

ORIGINAL ARTICLE

Brain Activity Tracks Population Information Sharing by Capturing Consensus Judgments of Value

B.P. Doré¹, C. Scholz¹, E.C. Baek¹, J.O. Garcia^{2,3}, M.B. O'Donnell¹,
D.S. Bassett^{1,3,4,5,6}, J.M. Vettel^{2,3,5,7} and E.B. Falk¹

¹Annenberg School for Communication, University of Pennsylvania, PA 19104, USA, ²US Army Research Laboratory, Adelphi, MD 20783, USA, ³Department of Bioengineering, University of Pennsylvania, PA 19104, USA, ⁴Department of Neurology, University of Pennsylvania, PA 19104, USA, ⁵Department of Physics & Astronomy, University of Pennsylvania, PA 19104, USA, ⁶Department of Electrical & Systems Engineering, University of Pennsylvania, PA 19104, USA and ⁷Department of Psychological & Brain Sciences, University of California Santa Barbara, Santa Barbara, CA 93106, USA

Address correspondence to Bruce Doré and Emily Falk, Annenberg School for Communication, 3620 Walnut St, Philadelphia, PA 19104, USA.
Email: brucedore@gmail.com; falk@asc.upenn.edu

Abstract

Information that is shared widely can profoundly shape society. Evidence from neuroimaging suggests that activity in the ventromedial prefrontal cortex (vmPFC), a core region of the brain's valuation system tracks with this sharing. However, the mechanisms linking vmPFC responses in individuals to population behavior are still unclear. We used a multilevel brain-as-predictor approach to address this gap, finding that individual differences in how closely vmPFC activity corresponded with population news article sharing related to how closely its activity tracked with social consensus about article value. Moreover, how closely vmPFC activity corresponded with population behavior was linked to daily life news experience: frequent news readers tended to show high vmPFC across all articles, whereas infrequent readers showed high vmPFC only to articles that were more broadly valued and heavily shared. Using functional connectivity analyses, we found that superior tracking of consensus value was related to decreased connectivity of vmPFC with a dorsolateral PFC region associated with controlled processing. Taken together, our results demonstrate variability in the brain's capacity to track crowd wisdom about information value, and suggest (lower levels of) stimulus experience and vmPFC–dlPFC connectivity as psychological and neural sources of this variability.

Key words: fMRI, information sharing, valuation, vmPFC

From the latest word on the benefits of high-intensity exercise to breaking news about the cause of rising health care costs, information that diffuses widely in the media environment can have a profound impact on society. We sought to better understand this phenomenon, building multilevel models that link the brain responses of individual people to the information sharing behavior of a population, and investigating sources of person-to-person differences in these links. Beyond enriching our scientific understanding of how and why information

sharing occurs, models of this kind could ultimately be used to forecast and enhance the impact of persuasive communication at population scale.

Neuroimaging methods provide a noninvasive means of monitoring the mechanisms that underlie how people perceive and evaluate stimuli. Recent studies have shown that functional brain responses can track future behavior above and beyond tracking afforded by traditional measures of intentions and attitudes (Berns and Moore 2012; Falk et al. 2010; Chua

et al. 2011; Falk et al. 2012; Wang et al. 2013). In particular, value-related responses in the ventromedial prefrontal cortex (vmPFC) within a small sample of people tend to track with the behavior of a larger population health campaign effectiveness (Falk, Berkman, and Lieberman 2012), and news article sharing (Scholz et al. 2017). Other work has shown similar relationships for ventral striatal activity, showing that ventral striatal responses tend to track with population-level music sales (Berns and Moore 2012) micro-loan funding decisions (Genevsky and Knutson 2015) and ad-related sales (Venkatraman et al. 2015).

In cognitive neuroscience, contemporary models of vmPFC implicate its activity in neural computations related to the value of a stimulus for oneself (Rangel and Hare 2010; Roy et al. 2012; Bartra, McGuire, and Kable 2013). However, these recent population prediction studies suggest that vmPFC responses can also track stimulus effects in a larger population of people. One possible explanation is that vmPFC activity reflects how a stimulus tends to be valued by other people. That is, the ability of vmPFC activity to track population behavior may be linked to how closely this activity tracks with social consensus, or crowd wisdom, about information value. Here, we refer to the extent to which vmPFC activity shows a linear relationship with population behavior as vmPFC population behavior tracking. We refer to the extent to which vmPFC shows a linear relationship with people's average ratings of information value as vmPFC consensus value tracking.

Because vmPFC integrates different inputs into a summary signal that is sensitive to personal context and motivation (Rangel and Hare 2010; Roy et al. 2012), a tendency to show vmPFC activity that tracks consensus value could emerge from individual differences in the information people are exposed to in daily life. At a mechanistic level, vmPFC value tracking could also emerge from underlying differences in how vmPFC communicates with brain systems that provide inputs into the computation of value.

We sought to address these gaps in knowledge by building multilevel models that used brain responses to New York Times articles as predictor variables in models where the outcome variables reflected how these articles tended to be valued by participants in our study, and to what extent they were shared in the broader population (Fig. 1). Our primary goals were 1) to estimate person-to-person differences in how closely vmPFC activity tracked consensus value judgments, and 2) to ask if these differences could explain how closely vmPFC activity tracked with population article sharing. Further, we investigated potential sources of this variability by asking 3) how vmPFC value tracking related to daily life news experience, and 4) how vmPFC value tracking related to functional connectivity between vmPFC and other brain regions. Collectively, our results demonstrate variability in the brain's capacity to track crowd wisdom about information value, and suggest daily life experience and connectivity between vmPFC and lateral prefrontal cortex as psychological and neural sources of this variability.

Materials and Methods

Participants

We recruited 43 adults and screened them in an initial session to confirm that they were right-handed, could read and speak fluently in English, had normal or corrected-to-normal vision, had never been diagnosed with a psychiatric or neurological disorder, were not currently using psychiatric medication or

legally prohibited drugs, were not currently pregnant or breast-feeding, and had no conditions that contraindicated MRI. Informed consent was obtained in writing according to procedures approved by the Institutional Review Board of the University of Pennsylvania. Two participants were excluded from analysis due to data corruption (one due to errors in stimulus presentation, and one due to poor normalization to the template brain), leaving a final sample of 41 (29 females) adults (mean age = 20.6 years, SD = 2.1 years, range: 18–24 years). The median level of education in this sample was some college (degree not attained), with a range from high school to doctoral degree. In this sample, 31% identified as Asian, 10% identified as Black, 7% identified as Hispanic or Latino, and 51% identified as White. This dataset has been reported on in previous papers focusing on brain activity apparent during decisions to share versus consume information (Baek et al. 2017) and a stimulus-to-stimulus analysis of self-, social-, and value-related brain activity tracking with article virality (Scholz et al. 2017). Here, we report only novel analyses aiming to enhance our mechanistic understanding of the relationship between vmPFC activity and population sharing behavior.

Image Acquisition

Data were acquired on a 3 T Siemens Magnetom TimTrio scanner with a 32-channel RF head coil for 39 participants, and a 3 T Siemens Prisma scanner with a 64-channel head/neck array for 2 participants. Structural volumes were acquired using a high-resolution T1-weighted axial MPRAGE sequence yielding 160 slices with a 0.9 by 0.9 by 1.0 mm³ voxel size. Functional volumes were acquired using a T2*-weighted image sequence with a repetition time (TR) of 1500 ms, an echo time (TE) of 25 ms, a flip angle of 70°, and a 20 cm FOV consisting of 54 slices (52 on Prisma scanner) with 3 mm thickness acquired at a negative 30° tilt to the AC–PC axis, with a 3 mm³ isotropic voxel size. Finally, we collected an in-plane structural T2-weighted image consisting of 176 axial slices with 1 mm thickness and 1 mm³ isotropic voxel size to implement a 2-stage coregistration procedure between functional and anatomical images.

Design

New York Times Article Viewing Task

Participants completed an in-scanner task in which they viewed summaries of articles (headlines and abstract) from the Health section of the “New York Times” website (www.nytimes.com) (Fig. 1). The articles were chosen from a census of articles ($N = 760$) published online in the 7 1/2 months between 11 July 2012 and 28 February 2013 (Kim 2015). Articles for the viewing task were chosen from this broader census to maximize comparability in content (i.e., healthy living and physical activity) and length (i.e., word count of title and abstract). To control for reading speed, we produced audio files where a female voice read each of the article headlines (audio files were 8, 10, or 12 s in duration). Our analyses focused on trials from this task (20 total) during which participants were asked to consider whether to read the full text of the article on the basis of the headline and abstract and, at the end of the trial, to indicate whether they were likely to read the article (1: “very unlikely” to 5: “very likely”). These ratings were used to compute the group consensus value used in our analyses. Here, we use the term consensus value to refer to the average group judgment (i.e., the nature of the consensus), and not to

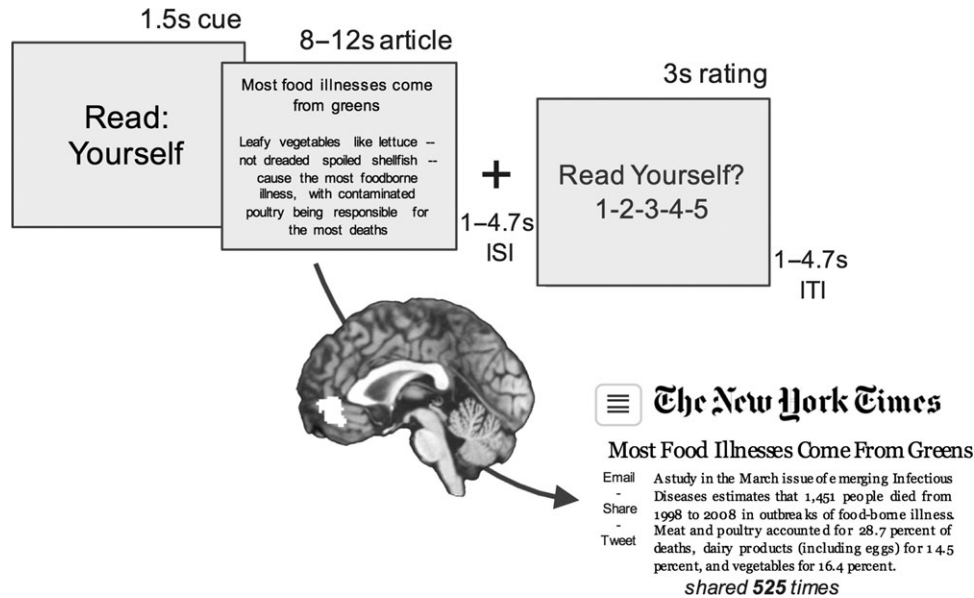


Figure 1. In-scanner New York Times article viewing task and tracking of population information sharing. Brain activity was measured as people read and listened to headlines and abstracts of New York Times articles focusing on health and fitness. Population-level counts of the number of times each article was shared online within the first 30 days after publication (via email or social media) were collected from the New York Times website. Multilevel models used brain responses from the viewing task to as a predictor variable, and counts of population article sharing as an outcome variable.

the variance of the judgments (i.e., the degree of the consensus). Within other trials of the task (not analyzed here), participants were asked to look at different article stimuli and indicate: whether the content of the article focused on a specified topic, whether they would share the article on their Facebook wall, and whether they would share the article with one friend via Facebook (Baek et al. 2017; Scholz et al. 2017). Population-level data on the number of shares of each article within 30 days of publication were collected via the NYTimes Application Program Interface (Kim 2015). Before the scan session, participants completed a survey in which they rated how often they typically: read news generally, read online newspapers, read news on the New York Times website, read articles about healthy living, and read articles about physical activity (1: less than once a month, 2: once a month, 3: 2–3 times a month, 4: once a week, 5: 2–3 times a week, 6: daily, 7: several times a day). Responses to each of these questions were averaged into an overall news reading frequency score, with a mean of 3.6 and a standard deviation of 1.1.

fMRI Analyses

Preprocessing and General Linear Model

Data were preprocessed with SPM8, incorporating tools from AFNI and FSL, and consisted of despiking, slice-time correction, realignment, coregistration of functional and structural images, and normalization to the standard Montreal Neurological Institute (MNI) brain by segmentation of the structural image. Normalized images were smoothed with an 8 mm full width at half maximum Gaussian.

First-level (individual participant) GLM analyses were implemented in SPM8. Analyses used a β -series approach in that each article viewed in 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 article viewing period, for each participant. A single regressor for all cue and behavioral

response periods, 6 rigid-body motion parameters, and a high-pass filter for 128 s were included as regressors of no interest.

Regions of Interest

We constructed a region of interest (ROI) in order to extract estimates of brain activity (and connectivity) from a region of vmPFC (MNI center of mass 1, 46, -7; 133 3 mm isotropic voxels) identified via meta-analysis as carrying a monotonic, modality-independent signal for subjective reward value (Bartra, McGuire, and Kable 2013). We also conducted follow-up analyses focusing on ventral striatum, as well as brain regions associated with self-related processing, and brain regions associated with social cognition (see Supplementary Materials). For analyses assessing vmPFC connectivity with the rest of the brain, we constructed spheres of 8 mm radius (79 3 mm isotropic voxels) at 264 locations spanning the entire cortex defined on the basis of a large-scale study (Power et al. 2011). By defining this whole-brain set of ROIs, we were able to take a network approach to the analysis of brain connectivity.

Multilevel Modeling

We used R (cran.r-project.org; ver 3.3.1), Stan (mc-stan.org; rstan ver 2.10), and the “brms” package (Bayesian Regression Models using Stan ver 0.10.0) to fit hierarchical Bayesian regression models. For each article viewed in the task, the consensus rating of article value (i.e., group mean rating) was defined as the average of all fMRI participant’s ratings of how likely they would be to read that article (after excluding the participant currently being considered). We fit models that used vmPFC activity as a predictor and used these consensus ratings of article value as an outcome, as well as analogous models that used vmPFC activity as a predictor and the logarithm of the number of times each New York Times article was shared online (from 34 to 12 740) as an outcome. The log transformation addresses the right skew in the distribution of raw sharing counts and yields a multiplicative model in which a unit difference in brain

activity shown in response to an article is associated with a percent difference in population sharing. All models incorporated “random effect” terms allowing model coefficients to vary from person to person, resulting in person-specific posterior estimates of the relationship between vmPFC activity and the outcome variable (sample consensus value or population article sharing). To ask whether vmPFC consensus value tracking and population behavior tracking varied as a function of news consumption, we tested for an interaction of (within-person) differences in vmPFC activity and (between-person) differences in news engagement in generating expected values of consensus value and population sharing. For all models, outcome and predictor variables were standardized, yielding standardized coefficients. Predictors that varied within-person were person-mean centered, yielding standardized coefficients indicating the average within-person relationship between the predictor and the outcome. This is distinct from an analytic approach that first averages brain and behavioral responses to each article (Scholz et al. 2017), which cannot be used to estimate the mean or variance of within-person relationships relating brain responses to individual or population behavior. We used 95% Bayesian credibility intervals (central posterior intervals) to convey a plausible range of values that a given effect could take in light of the observed data. All analyses took an estimation approach in that the goal was to generate plausible ranges for population parameters (“effect sizes”) and not to accept or reject point hypotheses. Further, analyses aimed to estimate population parameters but not to generate point predictions for an external set of new participants or new articles.

Because vague priors centered at zero yield inferences that are similar to traditional maximum likelihood estimates, we used a vague normal prior (location zero, standard deviation 1000) on β -coefficients (overall “fixed” terms for model intercepts and/or slopes) and a vague positive half-normal prior (location zero, standard deviation 1000) standard deviations (varying “random” terms for coefficient variation) (Stan Development Team 2016). Models were estimated with Markov Chain Monte Carlo (MCMC) sampling, running 4 parallel chains for 1000 iterations each (the first 500 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 modeling, this Bayesian estimation approach offers posterior inference, more accurate estimation of hierarchical variance parameters, better rates of convergence, and diagnostics for assessing the validity of the MCMC-based statistical inferences (Stan Development Team 2016).

Functional Connectivity Analyses

We conducted functional connectivity analyses to ask whether the capacity of vmPFC to track consensus value was related to connectivity between vmPFC and other regions of the brain. We used the Nilearn package (Abraham et al. 2014) to extract time-series during the period of the task in which participants were exposed to the articles. Data were detrended, standardized, and extracted from 8-mm radius spheres around the nodes defined above. Next, timeseries were wavelet transformed in Field-Trip (Oostenveld et al. 2011) and average wavelet coherence (0.0635–0.1562 Hz) was calculated in MATLAB (Mathworks, Inc.) to assess connectivity within a short timescale. Because person-to-person differences in head motion can artifactually influence measures of functional connectivity, we used mean

framewise head displacement as a covariate of no interest within all group-level connectivity analyses (Power et al. 2012; Van Dijk, Sabuncu, and Buckner 2012; Ciric et al. 2017).

To assess connectivity between vmPFC and other brain regions, we identified the 2 nodes of the Power atlas (nodes 107 and 109) closest in space to the meta-analytically identified vmPFC valuation peak (Bartra, McGuire, and Kable 2013) and computed estimates of connectivity of these nodes with the other 262 nodes. Next, we asked if there were any brain regions for which connectivity with the vmPFC nodes was associated with individual differences in 1) news reading, and 2) how closely vmPFC activity tracked with group consensus value ratings. To avoid overfitting, we used sparsity-promoting priors to regularize estimates toward zero (i.e., to make these analyses more conservative and increase the accuracy of the resulting estimates). Specifically, we used a Bayesian LASSO prior (a Laplace distribution centered at zero), with the scale of the prior estimated during model fitting as a hyperparameter (Park and Casella 2008; Stan Development Team 2016).

Results

Activity in vmPFC Tracked With Consensus Value and Population Article Sharing, With Variability From Person to Person

In an initial analysis, we fit a model that used vmPFC responses to article summaries as a predictor of group consensus ratings of article value, incorporating terms allowing the magnitude of this relationship to vary from person to person. We defined consensus value as the average of all participant ratings of reading intentions for that article (except for the participant currently being considered). This model revealed that, on average, within-person variation in vmPFC activity tended to track with consensus value ratings of news articles, $\beta = 0.16$, 95% CI [0.08, 0.24], and this relationship held after controlling for the participant’s own personal ratings of likelihood of reading the article, $\beta = 0.11$, 95% CI [0.04, 0.18]. Further, there was variability in this relationship from person to person, SD = 0.10, 95% CI [0.01, 0.22], such that some people showed vmPFC responses that more closely tracked consensus value whereas others showed less correspondence (Fig. 2A).

We next turned to the relationship between vmPFC activity and population article sharing. Consistent with our previously reported analyses (Scholz et al. 2017), vmPFC responses to New York times articles showed a within-person relationship with population sharing behavior on average, $\beta = 0.11$, 95% CI [0.03, 0.18], and this relationship held when controlling for trial-by-trial ratings of reading intentions, $\beta = 0.09$, 95% CI [0.001, 0.18]. (For further consideration of the relationship between behavioral ratings and population sharing, see Supplementary Materials.) There was also person-to-person variability in this relationship, SD = 0.10, 95% CI [0.01, 0.22] (Fig. 2B), indicating that some people showed vmPFC responses that tracked with population behavior more strongly than others.

To understand this variability, we asked whether individual differences in vmPFC consensus value tracking were related to individual differences in vmPFC population behavior tracking. Indeed, there was a positive relationship between these individual difference metrics: participants who showed vmPFC activity that corresponded with consensus value also tended to show vmPFC activity that tracked more closely with population sharing, $\beta = 0.40$, 95% CI [0.11, 0.64] (Fig. 2C). This indicates that how

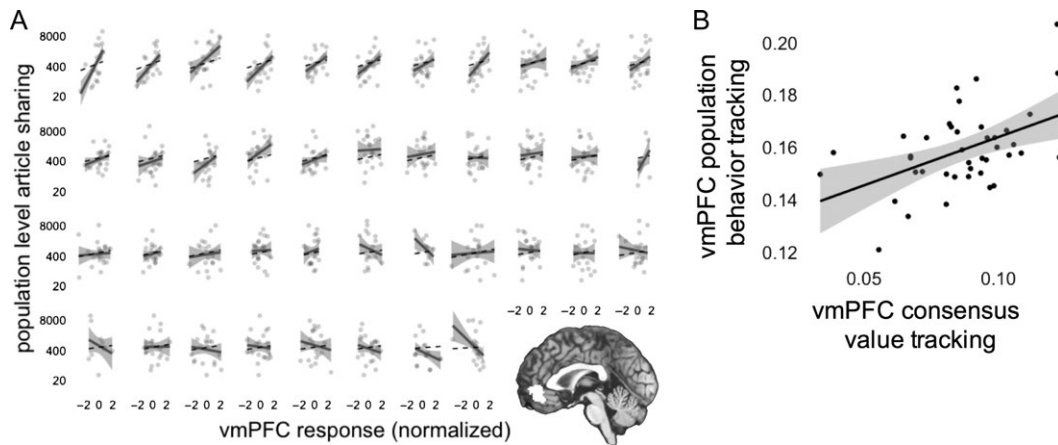


Figure 2. Individual differences in vmPFC tracking of consensus value and population behavior. (A) Multilevel models revealed person-to-person variability in how closely vmPFC responses tracked with population-level news article sharing (light gray lines and bands reflect models fit to each participant separately, with 80% CI; dotted black lines reflect multilevel estimates regularized toward group average). (B) Participants who tended to show vmPFC activity that corresponded more closely to consensus value (i.e., group average ratings of article reading intentions) tended to show vmPFC activity that more closely tracked with of population article sharing (data points reflect unstandardized person-specific coefficients estimated from a multilevel model).

well vmPFC activity tracked consensus value was strongly related to how closely it corresponded with population sharing behavior.

Infrequent News Readers Showed Lower Self-reported Reading Intentions and Lower Average vmPFC Activity

Our initial analyses revealed individual differences in how accurately vmPFC responses tracked consensus value, and further showed that these differences were related to how accurately vmPFC activity tracked with population behavior. However, these analyses could not speak to the question of what psychological factors are responsible for bringing these differences about. To address this, we next asked whether vmPFC population behavior tracking related to daily life experience with news content similar to the articles viewed in the scanner task. We found that participants who reported more frequently reading news tended to show greater average vmPFC responses, $\beta = 0.35$, 95% CI [0.05, 0.65] (Fig. 3, top left). Similarly, news reading frequency was associated with higher intentions to read the articles viewed in the scanner, $\beta = 0.31$, 95% CI [0.02, 0.60]. Together, these results suggest that people who reported more frequently engaging with similar kinds of articles in daily life showed greater neural and behavioral valuation of the news article stimuli shown in the scanner.

Infrequent News Readers Showed Greater vmPFC Population Behavior Tracking, Mediated by Superior Tracking of Consensus Value

Having seen that frequent news readers showed greater valuation of these articles, we next asked whether news reading was also related to how closely vmPFC activity tracked with population article sharing and consensus value ratings. We found an interaction of (within-person variation in) vmPFC activity with (between-person variation in) news reading, $\beta = -0.12$, 95% CI [-0.22, -0.03], such that vmPFC responses tended to most closely track population sharing behavior for participants who reported less frequently reading similar news articles in daily life. Put another way, the estimated person-specific slopes for the relationship between vmPFC activity and population sharing were negatively correlated with frequency of engagement

with similar news articles, $\beta = -0.41$, 95% CI [-0.64, -0.12] (Fig. 3, top right).

To visualize and communicate this continuous interaction, we estimated vmPFC activity and tracking effects separately for participants in the top 33% of news engagement (i.e., frequent readers) and those in the bottom 33% (i.e., infrequent readers). This revealed that frequent readers showed high vmPFC activity to the articles overall (relative to baseline fixation), +0.06% signal change, 95% CI [-0.09, 0.20], but showed article-to-article differences in vmPFC activity that were not strongly related to population sharing, $\beta = -0.003$, 95% CI [-0.14, 0.13]. Infrequent readers, however, showed low valuation to articles in general, -0.16% signal change, 95% CI [-0.25, -0.07], but showed article-to-article differences that were more strongly related to population sharing, $\beta = 0.22$, 95% CI [0.06, 0.38]. Overall, these results indicate that more frequent readers tended to show generally high valuation activity to all articles, whereas more infrequent readers tended to show high valuation activity only to highly shared articles (Fig. 3, bottom).

If frequent news readers show vmPFC activity that is less related to population behavior, this might reflect the fact that their brain valuation responses are not closely tracking consensus value. Supporting this idea, we also found an interaction between news reading and vmPFC activity, $\beta = -0.09$, 95% CI [-0.18, -0.01], such that article-to-article differences in vmPFC activity most closely corresponded with consensus value for less frequent readers. Intuitively, this result means that person-specific estimates for how closely vmPFC activity tracked with of consensus value were also negatively correlated with news reading frequency, $\beta = -0.39$, 95% CI [-0.63, -0.08]. Moreover, as shown in Figure 3, the relationship between higher news reading and lower vmPFC population tracking capacity was mediated by lower vmPFC tracking of consensus value, a'b path = -0.12, 95% CI [-0.29, -0.02]. When adjusting for how well vmPFC tracked consensus value, the relationship between news reading and vmPFC population behavior tracking dropped in magnitude, c' path = -0.29, 95% CI [-0.59, 0.00]. Thus, the data were consistent with a model in which the between-person relationship between news reading frequency and lower population behavior tracking is mediated by lower vmPFC tracking of consensus value.




Figure 3. News reading frequency moderated vmPFC tracking of population behavior. (A) Participants who reported frequently reading news in daily life tended to have higher vmPFC responses to the news articles on average (y axis reflects average vmPFC activity). (B) Infrequent readers tended to have better vmPFC tracking of population article sharing (y axis reflects vmPFC population behavior tracking). (C) The overall pattern of data indicated that news consumption frequency showed a moderating effect; frequent news readers (top 33%) tended to show high ventromedial prefrontal cortex (vmPFC) responses across all articles, whereas infrequent readers (bottom 33%) tended to show high vmPFC only to articles that would be heavily shared. (D) The person-to-person relationship between news reading and better vmPFC tracking of population sharing was mediated by better vmPFC tracking of consensus value ratings. See also Supplemental Figure S4.

Infrequent News Readers Tended to Show Lower vmPFC–dlPFC Connectivity, Which was Associated With Superior Consensus Value Tracking

Our analyses revealed that vmPFC tracking of population behavior was disproportionately driven by infrequent news readers, whose vmPFC activity was more representative of group-level consensus judgments of article value. Because vmPFC activity is thought to integrate diverse sources of information into a summary valuation, we next examined whether news reading experience was related to functional connectivity between vmPFC and other brain regions.

Using a whole-brain atlas (Power et al. 2011), we estimated functional connectivity between vmPFC and all other brain regions. We then asked if any brain regions showed individual differences in connectivity with vmPFC that were related to both less frequent news reading, and greater vmPFC tracking of consensus value. We thresholded our results such that we report only relationships for which the 95% credibility interval under a Bayesian LASSO prior excluded zero. This analysis revealed that participants who reported less frequent news reading tended to show lower connectivity of vmPFC with several brain regions, including dorsolateral PFC (dlPFC), medial PFC, lateral temporal cortex, premotor cortex, and insula (Fig. 4A). It also revealed that lower connectivity of vmPFC with dlPFC and occipital cortex was associated with better vmPFC tracking of consensus value (Fig. 4B). The overlap of these analyses identified a single node within dlPFC for which lesser connectivity with vmPFC was associated with both less frequent news reading and better vmPFC consensus value tracking. Overall, this pattern of results indicates that infrequent news readers showed vmPFC activity that was more independent

from activity within a lateral PFC region associated with top-down control of cognition and emotion (Miller and Cohen 2001; Ochsner, Silvers, and Buhle 2012), a pattern of connectivity that was also associated with superior vmPFC tracking of consensus value.

News Reading, vmPFC Population Behavior Tracking, and Consensus Value Tracking Were not Strongly Related to Other Demographics

Finally, we conducted follow-up analyses to ask whether vmPFC tracking of consensus value tended to vary with other person-level demographic characteristics. We did not find clear evidence for a moderating influence of gender, education, or race on how closely vmPFC activity tracked consensus value or tracked population article sharing (95% CIs showed substantial overlap with zero). However, we saw some evidence for a moderating influence of age, $\beta = 0.09$, 95% CI [0.01, 0.17], on how closely vmPFC activity tracked population article sharing: older participants (those closer to the upper limit of age 24) showed somewhat stronger vmPFC population behavior tracking. However, news consumption was not strongly related to age, $\beta = -0.06$, 95% CI [-0.38, 0.25], and the estimated moderating influence of news engagement held, $\beta = -0.11$, 95% CI [-0.18, -0.03], when adjusting for potential moderating influences of age, gender, race, and education.

Discussion

Information sharing can profoundly shape society, but the brain processes that underlie this phenomenon are not well understood. Here we asked 1) if there are individual differences




Figure 4. News reading and vmPFC consensus value tracking were associated with lower levels of vmPFC–dlPFC connectivity. (A) Less frequent news reading was associated with lower connectivity of vmPFC with several brain regions, including dlPFC, medial PFC, lateral temporal cortex, premotor cortex, and insula. (B) Higher vmPFC value tracking was associated with lower connectivity of vmPFC with dlPFC and occipital cortex. (C) vmPFC and dlPFC nodes for which lower connectivity was associated with both less news reading and higher vmPFC value tracking. (Regions visualized reflect those for which 95% credibility intervals under a LASSO prior excluded zero.)

in how closely vmPFC activity tracks consensus judgments of information value, and 2) whether these differences relate to how well vmPFC activity corresponds with population-level information diffusion. Further, we investigated potential sources of this variability by asking 3) if variability in vmPFC value tracking could be explained by daily life news experience, and 4) if vmPFC value tracking related to functional connectivity of vmPFC with other brain regions. To address these questions, we built multilevel models leveraging variability in brain responses to New York Times articles and relating this variability to sharing of these articles within the broader population of readers.

Our results can be summarized with 4 key findings. First, there were substantial individual differences in how closely vmPFC activity tracked consensus judgments of article value. Second, these differences in consensus value tracking were closely related to how accurately vmPFC activity corresponded with population article sharing. Third, vmPFC consensus value tracking was related to daily life news reading: frequent news readers tended to show generally high valuation activity to all articles, whereas infrequent news readers tended to show high valuation activity only to articles that were highly valued and heavily shared by others. Fourth, both less frequent news reading and more accurate vmPFC tracking of consensus value were associated with lower vmPFC connectivity with a dlPFC region implicated in the top-down control of cognition and emotion (Miller and Cohen 2001; Ochsner, Silvers and Buhle 2012).

Implications for Models of Information Value and Viral Sharing

The results of this study extend prior models of information sharing by demonstrating the population relevance of individual differences in valuation-related brain activity. When a news article is at the highest levels of population retransmission (i.e., it has “gone viral”), it has by definition reached and been shared by a larger audience than the people who typically read news of its kind (Berger and Iyengar 2013). Previous studies of viral sharing have described and predicted the diffusion of information (Goel et al. 2015; Heimbach et al. 2015; Kim

2015), but have not focused on the psychological and brain mechanisms underlying this phenomenon.

Within the brain-as-predictor framework, neuroimaging studies have asked whether aggregate brain activity from a small group of perceivers can track with population behavior (for a review, see: Falk and Scholz 2018; Genevsky and Knutson 2018) but have not considered diversity across individuals in brain activity or processing dynamics. Further, previous work from our lab shows that activity in brain regions associated with self-, social-, and value-related processing is increased when people make sharing judgments and tracks with decisions to share and consume information (Baek et al. 2017). The current investigation reveals that, in response to the same media information, perceivers show variability in evoked brain activity that affects how well their brain activity can track with population sharing. Furthermore, people with less experience with the class of media we focused on tended to show better vmPFC tracking of consensus value and better tracking of media effects in the population. Moreover, the same people also showed lower connectivity of vmPFC with brain regions involved in controlled processing. This pattern suggests that vmPFC is most related to population behavior when it acts more independently from regions involved in top-down cognitive processing. Future work could extend these results by asking whether vmPFC population behavior tracking can be enhanced with specific kinds of training or instructions, or if it varies according to individual differences in cognition or social network-based variables (Falk and Bassett 2017). A further implication of this work is that whereas a traditional marketing approach might seek out participants who are target experts in a particular domain (e.g., frequent New York Times readers), we show that the brains of such experts were least related to population level outcomes. This may be due to the fact that by definition, to go viral, information has to appeal to readers who are not the most frequent readers or experts. An approach that exclusively sampled these experts might be able to provide specialized insights with some methods, but might diverge from the population that defines virality in important ways (e.g., these experts may have tastes or manners of approaching the material that are idiosyncratic with respect to the broader population).

Implications for Integrative Valuation Theories of vmPFC

In the cognitive neuroscience literature, theories of vmPFC activity have emphasized this region's role in computing the value of stimuli for oneself (Rangel and Hare 2010; Bartra, McGuire, and Kable 2013). However, previous studies have shown that vmPFC activity can be influenced by shifting goal states, as when asked to focus on the health versus taste value of unhealthy foods (Hare et al. 2011), and by the influence of an immediate peer group, as when learning that your peers tend to value unhealthy food differently than you do (Nook and Zaki 2015). Here we show that vmPFC activity can track how stimuli tend to be valued by other people, and, to the extent that it does so, it can track stimulus effects in the population. Moreover, vmPFC consensus value tracking was moderated by daily life news experience, consistent with models positing that the vmPFC acts to integrate diverse neural inputs into a summary value signal that is sensitive to life history and motivational context (Roy et al. 2012). In follow-up analyses investigating ventral striatum, another core region implicated in computing the expected value of stimuli (Haber and Knutson 2010), we found that ventral striatal activity also tracked with population behavior and consensus ratings of value, showing variance from person-to-person in the extent of these relationships. However, this striatal person-to-person variance did not clearly correspond with news reading frequency as did vmPFC (see Supplementary Materials).

Broadly, these data support the notion that value is not exclusively an inherent property of a stimulus, but the result of an appraisal process that relies on communication across brain systems associated with controlled processing, perceptual representation, and integrative valuation (Cunningham and Zelazo 2007; Rangel et al. 2008; Scherer et al. 2001). In particular, vmPFC activity corresponded more closely with consensus value ratings when it exhibited lower connectivity with dlPFC, suggesting that top-down influences on the value signal may push valuation to be more person-specific and less generalizable. Previous work has shown that vmPFC-dlPFC connectivity tends to be higher during decisions that require self-control, suggesting that dlPFC plays a role in representing abstract or goal-driven inputs to valuation (Hare et al. 2009; Kable and Glimcher 2010; Hare et al. 2011). For frequent news readers, it may be that vmPFC-dlPFC connectivity reflects top-down control of the value signal, whereas less frequent readers may show valuation that is based on more immediate and less controlled responses that are more generalizable and less person-specific. Future work could ask whether cognitive goal-based manipulations or direct physical manipulation of brain activity can alter the extent to which brain activity tracks with population behavior, and ask whether such manipulations differ in their effects across vmPFC, ventral striatum, and other brain systems. Further, future studies could also attempt to characterize heterogeneity within a measured population in order to ask whether sampling from specific subpopulations of individuals (with specific kinds of experience or interests) can lead to more efficient prediction of the behavior of those subpopulations.

Conclusion

If we can use brain activity as a window into large-scale behavior, whose brains provide the clearest view? Here we suggest that population information sharing is seen for stimuli eliciting

greater brain valuation responses amongst people who are not frequent consumers of the kind of information being shown. Our data suggest this may arise because the brain valuation responses of these individuals tend to reflect broader consensus about information value, which in turn is linked to lower levels of connectivity between valuation- and control-related brain systems. These findings lay the foundation for a mechanistic and prospectively predictive understanding of how and why information can diffuse (or fail to diffuse) across a population of individuals.

Supplementary Material

Supplementary material is available at *Cerebral Cortex* online.

Funding

Defense Advanced Research Projects Agency (D14AP00048; P.I.F.), the National Institutes of Health in form of a New Innovator Award (NIH 1DP2DA03515601; P.I.F.), a National Cancer Institute and FDA Center for Tobacco Products (CTP) grant (P50CA179546), the US Army Research Laboratory (including work under Cooperative Agreement Number W911NF-10-2-0022) and HopeLab. D.S.B. would also like to acknowledge support from the Alfred P. Sloan Foundation and the John D. and Catherine T. MacArthur Foundation. The views, opinions, and/or findings contained in this article are those of the authors and should not be interpreted as representing the official views or policies, either expressed or implied, of the Defense Advanced Research Projects Agency, the Department of Defense, the NIH, or the Food and Drug Administration.

Notes

The authors thank Hyun Suk Kim and Joseph Cappella for provision of population-level data. *Conflict of Interest:* None declared.

References

- Abraham A, Pedregosa F, Eickenberg M, Gervais P, Mueller A, Kossaiji J, Gramfort A, Thirion B, Varoquaux G. 2014. Machine learning for neuroimaging with scikit-learn. *Front Neuroinform.* 8:14.
- Baek EC, Scholz C, O'Donnell MB, Falk EB. 2017. The value of sharing information: a neural account of information transmission. *Psychol Sci.* 28:851–861.
- 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.
- Berger J, Iyengar R. 2013. Communication channels and word of mouth: how the medium shapes the message. *J Consum Res.* 40:567–579.
- Berns GS, Moore SE. 2012. A neural predictor of cultural popularity. *J Consum Psychol.* 22:154–160.
- Chua HF, Ho SS, Jasinska AJ, Polk TA, Welsh RC, Liberzon I, Strecher VJ. 2011. Self-related neural response to tailored smoking-cessation messages predicts quitting. *Nat Neurosci.* 14:426–427.
- Ciric R, Wolf DH, Power JD, Roalf DR, Baum GL, Ruparel K, Shinohara RT, Elliott MA, Eickhoff SB, Davatzikos C, et al. 2017. Benchmarking of participant-level confound regression strategies for the control of motion artifact in studies of functional connectivity. *Neuroimage.* 154:174–187.

- Cunningham WA, Zelazo PD. 2007. Attitudes and evaluations: a social cognitive neuroscience perspective. *Trends Cogn Sci*. 11:97–104.
- Falk EB, Bassett DS. 2017. Brain and social networks: fundamental building blocks of human experience. *Trends Cogn Sci*. 21:674–690.
- Falk EB, Berkjman ET, Lieberman MD. 2012. From neural responses to population behavior: neural focus group predicts population-level media effects. *Psychol Sci*. 23:439–445.
- Falk EB, Berkman ET, Mann T, Harrison B, Lieberman MD. 2010. Predicting persuasion-induced behavior change from the brain. *J Neurosci*. 30:8421–8424.
- Falk EB, Cascio CN, Coronel JC. 2015. Neural prediction of communication-relevant outcomes. *Commun Methods Meas*. 9:30–54.
- 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.
- Goel S, Anderson A, Hofman J, Watts DJ. 2015. The structural virality of online diffusion. *Manage Sci*. 62:180–196.
- Haber SN, Knutson B. 2010. The reward circuit: linking primate anatomy and human imaging. *Neuropsychopharmacology*. 35:4–26.
- Hare TA, Camerer CF, Rangel A. 2009. Self-control in decision-making involves modulation of the vmPFC valuation system. *Science*. 324:646–648.
- 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.
- Heimbach I, Schiller B, Strufe T, Hinz O. 2015. Content virality on online social networks: empirical evidence from Twitter, Facebook, and Google+ on German News Websites. In: *Proceedings of the 26th ACM Conference on Hypertext & Social Media*. HT'15. New York, NY, USA: ACM. p. 39–47.
- Kable JW, Glimcher PW. 2010. An “As Soon As Possible” effect in human intertemporal decision making: behavioral evidence and neural mechanisms. *J Neurophysiol*. 103:2513–2531.
- Kim HS. 2015. Attracting views and going viral: how message features and news-sharing channels affect health news diffusion. *J Commun*. 65:512–534.
- Miller EK, Cohen JD. 2001. An integrative theory of prefrontal cortex function. *Annu Rev Neurosci*. 24:167–202.
- Nook EC, Zaki J. 2015. Social norms shift behavioral and neural responses to foods. *J Cogn Neurosci*. 27:1412–1426.
- 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.
- Oostenveld R, Fries P, Maris E, Schoffelen J-M. 2011. FieldTrip: Open source software for advanced analysis of MEG, EEG, and invasive electrophysiological data. *Comput Intell Neurosci*. 2011:156869.
- Park T, Casella G. 2008. The Bayesian Lasso. *J Am Stat Assoc*. 103:681–686.
- Power JD, Barnes KA, Snyder AZ, Schlaggar BL, Petersen SE. 2012. Spurious but systematic correlations in functional connectivity MRI networks arise from subject motion. *Neuroimage*. 59:2142–2154.
- Power JD, Cohen AL, Nelson SM, Wig GS, Barnes KA, Church JA, Vogel AC, Laumann TO, Miezin FM, Schlaggar BL, et al. 2011. Functional network organization of the human brain. *Neuron*. 72:665–678.
- Rangel A, Camerer C, Montague PR. 2008. A framework for studying the neurobiology of value-based decision making. *Nat Rev Neurosci*. 9:545–556.
- Rangel A, Hare T. 2010. Neural computations associated with goal-directed choice. *Curr Opin Neurobiol*. 20:262–270.
- Roy M, Shohamy D, Wager TD. 2012. Ventromedial prefrontal-subcortical systems and the generation of affective meaning. *Trends Cogn Sci*. 16:147–156.
- Scherer KR, Schorr A, Johnstone T. 2001. *Appraisal Processes in Emotion: Theory, Methods, Research*. Oxford University Press.
- Scholz C, Baek EC, O'Donnell MB, Kim HS, Cappella JN, Falk EB. 2017. A neural model of valuation and information virality. *Proc Natl Acad Sci U S A*. 114:2881–2886.
- Team SD, Others. 2016. RStan: the R interface to Stan. R package version 2.14.1.
- Van Dijk KRA, Sabuncu MR, Buckner RL. 2012. The influence of head motion on intrinsic functional connectivity MRI. *Neuroimage*. 59:431–438.
- Venkatraman V, Dimoka A, Pavlou PA, Vo K, Hampton W, Bollinger B, Hershfield HE, Ishihara M, Winer RS. 2015. Predicting Advertising Success Beyond Traditional Measures: New Insights from Neurophysiological Methods and Market Response Modeling. *J Mark Res*. 52:436–452.
- Wang A-L, Ruparel K, Loughhead JW, Strasser AA, Blady SJ, Lynch KG, Romer D, Cappella JN, Lerman C, Langleben DD. 2013. Content matters: neuroimaging investigation of brain and behavioral impact of televised anti-tobacco public service announcements. *J Neurosci*. 33:7420–7427.