookie-type) technology. Thus, any advertiser who has allowed the third party to include Facebook Pixel can seed for lookalike targeting without access to any of its own first party data. However, seeding using data on loyalty (repeat purchases and engagement across other channels etc.) can be done only with first party data. Given that not all advertisers have their own extensive customer lists and rich information on existing or prospective customers, comparing the relative benefits from HelpAge's first party data as seeds relative to seeds along the donor journey using third party data is of general interest to advertisers. This is especially relevant for young brands without extensive customer history data and whose initial goals are often to raise brand interest and upstream behaviors. Whether late stage and loyalty-based seeding adds value for targeting can provide practical guidance on data investment for seeding purposes. Further, given the increasing privacy challenges to the use of third party cookies, these questions about the value of investments in first party data also gain increased urgency and practical importance.

Figure 4 Cost Per Acquired Donor

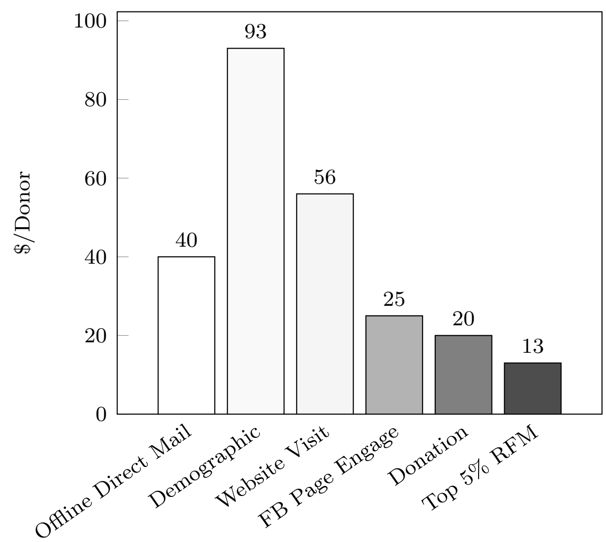

Finally, along the line of increase in donation rate based on seeding along the customer journey, the cost per donor acquisition also falls. See Figure 4. In the graph, as a benchmark we include the average cost of traditional offline targeting campaigns by direct mail. Our results show that online lookalike targeting is not always more cost-effective than offline targeting. Specifically, here we find that demographic targeting (as done using the variables used by this non-profit) and even visit based targeting is less cost-effective than doing offline direct mail targeting, for performance marketers focused primarily on donor acquisition.

5. Moderating Effect of Match Rank on Journey Stage Seeding

Having demonstrated the effectiveness of using journey stage as seeds for lookalike targeting, we next address the advertiser's choice of the second strategic variable, match rank range, and its interaction with journey stage seeding.

5.1. Research Questions

The seed set and the match rank range specified by the advertiser jointly determines the "search space" for Facebook to look for targeting prospects. Should an advertiser choose lookalikes among those with the highest match scores (and match ranks)? Or should the advertiser consider expanding the search to also consider those with lower match scores (and match ranks)? Answering these questions involves an exploration-exploitation tradeoff for the advertiser; if the lookalike targeting algorithm's match score (and rank) with the seed is very highly predictive of the lookalike's desired behaviors, then an exploitation strategy would be preferable. As the predictive accuracy of the match score with the seed declines in predicting the lookalike's desired behaviors, relatively more exploration over lower ranks may be fruitful.

To the extent the predictive accuracy of the match scores is a function of the seed quality, we ask the following research questions: How should an advertiser choose match rank range based on journey stage of the seeds? How sensitive is ad performance to the reducing match rank for upstream journey stage seeding versus downstream journey stage seeding? How do these effects vary based on desired advertising outcomes i.e., upstream behaviors such as clickthrough for brand marketing versus downstream behavior such as donation for performance marketing?

We expect less information content in upstream journey stage seeds, and therefore lower predictive power of seed-lookalike match score in predicting desired targeted behaviors. We therefore conjecture that it should be optimal to explore the space of lookalikes more by choosing a wider range of ranks when seeding by upstream journey stages. But as one moves along the journey, selection should lead to more information in downstream journey stage seeds and therefore higher accuracy for match score between seeds and lookalikes in predicting desired behaviors. So we conjecture that it would be optimal for advertisers using downstream journey stage seeding to be more conservative in exploring for lookalikes and exploit the information embedded through selection in these seeds by focusing on a narrow range of highest ranked matches.

Further, we conjecture that the gains from exploration by searching among lower ranked lookalikes will be smaller for downstream behaviors (e.g., purchases, donations) than for upstream behaviors (e.g., clicks). This is because marketing campaigns for downstream behaviors need more focused targeting, and given the inherent selection and information embedded in the movement down the journey, while upstream behaviors may be induced by more transient and contextual factors. We therefore conjecture that it would be optimal for advertisers to choose a more narrow range of highest ranked matches when targeting for downstream behaviors, but explore more among lower ranked matches when targeting for upstream behaviors.

We note that extant research on the issue of identifying similar audiences has mainly focused on the engineering aspects of designing optimal algorithms to identify similar audiences (e.g., Cotta et al. 2019, Popov and Iakovleva 2018). In contrast, we focus on the effects of interaction between seed journey stage and match rank from an exploration-exploitation perspective for effective acquisition and audience reach. Given the black box nature of lookalike targeting offered by platforms, it is important from an advertiser perspective to empirically assess how sensitive the donor acquisition performance is with respect to the chosen lookalike search space in terms of the chosen seed set and match rank.

5.2. Experiment Design

We address the exploration-exploitation tradeoff discussed above in the choice of match rank range by journey stage seed with two sets of field experiments that examine the interaction of match rank range by different journey stages. Specifically, we conduct two sets of A/B tests to assess the effects of high match rank (top 0-1%, exploitation of seed-lookalike correlation) and lower rank (top 1-2%, exploration of different lookalike search space) with upstream and downstream journey seeding strategy, as shown in Table 4 and 5. In our India context, the top 1% match rank leads to a Lookalike Audience set of around 4M individuals (highest ranked by similarity) on the Facebook platform. With the next level of match ranks (top 1-2%), the Lookalike Audience set gives the next 4M individuals ranked in similarity.

For upstream journey stage seeding, we use (i) Website Visits and (ii) Facebook Page Engagement. For downstream stage journey stage seeding, we use the high value donors (in terms of

Table 4 Match Rank Range: Interaction with Upstream Journey Stage Seeds

Journey Stage	Seed Data	Match Rank Range	Seed Data Size	Lookalike Audience Size
Upstream	Website Visit	Top 0-1%	46,000	4,100,000
	Website Visit	Top 1-2%		3,900,000
	Facebook Engagement	Top 0-1%	15,000	4,000,000
	Facebook Engagement	Top 1-2%		4,000,000

Note: 1) *Website Visit*: individuals who visited the firm's website in the past 120 days tracked by Facebook pixel.
2) *Facebook Page Engagement*: individuals who engaged in the firm's Facebook page in the past 120 days tracked by Facebook pixel.

Table 5 Match Rank Range: Interaction with Downstream Journey Stage Seeds

Journey Stage	Seed Data	Match Rank Range	Seed Data Size	Lookalike Audience Size
Downstream	Top 5% RFM	Top 0-1%	1829	3.7-3.9M
	Top 5% RFM	Top 1-2%		
	Top 10% RFM	Top 0-1%	3659	
	Top 10% RFM	Top 1-2%		

Note: During this experiment period, Facebook did not provide the exact audience size, but provided a range with the following message: "To protect the privacy of people on our platforms, we aren't showing the audience size." The suggested range for Lookalike Audience size is between 3.7-3.9 million, 1% of the population of Facebook users in India at the time of the experiment.

RFM) present in the nonprofit's CRM data. Along with this, we further assess the robustness of downstream stage seeding effectiveness by diluting the first party seed quality from highest value (top 5% RFM) to top 10%, doubling the seed size.[14] Finally, similar to our first experiment on seeding, same ad budgets are allocated for each ad set, and Facebook's audience A/B test feature ensures each lookalike audience set has an equal chance of winning the bid.

5.3. Results

The interaction effects of match rank range and journey stage seeding are presented in Figure 5a and 5b. First, we show that the role of higher match rank range become more critical going down the journey stages of the seeds (i.e., from website visits, Facebook page engagement, to latest stage of "having donated" history). That is, the performance gap between the highest 1% and subsequent 1-2% match rank widens going down the journey stages of the seeds, both in terms of clicks and

[14] To be consistent with our first experiment on journey stage seeding, we wanted to conduct the experiment with website donations and Top 5% RFM to assess downstream journey stage seeding. But the non-profit executives were keen on understanding differences arising from loyalty information the CRM data (Top 5% vs Top 10% RFM).

donations. When seeded on upstream stages of the journey, we find that ad performance is not very sensitive to the choice of match rank range and the clickthrough rates and donation rates are not statistically significantly different for web visits and FB engage seeding. This suggests that the "optimal" (i.e., in terms of ad campaign objective) lookalike candidates may be more dispersed over the different match rank range in upstream stage seeding, when the correlation between seed and desired lookalike behavioral outcomes is more noisy. Hence, the match rank range decision becomes less important under upstream stage seeding, and advertisers can explore "suboptimal" (in terms of third party match rank range criterion) lookalike search spaces. That is, the cost of exploration for Lookalike audience expansion is lower for upstream seeds. The moderating effect of match rank range on ad effectiveness becomes more critical for downstream seeds where selection helps to increase the predictive accuracy of lookalike's desired behavior.

When seeded on downstream stages of the journey, however, targeting performance drastically decreases as Lookalike match is reduced from top 0-1% to the next 1-2% in terms of match rank. Reduction of match rank resulted in 39% (31%) reduction in click-through rates and 75% (63%) drop in donations of Lookalikes of Top 5% (10%) seed individuals. Second, ad performance differences due to reducing seed quality by doubling the seed sample size from top 5% to 10% seems negligible. While there appears a slight reduction in outcomes for both clicks and donations, the differences are not statistically significant. Hence, we find that clicks and donations are not very sensitive to seed audience quality reduction but can sharply decline with lower match rank. This suggests that when seeding on downstream stages of the journey, exploitation of match value by focusing on the highest rank works better, and ad performance relies critically on the third party audience profiling quality (in terms of ability to predict desired behaviors by lookalikes), relative to the quality of the first party seed data.[15]

[15] We conducted an additional experiment on seed data quality, where we considered seed quality to be Top 6-10%. The clickthrough and donation rates continued to remain insensitive to the change in seed quality.

Figure 5 Moderating Role of Match Rank on Journey Stage Seeding

(a) Clickthrough Rate (%)

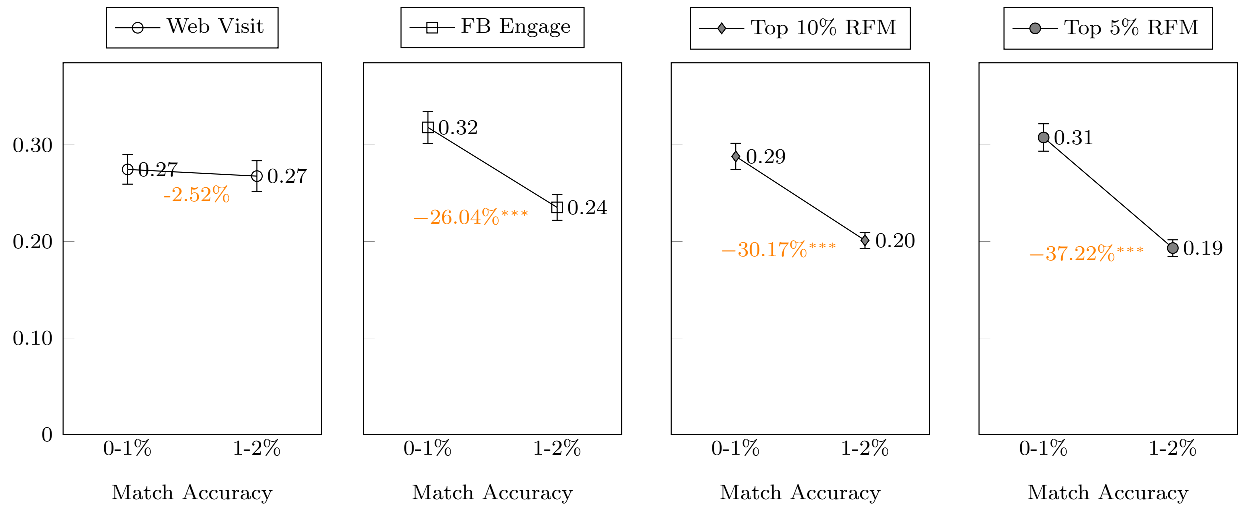

(b) Donation Rate (%)

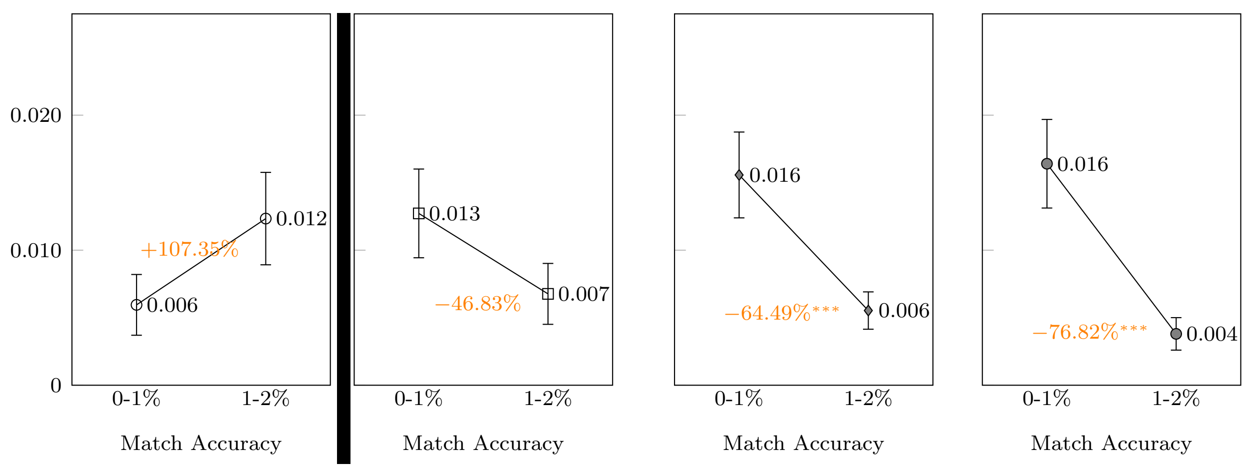

Note: * p<0.05, ** p<0.01, *** p<0.001. Error bars represent standard error.

6. Can Targeting Salience Substitute for Lower Match Ranks?

Thus far, we have evaluated the effects of two key decision variables for lookalike targeting: journey seed stage and match rank range. Our findings showed a significant reduction in ad performance both with clickthrough and donation rates with lower match rank especially with more downstream journey stages. This suggests that in these settings, marketers may be better off by using an exploitation strategy within a high match rank range, rather then exploring in a lower match rank range.

In many kinds of narrowly targeted markets, expanding the rank range of lookalikes might be the only feasible way to get sufficiently large reach for customer acquisition. In this section, we therefore consider whether increasing an ad's targeting salience can help improve ad performance even if a marketer needs to explore lower rank ranges. However theoretically, this could be a double-edged sword as targeting salience may help or hurt ad effectiveness.

6.1. Does Targeting Salience Help or Hurt Ad Effectiveness?

Unlike retargeting, in which an individual is more likely to infer that they are seeing a behaviorally targeted ad based on their recent browsing history, the lookalike targeted ad's relevance need not be evident to the targeted individual. Digital ad platforms often embed messages in advertisements that explicitly let individuals know they are viewing a targeted ad or a recommendation. Examples include Amazon's "Recommended for you" product section or Netflix's "Because you watched..." content recommendation section. Platforms expect that such messages can increase salience, make the product appear more relevant and tailored to the individuals' needs, ultimately enhancing effectiveness.

Specifically, we examine the differential role of targeting salience on higher versus lower lookalike rank ranges. Past research reports mixed findings on the effects of personalized targeting messages. Several papers have highlighted the negative effects of making targeting salient due to privacy concerns (e.g., White et al. 2008, Goldfarb and Tucker 2011). But others note positive effects on ad performance (e.g., Summers et al. 2016, Shin and Yu 2019). Kim et al. (2018) explore the effects of Facebook's "Why am I Seeing this Ad" disclosure feature on ad performance and finds that ad transparency can have negative effects when the firm's information usage violates consumers' norms about information flows. In contrast to the previous work on targeting disclosure, we directly manipulate audience match relevance and investigate the moderating role of targeting salience on low versus high Lookalike rank ranges.

In our context, advertisers have control over the congruence between the ad and the individual through the choice of match rank range. Whether a lower match rank range reduces effectiveness or mitigates privacy concerns is an empirical question. In fact, both effects may exist, and advertisers have to evaluate net effects in making their choice of match rank range. Our objective is to provide actionable insights on whether to make targeting salient, and whether the effect varies by different match ranks of the Lookalike Audiences.

6.2. Experiment Design

Table 6 summarizes the experimental treatments. We used the top 5% of donors based on RFM scores as seeds across all treatments. The experiment varied the match rank range (Top 1% and 1-2%) and whether targeting was made salient or was not mentioned. See Figure 6 for how targeting salience was operationalized. Specifically, to make targeting salient, we mentioned "Recommended for you" on the top of the ad, and "people like you have donated to HelpAge" on the headline. For controls, we were silent on targeting, in that we do not mention these statements that highlight targeting. The rest of the advertising message remained identical across treatments. The complete messages in the ad are presented in Appendix C.

Table 6 Match Rank Range and Targeting Salience: Experiment Design

	Match Rank Range	Message
1	Top 0-1%	targeting silent
2	Top 1-2%	targeting silent
3	Top 0-1%	targeting salient
4	Top 1-2%	targeting salient

Note: Seeding on Top 5% of customers based on RFM in all conditions

6.3. Results

Figure 7 illustrates the interaction effects between targeting salience and match rank range. For low match rank range (1-2%), we find that targeting salience increases performance. Specifically, making targeting salient increased clicks by 19% ($p < 0.01$), and donations by 93% ($p < 0.05$). Thus the increased attention effects are stronger than reactance effects arising from privacy concerns. In contrast, for high match rank range (top 1%), there is little difference between targeting salience

Figure 6 Experiment Design Sample Screenshots

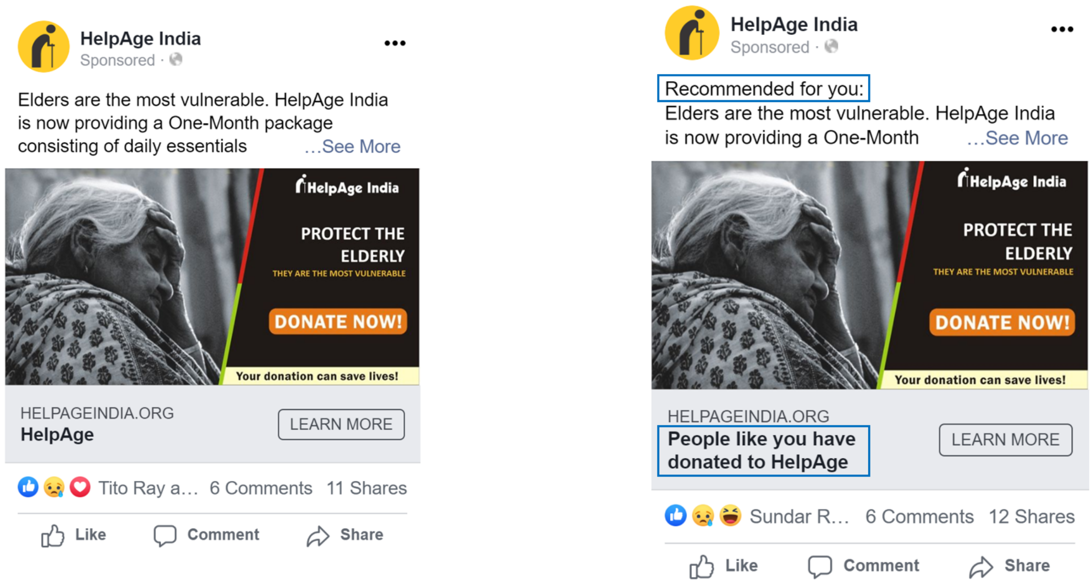

(a) Control **(b) Treatment:Targeting Salient**

and silence treatments. This is perhaps because the privacy concerns are stronger when indeed the ads are more relevant (given the higher match rank), and this neutralizes any increased positive attention or persuasive effects. Overall, our results suggest that marketers can enhance ad performance by making targeting salient when they use lower match rank ranges. As an aside, we note that this experiment replicates the results from Study 2, that there is a sharp reduction in clicks and donation rates when using downstream journey stage for seeding (in this case Top 5% in terms of RFM).

6.4. Cost-Benefit Analysis

We explore the tradeoff from targeting using lower match ranks, given that making targeting more salient improves ad effectiveness for lower match ranks. Table 7 presents the cost per thousand impressions, cost per click and cost per acquisition under the four conditions. As expected, the cost per thousand impressions are indeed much lower for lower match rank in both the targeting silent and salient conditions. Cost per click also drops with lower match rank.

What is particularly interesting is the interaction effect of match ranks and targeting salience on the cost of acquisition. These results are presented in Figure 8. Even though the cost per impression

Figure 7 Interaction Effects of Targeting Salience on Match Rank Range

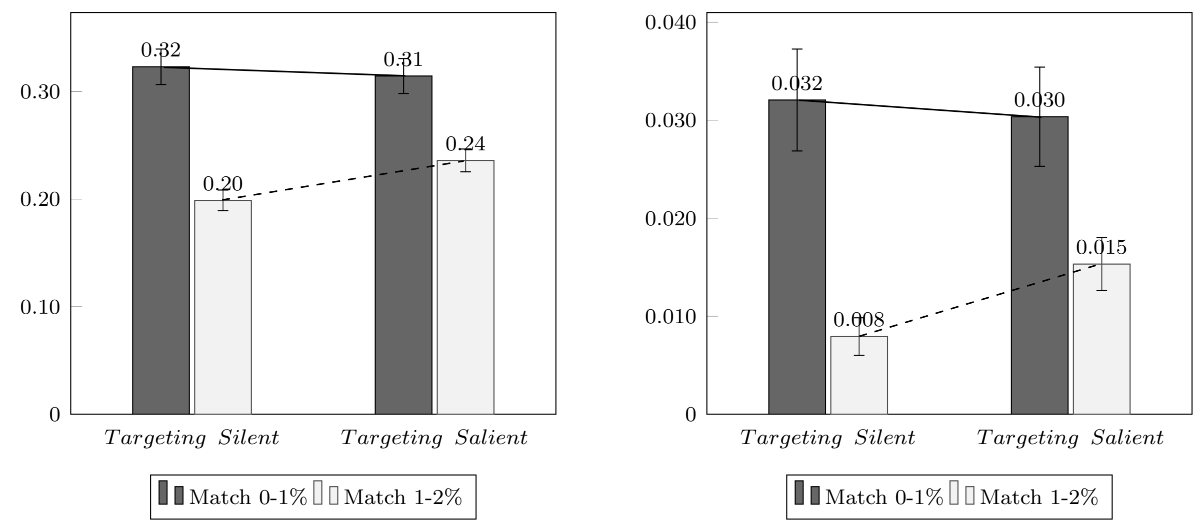

(a) Click-through Rate (%) **(b) Donation Rate (%)**

Note: Error bars represent standard error.

Table 7 Advertising Costs: Match Rank Range and Targeting Salience

Seed	Match Rank Range	CPM ($) (Cost per 1,000 Impressions)	CPC ($) (Cost per Link Click)	CAC ($) (Cost per Acquisition)
Targeting Silent	0-1%	2.35	0.73	7.33
Targeting Silent	1-2%	1.30	0.65	16.40
Targeting Salient	0-1%	2.35	0.75	7.74
Targeting Salient	1-2%	1.33	0.56	8.71

Note: Downstream journey stage seeding: customer-value based, i.e., Top 5% RFM.

is lower with low match rank range (1-2%), the cost per acquisition of customers is more than double ($16.40 versus $7.33) when the advertiser simply targets, but does not make it salient. In contrast, once targeting is salient, the cost of acquisition is much more comparable ($8.71 versus $7.74). This means that one can expand customer acquisition at comparable acquisition costs by making targeting silent when the targeting is more accurate with high match rank (top 1%), but make targeting salient when the target matching is lower (top 1-2%). And since differences in match rank brings in different lookalike customers, pursuing such a segmented messaging strategy for different match rank ranges is entirely additive (across two separate segments) for new customer acquisition. This is a managerially useful insight.

Figure 8 Cost of Acquisition: Match Rank and Targeting Salience

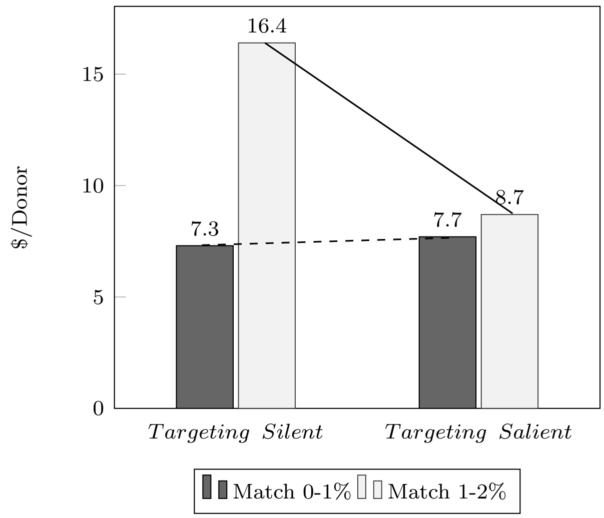

7. Conclusion

Lookalike targeting has emerged as an important ad targeting offering on most major digital advertising platforms. Unlike much of the focus in digital targeting based on one's own behaviors, lookalike targeting is based on similarity in behaviors (and descriptors) with "seeds" chosen by the advertisers. While engineering aspects of designing effective Lookalike profiling algorithms had been explored in the fields of computer science and engineering, there is little academic research in marketing from the perspective of the advertiser. To the best of our knowledge, this is the first paper to empirically test the effectiveness of lookalike targeting and provide guidelines for how brand marketers who seek upstream journey outcomes (e.g., clickthrough) and performance marketers who seek downstream journey outcomes (e.g., sales, donations) can effectively use lookalike targeting. The paper focused on two critical choices faced by advertisers in lookalike targeting: (i) seeding based on journey stage and (ii) seed-lookalike match rank. We also assess how making the targeting salient to the advertiser moderates advertising effectiveness as a function of journey stage of the seed and the match rank.

Overall, lookalike targeting using other's journey stages can be effective; advertisers can exploit the information embedded through selection in the journey stages of others, captured in third party data to improve customer acquisition. We highlight that though previous research has suggested that past purchase behaviors are the best predictors of future behaviors for the same person (e.g., Rossi et al. 1996, Pancras and Sudhir 2007), the current paper demonstrates that similarities with others, i.e., seeds exhibiting desirable behaviors (movement along journey), can be good predictors for marketers seeking those desirable behaviors. Importantly, we also show that the value of the information embedded in the selection from using seeds further down the journey is only relevant for a performance marketer seeking similar downstream outcomes. For brand marketers, seeking upstream outcomes, there is little incremental value from the refined information through selection.

Upstream seeds typically have less information content than downstream seeds because upstream behaviors are more prevalent and there is no embedded information through selection. As such, match scores between lookalikes and seeds have less predictive power for desirable behaviors than for downstream behaviors. Hence we find that it is best for performance marketers using downstream seeds and seeking downstream behaviors to "exploit" high match ranks, while it is more valuable for brand marketers using upstream seeds and seeking upstream behaviors to "explore" among lower ranked lookalikes. However, we find a cost-effective way for performance marketers to expand their reach to lookalikes with low match ranks by making the targeting more salient. Increasing salience increases ad effectiveness for lookalikes of downstream seeds, only among lower match rank lookalikes, but it has little impact on higher match rank lookalikes. This makes it feasible for performance marketers to acquire new donors at comparable cost from both high and low match ranks. Because high and low rank lookalikes belong to different segments, the effect of targeting salience is entirely additive.

7.1. Limitations and Future Research

We conclude with a discussion of limitations and suggestions for future work. First, research should consider generalizability and boundary conditions for lookalike targeting across categories; specifically, are there categories where similarities among individuals who exhibit desired behaviors serve more or less effectively as a targeting tool? Also while we found final donation behavior to be most effective for targeting, are there potentially other sweet spots along the customer journey that could be practically effective for advertisers? For example, for products with a clear time component for purchase (e.g., tickets to sporting events), would consumers who begin the search journey for a particular event be more informative than past buyers of tickets to similar sporting events?

Second, while our targeting experiments were done using Facebook Lookalike Audiences, it would be useful and important to see how the results replicate across other platforms such as Google, Twitter and LinkedIn. Since the data available for Lookalike modeling can be different across these platforms, framing research questions to test effectiveness based on the kind of data available to these platforms to inform when each of these platforms should be used can be valuable. In that spirit, though the emphasis in this paper is based on Lookalike targeting with third party data, similar research questions could be considered with lookalike targeting using second party data, where advertisers seek cooperation with specific firms who have data on consumption in related categories. It is possible that sensitivity to match rank may be lower with second party data in closely related categories.

Third, we measured the effectiveness of lookalike targeting through the metric of immediate donor acquisition. Future research should explore longer-term effects in terms of ongoing donations/purchases and lifetime value. For example, even though we found that conditional on recent donations, donation rates are insensitive to seed loyalty (in terms of RFM), it is possible that seeding based on CLV may matter for acquiring longer-term higher CLV customers. While this issue may be less salient in a nonprofit donation setting, where all donations are incrementally valuable, in settings involving high acquisition and ongoing maintenance and retention costs, seed quality (in terms of CLV) may be more critical.

Fourth, our focus of this paper is on donor acquisition at a nonprofit; it would be natural to study whether these results generalize to customer acquisition in for-profit settings. As a specific example, it is possible that individuals perceive the ad as less intrusive when coming from a nonprofit than from a for-profit. Hence the results about privacy concerns due to targeting disclosure may differ in for-profit settings.

Finally, there are some threats to the use of lookalike targeting through third parties as browsers are increasingly making data collection through cookies challenging. To the extent that Facebook itself can identify purchases on the website through its Facebook Pixel, it is interesting in this setting that Lookalike targeting based on purchases does not require even uploads to Facebook of customer data for identifying seeds who have made donations. But as third party data collection becomes more challenging, it would be useful to consider alternative approaches to execute lookalike targeting. For instance advertisers with first party seeding data, may need to partner/contract with specific second parties with the most relevant "lookalike" behavioral information to execute lookalike targeting. Our work can serve as a useful empirical framework to assess the value of such partnerships and contracts. Overall, we hope our initial investigation into Lookalike Targeting will be an impetus to explore a variety of related questions—given its extensive adoption across a range of digital platforms, and its critical importance for new customer acquisition.

References

Agarwal A, Hosanagar K, Smith MD (2011) Location, location, location: An analysis of profitability of position in online advertising markets. *Journal of marketing research* 48(6):1057–1073.

Anand BN, Shachar R (2009) Targeted advertising as a signal. *QME* 7(3):237–266.

Aral S, Muchnik L, Sundararajan A (2013) Engineering social contagions: Optimal network seeding in the presence of homophily. *Network Science* 1(2):125–153, publisher: Cambridge University Press.

Aral S, Walker D (2012) Identifying influential and susceptible members of social networks. *Science* 337(6092):337–341, publisher: American Association for the Advancement of Science.

Cotta R, Hu M, Jiang D, Liao P (2019) Off-Policy Evaluation of Probabilistic Identity Data in Lookalike Modeling. *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining*, 483–491.

Domingos P, Richardson M (2001) Mining the network value of customers. *Proceedings of the seventh ACM SIGKDD international conference on Knowledge discovery and data mining*, 57–66.

Fader PS, Hardie BG, Lee KL (2005) RFM and CLV: Using iso-value curves for customer base analysis. *Journal of marketing research* 42(4):415–430.

Fong NM, Fang Z, Luo X (2015) Geo-conquesting: Competitive locational targeting of mobile promotions. *Journal of Marketing Research* 52(5):726–735, publisher: SAGE Publications Sage CA: Los Angeles, CA.

Ghose A, Li B, Liu S (2019) Mobile targeting using customer trajectory patterns. *Management Science* 65(11):5027–5049.

Goldenberg J, Han S, Lehmann DR, Hong JW (2009) The role of hubs in the adoption process. *Journal of marketing* 73(2):1–13, publisher: SAGE Publications Sage CA: Los Angeles, CA.

Goldfarb A, Tucker C (2011) Online display advertising: Targeting and obtrusiveness. *Marketing Science* 30(3):389–404.

Hinz O, Skiera B, Barrot C, Becker JU (2011) Seeding strategies for viral marketing: An empirical comparison. *Journal of Marketing* 75(6):55–71, publisher: SAGE Publications Sage CA: Los Angeles, CA.

Jiang Z, Chan T, Che H, Wang Y (2021) Consumer search and purchase: An empirical investigation of retargeting based on consumer online behaviors. *Marketing Science* .

Kim T, Barasz K, John LK (2018) Why am I seeing this ad? The effect of ad transparency on ad effectiveness. *Journal of Consumer Research* 45(5):906–932.

Kitsak M, Gallos LK, Havlin S, Liljeros F, Muchnik L, Stanley HE, Makse HA (2010) Identification of influential spreaders in complex networks. *Nature physics* 6(11):888–893, publisher: Nature Publishing Group.

Kumar V, Sudhir K (2019) Can friends seed more buzz and adoption? Publisher: Cowles Foundation Discussion Paper.

Lambrecht A, Tucker C (2013) When does retargeting work? Information specificity in online advertising. *Journal of Marketing research* 50(5):561–576.

Li C, Luo X, Zhang C, Wang X (2017) Sunny, rainy, and cloudy with a chance of mobile promotion effectiveness. *Marketing Science* 36(5):762–779.

Liu H, Pardoe D, Liu K, Thakur M, Cao F, Li C (2016) Audience expansion for online social network advertising. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 165–174.

Luo X, Andrews M, Fang Z, Phang CW (2014) Mobile targeting. *Management Science* 60(7):1738–1756, publisher: INFORMS.

Manchanda P, Dubé JP, Goh KY, Chintagunta PK (2006) The effect of banner advertising on internet purchasing. *Journal of Marketing Research* 43(1):98–108, publisher: SAGE Publications Sage CA: Los Angeles, CA.

McPherson M, Smith-Lovin L, Cook JM (2001) Birds of a feather: Homophily in social networks. *Annual review of sociology* 27(1):415–444, publisher: Annual Reviews 4139 El Camino Way, PO Box 10139, Palo Alto, CA 94303-0139, USA.

Nejad MG, Amini M, Babakus E (2015) Success factors in product seeding: The role of homophily. *Journal of Retailing* 91(1):68–88, publisher: Elsevier.

Neumann N, Tucker CE, Whitfield T (2019) Frontiers: How effective is third-party consumer profiling? Evidence from field studies. *Marketing Science* 38(6):918–926.

Pancras J, Sudhir K (2007) Optimal marketing strategies for a customer data intermediary. *Journal of Marketing research* 44(4):560–578.

Popov A, Iakovleva D (2018) Adaptive look-alike targeting in social networks advertising. *Procedia computer science* 136:255–264.

Rafieian O, Yoganarasimhan H (2020) Targeting and privacy in mobile advertising. *Available at SSRN 3163806* .

Rossi PE, McCulloch RE, Allenby GM (1996) The value of purchase history data in target marketing. *Marketing Science* 15(4):321–340, publisher: INFORMS.

Rutz OJ, Bucklin RE (2011) From generic to branded: A model of spillover in paid search advertising. *Journal of Marketing Research* 48(1):87–102, publisher: SAGE Publications Sage CA: Los Angeles, CA.

Sahni NS, Narayanan S, Kalyanam K (2019) An experimental investigation of the effects of retargeted advertising: The role of frequency and timing. *Journal of Marketing Research* 56(3):401–418.

Seiler S, Yao S (2017) The impact of advertising along the conversion funnel. *Quantitative Marketing and Economics* 15(3):241–278, publisher: Springer.

Shin J, Yu JJ (2019) Targeted advertising and consumer inference. *Quantitative Marketing and Economics Conference*.

Statista (2020) Number of monthly active Facebook users worldwide as of 1st quarter 2020. URL `https://www.statista.com/statistics/264810/number-of-monthly-active-facebook-users-worldwide/`, (Accessed July 17, 2020).

Summers CA, Smith RW, Reczek RW (2016) An audience of one: Behaviorally targeted ads as implied social labels. *Journal of Consumer Research* 43(1):156–178.

Trusov M, Ma L, Jamal Z (2016) Crumbs of the cookie: User profiling in customer-base analysis and behavioral targeting. *Marketing Science* 35(3):405–426.

Tucker CE (2014) Social networks, personalized advertising, and privacy controls. *Journal of Marketing Research* 51(5):546–562.

Van Doorn J, Hoekstra JC (2013) Customization of online advertising: The role of intrusiveness. *Marketing Letters* 24(4):339–351, publisher: Springer.

White TB, Zahay DL, Thorbjørnsen H, Shavitt S (2008) Getting too personal: Reactance to highly personalized email solicitations. *Marketing Letters* 19(1):39–50.

Zhang Y, Bradlow ET, Small DS (2014) Predicting customer value using clumpiness: From RFM to RFMC. *Marketing Science* 34(2):195–208.

Appendix

A. Facebook Custom and Lookalike Audience Generation

A.1. Firm-provided Demographic Criterion

Prior to this study, the nonprofit has been using the following demographic criterion to request and purchase mailing lists from the third-party data brokers for new donor acquisition. We hence use the desired demographic specification as benchmark to assess the incremental value of Lookalike targeting. We outline the demographic criterion as below:

The Nonprofit's Cold Mailing List Criterion:

1. High networth working individuals: Individuals with at least ₹500,000 in investments in the stock market or mutual funds. Also, the annual income needs to be at least ₹500,000 and preferably ₹1,000,000 or more.
2. Region: Metropolitan and Tier 2 cities preferred, excluding eastern India regions.
3. Age : 35 +
4. Gender: Female (Working women)

We applied the above demographic targeting criterion on Facebook audience creation tool to the extent possible. We used higher level education as proxies for high networth individuals and excluded Eastern India regions for targeting, following the nonprofit's suggestion. We detail the specific demographic targeting used on Facebook Audience Creation tool as follows. We also present the demographic-based seed generation process in Figure A.1.

Facebook Demographic Target Audience Criterion:

- Potential Audience:
  - Potential Reach: 660,000 individuals
- Audience Details:
  - Location - Living In: India
    - Exclude Eastern India Regions: Assam; Manipur; Meghalaya; Nagaland; Odisha; Tripura; West Bengal; Sikkim; Arunachal Pradesh; Mizoram; Bihar; Jharkhand
  - Age: 35 - 65+

—Gender: Female

—Education Level: Master's degree, Professional degree or Doctorate degree

—Industry: Administrative Services, IT and Technical Services, Legal Services, Sales, Education and Libraries, Business and Finance, Management, Arts, Entertainment, Sports and Media, Architecture and Engineering, Food and Restaurants, Construction and Extraction, Production, Healthcare and Medical Services, Installation and Repair Services, Life, Physical and Social Sciences, Computation and Mathematics, Community and Social Services, Protective Services, Farming, Fishing and Forestry, Cleaning and Maintenance Services, Military (Global), Transportation and Moving or Government Employees (Global)

A.2. Facebook Seed Dataset Creation

We consider four alternative seed data creation options in Facebook platform that enable marketers to seed on individuals without the need of providing any proprietary data: i) demographic targeting, ii) website visit, iii) Facebook engagement, and iv) website donation tracked using Facebook Pixel. Facebook provides a straightforward interface to generate the following custom seeds, as demonstrated in the following screenshots.

Figure A.1 Audience Creation using Demographic Criterion on Facebook

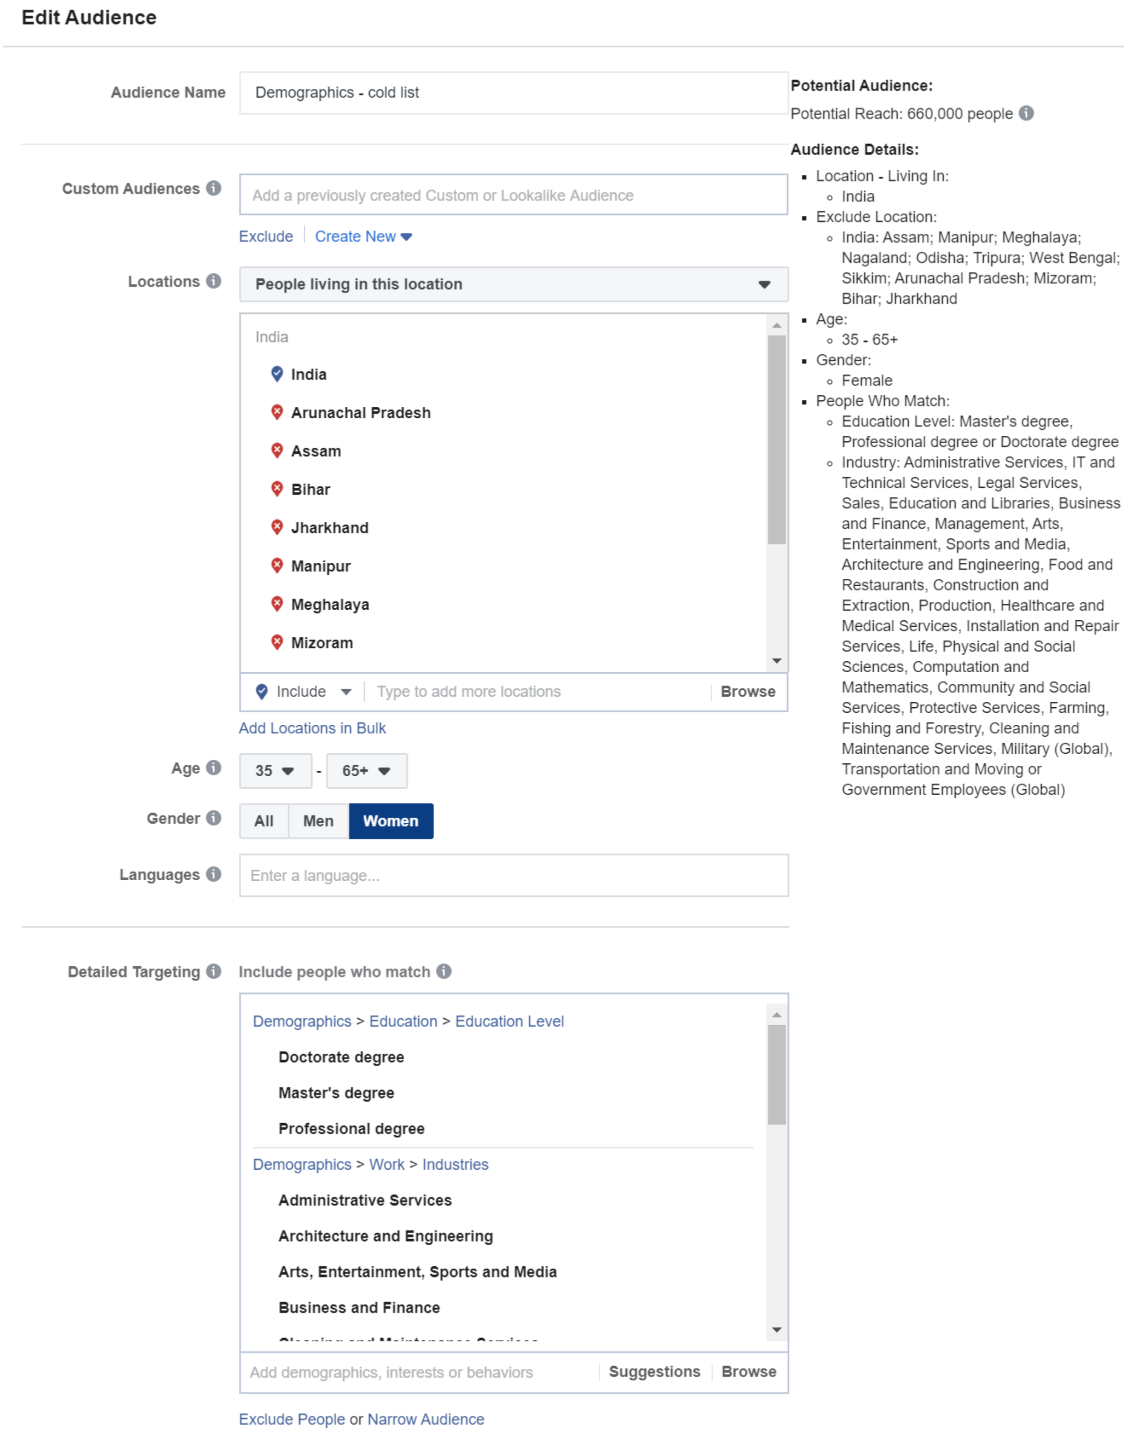

Figure A.2 Seed Creation: Facebook Page Engagment

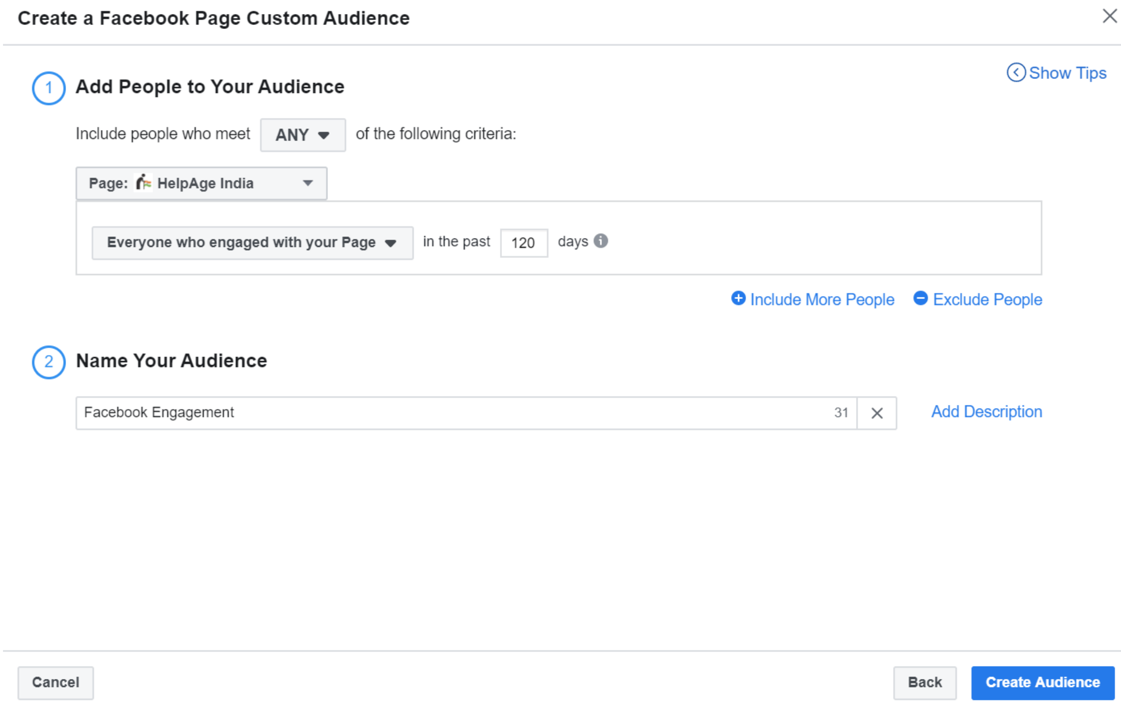

Figure A.3 Seed Creation: Website Visits and Purchase using Facebook Pixel

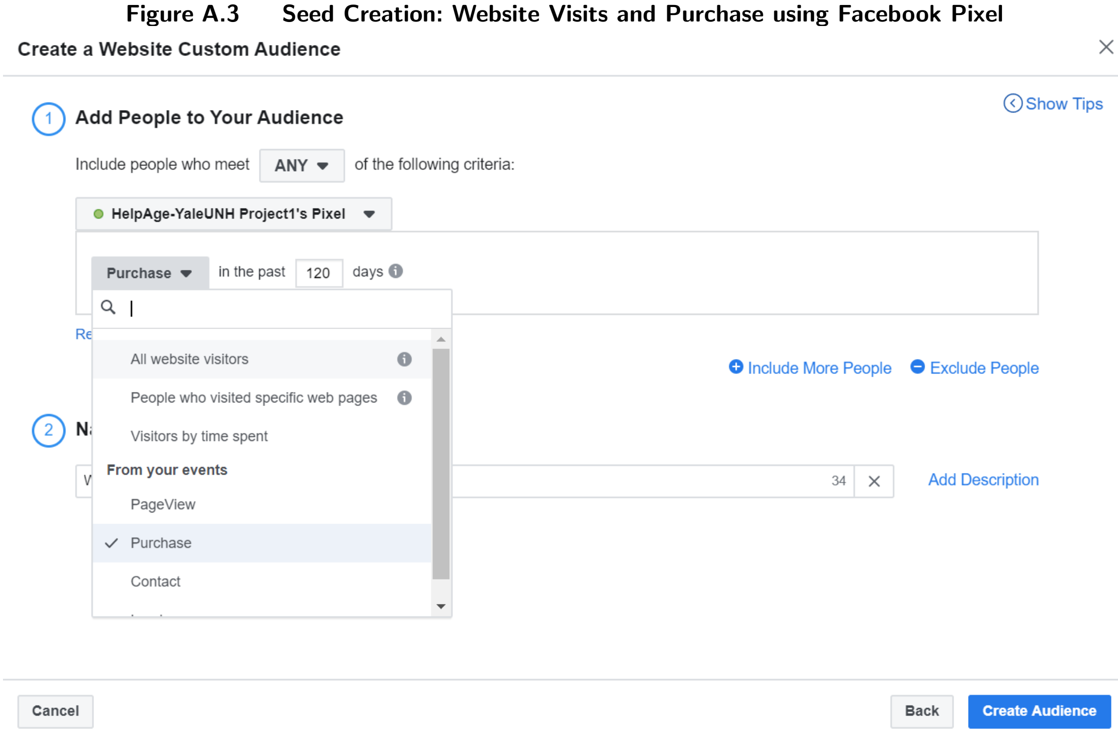

B. Experiment Details

In this section, we provide further information on the experiment setup as follows:

- Ad Placement: Facebook newsfeed (mobile & desktop)
- Objective: Conversion campaign
- Audience: Seed-based custom Lookalike Audience
- Bidding: Dynamic (minimum cost) -Facebook's A/B testing feature ensures that each treatment condition has an equal chance of winning the bid using the lowest cost bid strategy at any given time under a dynamic competition setting. Same ad budgets are allocated for each ad set.
- Experiment Duration: Each experiment was conducted for a duration of 7 days. The entire set of experiments were conducted over the period April 2020–March 2021.
- Attribution window: 1-day view, 28-day click

Figure B.1 Sample Ad Design

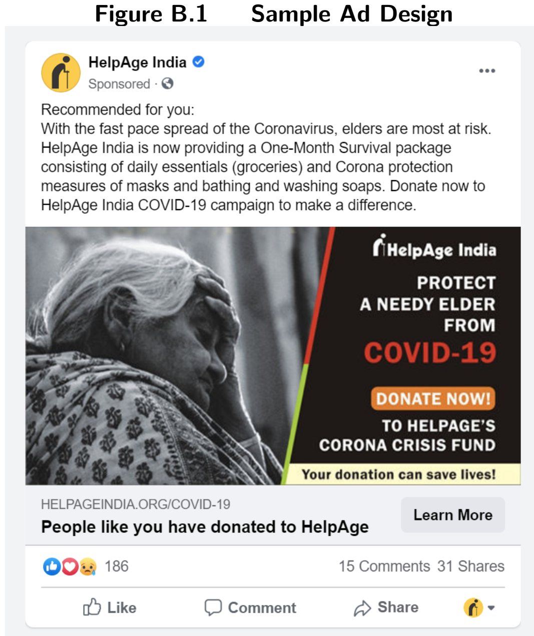

C. Targeting Salience

We provide further details on the targeting salient message conditions.

Table C.1 Ad Message in Different Targeting Salience Conditions

Control: Silent on targeting	*Primary text*: Elders are the most vulnerable. HelpAge India is now providing a One-Month package consisting of daily essentials (groceries), masks and bathing and washing soaps. Donate now to HelpAge India to protect the elderly. *Headline*: HelpAge
Treatment: Targeting salient	*Primary text*: Recommended for you: Elders are the most vulnerable. HelpAge India is now providing a One-Month package consisting of daily essentials (groceries), masks and bathing and washing soaps. Donate now to HelpAge India to protect the elderly. *Headline*: People like you have donated to HelpAge