# Why Do Some Advertisements Get Shared More than Others? Quantifying Facial Expressions To Gain New Insights

#### DANIEL MCDUFF

Microsoft Research damcduff@microsoft.com

#### JONAH BERGER

The Wharton School of the University of Pennsylvania jberger@wharton.upenn. edu Why do some advertisements get shared more than others? Using scalable automated facial coding algorithms, the authors quantified the facial expressions of thousands of individuals in response to hundreds of video advertisements. Results suggest that not all emotions increase sharing and that the relationship between emotion and transmission is more complex than mere valence alone. Facial actions linked to positive emotions, (*e.g.*, smiles) were associated with increased sharing. But while some actions associated with negative emotion (*e.g.*, lip depressor associated with sadness) were linked to decreased sharing, others (*e.g.*, nose wrinkles) associated with disgust were linked to increased sharing.

#### INTRODUCTION

The value of earned media is undeniable. Rather than having to pay per impression, word of mouth and social shares are free, so more consumers sharing an advertisement or message means greater impact at lower costs. Consequently, across industries, and both online and offline, companies are working to harness the power of social sharing.

The challenge, however, is getting consumers to share. Some content goes viral, getting millions of shares, whereas other content barely gets shared at all.

Why are some advertisements shared more than others? And how, by understanding the answer to this question, how can companies design advertisements that are more likely to be shared?

The authors examined this question in the context of online video advertisements; specifically, the link between consumer facial expressions and sharing. The face plays a key role in emotional expression (Ekman and Rosenberg, 1997). People who are happy, for example, have markedly different facial expression from those who are sad. Consequently, facial analysis provides an unobtrusive way of measuring emotional reactions that can be performed quickly, easily, and cheaply at scale. The authors designed a custom facial action detection algorithm and demonstrated that it can reliably predict emotional expression. Using this algorithm, the authors quantified the facial expressions of thousands of individuals in response to hundreds of video advertisements and tested which emotional expressions are most linked to sharing.

This research makes four key contributions. First, it helps managers boost sharing. By understanding why people share, content creators can design content that is more likely to be shared.

# Management Slant

- Not all emotions increase sharing, and the relationship between emotion and transmission is more complex than mere valence alone.
- Arousal seems to drive the relationship between emotions and sharing so that high-arousal emotions increase sharing and low-arousal emotions decrease it.
- Facial expressions provide a valuable tool to predict and understand consumer behavior and can be measured in a scalable manner using an Internet-based framework.

Submitted April 14, 2020; revised June 10, 2020; accepted July 31, 2020.

.....

Second, it deepens understanding around drivers of sharing. The little work that has been done on facial expressions (Teixeira, Picard, and el Kaliouby, 2014; Teixeira, Wedel, and Pieters, 2012) has examined only one or two emotions at a time and has almost exclusively examined positive emotions. Thus, researchers know relatively little about whether negative emotions might increase the sharing of advertisements or inhibit it. By simultaneously investigating multiple emotions, both positive and negative, researchers can more precisely identify the effect of each on what people share.

Third, most previous approaches to facial coding have required participants to come into a laboratory and sit in front of a computer loaded with special hardware and software. This is somewhat restrictive, unnatural, and costly, and it limits the breadth of data that can be collected. In contrast, the authors utilized an approach that requires no specialized hardware or software, just a webcam.

Finally, this approach is highly scalable. Studies of human behavior and emotion such as this one are almost always limited because of the prohibitive cost of manual video coding by expert human coders. Automated coding, however, enables collection of a large amount of data from participants across a range of countries in a quick, cost-effective manner. Given that expressions of emotions in response to mundane media content can be sparse and that large interpersonal variability exists in nonverbal behaviors, it is important to sample a broad number of people. Performing this type of analysis at scale is nontrivial and the current work shows that these automated techniques are valid methods for scientific research.

#### LITERATURE REVIEW

#### **Social Sharing**

Researchers have become increasingly interested in what drives people to talk and share (Berger, 2014). Some work suggests that impression management shapes what gets passed along (Packard and Wooten, 2013). People care about self-presentation, or how they look to others, so the better something makes them look, the more likely they are to share it. Consistent with this notion, people are more likely to share their consumption experiences if they go well, rather than badly (Deangelis, Bonezzi, Peluso, Rucker, and Costabile, 2012), and news articles (Berger and Milkman, 2012) or brands (Berger and Schwartz, 2011) that are more interesting.

Other work suggests that emotion may increase sharing. When people experience emotions, they often turn to others to regulate those emotions and help them make sense of what they are feeling (Rimé, 2009). Talking with others can help people understand what they feel and why (Rimé, Mesquita, Philippot, and Boca, 1991) and can help them relive positive experiences and receive People care about self-presentation, or how they look to others, so the better something makes them look, the more likely they are to share it. Consistent with this notion, people are more likely to share their consumption experiences if they go well, rather than badly

social support around negative ones. Consistent with this perspective, people report greater willingness to share urban legends that they feel evoke more emotion (Heath, Bell, and Sternberg, 2001), and news articles that contain more emotional words are more likely to make the "most e-mailed" list (Berger and Milkman, 2012).

Properly measuring emotion, however, is challenging. First, although one can ask someone how they feel, self-reports are often inaccurate. People do not always have the best insight into their own emotional states. Someone may report feeling negative, for example, without knowing whether they feel angry or anxious. Because they rely on cognitive elaborations of experienced emotions, self-reports also have trouble picking up quick moment-to-moment emotional shifts over time. The mere act of introspection can alter emotional experience (Kassam and Mendes, 2013). Reporting or having to report how one is feeling can affect the physiological state and thus bias responses.

Second, although automated methods exist to extract emotion from text, this is less feasible with video. Recent advances in natural language processing and sentiment scoring have allowed researchers and practitioners to estimate the valence and emotionality of text (Berger, Humphreys, Ludwig, Moe, *et al.*, 2020). The Linguistic Inquiry and Word Count application (Pennebaker, Mehl, and Niederhoffer, 2003), for example, estimates positivity and negativity by counting the number of positive and negative words in each document.

But while advances in computer vision have begun to allow automated recognition of images, but inferring likely emotional reactions to advertisements is far more challenging. Systems may be able to recognize that an advertisement contains a dog, for instance, that alone is not enough information to determine what emotions viewers will feel when watching that advertisement. Stimulus analysis assumes that all individuals will respond in the same way to an advertisement and does not account for the individual differences in how an advertisement will be appraised.

#### **Facial Responses**

To address these measurement challenges, the authors used facial responses. The face plays a key role in emotional expression (Ekman and Rosenberg, 1997) and can be an outward manifestation or signal of these often otherwise internal states. Decades of quantitative research have revealed reliable patterns in the ways in which emotions are expressed on the face. The Facial Action Coding System (FACS), for example, has found consistent facial behaviors associated with anger, fear, joy, surprise, pain, and deceit (Ekman and Rosenberg, 1997). Of the signals that communicate affective information—physiology, nonverbal behaviors, and brain activity—the face is one of the more easily interpreted.

As a result, facial analysis provides an unobtrusive method of passively measuring in-the-moment behavior with a set of basic emotional states linked to prototypic expressions (Ekman, Friesen, O'Sullivan, Chan, *et al.*, 1987). When experiencing joy, for example, people tend to make facial expressions that differ from those of people experiencing disgust. These different expressions involve different movements of individual facial muscles, which can be coded as specific action units (AUs). Consistent with the value of this approach, marketing scientists have recently started to use facial coding as a measurement tool (Xiao and Ding, 2014). Researchers have examined the relationship between smiling and purchase intent (Teixeira *et al.*, 2014), for example, and the link between positive emotion and engagement (Teixeira *et al.*, 2012).

Examining facial expressions avoids some common challenges in studying expressions of emotion. Rather than relying on selfreport, recording facial responses to advertisements allows an unobtrusive measurement of expressed emotion response that is not disturbed by cognitive elaboration or reflection. It is difficult for an individual to introspect on their own expressions of emotion and report these in real time. Rather than simply relying on aggregate measures, such as one advertisement being more positive than another, coding facial responses allows researchers to examine variation in individual responses-that is, whether one person reacts more positively than another to the same advertisement-and whether such variation is linked to sharing. Finally, the advent of automated systems that can accurately detect facial muscle movements (De la Torre and Cohn, 2011) allows for increased scalability, repeatability, and observation in naturalistic environments. Rather than examining a small set of laboratory stimuli, researchers can investigate how thousands of people naturally react to hundreds of advertisements and code

these reactions in a consistent, reliable manner. Rather than studying a single emotion in isolation, as in most previous work, this approach allows researchers to examine how multiple emotions, experienced at various points in watching an advertisement, might impact sharing.

#### THE CURRENT RESEARCH

The authors measured a range of facial actions representing both positive and negative emotions to examine whether facial actions are linked to sharing and which actions are most linked to sharing.

The simplest way that emotion might relate to sharing is that any emotion increases transmission. As noted previously, people often share emotional experiences with others. Psychological research on the social sharing of emotion argues that 90 percent of emotional experiences are passed on (Rimé, 2009) to others. Consistent with the notion that emotionality increases sharing, movies and news articles that evoke more emotion are more likely to be shared (Berger and Milkman, 2012; Luminet, Bouts, Delie, Manstead, and Rimé, 2000). If all emotions increase sharing, then videos that evoke any facial expressions should be shared more.

A second possibility is that emotional valence drives sharing. Emotions such as happiness are perceived as more positive, whereas others such as anger are perceived as more negative. Given that impression management shapes what people share, one could argue that people might avoid sharing negative things to avoid communicating negative identities (Berger, 2014; Tesser and Rosen 1975). People prefer interacting with positive rather than negative others (Kamins, Folkes, and Perner, 1997), so consumers may share positive things to avoid seeming like a negative person or a "Debbie Downer." Similarly, most people would prefer to put others in a good mood rather than a bad one. Research finds that positive news is more likely to be shared than negative news (Berger and Milkman, 2012; Tesser and Rosen, 1975). Overall, these perspectives suggest that content that evokes facial expressions linked to positive, rather than negative, emotions may be shared more.

The authors hypothesized that a third, more complex, possibility was more likely; specifically that different specific emotions have different effects on sharing. In addition to valence, specific emotions differ on various other dimensions such as arousal or certainty (Lerner and Keltner, 2000; Smith and Ellsworth, 1985). This approach suggests that even though both disgust and sadness are negative emotions, they may have different effects on whether something is shared.

To test these possibilities, the authors used a new data collection framework to examine the naturalistic facial responses -i.e.,

# Each participant was shown a set of 10

# of the 230 advertisements; video order

### within that set was randomized

viewers in their home environment-of thousands of people watching hundreds of videos. Computer algorithms (Senechal, McDuff, and el Kaliouby, 2015) built specifically for this task automatically coded facial responses. These algorithms leverage significant advances in computer science, specifically in the field of machine learning, which have enabled the accurate measurement of subtle facial expression in situ. Given the very challenging nature of coding AUs, computer algorithms are still not capable of accurately coding all actions. The authors therefore selected five AUs that are typically associated with expression of different emotions and for which the algorithms had a high degree of accuracy, precision, and recall. The authors examined whether facial expressions reliably predict sharing and, if so, which of five facial actions are most predictive. The current study is the largest of its kind using an online framework and the first to investigate the link between emotional responses to video advertisements and sharing.

The authors examined both positive and negative emotions, but one may wonder whether advertisements would ever aim to elicit negative emotions. After all, if the goal is to increase consumer evaluations or encourage purchases, why would a brand ever want to associate themselves with negativity?

A few points are worth noting. First, some advertisements seem to purposely evoke negative emotions. Advertisements for nonprofits or related to corporate social responsibility, for example, are often classified as heartfelt tearjerkers. Similarly, paper towel advertisements may show a disgusting mess to demonstrate how well the paper towels can clean it up. Although positive emotions certainly seem more prevalent, negative emotions do exist in advertisements, and the authors examined their link to sharing.

Second, without intending to, some advertisements may induce negative emotion. One person might find an advertisement funny, whereas another finds it disgusting or possibly offensive. Someone might find an advertisement clever, whereas another finds it confusing. Thus, examining individual heterogeneity in facial responses provides valuable insight.

Third, cultures vary in "display rules," or when and where it is acceptable to express certain emotions (Matsumoto, 1990; Tsai and Chentsova-Dutton, 2003). In some cultures, expressing strong emotion is supported, but in others, it is seen as inappropriate and discouraged (Matsumoto, Yoo, Fontaine, Anguas-Wong, *et al.*, 2008). Consequently, one might expect that responses to advertisements would also vary across cultures. An exploratory analysis examines this possibility.

#### METHOD

## Materials

The authors recorded emotional reactions to 230 video advertisements from a variety of product categories (*e.g.*, instant foods, confectionary, pet care, and beauty products). The videos were between 20 and 120 seconds in length (mean duration = 27.3 seconds; SD =8.65 seconds) and from five different countries: Germany (70 advertisements), the United States (60 advertisements), France (40 advertisements), the United Kingdom (40 advertisements), and China (20 advertisements). Participants viewed advertisements from their own country in their native language. The advertisements were all recent (aired in the past 10 years) and from major brands.

#### **Participants**

Participants (N = 2,106; mean age = 33.6 years; 51 percent male) were recruited from an online market research panel. Most participants completed the survey from home. The participants were compensated with the equivalent of \$8.10 in their local currency. This sample size is an order of magnitude larger than most studies of nonverbal responses to advertisements. Given the variability in nonverbal behavior between individuals, it is important to consider large populations.

#### Procedure

Figure 1 summarizes the study's online data collection framework. Participants were contacted via e-mail. They were told that they were taking part in a study to evaluate video advertisements. Participants simply clicked on a link and opted in through a browser-based survey. They were asked for consent to use their webcam to record while they took part in the study. Participants only needed an Internet connection and a webcam to take part. There was no requirement for specialized hardware or to download or install software. Consequently, their experiences while taking the survey were similar to watching online content during everyday life.

Each participant was shown a set of 10 of the 230 advertisements; video order within that set was randomized. After watching a given advertisement, they indicated their willingness to share ("If you watched this advertisement on a website such as YouTube, how likely would you be to share it with someone else?") on a scale ranging from 1 (very unlikely) to 5 (very likely).



**Figure 1** The Study Used an Online Platform to Record Facial Responses to Advertisements Note: fos = frames per second

Participants completed the study from home, and the average time to finish the survey, watch all 10 advertisements, and complete the questions was 26.3 minutes.

#### **Automated Facial Coding**

Automatically detecting spontaneous facial actions in everyday environments is challenging. Designing automated facial coding algorithms is difficult, and it is not possible to detect every single facial AU with the requisite performance in everyday settings because of the subtlety and variability in how they appear in videos. The authors therefore designed custom facial action detection algorithms focusing on five of the most commonly occurring, informative, and reliably detected actions: smiles, outer eyebrow raises, brow furrows, lip corner depressors, and nose wrinkles (See Figure 2 for examples). The selection of these actions in particular was based on the data collected; these actions were found to occur most frequently and capture a range of different expression responses to the content. Although anxiety is a common negative emotion more generally, it is not a common reaction to advertisements.

Smiles are defined as contractions of the *zygomaticus major* muscle, which pulls the lip corners toward the ears (AU 12;

**374** JOURNAL OF ADUERTISING RESEARCH December 2020

for details on smiles and other facial expressions, see Ekman, Friesen, and Hager, 2002). AU 12 is often associated with positive affect. Outer brow raises are defined as contractions of the frontalis (pars lateralis) muscle, which pulls the eyebrows upward (AU 2). AU 2 is often associated with surprise. Brow furrows are defined as contractions of the corrugator supercilii muscle, which pulls the eyebrows down and together to form vertical wrinkles on the inner brow (AU 4). Researchers commonly interpret AU 4 (brow knitting) as a signal of mental effort (Oster, 2017); confusion, worry, and concentration are specific affective states associated with it (Rozin and Cohen, 2003). Lip corner depressors are defined as contractions of the depressor anguli oris muscle, which pulls the lip corners down (AU 15). AU 15 is often associated with sadness. Nose wrinkles are defined as contractions of the levator labii superioris alaeque nasi muscle, which pulls the eyebrows down and the nose corners upward (AU 9). AU 9 is often associated with disgust (Kassam, 2010).

The automated software is designed to detect these facial actions. The software has two principal constituent components, both of which are created using supervised learning (McDuff, 2014). Supervised learning is a machine learning approach that involves training a mathematical model *via* a set of labeled examples. The first



**Figure 2** Posed and Spontaneous Examples of Facial Actions and the Emotions Commonly Associated with Them Note: AU = action unit, from the Facial Action Coding System

component involves face tracking: It identifies landmarks on the face-specific locations around the eyes, mouth, and nose-using a computer vision technique known as supervised descent. Once the face is located within each video frame, a region of interest is identified including the mouth, nose, eyes, and eyebrows. The second component analyzes how the texture of the face region of interest changes to identify which actions, such as a smile, are present; this component is implemented using computer vision algorithms known as support vector machines. The output is a probability score for each action. When no expression is present, these probabilities are all 0, and when one or more actions fire, the corresponding action probabilities will rise. The software computes these probabilities for every frame of video (14 times per second). As people may exhibit an expression for a short time and then return to a neutral expression, it will capture these dynamics with a resolution of approximately 70 milliseconds, thus avoiding problems with dialbased methods of moment-to-moment measurement where participants may forget to turn the dial for several seconds.

The algorithms also leverage an estimate of each subject's neutral face. A moving time window is used to baseline estimates on the basis of the temporal dynamics of normal facial responses.

Automatic detectors were trained using example videos manually coded by expert human observers. Twenty coders were given training from the FACS manual (Ekman *et al.*, 2002) and were certified by passing the FACS test. They were shown similar participant videos collected from a previous study and asked to code the presence of each of the five facial actions as defined by the FACS in each frame of the video. Agreement among expert human coders regarding the presence of a given action was high (framelevel  $\kappa$ s = 0.74–0.92) and similar to agreement between humans (frame-level  $\kappa$ s = 0.55–0.84).

Using these expert-coded data, the detectors were then trained using supervised learning. The manually coded videos were partitioned into training (80,000 labeled video frames), validation (10,000 labeled video frames), and testing (900,000 labeled video frames) sets. The training set was a random sample of 4,000 unique individuals, and the validation set was a random sample of 2,500 unique individuals. The training and validation videos were collected during previous market research studies. Using the training set, the authors trained a two-class support vector machine with a Nystrom-approximated radial basis function kernel (Senechal et al., 2015). Using the validation data, the authors optimized the number of samples used in the Nystrom approximation ( $N_s = \{200,$ 500, 1,000, 2,000}), the support vector machine cost parameter (C =  $\{0.01, 0.1, \dots 100\}$ ), and the radial basis function spread parameter  $(\gamma = \{0.01, 0.1, ..., 100\})$ . More detail about implementation and validation of this approach is reported by Senechal et al. (2015). The most commonly observed expressions were smiles (5.95 percent), followed by lip depressors (3.38 percent), eyebrow furrows (3.31 percent), eyebrow raises (1.80 percent), and nose wrinkles

(0.45 percent). Although these numbers might seem low, they are similar to normal levels of responses to everyday television content and similar to the base rates observed in other studies of facial responses to television advertisements (McDuff, Girard, and el Kaliouby, 2016a; Teixeira *et al.*, 2014).

#### Model

The authors used a linear mixed-effects model to examine the relationship between facial actions and sharing. This model captures the effect of facial actions on sharing, accounting for the fact that some advertisements might be more likely to be shared or that people from certain countries might have a higher propensity to share.

$$Sharing = \beta_0 + \beta_1 Smiles + \beta_2 Eyebrow Raise + \beta_3 Eyebrow Furrow + \beta_4 Lip Depressor + \beta_5 Nose Wrinkle + Z_1 Subject + Z_2 Ad + Z_3 Country + \varepsilon$$
(1)

Here,  $\beta_0$  is an intercept, and  $\beta_1$ ,  $\beta_2$ ,  $\beta_3$ ,  $\beta_4$ , and  $\beta_5$  are the parameters that estimate the marginal linear effects of smiles, eyebrow raises, eyebrow furrows, lip depressors, and nose wrinkles, respectively, on sharing. The score for each of the actions was calculated by first thresholding the classifier output, at the threshold determined in the classifier validation process, and then calculating the fraction of frames in which the output was above the threshold (this can be described as the action "base rate").  $Z_{1/2} Z_{2/2}$  and  $Z_{2}$  are parameters describing the variance in sharing that can be explained by the differences among subjects, advertisements, and countries, respectively. E is an error term. Modeling subject, advertisement, and country as random effects means that the authors were not interested in the specific effect of any one subject, advertisement, or country but rather wanted to account for the overall variability they exert on sharing. This allowed the authors to control for content-related factors unrelated to facial expressions that might impact sharing. Including interactions between facial actions did not improve model fit, so these were left out for the sake of simplicity.

#### RESULTS

Results indicate that smiles (AU 12) were positively and most strongly associated with sharing ( $\beta$  = 1.45, *SE* = 0.03, *p* < 0.001; See Figure 3). A 30 percent increase in smiling is associated with a 10 percent increase in willingness to share.

Some negative emotions seemed to decrease sharing: Lip depressor (AU 15, associated with sadness) and brow furrow (AU 4, associated with confusion) were both negatively associated with sharing ( $\beta = -0.17$ , *SE* = 0.05, *p* < 0.01; and  $\beta = -0.18$ , *SE* = 0.05, *p* < 0.01).

Other negative emotions, however, seemed to increase sharing: Nose wrinkles (AU 9, often associated with disgust) were positively associated with sharing ( $\beta = 0.22$ , *SE* = 0.11, *p* < 0.05).



**Figure 3** Link between Emotional Expressions and Sharing Notes: Examples of facial actions and emotions commonly associated with them are shown. Standard errors are given in brackets. AU = action unit. \*p < 0.05; \*\*p < 0.001.

#### **Effects Over Time**

One may wonder whether experiencing certain emotions early versus late would have different impacts. To test this, the authors also built a time-varying model. Advertisements vary in length and structure; as a result, modeling time as a continuous variable, even with a large amount of observations, is challenging. Most advertisements, however, do have some high-level structure. The simplest model to describe this would be a three-act composition: setup (beginning), confrontation (middle), and resolution (end). The authors therefore segmented them into thirds (beginning, middle, and end) and measured the presence of facial action variables in each segment. The authors estimated the following model:

Sharing = 
$$\beta_0 + \beta_{11}Smiles + \beta_{12}(Smiles \times time)$$
  
+ $\beta_{21}EyebrowRaise + \beta_{22}(EyebrowRaise \times time)$   
+ $\beta_{31}EyebrowFurrow + \beta_{32}(EyebrowFurrow \times time)$   
+ $\beta_{41}LipDepressor + \beta_{42}(LipDepressor \times time)$   
+ $\beta_{51}NoseWrinkle + \beta_{52}(NoseWrinkle \times time)$   
+ $Z_1Subject + Z_2Ad + Z_3Country + \varepsilon$  (2)

Here,  $\beta_0$  is an intercept, and  $\beta_{11'}$ ,  $\beta_{21'}$ ,  $\beta_{31'}$ ,  $\beta_{41'}$  and  $\beta_{51}$  are the parameters that estimate the marginal linear effects of smiles, eyebrow raises, eyebrow furrows, lip depressors, and nose wrinkles on sharing.  $\beta_{12'}$ ,  $\beta_{22'}$ ,  $\beta_{32'}$ ,  $\beta_{42'}$  and  $\beta_{52}$  are the parameters that estimate the time-dependent effect of each action on sharing. As in the previous model, the authors treated subject, advertisement, and country as random effects. Time is treated as a discrete variable in which the advertisements were divided into evenly sized temporal bins (n = 3) to capture the effect of expressions during the beginning, middle, and end of the advertisements. As before, adding interactions between the facial actions did not increase model fit.

Results indicate that the relationship between smiling (AU 12) and sharing increased over time ( $\beta = 0.05$ , *SE* = 0.02, *p* < 0.01). Although smiles at the beginning of the video were positively linked to sharing, consistent with the notion that the end of emotional experiences has a strong impact (Kahneman, Fredrickson, Schreiber, and Redelmeier 1993), smiles at the end had an even more positive effect. Although directionally similar, time effects for other actions did not reach significance.

#### **Role of Culture**

Given that the Internet-based framework that the authors utilized allowed them to collect data across multiple countries, they also explored cultural variation. Building on previous work analyzing cultural difference in facial expressivity, the authors performed an exploratory analysis testing whether the relationship between facial expressions and sharing varied cross-culturally. They focused on smiles, as most previous work studying cultural differences has analyzed smiling (Girard and McDuff, 2017; McDuff *et al.*, 2016a) and the prevalence of other emotional expressions are too infrequent for cross-cultural comparisons to be meaningful.

Smiles were always positively associated with sharing, but the exact magnitude varied cross-culturally. The results mirror crosscultural differences in individualism and collectivism (Hofstede, 2001). The relationship between smiling and sharing was largest in the United States ( $\beta_1$  = 1.54) and the United Kingdom ( $\beta_1$  = 1.54), the two countries with the largest individualism indices (91 and 89, respectively), and smallest in China ( $\beta_1$  = 0.58), the country with the smallest individualism index (20). France ( $\beta_1$  = 1.50) and Germany ( $\beta_1$  = 1.29), which have intermediate individualism indices (71 and 67, respectively), also showed intermediate relationships between smiling and sharing.

The frequency of smiling (smile base rate) also varied between countries, but this does not explain the differences in magnitude of the effects described earlier. The smile base rate was largest in the United States (5.6 percent of video frames features a smile), followed by Germany (4.5 percent), France (4.4 percent), the United Kingdom (4.3 percent), and China (1.5 percent).

Individuals in different countries saw different videos, and smiles may mean different things in different cultures, making it difficult to infer too much from these differences. That said, they are suggestive and highlight that examining cross-cultural differences in emotion and sharing is a valuable direction for future work.

#### DISCUSSION

A large-scale investigation—*i.e.*, thousands of participants and hundreds of pieces of content—suggests that facial responses are linked to sharing. Although smiles (AU 12) and nose wrinkles (AU 9) were associated with increased sharing, brow furrows (AU 4) and lip depressors (AU 15) were associated with decreased sharing.

These results suggest that not all emotions increase sharing. The fact that some facial actions associated with negative emotions increased sharing (*i.e.*, AU 9, linked to disgust), whereas other facial actions associated with negative emotions decreased sharing (*i.e.*, AU 4 and AU 15, linked to confusion and sadness) suggests that sharing is driven by more than mere valence alone.

Instead, results are more consistent with specific emotions and, potentially, arousal (Berger and Milkman, 2012). There is ambiguity around arousal associated with discrete emotion states, but emotions characterized by high arousal (*e.g.*, joy) seem to be associated with increased sharing, whereas emotions characterized by low arousal (*e.g.*, sadness) seem to decrease sharing. This suggests that advertisements that evoke high-arousal emotions (e.g., anger,

 Table 1 Regression Coefficients for the Time-Varying Liking

 Model (Overall Base Rates) and the Difference between the

 Sharing and Liking Models

Variable	Time-Varying Liking Model	Difference between Sharing and Liking Models
Intercept	3.45*** (0.07)	0.66*** (0.01)
Smile	1.56*** (0.04)	-0.06 (0.04)
Eyebrow furrow	-0.18^ (0.10)	-0.38*** (0.10)
Lip depressor	-0.36*** (0.04)	-0.14 (0.05)
Nose wrinkle	-0.19* (0.05)	0.06 (0.05)
Eyebrow raise	-0.37* (0.12)	-0.80*** (0.13)

Notes: Standard errors in brackets;  $^p < 0.01$ ;  $^p < 0.05$ ;  $^{***}p < 0.01$ .

anxiety, disgust, and surprise) may be shared more. The authors hope that demonstrating these effects in an advertising context and using a range of specific negative emotions are part of the contributions of this research.

Ancillary results are also consistent with work on recency or end effects in emotional experiences (Kahneman *et al.*, 1993), suggesting that latter parts of videos and other content may have more impact on sharing. Future work might examine this in text content, analyzing whether articles that evoke high-arousal emotions toward the end of the piece are more likely to be shared.

#### **Alternative Explanations**

One could argue that, rather than reflecting sharing, participants sharing responses simply indicated their reactions to the videos. Three things cast doubt on this possibility. First, such an account has difficulty explaining why some negative emotions, such as disgust, seem to be positively linked to sharing, whereas others, such as sadness, seem to be negatively linked to sharing.

Second, results for liking are different than for sharing (See Table 1). Participants also rated how much they liked each video (item: "How much did you like the advertisement you just watched?") on a scale ranging from 1 (not at all) to 5 (very much). Running the same analyses predicting liking rather than sharing showed different effects (See Table 1). Facial actions commonly associated with positive emotions increased liking, and facial actions commonly associated with negative emotions decreased liking. Although nose wrinkles (often associated with disgust) increased sharing, for example, they decreased liking ( $\beta = -0.37$ , *SE* = 0.12, *p* < .05).

Third, the study's main results persist even when controlling for liking. The fact that facial expressions have different relationships with liking and sharing cast doubt on the possibility that the sharing measure is merely picking up liking of the videos.

#### Implications

These findings offer an important methodological contribution for people interested in studying facial expressions. A simple online tool allows researchers to collect facial responses from a range of individuals in their natural environment. The fact that automatic coding is as reliable as individuals' manual coding means that data can be collected and analyzed at scale. This opens up a range of avenues for further research.

Trained models produced with the approach used in this research have been shown to be robust to changes in ethnicity (McDuff *et al.*, 2016a), gender (McDuff, Kodra, el Kaliouby, and LaFrance, 2017), and age (McDuff, 2017), making them widely applicable. The trained classifiers the authors used are the same as those used in the AFFDEX software development kit (McDuff *et al.*, 2016b), which is publicly available. This allows other researchers to use the same system for future experiments.

The findings also have a number of practical implications. First, the results provide suggestions for boosting shares. Most advertisements already try to make people smile, but the current findings suggest that certain negative emotions, such as disgust, may boost transmission as well. Consequently, content creators need not avoid all negative emotions; in fact, some negative emotions may help content propagation (Berger and Milkman, 2012).

Second, the tools used here can be useful in advertisement design and copy testing more broadly. Rather than relying on evaluations of the entire advertisement as a whole, marketers can use moment-to-moment analysis of facial expressions to determine which particular components may be working as desired and which should be replaced. These tools allow companies to determine whether one character can be dropped, for example, or a certain scene shifted, without having to replace the whole piece of content. Although we have applied these measures to sharing, they can be applied just as easily to evaluation and other outcomes.

Third, even once content has been created, these methods may be useful in determining resource allocation. Should more resources be put behind seeding and showing advertisement A or advertisement B? By estimating the likelihood of sharing, facial responses can help determine the likely impact of an advertisement. It also can help determine which advertisements might be better suited for television and which are better suited for social media on the basis of the relative expectation of sharing versus other downstream outcomes.

#### **Limitations and Future Research**

As with many methods, there are also limitations. When specific facial movements are often tied to emotions, the link is not always simple. Both envy and disgust, for example, include nose wrinkles. Therefore, when applying these tools, researchers and practitioners should be careful to understand what different facial movements indicate.

Another question is how these results translate into sales. As is often the case in research on word of mouth, the outcome measure in this study was sharing intentions, but one could wonder whether this actually links to sharing. Just because someone shares an advertisement, that does not always lead to sales. That said, word-of-mouth intentions are a reasonably good proxy of actual sharing (Berger and Schwartz, 2011), and more word of mouth tends to increase sales (Godes and Mayzlin, 2009). Future work might further test these relationships.

Although the authors focused on video advertisements, future work may apply these methods to advertising that relies on still images. By combining eye-tracking measures with facial responses, one may gain insight into which features of an advertisement are generating which reactions. Similar approaches may even be applied to text advertisements.

In conclusion, facial expressions provide a valuable tool to predict and understand consumer behavior. The authors hope that this work will encourage more researchers and practitioners to utilize emerging tools in this area.

#### ACKNOWLEDGMENT

This work benefited from helpful conversations with Kristen Lindquist and Jordan Etkin.

#### ABOUT THE AUTHORS

**DANIEL McDurF** is a principal researcher at Microsoft where he is working on scalable artificial intelligence tools for understanding human behavior, health, and well-being. He completed his PhD at the MIT Media Lab, where his research helped spawn a new field of imaging-based physiological measurement. Previously, McDuff was director of research at Media Lab spinout Affectiva, where he led analysis of the largest facial expression dataset in the world. His work has received awards from *Popular Science*, SXSW, The Webby Awards, and ESOMAR and has been reported in *The Times*, *The New* York *Times*, *The Wall Street Journal*, BBC News, *Scientific American* and *Forbes*.

JONAH BERGER is a marketing professor at the Wharton School of the University of Pennsylvania and an expert on word-of-mouth, natural-language processing and how products, ideas, and behaviors catch on. He has published over 50 articles in academic journals such as the Journal of Consumer Research, Journal of Marketing Research, Journal of Marketing, and Proceedings of the National Academy of Sciences. Popular accounts of his work often appear in places such as The New York Times, The Wall Street Journal, and Harvard Business Review. Berger is the author of books including Contagious, Invisible Influence, and The Catalyst: How to Change Anyone's Mind.

------

#### REFERENCES

BERGER, J. "Word of Mouth and Interpersonal Communication: A Review and Directions for Future Research." *Journal of Consumer Psychology* 24, 4 (2014): 586–607.

BERGER, J. and K. L. MILKMAN. "What Makes Online Content Viral?" Journal of Marketing Research 49, 2 (2012): 192–205.

BERGER, J. and E. M. SCHWARTZ. "What Drives Immediate and Ongoing Word of Mouth?" *Journal of Marketing Research* 48, 5 (2011): 869–880.

BERGER, J., A. HUMPHREYS, S. LUDWIG, W. W. MOE *et al.* "Uniting the Tribes: Using Text for Marketing Insight." *Journal of Marketing* 84, 1 (2020): 1–25.

DEANGELIS, M., A. BONEZZI, A. M. PELUSO, D. D. RUCKER, and M. COSTABILE. "On Braggarts and Gossips: A Self-Enhancement Account of Word-of-Mouth Generation and Transmission." *Journal of Marketing Research* 49, 4 (2012): 551–563.

DE LA TORRE, F. AND J. F. COHN. "Facial Expression Analysis." In *Visual Analysis of Humans*, T.B. Moeslund, A. Hilton, V. Krüger, and L. Sigal, eds. London: Springer, 2011.

EKMAN, P. and E. L. ROSENBERG. What the Face Reveals: Basic and Applied Studies of Spontaneous Expression Using the Facial Action Coding System (FACS). Oxford: Oxford University Press, 1997.

EKMAN, P., W. V. FRIESEN, M. O'SULLIVAN, A. CHAN *et al.* "Universals and Cultural Differences in the Judgments of Facial Expressions of Emotion." *Journal of Personality and Social Psychology* 53, 4 (1987): 712–717.

EKMAN, P., W. V. FRIESEN, and J. HAGER. Facial Action Coding System: A Technique for the Measurement of Facial Movement. Salt Lake City, UT: Research Nexus, 2002.

GIRARD, J. M. and D. McDUFF. "Historical Heterogeneity Predicts Smiling: Evidence from Large-Scale Observational Analyses." In 2017 12th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2017). Washington, DC: IEEE, 2017.

HEATH, C., C. BELL, and E. STERNBERG. "Emotional Selection in Memes: The Case of Urban Legends." *Journal of Personality and Social Psychology* 81, 6 (2001): 1028–1041.

HOFSTEDE, G. H. Culture's Consequences: Comparing Values, Behaviors, Institutions and Organizations Across Nations, 2nd ed. Thousand Oaks, CA: Sage Publications, 2001.

KAHNEMAN, D., B. L. FREDRICKSON, C. A. SCHREIBER, and D. A. REDELMEIER. "When More Pain Is Preferred to Less: Adding a Better End." *Psychological Science* 4, 6 (1993): 401–405.

KAMINS, M. A., V. S. FOLKES, and L. PERNER. "Consumer Responses to Rumors: Good News, Bad News." *Journal of Consumer Psychology* 6, 2 (2008): 165–187. KASSAM, K. S. Assessment of Emotional Experience Through Facial Expression. Cambridge, MA: Harvard University, 2010.

KASSAM, K. S. and W. B. MENDES. "The Effects of Measuring Emotion: Physiological Reactions to Emotional Situations Depend on Whether Someone is Asking." *PloS ONE* 8, 6 (2013): e64959.

LERNER, J. S. and D. KELTNER. "Beyond Valence: Toward a Model of Emotion-Specific Influences on Judgement and Choice." *Cognition & Emotion* 14, 4 (2000): 473–493.

LUMINET, O., P. BOUTS, F. DELIE, A. S. R. MANSTEAD, and B. RIMÉ. "Social Sharing of Emotion Following Exposure to a Negatively Valenced Situation." *Cognition and Emotion* 14, 5 (2000): 661–688.

MATSUMOTO, D. "Cultural Similarities and Differences in Display Rules." Motivation and Emotion 14, 3 (1990): 195–214.

MATSUMOTO, D., S. H. YOO, J. FONTAINE, A. M. ANGUAS-WONG *ET AL.* "Mapping Expressive Differences Around the World: The Relationship Between Emotional Display Rules and Individualism Versus Collectivism." *Journal of Cross-Cultural Psychology* 39, 1 (2008): 55–74.

McDuFF, D. Crowdsourcing Affective Responses for Predicting Media Effectiveness. (Doctoral dissertation), Massachusetts Institute of Technology, 2014.

McDUFF, D. "Smiling from Adolescence to Old Age: A Large Observational Study." In 2017 Seventh International Conference on Affective Computing and Intelligent Interaction (ACII). Washington, DC: IEEE, 2017.

McDuff, D., J. M. GIRARD, and R. EL KALIOUBY. "Large-Scale Observational Evidence of Cross-Cultural Differences in Facial Behavior." *Journal of Nonverbal Behavior* 41, 1 (2016a): 1–19.

McDuff, D., MAHMOUD, A., MAVADATI, M., AMR, M., TURCOT, J., and EL KALIOUBY, R. "AFFDEX SDK: A Cross-Platform Real-Time Multi-Face Expression Recognition Toolkit." In *Proceedings of the 2016 CHI Conference: Extended Abstracts on Human Factors in Computing Systems*, J. Kaye, ed. New York: Association for Computing Machinery, 2016b.

McDuff, D., E. Kodra, R. El Kaliouby, and M. LaFrance. "A Large-Scale Analysis of Sex Differences in Facial Expressions." *PloS ONE* 12, 4 (2017): e0173942.

OSTER, H. "Facial Expression and Affect Development." In *The Development* of *Affect*, M. Lewis, ed. New York: Springer US, 2017.

PACKARD, G. and D. B. WOOTEN. "Compensatory Knowledge Signaling in Consumer Word of Mouth." *Journal of Consumer Psychology* 23, 4 (2013): 434–450.

PENNEBAKER, J. W., M. R. MEHL, and K. NIEDERHOFFER. "Psychological Aspects of Natural Language Use: Our Words, Our Selves." *Annual Review of Psychology* 54, 1 (2003): 547–577.

RIMÉ, B. "Emotion Elicits the Social Sharing of Emotion: Theory and Empirical Review." *Emotion Review* 1, 1 (2009): 60–85.

RIMÉ, B., B. MESQUITA, P. PHILIPPOT, and S. BOCA. "Beyond the Emotional Event: Six Studies on the Social Sharing of Emotion." *Cognition and Emotion* 5, 5–6 (1991): 435–465.

ROZIN, P. and A. B. COHEN. "High Frequency of Facial Expressions Corresponding to Confusion, Concentration, and Worry in an Analysis of Naturally Occurring Facial Expressions of Americans." *Emotion* 3, 1(2003): 68–75.

SENECHAL, T., D. J. McDuff, and R. EL KALIOUBY. "Facial Action Unit Detection Using Active Learning and an Efficient Non-Linear Kernel Approximation." In 2015 IEEE International Conference on Computer Vision Workshop (ICCVW). Washington, DC: IEEE, 2015.

SMITH, C. A. and P. C. ELLSWORTH. "Patterns of Cognitive Appraisal in Emotion." *Journal of Personality and Social Psychology* 48, 4 (1985): 813–838.

TEIXEIRA, T., M. WEDEL, and R. PIETERS. "Emotion-Induced Engagement in Internet Video Advertisements." *Journal of Marketing Research* 49, 2 (2012): 144–159.

TEIXEIRA, T., R. PICARD, and R. EL KALIOUBY. "Why, When, and How Much to Entertain Consumers in Advertisements? A Web-Based Facial Tracking Field Study." *Marketing Science* 33, 6 (2014): 809–827.

TESSER, A. and S. ROSEN. "The Reluctance to Transmit Bad News." *Advances in Experimental Social Psychology* 8 (1975): 193–232.

TSAI, J. L. and U. CHENTSOVA-DUTTON. "Variation Among European Americans in Emotional Facial Expression." *Journal of Cross-Cultural Psychology* 34, 6 (2003): 650–657.

XIAO, L., and M. DING. "Just the Faces: Exploring the Effects of Facial Features in Print Advertising." *Marketing Science* 33, 3 (2014): 338–352.