3-2019

Do mobile banner ads increase sales? Yes, in the offline channel

Ernst C. OSINGA
Singapore Management University, ecosinga@smu.edu.sg

Menno ZEVENBERGEN
Greenhouse Group

Mark W. G. VAN ZUIJLEN
Essent/Innogy

Do mobile banner ads increase sales? Yes, in the offline channel

Ernst C. Osinga, Lee Kong Chian School of Business, Singapore Management University, 50 Stamford Road, 178899 Singapore, Singapore

Menno Zevenbergen, Greenhouse Group, Emmasingel 25, 5611 AZ Eindhoven, the Netherlands

Mark W.G.van Zuijlen, Essent/Innogy, Willemsplein 4, 5211 AK 's-Hertogenbosch, the Netherlands

Published in International Journal of Research in Marketing, March 2019, Advance online

https://doi.org/10.1016/j.ijresmar.2019.02.001

Abstract

Firms allocate increasingly large budgets to mobile banner advertising. Yet, existing research paid only scant attention to the sales effects of mobile banner ads. In this paper, we fill this gap by determining the offline and online sales impact of a large-scale mobile banner advertising campaign. As part of a geographical field experiment, over 3.5 million mobile banner ads were served to a predetermined geographical area. We determine the offline and online sales effects of the mobile banner ad campaign by analyzing twenty months of sales data for regions covering the entire country of the Netherlands. Relying on a difference-in-difference approach and two matching methods, we demonstrate an offline sales increase of around 2%. The online sales effect is not significant. We conclude that firms can use mobile banner advertising to boost offline sales. We find no evidence for cross-channel sales cannibalization.

Keywords

Mobile banner advertising, Mobile marketing, Digital advertising, Geographical field experiment, Propensity score matching, Coarsened exact matching

1. Introduction

The worldwide number of smartphone users is expected to reach 2.53 billion in 2018, or over 53% of worldwide mobile phone users (eMarketer, 2016b). In the U.S., nearly 200 million consumers, or about 80% of the mobile market, own a smartphone (comScore, 2016). Already in 2014, U.S. consumers spent more time on smartphones than on desktop computers (comScore, 2015). The significant amount of time that consumers spend on smartphones does not go unnoticed by managers. Mobile advertising is an increasingly important component of a firm's digital strategy. Investments in mobile display advertisements are expected to surpass 57 billion dollars in 2018 in the U.S., an increase of over 90% from 2015 (eMarketer, 2016a). The mobile advertising share of digital advertising will approach 70% in 2018 (eMarketer, 2016a) and Facebook already reports that a staggering 84% of its advertising revenue is generated by mobile advertising sales (New York Times, 2016).

Mobile *advertising* includes mobile banner advertising, the focus of this study, and mobile video advertising and should be distinguished from mobile *promotions* which are mostly in the form of SMS or in-app coupons (Andrews, Goehring, Hui, Pancras, & Thornswood, 2016).[1] Existing research on mobile marketing mostly looked at the effects of mobile promotions (e.g., Barwise & Strong, 2002; Fong, Fang, & Luo, 2015; Luo, Andrews, Fang, & Phang, 2014; Zubcsek, Katona, & Sarvary, 2017) and paid only scant attention to the impact of mobile banner advertising. Bart, Stephen, and Sarvary (2014) provide a notable exception of research on mobile banner advertising. They analyze the impact of 54 mobile banner advertising campaigns on attitude and intention, where each campaign was set up as a randomized field experiment. They show that mobile banner advertising positively impacts attitude and intention for high-involvement utilitarian products. While Bart, Stephen, and Sarvary analyze the impact of a single mobile banner ad exposure on intermediate measures directly after the exposure, the effects on offline and online purchase behavior are largely unknown. For managers, existing literature thus provides little guidance. Existing literature leaves several questions unanswered.

First, it is unclear whether mobile banner advertising impacts consumers' purchase behavior. Mobile devices have been found to be more often used strictly for browsing (Einav, Levin, Popov, & Sundaresan, 2014), however, recently, they are increasingly being used to make transactions (The Telegraph, 2016). Second, it is unclear whether the sales impact, if any, manifests in the offline or online channel, or even in both channels (also see Kannan, Reinartz, & Verhoef, 2016). In addition, it is of great importance to understand whether a possible effect in one channel goes at the cost of sales in the other channel, i.e., do mobile banner ads boost overall sales or do they just shift sales from one channel to the other? And if so, from and to which channel do sales shift?

With this research, we aim to address the questions raised above, hereby addressing the call from Grewal, Bart, Spann, and Zubcsek (2016) for more research into the effects of mobile banner advertising. Specifically, in this study, we analyze the offline and online sales effects of mobile banner advertising. We rely on a geographical field experiment (Blake, Nosko, & Tadelis, 2015) combined with an analysis of offline and online sales data. We observe offline and online sales from October 2013 to May 2015 for separate regions that cover the entire country of the Netherlands. From January 6, 2015, to January 10, 2015, we exclusively target mobile banner advertising to consumers that are located in one of three selected Dutch provinces. Consumers that are in any of the remaining nine provinces are not exposed to mobile banner advertising. We use a difference-in-difference approach, cast in a fixed effects regression model specification, to determine the offline and online sales impact of the mobile banner ads. Our results indicate a 2% offline sales uplift due to the mobile banner ad campaign. We find no evidence for an online sales impact, i.e. we find no evidence for a shift in sales from the online to the offline channel. In a series of robustness checks, we confirm that the key assumption of the difference-in-difference approach is met and show that our results hold when matching the control and treatment group using the propensity scores and coarsened exact matching method. Also, we find evidence for a positive lagged offline sales effect of the mobile banner ad campaign.

Our findings provide an important contribution to the current literature on mobile marketing. In line with the increased importance of mobile devices for consumer decision making (De Haan, Kannan, Verhoef, & Wiesel, 2018; The Telegraph, 2016), we find that mobile banner ads do impact consumers' purchase behavior, yet, we find that the sales effect manifests in the *offline* channel. The findings of this study are highly relevant for managers. Particularly managers of retail firms with a semi-digital business model based on both offline and online sales, such as fashion retailers and department stores, can consider mobile banner ads to boost offline sales. In addition, managers of fast-food chains, eateries, cafes, and manufacturers of products such as beverages and ice cream, can generate additional offline sales through mobile banner ads. Our results are also of interest to digital-native firms that wish to grow their offline business. Perhaps counterintuitively, we find no evidence for a shift in sales from the offline to the online channel due to mobile banner ads. Firms that wish to move their business from the offline to the online channel to adopt an all-digital business model should not blindly rely on mobile banner ads.

2. Banner advertising: mobile versus desktop

Banner ads are served on mobile and desktop devices. In 2016 in the U.S., spending on *mobile* banner advertising accounted for about 78% of total spending on banner advertising, the remaining 22% of banner ad spending being allocated to *desktop* banner ads (eMarketer, 2016c).[2] Despite mobile banner advertising enjoying larger investments than desktop banner advertising, the majority of existing studies on banner advertising focus exclusively on the desktop channel (e.g., Hoban & Bucklin, 2015; Lobschat, Osinga, & Reinartz, 2017) or do not specify whether mobile ad exposures and purchases are included. Importantly, the effects of mobile and desktop banner ads may be different due to innate differences between the two device types.

First, mobile devices are used at different moments and locations than desktops (Ghose, Goldfarb, & Han, 2013). Mobile devices are used throughout the entire day and in different geographical locations such as while traveling on subway trains (Andrews, Luo, Fang, & Ghose, 2016) or while shopping offline. Consumers enjoy browsing on their mobile devices, yet, they may be hesitant to conduct mobile purchases due to the perceived privacy and security risks (Chin, Felt, Sekar, & Wagner, 2012; De Haan et al., 2018). Although the number of mobile purchases is on the rise (The Telegraph, 2016), consumers may still prefer to switch to either less mobile devices (De Haan et al., 2018) or an offline store for their purchases. Hence, while desktop banner ads reach consumers on a device associated with purchasing, and thus may generate online purchases, the effects of mobile banner ads are more likely to manifest in the offline channel, particularly when the consumer is exposed to mobile banner ads while being closer to an offline outlet than to a desktop device.

Second, apart from the different usage situations of mobile and desktop devices, tangible differences between the two types of devices may affect the impact of ads. Mobile devices are typically equipped with smaller screens than desktop devices (Ghose et al., 2013). The smaller screen sizes make that, usually, only one banner ad is displayed at a time on mobile devices, whereas multiple banner ads may compete for the consumer's attention on desktop screens (Grewal et al., 2016). Mobile banner ads may thus grab a larger part of the consumer's attention than desktop banner ads. Also, consumers typically interact with mobile devices through touchscreens instead of a keyboard and mouse. Touchscreens may increase consumers' psychological ownership and product valuation levels due to their haptic nature (Brasel & Gips, 2014). Thus, the purchase effect of mobile banner ads displayed on touchscreens may be stronger than the effect of desktop banner ads displayed on devices with keyboard and mouse interface. Moreover, the sense of ownership generated by mobile banner ads may trigger consumers to purchase in the offline channel to quickly obtain the physical product.

Finally, compared with desktop devices, mobile devices are more personal and are typically kept closer to the user throughout the entire day (Grewal et al., 2016; Shankar, Venkatesh, Hofacker, & Naik, 2010). Compared with desktop banner ads, mobile banner ads may thus be viewed as being more intrusive as they enter a consumer's personal space. As a mobile device typically displays only one banner ad at a time, this ad will bear the brunt of the negative feelings. Compared to desktop banner ads, mobile banner ads may thus generate smaller within-device effects, cf. Bart et al. (2014) who find no evidence for an immediate impact of mobile banner ads on consumer attitudes for low-involvement products.

We conclude that the knowledge regarding desktop ad effects does not necessarily generalize to mobile banner ads due to key differences between mobile and desktop devices and their usage situations.

3. Methodology

3.1. Overall approach

To test the impact of mobile banner advertising on offline and online sales, we employ a geographical field experiment in the style of Blake et al. (2015). Field experiments have enjoyed increasing popularity for testing the effects of online advertising (e.g., Bart et al., 2014; Lambrecht & Tucker, 2013) as they offer several important advantages. Most notably, consumers are unaware of being studied (Gneezy, 2017) and variation in ad exposure between the treatment and control group is controlled by the researcher instead of being due to endogenous targeting practices.

In our field experiment, for a period of five days, we expose consumers to mobile banner ads only in a geographically-defined treatment area. More specifically, consumers in three Dutch provinces are exposed to mobile banner ads (the treatment group). Consumers in the remaining nine provinces are not exposed to mobile banner ads for the focal firm (the control group). Apart from the field experiment, the focal firm does not set advertising budgets geographically. We rely on a geographical experimental setup, as opposed to a fully randomized design at the individual consumer level (e.g., Bart et al., 2014), to overcome two important practical difficulties associated with the latter approach in the context of mobile banner advertising. First, at the individual consumer level, it is impossible to link offline sales to mobile ad exposure without the use of loyalty cards, coupons, and/or surveys (cf. Danaher & Dagger, 2013; Lewis & Reiley, 2014; Lobschat et al., 2017). Second, cross-device sales, e.g., a customer is exposed to a mobile banner ad but buys on a desktop computer, are impossible to track at the individual consumer level without a full cross-device cookie implementation or the requirement that consumers are logged in to a browser or retailer's website on both devices. We identify the sales impact of the mobile banner ads by applying a difference-in-difference approach to twenty months of region-level offline and online sales data for the focal firm. For each region, we know whether it is located in the treatment or control area.

3.2. Data collection details

The field experiment is carried out in collaboration with online marketing agency Blue Mango Interactive and the Dutch State Lottery. The Dutch State Lottery is one of the largest lotteries in the Netherlands. Monthly draws are on the 10th day of each month. Consumers can buy tickets for a single draw or subscribe to the lottery. In our analyses, we focus on non-subscription lottery ticket sales. Tickets can be purchased offline and online. The lottery thus adopts a semi-digital business model. We define online sales as the sum of mobile, tablet, and desktop sales, i.e. we focus on the broader online sales channel. The lottery offers full tickets, at a price of 15 euro, as well as one-fifth tickets. The price of one-fifth tickets are 20% of a full ticket and prizes are prorated accordingly. Moreover, a consumer can pay a small extra amount to compete for large, so-called “XL”, prizes. Ticket prices are fixed across draws. The majority, about 90% of tickets, are purchased through the offline channel, but a slight negative trend in offline sales can be observed, whereas online sales show a modest increase over time, i.e. consumers are gradually switching to the online channel. For the lottery, shifting sales from the offline to the online channel is beneficial as the costs of serving customers is lower in the online channel. About 2.6 million tickets are sold per draw (Staatsloterij, 2016). A large percentage of these tickets are sold on a non-subscription basis. For confidentiality reasons, we cannot disclose more detailed sales figures.

We obtain offline and online euro sales for all 20 monthly draws from October 10, 2013, to May 10, 2015, from The Dutch State Lottery. The sales data cover the entire country of the Netherlands. Due to offline and online sales data being drawn from separate databases (internal sales data versus Google Analytics), the definition of a region differs across these two datasets. Specifically, offline sales are at the level of regions, typically representing a city or (collection of) town(s) or village(s). For example, the city of Haarlem is a single region, and this region may include very small villages that are geographically connected to the city of Haarlem. For each region, we observe the four-digit zip codes of the stores contributing to the region's sales. Online sales are also at the region level, where we observe the name of the largest city or town in that region. Since online sales figures are smaller than offline sales, we decide to analyze online sales at a higher level of aggregation: the two-digit zip code level. We use the online sales regions' identifying cities or towns to allocate these regions to two-digit zip code regions. The smaller number of regions in the online sales data may make it harder to detect an online sales effect. We emphasize that the offline and online sales data cover the same consumers, i.e. *all* consumers in the Netherlands eligible to buy lottery tickets (all consumers aged 18 and over). Moreover, the treatment and control group sales pertain to the same groups of consumers in the offline and online data because treatment is linked to the province the region is in, and we observe the province for each region in the offline and online data.[3] A total of 782 regions are distinguished in the offline sales data, whereas the online sales data pertain to each of the 90 two-digit zip code areas in the Netherlands. Hence, our offline and online sales data are panel data sets that comprise 15,640 observations (782 regions, 20 draws) and 1800 observations (90 regions, 20 draws), respectively.

Mobile banner ads are geo-targeted to consumers using a mobile device in any of the three selected provinces, where tablets are not classified as mobile devices. We target the three selected provinces with mobile banner ads from January 6 until January 10, 18:00 CET, i.e., the days leading up to our focal draw, the 16th draw in our dataset, on January 10, 2015.[4] We obtain the total number of mobile banner ads served from AppNexus, the platform that loads the ads into the webpages.

The three provinces that are targeted with mobile banner ads, Drenthe, Overijssel, and South-Holland, are home to about 4.2 million inhabitants aged 18 years and older, or 31% of the total population of the Netherlands (all reported demographic information pertains to 2014 and is obtained from interactive tables available from Statistics Netherlands (CBS, www.cbs.nl)). A total of approximately 9.4 million people inhabit the remaining nine provinces (69% of the population). The province of South-Holland is non-adjacent to the other two provinces, which are adjacent to each other. More specifically, the province of South-Holland is in the densely populated west of the Netherlands. Drenthe lies in the more rural north, and Overijssel is in the medium-populated east of the Netherlands.[5] The decision to target non-adjacent provinces helps us to rule out regional events as a confounding factor in one of our robustness checks. However, the targeting of non-adjacent provinces also increases the probability of consumers traveling between the treatment and control group, for example, because they live in a region that is in the treatment area while working in a region in the control area. We note that potential consumer movement between the treatment and control area gives us conservative results. Assuming a positive impact of mobile banner ads on sales, treatment group consumers that travel into the control area “import” the impact of the mobile banner ads into control area sales.

We acknowledge that our allocation of regions to the treatment and control groups is non-random. By non-randomly allocating geographical areas to the treatment and control group, we deviate from Blake et al. (2015) who apply a treatment to a random selection of DMAs. Random allocation of geographical areas is complicated in our study by two important differences between our and Blake, Nosko, and Tadelis' setting. First, Blake, Nosko, and Tadelis run their experiment in the U.S. which has, due to having nearly 20 times the population of the Netherlands, far more geographical areas of substantial size. Random allocation of a small number of geographical areas of substantial size, e.g., the 12 provinces, would not provide any benefits. Second, Blake, Nosko, and Tadelis study the effects of paid search advertising, which are likely to be immediate: search results directly influence the next browsing step.

The sales effects of mobile banner advertising, our focus, may not be immediate. Due to the delay between ad exposure and sales response, we cannot randomly allocate small geographical areas as this would greatly increase the probability of consumers traveling between the treatment and control area during this delay. We emphasize that we identify the mobile banner advertising effect by analyzing panel data instead of a cross-sectional comparison of the treatment and control group, which would require near-identical groups in terms of population characteristics. Nevertheless, the non-random allocation of regions may result in different demographics in the treatment and control group, thus opening the door for competing explanations for our results. We rule out such competing explanations in our robustness checks.

3.3. Data description

Table 1 provides an overview of the treatment and control group. We observe that offline and online sales percentages represented by the treatment and control group are roughly in line with population figures. The treatment group was exposed to over 3.5 million mobile banner ads while the control group was not exposed to mobile banner advertising. For all draws other than the focal draw, offline and online advertising budgets were not set geographically.

Table 1. Overview treatment and control group characteristics.

Group	Treatment	Control
Provinces	Drenthe, Overijssel, South-Holland	Flevoland, Friesland, Gelderland, Groningen, Limburg, North-Brabant, North-Holland, Utrecht, Zeeland
Population 18 years and older	4.2 M (30.76%)	9.4 M (69.24%)
Number of regions (offline sales)	202 (25.83%)	580 (74.17%)
Number of regions (online sales)	22 (24.44%)	68 (75.56%)
Offline euro sales excl. focal draw (%)	31.39%	68.61%
Online euro sales excl. focal draw (%)	29.31%	70.69%
Mobile banner ads focal draw	3.5 M	0

Due to a confidentiality agreement, we cannot disclose absolute sales figures.

Next, we compare the demographics of the regions in the treatment and control area. We focus, at the region level, on (log-transformed) population size, average household size, proportion of retirees, and average income (in thousands of euros). We focus on these variables because they are important drivers of a consumer's buying decisions (Kotler & Armstrong, 2016). Population size informs about whether consumers follow a big-city or a village lifestyle. Household size and the percentage of retirees inform about consumers' age and lifecycle stage. Finally, average income captures a consumer's economic situation and also relates to occupation. For the offline data, demographics data pertain to regions' identifying cities or towns and demographics data for the online data are at the two-digit zip code level.

Table 2. Summary statistics of demographic variables in treatment and control area.

Variable	Group	M	ΔM	SD	Min	Median	Max
ln(Population)	T	9.331	0.282	1.063	7.432	9.083	13.240
	C	9.050		1.120	1.609	8.936	13.511
Household size	T	2.393	0.060	0.188	1.772	2.378	3.011
	C	2.333		0.188	1.143	2.346	3.371
Proportion 65+	T	0.198	−0.008	0.040	0.104	0.195	0.323
	C	0.206		0.040	0.000	0.204	0.385
Income	T	3.655	0.061	0.377	2.923	3.620	5.271
	C	3.593		0.428	1.532	3.558	5.960

T = treatment area, C = control area, ΔM indicates the difference in means between the treatment and control area for the relevant variable. Numbers pertain to the offline dataset. Numbers for the online dataset are presented in Table A1 in the Online Appendix.
T = treatment area, C = control area, ΔM indicates the difference in means between the treatment and control area for the relevant variable. Numbers pertain to the offline dataset. Numbers for the online dataset are presented in Table A1 in the Online Appendix.

In Table 2 and Fig. A1 in the Online Appendix, we compare the distributions of these four variables for regions in the treatment and control area in our offline data. We observe that the differences in means between the treatment and control area are relatively small. Moreover, the univariate distributions of the variables for the treatment and control area show a large degree of overlap. In Table A1 in the Online Appendix, we provide the distributions of the four demographic variables for regions in the treatment and control area in the online data. Again, we observe small differences in means between the treatment and control area.

3.4. Mobile banner ad characteristics

All banner ads for the field experiment are purchased through real-time-bidding (RTB). Under RTB, ad slots are sold in a Vickrey auction where the winner pays the second highest bid (Vickrey, 1961). Advertisers automatically place bids in a small fraction of time, before the webpage is fully loaded. The placed bid is based on an algorithm, which adjusts bids based on, among others, the ad position and the frequency and recency of mobile banner ad exposures at the user level. Bids are optimized based on whether a consumer spends at least 15 s on the lottery's website, and consumers are exposed to a maximum of ten ads per day and a maximum of one ad per minute. The maximum bid price is around euro 0.002 per banner (or a CPM of euro 2) and the maximum price paid is close this number, where we note that this is the price paid by the agency that purchases the ad; the price paid by the lottery is higher due to both variable and fixed campaign costs. Only high-quality websites such as sports and news websites from well-known publishers are included when bidding, preventing exposures on low-quality websites such as illegal download or scam sites. A single algorithm is used for the purchase of all mobile banner ads. Importantly, during the field experiment, the bid is only placed when a user is located within the geographical boundaries of the treatment area.

The mobile banner ads used in the field experiment are designed to generate sales, as opposed to image campaign ads aimed at brand building. Based on previous literature, the banner ads are animated and include a call to action (Chandon, Chtourou, & Fortin, 2003). The call to action, “Buy your State Lottery ticket now” (translated from Dutch) does not mention the offline or online sales channel. A single mobile ad design is used, i.e., all consumers in the three selected provinces are exposed to the same mobile banner ad. The mobile banner ads do not mention the price of a lottery ticket or offer a discount. A click on a mobile banner ad takes the consumer to an overview page on the lottery's website, showing the different types of lottery tickets available. After selecting one of these lottery tickets, the consumer is taken to the product detail page for the specific ticket. We stop targeting consumers that reached a product detail page because these consumers either purchased a lottery ticket using their mobile device or, after looking at the ticket's details, decided not to purchase using this device (possibly they did purchase on a desktop device or in an offline store). Consumers that visited the product detail page but, after deliberation, decided not to purchase in any channel are unlikely to change their mind due to banner ads, given that lottery tickets can be characterized as impulse purchase products. Continued targeting of these consumers thus yields little to no additional sales and only increases costs of the campaign. We expect no noticeable effect on our results from the decision to stop targeting consumers that visited the product detail page on a mobile device.

3.5. Model specification

We obtain our base estimate of the effect of mobile banner advertising ($Mobile_{it}$) on offline and online euro sales ($EuroSalesOffline_{it}$ and $EuroSalesOnline_{it}$, respectively) by applying the difference-in-difference method cast in fixed effects regression models (we refer to Angrist & Pischke, 2008, Chapter 5, for an introduction to the differences-in-differences identification strategy). We specify the fixed effects regression models at the region-draw level and denote a region by *i* and a draw by *t*. We note that the definition of a region varies in the offline and online sales data, but, for ease of exposition, denote both by *i*. We log transform the dependent variables to ensure strictly positive euro sales outcomes, and specify $Mobile_{it}$as a dummy variable that is one for the focal draw, i.e. the draw supported by the geographically targeted mobile banner advertising campaign ($t = 16$), and when region i is in any of the three provinces that make up the treatment group; otherwise $Mobile_{it}$ is zero. Apart from region fixed effects, we include several covariates to rule out competing explanations for our results.

First, we include draw fixed effects (β_{2t} and γ_{2t} in the offline and online sales equation, respectively) to control for the overall attractiveness of a draw. Second, we include a linear time trend for regions in the treatment area ($\beta_3 Treatment_i t$ and $\gamma_3 Treatment_i t$ in the offline and online sales equation, respectively). We obtain $Treatment_i t$ by multiplying the binary variable $Treatment_i$, which is one (zero) when region i is in the treatment (control) area, and our draw number variable t. The linear time trend for regions in the treatment area allows for the possibly changing popularity of a sales channel in the treatment area relative to the control area.[6] An important assumption of the difference-in-difference approach is that the treatment and control area do not show differential trends prior to the experiment. Hence, we expect the time trend parameters to be insignificant and discuss the non-differential trend assumption in more detail in one of our robustness checks. Finally, we control for the influence of demographic differences in the treatment and control area by allowing the focal draw's popularity to vary according to the demographics introduced earlier. We thus specify, for region i and draw t,

$$\ln(EuroSalesOffline)_{it} = \beta_{0i} + \beta_1 Mobile_{it} + \beta_{2t} + \beta_3 Treatment_i t + I(t = 16) \times (\beta_4 \ln(Population_i) + \beta_5 HHSize_i + \beta_6 Retiree_i + \beta_7 Income_i) + \varepsilon_{it} \quad (1)$$

And

$$\ln(EuroSalesOnline)_{it} = \gamma_{0i} + \gamma_1 Mobile_{it} + \gamma_{2t} + \gamma_3 Treatment_i t + I(t = 16) \times (\gamma_4 \ln(Population_i) + \gamma_5 HHSize_i + \gamma_6 Retiree_i + \gamma_7 Income_i) + \eta_{it}, \quad (2)$$

where $I(t = 16)$ is an indicator variable that takes the value one when $t = 16$ and zero otherwise, $Population_i$ indicates the population, $HHSize_i$ the average household size, $Retiree_i$ the proportion of the population over 65 years of age and $Income_i$ the average income, all for region i. We note that the main effects of the (time-invariant) demographic variables do not need to be included because we specify region fixed effects.

We estimate Eqs. (1), (2) by OLS and assume clustered robust standard errors, where we use region as the clustering variable. These standard errors are robust to heteroscedasticity and autocorrelation (Wooldridge, 2002, pp. 152–153, Vogelsang, 2012). We note that Eqs. (1), (2) share the same set of right-hand side variables. Allowing for correlated errors across equations thus does not change the results (Johnston & DiNardo, 1997, pp. 319–320).

For regions in the control area, the overall difference in (ln) euro offline sales between the tth and the 16th draw, our focal draw, is indicated by $\beta_{2,16} - \beta_{2,t}$, where we omit the demographics' influence for ease of exposition. For regions in the treatment area, the difference in (ln) euro offline sales between the tth and the 16th draw is different. For these regions the difference is $(\beta_1 + \beta_{2,16} + \beta_3 16) - (\beta_{2,t} + \beta_3 t)$. Thus, the difference in the difference between treatment and control area regions is $\beta_1 + \beta_3(16 - t)$, which reduces to β_1 if the treatment and control area do not show differential trends. A positive and significant estimate for β_1 indicates an uplift in sales due to mobile banner ads. Similarly, γ_1 indicates the impact of mobile banner ads on online euro sales. A shift in sales from one channel to the other is indicated by a positive coefficient in one equation and a negative in the other.

4. Results

4.1. Naïve analysis

Before presenting the estimation results of Eqs. (1), (2), we perform a naïve analysis of the effect of mobile banner advertising on offline and online euro sales. For each region, we determine the ratio of the focal draw sales to the average monthly draw sales in the 12 months before the focal draw. Next, we compare the distribution of the region-level ratios in the treatment and control group. We perform this naïve analysis for offline and online sales separately. With regard to offline sales, we obtain an average sales ratio of 1.123 for the treatment group (standard deviation of the mean equals 0.008). The control group shows a lower average offline sales ratio of 1.088 (standard deviation of the mean equals 0.005). An independent samples *t*-test shows that the difference between the two ratios is significant ($p < 0.001$). The online sales ratios of the treatment and control group do not differ significantly ($p = 0.340$). The average online sales ratio for the regions that are exposed to mobile banner advertising is 1.567 (standard deviation of the mean equals 0.031). A slightly lower online sales ratio of 1.531 (standard deviation of the mean equals 0.021) is obtained for the regions in the control group.[7] We visualize these results in Fig. 1. Our naïve analysis thus provides initial evidence for a positive impact of mobile banner ads on offline sales, in line with our expectations. We do not find evidence for an effect on online sales, neither for a sales uplift, nor for a negative cannibalization effect.

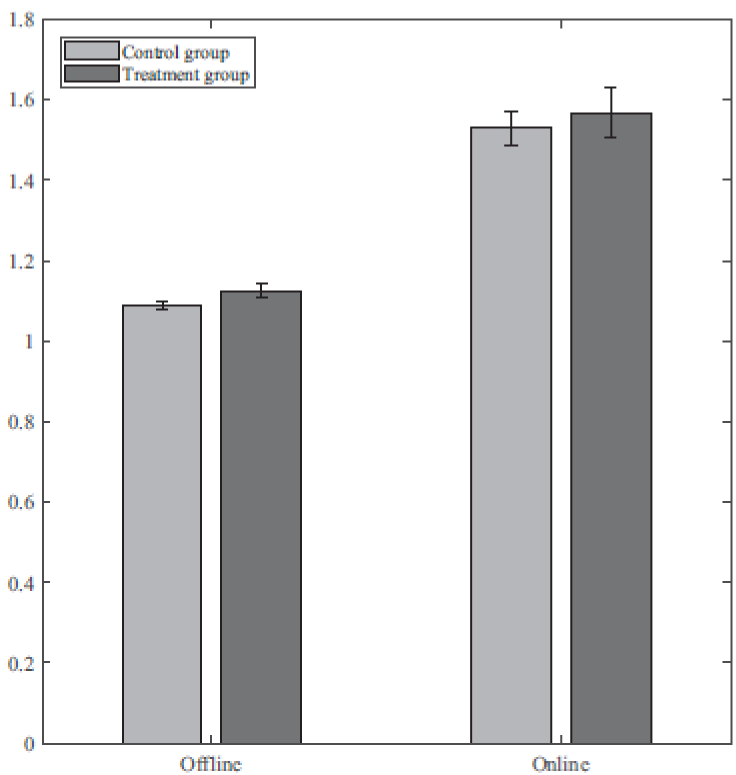

Fig. 1. Ratio of focal draw sales to average monthly draw sales. (Error bars represent twice the standard deviation of the mean.)

4.2. Estimation results

We proceed by presenting the estimation results of Eqs. (1), (2). We estimate our offline sales model, Eq. (1), on 15,640 observations (782 regions, 20 draws). In line with our expectations and naïve analysis, we obtain a significant and positive mobile advertising effect ($\beta_1 = 0.020$, $p = 0.011$). Based on this estimate, the mobile advertising campaign thus gave an offline sales uplift of 2%. The linear trend is not significant ($\beta_3 = 0.000$, $p = 0.727$). We thus find no evidence for differential trends in the treatment and control area. Two of the four demographic variables are significant. Household size is associated with a higher offline popularity of the focal draw ($\beta_5 = 0.113$, $p < 0.001$), whereas the focal draw was less popular offline in regions with a higher proportion of retirees ($\beta_6 = -0.466$, $p < 0.001$). Our fixed effects model explains a reasonable proportion of within-region variation as judged by a within-R^2 of 0.587.

We estimate the online sales model, Eq. (2), on 1800 observations (90 regions, 20 draws). The effect of mobile banner advertising on online sales is close to zero and, in line with our naïve analysis, not significant ($\gamma_1 = -0.001$, $p = 0.979$). Again, we find no evidence for differential trends in the treatment and control area ($\gamma_3 = 0.001$, $p = 0.430$). The demographic variables are not statistically significant. With a within-R^2 of 0.897, the online sales model shows a better fit than the offline sales model, likely due to the higher level of aggregation.

Our results confirm our initial findings that the effect of mobile banner advertising manifests in the offline sales channel, where we note that the offline and online sales effects do not significantly differ ($z = 0.859$, $p = 0.390$). Moreover, we find no evidence for cannibalization of online sales as the online sales impact of the mobile banner ad campaign is close to zero and not significant. We provide an overview of our estimation results in Table 3.

Table 3. Estimation results.

	Offline sales		Online sales	
	Coefficient	SE	Coefficient	SE
Mobile banner ad campaign	0.020**	0.008	−0.001	0.023
Treatment area trend	0.000	0.001	0.001	0.002
Interactions: Focal draw ×				
ln(Population)	0.003	0.005	0.042	0.027
Household size	0.113***	0.027	−0.012	0.098
Proportion 65 +	−0.466***	0.117	−0.542	0.409
Income	−0.007	0.010	0.023	0.051
Number of observations	15,640		1800	
Number of regions	782		90	
Number of draws	20		20	
Within-R^2	0.587		0.897	

Region fixed effects omitted due to space constraints, draw fixed effects omitted for confidentiality reasons. Reported standard errors are cluster robust.

** $p < 0.05$.

*** $p < 0.01$.

5. Robustness checks

We perform several robustness checks to enhance confidence in our finding that mobile banner advertising increases offline sales. Specifically, we demonstrate that our results cannot be explained by local anomalies or differential time trends for the treatment and control group. Moreover, we show robustness of our results when matching the treatment and control group on demographic variables using propensity score and coarsened exact matching. Finally, we discuss why self-selection effects do not apply to our results and show that the positive impact of mobile banner ads is not undone by a negative post-campaign effect.

5.1. Local anomalies

First, we re-estimate Eqs. (1), (2) three times, while each time omitting all regions in one of the three provinces that make up the treatment group from the analysis. The aim of this robustness check is to rule out that the results are due to a local anomaly in one or more of the regions in one of the three provinces in the treatment area. Our results, presented in Table 4, are highly robust. We obtain a significant offline sales effect in each of the three sets of results. Leaving out all regions in Drenthe, Overijssel, and South-Holland, we obtain mobile ad campaign effects (p-values) of 0.017 (0.049), 0.018 (0.038), and 0.031 (0.014), respectively. We note that these offline sales effects are around the initially obtained effect of

0.020. The online sales effect does not become significant when omitting all regions in one of the three treatment group provinces.

Table 4. Mobile banner ad campaign effect when leaving out all regions in one of the treatment provinces.

	Offline sales		Online sales	
	Coefficient	SE	Coefficient	SE
Mobile banner ad campaign				
All regions in Drenthe omitted	0.017**	0.009	−0.003	0.025
All regions in Overijssel omitted	0.018**	0.008	−0.005	0.025
All regions in South-Holland omitted	0.031**	0.012	0.019	0.030

** $p < 0.05$, reported standard errors are cluster robust.

5.2. Differential trends

A key assumption of our difference-in-difference estimator is that, in the periods *prior* to the experiment, the treatment and control group do not show differential trends (Gertler, Martinez, Premand, Rawlings, & Vermeersch, 2011, pp. 100–101). Any differential trends not accounted for in the model would provide a competing explanation for our results. We note that this assumption can only be tested when observations are available for multiple time periods before the treatment. When only one before-observation is available, this single observation is used to identify the baseline difference between the treatment and control group; no trends can be derived from this single observation. We are fortunate to observe data for fifteen draws preceding the focal draw.

We have already established that the linear trend parameters (β_3 and γ_3) are not significant. However, these parameters were estimated on data for all draws, including the focal and subsequent draws. To test the assumption that trends in the treatment and control area do not differ *prior* to the experiment, we re-estimate Eqs. (1), (2) using only data for the first fifteen draws, where we note that the variables involving the focal draw now drop out of the equations. We find that the linear trend parameters in both Eqs. (1), (2) are not significant ($\beta_3 = -0.002$, $p = 0.174$ and $\gamma_3 = 0.003$, $p = 0.298$, respectively). We thus find no evidence for differential linear trends prior to the experiment.

We also consider the possibility of nonlinear differential trends in the treatment and control area. To test for differential nonlinear trends, we again estimate Eqs. (1), (2) using only data for the first fifteen draws but now include a quadratic trend in addition to the linear trend. The linear and quadratic trend parameters (*p*-values) are 0.004 (0.241) and −0.000 (0.135), respectively for the offline data and 0.004 (0.597) and −0.000 (0.848) for the online data, respectively. We thus conclude that no evidence is found for differential nonlinear trends in the treatment and control area.

5.3. Matching estimators

In Eqs. (1), (2), we controlled for the potential influence of four demographic variables. In doing so, we made the implicit, and debatable, assumption of a linear effect of the demographic variables on log-transformed sales. An approach superior to using control variables is to create treatment and control groups that are highly similar in terms of the considered demographic variables, *prior* to model estimation, thus reducing the influence of the chosen functional form of the sales model. As Gelman and Hill (2007, p. 199) write: “Imbalance and lack of complete overlap are issues for causal inference largely because they force us to rely more heavily on model specification and less on direct support from the data.” A highly similar treatment and control group may be obtained by using matching methods. We apply two matching approaches, propensity score matching (PSM) and coarsened exact matching (CEM), where the latter approach is particularly good at aligning the multivariate distributions of the demographic variables for the treatment and control group.

To facilitate comparison of the multivariate distribution of the demographic variables for unweighted versus matching-based weighted samples, we provide the L1-statistic suggested by Iacus, King, and Porro (2011). The L1-statistic is defined as

$$\mathcal{L}_1(rf_{EG}, rf_{CG}; H) = \frac{1}{2}\sum_{h\in H}\left|rf_{EG_h} - rf_{CG_h}\right| \quad (3)$$

where rf_{EG} and rf_{CG} represent the relative frequencies for observations of the treatment and control group, respectively. H denotes the set of multivariate bins (strata) formed by first dividing each of the four demographic variables into univariate bins and then identifying all cells h that follow from the multivariate cross-tabulation of these univariate bins. Thus, rf_{EGh} indicates the relative frequency of observations of the treatment group that falls in stratum h. Perfect overlap of the treatment and control group distributions gives a L1-statistic of 0, whereas perfect separation is indicated by a value of 1. As noted by Iacus et al. (2011), what size of H to use, may seem like an arbitrary decision: applied to the same dataset, a large number of multivariate bins gives a high L1-statistic, a small number yields a low score. However, Iacus et al. (2011) find that the choice of H is typically unimportant, provided that the same set of multivariate bins is used for scoring different matching methods. The L1-statistic provides a *relative*measure, with the better matching method indicated by the lower L1 value (Iacus, King, & Porro, 2012). For each of the four demographic variables, we define a set of multivariate bins, H, of size 10^4 for the offline and of size $6^4$for the online data. In line with the lower number of regions, we opt for a lower number of multivariate bins for the online data. Per demographic variable in the offline (online) dataset, we define 10 (6) univariate bins based on equidistant cutoff-points, from the minimum to the maximum value of the respective variable. We use these bins for all reported L1-statistics. Our baseline L1 value is 0.428 and 0.868 for the offline and online data, respectively.

5.3.1. Propensity score matching

The goal of PSM is to match regions in the treatment and control group based on the probability of being allocated to the treatment area (we refer to Ghose & Todri-Adamopoulos, 2016 for a recent application of PSM to digital advertising data).[8] The probability of being allocated to the treatment area is referred to as the propensity score. Propensity scores are typically obtained from a binary logistic regression. The dependent variable in this regression is the binary variable indicating whether the region is in the treatment area (1) or not (0). Demographic variables, or other background characteristics, are used as independent variables. In our application of PSM, we use the four aforementioned demographic variables to predict membership of the treatment group. We find that, for the offline and online data, two variables are significantly associated with treatment: a region's population and average household size. Both variables are positively associated such that regions with a larger population and larger average household size are more likely to be in the treatment area. We proceed by selecting, without replacement, those regions from the control area that have a similar propensity score to the regions in the treatment area. We note that the treatment and control area now include the same number of regions.

We show a comparison of the demographics of the treatment and control area in the PSM-based offline data in Table 5—for easy reference we again provide the information from Table 2 in this table—and provide histograms of the respective distributions in Fig. A2 in the Online Appendix. Table A1 in the Online Appendix gives a comparison of the demographics of the treatment and control area in the online data. We observe that the distributions of the demographic variables for the treatment and control area now are even more comparable than before. The multivariate distribution, as judged by the L1 value of 0.450 (0.864) for the offline (online) data, shows slightly less (more) overlap than before applying PSM. The reduction in overlap after applying PSM observed for the offline data is in line with Iacus et al. (2011).

Table 5. Summary statistics of demographic variables in treatment and control area (offline data).

Variable	Matching method	Group	M	ΔM	SD	Min	Median	Max
ln(Population)	None	T	9.331	0.282	1.063	7.432	9.083	13.240
		C	9.050		1.120	1.609	8.936	13.511
	PSM	T	9.331	−0.130	1.063	7.432	9.083	13.240
		C	9.461		1.043	6.975	9.306	12.316
	CEM	T	9.307	0.065	1.030	7.432	9.075	13.152
		C	9.242		1.022	6.975	9.078	12.316
Household size	None	T	2.393	0.060	0.188	1.772	2.378	3.011
		C	2.333		0.188	1.143	2.346	3.371
	PSM	T	2.393	0.020	0.188	1.772	2.378	3.011
		C	2.373		0.180	1.806	2.383	3.204
	CEM	T	2.383	0.021	0.172	1.772	2.371	2.845
		C	2.362		0.179	1.758	2.363	2.968
Proportion 65+	None	T	0.198	−0.008	0.040	0.104	0.195	0.323
		C	0.206		0.040	0.000	0.204	0.385
	PSM	T	0.198	0.003	0.040	0.104	0.195	0.323
		C	0.195		0.036	0.096	0.195	0.295
	CEM	T	0.200	0.000	0.038	0.119	0.196	0.323
		C	0.200		0.039	0.096	0.197	0.333
Income	None	T	3.655	0.061	0.377	2.923	3.620	5.271
		C	3.593		0.428	1.532	3.558	5.960
	PSM	T	3.655	0.047	0.377	2.923	3.620	5.271
		C	3.608		0.406	1.532	3.594	5.040
	CEM	T	3.638	0.003	0.344	2.966	3.620	4.549
		C	3.635		0.334	2.729	3.620	4.861
$\mathcal{L}_1$	None		0.428					
$\mathcal{L}_1$	PSM		0.450					
$\mathcal{L}_1$	CEM		0.370					

T = treatment area, C = control area, ΔM indicates the difference in means between the treatment and control area for the relevant variable and matching method. Numbers pertain to the offline dataset.

Next, we test the offline sales effect of mobile banner advertising by estimating Eq. (1) on the sample given by PSM. The updated offline (online) estimation sample pertains to 404 (44) regions and contains a total of 8080 (880) observations. We present the estimation results in Table 6. These results confirm our previous results. We find a positive and significant impact of the mobile banner ad campaign on (ln) offline euro sales of 0.022 ($p = 0.012$), or an offline sales uplift of just over 2%. Also, we find that focal draw offline sales are higher in areas with a larger population and lower in areas with a larger proportion of consumers aged 65 or over and in areas with lower income levels. Again, the online sales effect is not significant ($p = 0.592$).

Table 6. Estimation results for matched samples.

Variable	Offline		Online	
	PSM	CEM	PSM	CEM
Mobile banner ad campaign	0.022**	0.022***	0.012	−0.001
	(0.009)	(0.008)	(0.022)	(0.025)
Treatment area trend	−0.000	0.000	0.002	0.003
	(0.001)	(0.001)	(0.004)	(0.003)
Interactions: Focal draw ×				
ln(Population)	−0.014**	−0.007	0.062	0.055
	(0.006)	(0.006)	(0.045)	(0.044)
Household size	0.047	0.044*	0.104	0.153
	(0.036)	(0.026)	(0.168)	(0.172)
Proportion 65+	−0.747***	−0.513***	−0.537	−1.997**
	(0.160)	(0.137)	(0.725)	(0.984)
Income	−0.037***	−0.000	−0.030	−0.075
	(0.013)	(0.012)	(0.076)	(0.071)
Number of observations	8080	13,880	880	1020
Number of regions	404	694	44	51
Number of draws	20	20	20	20

Region fixed effects omitted due to space constraints, draw fixed effects omitted for confidentiality reasons. Cluster robust standard errors in brackets.

* $p < 0.1$.

** $p < 0.05$.

*** $p < 0.01$.

5.3.2. Coarsened exact matching

Although PSM ensures that regions in the treatment and control area have a similar probability of being subjected to the treatment based on their demographics, differences between the treatment and control area may persist. For example, suppose that both household size and income are positively correlated with allocation to the treatment group, then a region with large household size and low income may have the same propensity score as a region with small household size and high income. PSM thus does not ensure overlap between the *multivariate* distributions of the demographic variables for the treatment and control group. To overcome this problem, Iacus et al. (2011) suggest a matching procedure called coarsened exact matching (CEM). We refer to Zervas, Proserpio, and Byers (2017) for a recent application of CEM in marketing.

CEM first bins the demographic variables. Then, for each multivariate combination of variable bins, i.e. for each stratum, we determine the number of regions in that stratum in the control and in the treatment group. Finally, we only include strata that are represented in both the treatment and control group and assign weights based on the number of regions in each group in that stratum. To illustrate this method, let us assume a focus on two demographic variables, household size and income. We would first bin both variables, for example by distinguishing above and below median household size and above and below median income. This gives us four strata. Now, we would, for each of the four strata, determine the number of regions in the treatment and in the control group. Suppose that, none of the regions in the treatment group falls in the stratum of small household size and low income, then this stratum would not be included in further analyses. The other three strata will be assigned weights based on the number of regions in each stratum in the treatment and control group.

We apply the CEM method to our four demographic variables and distinguish six (four) bins per variable for our offline (online) data.[9] We distinguish a different number of bins in our offline and online application in line with the different levels of aggregation. The number of bins give us a total of 6^4 (4^4) or 1296 (256) strata. In the offline (online) data, 85 (31) strata are represented. Moreover, we observe that 36 (9) strata are represented in both the treatment and the control group. We continue our analysis with these strata and assign weights to each stratum depending on the number of regions in a stratum in the treatment and control group. This approach leaves us with 13,880 (1020) observations pertaining to 694 (51) regions.

In Table 5 and Fig. A3 in the Online Appendix, we present the distributions of the resulting offline sample, while applying the weights. Table A1 in the Online Appendix provides summary statistics for the distributions of the resulting online sample. We observe that, in the offline sample, the differences in demographic variable means are mostly even smaller than after applying PSM and also the univariate histograms for the treatment and control group highly overlap. Importantly, we also observe a reduction in the offline (online) L1-statistic to 0.370 (0.803). This reduction is not surprising, given that the aim of CEM is to enhance the overlap of the multivariate distributions of the treatment and control group.

We re-estimate Eqs. (1), (2) on the weighted CEM sample. Table 6 displays the estimation results. Again, we obtain a positive and significant effect of the mobile banner ad campaign on offline euro sales. The estimate of 0.022 ($p = 0.009$) is near identical to the PSM result. Similar to the PSM results, we find that focal draw offline sales are lower in regions with a higher proportion of consumers aged 65 and above; the effects of population size and income are no longer significant. The online sales effect is not significant ($p = 0.962$) and focal draw online sales are lower in areas with a larger proportion of consumers aged 65 and over.

5.4. Selection

Selection is a common problem in studies on online advertising effectiveness, particularly when algorithms are used for ad targeting (Hoban & Bucklin, 2015). For example, when car ads are targeted to consumers that have shown an interest in cars, one cannot simply compare outcome metrics for users that have and have not been exposed to ads: users that are exposed to ads are by definition more interested in the advertised product because they were selected based on this characteristic. In our study, mobile banner ads are purchased through RTB, where we again note that a single algorithm is used to purchase all mobile banner ads. Certain consumers will thus be more likely to be exposed to ads than others. However, we do not compare exposed consumers to unexposed consumers. Instead, we compare those consumers in the treatment area to those that are not. We thus analyze data from *all* consumers in the treatment area, i.e., also purchases by consumers in the treatment area that are not exposed to mobile banner ads are included. Consumers do not enter the treatment group based on the RTB algorithm but are assigned based on their geographical location; variables such as browsing behavior and consumer interests have no influence whatsoever on the probability of being in the treatment or control group. Hence, we conclude that our results are not driven by selection.

5.5. Post-campaign dip

Finally, we test whether the mobile banner advertising campaign gives a net sales increase when we take a temporal perspective. Possibly, the higher sales of the focal draw are undone by lower sales in the subsequent draw, comparable to the dip commonly observed after a price promotion (Van Heerde, Leeflang, & Wittink, 2000). A post-campaign dip may occur when consumers that were planning to buy a lottery ticket for the draw following the focal draw, decided to purchase earlier because of the mobile banner ads.

To shed light on a potential effect of the mobile banner advertising campaign on subsequent draws, we extend Eqs. (1), (2) with three lagged effects of $Mobile_{it}$, i.e., $Mobile_{it-1}$, $Mobile_{it-2}$, and $Mobile_{it-3}$, and the interactions of the respective draws with the four demographic variables. Moreover, to distinguish possible lagged effects of the mobile banner ad campaign from effects due to autocorrelation, we include the one-period lagged dependent variable as a covariate. We thus obtain

$$\begin{aligned}\ln(EuroSalesOffline)_{it} = {} & \beta_{0i} + \sum_{k=0}^{3} \beta_{1+k} Mobile_{it-k} + \beta_{5t} + \beta_6 Treatment_i t + \sum_{k=0}^{3} I(t = (16+k)) \\ & \times (\beta_{7+k} \ln(Population_i) + \beta_{11+k} HHSize_i + \beta_{15+k} Retiree_i + \beta_{19+k} Income_i) \\ & + \beta_{23} \ln(EuroSalesOffline)_{it-1} + \varepsilon_{it}\end{aligned} \tag{4}$$

as our updated offline sales model and use an analogous model to explain online sales. We estimate these updated equations using OLS which gives biased results due to the presences of both fixed effects and a lagged dependent variable. To obtain unbiased results, we estimate the updated equations using the approach of Arellano and Bond (1991), where we use the first four lags as instruments. We present our results in Table 7.

We find evidence for a lagged effect of the mobile banner ad campaign on offline sales. The effect, however, is positive instead of negative. More precisely, based on both OLS and Arellano-Bond estimation results, offline sales in the draw directly following the focal draw are significantly higher. The Arellano-Bond offline sales coefficients for the focal draw, and the first, second, and third lag are 0.028 ($p = 0.002$), 0.021 ($p = 0.021$), 0.009 ($p = 0.253$), and −0.001 ($p = 0.929$), respectively. These results further demonstrate that the offline sales effect holds in specifications including a lagged dependent variable. Lagged online sales effects are not significant—only the first lag is marginally significant when using the Arellano-Bond estimator—in line with the insignificant immediate effect.

Table 7. Estimation results lagged effects.

Variable	Offline		Online	
	OLS	AB	OLS	AB
Mobile banner ad campaign				
Focal draw ($k = 0$)	0.032***	0.028***	0.002	0.009
	(0.009)	(0.009)	(0.024)	(0.023)
One-period lag ($k = 1$)	0.023**	0.021**	0.021	0.033*
	(0.009)	(0.009)	(0.018)	(0.019)
Two-period lag ($k = 2$)	0.015*	0.009	0.005	0.026
	(0.009)	(0.008)	(0.022)	(0.024)
Three-period lag ($k = 3$)	0.010	−0.001	−0.020	0.004
	(0.009)	(0.008)	(0.031)	(0.035)
Treatment area trend	−0.001	0.000	0.000	−0.002
	(0.001)	(0.001)	(0.001)	(0.002)
Lagged dependent variable	0.313***	0.162***	0.345***	0.154**
	(0.020)	(0.028)	(0.096)	(0.086)
Number of observations	14,858	14,076	1710	1620
Number of regions	782	782	90	90
Number of draws	19	18	19	18

AB = Arellano-Bond, region fixed effects and interactions omitted due to space constraints, draw fixed effects omitted for confidentiality reasons. Cluster robust standard errors in brackets.

* $p < 0.1$.

** $p < 0.05$.

*** $p < 0.01$.

6. Discussion

6.1. Conclusion and contribution

We studied the effects of mobile banner advertising on offline and online sales using a large-scale geographical field experiment and region-level time-series data. Our results demonstrate that the studied mobile banner ad campaign gave a 2% uplift in offline euro sales. The effect of mobile banner ads on online euro sales is not significant. Moreover, we find evidence for a positive post-campaign effect on offline euro sales, i.e., we find no evidence for cross-channel or cross-period sales cannibalization.

We contribute to the literature in several ways. First, we believe to be among the first to provide evidence for a sales effect of mobile banner ads, thus extending the findings of Bart et al. (2014) who focus on two intermediate metrics: attitude and intention. Also, Bart et al. (2014) measure consumer response immediately after a single ad exposure. We show that the impact of mobile banner ads on consumers is longer-lasting. In addition, we add to the broader stream of literature on mobile marketing. Previous studies mostly analyzed the effects of mobile promotions, which typically require prior consumer consent, appear as a distinct message, and offer a price discount to the consumer. We show that sales effects can also be obtained for ads that are more generically targeted, are displayed in conjunction with web content, and do not offer a price discount. Finally, we add to the growing body of literature on the effects of online advertising on offline sales (e.g., Danaher & Dagger, 2013; Johnson, Lewis, & Reiley, 2017; Lobschat et al., 2017) by showing that mobile banner ads impact offline sales.

6.2. Discussion of empirical findings

We obtain a positive effect of mobile banner ads on offline sales while we find no evidence for an online sales impact. The null result for online sales may be due to several reasons, two of which we discuss below.

A first explanation is that there simply is no online sales effect. Consumers may avoid mobile purchases due to high perceived risk (Chin et al., 2012). Instead of switching to desktop devices to conduct their purchases, consumers may prefer to buy offline, particularly when offline stores are close by or more familiar to the consumer. Moreover, mobile banner ads may increase perceived psychological ownership of the advertised product (Brasel & Gips, 2014) which may trigger consumers to obtain the physical product in the offline channel. We further note that mobile banner ads may initially create negative feelings by invading consumers' personal space, i.e. their mobile devices. However, the ad's positive influence may recover with the passage of time. This effect, referred to as the sleeper effect, occurs when the mobile banner ad's message, the positive force, becomes dissociated from the discounting cue, the negative force (Hovland, Lumsdaine, & Sheffield, 1949, Chapter 7, Kelman & Hovland, 1953). A brand cue such as out-of-store advertising or a desktop banner ad may reactivate the positive information from the mobile banner ad but not the discounting cue, resulting in an overall positive effect on attitude and purchase incidence. The positive effect may be larger in the offline channel because it is more dissociated from the context of the mobile device.

Second, mobile banner ads may significantly increase online sales, but the effect may be too small to identify (also see Bart et al., 2014) using our approach. Since the number of regions in the online dataset is smaller, the online sales effect may be harder to detect. We test whether we can pick up the offline sales effect when analyzing the offline data at the two-digit zip code level, i.e. the same level of aggregation as the online data. As shown by a significant offline sales effect (*p*-value) of 0.018 (0.046), analysis at the two-digit zip code level does not preclude detection of a sales effect. Having both datasets at the same level of aggregations also allows us to test the overall sales impact. Using (ln) total sales as the dependent variable, we obtain a (marginally) significant mobile ad campaign effect of 0.016 ($p = 0.084$). Next, we perform a formal power analysis (Burlig, Preonas, & Woerman, 2017) to determine the minimum detectable online sales effect size for our baseline online sales model (Eq. (2)). Assuming a significance level of 5% (10%) and 80% (70%) power, our model can detect an online sales effect of about 6% (below 5%). We note that these effect sizes are below the median and average conversion lift estimates of 8.1% and 19.9%, respectively, for online display ads as reported in Johnson, Lewis, and Nubbemeyer (2017). Given that offline sales are about 90% of total sales, we would (likely) have detected an online sales bump due to the mobile banner ad campaign if about 25% of the additional offline revenues would have been generated in the online channel.

6.3. Managerial implications

Our study has several implications for managers. First, managers can use mobile banner ads to influence offline consumer behavior. Our results thus are highly relevant for firms that generate (part of) their sales in the offline channel, such as lotteries, fashion retailers, department stores, fast-food chains, eateries, cafes, and more, where we caution that our results pertain to a single industry and campaign. Although we cannot reveal the costs of the mobile banner ad campaign, we can share that the uplift in euro sales outweighs the costs. Assuming a 2% uplift in euro offline sales, close to 70% of revenues being returned to lottery players in prizes, including gaming tax (Netherlands Gaming Authority, 2016), and considering other variable costs, every euro invested in the mobile advertising campaign gives about euro 0.97 in profits. This result is welcome news for managers struggling to justify or measure their mobile advertising investments—only 18% of marketers indicate to measure in-store sales lifts resulting from mobile ads (Forrester Research, 2014). Our study shows that firms that do not exclusively sell online should include mobile banner advertising in their digital media strategy. However, rising prices of mobile banner ads put the profitability of mobile banner ad campaigns under pressure. Internal ad agency data show an increase of average mobile banner ad prices of over 250% from January 2015 to January 2018. Considering this increase in banner ad prices (the variable costs of the campaign), and keeping all other numbers equal, every euro invested in the mobile advertising campaign in 2018 would give “a mere” euro 0.31 in profits. As a benchmark, Nielsen (2016) reports an average of USD 2.45 in additional sales for every dollar invested in mobile advertising for consumer packaged goods (CPG) products. Combined with an approximate operating margin of CPG manufacturers of 17% (Medeiros, Lauster, Veldhoen, & Soundararajan, 2017), this amounts to euro 0.42 in profits for every euro invested mobile advertising, a number that is in the same ballpark as our ROI estimates. Managers are advised to closely monitor the profitability of their mobile banner ad campaigns as ad prices keep rising.

Second, our results are important for digital-native firms that wish to grow their offline business. For example, Amazon could use mobile banner ads to boost Amazon Go store sales. Since digital-native firms typically have access to rich consumer data, they might be able to customize mobile banner ads to appeal to consumers' interests. In addition, our results are of interest to firms that wish to shift their sales from the offline to the online channel, for example, to reduce costs. Intuitively, managers of such firms may turn to digital advertising believing that online ads generate online sales. Counter to this belief, our results demonstrate that mobile banner ads boost offline sales. Mobile banner ads may not help in transferring sales to the online channel.

Third, while we used geographical targeting to establish the effects of mobile banner ads, this targeting approach can also be used to target the ads to specific areas where a sales boost is most desired and/or where mobile banner ads are most effective. Managers can thus easily update their advertising budgets for different regions in response to sales information or other signals.

6.4. Limitations and future research

We believe to provide an important contribution to the literature on mobile marketing and mobile banner advertising in particular. Nevertheless, our study suffers from a number of limitations which, at the same time, give rise to exciting future research avenues.

First, our results are in line with our expectation that consumers exposed to mobile banner ads are likely to purchase offline, however, future research should attempt to explicitly test the underlying mechanism. Also, future research should try to uncover factors that moderate the impact of mobile banner ads on offline and online sales. Our results reported in Table 4provide initial evidence that areas with lower population density show a stronger response to mobile banner ads—the offline and online sales effects are larger when omitting all regions in the province with the highest population density, South-Holland. In addition, we considered the role of distribution intensity. Overall distribution intensity figures for the treatment and control area are comparable at 0.274 and 0.271 offline stores per 1000 inhabitants aged 18 and over, respectively. When we interact our mobile ad campaign dummy variable ($Mobile_{it}$) with distribution intensity per region, we obtain insignificant interaction coefficients while our main effects are robust. Hence, our results indicate that future research should focus on variables related to population density or associated lifestyles, and less so on distribution intensity.

Second, our results pertain to a single country and a single industry. The Netherlands is characterized by a high smartphone penetration rate, about 80% (GfK, 2015), and lottery tickets are relatively cheap and may be purchased in an impulse. Future research should study the impact of mobile banner advertising in different countries and industries to better understand the drivers of mobile banner advertising effectiveness. In doing so, researchers are encouraged to analyze campaigns with different characteristics for firms with varying online-to-offline sales ratios.

Finally, future research should ideally conduct field experiments randomized at the individual consumer level to rule out unobserved time-varying influences. Random allocation of mobile users to the treatment and control group is straightforward, but linking the allocation to offline and cross-device sales is not. To link offline and cross-device sales to mobile ad exposure, loyalty cards or surveys may be used, where loyalty cards only give access to (a subset of) existing customers and the use of surveys likely gives a relatively small sample size.

References

Andrews, M., Goehring, J., Hui, S., Pancras, J., & Thornswood, L. (2016). Mobile promotions: A framework and research priorities. Journal of Interactive Marketing, 34, 15–24.

Andrews, M., Luo, X., Fang, Z., & Ghose, A. (2016). Mobile ad effectiveness: Hyper-contextual targeting with crowdedness. Marketing Science, 35(2), 218–233.

Angrist, J. D., & Pischke, J. S. (2008). Mostly harmless econometrics: An empiricist's companion. Princeton University Press. New Jersey: Princeton.

Arellano, M., & Bond, S. (1991). Some tests of specification for panel data: Monte Carlo evidence and an application to employment equations. The Review of Economic Studies, 58(2), 277–297.

Bart, Y., Stephen, A. T., & Sarvary, M. (2014). Which products are best suited to mobile advertising? A field study of mobile display advertising effects on consumer attitudes and intentions. Journal of Marketing Research, 51(3), 270–285.

Barwise, P., & Strong, C. (2002). Permission-based mobile advertising. Journal of Interactive Marketing, 16(1), 14–24.

Blake, T., Nosko, C., & Tadelis, S. (2015). Consumer heterogeneity and paid search effectiveness: A large-scale field experiment. Econometrica, 83(1), 155–174.

Brasel, S. A., & Gips, J. (2014). Tablets, touchscreens, and touchpads: How varying touch interfaces trigger psychological ownership and endowment. Journal of Consumer Psychology, 24(2), 226–233.

Burlig, F., Preonas, L., & Woerman, M. (2017). Panel data and experimental design. Working paper University of Chicago.

Chandon, J. L., Chtourou, M. S., & Fortin, D. R. (2003). Effects of configuration and exposure levels on responses to web advertisements. Journal of Advertising Research, 43(2), 217–229.

Chin, E., Felt, A. P., Sekar, V., & Wagner, D. (2012). Measuring user confidence in smartphone security and privacy. Proceedings of the eighth symposium on usable privacy and security.

comScore (2015). Mobile internet usage skyrockets in past 4 years to overtake desktop as most used digital platform. https://www.comscore.com/Insights/Blog/

Mobile-Internet-Usage-Skyrockets-in-Past-4-Years-to-Overtake-Desktop-as-Most-Used-Digital-Platform, Accessed date: 23 October 2016.

comScore (2016). comScore reports February 2016 U.S. smartphone subscriber market share. https://www.comscore.com/Insights/Rankings/comScore-Reports-February-2016-US-Smartphone-Subscriber-Market-Share, Accessed date: 23 October 2016.

Danaher, P. J., & Dagger, T. S. (2013). Comparing the relative effectiveness of advertising channels: A case study of a multimedia blitz campaign. Journal of Marketing Research, 50(4), 517–534.

De Haan, E., Kannan, P. K., Verhoef, P. C., &Wiesel, T. (2018). Device switching in online purchasing: Examining the strategic contingencies. Journal ofMarketing, 82(5), 1–19.

Einav, L., Levin, J., Popov, I., & Sundaresan, N. (2014). Growth, adoption, and use of mobile E-commerce. American Economic Review, 104(5), 489–494.

eMarketer (2016a). Digital ad spending to surpass TV next year. http://www.emarketer.com/Article/Digital-Ad-Spending-Surpass-TV-Next-Year/1013671, Accessed date: 23 October 2016.

eMarketer (2016b). Slowing growth ahead for worldwide internet audience. http://www.emarketer.com/Article/Slowing-Growth-Ahead-Worldwide-Internet-Audience/1014045, Accessed date: 23 October 2016.

eMarketer (2016c). US digital display ad spending to surpass search ad spending in 2016. https://www.emarketer.com/Article/US-Digital-Display-Ad-Spending-Surpass-Search-Ad-Spending-2016/1013442, Accessed date: 23 February 2017.

Fong, N. M., Fang, Z., & Luo, X. (2015). Geo-conquesting: Competitive locational targeting of mobile promotions. Journal of Marketing Research, 52(5), 726–735.

Forrester Research (2014). Master mobile measurement to unleash true cross-channel advertising. Available at https://www.4info.com/4Info/media/4Info/Resources/Whitepapers/4INFO_MasterMobileMeasurement_Aug2014.pdf?ext=.pdf.

Gelman, A., & Hill, J. (2007). Data analysis using regression and multilevel/hierarchical models. Cambridge, New York: Cambridge University Press.

Gertler, P. J., Martinez, S., Premand, P., Rawlings, L. B., & Vermeersch, C. M. J. (2011). Impact evaluation in practice. Washington DC, United States: World Bank.

GfK (2015). Geen groeimeer in bezit (mobiele) devices. http://www.gfk.com/nl/insights/press release/geen-groei-meer-in-bezit-mobiele-devices/, Accessed date: 31 October 2016.

Ghose, A., Goldfarb, A., & Han, S. P. (2013). How is the mobile internet different? Search costs and local activities. Information Systems Research, 24(3), 613–631.

Ghose, A., & Todri-Adamopoulos, V. (2016). Towards a digital attribution model: Measuring the impact of display advertising on online consumer behavior. MIS Quarterly, 40(4), 889–910.

Gneezy, A. (2017). Field experimentation in marketing research. Journal of Marketing Research, 54(1), 140–143.

Grewal, D., Bart, Y., Spann, M., & Zubcsek, P. P. (2016). Mobile advertising: A framework and research agenda. Journal of Interactive Marketing, 34, 3–14.

Guo, T., Sriram, S., &Manchanda, P. (2018). The effect of information disclosure on industry payments to physicians. Working paper. Available at SSRNhttps://ssrn.com/abstract=3064769.

Hoban, P. R., & Bucklin, R. E. (2015). Effects of internet display advertising in the purchase funnel:Model-based insights froma randomized field experiment. Journal of Marketing Research, 52(3), 375–393.

Hovland, C. I., Lumsdaine, A. A., & Sheffield, F. D. (1949). Experiments on mass communication, Vol. 3..

Iacus, S. M., King, G., & Porro, G. (2011). Multivariate matching methods that are monotonic imbalance bounding. Journal of the American Statistical Association, 106(493), 345–361.

Iacus, S. M., King, G., & Porro, G. (2012). Causal inference without balance checking: Coarsened exact matching. Political Analysis, 20(1), 1–24.

Johnson, G. A., Lewis, R. A., & Nubbemeyer, E. I. (2017). The online display ad effectiveness funnel & carryover: Lessons from 432 field experiments. Working paper. Available at SSRN: https://ssrn.com/abstract=2701578.

Johnson, G. A., Lewis, R. A., & Reiley, D. H. (2017). When less is more: Data and power in advertising experiments. Marketing Science, 36(1), 43–53.

Johnston, J., & DiNardo, J. (1997). Econometric methods (4th ed.). Singapore: McGraw-Hill.

Kannan, P. K., Reinartz,W., & Verhoef, P. C. (2016). The path to purchase and attribution modeling: Introduction to special section. International Journal of Research in Marketing, 33, 449–456.

Kelman, H. C., & Hovland, C. I. (1953). "Reinstatement" of the communicator in delayed measurement of opinion change. The Journal of Abnormal and Social Psychology, 48(3), 327.

Kotler, P., & Armstrong, G. (2016). Principles of marketing (Sixteenth ed.). Harlow, England: Pearson Education.

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

Lewis, R. A., & Reiley, D. H. (2014). Online ads and offline sales:Measuring the effect of retail advertising via a controlled experiment on Yahoo! Quantitative Marketing and Economics, 12(3), 235–266.

Lobschat, L., Osinga, E. C., & Reinartz,W. J. (2017).What happens online stays online? Segment-specific online and offline effects of banner advertisements. Journal of Marketing Research, 54(6), 901–913.

Luo, X., Andrews, M., Fang, Z., & Phang, C. W. (2014). Mobile targeting. Management Science, 60(7), 1738–1756.

Medeiros, A., Lauster, S., Veldhoen, S., & Soundararajan, R. (2016). 2017 consumer packaged goods trends: Cost cutting is not the only strategic choice for profit growth.

Strategy&, PwC Report https://www.strategyand.pwc.com/media/file/2017-Consumer-Packaged-Goods-Trends.pdf, Accessed date: 11 October 2018.

Netherlands Gaming Authority (2016). Three years of gaming under a new regulator, annual report 2015. https://kansspelautoriteit.nl/publish/pages/5163/annual_report_gaming_authority_2015.pdf, Accessed date: 9 September 2018.

New York Times (2016). Facebook profit nearly triples on mobile ad sales and new users. https://nyti.ms/2aiwn71, Accessed date: 25 February 2017.

Nielsen (2016). From ad to aisle: The CPG advertising benchmark report. https://www.ncsolutions.com/wp-content/uploads/2016/09/Multimedia-CPG-Benchmarks.pdf, Accessed date: 11 October 2018.

Shankar, V., Venkatesh, A., Hofacker, C., & Naik, P. (2010). Mobile marketing in the retailing environment: Current insights and future research avenues. Journal of Interactive Marketing, 24(2), 111–120.

Staatsloterij (2016). Je Winkans bij de Staatsloterij. https://www.staatsloterij.nl/spel/winkans, Accessed date: 23 October 2016.

The Telegraph (2016). Are mobiles changing how we shop? https://www.telegraph.co.uk/news/shopping-and-consumer-news/12172230/Are-mobiles-changinghow-we-shop.html, Accessed date: 10 October 2018.

Van Heerde, H. J., Leeflang, P. S., &Wittink, D. R. (2000). The estimation of pre- and postpromotion dips with store-level scanner data. Journal ofMarketing Research, 37(3), 383–395.

Vickrey, W. (1961). Counterspeculation, auctions, and competitive sealed tenders. The Journal of Finance, 16(1), 8–37.

Vogelsang, T. J. (2012). Heteroskedasticity, autocorrelation, and spatial correlation robust inference in linear panel models with fixed-effects. Journal of Econometrics,166(2), 303–319.

Wooldridge, J. M. (2002). Econometric analysis of cross section and panel data. Cambridge, Massachusetts: The MIT Press.

Zervas, G., Proserpio, D., & Byers, J.W. (2017). The rise of the sharing economy: Estimating the impact of Airbnb on the hotel industry. Journal ofMarketing Research, 54(5), 687–705.

Zubcsek, P. P., Katona, Z., & Sarvary, M. (2017). Predicting mobile advertising response using consumer colocation networks. Journal of Marketing, 81(4), 109–126.

The following are the supplementary data related to this article.

Table A1

Summary statistics of demographic variables in treatment and control area (online data).

Variable	Matching method	Group	M	ΔM	SD	Min	Median	Max
ln(Population)	none	T	12.232	0.319	0.520	10.987	12.294	13.240
		C	11.913		0.576	10.487	12.037	13.511
	PSM	T	12.232	-0.003	0.520	10.987	12.294	13.240
		C	12.234		0.491	11.084	12.284	13.511
	CEM	T	12.119	-0.023	0.515	10.987	12.127	13.240
		C	12.142		0.414	11.422	12.134	13.511
Household size	none	T	2.283	0.023	0.149	1.961	2.308	2.488
		C	2.260		0.164	1.761	2.297	2.511
	PSM	T	2.283	0.006	0.149	1.961	2.308	2.488
		C	2.276		0.196	1.761	2.354	2.481
	CEM	T	2.302	-0.003	0.138	1.961	2.316	2.472
		C	2.305		0.138	1.817	2.319	2.511
Proportion 65+	none	T	0.182	-0.006	0.022	0.142	0.186	0.223
		C	0.188		0.027	0.099	0.192	0.240
	PSM	T	0.182	0.004	0.022	0.142	0.186	0.223
		C	0.178		0.023	0.120	0.177	0.219
	CEM	T	0.187	0.002	0.020	0.144	0.188	0.223
		C	0.185		0.021	0.120	0.182	0.221
Income	none	T	3.489	0.042	0.271	2.954	3.487	3.944
		C	3.447		0.292	2.875	3.444	4.097
	PSM	T	3.489	-0.029	0.271	2.954	3.487	3.944
		C	3.518		0.297	2.905	3.500	4.097
	CEM	T	3.462	-0.071	0.228	2.954	3.479	3.780
		C	3.533		0.244	2.905	3.505	4.097
$\mathcal{L}_1$	none		0.868					
$\mathcal{L}_1$	PSM		0.864					
$\mathcal{L}_1$	CEM		0.803					

T=treatment area, C=control area, ΔM indicates the difference in means between the treatment and control area for the relevant variable and matching method. Numbers pertain to the online dataset.

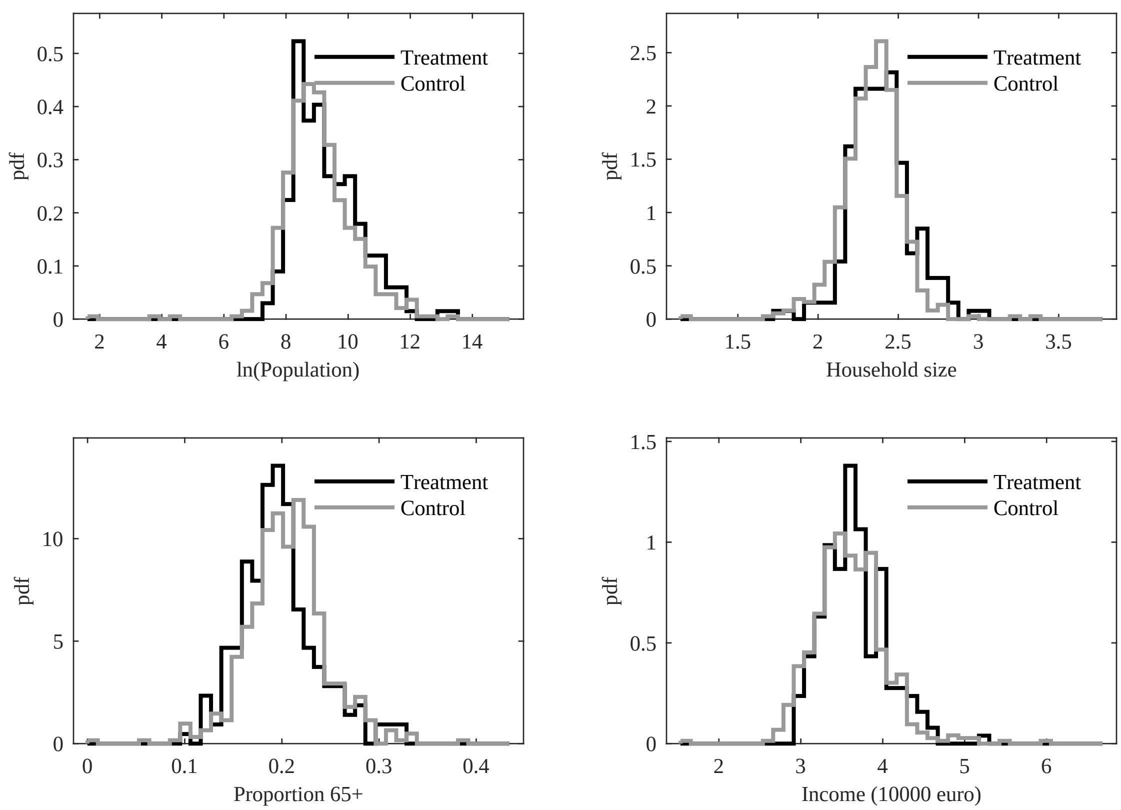
Treatment
Control
pdf
0.5
0.4
0.3
0.2
0.1
0
2
4
6
8
10
12
14
ln(Population)
Treatment
Control
pdf
2.5
2
1.5
1
0.5
0
1.5
2
2.5
3
3.5
Household size
Treatment
Control
pdf
10
5
0
0
0.1
0.2
0.3
0.4
Proportion 65+
Treatment
Control
pdf
1.5
1
0.5
0
2
3
4
5
6
Income (10000 euro)

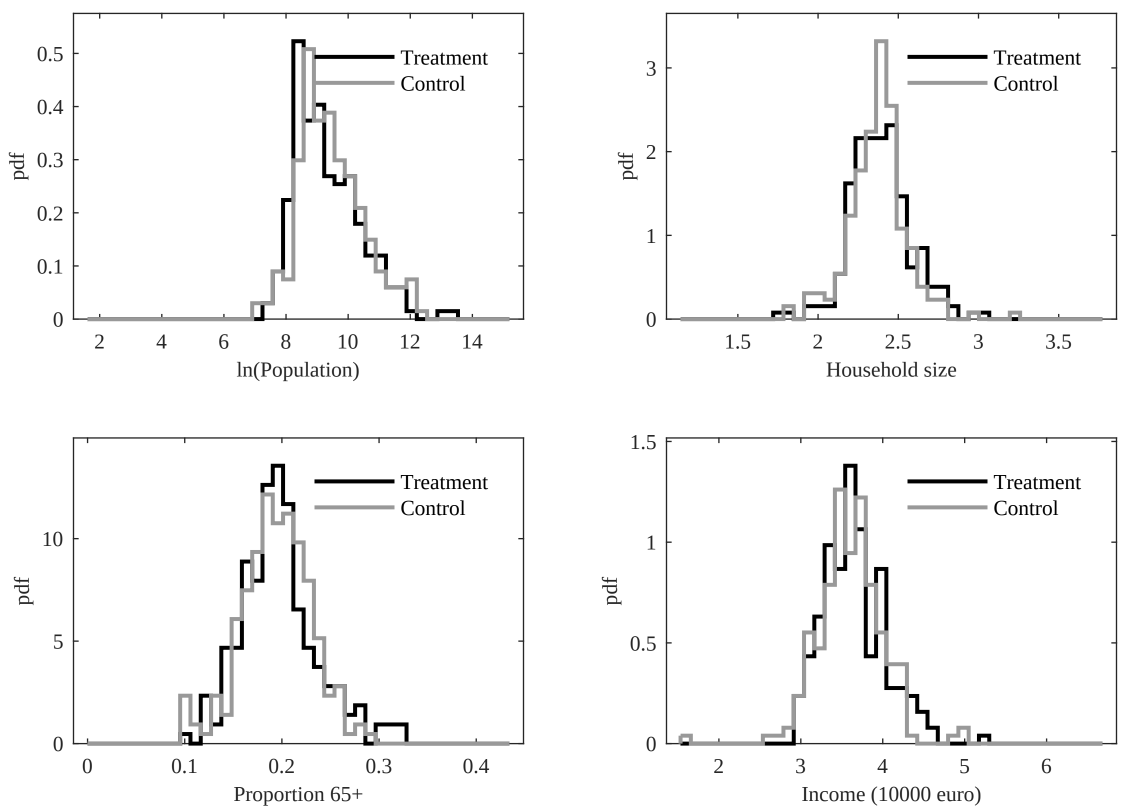
Treatment
Control
pdf
ln(Population)
Household size
Proportion 65+
Income (10000 euro)
0.5
0.4
0.3
0.2
0.1
0
2
4
6
8
10
12
14
3
2
1
1.5
2.5
3.5
10
5
0.1
0.2
0.3
0.4
1.5
1
0.5
2
3
4
5
6

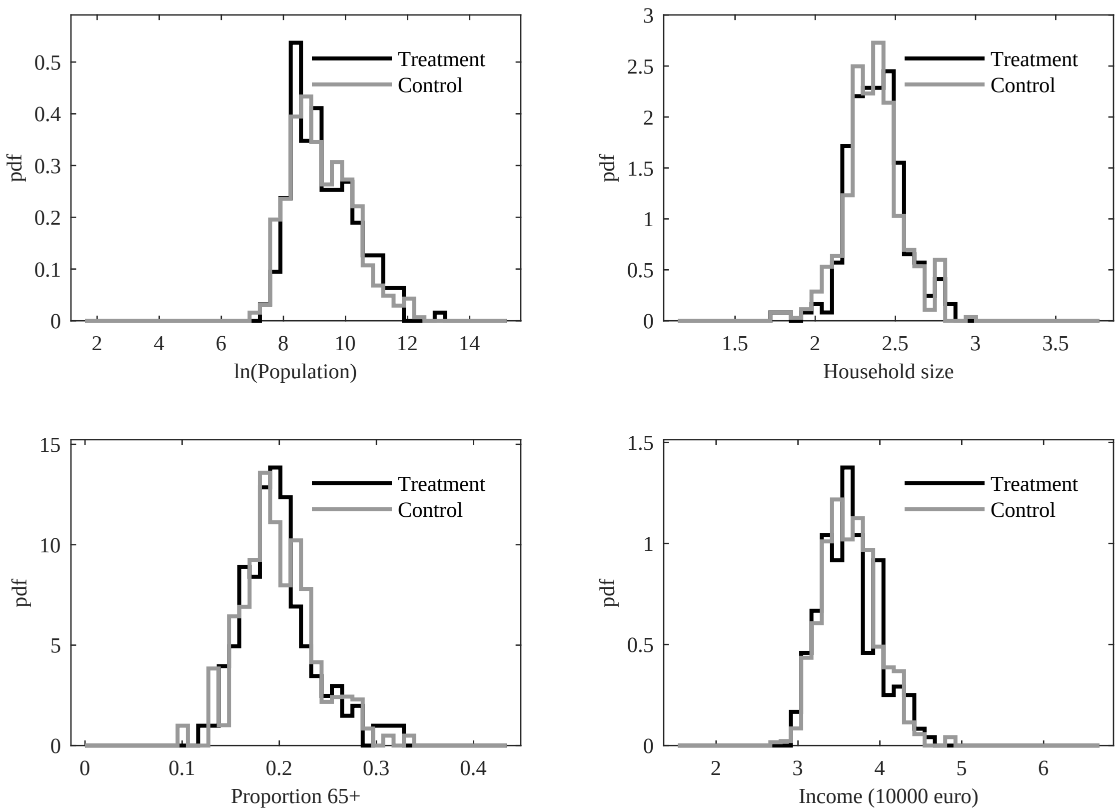

Treatment
Control
pdf
ln(Population)
Household size
Proportion 65+
Income (10000 euro)