

YAKOV BART, ANDREW T. STEPHEN, and MIKLOS SARVARY*

Mobile advertising is one of the fastest-growing advertising formats. In 2013, global spending on mobile advertising was approximately \$16.7 billion, and it is expected to exceed \$62.8 billion by 2017. The most prevalent type of mobile advertising is mobile display advertising (MDA), which takes the form of banners on mobile web pages and in mobile applications. This article examines which product characteristics are likely to be associated with MDA campaigns that are effective in increasing consumers' (1) favorable attitudes toward products and (2) purchase intentions. Data from a large-scale test-control field experiment covering 54 U.S. MDA campaigns that ran between 2007 and 2010 and involved 39,946 consumers show that MDA campaigns significantly increased consumers' favorable attitudes and purchase intentions only when the campaigns advertised products that were higher (vs. lower) involvement *and* utilitarian (vs. hedonic). The authors explain this finding using established theories of information processing and persuasion and suggest that when MDAs work effectively, they do so by triggering consumers to recall and process previously stored product information.

Keywords: mobile advertising, field experiments, advertising effectiveness, digital advertising, persuasion

Which Products Are Best Suited to Mobile Advertising? A Field Study of Mobile Display Advertising Effects on Consumer Attitudes and Intentions

The penetration of mobile phones, including smartphones, continues to rise in the United States and around the world. According to a May 2013 survey, 91% of the U.S. adult population uses some type of mobile phone, and 61%

of U.S. adult mobile users have a smartphone (Smith 2013). Smartphone penetration is even greater in younger age groups (e.g., 80% among 18- to 34-year-olds) and in higher-income households (e.g., 90% among members of households with incomes greater than \$75,000). Consumers also spend substantial amounts of time on mobile devices. For example, in 2013, the average U.S. adult spent approximately 20% of his or her daily media time on mobile devices (eMarketer 2013c). It is therefore not surprising that marketers' spending on mobile advertising has been rapidly increasing. Annual spending on mobile advertising in the United States rose from \$770 million in 2010 to \$4.1 billion in 2012, and it is forecasted to reach \$32.2 billion by 2017 (eMarketer 2012, 2013a). Globally, annual mobile advertising spending in 2013 was \$16.7 billion and is forecasted to reach \$62.8 billion by 2017 (eMarketer 2013a). Notably, most of the forecasted growth in global digital advertising spending over the next few years is due to expected

*Yakov Bart is Assistant Professor of Marketing, INSEAD (e-mail: Yakov.Bart@insead.edu). Andrew T. Stephen is Assistant Professor of Business Administration and Katz Fellow in Marketing, Joseph M. Katz Graduate School of Business, University of Pittsburgh (e-mail: AStephen@katz.pitt.edu). Miklos Sarvary is Carson Family Professor of Business and Co-Director of the Media Program, Columbia Business School, Columbia University (e-mail: miklos.sarvary@columbia.edu). All authors contributed equally to this work, and their names are listed in random order. This research was generously funded by a Google-WPP Marketing Research Award as well as the INSEAD Alumni Fund. The authors are grateful to Dan Bartels, Bill Havlena, Aaron Katz, Don Lehmann, Kara Manatt, George Pappachen, Michel Pham, Ali Rana, Keith Wilcox, members of the Social Networks and Media (SNM) Research Lab at the University of Pittsburgh, and the two anonymous *JMR* reviewers for their assistance and feedback. Avi Goldfarb served as associate editor for this article.

increases in mobile advertising, which is anticipated to constitute approximately 36% of global digital advertising expenditures by 2017 (eMarketer 2013a).

Despite strong interest, marketers' beliefs about the effectiveness of mobile advertising seem to be at best mixed, if not negative. For example, the CMO Council's (2012) survey of global marketing executives revealed that only 14% of surveyed marketers were satisfied with how they were leveraging mobile advertising channels. Instead, 43% of respondents reported that they were not satisfied with their mobile advertising efforts, and 46% reported that they were reviewing the role of mobile advertising in their organizations. Marketers nevertheless intend to keep searching for ways to use mobile advertising effectively. For example, a survey of brand marketers revealed that 69% of respondents expect to increase their use of mobile advertising in the near future (Nielsen and CMO Council 2013). Many companies, however, approach mobile advertising with a "spray-and-pray" mentality—that is, placing advertisements without any sense of how effective they will be (Patel, Schneider, and Surana 2013). Given consumers' widespread use of mobile devices and marketers' continued interest in advertising through this nascent medium, a better understanding of factors affecting mobile advertising campaign performance is needed.

This article aims to address this need by considering which types of *products* are best suited to mobile advertising. Specifically, our central research question is: Under what product-related conditions are mobile display advertisements (MDAs) effective in changing consumers' product-related attitudes and purchase intentions? Mobile display advertisements are small banner images displayed on a mobile phone's screen either in a web browser or in an application (for examples, see Figure 1). Although mobile advertisements can come in several formats (e.g., video, rich media, SMS/text message), MDAs are the most common (Zoller and Oliver 2011). Moreover, whereas the popularity of other formats has exhibited no growth or has declined (e.g., SMS/text), the popularity of MDAs has increased (eMarketer 2013b).

In this research, we focus on two psychological measures of advertising effectiveness: (1) how favorable consumers' attitudes are toward advertised products and (2) consumers' intentions to purchase or use advertised products. Although other measures of advertising effectiveness exist, attitudes and intentions are frequently used in advertising research (e.g., Grewal et al. 1997; Vakratsas and Ambler 1999) and are common campaign objectives in practice (e.g., increase favorable product attitudes, increase purchase intentions). This is also consistent with related research on the effectiveness of online (nonmobile) advertisements (Goldfarb and Tucker 2011). We do not examine effectiveness in terms of actual behaviors (e.g., purchases, product choice) for two reasons. First, in the mobile context, it is still overly ambitious to expect that MDAs trigger behaviors that are traceable and directly attributable to the MDA exposure because the majority of consumers do not yet purchase products through mobile devices (Mojiva 2012), with the exception of a few categories (e.g., mobile apps, music). Second, the campaigns included in our data set were not focused on directly inducing behaviors but rather on changing consumers' attitudes and/or intentions, and therefore, we do not

Figure 1
EXAMPLES OF MOBILE DISPLAY ADVERTISEMENTS

A: Nike Advertisement in Mobile WAP Browser



B: Match.com Advertisement in The Weather Channel App for iPhone



have measures of relevant product-related behaviors in our data set.¹

To address our research question, we use a novel data set featuring responses from 39,946 consumers across 54 MDA campaigns that ran over a three-year period (July 2007 to June 2010). Each campaign was conducted as a test-control experiment and advertised products from a variety of industries. This enables us to estimate how exposure to MDAs affected consumers' attitudes and intentions and to test whether product-related aspects can at least partially explain differences in campaign effectiveness on these metrics. To the best of our knowledge, our study is the first to use extensive field data to study MDA effectiveness.

We contribute to the nascent mobile marketing literature stream by examining a meta-analysis of campaigns to determine which types of products are better suited to being advertised with MDAs. This approach has been used in prior work on other advertising channels (e.g., Goldfarb and Tucker 2011; Lodish et al. 1995) and has the advantage of enabling us to quantify the extent to which campaign outcomes differ. This is particularly useful in a relatively new advertising channel because it helps identify the ranges of possible outcomes marketers can expect given a particular type of product or category. In addition, because advertising treatment effects tend to be small and digital advertising field experiments often are statistically underpowered (Lewis and Rao 2013), a meta-analytic approach that pools data from multiple campaigns can be helpful. However, a potential disadvantage of our approach is that it assumes that marketers have already decided to invest in MDA. Although this is a reasonable assumption given the aforementioned spray-and-pray approach, it means that the current research does not compare MDA with other forms of advertising (due to data limitations). We therefore do not consider questions regarding the relative effectiveness of MDA, how marketers should allocate advertising budgets between MDA and other channels, and when MDA should be used in combination with (or instead of) other channels. These worthwhile questions are beyond the scope of the current research but would be promising avenues for future studies. We note that multiple field studies on advertising effectiveness that focus on a single medium and use an approach similar to ours have faced the same limitation (e.g., Goldfarb and Tucker 2011; Hu, Lodish, and Krieger 2007; Johnson, Lewis, and Reiley 2013; Lewis 2010; Lewis and Reiley 2010; Lodish et al. 1995; Sahni 2012).

To preview our results, we find that most of the MDA campaigns we examined have no significant effect on consumers' product-related attitudes and intentions. However, campaign effectiveness seems to vary considerably according to the advertised product. Using two well-established general dimensions for classifying products, we show that MDA campaigns tend to be effective only for products that are both utilitarian (vs. hedonic) *and* higher involvement (vs. lower involvement). We offer a theory-based explanation for this finding using the elaboration likelihood model (ELM) from the persuasion literature (e.g., Petty and Cacioppo 1981). Briefly, we propose that when MDAs

affect attitude and intention, they do so by reminding consumers of previously encoded product-specific information, which the consumers subsequently process. Thus, the effectiveness of an MDA depends on how likely it is that it triggers memory recall that leads to central-route elaborative information processing. We argue that this is more likely for products that are both utilitarian and higher involvement.

Our findings are helpful to marketers on at least two fronts. First, given a product, marketers can have a better understanding of whether an MDA campaign is likely to work (i.e., conditional on product type, should they invest in MDA?). Second, given a decision to invest in MDA, marketers can have a better sense of how their product should be positioned in a campaign to maximize effectiveness (i.e., conditional on investing in MDA, which features of their products should they emphasize in this medium?).

The remainder of the article is organized as follows. In the next section, we consider whether MDAs are effective by reviewing the limited prior research from both academic and industry sources and then reporting the average advertising treatment effects on attitudes and intentions from the 54 campaigns represented in our data set. Then, we develop a theory to explain the differences in observed MDA campaign effectiveness on the basis of two general product classification dimensions (utilitarian/hedonic and higher/lower involvement). We then report the results of the statistical analyses used to test this theory. Finally, we conclude with a discussion of the implications of our results for theory and practice, limitations of our study, and directions for further research.

ARE MDAs EFFECTIVE?

Many marketers report achieving only moderate success or inconsistent results with mobile advertising campaigns (CMO Council 2012). Furthermore, some doubt that mobile advertising can be effective at all (Del Rey 2012; Ovide and Bensinger 2012), in large part due to some of the significant technical limitations MDAs face (Grobart 2012). Because MDAs are very small banners that are displayed on small screens, this format is typically unable to deliver information-rich messages to consumers through text, images, audio, video, or a combination of these media types (Shankar and Balasubramanian 2009). These advertisements usually contain minimal information such as a logo or a very short message or slogan (e.g., Figure 1) so that they render properly on a variety of mobile devices and display quickly over even the slowest connections. Although advertisers could create larger-sized advertisements that occupy more of the screen, this is not typically done because it is believed to irritate consumers (Patel, Schneider, and Surana 2013). Furthermore, because consumers are exposed to these advertisements on mobile devices, it is probably the case that they do not pay much attention to MDAs because they see them while they are "on the move," distracted, or attending to other environmental stimuli. Given these constraints, it is reasonable not to expect much from MDAs in terms of affecting consumers' attitudes and intentions.

Some industry studies suggest that MDAs can potentially work, albeit with relatively small effect sizes. We note, however, that it is typical for digital advertising campaigns to have small effects (Lewis and Rao 2013). For example, Pappachen and Manatt (2008) find that MDAs generated

¹Although purchase intention is an imperfect predictor of actual purchasing, it is widely used and is often a reasonable predictor subject to appropriate calibration (Sun and Morwitz 2010).

average increases in favorable brand attitudes and purchase intentions of 3.4% and 4.0%, respectively, across a variety of categories. Insight Express (a research firm) reports that MDA campaigns were up to five times more effective than comparable online display campaigns in shifting brand awareness, attitudes, and intentions (Butcher 2010). Similarly, Nielsen (2012) identifies that a substantial proportion of consumers exposed to mobile advertisements reported being more likely to subsequently purchase the advertised brand. Although the academic research on mobile marketing is scant (Shankar 2012) and MDAs in particular have not been studied, the few published studies on other limited-information mobile advertising formats (e.g., advertisements delivered through SMS/text messages) have reported positive effects on various consumer attitudes and behaviors (Barwise and Strong 2002; Drossos et al. 2007; Luo et al. 2013; Tsang, Ho, and Liang 2004).

Taking into account both the technical limitations of MDAs that constrain their information-carrying capacity and marketers' increased investments in MDA campaigns, it seems plausible that MDAs could be effective under some, but not all, conditions. We consider these conditions in relation to products after initial examination of MDA campaign effectiveness across all campaigns in our data set.

Data Set Description and Data Collection Process

Our data set is from a large U.S. market research agency that specializes in measuring the effectiveness of digital advertising, including MDAs. As mentioned previously, our data comprise 54 MDA campaigns covering a three-year period (July 2007 to June 2010).² The 54 campaigns include brands from 10 broad industries (e.g., consumer packaged goods, entertainment, finance, health).³ The agency worked with advertising networks and mobile service providers to place clients' advertisements on specific mobile web pages on targeted consumers' mobile devices. In total, data from 39,946 people who participated in these field tests were used. Table 1 reports campaign summary statistics. Mean number of participants per campaign was 739.74 (SD = 683.76), and mean campaign length was 54.80 days (SD = 32.79). The most represented industry was consumer packaged goods (33.33% of the campaigns), followed by financial services (16.67%) and automotive (12.96%). The majority of the campaigns were for products (85.19%; vs. services), business-to-consumer (B2C) brands (96.30%, vs. business-to-business [B2B]), and existing products or brand extensions (92.59%, vs. new products). Within each campaign, participants did not receive multiple exposures (but it is possible that a given participant was exposed to multiple independent campaigns).⁴

Each campaign ran as a standard digital advertising testing field study, similar to the online display advertisements by Goldfarb and Tucker (2011) investigate. Participants who browsed to one of the many mobile web pages in the

Table 1
CAMPAIGN SUMMARY STATISTICS

Number of campaigns	54
Mean number of participants per campaign (SD)	739.74 (683.76)
Mean campaign length in days (SD)	54.80 (32.79)
Percentage of campaigns for B2C brands (vs. B2B brands)	96.30%
Percentage of campaigns for new products (vs. existing products or extensions)	7.41%
Percentage of campaigns for products (vs. services)	85.19%
<i>Percentage of Campaigns by Industry</i>	
Alcohol	3.70%
Automotive	12.96%
Consumer packaged goods	33.33%
Entertainment	9.26%
Finance	16.67%
Government and nonprofit	3.70%
Health and pharmaceutical	5.56%
Restaurant	3.70%
Retail	1.85%
Technology and communications	9.26%

advertising network were randomly assigned to one of two conditions: exposed, in which an advertisement was displayed on their screen, or control, in which they were not exposed to an advertisement. In total, 19,695 (49.3%) participants were in the exposed conditions and 20,251 (50.7%) participants were in the control conditions across the 54 campaigns. Neither the agency nor their clients had direct control over the specific websites on which advertisements were placed (they were determined by the advertising network, usually by an automatic algorithm). General targeting rules were applied, however, so that a website's typical users were likely to fit the product, making it less likely that any observed campaign underperformance is due to poor targeting.

In both conditions, target web pages included a prominently placed banner the same size as the advertisement in the exposed condition that invited participants to complete a short survey (for an example, see Figure 2).⁵ We used this survey to measure attitude and intention (both on five-point scales). We measured attitude as the extent to which participants had a favorable attitude toward the product ("How would you describe your overall opinion of [product]?" 1 = "very unfavorable," and 5 = "very favorable"). We measured intention as the extent to which participants said they were likely to purchase the product the next time they were shopping in that product category ("Next time you are looking to purchase [product category], how likely are you to purchase [product]?" 1 = "very unlikely," and 5 = "very likely"). An obvious limitation is the use of single-item scales; however, the use of multi-item scales was not possible because of challenges associated with administering mobile surveys (e.g., maintaining attention, technical limi-

²The agency and its advertiser clients ran these campaigns. We did not influence any aspects of the campaigns or the field testing methodology that was used.

³For confidentiality reasons, we cannot disclose the names of the brands. All were national brands. In addition, we cannot disclose the name of the market research agency with which we collaborated.

⁴To the best of our knowledge which is based on information from the agency, there were no repeat exposures within campaigns.

⁵As Figure 2 illustrates, in the exposed condition the survey invitation was displayed next to the treatment advertisement. It is possible that this placement, although common in practice, could lead to higher treatment effect sizes because the survey conditions on participants who pay attention to the survey invitation banner. In addition, there is minimal time between the advertisement exposure and survey, which reduces the likelihood of forgetting.

Figure 2
EXPOSED VERSUS CONTROL CONDITIONS AND SURVEY INVITATION FOR FIELD TESTS



tations) and the need to keep these surveys very short.⁶ The survey included three extra items: prior awareness, recency, and category usage. Prior awareness was an aided recall question in which participants indicated whether they had heard of the product before (no/maybe/yes). Recency measured whether participants recalled having been recently (30 days before the survey) exposed to advertising for the product in other channels (no/yes). Category usage measured whether participants were light, moderate, or heavy category users.

In addition to the survey data, the agency provided non-survey data on participants and, more extensively, on campaigns. At the participant level, we knew whether each participant was deemed to match one of the campaign's target segments. The majority of participants were in a target seg-

ment (88.09% and 89.12% in the exposed and control conditions, respectively). At the campaign level, the agency provided measures for (1) product novelty (existing product, brand extension, or completely new product), (2) whether the campaign was advertising a physical good or a service, (3) whether the product was B2C or B2B, (4) campaign length (in days), (5) whether the campaign had a general focus (e.g., overall brand image) or specific focus (e.g., a specific promotion or feature), (6) whether the advertiser's goal was to improve attitude or intention, (7) industry, and (8) year.

In summary, each campaign's procedure was as follows: (1) a mobile Internet user requested a mobile web page that was in one of the agency's advertising networks, (2) the server randomly assigned the user to either the exposed or control condition, (3) the web page was loaded in that user's mobile web browser with (exposed) or without (control) the mobile display advertisement, (4) the web page also displayed a banner inviting people to complete the survey (see Figure 2), and (5) users completed the survey in their mobile web browser on a separate web page. The participants in our study are the people who completed all steps in this procedure (39,946 users). An additional 11,565 people were assigned to a condition but did not complete the survey (we address the possibility of selection bias subsequently).


Assessment of Overall Campaign Effectiveness

We quantified how well each of the 54 MDA campaigns performed with respect to improving consumers' attitudes and intentions by computing, for each campaign and metric, the average treatment effect (ATE), which is the difference between the mean attitude or intention ratings in the exposed and control conditions. Given that participants were randomly assigned to conditions, a significant positive ATE for a campaign on a given metric indicates that the campaign was effective in improving that metric. Figure 3 plots the campaign-level ATEs for attitude and intention (on their original five-point scales) from left to right in ascending order of ATE size (error bars correspond to the 95% confidence intervals).

Expressed as percentage differences between the control and exposed groups, the mean ATE across campaigns for attitude was 4.86% (SD = 9.00%, min = -5.85%, max = 43.80%, mean weighted by respondents per campaign = 6.20%), and the mean ATE for intention was 7.87% (SD = 13.18%, min = -5.35%, max = 61.18%, mean-weighted by respondents per campaign = 8.86%). Campaign-level results are less encouraging when we consider the proportions of campaigns that had significant positive, null, or significant negative ATEs on each metric, which we report in Table 2. Because we compared multiple campaigns, we determined significance on the basis of a multiple comparisons correction, for which we used Hochberg and Tamhane's (1987) linear step-up procedure (for details, see Hu, Lodish, and Krieger 2007, p. 345). Only one-third of the campaigns had significant positive ATEs on either one or both effectiveness metrics. Only 16.7% of campaigns (9) had significant positive effects on both attitude and intention, and 16.7% (9) had a significant positive effect on one but not the other. Among those campaigns with a significant positive ATE on attitude, the exposed group's attitude was on average 17.67% higher than that of the control group. Similarly,

⁶The agency that ran these campaigns and developed the field-test methodology was one of the agencies examined by Lavrakas (2010) in an evaluation of methods used to assess Internet advertising effectiveness. In general, Lavrakas finds the survey-based measurement approach to be sound given the constraints of the online medium.

Figure 3
AVERAGE CAMPAIGN TREATMENT EFFECTS



Notes: For each of the 54 campaigns, the ATE is the difference between the exposed and control groups' mean attitude or intention scores, which were measured on five-point scales. Error bars indicate the 95% confidence intervals for each ATE.

Table 2
OVERALL CAMPAIGN EFFECTIVENESS

	Attitude			Totals
	Negative	Null	Positive	
<i>Intention</i>				
Negative	0%	0%	0%	0%
Null	0%	66.7%	3.7%	70.4%
Positive	1.9%	11.1%	16.7%	29.6%
Totals	1.9%	77.8%	20.4%	100.0%

Notes: Numbers are the percentages of campaigns (out of 54) that had significant negative, null, or significant positive ATEs.

among those campaigns with a significant positive ATE on intention, the exposed group's intention was on average 21.11% higher than that of the control group.

Conversely, we found that 66.7% of campaigns (36) had null or significant negative effects on both measures. The finding that two-thirds of the MDA campaigns in our data set did not achieve positive treatment effects is concerning. The small, near-zero ATEs and wide confidence intervals for some campaigns in Figure 3 raise the possibility that some of these tests lacked sufficient statistical power. This

observation is consistent with previous findings on assessing advertising effectiveness. In particular, Lewis and Rao (2013) use 25 large online (nonmobile) advertising experiments to demonstrate the difficulty of estimating advertising treatment effects because of underpowered tests. They argue that this problem may be inherent to the estimation of advertising effectiveness because advertising treatment effects are typically small and could be caused by either high variability in dependent variables or insufficient sample sizes. In our case, high variability in attitude and intention is unlikely because both were measured on five-point scales and their standard deviations were not large (they ranged between .64 and 1.62 across campaigns and conditions). It is likely, however, that at least some of the campaigns (i.e., those with near-zero ATEs) were underpowered because of insufficient sample sizes. We further expound on this possibility in the "Discussion" section.

WHEN WILL MDA CAMPAIGNS BE EFFECTIVE?

Next, we attempt to explain the variance in campaign ATEs. The variability in MDA campaign effectiveness we found is similar to variation in advertising effectiveness found in other media, such as television (e.g., Hu, Lodish, and Krieger 2007) and online display (e.g., Goldfarb and Tucker 2011; Lewis 2010). Why some campaigns were effective and others were not, however, is unclear. Extant mobile advertising literature has provided little guidance on which factors are likely to affect MDA campaign performance. The academic literature on mobile advertising is sporadic (Shankar 2012), and the few empirical mobile advertising studies do not examine multicampaign, multi-industry MDA data (e.g., Danaher et al. 2011; Han, Ghose, and Park 2013; Hui et al. 2013). For guidance, we turn to the persuasion and information processing literature streams, which have been used to understand advertising effects in other media.

An Information Processing Perspective on MDA Effectiveness

Because MDAs do not contain much information, to understand how they work, we were tempted to draw on research covering incidental advertising exposures (Shapiro, MacInnis, and Heckler 1997), low-involvement learning about products through memory and repeated exposures (Hawkins and Hoch 1992), and preattentive mere-exposure effects (Janiszewski 1993). The mechanisms proposed in these studies involve marketing communications in which relatively little *new* information is conveyed to consumers by any particular message or advertisement (similar to MDA campaigns). Nevertheless, persuasion can occur in these situations despite the lack of information. For example, studies on low-involvement learning and, in particular, repeated message exposures (in which no new information is provided with the provision of each message) indicate that repeating claims (vs. making new claims) can be sufficient to induce attitude changes and that the effect operates through memory recall (e.g., Arkes, Boehm, and Xu 1991; Arkes, Hackett, and Boehm 1989; Hasher, Goldstein, and Toppino 1977; Hawkins and Hoch 1992; Hawkins, Hoch, and Myers-Levy 2001; Krugman 1965). In our case, however, a *direct* repeat-exposure mechanism is unlikely because in the campaigns in our data set, consumers were not repeatedly exposed to the same mobile advertising messages.

However, the aforementioned research on incidental advertising effects and repeat exposures is useful because it suggests that messages or advertisements that contain little or no new information can be effective by jogging consumers' memories. Thus, our central theoretical conjecture is that an MDA can be persuasive (and therefore effective in increasing consumers' favorable product attitudes and purchase intentions) by serving as a cue that prompts them to *recall* previously encoded product information and *process* it in a deliberate manner. Put simply, we expect that when MDAs work, they do so because they trigger memory recall and elaborative processing of the stored information. This differs from both traditional models of advertising persuasion, in which the advertisement delivers new information or makes a novel claim that consumers then process, as well as models of low-involvement learning, in which repetition leads to persuasion over time. Instead, for MDAs, we anticipate that the only way cognitive processing will lead to persuasion is if the advertisement triggers recall and processing of information from previous experiences, product encounters, or advertising exposures with the focal product or brand.

Various factors could trigger memory recall and elaborative processing. For example, the advertisement itself (i.e., copy or message characteristics) as well as how the campaign is executed (i.e., media channel, placement) have been shown to affect advertising effectiveness in traditional media channels (e.g., Arnold et al. 1987; Kirmani 1990; Stewart and Furse 1985; Vakratsas and Ambler 1999). However, MDAs are very limited in format and, in general, show minimal variation in copy, which tends to be extremely basic (e.g., a logo and short sentence or slogan). Execution options are also relatively limited for MDA campaigns given technical constraints on banner size, placement, and targeting. Because we observed large differences in MDA campaign effectiveness, but copy and execution typically exhibit minimal variance across campaigns, we do not expect these marketer-controlled factors to have major roles in driving MDA effectiveness.⁷

Another factor that might affect campaign performance is the product itself. Some products may have certain characteristics that make them more suitable for MDA campaigns. Given that we have substantial variation in the types of products in our data set (e.g., allergy medication, family cars, movies, retail banking services), it is possible that something about certain products makes them more suitable for an MDA campaign. Under the right product-related conditions, an MDA might trigger consumers to recall product information from their memories and, importantly, think more about the product in light of the recalled information. If this happens, it should be possible for consumers' attitudes and intentions to be affected.

Our central claim is that MDAs affect attitude and intention by reminding consumers of product-specific information, which they subsequently process. As such, we use the ELM (Petty and Cacioppo 1981, 1983, 1986; Petty, Cacioppo, and Schumann 1983) to develop predictions for the types of

products that are likely to trigger this process.⁸ The ELM is an information-processing theory of how beliefs (i.e., attitudes and intentions) change in response to persuasive stimuli such as advertisements. The ELM postulates two routes to persuasion: central and peripheral. People process information in a more cognitive manner under the central route but rely more on affective evaluation of stimuli under the peripheral route. Persuasion through the central route implies that consumers deliberately process message-relevant information. In contrast, under peripheral route processing, beliefs are determined by incidental environmental cues and not necessarily by the information itself. Because consumers are exposed to MDAs in environments that are likely to be highly varied, distracting, and noisy (e.g., while "on the move" or doing other things in a variety of physical settings), if an MDA is to have any systematic persuasive effect, we do not expect it to be through the peripheral route (Petty, Ostrom, and Brock 1981). Thus, we concentrate on how MDAs can trigger memory recall and central route elaborative information processing.

Petty and Cacioppo (1979, 1986) suggest that the likelihood of central route processing depends on processing motivation and processing ability, both of which have been linked to consumers' responses to advertising (MacInnis and Jaworski 1989; Petty and Cacioppo 1986; Vakratsas and Ambler 1999). Processing motivation is the extent to which a person is intrinsically motivated to engage in deliberate and thoughtful cognitive processing. A consumer with high processing motivation will have a higher elaboration likelihoods, and therefore central route processing is more likely (Mackenzie and Spreng 1992). Processing ability is the relative ease or difficulty of elaborating on (i.e., thinking about) stimuli when exposed to a persuasive message. Central route processing is also more likely with higher processing ability.

Product-Related Characteristics and Central Route Persuasion

Because we posit that MDAs are only effective when they cue memory retrieval of product information that is cognitively processed along the central route of the ELM, we next consider which types of products are associated with high processing motivation and high processing ability. We consider two general product-related characteristics that can be used to classify products (or services) across different industries or product categories and that are related to processing motivation and ability. The first is whether a product is more utilitarian or hedonic (Dhar and Wertenbroch 2000; Khan, Dhar, and Wertenbroch 2005). Utilitarian products are those for which consumption is cognitively driven, instrumental, and goal oriented and accomplishes a functional or practical task (Strahilevitz and Myers 1998). In contrast, hedonic products are those for which consumption is associated with an affective and sensory experience of pleasure, fantasy, and fun (Holbrook and Hirschman 1982). The second product-related characteristic is whether a product is higher or lower involvement with respect to how much active engagement is needed for product-focused judgment

⁷We do not suggest that copy and execution are unimportant for *other* forms of mobile advertising (e.g., rich media mobile advertisements on smartphones or full-screen mobile interstitials). Rather, in the context of MDAs and the multiple campaigns in our study, it seems unlikely that these factors play important roles.

⁸The ELM has been used extensively in the advertising literature, including research on online advertising (e.g., Shamdasani, Stanaland, and Tan 2001).

and decision-making processes (Petty, Cacioppo, and Schumann 1983; Zaichkowsky 1985). Higher-involvement products have more important personal consequences (Apsler and Sears 1968), and decisions about them carry greater risk (Bloch and Richins 1983). As a result, consumers tend to think more deliberately, carefully, and deeply about higher-involvement products when evaluating them.

Our proposed mechanism for MDAs to work relies first on memory recall. Recall, which requires sufficient attention, is more likely for higher-involvement products. Consumers pay more attention to advertisements for higher-involvement products (Holbrook and Lehmann 1980) and examine them in more detail and with greater care and attention (Celsi and Olson 1988). This is not surprising, because higher-involvement products are more personally relevant and intrinsically important to consumers (Sherif and Hovland 1961). Under the ELM, processing motivation is higher when the stimulus (product) is more personally relevant (Petty and Cacioppo 1979), which is the case for higher-involvement products. Thus, we expect higher-involvement products to engender higher levels of processing motivation in consumers, which in turn means that memory-cuing MDAs should be more likely to lead to central route processing (Mackenzie and Spreng 1992; Petty, Cacioppo, and Schumann 1983). In summary, higher-involvement products will receive more attention from consumers, who will be more motivated to process the recalled product information along the ELM's central route to persuasion.

Central route processing implies that information is processed in a more cognitive (vs. affective/emotional) manner (Petty and Cacioppo 1981). Because processing ability is also important under the ELM, we argue that MDAs for products that are easier to process cognitively will more likely lead to positive persuasion outcomes. Prior research has suggested that utilitarian (vs. hedonic) products are more likely to be processed in a more rational manner (Geuens, Pham, and De Pelsmacker 2011)—that is, using cognitive resources. As such, because it involves cognitive elaboration, central route processing is expected to be a better fit with utilitarian products than with hedonic products.⁹ In summary, under high involvement, product information recalled from memory is more likely to be processed through the central route (Petty, Cacioppo, and Schumann 1983); as a result, we expect it to have a stronger effect for utilitarian products because these products better fit the cognitive central route processing style. This reasoning suggests an interaction between product involvement (higher vs. lower) and product type (utilitarian vs. hedonic) in the sense that we only expect an MDA to have significant positive treatment effects on attitudes and intentions when the advertised product is both higher involvement and utilitarian.

AN EMPIRICAL TEST OF PRODUCT-RELATED EFFECTS ON MDA EFFECTIVENESS

We test our product-related predictions with the multi-campaign, multi-industry MDA data set described previously.

Specifically, we estimate ATEs for attitude and intention, moderated by product involvement (higher vs. lower) and product type (utilitarian vs. hedonic). We expect to find positive ATEs only for higher-involvement, utilitarian products.

Campaign Classification Procedure

As a first step, we classified the 54 campaigns' products in terms of product involvement (higher vs. lower) and product type (utilitarian vs. hedonic). We collected product-rating data from independent judges who were consumers recruited from a large U.S. online panel. Each campaign's product was considered by 19 or 20 judges ($M = 19.80$, $SD = .41$) who were each presented with the product/brand name,¹⁰ Dhar and Wertenbroch's (2000) definitions of utilitarian and hedonic products, a nine-point bipolar scale for product type (1 = "mostly hedonic," and 9 = "mostly utilitarian"), and five items on five-point Likert scales (1 = "strongly disagree," and 5 = "strongly agree") designed to measure product involvement (e.g., "Deciding whether to purchase [product] is an important decision," "Whether or not [product] turns out to be good would matter a lot to me"; $\alpha = .86$).

Although it was possible to use judges' ratings to classify each campaign's product directly as higher or lower involvement and as utilitarian or hedonic, we did not do this because of the potential for error in the judges' product ratings resulting from our inability to provide them with complete information about each campaign's product, brand, and the advertising campaign itself. This is due to confidentiality restrictions imposed by the market research agency that provided the data. However, we had access to proprietary information that may have assisted with this classification task. The agency provided their classifications of each campaign's product as higher or lower involvement, and we carefully examined each product and all available campaign information to classify it as more hedonic or more utilitarian following Strahilevitz and Myers (1998) and Dhar and Wertenbroch (2000).¹¹ We used the judges' ratings to validate the agency's product involvement classification and our product type classification, both of which were based on more complete sets of information about the campaigns.

To validate these classifications, we estimated two random-effects logit models (with campaign random effects) to check for agreement between the classifications implied by the judges and the classifications we and the agency made. First, we averaged the five involvement items (from judges) to form a single measure of involvement from each judge. We regressed the agency-provided binary measure of product involvement on respondents' ratings of involvement. A strong positive relationship between the judge-measured rating of product involvement and the agency-provided classification confirmed the validity of the agency's classification ($p < .001$). Second, we regressed our binary measure of product type on the judges' bipolar hedonic-versus-utilitarian measure. Again, we found a strong positive relationship, which confirmed the validity of our classification ($p < .01$). In addition, we examined how often our classifications matched the judges' classifications. For involvement,

⁹Hedonic products are a better fit for peripheral route processing, which we do not expect to have a systematic positive ATE, as we argued previously. Furthermore, even if MDAs were peripherally processed, it is unlikely that they would be persuasive, because random environmental cues (i.e., extraneous "noise") would affect their message.

¹⁰The proprietary nature of the data meant that we could not make detailed information available to respondents.

¹¹We also took into account how each product was framed/positioned when making these judgments.

because judges' classifications were continuous (1–5) we assumed that a judge deemed a product high (low) involvement if his or her involvement score was greater than or equal to (less than) 3. Across campaigns and judges, there was 87% agreement for product involvement. For product type, judges' classifications ranged from 1 (hedonic) to 9 (utilitarian), and we assumed that a judge deemed a product utilitarian (hedonic) if their rating was greater than or equal to (less than) 5. Across campaigns and judges, there was 80% agreement for product type.

This two-stage procedure generated a 2 (product involvement: higher, lower) × 2 (product type: hedonic, utilitarian) campaign classification. We report the distributions of campaigns and participants across these cells in Table 3 and report the exposed versus control condition sizes for each cell in Table 4. Because they were field experiments and these factors were determined post hoc, we did not expect these distributions to be uniform. Note that some industries listed in Table 3 appear in more than one cell (e.g., automotive). This is because some advertised products within the same industry were classified differently (i.e., they belonged to different industry subcategories). For example, an automotive product under hedonic and higher involvement was a luxury sports car, whereas one under utilitarian and higher involvement was a functional family car.

Selection Considerations

We next considered potential selection biases in our data. In addition to the 39,946 people who fully participated (i.e., completed the survey), another 11,565 people were randomly assigned to one of the conditions and began the survey but, for unknown reasons, did not complete it (i.e., 77.55% completion rate). Our findings therefore reflect the stated attitudes and intentions of only the people who were willing to answer all the survey questions among those who initially responded to the survey request (see Figure 2). We do not know the total response rate (i.e., the proportion of people who were invited to complete the survey and who did so), although, in line with industry norms for survey-based digital advertising effectiveness research, it is likely to be relatively small (Lavrakas 2010). We acknowledge this as a methodological limitation, even though it is consistent with industry norms and prior research (e.g., Goldfarb and Tucker [2011] faced the same problem in their study of online advertising).

Nevertheless, we attempted to econometrically control for potential selection bias due to survey noncompletion.¹²

¹²However, this control does not account for general nonresponse, because the agency did not collect data on people who were invited to participate in but did not begin the survey.

Table 4

NUMBERS OF EXPOSED VERSUS CONTROL PARTICIPANTS BY CAMPAIGN PRODUCT TYPE AND INVOLVEMENT

Type	Involvement	Exposed	Control	Total
Hedonic	High	1,751	1,762	3,513
Hedonic	Low	5,381	5,764	11,145
Utilitarian	High	7,811	6,349	14,160
Utilitarian	Low	4,752	6,376	11,128
Total		19,695	20,251	39,946

Using the little information about the 11,565 nonrespondents that was recorded (i.e., whether their mobile device was a smartphone), we employed Heckman's two-step correction procedure (Heckman 1979). The first step involved estimating a binary probit model on all 51,511 people (study participants and incompletes). We regressed whether a person was in the study or was an incomplete on their device type (smartphone or not). It is conceivable that nonsmartphone users were less likely to complete the survey because mobile web browsing on nonsmartphone devices is more difficult and can be frustratingly slow. Therefore, our main findings may be biased toward smartphone users who responded to the survey invitation.¹³ From this probit model, we computed the inverse Mills ratio for each person that, in a second step, was added as an additional explanatory variable in the main models reported in the next section. The significance of the estimated parameters for the inverse Mills ratios in the models reported subsequently indicates that this correction was needed.¹⁴

A second selection consideration was the potential for exposed and control groups not to be equivalent on some relevant dimensions. The agency and its advertising network partners attempted to ensure that exposed and control groups were representative of adult U.S. mobile phone users and were demographically equivalent. Following several conversations with agency staff responsible for these campaigns, we are confident that the participant samples were representative of the general population of adult U.S. mobile phone users. However, it is possible (though unlikely because of random assignment) that the groups were not perfectly equivalent.

Demographic equivalence may not be enough to assuage concerns about this type of selection bias in the mobile context. For example, groups could have been equivalent

¹³Even if this were the case, on practical grounds it is not a major issue given the rapid penetration of smartphones and the increasing use of MDA in smartphone apps.

¹⁴Nevertheless, robustness checks without selection correction produced qualitatively identical results.

Table 3

NUMBERS OF PARTICIPANTS AND CAMPAIGNS BY CAMPAIGN PRODUCT TYPE AND INVOLVEMENT

Type	Involvement	Campaigns	Participants	Represented Industries
Hedonic	High	8	3,513	Automotive, technology and communications
Hedonic	Low	19	11,145	Alcohol, consumer packaged goods, entertainment, restaurants
Utilitarian	High	17	14,160	Automotive, finance, government and nonprofit, health and pharmaceuticals, technology and communications
Utilitarian	Low	10	11,128	Consumer packaged goods, government and nonprofit, retail
Totals		54	39,946	

demographically but not equivalent in their mobile phone usage. Because participant-level mobile usage data were not collected, we used whether a participant used a smartphone (vs. other mobile device) to check for equivalence. This is a reasonable proxy for mobile Internet usage because smartphone users are likely to browse more mobile web pages and use more mobile apps than nonsmartphone users. Accordingly, smartphone users are probably exposed to more MDAs than nonsmartphone users. The exposed and control groups were reasonably equivalent with respect to smartphone use: 45.18% of exposed- and 46.70% of control-condition participants were smartphone users.

As a further check, we examined two additional participant-level variables described previously: whether the agency considered a participant to be in a target segment for the advertised product and a participant's self-reported level of category usage. Both indicate how relevant a given advertisement is likely to be to a participant. Differences between groups with respect to relevance could bias results, but this was not the case. For the target market variable, 88.09% of exposed- and 89.12% of control-group participants were in a target segment. For the category usage variable, 92.54% of exposed- and 93.38% of control-group participants reported heavy category usage.

Model Specification and Estimation

Our data set is a large panel data set with individual participants grouped within campaigns. We estimated a random-effects regression model for each of the dependent variables (attitude and intention) with a specification that enabled us to identify ATEs for each of the four combinations of product involvement and product type in the 2 (product involvement: higher, lower) \times 2 (product type: hedonic, utilitarian) classification. The two campaign classification variables were dummy coded (involvement: higher = 1, lower = 0; type: utilitarian = 1, hedonic = 0), as was the treatment variable (exposed = 1, control = 0).

We entered several other variables into these regressions as covariates. We included three participant-level variables to help account for participant-level observed heterogeneity, eight campaign-level variables to account for campaign-level observed heterogeneity, a campaign-level random effect to account for campaign-level unobserved heterogeneity,¹⁵ and the inverse Mills ratio to correct for the aforementioned selection bias due to survey noncompletion. Note that an alternative estimation approach would have been to estimate fixed-effects models instead of random-effects models (i.e., with campaign fixed effects). We chose a random-effects specification on the basis of Hausman tests that indicated that it was an efficient estimator (attitude: $p = .40$, intention: $p = .06$). Nevertheless, results were unchanged when we estimated the models with campaign fixed effects instead of random effects (see the Appendix).

Results

Table 5 reports parameter estimates for the selection-corrected random-effects regressions for both dependent

variables. The first four parameters listed in Table 5 (b_1 – b_4) are used for computing the ATEs for the four classes of products because the ATE is the partial derivative of the regression equation with respect to the treatment dummy variable. Importantly, for both dependent variables, the involvement \times utilitarian \times exposed effect (b_4) is positive and significant, which is consistent with our prediction (attitude: $b_4 = .24$, $t = 4.76$, $p < .001$; intention: $b_4 = .34$, $t = 5.44$, $p < .001$). As we expected, the other parameters used for computing the ATEs (b_1 , b_2 , and b_3) were not significant, except for one, which was negative (the utilitarian \times exposed effect, b_3 , on intention). This results in a slightly negative ATE for products that were classified as utilitarian and lower involvement. Although not part of our conceptualization, this still fits with the ELM because lower involvement products are more likely to be processed along the peripheral route, which is incompatible with utilitarian products. Accordingly, MDAs for this combination of product type and product involvement could conceivably result in negative performance.

In Table 6, we report estimated means for attitude and intention in exposed and control groups and the corresponding ATEs. A positive ATE indicates that the average attitude or intention was significantly higher for exposed than for control participants for that particular class of product. We predicted a significant, positive ATE for campaigns classified as higher involvement and utilitarian, which is what we found for both attitude and intention. In Table 6, we observe that the ATE for these types of products is .16 (exposed is 4.49% higher than control) and .20 (exposed is 6.69% higher than control) for attitude and intention, respectively (both $ps < .001$). As we predicted, the MDA campaigns that featured high-involvement, utilitarian products seemed to achieve positive results with respect to improving consumers' attitudes and intentions.¹⁶

We next tested the robustness of these results to differences in the model specification by estimating a series of additional random-effects models based on various assumed functional forms for the dependent variables. For each dependent variable, we estimated three models (two binary logit, one ordered logit). For the binary logit models, we dichotomized the dependent variables similar to what is often done in market research practice. These models transformed each dependent variable into a binary variable: in one model, 1 = an attitude or intent score of 5 ("top box"), and 0 = otherwise; in the other model, 1 = an attitude or intent score greater than or equal to 4 ("top two boxes"), and 0 = otherwise. The ordered logit model treated attitude and intention as discrete, ordinal variables from 1 to 5. Tables 7 and 8 report parameter estimates for attitude and intention, respectively. Except for changes in the estimated values of the treatment effects parameters (b_1 – b_4), in line with changed scaling, the substantive results were unchanged.

Finally, we tested whether the treatment effects were further moderated by prior-to-MDA product/brand aware-

¹⁵Participants were not tracked across campaigns. Therefore, we could not account for participant-level unobserved heterogeneity because we did not have participant-level panel data (i.e., repeated observations across campaigns).

¹⁶Although not part of our conceptualization, the main treatment effects for utilitarian versus hedonic (pooling across involvement conditions) and higher versus lower involvement (pooling across utilitarian and hedonic) were positive and significant. A positive main effect for utilitarian requires $2b_3 + b_4 > 0$, and a positive main effect for involvement requires $2b_2 + b_4 > 0$. Both of these requirements hold for both dependent variables according to Wald tests ($ps < .03$).

Table 5
RANDOM EFFECTS MODELS FOR ATTITUDE AND INTENTION

	Attitude		Intention	
	Estimate	t-Statistic	Estimate	t-Statistic
<i>Treatment Effects Parameters</i>				
Exposed (b ₁)	-.01	-.53	.04	1.65
Involvement × exposed (b ₂)	-.03	-.79	-.08	-1.51
Utilitarian × exposed (b ₃)	-.04	-1.32	-.10**	-2.75
Involvement × utilitarian × exposed (b ₄)	.24***	4.76	.34***	5.44
<i>Other Parameters</i>				
Intercept	3.03***	5.84	3.75***	5.92
Utilitarian	.28	1.90	.48*	2.55
Involvement	-.29	-.56	-.79	-1.21
Involvement × utilitarian	.03	.10	.04	.11
Awareness	.39***	31.38	.22***	14.11
Recency of ad exposure	.38***	28.32	.54***	32.48
Level of product category usage	.28***	12.70	.42***	15.61
Matches target market segment(s) for campaign	.09**	3.21	.09*	2.45
Product novelty	-.21	-1.53	-.36*	-2.10
Physical product vs. service	.55**	2.58	.59**	2.18
B2C vs. B2B	.31	1.19	.69*	2.09
Campaign length in days	.00	-.32	.00	-.60
General vs. specific campaign focus	-.16	-.88	-.02	-.10
Advertiser's goal (attitude vs. intention)	.16	1.01	.21	1.04
Industry ^a	Included		Included	
Year of campaign ^a	Included		Included	
Inverse Mills ratio (selection correction)	-.94***	-4.68	-1.66***	-7.53
Campaign random effect	.08***	3.90	.14***	3.91
-2 log-likelihood full model (controls-only model)	117,650	(117,701)	133,922	(133,986)
Akaike information criterion full model (controls-only model)	117,654	(117,705)	133,926	(133,990)
Bayesian information criterion full model (controls-only model)	117,658	(117,709)	133,929	(133,994)

p* < .05.*p* < .01.****p* < .001.

^aVariables with more than two levels and therefore multiple dummy variables. Statistical significance reported for overall significance test across dummy variables' parameter estimates.

ness ("awareness," where 1 = "no," 2 = "maybe," and 3 = "yes"). According to our theory, the treatment for higher-involvement, utilitarian products should be stronger at

Table 6
ESTIMATED MEANS AND TREATMENT EFFECTS FOR
ATTITUDE AND INTENTION

A: Attitude				
	Exposed	Control	Δ	%
<i>Utilitarian</i>				
Lower involvement	3.77	3.82	-.05*	-1.31
Higher involvement	3.72	3.56	.16**	4.49
<i>Hedonic</i>				
Lower involvement	3.52	3.53	-.01	-.28
Higher involvement	3.20	3.25	-.05	-1.54
B: Intention				
<i>Utilitarian</i>				
Lower involvement	3.68	3.74	-.06*	-1.60
Higher involvement	3.19	2.99	.20**	6.69
<i>Hedonic</i>				
Lower involvement	3.30	3.26	.04	1.23
Higher involvement	2.44	2.47	-.03	-1.21

p* < .05.*p* < .001.

Notes: Δ = ATE = Exposed - Control. % = 100 × (Exposed - Control)/Control. Significance tests based on an F-test of Δ = 0.

higher levels of general product/brand prior awareness if our posited process (whereby prior information is recalled from memory) is correct. We should therefore observe a positive interaction between awareness and each of the other treatment effects parameters in the previous regression models. This was the case for both attitude and intention, and it thus provides further support for our theoretical explanation.

DISCUSSION

Given the ongoing growth in marketing spending on mobile advertising, alongside increasing penetration of mobile devices (particularly smartphones) among consumers, a better understanding of mobile advertising is needed. As an initial step in addressing this need, this study identifies which types of products seem to be well suited to MDA (one of the most popular forms of mobile advertising). We used a rich data set covering 54 MDA campaigns representing products and services from ten diverse industries, in which 39,946 consumers participated in test-control advertising effectiveness field experiments. The metrics of interest were consumers' favorable attitudes toward advertised products and purchase intentions. With this data set, we were able to examine (1) how effective MDA campaigns were in improving these two common and important brand metrics and (2) how campaign effectiveness could be accounted for by certain product characteristics (i.e., higher vs. lower involvement and utilitarian vs. hedonic product category).

Table 7
ALTERNATIVE RANDOM EFFECTS MODELS FOR ATTITUDE

	Binary Logit Attitude = 5		Binary Logit Attitude ≥ 4		Ordered Logit	
	Estimate	t-Statistic	Estimate	t-Statistic	Estimate	t-Statistic
<i>Treatment Effects Parameters</i>						
Exposed (b ₁)	-.07	-1.52	-.07	-1.55	-.05	-1.32
Involvement × exposed (b ₂)	.01	.11	-.03	-.33	-.02	-.31
Utilitarian × exposed (b ₃)	-.08	-1.24	-.04	-.60	-.06	-1.09
Involvement × utilitarian × exposed (b ₄)	.62***	5.37	.38***	3.59	.42***	4.70
<i>Other Parameters</i>						
Intercept (binary logit)	.47	.33	-.12	-.09		
Intercept 5 (ordered logit)					-1.63	-1.33
Intercept 4 (ordered logit)					-.72	-.58
Intercept 3 (ordered logit)					1.56	1.26
Intercept 2 (ordered logit)					2.20	1.78
Utilitarian	.48	1.43	.54	1.86	.53	1.81
Involvement	-1.32	-1.14	-.94	-.93	-.56	-.56
Involvement × utilitarian	.25	.41	.11	.20	.12	.22
Awareness	.61***	20.90	.74***	28.34	.67	30.15
Recency of ad exposure	.75***	26.70	.76***	26.85	.72	29.63
Level of product category usage	-.58***	-10.99	-.63***	-13.53	-.48	-12.49
Matches target market segment(s) for campaign	.25***	3.81	.38***	6.17	.18	3.54
Product novelty	-.47	-1.56	-.37	-1.41	-.39	-1.47
Physical product vs. service	1.17*	2.46	1.25**	3.01	1.21	2.91
B2C vs. B2B	1.08	1.80	.95	1.82	.69	1.35
Campaign length in days	.00	-.25	.00	-.65	.00	-.31
General vs. specific campaign focus	-.34	-.80	-.54	-1.47	-.33	-.90
Advertiser's goal (attitude vs. intention)	.49	1.34	.33	1.04	.34	1.08
Industry ^a		Included		Included		Included
Year of campaign ^a		Included		Included		Included
Inverse Mills ratio (selection correction)	-3.07***	-7.26	-2.38***	-5.96	-1.88***	-5.58
Campaign random effect	.42***	3.82	.32***	3.86	.32***	3.92
-2 log-likelihood		185,361		178,853		568,924
Akaike information criterion		185,429		178,921		568,992
Bayesian information criterion		185,497		178,989		569,060

* $p < .05$.

** $p < .01$.

*** $p < .001$.

^aVariables with more than two levels and therefore multiple dummy variables. Statistical significance reported for overall significance test across dummy variables' parameter estimates.

Our results suggest that MDA campaigns, in general, may be more likely to fail than to succeed if campaign success is indicated by a positive treatment effect (as is the case in practice). We observed that MDA campaign effectiveness, as measured by the sizes of the average campaign treatment effects, varied considerably, and the majority of the campaigns in our data set did not achieve significant positive ATEs. In general, these findings are consistent with prior research that has tested the effectiveness of other advertising media using similar test-control field experiment methods (e.g., Goldfarb and Tucker 2011; Hu, Lodish, and Krieger 2007; Lodish et al. 1995). Similar to, for example, television and online display advertising campaigns, the performance of MDA campaigns is also quite varied. Furthermore, when an MDA campaign is successful (indicated by a significant positive treatment effect), the ATE is typically small, in line with Goldfarb and Tucker's (2011) findings in their study of online display advertising. Therefore, marketers should not expect MDA campaigns to be extremely impactful on consumers' attitudes and intentions, if at all.

Notwithstanding this finding, we caution against generalizing from our study that MDAs are largely ineffective. Such a conclusion would be premature because the current

ATE estimates may suffer from insufficient statistical power, thus increasing the possibility of false-negative conclusions (i.e., Type II errors). As we mentioned previously, some of the campaigns with very small estimated ATEs could have had too few participants, thus resulting in underpowered tests for positive treatments.¹⁷ Our concern is that if some of our campaign-level tests were underpowered, we might have concluded that some campaigns had null effects when they actually had positive (albeit small) effects. While problematic, this possibility makes our general conclusions more conservative. Indeed, we believe that this is preferable to a Type I error (i.e., concluding that the campaigns had significant positive treatment effects when they did not), which could lead marketers to expect to achieve greater success with MDA campaigns than is probably warranted.

Nevertheless, it is clear that, as we expected, MDA treatment effects—if real—are probably small. Additional

¹⁷A way to mitigate this problem would be to recruit more participants for these campaigns, particularly if treatment effects are expected to be very small. However, as Lewis and Rao (2013) show, massive sample sizes are required for acceptable levels of power when digital advertising treatment effects are small. This may be infeasible because costs of running advertising field experiments linearly increase with sample size (Lewis, Rao, and Reiley 2013).

Table 8
ALTERNATIVE RANDOM EFFECTS MODELS FOR INTENTION

	Binary Logit Attitude = 5		Binary Logit Attitude ≥ 4		Ordered Logit	
	Estimate	t-Statistic	Estimate	t-Statistic	Estimate	t-Statistic
<i>Treatment Effects Parameters</i>						
Exposed (b ₁)	-.04	-.73	.02	.42	.05	1.23
Involvement × exposed (b ₂)	-.04	-.31	-.05	-.60	-.10	-1.47
Utilitarian × exposed (b ₃)	-.07	-.98	-.05	-.77	-.12*	-2.33
Involvement × utilitarian × exposed (b ₄)	.51***	3.98	.33***	3.14	.47***	5.47
<i>Other Parameters</i>						
Intercept (binary logit)	2.86	1.90	3.48*	2.72		
Intercept 5 (ordered logit)					.70	.55
Intercept 4 (ordered logit)					1.75	1.40
Intercept 3 (ordered logit)					1.56	1.26
Intercept 2 (ordered logit)					3.53***	2.81
Utilitarian	.54	1.56	.58*	1.98	.71*	2.38
Involvement	-1.56	-1.30	-2.25*	-2.22	-1.19	-1.16
Involvement × utilitarian	.36	.56	.40	.74	.13	.23
Awareness	.29***	9.52	.37***	14.30	.30***	14.18
Recency of ad exposure	.79***	27.04	.82***	29.85	.77***	32.79
Level of product category usage	.69***	12.17	.71***	14.86	.57***	15.00
Matches target market segment(s) for campaign	.22**	3.19	.31***	5.25	.13***	2.64
Product novelty	-.63*	-2.00	-.51	-1.90	-.52	-1.90
Physical product vs. service	.92	1.85	1.04*	2.50	1.05*	2.46
B2C vs. B2B	1.12	1.79	.78	1.50	1.06*	2.02
Campaign length in days	-.01	-.48	-.01	-1.03	-.00	-.57
General vs. specific campaign focus	-.20	-.46	-.17	-.48	-.08	-.22
Advertiser's goal (attitude vs. intention)	.32	.85	.28	.88	.35	1.06
Industry ^a		Included		Included		Included
Year of campaign ^a		Included		Included		Included
Inverse Mills ratio (selection correction)	-3.91***	-9.01	-3.11***	-7.99	-2.65***	-8.14
Campaign random effect	.46***	3.87	.32***	3.89	.34***	3.93
-2 log-likelihood		188,625		177,288		477,388
Akaike information criterion		188,693		177,356		477,462
Bayesian information criterion		188,761		177,424		477,524

* $p < .05$.

** $p < .01$.

*** $p < .001$.

^aVariables with more than two levels and therefore multiple dummy variables. Statistical significance reported for overall significance test across dummy variables' parameter estimates.

larger-scale field tests are needed to examine whether null effects are real or due to underpowered tests. If such tests are impractical, as Lewis and Rao (2013) suggest, further research could address this problem by running tests in which participants in the treatment condition receive multiple exposures to the same MDA, which may increase expected effect sizes and thus reduce concerns of underpowered tests and false-negative conclusions. Moreover, although it is unlikely a cause of low statistical power in our analysis, avoiding dependent variables with high variability may also help. For example, in the context of purchase decisions, binary dependent variables (e.g., buy vs. not buy) would be preferable to variables as purchase amounts.

Our analysis did find, however, that MDAs for certain kinds of products could be effective (i.e., when the product is higher involvement and utilitarian). We theorized that because higher involvement (which is necessary for product-specific information retrieval) promotes central route processing, utilitarian products are more likely to be elaborated on than hedonic ones. We also found that prior product awareness makes MDAs for higher-involvement, utilitarian products more effective. This supports our claim that, because they contain relatively little information and are used in distracting environments, MDAs are likely to oper-

ate as memory cues such that consumers process information stored in memory instead of information (if any) conveyed by the advertisement itself.

Note that although our findings suggest that MDAs are a potentially worthwhile investment for only a specific class of products, this need not be the case. Marketers with products that are more hedonic and/or lower involvement could position their products as more utilitarian and higher involvement when advertising in the mobile channel. This is similar to how advertisers tweak the framing of their products differently to fit the communication medium. This could be done by featuring product attributes that are more functional (i.e., utilitarian) and require more thought and consideration (i.e., higher involvement). For example, an MDA for a luxury car (inherently a more hedonic and higher involvement product) might be more effective if it emphasizes safety over sportiness. Furthermore, if our empirical results are indeed driven by the finding that MDAs operate as memory cues, an important practical implication is that marketers should use MDA campaigns as complements to other, "higher bandwidth" advertising campaigns. This ensures that consumers can recall relevant product information when reminded of products by MDAs.

Our findings build on extant research on advertising effectiveness, mobile marketing, and persuasion by identifying product-related conditions under which MDAs can have positive impacts on consumers' attitudes and intentions. Although we do not claim that these results generalize beyond MDA campaigns, our results on general campaign effectiveness (highly variable) and effect sizes (small) are broadly consistent with prior research on non-MDA campaigns. Similar results may therefore be found for other forms of advertising that, like MDAs, carry little (if any) information and therefore probably operate by jogging consumers' memories. Examining this possibility would be a worthwhile avenue for further research. An additional direction for further research would be to consider other factors that affect MDA effectiveness and other metrics. Our study is limited to two general product characteristics and two metrics, although other factors and metrics may also be

worth exploring. Another limitation is that our results are based on single exposures to MDAs. Future researchers might consider testing the impact of repeated exposures, which may be more likely to capture low-involvement learning or preattentive mere-exposure effects if they exist in this advertising domain. A final important direction for further research is to consider multimedia campaigns involving mobile and nonmobile advertising, which would allow for a comparative study of the effectiveness of mobile versus nonmobile advertisements.

The current research represents a first step in understanding what works (and what does not work) in a particularly fast-growing and increasingly popular form of digital advertising. We hope this article encourages additional research that examines other forms of mobile advertising and other fruitful questions pertaining to how marketers can develop effective mobile marketing strategies.

Appendix

PARAMETER ESTIMATES FROM FIXED-EFFECTS MODELS FOR ATTITUDE AND INTENTION

	Attitude		Intention	
	Estimate	t-Statistic	Estimate	t-Statistic
<i>Treatment Effects Parameters</i>				
Exposed (b ₁)	-.01	-.55	.04	1.62
Involvement × exposed (b ₂)	-.03	-.77	-.07	-1.48
Utilitarian × exposed (b ₃)	-.03	-1.34	-.10**	-2.77
Involvement × utilitarian × exposed (b ₄)	.23***	4.77	.33***	5.45
<i>Other Parameters^a</i>				
Awareness	.39***	31.39	.21***	14.11
Recency of ad exposure	.38***	28.33	.53***	32.51
Level of product category usage	.27***	12.77	.43***	15.69
Matches target market segment(s) for campaign	.08**	2.94	.08*	2.17
Inverse Mills ratio (selection correction)	-.92***	-4.59	-1.64***	-7.46
-2 log-likelihood full model		117,321		133,602
Akaike information criterion full model		117,449		133,730
Bayesian information criterion full model		117,999		134,280
Campaign fixed effects (dummy variables), joint significance test		F(54, 39,883) = 101.00, p < .001		F(54, 39,883) = 108.48, p < .001

*p < .05.

**p < .01.

***p < .001.

^aVariables that did not vary within campaign were excluded because they were collinear with the campaign fixed effects.

REFERENCES

- Apsler, Robert and David O. Sears (1968), "Warning, Personal Involvement, and Attitude Change," *Journal of Personality and Social Psychology*, 9 (2), 162-66.
- Arkes, Hal R., Lawrence E. Boehm, and Gang Xu (1991), "Determinants of Judged Validity," *Journal of Experimental Social Psychology*, 27 (6), 576-605.
- , Catherine Hackett, and Larry Boehm (1989), "The Generality of the Relation Between Familiarity and Judged Validity," *Journal of Behavioral Decision Making*, 2 (2), 81-94.
- Arnold, Stephen J., Tae H. Oum, Bohumir Pazderka, and Douglas W. Snetsinger (1987), "Advertising Quality in Sales Response Models," *Journal of Marketing Research*, 24 (February), 106-113.
- Barwise, Patrick and Colin Strong (2002), "Permission-Based Mobile Advertising," *Journal of Interactive Marketing*, 16 (1), 14-24.
- Bloch, Peter H. and Marsha L. Richins (1983), "A Theoretical Model for the Studying Product Importance Perceptions," *Journal of Marketing*, 47 (July), 69-81.
- Butcher, Dan (2010), "Mobile Ad Campaigns 5 Times More Effective Than Online: Insight Express Study," MobileMarketer.com, (February 5), (accessed February 27, 2014), [available at <http://www.mobilemarketer.com/cms/news/research/5308.html>].
- Celsi, Richard L. and Jerry C. Olson (1988), "The Role of Involvement in Attention and Comprehension Processes," *Journal of Consumer Research*, 15 (2), 210-24.
- CMO Council (2012), *Engage at Every Stage: The New Mobile Marketing Mandate*. Palo Alto, CA: Chief Marketing Officer Council World Wide.
- Danaher, Peter, Michael Smith, Kulan Ranasinghe, and Tracey Dagger (2011), "Assessing the Effectiveness of Mobile Phone Promotions," working paper, Monash University.
- Del Rey, Jason (2012), "Will Mobile's Massive Growth Ever Equal Real Revenue?" *Advertising Age*, (October 1), (accessed February 27, 2014), [available at <http://adage.com/article/digital/mobile-s-massive-growth-equal-real-revenue/237511/>].
- Dhar, Ravi and Klaus Wertenbroch (2000), "Consumer Choice Between Hedonic and Utilitarian Goods," *Journal of Marketing Research*, 37 (February), 60-71.
- Drossos, Dimitris, Geroge M. Giaglis, George Lekakos, Flora Kokkinaki, and Maria G. Stavragi (2007), "Determinants of Effective SMS Advertising: An Experimental Study," *Journal of Interactive Advertising*, 7 (2), 16-27.

- eMarketer.com (2012), "US Mobile Ad Spending Jumps to \$4 Billion," (December 18), (accessed February 27, 2014), [available at <http://www.emarketer.com/Article/US-Mobile-Ad-Spending-Jumps-4-Billion/1009548#U8PFWQfsCqDREjCc.99>].
- (2013a), "Mobile Expands Its Share of Worldwide Digital Ad Spend," (August 28), (accessed February 27, 2014), [available at <http://www.emarketer.com/Article/Mobile-Expands-Its-Share-of-Worldwide-Digital-Ad-Spend/1010170>].
- (2013b), "Mobile Gains Greater Share of Search, Display Spending," (August 21), (accessed February 27, 2014), [available at <http://www.emarketer.com/Article/Mobile-Gains-Greater-Share-of-Search-Display-Spending/1010148>].
- (2013c), "Worldwide, More Money Goes Mobile," (January 4), (accessed February 27, 2014), [available at <http://www.emarketer.com/Article/Worldwide-More-Money-Goes-Mobile/1009582>].
- Geuens, Maggie, Michel T. Pham, and Patrick De Pelsmacker (2011), "Product Involvement vs. Product Motives as Moderators of the Effects of Ad-Evoked Feelings: An Analysis of Consumer Responses to 1,100 TV Commercials," in *Advances in Consumer Research*, Vol. 38, Darren W. Dahl, Gita V. Johar, and Stijn M.J. van Osselaer, eds. Duluth, MN: Association for Consumer Research.
- Goldfarb, Avi and Catherine Tucker (2011), "Online Display Advertising: Targeting and Obtrusiveness," *Marketing Science*, 30 (3), 389–404.
- Grewal, Dhruv, Sukumar Kavanoor, Edward F. Fern, Carolyn Costley, and James Barnes (1997), "Comparative Versus Non-comparative Advertising: A Meta-Analysis," *Journal of Marketing*, 61 (October), 1–15.
- Grobart, Sam (2012), "Mobile Ads Are the Future. They're Also Lousy," *BusinessWeek*, (November 1), (accessed February 27, 2014), [available at <http://www.businessweek.com/articles/2012-11-01/mobile-ads-are-the-future-dot-theyre-also-lousy>].
- Han Sang Pil, Anindya Ghose, and Sung-Hyuk Park (2013), "Cross-Channel Synergies Between Web and Mobile Advertising: A Randomized Field Experiment," working paper, City University of Hong Kong.
- Hasher, Lynn, David Goldstein, and Thomas Toppino (1977), "Frequency and the Conference of Referential Validity," *Journal of Verbal Learning and Verbal Behavior*, 16 (1), 107–112.
- Hawkins, Scott A. and Stephen J. Hoch (1992), "Low-Involvement Learning: Memory Without Evaluation," *Journal of Consumer Research*, 19 (2), 212–25.
- , —, and Joan Meyers-Levy (2001), "Low-Involvement Learning: Repetition and Coherence in Familiarity and Belief," *Journal of Consumer Psychology*, 11 (1), 1–11.
- Heckman, James J. (1979), "Sample Selection Bias as a Specification Error," *Econometrica*, 47 (1), 153–61.
- Hochberg, Yosef and Ajit C. Tamhane (1987), *Multiple Comparison Procedures*. New York: John Wiley & Sons.
- Holbrook, Morris B. and Elizabeth C. Hirschman (1982), "The Experiential Aspects of Consumption: Consumer Fantasies, Feelings, and Fun," *Journal of Consumer Research*, 9 (2), 132–40.
- and Don R. Lehmann (1980), "The Role of Message Content Versus Mechanical Features in Predicting Recognition of Print Advertisements," *Journal of Advertising Research*, 20 (1), 53–62.
- Hu, Ye, Leonard M. Lodish, and Abba M. Krieger (2007), "An Analysis of Real World TV Advertising Tests: A 15-Year Update," *Journal of Advertising Research*, 47 (3), 341–53.
- Hui, Sam K., J. Jeffrey Inman, Yanliu Huang, and Jacob A. Suher (2013), "The Effect of In-Store Travel Distance on Unplanned Spending: Applications to Mobile Promotion Strategies," *Journal of Marketing*, 77 (March), 1–16.
- Janiszewski, Chris (1993), "Preattentive Mere Exposure Effects," *Journal of Consumer Research*, 20 (3), 376–92.
- Johnson, Garrett A., Randall A. Lewis, and David H. Reiley (2013), "Add More Ads? Experimentally Measuring Incremental Purchases Due to Increased Frequency of Online Display Advertising," working paper, Northwestern University.
- Khan, Uzma, Ravi Dhar, and Klaus Wertenbroch (2005), "A Behavioral Decision Theory Perspective on Hedonic and Utilitarian Choice," in *Inside Consumption: Frontiers of Research on Consumer Motives, Goals, and Desires*, S. Ratneshwar and David G. Mick, eds. London: Routledge, 144–65.
- Kirmani, Amna (1990), "The Effect of Perceived Advertising Costs on Brand Perceptions," *Journal of Consumer Research*, 17 (2), 160–71.
- Krugman, Herbert E. (1965), "The Impact of Television Advertising: Learning Without Involvement," *Public Opinion Quarterly*, 29 (3), 349–56.
- Lavrakas, Paul J. (2010), "An Evaluation of Methods Used to Access the Effectiveness of Advertising on the Internet," research report, Interactive Advertising Bureau, (May), (accessed February 27, 2014), [available at http://www.iab.net/media/file/Evaluation_of_Internet_Ad_Effectiveness_Research_Methods.pdf].
- Lewis, Randall A. (2010), "Where's the 'Wear-Out?': Online Display Ads and the Impact of Frequency," doctoral dissertation, MIT Department of Economics, Massachusetts Institute of Technology.
- and Justin M. Rao (2013), "On the Near Impossibility of Measuring the Returns to Advertising," working paper.
- , —, and David H. Reiley (2013), "Measuring the Effects of Advertising: The Digital Frontier," NBER Working Paper No. 19520, National Bureau of Economic Research.
- and David H. Reiley (2010), "Does Retail Advertising Work? Measuring the Effects of Advertising on Sales via a Controlled Experiment on Yahoo," working paper, MIT Department of Economics, Massachusetts Institute of Technology.
- Lodish, Leonard M., Magid Abraham, Stuart Kalmenson, Jeanne Livelsberger, Beth Lubetkin, Bruce Richardson, et al. (1995), "How T.V. Advertising Works: A Meta-Analysis of 389 Real World Split Cable T.V. Advertising Experiments," *Journal of Marketing Research*, 32 (May), 125–39.
- Luo, Xueming, Michelle Andrews, Zheng Fang, and Chee Wei Phang (2013), "Mobile Targeting," *Management Science*, (published electronically December 20), [DOI: 10.1287.mnsc.2013.1836].
- MacInnis, Deborah J. and Bernard J. Jaworski (1989), "Information Processing from Advertisements: Toward an Integrative Framework," *Journal of Marketing*, 53 (October), 1–23.
- MacKenzie, Scott B. and Richard A. Spreng (1992), "How Does Motivation Moderate the Impact of Central and Peripheral Processing on Brand Attitudes and Intentions?" *Journal of Consumer Research*, 18 (4), 519–29.
- Mojiva (2012), "Mobile Audience Guide (MAG)," industry report, (May 23).
- Nielsen (2012), "State of the Media: U.S. Digital Consumer Report, Q3-Q4 2011," research report, (accessed February 27, 2014), [available at <http://www.slideshare.net/blueeyepathrec/nielsen-digital-consumer-report-q3q4-2011>].
- and CMO Council (2013), "2013 Online Brand Advertising Outlook Report," research report, (accessed February 27, 2014), [available at <http://www.nielsen.com/us/en/reports/2013/2013-online-advertising-performance-outlook.html>].
- Ovide, Shira and Greg Bensinger (2012), "Mobile Ads: Here's What Works and What Doesn't," *The Wall Street Journal*, (September 27), (accessed February 27, 2014), [available at <http://online.wsj.com/news/articles/SB10000872396390444083304578016373342878556>].
- Pappachen, George and Kara Manatt (2008), "The Mobile Brand Experience: Measuring Advertising Effectiveness on the Mobile Web," *ESOMAR*, (accessed February 27, 2014), [available at http://www.esomar.org/web/research_papers/Advertising-Research_1828_The-Mobile-Brand-Experience.php].

- Patel, Dev, Jeremy Schneider, and Kushan Surana (2013), "Cracking the Mobile Advertising Code," *Harvard Business Review Blog Network*, (September 11), (accessed February 27, 2014), [available at <http://blogs.hbr.org/2013/09/cracking-the-mobile-advertisin/>].
- Petty, Richard E. and John T. Cacioppo (1979), "Issue Involvement Can Increase or Decrease Persuasion By Enhancing Message-Relevant Cognitive Responses," *Journal of Personality and Social Psychology*, 37 (10), 1915–26.
- and ——— (1981), *Attitudes and Persuasion: Classic and Contemporary Approaches*. Dubuque, IA: William C. Brown.
- and ——— (1983), "Central and Peripheral Routes to Persuasion: Application to Advertising," in *Advertising and Consumer Psychology*, Larry Percy and Arch G. Woodside, eds. Lexington, MA: Lexington Books.
- and ——— (1986), *Communication and Persuasion: Central and Peripheral Routes to Attitude Change*. New York: Springer-Verlag.
- , ———, and David Schumann (1983), "Central and Peripheral Routes to Advertising Effectiveness: The Moderating Role of Involvement," *Journal of Consumer Research*, 10 (2), 135–46.
- , Thomas M. Ostrom, and Timothy C. Brock (1981), *Cognitive Responses in Persuasion*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Sahni, Navdeep (2012), "Effect of Temporal Spacing Between Advertising Exposures: Evidence from an Online Field Experiment," working paper, Graduate School of Business, Stanford University.
- Shamdasani, Prem N., Andrea J.S. Stanaland, and Juliana Tan (2001), "Location, Location, Location: Insights for Advertising Placement on the Web," *Journal of Advertising Research*, 41 (4), 7–21.
- Shankar, Venkatesh (2012), "Mobile Marketing Strategy," in *Handbook of Marketing Strategy*, Venkatesh Shankar and Gregory S. Carpenter, eds. Northampton, MA: Edward Elgar Publishing, 217–30.
- and Sridhar Balasubramanian (2009), "Mobile Marketing: A Synthesis and Prognosis," *Journal of Interactive Marketing*, 23 (2), 118–29.
- Shapiro, Stewart, Deborah J. MacInnis, and Susan E. Heckler (1997), "The Effects of Incidental Ad Exposure on the Formation of Consideration Sets," *Journal of Consumer Research*, 24 (1), 94–104.
- Sherif, Muzafer and Carl Hovland (1961), *Social Judgment: Assimilation and Contrast Effects in Communication and Attitude Change*. Oxford, UK: Praeger.
- Smith, Aaron (2013), "Smartphone Ownership—2013 Update," research report, Pew Research Center, (June 5), (accessed February 27, 2014), [available at http://www.pewinternet.org/files/old-media/Files/Reports/2013/PIP_Smartphone_adoption_2013_PDF.pdf].
- Stewart, David W. and David H. Furse (1985), "The Effects of Television Advertising Execution on Recall, Comprehension, and Persuasion," *Psychology and Marketing*, 2 (3), 135–60.
- Strahilevitz, Michal and John G. Myers (1998), "Donations to Charity as Purchase Incentives: How Well They Work May Depend on What You Are Trying to Sell," *Journal of Consumer Research*, 24 (4), 434–46.
- Sun, Baohong and Vicki G. Morwitz (2010), "Stated Intentions and Purchase Behavior: A Unified Model," *International Journal of Research in Marketing*, 27 (4), 356–66.
- Tsang, Melody M., Shu-Chun Ho, and Ting-Peng Liang (2004), "Consumer Attitudes Toward Mobile Advertising: An Empirical Study," *International Journal of Electronic Commerce*, 8 (3), 65–78.
- Vakratsas, Demetrios and Tim Ambler (1999), "How Advertising Works: What Do We Really Know?" *Journal of Marketing*, 63 (January), 26–43.
- Zaichkowsky, Judith L. (1985), "Measuring the Involvement Construct," *Journal of Consumer Research*, 12 (3), 341–52.
- Zoller, Eden and Mark Oliver (2011), "Marketing Perceptions of Mobile Advertising," research report, Ovum Consulting/Interactive Advertising Bureau, (July), (accessed February 27, 2014), [available at http://www.iab.net/media/file/Marketer-Perceptions-of-Mobile-Advertising-Ovum-Report_Final.pdf].

Copyright of Journal of Marketing Research (JMR) is the property of American Marketing Association and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.