Final version of this paper appears at International Conference on Affective Computing and Intelligent Interaction, 2009

Assessing the validity of a computational model of emotional coping

Stacy Marsella, Jonathan Gratch, Ning Wang, Brooke Stankovic
University of Southern California
13274 Fiji Way, Marina del Rey, CA 90405

Abstract

In this paper we describe the results of a rigorous empirical study evaluating the coping responses of a computational model of emotion. We discuss three key kinds of coping, Wishful Thinking, Resignation and Distancing that impact an agent's beliefs, intentions and desires, and compare these coping responses to related work in the attitude change literature. We discuss the EMA computational model of emotion and identify several hypotheses it makes concerning these coping processes. We assess these hypotheses against the behavior of human subjects playing a competitive board game, using monetary gains and losses to induce emotion and coping. Subject's appraisals, emotional state and coping responses were indexed at key points throughout a game, revealing a pattern of subject's altering their beliefs, desires and intentions as the game unfolds. The results clearly support several of the hypotheses on coping responses but also identify (a) extensions to how EMA models Wishful Thinking as well as (b) individual differences in subject's coping responses.

1. Introduction

Over the last several decades, there has been extensive work in computational models of human emotion. Emotion models have been proposed as a basic research methodology for exploring the dynamic properties of human cognition and emotion [1, 2]. In addition, they have been used in a range of applications to model users for human-computer interaction [3], in educational systems [4, 5] and to make intelligent systems or robots more robust and reactive [6]. Embodied agent research also extensively uses emotion models to create more life-like, expressive virtual characters for a variety of applications [7, 8].

The various computational models of emotion that have now been developed draw from a range of psychological theories of human emotions, animal behavior models and neuroscience theories. However, most computational models (e.g., [9-13]) have been heavily influenced by appraisal theory (e.g.[14]), a psychological theory of emotion that argues that emotion arises from a person's subjective interpretation of their relation to the environment, including whether a situation is desirable or not, how controllable it is and how expected or likely it is.

An emotion in turn leads to a range of influences on cognition and behavior. Two broad classes of coping response have been characterized: problem-focused-coping and emotion-focused coping, Problem-focused coping seeks to change the world while emotion-focused changes the self by adapting desires, intentions or beliefs. Additionally, research in the neurosciences and economics [15] has identified how emotion plays a central role in decision-making.

Although there has been extensive work in computational models of human emotional responses, little work has been done in validating these models, with a few exceptions [16, 17]. This article describes the results of a rigorous empirical study assessing the human behavior fidelity of a particular model of coping that is part of the EMA [18] computational model of emotion. Specifically, we study the model's predictions concerning three forms of emotion-focused coping response: resignation, distancing and wishful thinking. These responses are of particular interest as their impact on an agent's beliefs, desires and intentions is inconsistent with normative models of rationality. We assess the model by comparing its predictions with behavior of human subjects playing a competitive board game, using monetary gains and losses to induce emotion and coping responses. We indexed subjects' appraisals, emotional state and coping responses at key points throughout a game, revealing a coherent pattern in the dynamic relationship between these factors. The human responses are contrasted with a computational model of the game scenario constructed in EMA. The results provide support for the EMA model of emotion-focused coping as well as general guidance on how to improve computational models of emotion regulation.

2. Background

Appraisal theory argues that emotions arise from a subjective assessment of their relation to events in the environment, what Lazarus [19] calls the person-environment relationship. This appraisal occurs along several dimensions (good vs. bad; likely vs. unlikely; controllable vs. uncontrollable; etc.). Emotions lead to coping processes that act on a person's perceived relationship to their environment. These include "problem focused" strategies (e.g. planning and taking action) directed towards improving the world but also encompass "emotion-focused" strategies that influence epistemic or motivational states.

In this paper, we consider closely three coping strategies: resignation, wishful thinking and distancing. To model these strategies computationally we must give them precise definitions. Resignation is defined here as giving up an intention to achieving a desired goal. For example, one might become resigned to losing a game that is going badly. Wishful thinking entails altering one's beliefs about the likelihood of an undesired or desired outcome. One might, for example, believe the game is still winnable even in the face of disconfirming evidence. Distancing entails altering the desirability of a goal. For instance, as the game goes badly, one might view winning as less important. These three coping strategies impact three distinct aspects of cognition and decision-making as well as span all three parts of a belief-desire-intention (BDI) model of agency [20].

Coping models make very different predictions from decision theory concerning the relationship between probability and utility. In decision theory, these factors are independent (one's perception of the likelihood of a goal shouldn't influence its desirability, and vice versa). In contrast, the strategies of distancing and wishful thinking, creates this linkage (e.g., distancing reduces the desirability of an unlikely goal).

Such correlations are not unique to appraisal theories or emotion research. Related phenomena have been studied in a range of fields. Work in motivated inference and reasoning has explored the general question of how motivations influence belief systems [21]. In the attitude change literature, McGuire's System of Thoughts theory [22], poses a number of postulates about how desirability and likelihood *positively correlate*: The *wishful thinking postulate* states that a person's judgment of an event's desirability will affect the judgment of its likelihood – the higher the desirability, the greater the likelihood; The *rationalization postulate* posits that a person's expectation of an event's likelihood will positively correlate with its desirability. According to McGuire and McGuire [22], this encompasses both a sweet lemon rationalization (increase in likelihood leads to increases desirability) and sour grapes rationalization (a decrease in likelihood decreases desirability). Finally, research has also argued that the correlation between desirability and likelihood of an event is dependent on motivational involvement (e.g., [23]).

In contrast, work on the scarcity effect [24] argues for a *negative correlation* between likelihood and desirability, specifically a decrease in likelihood of an outcome would increase its desirability.

3. EMA and Appraisal Theory

Appraisal Theory provides high-level specification of the requirements for a computational model of emotion. To realize this model, we must make explicit choices about how to represent the person-environment relationship, how appraisal processes operate over this representation, how these appraisals lead to an evolving emotional state and how emotion leads to coping responses that in turn alter the person's/agent's actions and/or subjective assessments of their relation to the environment.

EMA (Emotion and Adaptation) is a computational model of emotion, including models of appraisal and coping processes. Details of the representations and process used in EMA have been described elsewhere[18, 25]. Here we provide a simplified, high-level illustration of how it works, sufficient for understanding the current study.

EMA consists of a set of processes that interpret an explicit representation of the person-environment relationship. This represents states and actions in the world, beliefs, desires and intentions of self and other, and the causal relationships between them.[1] From this it derives a set of posited appraisal variables, and a set of coping strategies that manipulate this representation in response to the appraised interpretation.

To illustrate how EMA works, consider a simple representation of Maria playing a game as it would be encoded in EMA, as depicted in top of Figure 1. Maria wants to win (subjective utility is 20) but thinks she is losing (subjective probability of winning is 30%). A standard decision-theoretic approach would simply use the expected value of this situation to inform its intention to act.[2] In contrast, EMA first appraises the situation, producing a set of emotions, and possibly emotionally copes with these emotions before acting. EMA appraises the same situation from multiple perspectives – Maria both hopes she will win and fears she will lose (see [26] for a discussion of the intensity of these responses). EMA selects the most intense emotion as the focus of coping (in this case fear), and adopts one of a number of coping strategies based on how the situation is appraised. Emotion-focused strategies are considered if, as in Maria's case, the appraisal is negative and there is little appraised sense of control (which EMA defines as likelihood that the agent can attain the goal).

Although EMA models twelve distinct forms of coping (see [18]), this study considers three: Wishful Thinking, Resignation and Distancing. We focus on these three, because, as they are modeled in EMA, (a) they are the mechanisms through which EMA creates correlations between likelihood and desirability and (b) they operate on distinct aspects of cognitive state in EMA (i.e., beliefs, desires and intentions).[3] As illustrated in the bottom of Figure 1, Wishful Thinking would alter Maria's belief about the likelihood of winning, Distancing would reduce her desire to win, and Resignation would lower her intention to try and win. Specifically these three strategies are realized in EMA as follows:

Wishful Thinking: Increase (lower) the probability

[1] EMA encoded domains using a domain-independent language of decision-theoretic plans, augmented by the explicit representation of intentions and beliefs.

[2] Typically, expected utility is used to rank alternative actions with the agent forming the intention to perform the action that maximizes expected utility. As there is only one action here, we can see expected utility as a measure of the agent's intention to play the game.

[3] Other strategies apply to social aspects of the situation (e.g., seek instrumental social support or shift blame), attention aspects (e.g., suppress information), or problem-directed activities (e.g., planning).

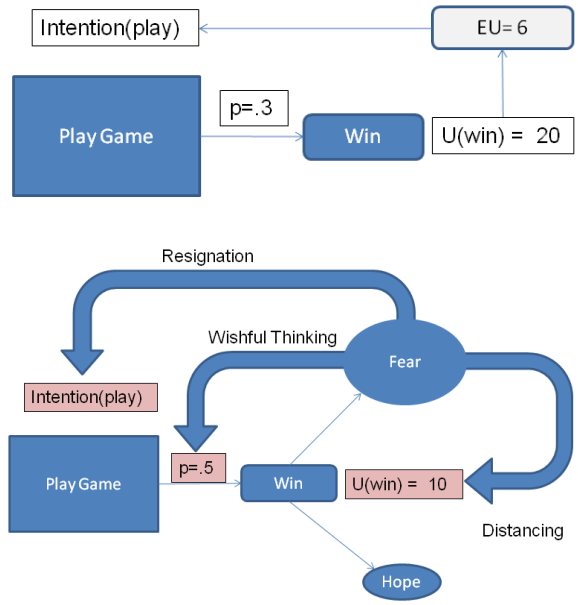

Figure 1. Top: Representation of Maria playing game. Bottom: Impact of appraisal and alternative coping responses.

of a pending desirable (undesirable) outcome or assume some intervening act or actor will improve desirability. For example, if the appraisal frame is associated with a future action with an undesirable outcome, wishful thinking will lower the perceived probability that this effect will occur, by a fixed percentage. Wishful thinking is considered as a response to a negative emotion and preferred if the appraised controllability of the outcome is low.

Distance/Mental disengagement: Lower utility attributed to a desired but threatened state. For example, if an agent's plan for achieving a goal has a low probability of success, the consequence of distancing is that the agent will come to care less about this goal. Specifically utility is reduced by a fixed percentage. Distancing is considered in response to a negative emotion and preferred if the appraised controllability of the appraised outcome is low.

Resignation: Drop an intention to achieve a desired state. For example, if a goal is appraised as essentially unachievable, the agent may abandon this goal. In EMA this intention is modeled as a binary choice, either the agent intends, or doesn't intend, to achieve the goal. Resignation strategy is considered in response to a negative emotion and is preferred if the agent has little appraised control over the state.

Note that in the case of these coping operations, EMA predicts a positive correlation between perceived utility and probability, but restricted to the case of negative emotions. Further the magnitude of change in utility for Distancing will depend on its initial magnitude, for three reasons. The reduction by fixed percentage will lead to a greater drop in utility magnitude in the case of high utility goals. Further, the focusing mechanism and threshold mechanisms may lead to other emotions being coped with or no coping response, respectively. On the other hand, for Wishful Thinking, the correlation between perceived utility and probability would be due only to the latter two mechanisms, the focus and threshold mechanisms.

Each of these strategies applies in situations where negative emotions are elicited and perceived control is low. The impact of these strategies is to reduce the negative emotion arising from subsequent appraisals. EMA makes no prediction about which of these strategies would be selected and one is selected and applied at random.

3.1. Hypotheses

The EMA model leads to the following hypotheses about an agent playing the game in which it has become apparent that they are going to lose (or win).

H1:Distancing:

H1: Perceived utility of winning will drop as player loses

H1b: The strength of the desire to win predicts the magnitude of the coping effect.

H2:Resignation:

H2: Willingness to play will drop as player loses

H2b: The strength of the desire to win will positively correlate with the effect

H3: Wishful thinking:

H3: Perceived probability of winning will be overestimated as player loses

H3b: The strength of the desire to win will positively correlate with the effect

Additionally, the study will explore two research questions. First, EMA provides no specific hypotheses concerning a game in which the player perceives they are winning and we wish to assess what subjects do in these circumstances. Second, we hope to identify some situational factors that distinguish when these different coping strategies will be selected.

4. Experiment

To evaluate EMA's coping model we need to systematically manipulate events concerning the probability of goal attainment and assess the subject's resulting responses, including their emotional state and perceptions of the probability and utility. We compare the predictions of the model with subjects' responses when playing a competitive board game called Battle-ship by the Milton Bradley Company. In the standard game, players secretly place ships on a small grid. Players then take turns shooting at squares in the grid in an attempt to sink their opponent's ships.

To induce emotions, subjects play for money (they can win or lose up to $10 US). To create a wide spread of positive and negative emotions we systematically

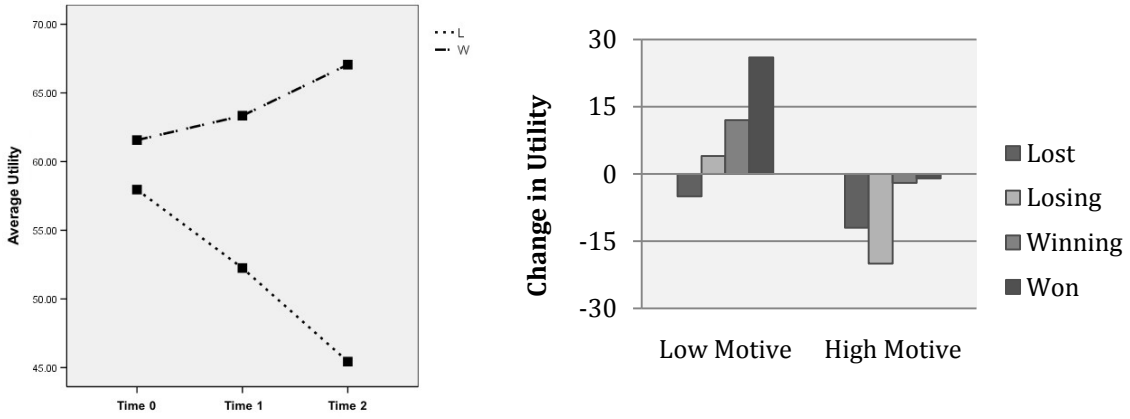

Figure 2: Distancing – subjects exhibited a tendency to distance as function of condition (wining vs. losing). Graph on the left illustrates that subjects attribute less utility to winning as they lose and more utility as they win. Graph on the right shows the pattern of distancing reverses depending on if subjects in were initially motivated to win or not.

manipulate perceptions about the likelihood of winning (both within and between subjects) by altering the sequence of hits and misses the subject obtains, and perceptions about the importance of winning/losing (between subject) by framing the game as an opportunity to win money or to lose money. We control the unfolding of the game by use of a confederate. Although subjects believe they are playing against another subject, in reality they are playing against a confederate that is watching their game play through a hidden camera and controls the series of hits and misses.

We also would like to assess how subjects' responses unfold over time. To explore dynamics, we use repeated measures to assess how subjects' appraisals, emotions and coping change within the game. We index subjects' subjective impressions at game start, middle and end.

4.1 Method

One hundred and seven people (41% women, 59% men) from the general Los Angeles area participated in this study. They were recruited using posts on Craigslist.com and were compensated $30 for one hour of their participation. On average, the participants were 36.2 years old (min = 18, max = 60, std = 11.9).

4.1.1 Design

The study is designed as a 2x2 between-subjects experiment. The two independent variables are framing and game play.

Framing. There were two alternative framings: positive incentive and negative incentive. In the positive incentive condition, participants are recruited using posters saying they will be paid $20. Upon arriving, they are then informed that they can win up to additional $10 if they win the game. In the negative incentive condition, the recruitment poster says the compensation is $30. At the lab, the experimenter then informs the participants that they can lose up to $10 if they lose the game. All participants are paid $30 in the end regardless of framing and game result.

Game play. There are two conditions for the game play: win game (n=53) and lose game (n=54). In the winning condition, participants win the game. In the losing condition, participants lose the game.

4.1.2 Procedure

Participant and the confederate enter the laboratory and are told they are participating in a study to play Battleship game. After they read the informed consent form, the experimenter explained to the participants in the positive incentive condition that the winning player can win up to additional $10. The participants in the negative incentive condition were told that the losing player could lose up to $10.

The confederate and the participant view a PowerPoint presentation about rules of the Battleship game and procedures of the experiment. They then fill out a pre-test questionnaire. Battle-ship game began after completion of the questionnaire. Experimenter leaves the room.

Game play is divided into three stages. T0 refers to the start of the game. The first stage continues until point T1 when one player is clearly winning: i.e., the participant has sunk 3 of the confederates 4 ships (in the winning condition) or lost 3 of 4 of their own ships (in the losing condition). Finally, T2 is the point when the game has been won by one of the players. At each time point the participant fills out an appraisal and emotion questionnaire.

Finally, participants were debriefed individually and probed for suspicion using the protocol from Aronson, Ellsworth, Carlsmith, and Gonzales [24]. No participants indicated that they believed their opponent was a confederate in the study. All subjects were allowed to retain the addition $10.

4.1.3 Equipment

The participant and confederate each sit in front of a desk that is placed opposite to each other. A white board is placed in the middle to separate the two desks. A Battleship board is place on each desk. A Dell desktop computer is place next to each of the Battleship board. The participant fills out the questionnaires on the computer. A hidden wireless camera is placed on the ceiling to record participant's moves on the Battleship board.

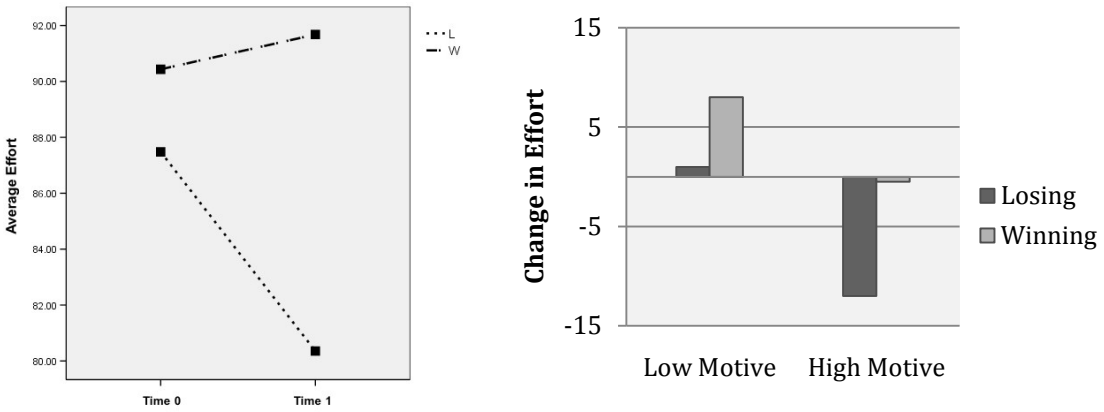

Figure 3: Resignation – subjects exhibited a tendency to moderate game playing effort as function of condition (wining vs. losing). Graph on the left illustrates that subjects reported trying less as they lose and more as they win. Graph on right shows pattern of resignation depends on whether subjects were initially motivated to win or not.

The camera video is sent to the confederate's computer so they can control precisely the subject's experience according to the condition (Win vs. Lose).

4.1.4 Measures[4]

Demographic/Dispositional information: At the beginning of the experiment we ask participants demographic information, board game and Battleship experience, and social value orientation [25], a measure of tendency to be cooperative or competitive.

Several items are repeatedly measured at time points T0 (start), T1 (middle), and T2 (end):

Emotions: All emotions were measured using a visual analog scale ranging from 0 to 100. We constructed a 5-item emotion scale measuring the intensity of emotional feeling experienced by the participant at a given time point. Emotions assessed include fear, joy, sadness, anger and hope.

Appraisal and Coping Scale: We developed a 12-item appraisal scale to measure participant's perceptions of winning utility and likelihood, ability to control the outcome, effort devoted to winning, as well as several measures related to importance and likelihood that the game was played fairly.

All scales are presented as an analog scale that ranges from zero (minimum value/intensity) to 100 (maximum value/intensity).

5. Results

Data from six sessions were excluded due to incomplete questionnaires or experiment procedure deviating from protocol. As a result, data from 101 participants were included in the analysis, 48 from the losing condition and 53 from the winning condition.

Our manipulation of subjective sense of winning was successful. Subjects perceived they have an approximately even chance of winning at the start of the game. Perceptions of winning increased in the wining conditions ($p<0.001$) and decreased in the losing condition ($p<0.001$). Perceived probability changed approximately linearly across the stages of the game: Pr(Losing)=0.27; Pr(Start)=0.55; Pr(Wining)=0.76. Our manipulation of incentive was unsuccessful – subjects' responses and behavior were largely indistinguishable across the two conditions.[5] However, both positive and negative emotions were successfully elicited and we collapse across the incentive conditions in all subsequent analysis.

5.1. Distancing

Hypothesis H1 predicts that subjects will cope with the distress of losing by attributing less utility to winning as the game unfolds. Figure 2a illustrates the average utility that subjects assigned to winning the game as a function of condition (Win vs. Lose) and time (start, middle, and end). Consistent with H1, subjects in Lose condition attribute less utility on average to winning as they begin to lose (T1) and after they have lost (T2). Subjects in the Win condition display an unpredicted trend of "inverse distancing" – i.e., attributing more utility to winning as they begin to win (T1) and after they have won (T2). We assessed significance of these trends using the general linear model (GLM) repeated measure test. The main effect of condition (Win vs. Lose) is significant ($p=.042$), indicating that subjects' coping style is significantly different depending on if they are in the winning or losing condition. The interaction between time and condition is also significant ($p=.003$) indicating that the trend over time is significantly different across conditions.

Hypotheses H1b argues that distancing will be stronger for subjects that are more motivated to win. To analyze this, we divided subjects into subgroups based on the utility they attributed to winning before the game begins (T0) and compared differences in their coping

[4] Several measures are included for completeness but not discussed as they apply to hypotheses in a companion article [26].

[5] A MANOVA showed no interaction between framing and game play on any independent variables (hope, fear, joy, sadness) except for a small significant interaction with fear when participants are losing/winning.

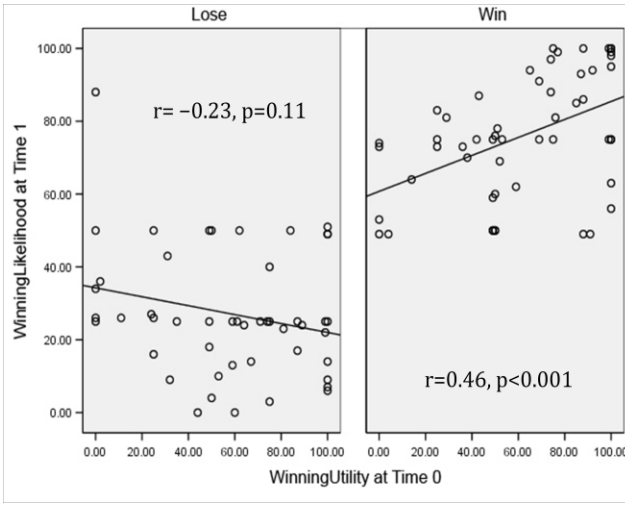

Figure 4: Wishful Thinking – subjects exhibited a correlation between their motives (utility at T0) and beliefs (subjective probability of winning at T1). Subjects in the winning condition were more optimistic about their chances of winning when they were more motivated to win. Subjects in the losing condition showed the opposite trend of being less optimistic when they were more motivated.

behavior. The *high-motive* group is the 1/3 of subjects within a condition that assigned the most utility to winning, the *low-motive* group is the 1/3 of subjects that assigned the least utility to winning. Figure 2b illustrates the average amount that these subgroups shifted their initial utility of winning as a function of condition and time period. As predicted, subjects in the high-motive group showed greater distancing in the losing condition than the low-motive group. The low-motive group also shows an "inverse distancing" (increased engagement) in the winning condition that EMA makes no prediction about. Subjects that didn't care about winning began to care if they started to win. Further, this low motive group seems to fully explain the upward trend in utility in the winning condition.

5.2. Resignation

Hypothesis H2 claims that subjects will cope with the prospect of losing by trying less hard to win. Figure 3a shows the average effort that subjects report at different stages of the game. Consistent with H2, subjects expend less effort on winning as they begin to lose. There is an unpredicted positive trend toward expending more effort when subjects begin to win. Using GLM repeated measure test, the main effect of condition (Win vs. Lose) play is marginally significant (p=.079) and the interaction between time and condition is also marginally significant (p=.082). This indicates that H2 was partially supported.

Hypothesis H2b asserted that the strength of resignation is proportional to the utility subjects' assign to winning. As in distancing, we divided subjects into high, medium and low utility groups based on the utility they assigned to winning before the start of the game (see Figure 3b). Consistent with H2b, when faced with losing, subjects reported more resignation on average in the high-motive group when compared with the low-motive group. As with distancing, we also observed an inverse in coping behavior that was not predicted: subjects with little motivation to win showed no resignation in the losing condition but showed "inverse resignation" (or engagement) when they began to win. Further, this low motive group seems to fully explain the observed upward trend in effort in the winning condition.

5.3. Wishful Thinking

Hypothesis H3 states that subjects will cope with the negative emotions of losing by denying the reality of their impending loss (i.e., they will overestimate their probability of winning in the Lose condition. We further predict that the strength of this effect (H3b) is proportional to the utility they attribute to winning. This argues that there will be a positive correlation between winning probability and winning utility in the Lose condition. Figure 4 illustrates scatter plots illustrating the relationship between utility of winning before the start of the game (T0) and perceived probability of winning at T1 in the losing and wining condition.

Contrary to H3 and H3b, we observed no significant positive correlation between utility and probability in the losing condition. Worse for the prediction, there was a near significant opposite trend (r= −0.23, p=0.11). More surprising, we observed a highly significant positive correlation in the Win condition: subjects that are motivated to win overestimated their chance of winning when compared with other subjects (r=0.46, p<0.001).

6. Discussion

Overall, the results provided support for EMA's predictions that, when their goals seem unachievable, people respond by distancing themselves from their goals (H1), resigning themselves to failure (H2) and that the strength of these tendencies is proportional to the utility they assign to goal achievement (H1b & H2b). The results did not support EMA's prediction that people will wish-away perceived threats (H3 and H3b). Indeed, we saw a slight trend in the opposite direction that was nearly significant. The results also give insight into how to model coping responses to positive situations and illuminates constraints on when coping strategies can be applied.

There are several possible explanations for failure of EMA's predictions concerning wishful-thinking (H3 and H3b). McGuire and McGuire [22] in a series of experiments found a wishful-thinking effect only for events that had high personal significance (i.e., only highly affectively charged events would lead to changes in likelihood). In terms of the model, this might suggest a higher threshold for wishful thinking, generally. It also is relevant to the issue discussed below concerning what coping strategies are chosen in response to an event. Of course, this does not explain the direction of the change. It has been argued that negative affective states lead to more systematic, data driven cognitive processing while

positive affective states lead to more heuristic processing [27], Thus, high motive subjects may have had a more affectively charged reaction to conditions in the game and that impacted their processing of likelihood information.

There were also several findings that point to ways to refine the EMA model. Recall that one research question concerned how people cope with positive emotions and the results give some insight into these processes. Subjects exhibited unpredicted shifts in desires and intentions in the Win condition, assigning increased utility to winning and trying harder to win as their likelihood of success increased. Further, subjects exhibited these coping behaviors depending on the initial utility they assigned to winning: low motive subjects became more engaged as they began to win, whereas high-motive subjects did not; high-motive subjects became less engaged as they began to lose, whereas low-motive subjects did not.

The low-motive versus high-motive differences in distancing and resignation suggest ways to model individual differences in the model. Specifically, one possible explanation of the results is that low-motive subjects reflect a pessimistic thinking style [28] with their low motivation at the beginning perhaps being a coping response to a negative emotional state prior to the start of the game: 'I am not going to win so I should not get too engaged.' In contrast, the fact that highly motivated subjects in the win condition engaged in wishful thinking at T1 may be due to an optimistic thinking style [28] that is less driven by estimates of likelihood [27].

Though confusing on the surface, these different coping effects can be seen as arising from a simple pattern. Regardless of high or low motive, coping serves to move subjects into a more positive emotional state: a) distancing and resignation in the losing high-motive subjects reduces the negative emotionality of a threat to an important goal; b) the over-optimistic wishful-thinking of winning high-motive subjects enhances the positive emotions of obtaining an important goal; and c) the engagement of winning low-motive subjects enhances the positive emotions arising from what was previously seen as an unattainable goal.

Another research question the study hoped to illuminate was the constraints that govern when different coping strategies will be selected the results provide some hints, though no conclusive answers. In EMA, wishful thinking, distancing and resignation, are applicable to situations where the appraised control of the situation is low. However, other factors might differentially impact the applicability of different strategies. For example, the evidence that one is losing the game may become too incontrovertible to allow wishful thinking, or a goal may be too centrally important to distance oneself from it (c.f., Lazarus's notion of emotion-focused coping potential[19]). Further, as the high-motive and low-motive groups suggest, there will be individual differences in preferences for coping strategies: scatter plots of the individual responses verified there are clear individual differences across subjects in which coping strategies they adopted – some subjects preferred just one strategy whereas others used a blend of strategies.

Several issues remain unaddressed by this study and will be the subject of future research on EMA. One prominent issue is the constraints governing belief change. Unchecked, the wishful thinking could allow the agent to wish away even the inevitable to an extreme that might lead to delusional behavioral responses that impact utility in other ways. Another source of constraint on belief change that needs to be considered is the relation to other beliefs, for beliefs are interconnected, influencing and constraining each other (e.g., see [29]).

These limitations notwithstanding, the results provide broad support for EMA's implementation of coping techniques and suggest concrete steps to further improve the fidelity of the model.

7. Conclusion

The last several decades has seen a rapid growth in research on modeling emotions computationally. A key question remains concerning the fidelity of these models. This paper has presented an approach to validating the behavioral fidelity of the EMA computational model of emotion and coping. Taken together with a companion study [26], the results provided support for many of the model's predictions and also clearly identified ways to improve the model. More broadly, the results of this study demonstrate the power of using controlled human subject experiments to improve a model.

ACKNOWLEDGMENTS

Several people contributed to this study. Anya Okhmatovskaia and Wendy Treynor helped with the early design. Rainer Reisenzein helped with many thoughtful suggestions. This work was sponsored by the U.S. Army Research, Development, and Engineering Command and the Air Force Office of Scientific Research under the grant #FA9550-06-1-0206. The content does not necessarily reflect the position or the policy of the Government, and no official endorsement should be inferred.

References

1. Gratch, J. and S. Marsella, *The Architectural Role of Emotion in Cognitive Systems*, in *Integrated Models of Cognitive Systems*, W. Gray, Editor. 2007, Oxford University Press.
2. Sander, D., D. Grandjean, and K.R. Scherer, *A systems approach to appraisal mechanisms in emotion.* Neural Networks, 2005. **18**: p. 317-352.
3. Picard, R.W., *Affective Computing*. 1997, Cambridge, MA: MIT Press.
4. Conati, C. and H. MacLaren. *Evaluating a probabilistic model of student affect*. in *7th International Conference on Intelligent Tutoring Systems*. 2004. Maceio, Brazil.
5. Paiva, A., J. Dias, and R. Aylett, *Learning by Feeling: Evoking Empathy with Synthetic Characters.*

Applied Artificial Intelligence special issue on "Educational Agents - Beyond Virtual Tutors", 2005. **19**(3-4): p. 235-266.
6. Scheutz, M. and A. Sloman. *Affect and agent control: experiments with simple affective states*. in *IAT*. 2001: World Scientific Publisher.
7. Neal Reilly, W.S., *Believable Social and Emotional Agents*. 1996, Carnegie Mellon University: Pittsburgh, PA.
8. Swartout, W., et al., *Toward Virtual Humans.* AI Magazine, 2006. **27**(1).
9. Neal Reilly, W.S., *Modeling what happens between emotional antecedents and emotional consequents*, in *Eighteenth European Meeting on Cybernetics and Systems Research*. 2006, Austrian Society for Cybernetic Studies: Vienna, Austria. p. 607-612.
10. Elliott, C., *The affective reasoner: A process model of emotions in a multi-agent system*. 1992, Northwestern University Institute for the Learning Sciences: Northwestern, IL.
11. Dias, J. and A. Paiva. *Feeling and Reasoning: a Computational Model for Emotional Agents*. in *Proceedings of 12th Portuguese Conference on Artificial Intelligence, EPIA 2005*. 2005: Springer.
12. El Nasr, M.S., J. Yen, and T. Ioerger, *FLAME: Fuzzy Logic Adaptive Model of Emotions.* Autonomous Agents and Multi-Agent Systems, 2000. **3**(3): p. 219-257.
13. Marinier, R., *A Computational Unification of Cognitive Control, Emotion, and Learning*, in *Computer Science*. 2008, University of Michigan: Ann Arbor, MI.
14. Scherer, K.R., A. Schorr, and T. Johnstone, eds. *Appraisal Processes in Emotion*. Affective Science, ed. R.J. Davidson, P. Ekman, and K.R. Scherer. 2001, Oxford University Press.
15. Loewenstein, G. and J.S. Lerner, *The role of affect in decision making.*, in *Handbook of Affective Science*, R. Davidson, H. Goldsmith, and K. Scherer, Editors. 2003, Oxford: Oxford University Press. p. 619-642.
16. Mao, W. and J. Gratch. *Evaluating a computational model of social causality and responsibility*. in *5th International Joint Conference on Autonomous Agents and Multiagent Systems*. 2006. Hakodate, Japan.
17. Gratch, J. and S. Marsella. *Evaluating a General Model of Emotional Appraisal and Coping*. in *AAAI Symposium on Architectures for modeling emotion: cross-disciplinary foundations*. 2004. Palo Alto, CA.
18. Marsella, S. and J. Gratch, *EMA: A Model of Emotional Dynamics.* Journal of Cognitive Systems Research, 2009. **10**(1): p. 70-90.
19. Lazarus, R., *Emotion and Adaptation*. 1991, NY: Oxford University Press.
20. Bratman, M., *Intention, Plans and Practical Reason*. 1987: Harvard University Press.
21. Kunda, Z., *The case for motivated reasoning.* Psychological Bulletin, 1990. **108**(3): p. 480-498.
22. McGuire, W.J. and C.V. McGuire, *The content, structure, and operation of thought systems*, in *Advances in social cognition*, R.S. Wyer and T.K. Srull, Editors. 1991, Lawrence Erlbaum: Hillsdale, NJ. p. 1-78.
23. Kay, A.C., M.C. Jimenez, and J.T. Jost, *Sour Grapes, Sweet Lemons, and the Anticipatory Rationalization of the Status Quo*. 2002. p. 1300-1312.
24. Cialdini, R.B., *Influence: Science and practice*. 4th ed. 2001, Boston: Allyn & Bacon.
25. Gratch, J. and S. Marsella, *A domain independent framework for modeling emotion.* Journal of Cognitive Systems Research, 2004. **5**(4): p. 269-306.
26. Gratch, J., et al., *Assess the validity of appraisal-based models of emotion*, in *International Conference on Affective Computing and Intelligent Interaction*. 2009, IEEE: Amsterdam.
27. Clore, G.L. and J. Storbeck, *Affect as Information about Liking, Efficacy, and Importance*, in *Affective Influences on Social Behavior*, J. Forgas, Editor. 2006, Psychology Press: New York.
28. Seligman, M., *Learned Optimism*. 1991, New WOrk: Simon & Schuster.
29. Thagard, P., *Why wasn't O. J. convicted: emotional coherence in legal inference.* Cognition and Emotion, 2002. **17**: p. 361-383.