Ad Virality and Ad Persuasiveness

December 9, 2011

Abstract

Many video ads are designed to go 'viral', so that the total number of views they receive depends on customers sharing the ads with their friends. This paper explores the relationship between 'earning' this reach and how persuasive the ad is at convincing consumers to purchase or adopt a favorable attitude towards the product. The analysis combines data on the total views of 400 video ads, and crowd-sourced measurement of advertising persuasiveness among 24,000 consumers. We measure persuasiveness by randomly exposing half of these consumers to a video ad and half to a similar placebo video ad, and then surveying their attitudes towards the focal product. Relative ad persuasiveness is on average 10% lower for every one million views an ad achieves. Ads that generated both views *and* online engagement in the form of comments did not suffer from the same negative relationship. We show that such ads retained their efficacy because they attracted views due to humor or visual appeal rather than because they were provocative or outrageous.

Keywords: Viral Advertising, Internet
JEL Codes: L86, M37

1 Introduction

In the past five years there has been a huge shift in digital marketing strategy away from an emphasis from 'paid' media, where a brand pays to advertise, to 'earned' media, where the customers themselves become the channel of delivery (Corcoran, 2009). Reflecting this shift, social video advertising is among the fastest-growing segments in advertising today. In 2010, social video advertising views increased 230%, over nine times the growth in online search and display advertising (Olenski, 2010). These video ads are crucially different from rich-media banner ads. Rather than the advertiser paying for placement, these ads are designed to be transmitted by consumers themselves either through the consumers posting them on their or social media site or sharing them directly with friends. This means that these video ads are commissioned and posted on websites such as `YouTube.com`, in the hope and expectation that consumers themselves will encourage others to watch the video. This is evidently attractive for firms, as it implies a costless means of transmitting advertising. However, in common with other forms of 'earned' media, the return on investment from views obtained in this manner is not clear.

This paper seeks to understand what the relationship is between the 'earning' of media and the persuasiveness of the media. The direction of the relationship is not clear. On the one hand, the very act of sharing a video ad suggests a degree of investment in the product and a liking of the ad that may speak well to its efficacy. On the other hand, advertisers may have to sacrifice elements of ad design in order to encourage people to share the ad that damage its persuasiveness.

We use historical data on the number of times that 400 different video ad campaigns that were posted on `YouTube.com` during 2010 were viewed. This data comes from a media metrics company that tracks major advertiser video ads and records the number of times these ads are viewed. The persuasiveness of these campaigns is then measured using techniques pioneered by media metrics agencies and previously used in data analysis by Goldfarb and Tucker (2011b,a,c). After recruiting 25,000 respondents through crowdsourcing, we measure the effect of exposure to the video ad on purchase intent, using a randomized treatment and control methodology for each campaign. Respondents are either exposed to a focal product video or to a placebo video ad of similar length for another product in our data. They are then asked questions about their purchase intent and brand attitudes towards the focal product.

The randomization induced by the field-test procedure means that econometric analysis is straightforward. First, we document the direction of the relationship between the number of times such an ad was viewed and traditional measures of advertising persuasiveness. We find that ads that achieved more views were less successful at increasing purchase intent. We show that this is robust to different functional forms and alternative definitions of the explanatory and dependent variable such as brand favorability and consideration. It is robust to controls that allow the effect of exposure to vary by video ad length, campaign length, respondent demographics, product awareness and category. It is also robust to excluding respondents who had seen or heard of the ad before, meaning that the results do not simply represent satiation.

We present estimates of the magnitude of this negative relationship and suggest that on average, ads that have received one million more views are 10% less persuasive. Of course, this drop in persuasiveness was compensated for by the increased number of views of the highly viral campaign, so we also present some rough projections to determine the point at which decreased persuasiveness outweighs the increased number of views in terms of the total persuasion exerted over the population. Our estimates suggest that this point occurs between 3-4 million views, a viewership achieved by 6% of campaigns in our data.

The crucial managerial question, though, is whether there are *categories* of ads for whom this negative relationship between virality and persuasiveness did not exist. Such cases would be a clear 'win-win' for advertising managers, where virality does not have to be costly in terms of the persuasiveness of the ad design. We found that viral ads that also induced consumers to comment on the ad, rather than just encouraging them to share it with others, did qualify as 'win-wins.' This has an important managerial implication. Marketing managers, as well as tracking total views for their ads, should also take into account other measures of online engagement such as the creation of user-generated content surrounding the ads. This should be used as an early indicator of successful engagement on the part of the ad, and its likely ability to be persuasive as well as viral.

We present evidence that these ads that did *not* exhibit this negative relationship between total views and persuasiveness were also more likely to be rated as being provocative or outrageous by participants. Instead, they were more likely to be rated as funny or visually appealing. This is in line with an older advertising research literature that has both emphasized that likability (such as produced by humor) is an important determinant of ad appeal (Biel and Bridgwater, 1990; Weinberger and Gulas, 1992; Vakratsas and Ambler, 1999; Eisend, 2009), and that intentional provocativeness or outrageousness is less likely to be effective (Barnes and Dotson, 1990; Vzina and Paul, 1997). Therefore, an explanation for our results is that videos are going 'viral' because they are intentionally provocative or outrageous, but that such ad design does not necessarily make the video ads more persuasive.

This paper contributes to three existing academic literatures.

The first literature is one on virality. Aral and Walker (2011) studies this question in the context of product design. They found that, using evidence from a randomized field trial for an application on Facebook, forcing a product to broadcast a message is more effective than allowing users to post more personalized recommendations at their discretion. There have also been a few studies of campaigns that were explicitly designed to go 'viral.' Toubia et al. (2009) presents evidence that a couponing campaign was more effective when transmitted using a 'viral' strategy on social media than when using more traditional offline methods.

Chen et al. (2011) has shown that such social influence is most important at the beginning of a product's life.

Some recent papers have modeled the determinants of whether or not a video ad-campaign goes 'viral.' This is increasingly important given that 71% of online adults now use video-sharing sites (Moore, 2011). Porter and Golan (2006) emphasize the importance of provocative content (specifically sexuality, humor, violence, and nudity) as a determinant of virality; Brown et al. (2010) echo the importance of comedic violence and argue that the provocative nature of these ads appears to be a key driver. Eckler and Bolls (2011) emphasize the importance of a positive emotional tone for virality. Outside of the video-ad sphere, Chiu et al. (2007) emphasized that hedonic messages are more likely to be shared by e-mail; Berger and Milkman (2011) emphasize that online news content is more likely to be shared if it evokes high or negative arousal as opposed to deactivating emotions such as sadness. Elberse et al. (2011) examined 12 months of data on popular trailers for movies and video games. They found evidence that their popularity was often driven by their daily advertising budget. Teixeira (2011) examines what drives people to share videos online and distinguishes between social utility and content utility in non-altruistic sharing behavior. Though these papers provide important empirical evidence about the drivers of virality, these papers did not actually measure how persuasive the video ads were and how this related to virality.

The second literature is on the persuasiveness of online advertising. Much of this literature has not considered the kind of advertising that is designed to be shared, instead focusing on non-interactive banner campaigns[1]. Generally, this literature has only considered the persuasiveness of video-advertising tangentially or as part of a larger study. For example, Goldfarb and Tucker (2011b) presented a result that video advertising is less persuasive when placed in a context which matched too closely the product being advertised. In the arena of video advertising, Teixeira et al. (2011) have shown that video ads that elicit

[1] Among many others Manchanda et al. (2006); Lambrecht and Tucker (2011); Tucker (2011, 2012)

joy or surprise are more likely to retain visual focus (as measured by eye-tracking) and are less likely to be fast-forwarded through. We believe, however, that this is the first study on the relationship between ad virality and ad persuasiveness, that is, how the ability of an ad to endogenously gain 'reach' is related to the ability of the ad to persuade.

Finally the paper contributes to an older managerial literature that argues that the internet has reduced the tradeoff between richness and reach in information delivery in the internet era (Evans and Wurster, 2000). For example, before the internet, firms had to choose between personal selling, which is an incredibly rich form of marketing communications but that has limited reach since there are no scale economies, and media like television advertising which achieves impressive reach but is not a rich form of marketing communications. They argue that the easy replication and personalization facilitated by the internet reduced this tradeoff. This paper suggests, however, that advertisers who try to achieve scale on the internet through the actions of internet users rather than their own efforts may still face tradeoffs in terms of the persuasiveness of ads that users can be persuaded to share.

2 Data

2.1 Video Viewing Data

We obtained data from a large video metrics company, Visible Measures. Data for movie campaigns provided by this company has also been used by Elberse et al. (2011) to study the effects of offline advertising budgets on video virality for movie previews. Visible Measures is an independent third-party media measurement firm for online video advertisers and publishers founded in 2005. It is the market leader in terms of tracking views and engagement for different types of social video ads. Visible Measures shared data with us for recent campaigns in the consumer goods category from 2010. We requested explicitly that they exclude from the data video ads for categories of products such as cars and other large ticket items, for which the majority of people were unlikely to be in the market. We also requested that they exclude video ads for entertainment products such as movies, video games, and DVDs whose ads have a short shelf-life. 29 percent of videos were for consumer packaged goods, 14 percent of videos were for electronics, 13 percent of videos were for apparel and 8 percent were for fast food. We allow persuasiveness to vary by these different 'product' categories as controls in subsequent robustness checks.

The videos of 396 of these campaigns were still live and online and consequently were included in this survey. Table 1a reports the campaign-level summary statistics that we received from Visible Measures. Since Visible Measures is primarily employed as a media measurement company, it does not have data on the design costs or the creative process that lay behind the ad it is tracking.

'Total views' captures the number of times these videos had been viewed by consumers. This encompasses both views of the original video as placed by the ad agency, and views that were generated by copies of the ad and derivatives of the ad. It is clear from the standard deviation that there is a high variance in the number of total views across the ad campaigns, which is one of the reasons that we use a logged measure in our regressions. We also show the robustness of our results to a raw linear measure. We use 'total views' as a proxy measure of virality, that is the number of times in total the ad was shared. This reflects a view that views of social video ads on pages such as `youtube.com` are gained by an organic process where people find such videos on blogs or social media sites and then share the video ad further with their friends. However, this process could be subject to manipulation[2], so we present robustness checks using both controls for firm interference in the process of achieving virality and an alternative measure of virality based on inter-day correlation of views in daily ad-viewing data. 'Total Comments' records the number of times that these videos had received a written comment from a consumer, typically posted below the ad on websites such as `Youtube.com`.

[2]See Wilbur and Zhu (2009) for a general discussion of manipulation of online ads

	Mean	Std Dev	Min	Max
Total Views	777996.53	2705048.25	57	37761711
Total Comments	1058.54	4382.75	0	64704
Length Ad (sec)	56.24	33.31	10	120
Observations	396			

(a) Campaign level

	Mean	Std Dev	Min	Max
Exposed	0.50	0.50	0	1
Purchase Intent	0.59	0.49	0	1
Intent Scale	3.63	1.12	1	5
Would Consider	0.60	0.49	0	1
Consideration Scale	3.67	1.10	1	5
Favorable Opinion Scale	3.75	0.99	1	5
Favorable Opinion	0.62	0.49	0	1
Aware of Product (Unexposed)	0.56	0.50	0	1
Age	29.57	9.44	18	65
Male	0.70	0.46	0	1
Income (000,USD)	35.53	24.22	20	100
Weekly Internet Hours	26.23	10.93	1	35
Lifetime tasks	6.18	33.68	0	251
Observations	24367			

(b) Survey participants

	Mean	Std Dev	Min	Max
Funny Rating	5.64	0.97	2	8
Provocative Rating	5.27	0.66	1	8
Outrageous Rating	5.13	0.74	1	8
Visual Appeal Rating	6.74	0.66	1	9
Observations	396			

(c) Campaign ratings from survey participants

Table 1: Summary Statistics

We wanted to gather data on advertising persuasiveness. Of course, a simple regression that correlated firm's sales and the virality of its ad campaigns is unlikely to be useful, since the decision to launch a viral ad campaign is confounded with many other factors. Direct measurement of consumer response rates for online video ads is also difficult. Typical 'direct response' methods of evaluating digital advertising, such as measuring click-throughs, are not appropriate. Many videos do not have embedded hyperlinks, and also many products that are advertised in the videos such as deodorant are not primarily sold online. As documented by Porter and Golan (2006); Golan and Zaidner (2008), viral advertising very rarely has a clear 'call to action', such as visiting a website, that is measurable. Therefore, we test advertising persuasiveness based on industry-standard techniques for measuring the persuasiveness of brand campaigns online. These techniques, developed by among others Dynamic Logic and Insight Express, combine a randomized control and exposure methodology with surveys on brand attitudes. Both major advertisers and major agencies use these same techniques for evaluating both banner campaigns and video campaigns. Therefore a conservative interpretation of our measure of ad persuasiveness is that it is a traditional metric for ad persuasiveness used by major advertisers.

Since such ad persuasiveness measures were not used as the campaigns were being rolled out, we have to collect this data retrospectively. Given the number of campaigns we want to evaluate, this requires a large number of participants. We obtain this large number using crowdsourcing techniques. Specifically, we recruited 25,000 separate individuals using the crowdsourcing platform Mechanical Turk. Similar crowdsourcing techniques have been used by Ghose et al. (2011) to design rankings for search results. Each of these participants visited a website that had been designed to resemble popular video sharing websites such as `Youtube.com`. The main difference between the study website and a traditional video-sharing website is that participants had no choice but to watch the video and that after watching the video, participants were asked to answer a series of questions concerning their brand attitudes.

For each campaign, we recruited on average 60 respondents. Half of the respondents are allocated to a condition where they are exposed to the focal video ad that we have virality data on. The other half of respondents (the control group) see a placebo video ad for another unrelated (random) product that was also part of our data. We also randomized the placebo ad shown among our control group to make sure that the choice of placebo ad did not influence our result.[3]

The randomization between whether someone saw the focal video ad or another, means that in expectation all our respondents are identical. Therefore we can causally attribute any differences in their subsequent attitudes towards the product to whether they were exposed to the video ad or not.

We record whether or not the respondent watches the video all the way through and exclude those who did not from our data. We also exclude participants who, despite the controls in place, managed to take the survey multiple times. This explains why we have 24,367 respondents, which is fewer than the original 25,000 respondents we recruited. We then ask participants a series of survey questions. Table 1b summarizes these responses. These include questions about their purchase intent towards the focal product and likelihood of consideration of the focal product. We also included decoy questions about another brand. All these questions are asked on a 5-point scale in line with traditional advertising persuasiveness questioning (Morwitz et al., 2007). Following Goldfarb and Tucker (2011b), we converted this 5-point scale to a binary purchase intent measure that captures whether someone is very likely or likely to purchase the product for our main analysis. As seen in Table 1b, average purchase intent was relatively high, reflecting the 'everyday' nature of the products in the ads. However, we show robustness to the use of the full scale in subsequent regressions.

[3]This could have occurred if the advertising was directly combative (Chen et al., 2009)

We acknowledge that our focus on purchase intent means that we focus on the effect of advertising at the later stages of the purchase funnel or traditional purchase decision process (Vakratsas and Ambler, 1999). Our methodological approach which necessitates forced exposure makes it hard for us to think about 'awareness' or other earlier stages of customer attitudes earlier in the purchase process. We do not have time-stamps for when consumers completed different parts of the survey, but in general the time they took to complete the task was only minutes more than it took to watch the video. This means that we are not able to collect measures on ad memorability. We do, however, control for heterogeneity in product awareness in subsequent regressions.

Survey responses are weaker measures of advertising persuasiveness than purchasing or profitability (as used by Reiley and Lewis (2009)), because though users may say they will purchase, they ultimately may not actually do so. However, as long as there is a positive correlation between whether someone intends to purchase a product and whether they actually do so, the directionality of our results should hold. Such a positive correlation between stated purchase intent and purchase outcomes has been broadly established (Bemmaor, 1995; Morwitz et al., 2007). However, a conservative view would be that our results reflect how total views is related to an established and widely-used measure of advertising persuasiveness that is used as an input when making advertising allocation decisions.

In addition to asking about purchase intent, the survey also asked participants about whether or not they recalled having seen the focal video ad before or had heard it discussed by their friends and media. We use this information in a robustness check to make sure that the fact that respondents are more likely to have seen viral videos before and that there may be less of an effect the second time around is not driving our results. We also asked participants to rate the video on a 10-point sliding scale based on the extent to which they found it humorous, visually appealing, provocative or outrageous. We then use the median ratings for the campaign when trying to explain whether ads with different characteristics display different relationships between virality and persuasiveness. Table 1c reports these ratings at the campaign level, based on the median response of our survey-takers.

The survey also asked respondents about their gender, income, age, and the number of hours they spent on the internet. These descriptives are reported in Table 1b. They are used as controls in the regression, though since respondent allocation to exposed and control group was random, they mainly serve to improve efficiency. However, they do serve also as a check on how representative our survey-takers were. It is clear that respondents are more male than the general population, are younger, earn less, and also spend more time online. The fact that there were more males than females reflects video-sharing site usage. Based on a survey conducted by Moore (2011), men are 28% more likely than women to have used a video-sharing site recently. However, we accept that since these participants were recruited via a crowdsourcing website, there is the possibility that they may differ in unobserved ways from the population.

The issue of how representative such respondents' answers are is faced by all research using survey-based evaluation techniques, as discussed in Goldfarb and Tucker (2011c). However, what is crucial is that there is no *a priori* reason to think that the kinds of ads that these participants would be favorably impressed by would differ from the more general video-sharing population, even if the magnitudes of their responses may differ. We also show that the magnitudes of the effects that we measure do match well to existing estimates of video-advertising efficacy that have been collected in less artificial settings.

3 Empirical Analysis

Our randomized procedure for collecting data makes our empirical analysis relatively straightforward.

For person i who was allocated to the testing cell for video ad for product j, their purchase intent reflects:

$$Intent_{ij} = I(\alpha Exposed_{ij} + \beta Exposed_{ij} \times LoggedViews^j + \theta X_{ij} + \delta_j + \epsilon_j > 0) \quad (1)$$

Therefore, α captures the main effect of being exposed to a video ad on purchase intent. Purchase intent is a binary variable for whether the respondent said they were likely or very likely to purchase the product. β captures the core coefficient of interest for the paper - whether exposure is more or less effective if the ad has proven to be viral; X_{ij} is a vector of controls for gender, age, income, and time online; δ^j is a series of 397 consumer good product level fixed effects that control for heterogeneity in baseline purchase intent for that product and includes the main effect of Ad Views ($LoggedViews_i^j$), which is why this lower-order interaction is not included in our specification. We use a logged measure of ad views, because we do not want our results to be biased by extreme values given the large variance in distribution of ad views. However, we show robustness to other measures subsequently. In our initial regressions, we assume that the ϵ_j is normally distributed, implying a probit specification, though we subsequently show robustness to other functional forms. Standard errors are clustered at the product level in accordance with the simulation results presented by Bertrand et al. (2004). This represents a conservative empirical approach, as in our setting we have randomization at the respondent level as well.

Table 2 shows our initial results that investigate the relationship between ad persuasiveness and virality where virality is measured by total views of the video. Column (1) reports an initial specification where we simply measure the main effect of *Exposed* on purchase intent. As expected, being exposed to the video ad has a positive and significant effect on the participant's purchase intent for that product.

Table 2: More viewed ads are less persuasive

	Probit	Probit	Probit	Probit	Probit	Probit	OLS	OLS
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
		$\geq$ Med Total View	$<$ Med Total View					
Exposed × Logged Views				-0.0153**	-0.0153**	-0.0164**	-0.00658**	
				(0.00738)	(0.00747)	(0.00752)	(0.00268)	
Exposed × Total Views (m)								-0.0153**
								(0.00722)
Exposed	0.181***	0.150***	0.212***	0.246***	0.250***	0.259***	0.0951***	0.0738***
	(0.0177)	(0.0259)	(0.0239)	(0.0363)	(0.0368)	(0.0370)	(0.0134)	(0.00691)
Age					-0.00316***			
					(0.000965)			
Income (000,USD)					0.00116***			
					(0.000342)			
Weekly Internet Hours					-0.0000646			
					(0.000797)			
Male					0.310***	0.247***	0.0893***	0.0893***
					(0.0199)	(0.0208)	(0.00747)	(0.00689)
Product Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Demo Controls	No	No	No	No	No	Yes	Yes	Yes
Observations	24367	12221	12146	24367	24367	24367	24367	24367
Log-Likelihood	-15353.9	-7531.9	-7820.3	-15351.7	-15193.8	-14896.6	-15687.3	-15688.4
R-Squared							0.121	0.121

Dependent variable is binary indicator for whether or not participant states that they are likely or very likely to purchase the product. Probit estimates Columns (1)-(6) and OLS estimates (7)-(8). * $p<0.10$, ** $p<0.05$,*** $p<0.01$. Robust standard errors clustered at the product level.

The estimate in Column (1) suggests that exposure to a video ad increases purchase probability by 6.6 percentage points, which is a similar to the average effect of exposure to 'in-stream' video ads reported by Goldfarb and Tucker (2011b). This is reassuring because that research used industry-sponsored data where survey-takers were people who had naturally come across the ad in the process of their web-browsing. This suggests that the recruitment method and forced exposure did not overly influence our measure.

Column (2) reruns this simple regression for the websites that had a below-median number of views. Column (3) reports results for the same regression for websites that have an above-median number of views. It is clear that on average the effect of exposure to the ad on purchase intent is greatest for video ads that have a below-median number of views. This is our first evidence that there may be a negative relationship between the virality of the ad and its persuasiveness at persuading a viewer to purchase the product.

To test this more robustly, Column (4) provides an explicit test of the apparent difference in size in the coefficients for *Exposed* in Column (2) and (3) by reporting the results of a basic version of (1). The key variable of interest, $Exposed_{ij} \times LoggedViews_i^j$, is negative and significant. This suggests that exposure to an ad which has received more views is less likely to be able to persuade an ad viewer to purchase the product.

This finding remains unchanged when we add linear controls for consumer characteristics in Column (5) which is as expected due to randomization. These linear controls suggest that richer, younger males who do more tasks are more likely in general to say they will purchase. Column (6) uses an alternative non-parametric set of controls for consumer characteristics which are simply indicators for six levels of income, age and internet usage. As can be seen in the log-likelihood, this non-parametric approach to controls is more efficient, which is why we use it for the rest of the specifications. In each case the use of such controls is indicated by a 'Yes' in the Demo Controls row at the bottom of the table.

An econometric concern is the interpretation of the main interaction terms. Research by

Ai and Norton (2003) suggests that the interaction in a non-linear model may not capture the true cross-derivative. In order to ensure that our results are not a function of the nonlinearity of the estimation function, we also show in Column (7) that a linear probability model gives qualitatively similar results, providing reassurance that the non-linear functional form does not drive our results. In Column (8), we show that the result is also robust if we use a linearized form of our key explanatory variable 'Total Views' rather than the logged form. The r-squared in each of these columns is relatively low, but this is very much in line with previous studies in this area such as Aral and Walker (2011).

To give an idea of the magnitude of these estimates, we used a probit model and the appropriate Ai and Norton (2003) correction to calculate different predicted values at different numbers of total (non-logged) views. Figure 1a presents the results. There is a sizeable loss of persuasiveness for ads that received a larger number of views, and it suggests that roughly for around every 1 million views an ad receives, it is on average 10% less persuasive.

However, this is not the whole story, as of course by definition the most viral videos had improved reach, meaning that they while they were less persuasive for any individual viewer, they also potentially were able to persuade more people. To take account of this, we did a rough simulation where we took account of the total 'expected' persuasion from a video ad. This is defined as '$Reach \times$ *Persuasiveness*' and reflects how persuasive the ad was multiplied by how many consumers it was viewed by. Figure 1b plots these rough estimates. Our simulation suggests that there are eventually decreasing returns to achieving virality overall, at 3-4 million total views. At this point the reduction in ad persuasiveness due to virality is large enough that incrementally more consumers viewing the ad achieves little. Only 6% of videos in our data achieved this level of virality, so our plot suggests that negative returns to virality are limited. We want to emphasize that Figure 1b is a very rough calculation. However, the existence of inverse-U-shaped returns from achieving virality in advertising is a new finding and one that deserves managerial attention.

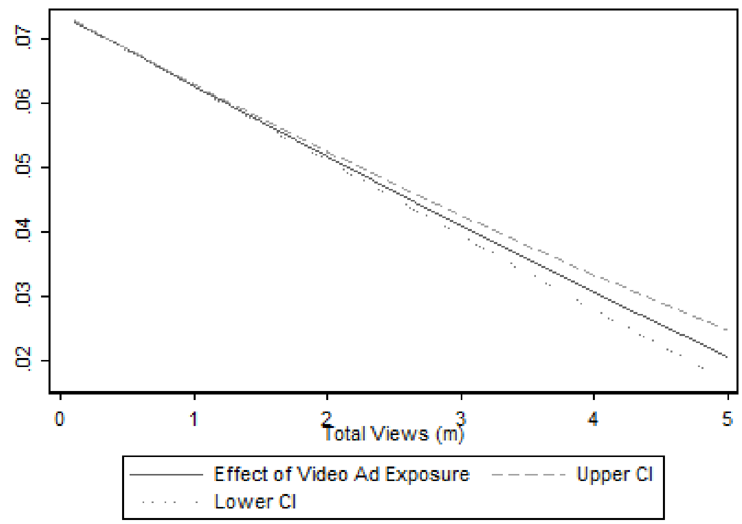

(a) Linear Prediction of Relationship

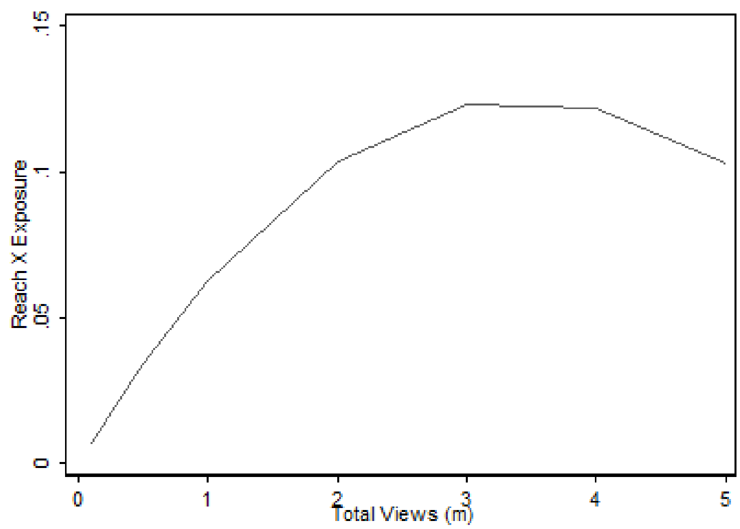

(b) Aggregate Effects

Figure 1: Predictions from Probit Model

3.1 Robustness

In section 4 we present evidence that some of the characteristics that make ads viral also make them less persuasive, and this can explain the empirical relationship. We explore alternative explanations to this characteristics-based explanation in this section.

3.2 Alternative definitions of explanatory variables

Table 3: Different definitions of explanatory variable

	Not Seen (1)	Adj Placements (2)	Adj Deriv. (3)	Virality (4)
Exposed × Logged Views	-0.0182**			
	(0.00780)			
Exposed × Placement Adjusted Views		-0.0231**		
		(0.0115)		
Exposed × Logged Non-Derivative Views			-0.0186**	
			(0.00758)	
Exposed × Daily Views Correlation				-0.126**
				(0.0613)
Exposed	0.228***	0.247***	0.267***	0.250***
	(0.0383)	(0.0337)	(0.0366)	(0.0349)
Daily Views Correlation				1.057
				(3.704)
Product Controls	Yes	Yes	Yes	Yes
Demo Controls	Yes	Yes	Yes	Yes
Observations	22298	24367	24367	24367
Log-Likelihood	-13655.7	-14896.7	-14895.9	-14996.9

In Column (1) all respondents who had seen or heard of the ad before are excluded. Probit estimates. Dependent variable is binary indicator for whether or not participant states that they are likely or very likely to purchase the product. Robust standard errors clustered at the product level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

One natural concern given our use of historical data is that our results may be biased because a general prior awareness of a campaign or its success may influence respondents' answers to questions about advertising persuasiveness. This would provide an alternative explanation of our findings, that the reason that more viral video ads are less effective is because the respondents have already been influenced by them, and repeated exposure is less effective (Tellis, 1988). We address this in Column (1) of Table 3 as we exclude our crowd-sourced field testers who stated they had seen or heard of the advertising campaign before. Our results are robust to excluding such observations. This suggests that the explanation of the measured negative relationship is not wearout among the general population.

In Column (2) we address another natural concern, which is that the number of placements (that is the number of websites) that the video was posted on drove the result. As discussed by Cruz and Fill (2008), the process whereby an ad agency determines the number of placements, commonly known as the number of 'seeds,' is highly strategic. Therefore an alternative interpretation of the measured negative relationship would simply be that videos with multiple placements got more views but the multiple placements themselves were in response to acknowledged ad ineffectiveness. Visible Measures collects data on the number of websites that each ad was placed on, though data on which websites these were. We use this data to create a measure of average views per placement. When we control for placements by using a measure of the average number of views per placement, the result holds.

Column (3) addresses the concern that our result could be an artifact of the fact that total views includes views of derivatives of the original ad. There is the possibility that if an ad were poorly executed, it could have invited scorn in the form of multiple parodic derivatives that could have artificially inflated the number of views. However, the robustness check shows that our results remain robust to excluding views that can be attributed to parodies.

Column (4) addresses the concern that total views is not an adequate measure of virality. In particular, there is a concern that the 'total views' measure may not truly capture a viral process whereby people share the video ad through their blogs or social media with their friends and acquaintances. Instead it could capture firm actions, for example if a firm has a popular website that has a link to the `youtube.com` url on it. Generally, virality is used to define a process whereby an ad is shared by people successively. To capture this we use a new measure of virality which is simply the inter-day correlation in views for that particular campaign. The idea is quite simply that ads whose views were the result of a successive sharing process are more likely to have daily views that are positively correlated. We want to emphasize that of course that this correlation is unlikely to be causal and highly likely to be biased upwards by as we have no exogenous shifter that allows us to actually identify causal network effects (Tucker, 2008; Ryan and Tucker, 2012). With this caveat, our results are similar when we use this alternative proxy which is likely to be related to virality.

3.3 Potential confounds

We then go on to explore whether other factors may potentially be confounding our results in Table 4.

One concern is that our results may simply be being driven by differences between the product category that the ads were advertising. For example, more aspirational or hedonic categories of products may receive more views (Chiu et al., 2007; Berger and Milkman, 2011), but also be less easy to persuade people to purchase via advertising. Column (1) of Table 4 addresses the concern and shows that the results are robust to our allowing the persuasiveness of the ad to vary by the category of product (for example, whether it is food or a personal care item). The results remain robust to the addition of these interactions between category-specific indicators and the indicator for exposure which would capture any differences in advertisers' potential ability to persuade respondents for that category.

Column (2) addresses the concern that the results are driven by differences in ad length. For example, it could be more likely that longer video ads are more persuasive but less likely to be viewed. To control for this, we included an interaction between exposure and ad length. Our results are robust to the inclusion of this control. They also suggest, interestingly, that ad length appears to have little relationship with the perceived persuasiveness of the ad.

Column (3) addresses the concern that the results are driven by differences in campaign length. For example, it could be more likely that longer campaigns gathered more views, but that the kind of products that tended to have long campaigns (perhaps those that were more traditional and less-fast paced) found it more difficult to persuade people to purchase the product. To control for this, we included an interaction between exposure and the number of days the campaign ran according to Visible Measures data. Our results are robust to the inclusion of this control. They also suggest, interestingly, that on average longer campaigns are more persuasive, which makes sense as it is more likely that ineffective campaigns would be withdrawn.

Column (4) addresses the concern that our results could be an artifact of the fact that workers may have different levels of experience with Mechanical Turk, and that perhaps its overly-sophisticated users were more likely to exhibit 'demand effects' and try and answer the questions in the way they thought that the questioner wanted, and that this might be driving the results if randomization failed. To control for this possibility, we allow our results to vary by the workers' number of previous tasks for other firms on Mechanical Turk. The results are again similar.

Column (5) addresses the concern that the result could be an artifact of the variation in ages of our survey-takers. For example, if video-ads are targeted at young people, and young people are more likely to share ads that 'older' people would disapprove or react poorly to, then this could explain our result. However, when we interact our main effect with a variable for 'age' then there is no change in our estimates, suggesting that age is not a moderating factor.

Table 4: Exploring different explanations

	Cat Int (1)	Ad Length (2)	Campaign Length (3)	Tasks (4)	Age (5)	Awareness (6)
Exposed × Logged Views	-0.0151**	-0.0159**	-0.0197**	-0.0160**	-0.0213***	-0.0159**
	(0.00758)	(0.00779)	(0.00775)	(0.00748)	(0.00789)	(0.00811)
Exposed × Ad Length		-0.000148				
		(0.000507)				
Exposed × Campaign Length			0.000142*			
			(0.0000753)			
Exposed	0.210***	0.265***	0.228***	0.258***	0.281***	0.275***
	(0.0590)	(0.0424)	(0.0409)	(0.0368)	(0.0398)	(0.0383)
Exposed × Age × Logged Views					0.0199	
					(0.0148)	
Exposed × Age					-0.0916	
					(0.0730)	
Age × Logged Views					-0.0159	
					(0.0121)	
Exposed × High Aware × Logged Views						0.0270
						(0.0240)
Lifetime Tasks				0.00322***		
				(0.000899)		
Exposed × Lifetime Tasks				0.000124		
				(0.00120)		
Logged Views × Lifetime Tasks				-0.000104		
				(0.000178)		
Exposed × Logged Views × Lifetime Tasks				-0.0000255		
				(0.000251)		
Exposed × High Aware						-0.275**
						(0.138)
Category Interactions	Yes	No	No	No	No	No
Product Controls	Yes	Yes	Yes	Yes	Yes	Yes
Demo Controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	24367	24367	24367	24367	24367	24367
Log-Likelihood	-14893.5	-14896.6	-14894.9	-14868.9	-14895.7	-14892.5

Probit estimates. Dependent variable is binary indicator for whether or not participant states that they are likely or very likely to purchase the product. Robust standard errors clustered at the product level. * p <0.10, ** p <0.05,*** p <0.01.

Another concern is that potentially the ads could be designed primarily to promote awareness for new products. If the ads which were most viral were also for the most 'new' products that were harder to persuade consumers to purchase, this could explain our results. To test this, we tried added an extra interaction with an indicator for whether the product had an above-average level of awareness as recorded among consumers who were not exposed to the ad. Column (6) reports the results. The interaction $Exposed_{ij} \times HighAwareness \times LoggedViews_i^j$ is insignificant, suggesting that awareness is not an important mediator of the effect we study.

3.4 Alternative definitions of dependent variables

In Table 5, we check the robustness of our results to alternative dependent variables. Columns (1) show robustness to using the entire purchase intent scale. In this OLS specification, the direction of the main effect of interest remains the same, which is to be expected given that the binary indicator for purchase intent was based on this scale.

Column (2) shows robustness to looking at an alternative measure of brand persuasiveness which is whether or not the consumer would consider the brand. This is an important check as most video advertising is explicitly brand advertising without a clear call to action. Therefore, it makes sense to see that our result applies to an earlier stage in the purchase process (Hauser, 1990). However, the results remain robust (both in significance and approximate magnitude) to a measure which attempts to capture inclusion in a consideration set. This suggests that the documented negative relationship holds across attempts to influence customer attitudes across different stages of the purchase cycle. In a similar spirit, Column (3) shows that are results to using as a dependent variable whether or not the respondent had a 'favorable' or 'very favorable' opinion of the brand.

Table 5: Checking robustness to different dependent variables

	OLS	Probit	Probit
	(1)	(2)	(3)
	Intent Scale	Would Consider	Favorable Opinion
Exposed × Logged Views	-0.00829**	-0.0145**	-0.0167**
	(0.00411)	(0.00737)	(0.00744)
Exposed	0.115***	0.274***	0.311***
	(0.0204)	(0.0359)	(0.0361)
Product Controls	Yes	Yes	Yes
Demo Controls	Yes	Yes	Yes
Observations	24367	24367	24367
Log-Likelihood	-25792.5	-14712.0	-14463.4
R-Squared	0.107		

OLS estimates in Column (1). Probit estimates Columns (2)-(3). Dependent variable is the full five-point purchase intent scale in Column (1). Dependent variable is whether or not the customer is likely or very likely to 'consider' purchasing the product in Column (2). Dependent variable is whether or not the customer is likely or very likely to have a 'favorable' opinion towards the product in Column (3). Robust standard errors clustered at the product level. * $p < 0.10$, ** $p < 0.05$,*** $p < 0.01$.

4 When is there no negative relationship?

So far, this paper has documented there is a negative relationship between the total views that ads achieve and their persuasiveness. However, of crucial interest to managers is when there is no such negative relationship, or what factors mitigate it. Therefore, one central aim of this research is to offer some practical guidance as to occasions when ads can both attract multiple views and be persuasive when inducing purchase intent.

4.1 Engagement

We do this by introducing an explicit measure of online engagement to our regressions. This is the 'total comments' that an ad receives. Total comments are 'user-generated content'. This is distinct from more general forms of online reputation systems (Dellarocas, 2003), and has been shown by Ghose and Han (2011); Ghose and Ipeirotis (2011) to correlate with product success. Moe and Schweidel (2011) have also show that comment ratings themselves may be subject to cascades and herding.

Figure A1 displays how comments usually appear below the ad on a video-sharing website. Of course, total comments are positively linked to the total number of views an ad receives, since without viewers there can be no comments, but it is conceptually distinct as well as requiring a different investment from the viewer. This definition of engagement, which is used by Visible Measures when promoting their provision of information on comments to video ads, is conceptually distinct from the kind of physical engagement measured by Teixeira et al. (2011) using eye-tracker technology.

Table 6: What mediates this negative relationship?

	Probit (1) Purchase Intent	Probit:Not Seen (2) Purchase Intent	Probit (3) Purchase Intent	OLS (4) Purchase Intent	OlS (5) Intent Scale	Probit (6) Would Consider	Probit (7) Favorable Opinion
Exposed × Logged Views	-0.0379***	-0.0397***		-0.0145***	-0.0202***	-0.0327***	-0.0416***
	(0.0143)	(0.0132)		(0.00510)	(0.00774)	(0.0126)	(0.0127)
Exposed × Logged Comments	0.0281**	0.0282**		0.0103**	0.0156**	0.0238*	0.0326**
	(0.0142)	(0.0139)		(0.00504)	(0.00765)	(0.0133)	(0.0133)
Exposed × Placement Adjusted Views			-0.0488***				
			(0.0178)				
Exposed × Placement Adjusted Comments			0.0375**				
			(0.0163)				
Exposed	0.420***	0.389***	0.473***	0.154***	0.205***	0.411***	0.498***
	(0.0931)	(0.0881)	(0.108)	(0.0334)	(0.0504)	(0.0844)	(0.0847)
Product Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Demo Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	24367	22298	24367	24367	24367	24367	24367
Log-Likelihood	-14894.3	-13653.6	-14893.7	-15684.8	-25790.0	-14710.4	-14460.4
R-Squared				0.121	0.107		

In Column (2) all respondents who had seen or heard of the ad before are excluded. Dependent variable is whether someone is likely or very likely to purchase the product in Columns (1)-(4). Dependent variable is the 5-point purchase intent scale in Column (5). Dependent variable is whether someone is likely or very likely to consider the product in Column (6). Dependent variable is whether someone is likely or very likely to have a favorable opinion of the product in Column (7). Robust standard errors clustered at the product level. // * $p < 0.10$, ** $p < 0.05$,*** $p < 0.01$.

Table 6 explores what occurs when we include this measure of engagement into our regressions. In Column (1), we show what happens when we add $Exposed_{ij} \times LoggedComments_i^j$ to our regression. The pattern for $Exposed_{ij} \times LoggedViews_i^j$ is similar if more precise than before. However, crucially, $Exposed_{ij} \times LoggedComments_i^j$ is both positive and significant. This suggests that video ads that are successful at provoking users to comment on them and engage with them directly are also the ads that are more successful at persuading consumers to purchase the product.

Column (2)-(6) show robustness to the various concerns explored in our earlier robustness checks. Again the result is robust to correcting for the potential for ad satiation (Column (2)), different definitions of the explanatory variable (Column (3)), different functional forms (Column (4)), and different definitions of the dependent variable (Columns (5)-(7))).

We then go on to explore what underlying ad characteristics drive this relationship between the effect of total views, ad persuasiveness and engagement. Table 7 indicates the ad characteristics that are linked both with high views and with this desirable high ratio between comments and views. It is clear that the ads that are both more likely to attract a large number of total views but less likely to attract a high ratio of comments to views are the ones that are intentionally provocative or outrageous in their ad design. On the other hand, the ads which are visually appealing and funny appear successful at eliciting more comments relative to views and, though successful at attracting more views, are less like to attract views than those that are provocative or outrageous.

4.2 Ad characteristics

To explore this further, we ran regressions where we looked at how ad persuasiveness varied with the total views that can be explained by the survey-takers' ratings of different ad characteristics. For each of the separate ad characteristics, we calculated the 'predicted total views' that can be attributed to variation in that characteristic for the campaign using

Table 7: Correlation of ad characteristics with total views and comments ratio

	Total Views	Total Comments:Total Views Ratio
Outrageous Rating	0.103***	-0.0191**
Provocative Rating	0.110***	-0.0381***
Funny Rating	0.0734***	0.0131*
Visual Appeal Rating	0.0384***	0.0203**

Raw Correlations shown between various Ad Characteristic Ratings and Total Views in Column (1) and the ratio of Total comments: Total Views in Column (2). * $p < 0.10$, ** $p < 0.05$,*** $p < 0.01$.

an ordinary least squared regression of total views on that characteristic. Table 8 presents the results.

We want to emphasize that we do not intend to estimate a simultaneous equations model where variables can be excluded from the first-stage regression and causality is ascribed. Ad characteristics jointly explain both ad persuasiveness and total views. Our aim is to explore this joint determination by recording a statistical relationship that results from the fact that both total views and persuasiveness can be explained by fundamental ad characteristics in the data.

Table 8: Linkage of ad characteristics with ad persuasiveness and total views

	Probit (1) Purchase Intent	OLS (2) Purchase Intent	OLS (3) Intent Scale	Probit (4) Would Consider	Probit (5) Favorable Opinion
Exposed × Views (Predicted Visual)	31.93***	10.25***	19.16***	32.17***	30.80***
	(4.341)	(1.418)	(2.292)	(4.187)	(4.519)
Exposed × Views (Predicted Funny)	10.28***	2.330**	3.329**	10.14***	11.51***
	(3.136)	(1.024)	(1.601)	(3.222)	(3.286)
Exposed × Views (Predicted Outrageous)	-5.018**	-1.707**	-2.688**	-5.840**	-7.358***
	(2.198)	(0.687)	(1.154)	(2.286)	(2.256)
Exposed × Views (Predicted Provocative)	-6.888**	-2.763***	-4.595***	-8.563***	-9.997***
	(2.814)	(0.907)	(1.448)	(2.843)	(2.930)
Exposed	-13.95***	-3.728***	-7.034***	-12.80***	-11.39***
	(1.845)	(0.543)	(0.852)	(1.753)	(1.718)
Product Controls	Yes	Yes	Yes	Yes	Yes
Demo Controls	Yes	Yes	Yes	Yes	Yes
Observations	23673	23673	23673	23669	23669
Log-Likelihood	-13096.3	-13843.8	-23584.7	-12991.8	-12606.4
R-Squared		0.219	0.213		

Probit estimates in Columns (1) and (4). OLS estimates in Columns (2) and (3). Dependent variable in Columns (1)-(2) is binary indicator for whether or not participant states that they are likely or very likely to purchase the product. Dependent variable in Column (3) is the full five-point purchase intent scale. Dependent variable in Column (4) is whether or not the customer is likely or very likely to 'consider' purchasing the product. Dependent variable in Column (5) is whether someone is likely or very likely to have a favorable opinion of the product. Explanatory variables are predicted total views based on linear regression on average 'rating' of characteristic for that ad on total views. Robust standard errors clustered at the product level. * $p < 0.10$, ** $p < 0.05$,*** $p < 0.01$.

Column (1) presents initial estimates for a probit model. Echoing Table 7, it suggests that if we look only at the variation in total views that can be explained by humor or visual appeal, then this is positively related to ad persuasiveness. On the other hand, variation in total views that can be attributed to the outrageous or provocative nature of the ad is actually less likely to be linked to persuasive advertising. Since there are obvious objections to putting a predicted value from a linear regression into a non-linear functional form (Wooldridge, 2000), we repeat our estimation with a linear probability model which does not raise these issues. Column (2) reports the results and shows similar results. Column (3)-(5) shows that our results are also robust to alternative definitions of the dependent variable either as the full-scale variable for purchase intent or as 'purchase consideration' and 'brand favorability'. Columns (4) and (5) have similar estimates if we use a linear probability model.

Both Tables 7 and 8 provides evidence about why there may be the measured negative relationship exists between advertising virality and advertising persuasiveness. Some video ads are purposely being designed to be outrageous or provocative in order to incite consumers to share the video with their friends (Porter and Golan, 2006; Brown et al., 2010; Moore, 2011). However, on average, they are neither provoking responses among viewers to the actual ad itself nor succeeding in persuading users to purchase the product. This is in line with existing research (Vzina and Paul, 1997). In other words, being outrageous is a reliable strategy for encouraging virality, but it reduces the persuasiveness of ads. On the other hand, ad characteristics such as humor appear to be successful at both promoting user response to the ad as well as virality. Again, this is in line with behavioral research into humor in ads which suggests that on average it does not harm the advertising message and can sometimes enhance it by increasing engagement (Weinberger and Gulas, 1992).

5 Implications

Firms online are increasingly switching their emphasis from 'paid media' such as online display advertising, to 'earned media' where consumers themselves transmit the message. This has been reflected in the growth of social video advertising, where video ads are now designed to go viral and achieve costless reach. This is a very different distribution system for advertising, compared to a typical placement process where an advertising manager simply decides on how many exposures they want and on what medium to purchase them. Instead, with viral advertising the advertising manager is responsible for designing ads that will generate their own exposures.

The aim of this paper is to quantify the empirical relationship in social advertising between ads that earn multiple views and ads that are persuasive. Combining historical data and a randomized treatment and control methodology among a large crowdsourced population of survey-takers, we measure this relationship empirically. We find evidence that there is a significant negative relationship between total ad views and ad persuasiveness. The ads that receive the most views are also the ones that are relatively less able to persuade consumers to purchase the product. We present evidence that after adjusting for the improved reach (that is, the larger number of people who view the ads) of ads that achieve many views, this negative relationship between views and persuasiveness only leads to negative consequences after an ad reaches 3-4 million views. We check the robustness of our results in a variety of ways.

We then provide some evidence about why this occurs. Videos that receive more comments alongside their views were more likely to be persuasive. In other words, ads that are successful not just at provoking consumers to share the ad with others but also to take time to respond to the ad itself appear more successful. The ads that do worst in terms of their comments to views ratio are ads that are viral by virtue of their being rated as outrageous or provocative. When we examine variation in total views that can be explained by ad characteristics, it is only the variation in total views that can be attributed to outrageousness and provocativeness that has this negative correlation with ad persuasiveness. The variation in total views that can be explained by humor or visual appeal is positively related to ad persuasiveness. Therefore, though provocative ad design is sufficient to induce participants to share an ad, it has a negative effect on the persuasiveness of the ad. On the other hand, ads that are viral by virtue of their humor or their visual design appear to have a positive relationship between their persuasiveness and how many times the ad was viewed.

There are of course limitations to this study. First, despite the extensive data collection, these results hold for 400 ad campaigns for the consumer goods category from 2010. It is not clear whether the results would hold for other products or across time. Second, the participants that we recruited may not be representative of the population. This is likely to mean that our estimates are not representative. However, unless this group responds very differently to different ads from the rest of the population, then our general conclusions should hold. Third, all ad design and consequently virality is exogenous to the study and was not explicitly manipulated. Last, since we study video ads for well-known consumer goods, we do not study the effects of viral video ads on product awareness. Notwithstanding these limitations, this study does document the potential for an empirical negative relationship between earned reach and ad persuasiveness for ad managers who are trying to exploit the new medium of video advertising.

References

Ai, C. and E. C. Norton (2003, July). Interaction terms in logit and probit models. *Economic Letters 80*(1), 123–129.

Aral, S. and D. Walker (September 2011). Creating social contagion through viral product design: A randomized trial of peer influence in networks. *Management Science 57*(9), 1623–1639.

Barnes, J. H. and M. J. Dotson (1990). An exploratory investigation into the nature of offensive television advertising. *Journal of Advertising 19*(3), 61–69.

Bemmaor, A. C. (1995). Predicting behavior from intention-to-buy measures: The parametric case. *Journal of Marketing Research 32*(2), 176–191.

Berger, J. and K. L. Milkman (2011). What makes online content viral? *Journal of Marketing Research, Forthcoming*.

Bertrand, M., E. Duflo, and S. Mullainathan (2004). How much should we trust differences-in-differences estimates? *The Quarterly Journal of Economics 119*(1), 249–275.

Biel, A. L. and C. A. Bridgwater (1990). Attributes of likable television commercials. *Journal of Advertising Research 30*(3), 38 – 44.

Brown, M. R., R. K. Bhadury, and N. K. Pope (2010). The impact of comedic violence on viral advertising effectivensss. *Journal of Advertising 39*(1), 49–65.

Chen, Y., Y. V. Joshi, J. S. Raju, and Z. J. Zhang (January/February 2009). A theory of combative advertising. *Marketing Science 28*(1), 1–19.

Chen, Y., Q. Wang, and J. Xie (2011). Online social interactions: A natural experiment on word of mouth versus observational learning. *Journal of Marketing Research 48*(2), 238–254.

Chiu, H.-C., Y.-C. Hsieh, Y.-H. Kao, and M. Lee (2007). The determinants of email receivers' disseminating behaviors on the internet. *Journal of Advertising Research 47*(4), 524.

Corcoran, S. (2009, December 16). Defining earned, owned and paid media. *Forrester Research*.

Cruz, D. and C. Fill (2008). Evaluating viral marketing: isolating the key criteria. *Marketing Intelligence & Planning 26*(7), 743 – 758.

Dellarocas, C. (2003). The digitization of word of mouth: Promise and challenges of online feedback mechanisms. *Management Science 49*(10), 1407–1424.

Eckler, P. and P. Bolls (2011). Spreading the virus: Emotional tone of viral advertising and its effect on forwarding intentions and attitudes. *Journal of Interactive Advertising 11*(2), 1–11.

Eisend, M. (2009). A meta-analysis of humor in advertising. *Journal of the Academy of Marketing Science 37*, 191–203. 10.1007/s11747-008-0096-y.

Elberse, A., C. Lee, and L. Zhang (2011). Viral videos: The dynamics of online video advertising campaigns. *Mimeo, HBS*.

Evans, P. and T. S. Wurster (2000). *Blown to bits: How the new economics of information transforms strategy*. Boston: Harvard Business School Press.

Ghose, A. and S. P. Han (September 2011). An empirical analysis of user content generation and usage behavior on the mobile internet. *Management Science 57*(9), 1671–1691.

Ghose, A. and P. G. Ipeirotis (2011). Estimating the helpfulness and economic impact of product reviews: Mining text and reviewer characteristics. *IEEE Trans. Knowl. Data Eng. 23*(10), 1498–1512.

Ghose, A., P. G. Ipeirotis, and B. Li (2011). Designing ranking systems for hotels on travel search engines by mining user-generated and crowd-sourced content. *Mimeo, NYU*.

Golan, G. J. and L. Zaidner (2008). Creative strategies in viral advertising: An application of taylor's six-segment message strategy wheel. *Journal of Computer-Mediated Communication 13*(4), 959–972.

Goldfarb, A. and C. Tucker (2011a, April). Advertising Bans and the Substitutability of Online and Offline Advertising. *Journal of Marketing Research*.

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

Goldfarb, A. and C. E. Tucker (2011c, January). Privacy regulation and online advertising. *Management Science 57*(1), 57–71.

Hauser, J. R. B. W. (1990, March). An evaluation cost model of consideration sets. *Journal of Consumer Research 16*, 393–408.

Lambrecht, A. and C. Tucker (2011). When does retargeting work? Timing information specificity. *Mimeo, LBS*.

Manchanda, P., J.-P. Dube, K. Y. Goh, and P. K. Chintagunta (2006). The effect of banner advertising on internet purchasing. *Journal of Marketing Research 43*(1), 98 – 108.

Moe, W. W. and D. A. Schweidel (2011). Online product opinions: Incidence, evaluation, and evolution. *Forthcoming, Marketing Science*.

Moore, K. (2011, July). 71 percent of online adults now use video-sharing sites. *Pew Internet Project*.

Morwitz, V. G., J. H. Steckel, and A. Gupta (2007). When do purchase intentions predict sales? *International Journal of Forecasting 23*(3), 347–364.

Olenski, S. (2010, December 16). Social video advertising: The explosion of 2010. *Social Media Today*.

Porter, L. and G. J. Golan (2006). From subservient chickens to brawny men: A comparison of viral advertising to television advertising. *Journal of Interactive Advertising 6*(2).

Reiley, D. and R. Lewis (2009). Retail advertising works! Measuring the effects of advertising on sales via a controlled experiment on Yahoo!". Working Paper, Yahoo! Research.

Ryan, S. and C. Tucker (2012). Heterogeneity and the dynamics of technology adoption. *Quantitative Marketing and Economics 10*(1).

Teixeira, T., M. Wedel, and R. Peters (2011). Emotion-induced engagement in internet video ads. *Journal of Marketing Research, Forthcoming*.

Teixeira, T. S. (2011). Video ads virality. *Mimeo, HBS*.

Tellis, G. (1988, May). Advertising exposure, loyalty, and brand purchase: A two-stage model of choice. *Journal of Marketing Research 25*(2), 134–144.

Toubia, O., A. T. Stephen, and A. Freud (2009). Viral Marketing: A Large-Scale Field Experiment. *Mimeo, Insead*.

Tucker, C. (2008). Identifying formal and informal influence in technology adoption with network externalities. *Management Science 54*(12), 2024–2038.

Tucker, C. (2011). Social networks, personalized advertising, and privacy controls. *Mimeo, MIT*.

Tucker, C. (2012). Social advertising. *Mimeo, MIT*.

Vakratsas, D. and T. Ambler (1999). How Advertising Works: What Do We Really Know? *The Journal of Marketing 63*(1).

Vzina, R. and O. Paul (1997). Provocation in advertising: A conceptualization and an empirical assessment. *International Journal of Research in Marketing 14*(2), 177 – 192.

Weinberger, M. G. and C. S. Gulas (1992). The impact of humor in advertising: A review. *Journal of Advertising 21*(4), pp. 35–59.

Wilbur, K. C. and Y. Zhu (March/April 2009). Click fraud. *Marketing Science 28*(2), 293–308.

Wooldridge (2000). *Econometric Analysis of Cross Section and Panel Data*. MIT Press.

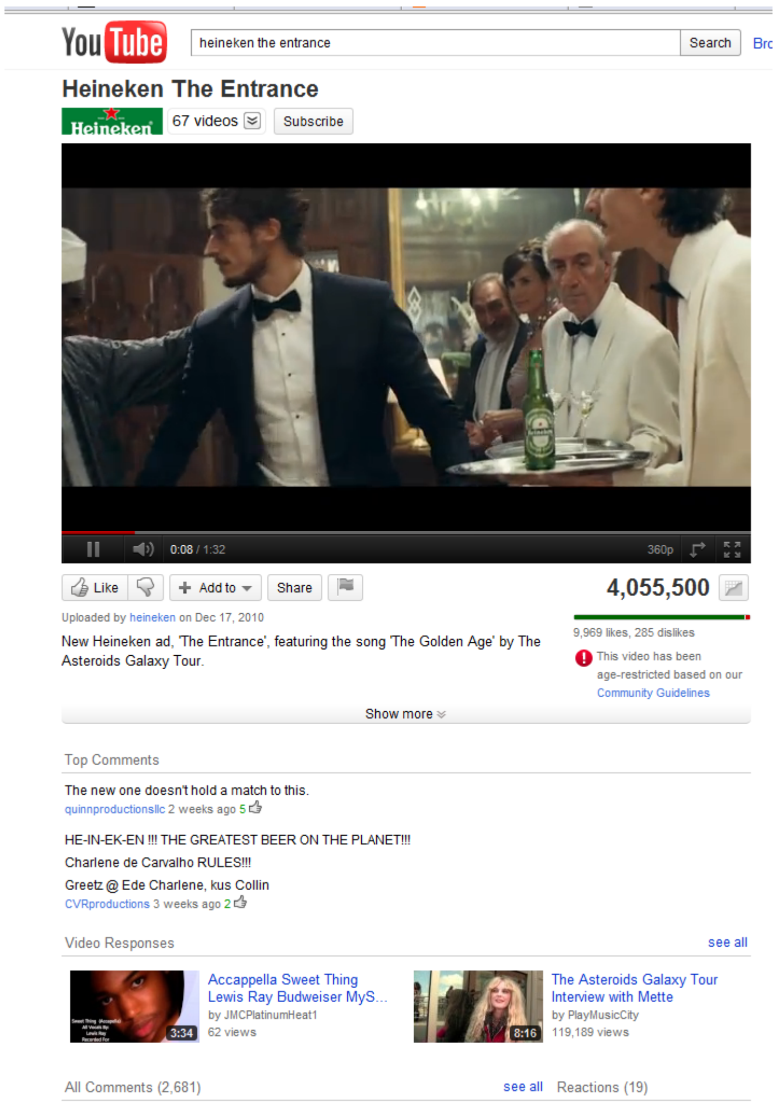

Figure A1: Screen Shot from Typical Video Ad Campaign Showing Comments and Total Views