

Highly accurate prediction of emotions surrounding the attacks of September 11,
2001 over 1-, 2-, and 7-year prediction intervals

Doré, B.P.¹, Meksin, R.², Mather, M.³, Hirst, W.², & Ochsner, K.N.¹.

¹Department of Psychology, Columbia University

²Department of Psychology, New School for Social Research

³Davis School of Gerontology, University of Southern California

Corresponding authors:

Bruce Doré

Kevin Ochsner

Department of Psychology

Columbia University

Schermerhorn Hall

New York, NY 10027, USA

brucedore@gmail.com

ochsner@psych.columbia.edu

Abstract

In the aftermath of a national tragedy, important decisions are predicated on judgments of the emotional significance of the tragedy in the present and future. Research in affective forecasting has largely focused on ways in which people fail to make accurate predictions about the nature and duration of feelings experienced in the aftermath of an event. Here we ask a related but understudied question: can people forecast how they will feel in the future about a tragic event that has already occurred? We found that people were strikingly accurate when predicting how they would feel about the September 11 attacks over 1-, 2-, and 7-year prediction intervals. Although people slightly under- or overestimated their future feelings at times, they nonetheless showed high accuracy in forecasting 1) the overall intensity of their future negative emotion, and 2) the relative degree of different types of negative emotion (i.e., sadness, fear, or anger). Using a path model, we found that the relationship between forecasted and actual future emotion was partially mediated by current emotion and remembered emotion. These results extend theories of affective forecasting by showing that emotional responses to an event of ongoing national significance can be predicted with high accuracy, and by identifying current and remembered feelings as independent sources of this accuracy.

In the wake of an unprecedented tragedy like the September 11, 2001 attacks on New York and Washington D.C., can we know for how long and in what way our future emotions about the event will unfold? Beyond deepening our understanding of the psychology of prospection, if people are able to make such predictions with accuracy, there are ramifications for the behavior of individuals, the allocation of community resources, and the broader movements of our society.

In the past 15 years, the process of making a prediction about a future emotional state has emerged as a major topic in psychology, referred to as *affective forecasting* (Gilbert & Wilson, 2009; Miloyan & Suddendorf, 2015). In a typical experiment, participants predict how they will feel after they experience a future emotional event. Predominantly, such studies have focused on ways such predictions can be inaccurate, noting that beyond simple judgments of pleasantness or unpleasantness, people are typically poor at predicting how they will feel in the future (Ayton, Pott, & Elwakili, 2007; Wilson & Gilbert, 2005). For example, a consistent finding is that people overestimate the magnitude of their reactions to future emotional events, particularly for events they have little previous experience with (Wilson & Gilbert, 2003). One explanation for this finding, *immune neglect*, posits that people can be unaware that their reactions to emotional events are often quickly diminished through emotion regulation processes (Gilbert et al., 1998; Hoerger, 2012). Another explanation, *focalism*, posits that people tend to overweigh the influence of an event they are currently

attending to when making a prediction of future global affect (Levine, Lench, Kaplan, & Safer, 2012; Wilson et al., 2000; Wilson & Gilbert, 2013).

Although this work has provided many insights into how and why we make inaccurate predictions about certain classes of affective events, relatively less attention has been paid to the question of whether people are able to reach high levels of forecast accuracy for their future emotions in some instances, and, if so, what mechanisms they engage to achieve this accuracy. While this is in part a matter of emphasis (highlighting inaccuracy rather than accuracy), with a few notable exceptions (e.g., Dunn, Brackett, Ashton-James, Schneiderman, & Salovey, 2010; Lam, Buehler, McFarland, Ross, & Cheung, 2004; Levine, Lench, Kaplan, & Safer, 2012) the psychological mechanisms promoting accurate forecasts have received far less attention than those promoting inaccurate forecasts (for general discussion why accuracy should be a focus of study, see Zaki & Ochsner, 2011).

Here, we examined these questions with data from a 10-year longitudinal study of memory (see Hirst et al., 2015; Hirst et al., 2009), emotion, and predictions of future affect concerning the attacks of September 11, 2001, one of the most impactful public events in recent American history. Although prior literature in affective forecasting has considered a variety of emotional events, most studies ask participants to forecast reactions to events they have not yet experienced and ask about their global affective state rather than their reaction to a specific focal stimulus (i.e., how they feel when thinking about a specific emotional event). Moreover, no studies have quantified the accuracy of affective

forecasts for an event of continuing national significance over a multiple year prediction interval.

Two fundamental questions guided our investigation. First, we asked if people can accurately predict how they will feel about the September 11 attacks in the future. To assess accuracy, we followed recent suggestions (Mathieu & Gosling, 2012) to quantify both the absolute accuracy of participants' forecasts at the level of participants as a group (i.e., the difference between average forecasted and experienced emotion), as well as relative accuracy of forecasts from person to person (i.e., the correlation between forecasted and experienced emotion). In addition to the overall intensity of emotional reactions we also assessed the strength of specific types of emotions — sadness, fear, and anger — that are commonly elicited by traumatic events, and that have distinct implications for behavior (e.g., Lerner, Li, Valdesolo, & Kassam, 2015; Scherer, Schorr, & Johnstone, 2001).

Second, we sought to identify psychological mechanisms that contributed to the accuracy of emotional predictions, asking in particular whether current feelings and memories of past feelings can act as sources of prediction accuracy. Here we were motivated by the idea that predicting one's future response to a focal event of ongoing relevance and impact entails extrapolating about sustained emotional responses that unfold over time. Drawing from theories of mental prospection, we predicted that these extrapolations would be influenced both by remembered past emotional responses and knowledge about present emotional responses (Szpunar, 2010; Tulving, 1985).

Method

Participants and Design

Participants were 202 US adults (146F; 56M) who completed all four time points of a 10-year study examining cognition and emotion surrounding the September 11 attacks. The four time points of data collection were September 2001, August 2002, August 2004, and August 2011. At the final time point of the study, the average age of participants was 48 (SD=18).

The study was designed such that recruitment materials were sent to as many potential participants as was feasible (for additional recruitment information, see Hirst et al., 2009 and Hirst et al., 2015). The initial 2001 survey was sent out 3-5 days after the September 11 attacks, and participants were asked to complete and return the survey within 1 day of receiving it. Of the 2117 participants who completed the 2001 survey, approximately 10% completed and returned all four surveys such that they could be linked together, yielding a sample size of 202 for the purposes of this investigation. Given the length of the project, the extensive nature of the survey, and the fact that we did not compensate participants for their efforts, this return rate is reasonable. For additional analyses of 324 participants who completed only the 2001 and 2002 surveys, and 186 participants who completed only the 2001, 2002, and 2004 surveys, see Supplementary Materials, and Supplementary Figure S1.

For the features of age, gender, residence at the time of the attack, student membership at the time of the attack, political affiliation, race/ethnicity, and religion, there were no significant differences between those who completed

all four surveys and (a) those who completed only the 2001 survey (see Hirst et al., 2015). Finally, although participants were recruited from across the United States, our sampling should not be viewed as representative of the American public.

Within each time point of the study, participants reported on three broad features of their reactions to the attacks relevant to the current investigation: 1) the intensity of their current feelings of sadness, fear, anger, confusion, shock and frustration, 2) predictions of the intensity of their future sadness, fear, anger, confusion, shock and frustration at approximately the next time point in the study (1, 2, and 7 years for surveys in 2001, 2002, and 2004 respectively), and 3) their memory of the intensity of their past feelings of sadness, fear, anger, confusion, shock and frustration in the 2 weeks of September 2001 immediately following the attack (for 2002, 2004, and 2011 studies only). All of these ratings were made on 5-point scales (1- low to 5-high). Participants also reported on a variety of other topics not considered here, including their specific memories for the events of the September 11 attacks and the circumstances in which they learned about it. Each survey was 17 pages long and took about 45 minutes to complete; copies can be found at <http://911memory.nyu.edu>.

Analyses

For group analyses of the absolute accuracy of affective predictions, we considered effects of study time point (1,2,3,4), report type (predicted, experienced), and emotion type (sadness, fear, anger). We focused on the emotion categories sadness, fear, and anger, which were of a priori interest as

common responses to tragedy with distinct implications for cognition and behavior (e.g., Lerner, Li, Valdesolo, & Kassam, 2015; Scherer, Schorr, & Johnstone, 2001). However, for purposes of comparison we also conducted parallel analyses for confusion, shock and frustration.

In person-to-person analyses of relative accuracy, we quantified overall intensity of negative affect by averaging across all categories of negative emotion that participants reported on: sadness, fear, anger, confusion, shock, and frustration. We also computed scores for reported and predicted sadness, fear, anger ratings that were centered on the mean of all reported negative emotions (within condition and time point), yielding the extent (in raw scale units) to which a participant reported relatively more of a particular negative emotion type in comparison to others. For purposes of comparison, we also conducted parallel analyses for confusion, shock, and frustration. This centering was done to separate accuracy in forecasting relative levels of different kinds of emotions from accuracy in forecasting overall negative affect intensity (see Supplementary Materials for additional information).

We used Stata 13 (<http://stata.com/stata13/>) to implement OLS single-level regression models (using *regress*) and REML multilevel models (using *mixed*) that incorporated random effect terms for participant intercepts. We built a multilevel path model using generalized structural equation modeling (*gsem*) and computed indirect paths via delta method estimation (*nlcom*). We used unstandardized betas with 95% confidence intervals (95%CI) to quantify effect sizes, inferring a nonzero effect when the 95%CI excluded zero. To evaluate

omnibus effects of time, emotion type, or their interaction, we used chi-square (X^2) likelihood ratio tests to compare models including particular main effects or interactions to reduced models not including them. To quantify evidence in favor of the null (e.g., that a particular main effect or interaction is equal to zero) we used Bayesian Information Criterion (BIC) differences to approximate Bayes Factors (Kass & Raftery, 1995).

Results

Question 1: Can people accurately forecast how they will feel in the future about the September 11 attacks?

Group-level forecasts of future negative affect are largely accurate.

First we examined forecasted and experienced sadness, fear, and anger, in order to quantify the absolute accuracy of affective forecasts at the level of participants as a group. Negative affect across time, emotion type, and report type (experienced versus forecasted) is displayed in Figure 1. Considering experienced negative affect, we observed a main effect of time, $X^2=207.0$, such that intensity of these negative emotions decayed from 2001 to 2002, $b=-.49$, 95%CI [-.61, -.39], was relatively stable from 2002 to 2004, $b=-.01$, 95%CI [-.12, .10], and then decayed again from 2004 to 2011, $b=-.32$, 95%CI [-.43, -.21]. We also observed a main effect of emotion type, $X^2=479.2$, such that sadness was reported at the highest intensity, mean=3.6, 95%CI [3.52, 3.76], followed by anger, mean=2.97, 95%CI [2.84, 3.09], and fear, mean=2.50, 95%CI [2.37, 2.62]. Qualifying these main effects, we observed a time point by emotion type interaction, $X^2=23.2$, such that anger decayed most slowly over the course of the

10 year study, $b=-.16$, 95%CI $[-.23, -.10]$, followed by fear, $b=-.25$, 95%CI $[-.31, -.19]$, and sadness, $b=-.34$, 95%CI $[-.40, -.28]$.

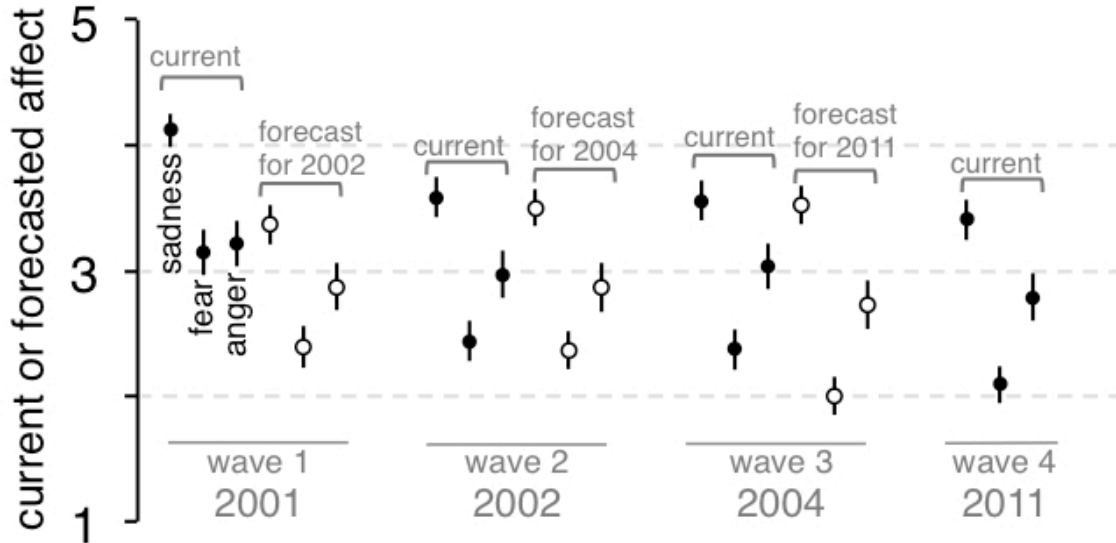


Figure 1. Current and forecasted intensity of sadness, fear, and anger in 2001, 2002, 2004, and 2011. Negative emotions decay slowly over time, and forecasts are largely accurate at the level of participants as a group. Group means with 95%CI.

Overall, affective forecasts were largely accurate. We observed no main effect difference between forecasted and experienced affect, $b=-.01$, 95%CI $[-.08, .05]$, standardized $\beta=-.004$, 95%CI $[-.03, .02]$, 176:1 odds in favor of the null of no difference. However, we did observe a time point by report type interaction, $X^2=15.0$. Follow-up analyses indicated that the difference between forecasted and experienced affect suggested slight underestimation of negative affect for 1-year forecasts (2001-2002), $b=-.13$ 95%CI $[-.24, -.01]$, no difference for 2-year (2002-2004) forecasts, $b=-.07$, 95%CI $[-.19, .05]$, and slight overestimation for 7-year (2004-2011) forecasts, $b=.16$, 95%CI $[.03, .29]$.

Participants were able to foretell different trajectories of change in anger, sadness, and fear. Within the data collected in 2001, we observed a condition (2001 experience versus 2001 prediction of 2002 experience) by emotion type interaction, $X^2=13.3$, such that people reported that they currently felt comparable levels of fear and anger, $b=-.07$, 95%CI [-.26, .11] but predicted that in 2002 they would feel higher levels of anger than fear, $b=.50$, 95%CI [.32, .68]. This foretold pattern was borne out in their 2002 reports of current sadness, fear, and anger, in that there was no evidence for an interaction between report type (2001 forecasted affect versus 2002 actual affect) and emotion type, $X^2=1.41$, 100:1 odds in favor of the null of no interaction, and experienced anger was indeed higher than experienced fear, $b=.52$, 95%CI [.39, .65]. Across all three forecasting intervals (1, 2, and 7 years), we observed no evidence for an interaction between report type (forecasted versus actual) and emotion type, $X^2=0.56$, 153:1 odds in favor of the null of no interaction. These results indicate that forecasts were accurate in that they corresponded with the overall decay of negative affect over time as well as different trajectories of decay for different types of emotion.

Of secondary interest, we also investigated absolute accuracy in forecasting confusion, shock, and frustration. First considering ratings of experienced emotion, we observed a main effect of time, $X^2= 90.1$, such that these emotions decayed from 2001 to 2002, $b=-.78$, 95%CI[-.96, -.60], stayed stable from 2002 to 2004, $b=-.01$, 95%CI[-.19, .17], and decayed again from 2004 to 2011, $b= -.17$, 95%CI[-.34, -.00]. We also observed a main effect of

emotion type, $X^2 = 161.0$, such that frustration was experienced at the highest intensity, mean = 2.76, 95%CI[2.62, 2.89], followed by shock, mean = 2.47, 95%CI[2.34, 2.60] and confusion, mean = 2.06., 95%CI[1.93, 2.19]. Qualifying these main effects, we also found a time point by emotion type interaction, $X^2 = 32.1$, such that shock decayed most quickly over the course of the 10 year study, $b = -.44$, 95%CI[-.51, -.37], whereas confusion, $b = -.26$, 95%CI[-.32, -.20], and frustration, $b = -.22$, 95%CI[-.28, -.15], decayed more slowly.

As with sadness, fear, and anger, forecasts were largely accurate. We observed no large differences between predicted and experienced affect for confusion, $b = .02$, 95%CI[-.09, .12], shock, $b = .04$, 95%CI[-.09, .17], or frustration, $b = -.09$, 95%CI[-.22, .03]. Collapsing across emotion type and time point, the data provided equivocal evidence for slight underestimation of these emotions (considered together), $b = -.07$, 95%CI[.00, -.15], 1.5:1 odds in favor of the null. As with sadness, fear, and anger, across all three prediction intervals (1, 2, and 7 years) we did not find an interaction between report type (forecasted versus actual) and emotion type (confusion, shock, frustration), $X^2 = 0.41$, 4675:1 odds in favor of the null of no interaction (i.e., forecasts and actual future affect show equivalent relative balance of these different emotion types) (see Supplementary Figure S2).

Forecasts are highly predictive of the overall intensity of future negative affect. Although our initial analyses indicated high absolute accuracy at the group level, they were not informative about whether, from person to person, forecasts of emotion are predictive of (correlated with) future emotion reports

(i.e., whether the forecasts show relative accuracy). In order to address this question, we conducted regression analyses to ask whether we could predict individual differences in the intensity of negative affect with differences in the predictions of intensity participants made 1, 2, or 7 years earlier. We found that

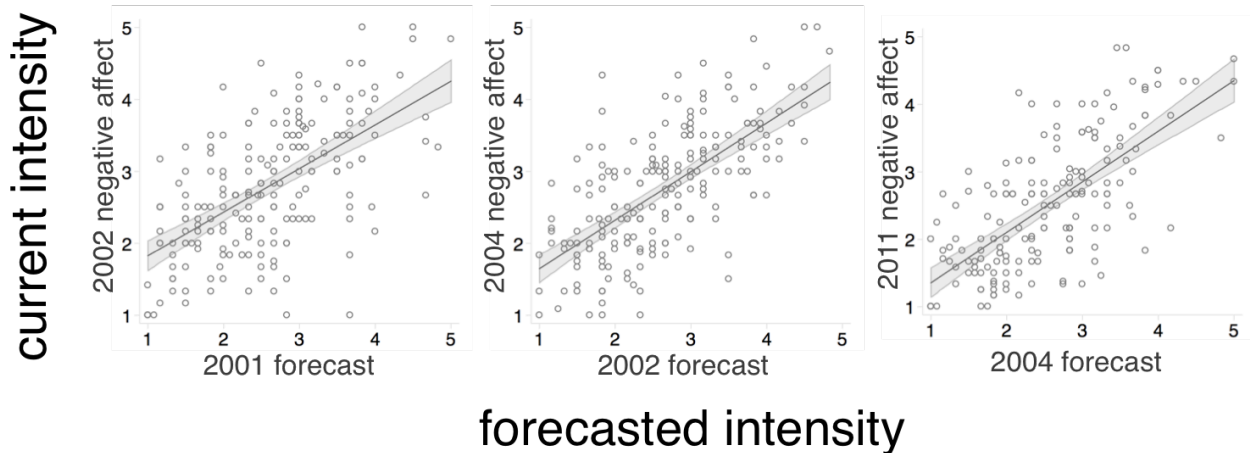


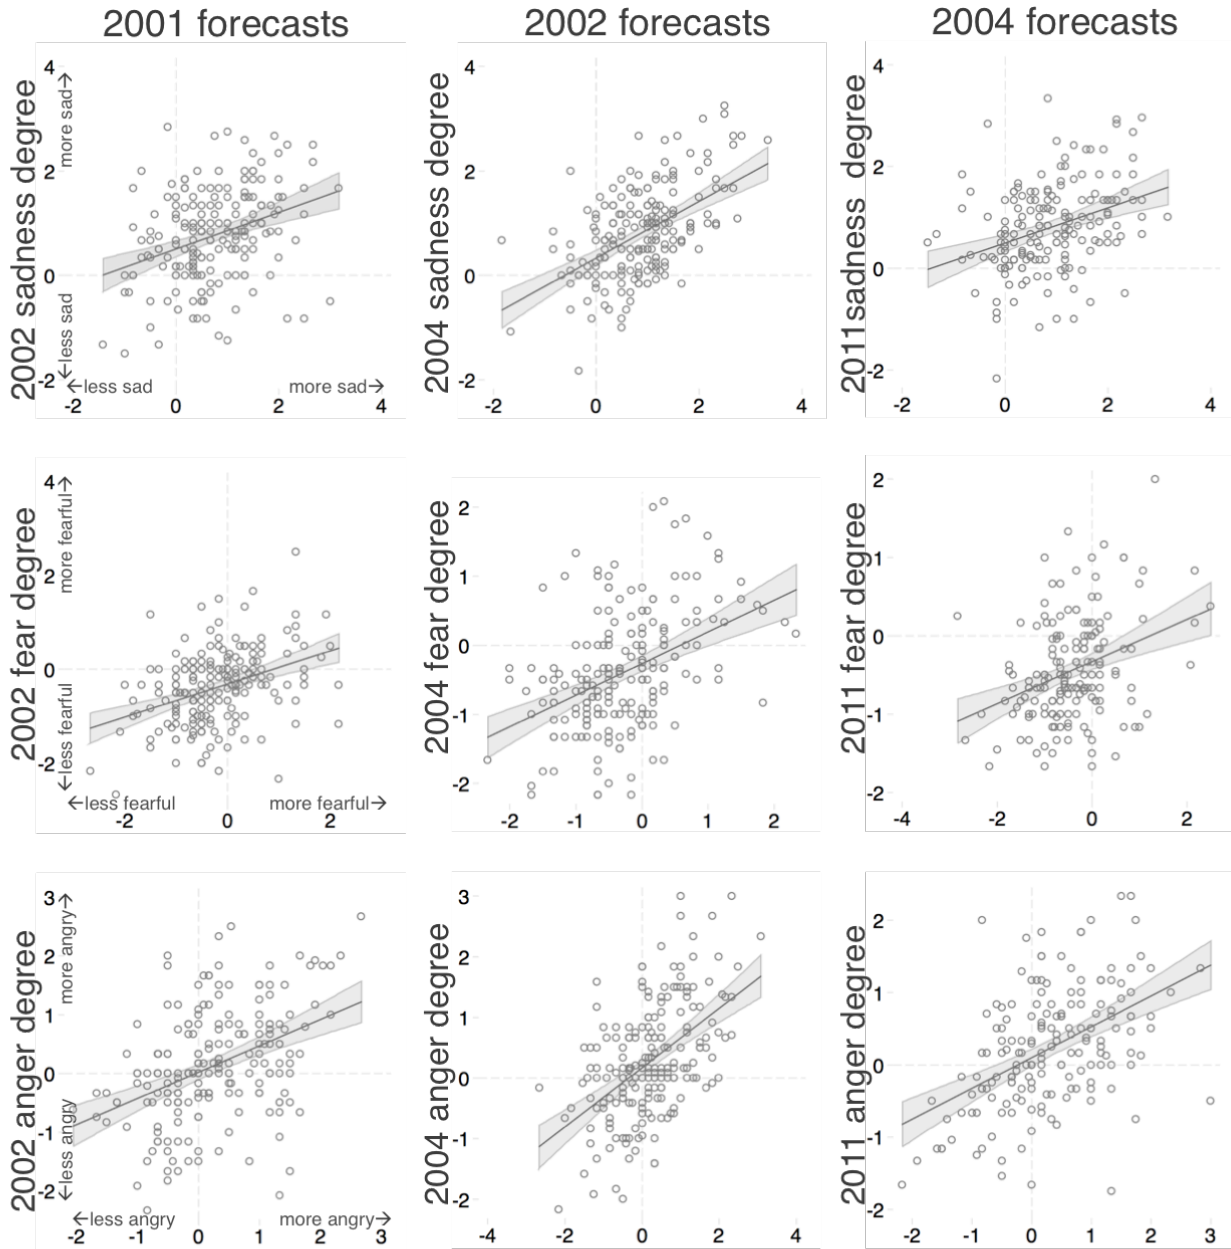
Figure 2. Relationship between the overall intensity of negative emotions reported in response to the attacks in 2002, 2004, and 2011 (on the y axes) and forecasts of negative emotion intensity made 1-, 2-, and 7-years earlier (on the x axes). Forecasts show a high degree of relative accuracy: people who predict they will feel the most negative tend to later report feeling the most negative.

forecasts of negative affect intensity were strongly predictive of experienced negative affect intensity for 1-year (2001 to 2002), $b=.61$, 95%CI [.49, .72], 2-year (2002 to 2004), $b=.68$, 95%CI [.57, .78], and 7-year (2004 to 2011) forecasting intervals, $b=.84$, 95%CI [.76, .92] (see Figure 2). To examine the possibility that the predictive value of affective forecasts could be attributed to a relationship between future affect and current affect, or between future affect and memories of past affect, we ran models that controlled for these variables. We found that the predictive relationship between forecasted and experienced affect decreased in magnitude but remained different from zero when adjusting for reports of

current affect and memories of affect experienced in September 2001 collected at the same time point as the predictions (see Supplementary Figure S3), suggesting that affective forecasts explain variance in future affect that is independent from the variance explained by current affect and remembered affect.

Forecasts are highly predictive of the relative degree of future sadness, fear, and anger. Even if people are able to accurately foretell the overall intensity of their future negative experience, they may be inaccurate in predicting the degree to which they will feel different specific types of negative emotion. To address this issue, we conducted another series of regression analyses that used forecasts of sadness, anger, and fear to predict future ratings of sadness, anger, and fear. Forecasts of sadness, anger and fear were centered on the mean of all predicted negative emotion categories (yielding the extent to which a participant predicted that he or she would feel relatively more sadness, anger, or fear than all other negative emotions) and used as predictor variables, and experienced sadness, anger, and fear were mean-centered in the same manner and used as outcome variables. This mean centering takes into account that a raw rating for a particular type of emotion (e.g., sadness rating = 3) has different relative meaning if other negative emotions are rated low (e.g., 2) or high (e.g., 4) on average (i.e., the 3 rating could indicate either low or high relative sadness) (see Supplemental Materials for additional information). For sadness, anger, and fear, we found that forecasted degree was highly predictive

current relative degree of specific emotion



forecasted relative degree of specific emotion

Figure 3. Relationships between the relative degree of specific types of emotion reported in response to the attacks in 2002, 2004, and 2011 (i.e. reports of particular categories centered on the mean of all reported negative emotion, on the y axes) and forecasts of relative degree made 1-, 2-, and 7-years earlier (i.e. forecasts mean-centered in the same manner, on the x axes). Forecasts show a high degree of relative accuracy – people who predict that they will feel predominantly sad, angry, or fearful tend to later end up doing so.

of future experienced degree of these emotions (see Figure 3). Forecasts of sadness were strongly predictive of experienced sadness for 1-year, $b=.35$, 95%CI [.22, .49], 2-year, $b=.54$, 95%CI [.42, .66], and 7-year forecasting intervals, $b=.34$, 95%CI [.20,.49]. Similarly, forecasts of fear were strongly predictive of experienced fear for 1-year, $b=.35$, 95%CI [.23, .47], 2-year, $b=.46$, 95%CI [.32, .60], and 7-year prediction intervals, $b=.27$, 95%CI [.15,.38]. Finally, forecasts of anger were also strongly predictive of experienced anger for 1-year $b=.45$, 95%CI [.31, .59], 2-year, $b=.49$, 95%CI [.37, .60], and 7-year forecasting intervals, $b=.42$, 95%CI [.31, .54]. These relationships decreased in magnitude but remained different from zero when adjusting for reports of current relative degree of emotion and remembered 2001 relative degree of emotion collected at the same time point as the affective predictions (see Supplementary Figure S4).

Of secondary interest, we also estimated relative accuracy for confusion, shock, and frustration. Forecasts of confusion were strongly predictive of experienced confusion for 1-year, $b=.41$, 95%CI [.30, .52], 2-year, $b=.56$, 95%CI [.44, .69], and 7-year intervals, $b=.33$, 95%CI [.20, .46]. Similarly, forecasts of shock were strongly predictive of experienced shock for 1-year $b=.32$, 95%CI [.21, .44], 2-year, $b=.40$, 95%CI [.27, .53], and 7-year intervals, $b=.44$, 95%CI [.31, .57]. Finally, forecasts of frustration were strongly predictive of experienced frustration for 1-year, $b=.38$, 95%CI [.25, .50], 2-year, $b=.42$, 95%CI [.29, .56], and 7-year prediction intervals, $b=.45$, 95%CI [.32,.58].

Question 2: Do current affect and memories of past affect act as sources of accuracy for forecasting future affect?

The relationship between forecasted and experienced negative affect is partially mediated by current and remembered negative affect. The observation that forecasts of future emotional experience are largely accurate leaves open the mechanistic question of how people are able to achieve this accuracy. We used a multilevel modeling approach to shed light on this question, building models that incorporate information about the relationship between forecasted and experienced negative affect at both within- and between-person levels. Across the entire dataset, we found a strong relationship, or *total path*, between predicted and experienced negative affect, $b=.59$, 95%CI [.53, .66], standardized $\beta=.57$, 95%CI [.50, .64] (see Figure 4).

Because theories of prospection implicate current and remembered experiences as sources of prospective ability (e.g., Szpunar, 2010; Tulving, 1985), we next ran a multilevel path model to test current negative affect and memory of negative affect felt in the immediate wake of the September 11 attacks as independent mediators of the predictive relationship between forecasted and experienced negative affect (see Figure 4). The model included forecasted negative affect as a predictor, current negative affect and remembered negative affect (available for time points 2, 3, and 4) as parallel mediators, and future experienced negative affect as an outcome variable. To account for the relationship between current and remembered affective experience, the model also included a path to estimate covariance between current and remembered affect ratings. The results of this model indicate 1) that

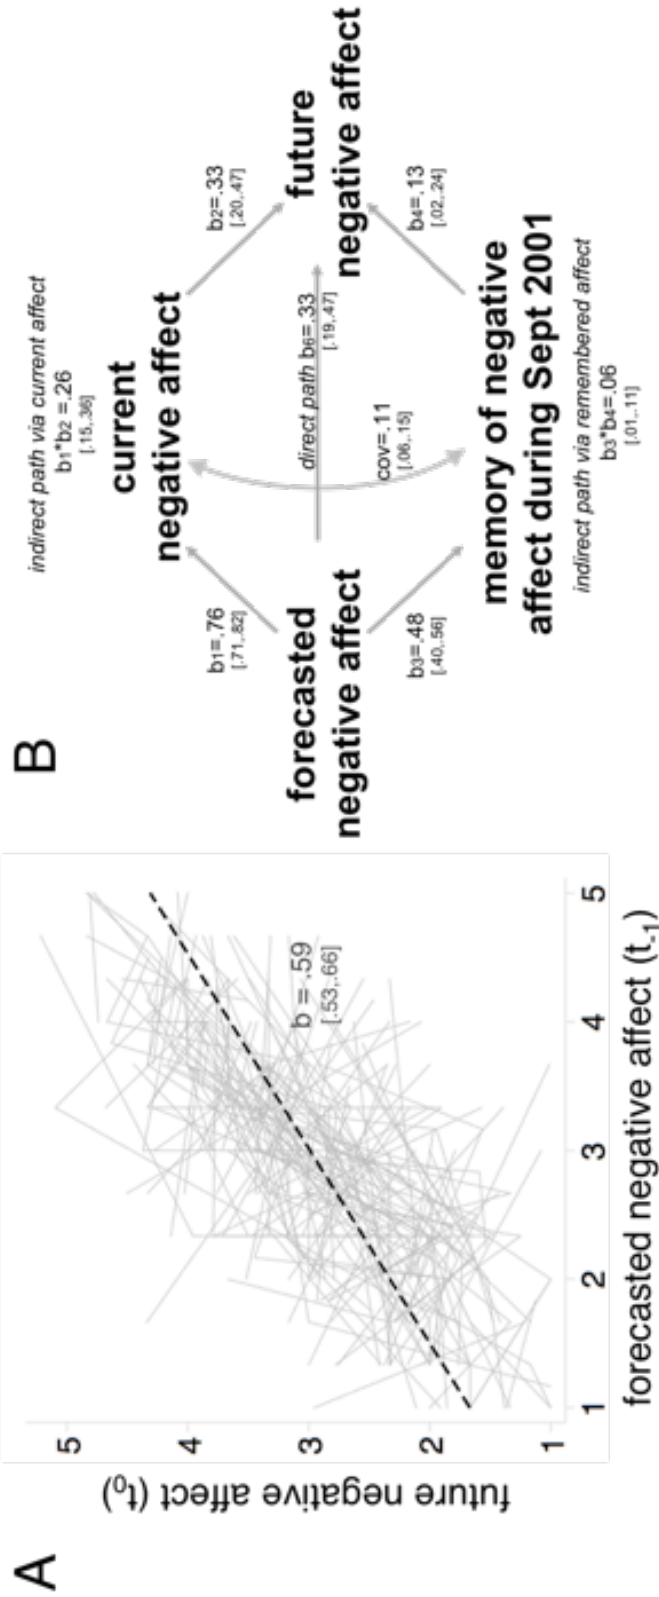


Figure 4. A) Total predictive relationship between forecasted negative affect (at a particular time point) and experienced negative affect (at the following time point, for which the prediction was made), shown for the group as a whole (dotted black line), and in terms of raw individual curves (light grey lines). **B)** Multilevel path model identifying mediators of the relationship between forecasted future negative affect and actual future negative affect. Results indicate that current feelings and memories of feelings from September 2001 independently mediate the relationship between forecasted and experienced negative affect. Unstandardized betas with 95%CI.

the relationship between forecasted negative affect and future negative affect is mediated via current negative affect, *indirect path* $b_1*b_2=.26$, 95%CI [.15, .36], as well as (to a lesser extent) remembered negative affect, *indirect path* $b_3*b_4=.06$ [.01, .11], and 2) after accounting for these variables, there is a nonzero direct path between predicted and future negative affect, *direct path* $b=.33$, 95%CI [.19, .47], representing 56% of the total predictive effect ($b=.59$) of forecasted affect on experienced affect. These results suggest that forecast accuracy can be partially explained by the extent to which forecasts correspond with (i.e., are anchored to) representations of current negative affect and (to a lesser extent) memories of past negative affect.

Discussion

We began this investigation by asking whether people can accurately predict their future emotional responses to a national tragedy, the attacks of September 11, 2001. Though prior work has highlighted ways in which we fail to forecast accurately, here we found that people could predict how they would feel about September 11 in 1, 2, or 7 years with striking accuracy. Predictions were accurate with respect to both the overall intensity as well as the dominant quality (e.g., sad, fearful, angry) of negative affect experienced in the future. Moreover, the relationship between participant's predictions and their future affective experiences was partially mediated by the intensity of current feelings and the intensity of remembered feelings experienced in the immediate wake of the attacks, indicating that prediction accuracy can be attributed in part to similarity of future feelings to current and remembered feelings.

Sources of accuracy and inaccuracy in emotional prediction

We had hypothesized that participants could reach high accuracy in emotional prediction by relying on knowledge of their feelings from the past and in the present, and this was borne out in our data. In general, our results are consistent with a theoretical model in which forecasters use representations of current and (to a lesser extent) past emotional responses as “anchors” guiding forecasts of how they will feel in the future. However, if people made forecasts based only on what they feel currently and felt in the past, they would not be able to reach the levels of accuracy we observed here. Our data thus suggest that forecasters also accurately “adjust” forecasts downward to correspond with the decay of negative affect over time (and different rates of decay for different kinds of emotion).

The finding that current and past feelings can be useful and accurate guides for predicting how we will feel in the future is important in two ways. First, it contrasts with previous findings that anchoring on remembered feelings can lead forecasters astray when predicting reactions to novel events (Morewedge, Gilbert, Wilson, 2005). Unlike this prior work, in which participants recalled atypical instances of a class of ordinary events (e.g., the worst time they ever missed a train), in this study people made forecasts of future feelings in response to an event that had already happened, for which they had representations of current and past feelings to draw on. Our data suggest that these current and past feeling are representative of how participants will feel when thinking about the attacks in the future, and thus relevant for making accurate forecasts. This

would not be the case for forecasting one's response to an event that has never occurred, especially if one has never experienced anything similar to the event. In such a case current and past feelings would be less relevant as an anchor for the forecast.

Second, consistent with our hypotheses that participants would be able to accurately predict their future feelings by relying, in part, on an extrapolation from their present and past feelings, the level of accuracy apparent in our data is appreciably higher than what has been observed in the prior literature. The standardized relationship between predicted and experienced negative affect in our dataset (0.57) is twice the magnitude of the average relationship reported in a recent meta-analysis (0.28) of 16 published studies (Mathieu & Gosling, 2012). The overall magnitude of absolute inaccuracy (the difference between predicted and experienced affect) in our dataset was approximately zero (i.e., no evidence for a directional bias), whereas the average effect size in the previous literature is an overestimation of about 0.55 standard deviations (Levine, Lench, Kaplan, & Safer, 2012; Mathieu & Gosling, 2012). Overall, our data indicate that participants are able to achieve extremely high accuracy in predicting their future feelings over multiple year intervals when predicting and reporting on emotional responses to a focal event that has already occurred and that is of ongoing impact and relevance.

Lay theories of emotional change and stability

That participants in this study were able to achieve this level of accuracy indicates that their predictions were sensitive to the overall decay of negative

affect over time as well as to different rates of decay for different categories of emotion. For example, participants correctly predicted that fear would decay more quickly than anger from 2001 to 2002. Moreover, from person to person, participants showed a high degree of relative accuracy in predicting how much and what kind of negative emotion they would feel in the future. Together, these data suggest that predictions are guided by an implicit theory of emotional change that incorporates different trajectories for different emotions in the context of a particular focal event, and which may be influenced by factors like personal experience of similar events, shared cultural knowledge, or social norms (see Molden & Dweck, 2006; Labroo & Mukhopadhyay, 2009). As such, the data suggest that immune neglect – the notion that people are unaware that their emotional reactions diminish over time through emotion regulatory processes – may not appear for certain classes of emotional prediction (see Gilbert et al., 1998; Wilson & Gilbert, 2013).

From main effect biases to context-dependent accuracy

Despite striking levels of overall accuracy, we observed examples of under- and overestimation of future emotion for the group as a whole, and for some individuals. In particular, predictions showed evidence for slight underestimation in 2001, were accurate in 2002, and showed evidence for slight overestimation in 2004, suggesting a protracted (3-year) shift from under- to over-estimation with the passing of time from a still impactful focal event. On another note, although there was strong evidence of relative forecast accuracy for every emotion type we asked about (sadness, fear, anger, confusion, shock

and frustration), the evidence for absolute forecast accuracy (i.e., equivalence across forecasted and actual affect) was not as compelling for frustration in particular or for confusion, shock and frustration considered as a group. In the context of the prior literature, which has demonstrated biases in forecasting as main effects, the present data suggest the importance of an approach that seeks to identify the factors that set the stage for accuracy and bias in emotional prediction (see Levine, Lench, Kaplan, & Safer, 2012; Gilbert & Wilson, 2009; Miloyan & Suddendorf, 2015; Wilson & Gilbert, 2012). More broadly, this approach converges with recent calls for increased focus on understanding contextual dependencies in the study of affective and social processes (Doré, Silvers, Ochsner, 2016; Gelman, 2015).

Implications for societal-level policy and decision making

Few events in recent memory have been as impactful on the United States' social and political sphere as the attacks of September 11, 2001. For an event of this nature, which can motivate political decisions with broad and long-lasting consequences, it matters a great deal whether predictions of how we will feel about them in the future are biased upwards or downwards, and whether they actually correspond with our future feelings as individuals or as a society. Our data indicate that predictions of future feelings contain information about the overall intensity of negative affect, as well as whether it will take the form of sadness, fear, or anger. At the individual level, having accurate knowledge that one's emotions will change in the future may influence patterns of appraisal, emotion, and behavior in the present. At the community level, the information in

predictions of this nature could be useful for allocating therapeutic resources, above and beyond reports of current and remembered feelings. At the societal level, it could be important for future research to consider the effects of current and predicted sadness, fear, and anger on policy decisions that unfold over months, years, and decades. Moreover, interventions that draw on people's ability to forecast their future feelings could enhance their ability to make decisions on the basis of their likely future, rather than their current, emotional states (see Lerner, Li, Valdesolo, & Kassam, 2015).

Limitations and Future Directions

Limitations of this report are worth noting, in part because they may provide direction for future research. We quantified accuracy in affective forecasting based on self-report ratings of affective experience, which are not isomorphic with psychophysiological, neural, or behavioral measures of affective responding (Mauss & Robinson, 2009), and which can be subject to demand characteristics. However, we find it unlikely that demand characteristics had a large impact on forecast accuracy in this sample in that 1) any demand would have to take effect over 1-, 2-, or 7-year delays (e.g., recalling one's previous forecast ratings when rating current affect), and 2) all within-subject studies of affective forecasting are subject to similar demands and yet most have yielded lower accuracy estimates (Mathieu & Gosling, 2012). Nonetheless, future work investigating the ability to forecast psychophysiological, neural, or behavioral indices of emotional responding may broaden and enrich current models of affective forecasting and theories of emotional awareness more generally.

Conclusion

Knowing how we will feel in the future helps us make decisions in the present. Although prior research has highlighted ways in which we fail to forecast with accuracy, our understanding of when and why emotional predictions can be made accurately is largely incomplete. Here we show that people make strikingly accurate forecasts about their future emotional states with reference to a tragic event of historical significance by relying, in part, on an extrapolation from their feelings in the past and present. We hope that future work will extend this research by seeking to further identify the conditions and mechanisms that underlie our ability to predict the emotional future.

Acknowledgments

We thank our editor and reviewers for their constructive feedback, and Dan Gilbert for helpful comments on an earlier version of this manuscript. We also thank Noam Zerubavel, Laura Braunstein, Greg Jensen, Seth Kallman, Rebecca Martin, Michael Gilead, and Chelsea Helion for discussion of this work; NIA grant R01AG043463-01 and Conte grant PAR-11-126 for support to KO; NIMH grant RO1-MH0066972, and the James S McDonnell Foundation for support to WH. Finally, we thank the 9/11 Memory consortium and the many student researchers who assisted with data collection and entry.

References

- Ayton, P., Pott, A., & Elwakili, N. (2007). Affective forecasting: Why can't people predict their emotions?. *Thinking & Reasoning*, 13(1), 62-80.
- Doré, B.P., Silvers, J.A., & Ochsner, K.N. (2016) Emotion regulation 2.0: Towards a personalized science of emotion regulation. *Social and Personality Psychology Compass*.
- Dunn, E. W., Brackett, M. A., Ashton-James, C., Schneiderman, E., & Salovey, P. (2007). On emotionally intelligent time travel: Individual differences in affective forecasting ability. *Personality and Social Psychology Bulletin*, 33(1), 85-93.
- Funder, D. C. (1995). On the accuracy of personality judgment: a realistic approach. *Psychological Review*, 102(4), 652-670.
- Gelman, A. (2015). The Connection Between Varying Treatment Effects and the Crisis of Unreplicable Research A Bayesian Perspective. *Journal of Management* (41), 632-443. .
- Gilbert, D. T., & Wilson, T. D. (2009). Why the brain talks to itself: Sources of error in emotional prediction. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 364(1521), 1335-1341.
- Gilbert, D. T., Pinel, E. C., Wilson, T. D., Blumberg, S. J., & Wheatley, T. P. (1998). Immune neglect: a source of durability bias in affective forecasting. *Journal of Personality and Social Psychology*, 75(3), 617-638.
- Hirst, W., Phelps, E. A., Buckner, R. L., Budson, A. E., Cuc, A., Gabrieli, J. D., ... & Vaidya, C. J. (2009). Long-term memory for the terrorist attack of

- September 11: flashbulb memories, event memories, and the factors that influence their retention. *Journal of Experimental Psychology: General*, 138(2), 161-176.
- Hirst, W., Phelps, E. A., Meksin, R., Vaidya, C. J., Johnson, M. K., Mitchell, K. J., ... & Olsson, A. (2015). A Ten-Year Follow-Up of a Study of Memory for the Attack of September 11, 2001: Flashbulb Memories and Memories for Flashbulb Events, *Journal of Experimental Psychology: General*, 144(3), 604-623.
- Hoerger, M. (2012). Coping strategies and immune neglect in affective forecasting: Direct evidence and key moderators. *Judgment and Decision Making*, 7(1), 86.
- Kass, R. E., & Raftery, A. E. (1995). Bayes factors. *Journal of the American Statistical Association*, 90(430), 773-795.
- Labroo, A. A., & Mukhopadhyay, A. (2009). Lay theories of emotion transience and the search for happiness: A fresh perspective on affect regulation. *Journal of Consumer Research*, 36(2), 242-254.
- Lam, K. C., Buehler, R., McFarland, C., Ross, M., & Cheung, I. (2005). Cultural differences in affective forecasting: The role of focalism. *Personality and Social Psychology Bulletin*, 31(9), 1296-1309.
- Lerner, J. S., Li, Y., Valdesolo, P., & Kassam, K. S. (2015). Emotion and decision making. *Annual Reviews of Psychology*, 66. 799-823.

- Levine, L. J., Lench, H. C., Kaplan, R. L., & Safer, M. A. (2012). Accuracy and artifact: reexamining the intensity bias in affective forecasting. *Journal of Personality and Social Psychology*, 103(4), 584-605.
- Mathieu, M. T., & Gosling, S. D. (2012). The Accuracy or Inaccuracy of Affective Forecasts Depends on How Accuracy Is Indexed A Meta-Analysis of Past Studies. *Psychological Science*, 23(2), 161-162.
- Mauss, I. B., & Robinson, M. D. (2009). Measures of emotion: A review. *Cognition and emotion*, 23(2), 209-237.
- Miloyan, B., & Suddendorf, T. (2015). Feelings of the future. *Trends in Cognitive Sciences*, 19(4), 196-200.
- Molden, D. C., & Dweck, C. S. (2006). Finding "meaning" in psychology: a lay theories approach to self-regulation, social perception, and social development. *American Psychologist*, 61(3), 192-203.
- Morewedge, C. K., Gilbert, D. T., & Wilson, T. D. (2005). The least likely of times how remembering the past biases forecasts of the future. *Psychological Science*, 16(8), 626-630.
- Scherer, K. R., Schorr, A., & Johnstone, T. (Eds.). (2001). *Appraisal processes in emotion: Theory, methods, research*. Oxford University Press: Oxford.
- Szpunar, K. K. (2010). Episodic future thought: an emerging concept. *Perspectives on Psychological Science*, 5(2), 142-162.
- Tulving, E. (1985). Memory and consciousness. *Canadian Psychology*, 26(1), 1.
- Wilson, T. D., & Gilbert, D. T. (2005). Affective forecasting: knowing what to want. *Current Directions in Psychological Science*, 14(3), 131-134.

Wilson, T. D., & Gilbert, D. T. (2013). The impact bias is alive and well. *Journal of Personality and Social Psychology*, 105, 740–748.

Wilson, T. D., Wheatley, T., Meyers, J. M., Gilbert, D. T., & Axson, D. (2000). Focalism: a source of durability bias in affective forecasting. *Journal of Personality and Social Psychology*, 78(5), 821-836.

Zaki, J., & Ochsner, K. (2011). Reintegrating the study of accuracy into social cognition research. *Psychological Inquiry*, 22(3), 159-182.

Figure Captions

Figure 1. Current and forecasted intensity of sadness, fear, and anger in 2001, 2002, 2004, and 2011. Negative emotions decay slowly over time, and forecasts are largely accurate at the level of participants as a group. Group means with 95%CI.

Figure 2. Relationship between the overall intensity of negative emotions reported in response to the attacks in 2002, 2004, and 2011 (on the y axes) and forecasts of negative emotion intensity made 1-, 2-, and 7-years earlier (on the x axes). Forecasts show a high degree of relative accuracy: people who predict they will feel the most negative tend to later report feeling the most negative.

Figure 3. Relationships between the relative degree of specific types of emotion reported in response to the attacks in 2002, 2004, and 2011 (i.e. reports of particular categories centered on the mean of all reported negative emotion, on the y axes) and forecasts of relative degree made 1-, 2-, and 7-years earlier (i.e. forecasts mean-centered in the same manner, on the x axes). Forecasts show a high degree of relative accuracy – people who predict that they will feel predominantly sad, angry, or fearful tend to later end up doing so.

Figure 4. A) Total predictive relationship between forecasted negative affect (at a particular time point) and experienced negative affect (at the following time point, for which the prediction was made), shown for the group as a whole (dotted black line), and in terms of raw individual curves (light grey lines). **B)** Multilevel path model identifying mediators of the relationship between forecasted future negative affect and actual future negative affect. Results indicate that current feelings and memories of feelings from September 2001 independently mediate the relationship between forecasted and experienced negative affect. Unstandardized betas with 95%CI.