Neural mechanisms of emotion regulation moderate the predictive value of affective and value-related brain responses to persuasive messages

Doré, B.P., Tompson, S.H., O'Donnell, M.B., An, L., Strecher, V., and Falk, E.B.

Abstract WC: 196
Significance WC: 119
Intro WC: 648
Discussion WC: 1500
Figures: 3

Corresponding authors:
Bruce Dore (brucedore@gmail.com)
Emily Falk (falk@asc.upenn.edu)

Abstract

Emotionally evocative messages can be an effective way to change behavior, but the neural pathways that translate messages into effects on individuals and populations are not fully understood. We used a human functional neuroimaging approach to ask how affect-, value-, and regulation-related brain systems interact to predict effects of graphic anti-smoking messages for individual smokers (both males and females) and within a population-level messaging campaign. Results indicated that increased activity in the amygdala, a region involved in affective reactivity, predicted both personal quit intentions and population-level information-seeking, and this was mediated by activity in ventromedial prefrontal cortex (vmPFC), a region involved in computing an integrative value signal. Further, the predictive value of these regions was moderated by expression of a meta-analytically defined brain pattern indexing emotion regulation. That is, amygdala and vmPFC activity strongly tracked with population behavior only when participants showed low recruitment of this brain pattern, which consists of regions involved in goal-driven regulation of affective responses. Overall, these findings suggest: i) that affective and value-related brain responses can predict the success of persuasive messages, and ii) that neural mechanisms of emotion regulation can shape these responses, moderating the extent to which they track with population-level message impact.

Significance Statement

People and organizations often appeal to our emotions in order to persuade us, but how these appeals engage the brain to drive behavior is not fully understood. We present an fMRI-based model that integrates affect-, control-, and value-related brain responses to predict the impact of graphic anti-smoking stimuli within a small group of smokers and a larger-scale public messaging campaign. This model indicated that amygdala activity predicted the impact of the anti-smoking messages, but that this relationship was mediated by vmPFC and moderated by expression of a distributed brain pattern associated with regulating emotion. These results suggest that neural mechanisms of emotion regulation can shape the extent to which affect and value-related brain responses track with population behavioral effects.

Persuasive messages succeed or fail depending not only on their content, but on the thoughts and feelings they elicit in people exposed to them. However, the neural pathways that translate messages into effects on individuals and populations are not fully understood. Here, we propose that during exposure to persuasive appeals, brain systems for emotional reactivity and stimulus valuation interact with brain systems involved in emotion regulation, and in concert these systems predict message effects.

Neuroimaging provides a powerful method for measuring mechanisms that underlie persuasive messaging effects (Berkman and Falk, 2013; Falk and Scholz, 2018; Genevsky and Knutson, 2018). Converging evidence suggests that brain activity associated with stimulus valuation can predict the population-level effects of stimuli like music clips, loan appeals, video advertisements, and persuasive health messages (Berns and Moore, 2012; Falk et al., 2012; Genevsky and Knutson, 2015; Genevsky, Yoon, and Knutson, 2017; Scholz et al., 2017; Venkatraman et al., 2015). However, although emotional processes feature centrally in psychological theories of persuasion (DeSteno et al., 2004; Dillard, 2000; Petty, Fabrigar, and Wegener, 2003; Witte, 1992), prior neuroimaging studies have not clarified how brain systems involved in generating and regulating emotions relate to population effects of emotionally arousing stimuli.

Further, behavioral work on emotion and persuasion has revealed variability in the effects of emotional content in changing behavior, but has not identified biological mechanisms that give rise to this variability (Peters, Ruiter, and Kok, 2013; Tannenbaum et al. 2015; Witte and Allen 2000). Together, models of the brain systems underlying emotional reactivity, emotion regulation, and integrative valuation provide an organizing framework for understanding how persuasive messages can exert their effects. Work in affective neuroscience implicates the amygdala in the generation of affective arousal (Phelps and Ledoux, 2005; Cunningham and Brosch, 2012). In tandem, other work has shown that cognitively regulating emotion recruits a network of prefrontal, parietal, and temporal brain regions that can evoke goal-driven modulation of subcortical regions like amygdala (Ochsner et al., 2012). Finally, a third line of research implicates the ventromedial prefrontal cortex (vmPFC) in integrating information from diverse brain regions into a summary signal of the subjective value of a stimulus (Barta et al., 2013; Haber and Knutson, 2010; Roy, Shohamy, and Wager, 2012). In order to understand how persuasive efforts can change attitudes and behaviors of individual people and populations, it is crucial to examine how these brain systems work together in the moment that persuasive messages are delivered.

In the present study, we test a neural account of how affective reactivity, integrative valuation, and emotion regulation interact during exposure to emotionally evocative health messages, and how this interaction relates to individual- and population-level message effects. Specifically, we hypothesized that amygdala reactivity serves as an input to a value signal reflected in vmPFC activity that more directly predicts behavior. Further, given that emotion regulation processes can modulate neural responses according to cognitive goals, we asked whether the pathway from amygdala to behavioral impact via vmPFC is moderated by recruitment of emotion regulation-related brain activity. That is, if recruitment of emotion regulation mechanisms reflects top-down modulation of neural reactivity in accordance with idiosyncratic goals, it may decrease the extent to which amygdala and vmPFC responses track with population behavior.

To address these questions, we used functional magnetic resonance imaging (fMRI) data collected from a group of smokers exposed to graphic anti-smoking messages, and data from a population-level email campaign that examined information seeking behavior evoked by the same graphic messages amongst likely smokers (Falk et al., 2016). Our analyses applied a multilevel mediation framework to ask: i) whether amygdala activity could predict personal quit intentions (for the smokers who were scanned) and population-level information seeking (in the email campaign), ii) whether relationships between amygdala activity and message effects were mediated by increased activity in vmPFC, and iii) whether this amygdala to vmPFC to message impact mediation pathway was moderated by emotion regulation pattern expression.

Method

Participants

Fifty smokers were recruited for the fMRI study. However, four participants were excluded due to excessive head motion (two participants), neurological abnormality (one participant), or data transfer difficulties (one participant). The remaining 46 subjects included 28 males and 18 females with a mean age of 31.9 (S = 12.5, range of 19–64 years old). To participate in the study, during a screening session subjects had to report smoking at least five cigarettes per day for the past month, having been a smoker for at least 12 months, and be between the ages of 18 and 65. Subjects also had to meet standard fMRI eligibility criteria, including no metal in their body, no history of psychiatric or neurological disorders, weight under 350 pounds, and currently not taking any psychiatric or illicit drugs. On the day of the scan, subjects reported smoking an average of 13 cigarettes per day (M = 13.1, SD = 6.8). Participants had smoked for an average of 15 years (M = 14.8, SD = 12.2). This dataset has been reported on in a previous paper focusing on the predictive value of activity within a region of mPFC associated with self-relevance (Falk et al., 2016). Here we report only novel analyses investigating interactions of emotional reactivity, valuation, and emotion regulation related brain activity in predicting within-individual and population effects of negatively-valenced messages.

For the population-level email campaign, likely smokers aged 18+ years within the State of New York were identified by a third party email messaging service (National Data Group; NDG). These participants had opted in to receive emails from NDG. NDG uses credit card data as one primary source of their direct email marketing lists and uses information such as purchases of cigarettes and disclosure of cigarette use in online marketing surveys or promotions as indicators of smoker status; NDG did not link to data from tobacco companies. This dataset has previously been used to study brain how self-related brain activity tracks stimulus click-through-rates virality in a stimulus-to-stimulus dependent manner (Falk et al., 2016). Here we report the results of novel analyses aiming to enhance our mechanistic understanding of how amygdala activity, vmPFC activity, and neural mechanisms of emotion regulation are related to quitting intentions and population sharing behavior.

Image acquisition

Neuroimaging data were acquired using a 3T GE Signa MRI scanner. The email campaign images task consisted of two functional runs (300 volumes total). Functional images were recorded using a reverse spiral sequence (TR = 2000 ms, TE = 30 ms, flip angle = 90, 43 axial slices, FOV ¼ 220 mm, slicethickness = 3 mm; voxel size = 3.44 x

3.44 x 3.0 mm). We also acquired in-plane T1-weighted images (43 slices; slice thickness = 3 mm; voxel size = 0.86 x 0.86 x 3.0 mm) and high-resolution T1-weighted images (SPGR; 124 slices; slice thickness = 1.02 x 1.02 x 1.2 mm) for use in coregistration and normalization.

Design

Scanner anti-smoking messages task. We examined neural activity as participants were exposed to each of 40 email-campaign images that were presented along with a tag-line from the body text of the email campaign: 'Stop Smoking. Start Living.'. The timing for each trial (see Figure 1) consisted of 4s of image presentation followed immediately with a 3s quit intention rating screen with the statement "This makes me want to quit" and a five-point rating scale (1: *definitely does not* to 5: *definitely does*) a jittered inter-trial interval (3–7.5 s). On other trials of this task, participants viewed additional images of scenes or faces which are not analyzed here.

Post-scanner anti-smoking messages rating task. After the scanner task, they rated each messages on 1 to 5 scales to indicate how negative the message was, how strong they found its argument to be, the extent to which it made them think about quitting, and self-relevance.

Population-level email campaign. We partnered with the New York State Smokers' Quit Line (NYSSQL) to launch an email campaign (n = 400 000 emails) in which the 40 images served as the basis for ads promoting internet-based quit-smoking resources (www.nysmokefree.com). Each target smoker of the email campaign received one of 40 images paired with text encouraging smokers to quit smoking. The images consisted of 20 negatively-valenced anti-smoking images that were modeled after proposed Food and Drug Administration graphic warning labels, as well as 20 neutral images. Because we aimed to better understand mechanisms underlying variability in the effectiveness of negatively-valenced messages, our analyses focused primarily on the 20 negatively valenced anti-smoking images. The emails also contained links to online quit-smoking resources managed by the NYSSQL. Email click-through rates in the email campaign (visits to the quit-website generate by clicking on the link in the email) were treated as the primary outcome of interest in this email campaign. The click-through rate for each image in the email campaign was calculated as the total number of clicks of links within the email divided by the number of emails opened for an image (clicks/opened). Consistent with our prior analyses, email click-through rates were rank-ordered from lowest to highest (see Falk et al., 2016).

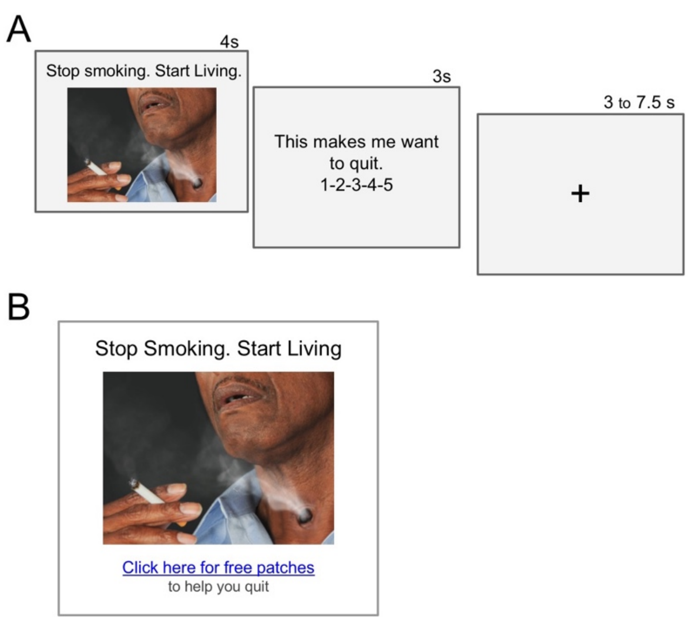

Figure 1 A) Scanner anti-smoking messages task trial structure B) Presentation of anti-smoking messages with population-level email campaign

Analysis

Preprocessing and general linear model. Functional data were preprocessed with Statistical Parametric Mapping (SPM8, Wellcome Department of Cognitive Neurology, Institute of Neurology, London, UK). Functional images were despiked using the 3dDespike program as implemented in the AFNI toolbox. Next, data were corrected for differences in the time of slice acquisition using sinc interpolation; the first slice served as the reference slice. Data were then spatially realigned to the first functional image. We then coregistered the functional and structural images using a two-stage procedure. First, in-plane T1 images were registered to the mean functional image. Next, high-resolution T1 images were registered to the in-plane image. After coregistration, high-resolution structural images were segmented to produce a grey matter mask, and then normalized to the skull-stripped MNI template provided by FSL. Finally, functional images were smoothed using a Gaussian kernel (8 mm FWHM).

We conducted first-level modelling in SPM8. Specifically, we constructed a single-trial GLM to quantify trial-level estimates of brain activity (Koyama et al., 2003; Rissman, Gazzaley, and D'Esposito, 2004), for use in multilevel modelling. Within this GLM, each stimulus (message-viewing) period of the task was modeled as a separate boxcar function convolved with the canonical hemodynamic response, generating separate estimates of brain activity (relative to implicit baseline) for each ad-viewing period, for each participant. Data were high-pass filtered with a filter of 128 seconds, and regressors for the behavioral response period and six rigid-body motion parameters were included as regressors of no interest.

Regions and pattern of interest. We defined regions and patterns of interest (ROIs and POIs) in order to estimate brain activity associated with core psychological processes of interest – including, affective reactivity, integrative valuation, and emotion regulation (coordinates refer to ROI center of mass in MNI space). Our models treated these a priori brain measures as predictor variables and the trial-by-trial message measures of message impact (quitting intentions and click-through rates) as outcome variables, and were designed to estimate the magnitude (with uncertainty) of these brain-behavior relationships using the data from the current study. We defined a bilateral amygdala ROI anatomically, using the Harvard-Oxford anatomical atlas thresholded at 25% probability (L -23, -5, -18; R 23, -4, -18) (5.3 cm^3). We defined a vmPFC ROI (1, 46, -7; 3.6 cm^3) via a meta-analysis identifying a monotonic, modality-independent signal for subjective reward value (Bartra, McGuire, and Kable, 2013). A pattern of interest is a generalization of the concept of an ROI wherein voxels are assigned continuous weights rather than a binary assignment of being included in an ROI or not. To index distributed neural processes related to emotion regulation, we used the unthresholded whole-brain results from a meta-analysis of 47 studies identifying neural regions activated for the contrast of emotion regulation versus looking naturally at emotionally evocative images (Buhle et al., 2014), without setting the predictive coefficient of any voxels (i.e., those that didn't meet a significance threshold) to exactly zero. The resulting meta-analytic pattern indicates that greater activity in a suite of brain regions implicated more generally in top-down control of cognition, as well as semantic and perceptual processing, is predictive of emotion regulation (Buhle et al., 2014; Doré, Weber, & Ochsner, 2017).

Pattern expression analysis. We conducted pattern expression analyses to test whether expression of our whole-brain pattern of interest moderated relationships between amygdala and vmPFC activity and population-level click-through rates. In order to calculate the extent to which trial-level beta images expressed the emotion regulation pattern, we treated the pattern as a vector of weights and calculated the dot product between this vector and each vectorized trial-level brain activation image, yielding a scalar value reflecting the extent to which the pattern of interest was expressed on each trial, for each participant.

Multilevel modelling. We used R (cran.r-project.org; ver 3.3.1), Stan (mc-stan.org; rstan ver 2.16.2), and the brms package (Bayesian Regression Models using Stan ver 2.1.0) to fit hierarchical Bayesian regression models that estimated the extent to which activity within our amygdala and vmPFC ROIs and expression of our emotion regulation pattern (the predictor variables) were predictive of smokers' behavioral ratings (an outcome variable) and, in separate models, email click-through-rates in the population-level email campaign (an outcome variable). Predictor and outcome variables were standardized and predictors that varied within-person were person-mean centered, yielding standardized beta coefficients indicating the average magnitude of the within-person relationship between the predictor variable and the outcome variable. Models incorporated variance and covariance parameters allowing for model intercepts and slopes to vary by person. We used posterior means and 95% Bayesian credibility intervals (central posterior intervals) to estimate the plausible range of values that a given relationship could take in light of the observed data. We did not conduct repeated tests of a universal null hypothesis (as in the mass univariate testing approach) that would require adjustment for multiple comparisons.

Because weakly informative priors centered at zero yield results that closely correspond with traditional maximum likelihood estimates (with the exception that they regularize extreme values toward zero), we used weakly informative priors on beta coefficients, variance parameters, and covariance parameters. Specifically, we used a zero-centered normal distribution with scale parameter 10 on beta coefficients, a positive half-normal distribution with scale parameter 10 for standard deviations, and an LKJ distribution with regularization parameter 1 for correlations between person-level intercepts and slopes (Stan Development Team, 2016). Models were estimated with Markov Chain Monte Carlo Sampling, running four parallel chains for 1000 iterations each (the first 500 warm-up samples for each chain were discarded). This number of iterations proved sufficient for convergence in that the Gelman-Rubin diagnostic reached a value between 0.95 and 1.05 for all parameters (Gelman and Rubin, 1992). In comparison to maximum likelihood based approaches to multilevel modelling, this approach offers: posterior inference, more accurate estimation of hierarchical variance and covariance parameters, better rates of convergence, and diagnostics for assessing the validity of the sampler-based statistical inferences (Stan Development Team, 2016).

Results

Amygdala activity during message exposure predicted personal quit intentions and population information seeking

Our initial question was whether message-evoked amygdala responses were predictive of personal quit intentions (rated within the scanner task), and population information seeking (collected via email click-through behaviors in a New York state email campaign). We found that higher amygdala activity was predictive of both in-scanner ratings of quit intentions, β=.078, 95%CI[.001, .157], and population information seeking, β=.099, 95%CI[.020, .174], across the negatively-valenced anti-smoking messages. In a next step, we additionally controlled for post-task ratings of message negativity, argument strength, self-relevance, and quit-related thoughts, finding that the predictive value of amygdala activity held for both quit intentions, β=.062, 95%CI[.007, .118], and population information seeking, β=.070, 95%CI[.005, .139]. Considering the neutral (non-negatively valenced) comparison messages, we did not see that amygdala activity was predictive of either in-scanner ratings of quit intentions, β=-.001, 95%CI[-.067, .069], or population information seeking, β=-.001, 95%CI[-.079, .072]. Overall, this pattern suggests that smokers' amygdala responses to negative anti-smoking messages predicted both their own message evoked quit intentions and the population-level impact of the same messages. Further, the value of a smoker's amygdala response in predicting population impact was partially independent of that smoker's own subjective response to the ad.

The predictive value of amygdala activity for personal quit intentions and population behavior was mediated by vmPFC activity

Next we used a multilevel mediation approach to ask whether the data were consistent with a model whereby the relationship between amygdala activity in response to emotionally evocative messages and the behavioral impact of those messages was mediated by increased vmPFC activity (associated with integrative valuation). To do this, we fit a mediation model with two outcome variables. In this model, we used estimates of message-evoked amygdala activity as the predictor variable, estimates of message-evoked vmPFC activity as the mediator variable, and in-scanner quit intentions and population information seeking as two parallel outcome variables. The results of this model (see Figure 2) indicated that the relationship between amygdala activity and quit intentions was mediated by vmPFC activity, $a*b_{quit\_int}$=.044, 95%CI[.017, .074], and, similarly, the relationship between amygdala activity and population information seeking was also mediated by vmPFC activity, $a*b_{pop\_beh}$=.038, 95%CI[.005, .073]. Within the neutral comparison messages, we did not see mediation by vmPFC for either quit intentions, a*b= .020, 95%CI[-.011, .051], or population information seeking, a*b= -.051, 95%CI[-.013, .025]. Overall, this pattern of results is consistent with a model whereby amygdala responses reflect bottom-up affective reactivity to an emotionally evocative anti-smoking message, and this bottom-up reactivity is predictive of individual- and population-level message effects to the extent that it tracks with corresponding valuation-related responses in vmPFC.

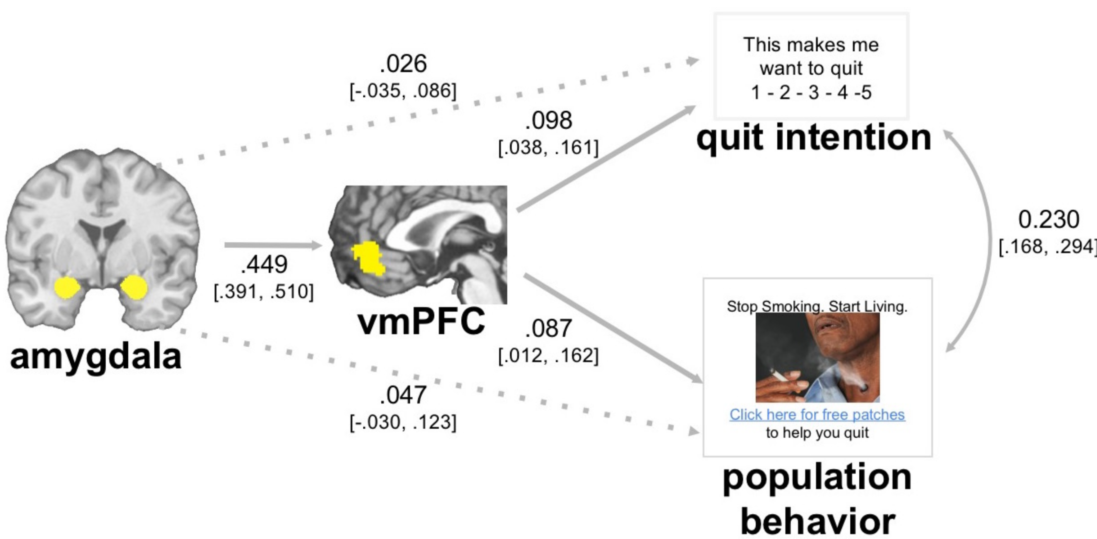

Figure 2 The relationship between amygdala activity and message impact was mediated by vmPFC activity, for both smoker's own ratings of message-evoked quit intentions within the scanner task and population information seeking behavior (click-through rates)

in the New York state email campaign that used the same messages. (Path coefficients reflect standardized betas with 95% credibility intervals, estimated with a multilevel structural equation model.)

Emotion regulation pattern expression moderated the predictive value of amygdala activity and vmPFC for population behavior

Our previous analyses indicated that amygdala responses were predictive of both smoker’s own intentions to quit smoking and information seeking within a population messaging campaign, and these effects were mediated by increased vmPFC activity. However, these analyses did not speak to the relevance of brain regions implicated in deliberately controlling one’s emotional response to an aversive stimulus. In order to address this, we asked whether spontaneous (i.e., uninstructed) recruitment of regions implicated in emotion regulation could moderate the predictive value of amygdala and vmPFC responses to these anti-smoking messages.

We first quantified trial-by-trial expression of a meta-analytically defined global brain pattern that is predictive of uninstructed down-regulation of negative emotional responses to aversive stimuli (Buhle et al., 2014; Doré et al., 2017). Next, we asked whether expression of this emotion regulation pattern moderated the predictive relationships linking amygdala and vmPFC activity with personal quit intentions and population-level information seeking. Considering smoker’s own personal quit intentions, we saw that emotion regulation pattern expression did not clearly moderate the predictive value of either amygdala activity, $\beta=.005$, 95%CI[-.037, .047], or vmPFC activity, $\beta=-.022$, 95%CI[-0.061, 0.019]. This result suggests that recruitment of neural mechanisms of emotion regulation did not clearly moderate the extent to which smokers’ affect- and valuation-related brain responses corresponded with their own message-evoked quit intentions. Turning to message-evoked population information seeking, we saw that emotion regulation pattern expression moderated the predictive value of both amygdala activity, β=-.039, 95%CI[-.078, -.001], and vmPFC activity, β=-.047, 95%CI[-.084, -.007]. This result suggests that recruitment of neural mechanisms of emotion regulation moderated the extent to which smoker's affect- and valuation-related brain activity tracked with population message impact. That is, the more the emotion regulation pattern was expressed, the less predictive the affect- and value-related responses were of the population-level impact of these negatively valenced anti-smoking messages.

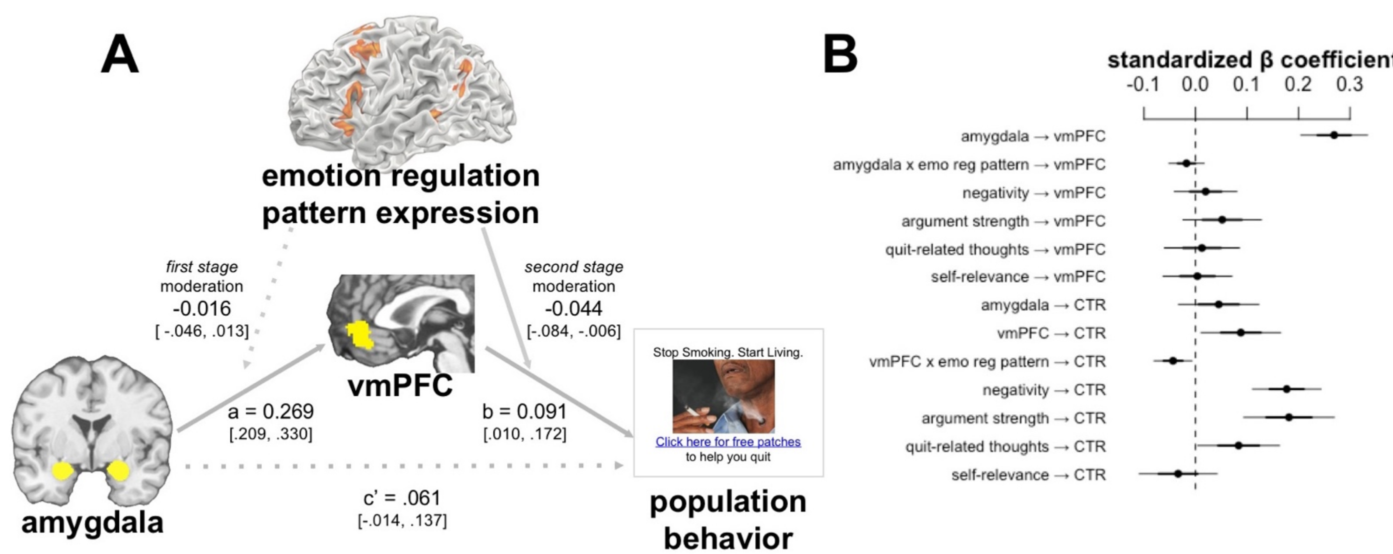

Figure 3 A) Emotion regulation pattern expression moderated the magnitude of the path relating vmPFC activity to population click-through rates. (Emotion regulation pattern is thresholded at FWE p<.05 for visualization, but in analyses all voxel weights were used; Path coefficients reflect standardized betas with 95% credibility intervals.) **B)** Coefficient plot (standardized betas with 68% and 95%CI) displaying parameters estimated in a moderated mediation model that including covariates for subjective ratings of message negativity, argument strength, quit-related thoughts, and self-relevance.

To better understand the nature of this moderation, we next modified our mediation model in order to incorporate trial-by-trial variation in emotion regulation pattern expression as a moderator of both the first stage (i.e., the extent to which amygdala activity showed a positive relationship with vmPFC activity), and the second stage (i.e., the extent to which vmPFC activity showed a positive relationship with population behavior, accounting for amygdala activity) of the amygdala to population behavior mediation pathway. This moderated mediation model (see Figure 3A) revealed evidence for second stage moderation – that is, emotion regulation pattern expression moderated the (b path) relationship between vmPFC activity and population behavior, β=-.048, 95%CI[-.087, -.008]. The model estimate for first stage moderation of the (a path) relationship between amygdala activity and vmPFC activity was β=-.013,95%CI[-.047, .020]. An analogous model fit to the data from the neutral comparison messages indicated that emotion regulation pattern did not moderate either the relationship between amygdala activity and vmPFC activity, β = -.008, 95%CI[-.042, .025], or the relationship between vmPFC and population click-through rate β = -.017, 95%CI[-.057, .022].

Overall, these results indicate that spontaneous (i.e., uninstructed) recruitment of brain mechanisms associated with emotion regulation disrupted the extent to which affect and valuation-related brain responses corresponded with population behavioral responses to the negative health messages. In an exploratory follow-up to the mediation model, we added subjective ratings of message negativity, argument strength, quit-related thoughts, and self-relevance as additional covariates within the mediation model, finding that the vmPFC mediation and emotion regulation pattern moderation effects held (see Figure 3B). In set of final exploratory follow-up analyses, we asked whether within-person amygdala, vmPFC, or emotion regulation pattern expression interaction effects in predicting quit intentions and population information seeking were moderated by individual differences in average number of cigarettes smoked per day, number of prior quit attempts, or nicotine dependence. These analyses did not reveal evidence for moderation by any of these individual difference variables.

Discussion

Persuasive messages often use emotional content in service of promoting particular attitudes or behaviors (Peters, Ruiter, and Kok, 2013; Tannenbaum et al. 2015; Witte and Allen 2000). In this study, we asked how neural mechanisms of emotional reactivity, valuation, and emotion regulation interact to determine the individual-level and population-level impact of emotionally evocative antismoking messages depicting a range of negative, smoking-related consequences (e.g., lung disease), a crucial step toward a mechanistic understanding of how emotionally-laden persuasive appeals can change behavior in individual people and across larger populations. Specifically, we examined brain activity within the amygdala and vmPFC, and quantified expression of a meta-analytically defined brain pattern that is predictive of decisions to implement emotion regulation. Next, we used these brain variables to predict effects of these messages on smokers' personal intentions to quit smoking and on large-scale information seeking behavior (email click-through rates) within an anti-smoking campaign. We obtained three key findings.

First, we found that higher amygdala activity during message exposure was predictive of greater message impact, both for smokers' personal quit intentions and for population email click-through rates to the same messages. Second, we found that the relationship between amygdala activity and message impact was mediated by increased activity in vmPFC. Third, we found that the strength of the mediation pathway relating amygdala responses to population behavior via vmPFC was moderated (i.e., dampened) by expression of the meta-analytically defined emotion regulation pattern. That is, when smokers spontaneously recruited brain regions involved in deliberately regulating one's response to an aversive image, their affect- and valuation-related brain responses tended to more weakly correspond with the population behavioral effects of the same anti-smoking messages. This overall pattern of results comports with theories positing that vmPFC reflects an integrative value signal that is sensitive to historical and motivational context (crucial in affective learning processes like extinction of a conditioned threat response), whereas lateral PFC regions are more directly involved in top-down control of cognition and emotion.

Implications for psychological and neural models of persuasion

These data are consistent with research proposing that affective responses are a core component of attitudes in general, and central to the effects of persuasive communications in particular (DeSteno et al., 2004; Dillard and Peck, 2000; Petty, Fabrigar, and Wegener, 2003; Witte, 1992). For example, health messages that highlight a person's risk can be effective (Witte and Allen 2000; Tannenbaum et al. 2015), particularly if receivers perceive themselves to be at risk (Brewer et al. 2004; Sheeran et al. 2014; Schmälzle et al. 2017). However, our data provide new insight into the underlying mechanisms of these kinds of effects, suggesting that whether an emotionally charged health message succeeds or fails can be predicted on the basis of both the affective reactivity that it elicits and the way that this affect is regulated.

Overall, these data indicate that theoretically-defined patterns of brain activity can interact to predict consequential behaviors of both the individuals whose brains were measured, and larger population of people exposed to the same stimuli. Specifically, these data provide novel insight into how mechanisms of emotional reactivity, integrative valuation, and regulation may interact to produce change in intentions and behaviors. Although amygdala activity was predictive of personal intentions to quit smoking and population quit-related information seeking behavior, these predictive relationships were largely mediated by vmPFC activity, suggesting that it may not be enough for persuasive communications to elicit emotional reactivity if they do not also contribute positively to an integrative value response.

Moreover, the pathway relating amygdala activity to population behavior via vmPFC activity was moderated by emotion regulation pattern expression. That is, when smokers showed recruitment of neural mechanisms involved in regulating negative affect during message exposure, their amygdala and vmPFC activity did not correspond as closely with population-level message impact. However, emotion regulation pattern expression did not moderate the extent to which amygdala and vmPFC activity tracked with smoker's own personal intentions to quit smoking. Importantly, the prefrontal and parietal brain regions that are reflected in the emotion regulation pattern are thought to be domain-general, in that they are implicated in regulating emotion via several strategies (including both up- and down-regulating emotion) as well as in top-down control of cognition more generally (Ochsner et al., 2012). One interpretation of this pattern of results is that the emotion regulation pattern expression observed here may reflect spontaneous up- or down-regulation of the affective impact of a message, leading to an evaluative response that is more driven by idiosyncratic top-down goals and less by stimulus-driven modes of responding that are more universal across the population of people exposed to these messages (Knutson and Genevsky, 2018; Doré et al., in press). That is, the observation that amygdala and vmPFC responses diverge from population behavior (but not one's own intentions) under instances where emotion regulation pattern expression is high suggests that affective and value-related brain responses can be shaped by neural mechanisms of emotion regulation such that they reflect more idiosyncratic, or population-atypical, ways of evaluating a stimulus. Equivalently, it may be that affective and value-related brain responses are most predictive of the population impact of persuasive messages when they are less influenced by top-down regulation (Doré et al., in press).

Implications for persuasive messaging interventions

The view that individual- and population-level message effects may emerge from affect-, value-, and control-related brain responses has several implications for theories and design of persuasive messaging interventions. First, messages that aim to evoke strong negative emotional reactivity (e.g., an image of a diseased lung) can be effective, but these data suggest an important boundary condition: this emotional reactivity occurs in tandem with value-related brain activity. Further, these data indicate that an important source of variance in the extent to which affect and value-related brain responses track with persuasive message effects is whether perceivers spontaneously recruit mechanisms of emotion regulation. In particular, the results suggest that neuroimaging measures may be most informative about the success of public health messages when the people being scanned approach the stimuli in a more reactive and less controlled manner (i.e., do not recruit neural mechanisms of emotion regulation). That is, this observation suggests that neural prediction of population behavior may be more efficient when recruitment of mechanisms of emotion regulation is experimentally manipulated or statistically accounted for. One reason for this finding could be that perceivers are recruiting cognitive control resources to up-or down-regulate their affective response such that it diverges from the typical population response. Drawing from models of the neural mechanisms predicting decisions to enact control of emotion (Doré, Weber, and Ochsner, 2017; Shenhav, Botvinick, and Cohen, 2013) future studies could focus on psychological and neural factors that could make recruitment of emotion regulation mechanisms more or less likely during the moment of message receipt. Overall, a crucial direction for future work is to quantify how bottom-up (stimulus-driven) and top-down (goal-driven) neural processes interact to determine the how people experience, recover from, and change their behavior in response to emotionally evocative messaging. In service of this overarching goal, future studies could ask whether psychological and neural variables are able to predict not only the typical effects of a message within a population (i.e., the mean), but also the presence of highly similar versus dissimilar effects across members of that population (i.e., the variance). Such an approach will build our understanding of specific persuasive strategies that give rise to relatively universal versus relatively idiosyncratic behavioral responses (Schmaelze et al., 2015).

Future Directions

In this study we used measurements of brain activity to predict immediate effects of graphic anti-smoking messages on smokers' quit-intentions and on information-seeking behavior in a population email campaign. A question for future work is whether stimuli that drive amygdala responses are more effective in evoking short-term change in behavior versus sustained longer-term behavior change. Further, this study focused on within-person variation in anti-smoking message impact. Whether similar neural mechanisms underlie within- versus between-person variation in message-evoked behavior change is not well understood, and is an important question for future work. Related, although vmPFC activity consistently tracks with subjective value, it is also involved in other processes like affective learning and self-relevance (Roy, Shohamy, & Wager, 2013). Future work could investigate how these processes, or another lower-level set of computations, are relevant for understanding responses to emotional health messages. Finally, future work examining neural responses to persuasive messages could also include psychological measures of emotion regulation abilities and tendencies (Silvers & Guassi Moreira, 2017), to examine how these individual differences may moderate prediction of message effects.

Conclusion

Whether a persuasive message succeeds or fails can hinge not only on its factual content, but on the affect it elicits and how this affect is regulated. Our data show that affect-related neural responses can predict the impact of persuasive health messages, and that this relationship is mediated by vmPFC and moderated by emotion regulation pattern expression. These findings suggest that spontaneous recruitment of neural mechanisms of emotion regulation can shape evaluations of persuasive messages and disrupt the extent to which these evaluations track with population-level message effects.

Acknowledgments

We acknowledge Richard Gonzalez and Sonya Dal Cin for collaboration on a larger study relevant to this work; Francis Tinney Jr., Kristin Shumaker, Li Chen, Nicolette Gregor, Becky Lau, Larissa Svintsitski, Cole Schaffer for assistance with data collection; Lynda Lin for assistance with data processing. This work was supported by grants from The Michigan Center of Excellence in Cancer Communication Research/ National Institutes of Health (grant number P50 CA101451, PI Strecher); National Institutes of Health New Innovator Award/ National Institutes of Health (grant number 1DP2DA03515601, PI Falk); U.S. Army Research Laboratory (including work under Cooperative Agreement Number W911NF-10-2-0022); and the National Cancer Institute at the National Institutes of Health and US Food and Drug Administration's Center for Tobacco Products (pilot grant from grant number P50CA179546). The content is solely the responsibility of the authors and does not necessarily represent the official views of any funding agencies.

References

Bartra O, McGuire JT, Kable JW (2013) The valuation system: a coordinate-based meta-analysis of BOLD fMRI experiments examining neural correlates of subjective value. Neuroimage 76:412–427.

Berkman, E.T., & Falk, E.B. (2013). Beyond brain mapping: Using neural measures to predict real-world outcomes. Current Directions in Psychological Science, 22(1), 45-50.

Berns GS, Moore SE (2012) A neural predictor of cultural popularity. J Consum Psychol 22:154–160.

Buhle JT, Silvers JA, Wager TD, Lopez R, Onyemekwu C, Kober H, Weber J, Ochsner KN (2014) Cognitive reappraisal of emotion: a meta-analysis of human neuroimaging studies. Cereb Cortex 24:2981–2990.

Carpenter B, Gelman A, Hoffman MD, Lee D, Goodrich B, Betancourt M, Brubaker M, Guo J, Li P, Riddell A (2017) Stan: A probabilistic programming language. J Stat Softw 76 Available at: https://www.osti.gov/biblio/1430202.

De Hoog N, Stroebe W, De Wit JBF (2007) The impact of vulnerability to and severity of a health risk on processing and acceptance of fear-arousing communications: A meta-analysis. Rev Gen Psychol 11:258.

DeSteno D, Petty RE, Rucker DD, Wegener DT, Braverman J (2004) Discrete emotions and persuasion: the role of emotion-induced expectancies. J Pers Soc Psychol 86:43–56.

Dillard JP, Peck E (2000) Affect and Persuasion: Emotional Responses to Public Service Announcements. Communic Res 27:461–495.

Doré BP, Boccagno C, Burr D, Hubbard A, Long K, Weber J, Stern Y, Ochsner KN (2017a) Finding Positive Meaning in Negative Experiences Engages Ventral Striatal and Ventromedial Prefrontal Regions Associated with Reward Valuation. J Cogn Neurosci 29:235–244.

Doré BP, Weber J, Ochsner KN (2017b) Neural Predictors of Decisions to Cognitively Control Emotion. J Neurosci 37:2580–2588.

Falk E, Scholz C (2018) Persuasion, Influence, and Value: Perspectives from Communication and Social Neuroscience. Annu Rev Psychol 69:329–356.

Falk EB, O'Donnell MB, Tompson S, Gonzalez R, Dal Cin S, Strecher V, Cummings KM, An L (2016) Functional brain imaging predicts public health campaign success. Soc Cogn Affect Neurosci 11:204–214.

Gelman A, Rubin DB (1992) Inference from Iterative Simulation Using Multiple Sequences. Stat Sci 7:457–472.

Genevsky A, Knutson B (2015) Neural Affective Mechanisms Predict Market-Level Microlending. Psychol Sci 26:1411–1422.

Hare TA, Malmaud J, Rangel A (2011) Focusing attention on the health aspects of foods changes value signals in vmPFC and improves dietary choice. J Neurosci 31:11077–11087.

Knutson B, Genevsky A (2018) Neuroforecasting Aggregate Choice. Curr Dir Psychol Sci 27:110–115.

Ochsner KN, Silvers JA, Buhle JT (2012) Functional imaging studies of emotion regulation: a synthetic review and evolving model of the cognitive control of emotion. Ann N Y Acad Sci 1251:E1–E24.

Peters G-JY, Ruiter RAC, Kok G (2013) Threatening communication: a critical re-analysis and a revised meta-analytic test of fear appeal theory. Health Psychol Rev 7:S8–S31.

Petty RE, Fabrigar LR, Wegener DT (2003) Emotional factors in attitudes and persuasion. Handbook of affective sciences 752:772.

Roy M, Shohamy D, Wager TD (2012) Ventromedial prefrontal-subcortical systems and the generation of affective meaning. Trends Cogn Sci 16:147–156.

Schmälzle R, Häcker FEK, Honey CJ, Hasson U (2015) Engaged listeners: shared neural processing of powerful political speeches. Soc Cogn Affect Neurosci 10:1137–1143.

Shenhav A, Botvinick MM, Cohen JD (2013) The expected value of control: an integrative theory of anterior cingulate cortex function. Neuron 79:217–240.

Silvers JA, Guassi Moreira JF (2017) Capacity and tendency: A neuroscientific framework for the study of emotion regulation. Neurosci Lett.

Tannenbaum MB, Hepler J, Zimmerman RS, Saul L, Jacobs S, Wilson K, Albarracín D (2015) Appealing to fear: A meta-analysis of fear appeal effectiveness and theories. Psychol Bull 141:1178–1204.

Witte K (1992) Putting the fear back into fear appeals: The extended parallel process model. Commun Monogr 59:329–349.

Witte K, Allen M (2000) A meta-analysis of fear appeals: implications for effective public health campaigns. Health Educ Behav 27:591–615.