

REPORT DOCUMENTATION PAGE

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ADDRESS.

1. REPORT DATE (DD-MM-YYYY) 26-07-2009		2. REPORT TYPE Final Report		3. DATES COVERED (From - To) 01-06-2007 - 30-11-2008	
4. TITLE AND SUBTITLE Affect, Risk and Uncertainty in Decision-Making An Integrated Computational-Empirical Approach				5a. CONTRACT NUMBER FA9550-07-C-0055	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) Eva Hudlicka Gerald Matthews				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Psychometrix Associates, Inc. 1805 Azalea Drive Blacksburg, VA 24060				8. PERFORMING ORGANIZATION REPORT NUMBER Psychometrix Associates Technical Report 0904	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) USAF AFRL/NL AF OFFICE OF SCIENTIFIC RESEARCH 875 NORTH RANDOLPH STREET ARLINGTON VA 22203				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION / AVAILABILITY STATEMENT UU <i>Distribution A: Approved for Public Release</i>					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT We summarize a cross-disciplinary effort exploring affective biases in decision-making. The work consisted of an empirical and a computational modeling study, within the same synthetic task: a search-and-rescue task. The empirical study assessed effects of anxiety on decision-making (route selection). Participants were more sensitive to probabilities of costs and benefits, than to their quantitative values. Both threat and anxious mood induction (under low threat) appeared to increase sensitivity to loss. With a neutral emotion-induction, trait anxiety was associated with a classic selective attention bias. Anxious individuals sampled information on potential costs more frequently than information on potential gains. This bias was eliminated in the anxious emotion-induction condition. In the neutral condition, anxious subjects may frame decisions as requiring vigilance to threat (i.e., elevated attention and analysis), whereas in the anxious condition, the frame is one of escape (requiring less analysis). Computational modeling studies used the MAMID cognitive-affective architecture to construct a process model of anxiety effects: attentional threat and self-bias, and interpretive threat bias. Different levels of anxiety intensities were encoded in different values of architecture parameters, which controlled processing within the architecture modules, yielding results consistent with existing empirical data. The model was also used to construct alternative mechanisms capable of explaining the observed effects, thereby providing a means of generating candidate hypotheses regarding the nature of the processes mediating the biases. Findings make a methodological contribution in demonstrating how experimental emotion-induction can be successfully employed in a task that is longer, more complex and more demanding than those typically used in affective bias research. The data support the validity of the empirical-computational approach of this project. The biasing effects of anxiety cannot be characterized as a global bias towards prioritizing processing of threat. Instead, anxious emotion has several independent effects, tentatively assigned to selective attention, framing and weighting of probabilistic information, that requires modeling within a cognitive architecture comprised of multiple processing modules. The biases revealed in the study suggest that decision-makers may be vulnerable to a variety of potentially damaging biases in conditions characterized by uncertainty and threat, including neglect of the magnitudes of outcome values, and over-attention to costs over benefits.					
15. SUBJECT TERMS Affective biases, anxiety, decision-making, empirical study, computational modeling, mechanisms mediating decision biases					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	18. NUMBER OF PAGES 61	19a. NAME OF RESPONSIBLE PERSON Eva Hudlicka, Ph.D.
a. REPORT UU	b. ABSTRACT UU	c. THIS PAGE UU			19b. TELEPHONE NUMBER (include area code) 540 257 3889

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std. Z39.18

20090904440

Abstract

We summarize a cross-disciplinary effort exploring affective biases in decision-making. The work consisted of an empirical and a computational modeling study, within the same synthetic task: a search-and-rescue task.

The empirical study assessed effects of anxiety on decision-making (route selection). Participants were more sensitive to probabilities of costs and benefits, than to their quantitative values. Both threat and anxious mood induction (under low threat) appeared to increase sensitivity to loss. With a neutral emotion-induction, trait anxiety was associated with a classic selective attention bias. Anxious individuals sampled information on potential costs more frequently than information on potential gains. This bias was eliminated in the anxious emotion-induction condition. In the neutral condition, anxious subjects may frame decisions as requiring vigilance to threat (i.e., elevated attention and analysis), whereas in the anxious condition, the frame is one of escape (requiring less analysis).

Computational modeling studies used the MAMID cognitive-affective architecture to construct a process model of anxiety effects: attentional threat and self-bias, and interpretive threat bias. Different levels of anxiety intensities were encoded in different values of architecture parameters, which controlled processing within the architecture modules, yielding results consistent with existing empirical data. The model was also used to construct alternative mechanisms capable of explaining the observed effects, thereby providing a means of generating candidate hypotheses regarding the nature of the processes mediating the biases.

Findings make a methodological contribution in demonstrating how experimental emotion-induction can be successfully employed in a task that is longer, more complex and more demanding than those typically used in affective bias research. The data support the validity of the empirical-computational approach of this project. The biasing effects of anxiety cannot be characterized as a global bias towards prioritizing processing of threat. Instead, anxious emotion has several independent effects, tentatively assigned to selective attention, framing and weighting of probabilistic information, that requires modeling within a cognitive architecture comprised of multiple processing modules. The biases revealed in the study suggest that decision-makers may be vulnerable to a variety of potentially damaging biases in conditions characterized by uncertainty and threat, including neglect of the magnitudes of outcome values, and over-attention to costs over benefits.

Table of Contents

1.0 Introduction, Objectives and Significance of Proposed Research	1
1.1 Objectives and Research Questions	2
1.2 Summary of Approach	5
1.3 Guide for the Reader	5
2.0 Background Information: Psychological Research in Decision-Making	6
2.1 Affective Biases in Decision-Making and Information-Processing	6
2.2 Uncertainty, Risk, and Stress	7
3.0 Background Information: Computational Modeling of Decision-Making	9
3.1 Modeling Decision-Making in Cognitive Architectures: Approach and Benefits	9
3.2 MAMID Cognitive-Affective Architecture	10
4.0 Task Context: Search-and-Rescue Synthetic Task	14
5.0 Empirical Studies	17
5.1 Search-and-Rescue Task Vignettes for Assessing Effects of Affective Biases on Tactical Decision-Making	17
5.2 Software Developed for Administration of the Empirical Studies	18
5.3 Design and Administration of the Empirical Studies	21
5.3.1 Pilot Study	22
5.3.2 Tactical Decision-Making Study	32
6.0 Computational Modeling	36
6.1 Search-and-Rescue Task 'Vignettes' Used for the Simulation Studies	36
6.2 Simulation Studies Aimed at Identifying Candidate Hypotheses Regarding Affective Bias Mechanisms	37
6.2.1 Modeling Anxiety-Associated Biases on Decision-Making	37
6.4.2 Modeling Anger-Associated Biases on Decision-Making	44
6.4.3 Modeling Anxiety-Associated Biases on Decision-Making	47
7.0 Summary and Conclusions	52
7.1 Summary of Empirical Findings	52
7.2 Summary of Modeling Results	53
7.3 Future Work	53
8.0 References	55
Appendix A: Detailed Description of the Empirical Study Procedures	56

1.0 Introduction, Objectives and Significance of Proposed Research

Current military operational environments are characterized by high information load, uncertainty in both the information and the course-of-action outcomes, and the need for rapid-tempo, high-stakes decisions. These conditions exist at the individual, the team and the organizational levels, and contribute to the decision-maker's stress level, high workload, and mental and physical fatigue, which may adversely bias the decision-making process. Decision-making research over the past two decades has identified affect (emotion) as a key factor in decision-making (Mellers et al., 1998; Loewenstein et al., 2001). Case studies (see Driskell & Salas, 1996) have implicated affective factors, including stress, anxiety and anger, in operator errors across a range of human-machine system contexts, both individual and team, involving the need for rapid action selection under conditions of limited time, high information load, and high uncertainty. These are precisely the conditions that characterize typical Air Force C2 operations.

Decision-maker misperceptions and errors can have disastrous consequences under these conditions (e.g., the USS Vincennes incident). Specifically, the range of affect-induced biases associated with stress may adversely affect the ability to detect the relevant cues, accurately assess the situation and predict its likely course, and interfere with accurate assessment of the tradeoffs involved among the available courses of action. The military has expended considerable effort to better understand decision-making under stress (e.g., the TADMUS project (Cannon-Bowers & Salas, 1998)).

While decision-biases in general, and affective biases in particular, have been studied for decades (e.g., Loewenstein et al., 2001; Mellers et al., 1998; Kahneman et al., 1982), we still lack an understanding of the cognitive and affective mechanisms involved. In-depth understanding of these mechanisms would allow the identification of the individual and contextual attributes that contribute to decision-errors in both individual and team contexts. This would in turn enable the design of more effective human-machine systems for operational contexts, and more effective training environments. For example, understanding the effects of stress and anxiety on the fundamental attentional processes mediating cue detection (bias for threatening cues, neglect of non-threatening cues) can contribute to the design of user interfaces and decision-support systems that can help counteract these deleterious effects (Hudlicka & McNeese, 2002). In-depth understanding of the interpretational threat bias associated with anxiety, and the higher-risk behavioral bias associated with anger, can help improve assessment and training environments, by (a) identifying individuals particularly susceptible to these types of biases, and (b) developing training protocols to counteract them.

The multidisciplinary research described in this final report integrated methods from Cognitive Science and Artificial Intelligence (computational cognitive and affective modeling), and experimental and cognitive psychology. Its aim was to develop a computational model of affective biases, and begin to characterize the mechanisms of affective influences on the structures and processes mediating decision-making, as well as the individual and contextual attributes that contribute to degraded performance associated with anxiety and anger-induced biases.

1.1 Objectives and Research Questions

This report describes the first phase a broadly conceived multi-phase research program, whose objective was to use a combined empirical – computational modeling cross-disciplinary approach to study affect-induced biases in tactical and strategic decision-making. The objective of this program was to develop a comprehensive model of the influence of affective factors on decision-making processes, using both *computational modeling* and *experimental psychological* methods. The *primary aim* was to develop a computational model of affective biases based on empirical data, and outline the requirements to establish its predictive validity. The focus was on the effects of anxiety, frustration and anger, as the primary components of stress. (As there is no empirical work that rigorously distinguishes the constructs of frustration and anger, we focused on the basic emotion of anger to reflect both frustration and anger.) A *secondary aim* was to identify the mechanisms of these biases, across multiple stages of the decision-making process; e.g., attention, situation assessment, expectation generation, goal prioritization, and action selection, as well as biases in working and long-term memory (encoding and recall). We also expected to contribute to the characterization of the mutual influence between affect, and the perception, assessment and management of uncertainty and risk, and begin to identify the mechanisms that mediate these processes. Associated objectives included:

- (1) evaluation of the integrated computational-empirical approach as a means of identifying mechanisms mediating decision-biases;
- (2) exploration of the effectiveness of using an interactive search-and-rescue synthetic task as a vehicle for decision-making research; and
- (3) development of productive, empirically-justified and mechanistically-oriented definitions of stress and risk, and identification of their effects on decision-making.

To meet these objectives we proposed to conduct symbolic computational modeling studies as well as a series of empirical studies with human subjects, aimed at establishing the degree of predictive validity of the computational model, and at an iterative refinement approach to the development and validation of specific hypotheses regarding the mechanisms of affective influences on decision-making. In this iterative refinement approach, the data from the empirical studies would drive the development and fine-tuning of computational models of the hypothesized decision mechanisms, and help quantify the influence of specific affective factors. The resulting models would then generate specific hypotheses regarding the operation of particular decision-biases, and the effects of a range of behavior moderators on these biases (e.g., stress, risk, uncertainty of information), which would then be evaluated and validated in further targeted empirical studies (refer to figures 1-1 and 1-2).

The computational modeling component was built upon an existing cognitive-affective architecture, MAMID (Methodology for Analysis and Modeling of Individual Differences), developed by Hudlicka (2002; 2003). MAMID was designed with the explicit purpose to model the effects of affective states and personality traits on decision-making. It implements a novel method for modeling the interacting effects of multiple affective factors in terms of a set of parameters that control the cognitive processes mediating decision-making. MAMID is distinct from existing cognitive architectures (e.g., Soar, ACT, COGNET) in its emphasis on psychologically-principled, flexible models of the effects of a broad range of interacting affective factors. It is distinct from most current computational models of emotion (e.g., Gratch & Marsella, 2004), in its focus on, and elaboration of, the effects of emotions on cognition, rather than limited to models of appraisal.

MAMID's ability to model affective biases was successfully demonstrated in two domains: an *Army peacekeeping scenario*, where MAMID models different types of commanders and demonstrates distinct behaviors associated with different affective state and trait profiles (Hudlicka, 2003), and a *search-and-rescue team task*, where MAMID models individual team members and demonstrates differences in individual and team performance, as a function of distinct trait and state profiles of the individual players (Hudlicka, 2006b).

Both the computational modeling and the empirical studies components were conducted within the context of the search-and-rescue task, which provided a complex, yet constrained, decision-making environment, with opportunities for a range of decision-types, under varying conditions of risk, uncertainty, and complexity. The team configuration of this task also allows both individual and team focus, in both the modeling and the empirical studies.




Figure 1-1: Overview of the Relationship Between the Empirical Studies and the Computational Modeling Components in the Proposed Research Program




Figure 1-2: Overview of the Proposed Iterative Refinement Approach for Identifying Mechanisms of Affective Factors' Influence on Decision-Making, and the Relationship Among Key Affective Factors and Processes

This longer-term research program was aimed at addressing several research questions, including:

- What are the possible causal mechanisms of affect-induced decision-biases and heuristics, and how are they influenced by risk and uncertainty?
- How do personality traits and affective states facilitate or prevent the expression of particular types of decision heuristics or biases (e.g., framing), for different decisions (e.g., tactical vs. strategic), and under varying conditions of risk and uncertainty?
- How can the improved understanding of affective bias mechanisms contribute to the design of more effective human-machine systems, and training environments for real-time, high-stakes decision-making involving complex tradeoffs?
- What role does affect play in mediating the influence of uncertainty and risk on decision-making, and in decisions involving complex tradeoffs?
- What aspects of the decision-making structures and processes change over time as a function of bias operation?
- How do chronic states of stress contribute to these changes?

1.2 Summary of Approach

The originally envisioned research program consisted of the following goals:

- 1 Develop tasks assessing tactical and strategic decision-making, within the search-and-rescue synthetic task.
- 2 Develop software for administration of empirical studies and performance assessment.
- 3 Augment MAMID cognitive-affective architecture to model tactical and strategic decision-making within search-and-rescue task context.
- 4 Augment MAMID testbed to facilitate model development and 'tuning'.
- 5 Conduct empirical studies assessing affective biases in tactical and strategic decision-making contexts.
- 6 Incorporate findings into MAMID architecture.
- 7 Use MAMID to generate hypotheses regarding bias mechanisms.
- 8 Conduct further targeted empirical studies to validate hypotheses.

This final report summarizes the work conducted to meet goals 1 - 5 and 7 above, with a focus on tactical decision-making.

1.3 Guide for the Reader

This document is organized as follows. Section 2 provides background information on research in experimental and cognitive psychology on the nature of affective biases in decision-making and information processing. Section 3 provides background information on computational modeling of decision-making (section 3.1), as well as a brief description of the MAMID architecture (section 3.2). (Additional information about relevant emotion research in psychology, and the MAMID cognitive-affective architecture, can be found in a related document (Hudlicka 2008). Section 4 describes the task context used to conduct this research, a synthetic search-and-rescue game task, which was used for both the empirical studies and the computational modeling. Section 5 discusses the empirical studies. Section 6 discusses the computational modeling. Section 7 provides a summary and conclusions, highlighting relevance of this research to the Air Force.

2.0 Background Information: Psychological Research in Decision-Making

Below we provide a brief summary of research from experimental and cognitive psychology related to decision-making biases, and decision-making under stress and uncertainty. A more extensive summary of the relevant emotion research in experimental and cognitive psychology can be found in an earlier report prepared during this project (Hudlicka 2008).

2.1 Affective Biases in Decision-Making and Information-Processing

Influence of specific emotions on decision-making has been studied in gambling, social judgments, vehicle operation, medical decision-making and military tactical decisions (e.g., Cannon-Bowers & Salas, 1998). Emotions appear to influence multiple processing components, including encoding of information, reasoning, retrieval of information from memory, and response selection. Four broad types of affective influences may be differentiated. First, real-life decision-makers typically operate in stimulus-rich environments, within which it is easy to neglect critical information. Affective factors influence these *encoding processes*, through narrowing the focus of attention, or through biasing appraisal of risk, threat, and uncertainty. Second, affect may relate to *content biases* that derive from the contents of the cognitive schemata mediating decision-making processes (in contrast to the inferencing processes using these schemata), and represent the values and beliefs influencing perception, situation assessment, and goal and behavior selection. These biases are reflected in the knowledge structures influencing the decision-making process, both static (e.g., schemata in long-term memory), and dynamic (e.g., temporarily activated schemata reflecting current situation assessments and expectations). Third, affect may influence the type and magnitude of biases in *inferencing processes*; e.g., negative emotion influences risk estimation (Johnson & Tversky, 1983). Fourth, emotions relate to *action tendencies* (Frijda, 1987), i.e., preferred styles of response, such as aggressive behaviors in states of anger. In general, it is important to investigate how affect may bias not just the core decision-making processes identified by Kahneman et al. (1982), but also the inputs to decision-making, and preferred choices of action. In studies of reasoning and inferencing, associations have been found between positive emotions and 'assimilative' processing in problem-solving tasks, elaboration of information and creative thought, and between negative moods and 'accommodative' mode of processing, that promotes careful stimulus analysis (Fiedler, 2001). Mood-congruent biases in memory associated with both positive and negative affect have also been observed (Bower, 1981; Isen 1993). Both negative and positive affect have robust mood-congruent effects on self-evaluations and predictions of future benefits and losses (Lerner & Keltner, 2001; Wells & Matthews, 1994).

Specific negative emotions appear to have distinct effects on decision-making (e.g., Nabi, 2003). These include an anxiety-linked threat bias in attention (Williams et al., 1997) leading to a neglect of critical cues (Hartley, 1989), biases in later inferencing processes (e.g., making predictive inferences from threatening material (Calvo & Castillo, 2001)), and apparent promotion of behavioral avoidance (Wells, 2000). Anxiety can also generally degrade attention and performance, by diverting resources from task- to self-related processing. Anger is linked to misappraisal of others' intentions, and false attributions of hostility (Matthews et al., 2000a).

Anger is also linked to impulsive response in confrontational situations. For example, Kassino et al. (2002) modified the Prisoner's Dilemma game to simulate wartime confrontations. Angry players committed more 'competitive attack responses' even when aware that the strategy would lead to losses. Anger and fear are associated with different framing effects; retribution terms for angry individuals, and self-protective terms for anxious individuals (Nabi, 2003).

We focus here primarily on short-duration *state* factors, corresponding to the immediate experience of stress, but stable personality *traits* will also be investigated; neuroticism is associated with a vulnerability to stress and negative affect, and relates to a heightened awareness of danger and a depressed sense of self-efficacy, leading to cautious decision-making in threatening situations (Matthews et al., 2000a). Other traits too may be linked to biases in fundamental processes or social beliefs, and associated functional or maladaptive cycles of interaction with the environment (Matthews et al. 2003).

Although significant progress has been made, typical laboratory studies provide only a limited basis for predicting how affective factors relate to real-life decision-making, in part reflecting the greater sensitivity of complex decision-making to context and domain factors, compared with the simple tasks typically used in laboratory studies. Existing studies have also typically failed to explore dynamic aspects of the inter-relationships between affect and risky decision-making, including the effects of feedback processing, as the decision-maker evaluates the outcomes of prior choices. A key insight of recent research on cognitive architectures capable of modeling affect is that emotions relate to multiple component processes, represented at different levels and stages of information-processing (Ortony et al., 2005). Existing empirical research is not well-suited to exploring the interactions of these multiple processes, which may have synergistic effects that cannot be predicted from a linear summation of the various individual bias effects. Simulation of the operation of multiple biases at different levels and stages of a model that explicitly represents the cognitive architecture may be the most effective means for developing more powerful predictive models of decision-making. It is our hope that a systematic exploration of the different external risk and uncertainty conditions, along with differences in the decision-maker trait and state profiles, using the MAMID cognitive-affective architecture, will contribute towards consistent explanations for the observed empirical data, predictive models, and descriptions of causal mechanisms.

2.2 Uncertainty, Risk, and Stress

A computational modeling approach to decision-making requires precise definitions of the key constructs of interest: *uncertainty*, *risk*, and *stress*. *Uncertainty* plays a large role in real-life decision-making, because the decision-maker lacks knowledge about which loss categories are possible, the probabilities of specific losses occurring, and evidence indicating the likelihood of loss outcomes (Yates & Stone, 1992). Again, modeling may introduce uncertainty into both the simulated environment (e.g., the extent to which outcomes of actions are probabilistically determined), and into internal representations; e.g., as an output of appraisal ("I don't know how severe a threat this is") or in weighting uncertainty information in decision-making ("I will choose the action whose outcomes are most predictable, other things being equal").

Yates and Stone (1992) suggest that *risk* may refer to three, inherently subjective, elements: losses, the significance of losses, and uncertainty associated with those losses. In computational modeling, risk may be associated both with the "objective" simulated environment in which the model operates (i.e., likelihood of some harmful event occurring), and with internal

representations of losses. Such representations may be supported by multiple processing components, including threat appraisals, beliefs about the likely costs and benefits accruing from events, and beliefs about the consequences of actions.

Important real-world decisions are often made under some level of *stress*, e.g., because high stakes attach to the outcome of the decision. Stress may be broadly defined as a relationship between the person and situational demands that taxes or overloads the decision-maker (e.g., Lazarus, 1999), producing negative affect. On the one hand, stress influences judgments of risk, and decisional choice. Case studies suggest that stress and emotion may bias decision-making and willingness to engage in risk-taking behavior (e.g., anxiety may have contributed to Admiral H.E. Kimmel's reluctance to take precautions against a possible Japanese attack on Pearl Harbor: Mann, 1992). On the other hand, decision-making under risk and uncertainty may itself be a source of stress (Loewenstein et al., 2001). Importantly, stress and risk are *dynamically* related: the decision-maker's efforts to cope with stress may influence future external risk, which in turn feeds back to influence stress.

Research on the interplay between stress and risk is hindered by the multi-faceted nature of the stress process, encompassing multiple mechanisms and state and trait factors. We plan to operationalize stress factors primarily as the negative emotional states of anxiety and anger, that may mediate effects of stress on decision biases. Terminology in this area may be confusing due to overlap of terms including affect, emotion, mood and feelings. We will use *affect* as an umbrella term for the field of emotion and subjective feeling states, and *emotion* to refer to coordinated changes in feeling state, and cognitive and psychophysiological functioning elicited by specific events, such as anxiety and anger.

3.0 Background Information: Computational Modeling of Decision-Making

Below we provide a brief introduction to architecture-based models of decision-making (section 3.1) and a description of the MAMID cognitive-affective architecture (section 3.2). A more extensive description of the MAMID architecture, both its structure and functionality, can be found in an earlier report prepared during this project (Hudlicka 2008).

3.1 Modeling Decision-Making in Cognitive Architectures: Approach and Benefits

Mathematical and computational models of decision-making have changed dramatically over the past 40 years (Hudlicka, 2006a). Both the methodologies, and the underlying assumptions about the decision-maker (e.g., 'optimal' vs. 'satisficing'), have evolved, as technological developments became capable of supporting increasingly computationally-intensive, differentiated, and highly-structured models. These developments have led both to advancements in the earlier utility-theory decision models (e.g., Busemeyer, 2007) and to the development of simulation-based, causal computational models. By attempting to emulate the actual cognitive processes and structures mediating decision-making (e.g., attention, situation assessment, goal management, memory), these dynamic models are well-suited for the development of causal mechanisms of decision-making. Depending on the level of resolution and complexity, a given model may represent a single function (e.g., attention, situation assessment), or the entire 'end-to-end' decision-making sequence. These latter models are referred to as *cognitive architectures* (also agent architectures) (see Pew & Mavor (1998) for an overview of many existing cognitive architectures such as Soar, ACT, EPIC, COGNET, etc).

Cognitive architectures have been used both to improve our understanding of human cognition (e.g., Anderson, 1993) and its interaction with emotion (Sloman et al., 2005; Ortony et al., 2005), and for a variety of applications, including user interface design, human-machine system risk assessment, and training (e.g., Kieras et al., 1997; Deutsch & Pew, 2001; Pew & Mavor, 1998; Corker et al., 2000; Dautenhahn et al., 2002). The key benefit of the cognitive architecture approach to modeling decision-making is the associated necessity to operationalize the theoretical hypotheses in terms of detailed specifications of the *structures* (e.g., long-term memory, schemas representing situations, expectations, goals) and *processes* (attention, situation assessment, goal management) mediating decision-making. The development of such detailed, simulation-based models provides opportunities for development and validation of the causal mechanisms of the associated processes, and the factors that influence them, and frequently identifies gaps in knowledge, which can be explored in focused empirical studies. These models also enable the generation of hypotheses regarding specific causal mechanisms, which can then be evaluated in further empirical studies. Computational models thus serve both to *validate existing hypotheses* regarding the causal mechanisms of decision processes and decision biases, and *generate refined or alternative hypotheses* for further empirical exploration (refer to figures 1-1 and 1-2).

3.2 MAMID Cognitive-Affective Architecture

The MAMID cognitive-affective architecture served as the computational model used to conduct simulation studies of affective biases. Its capabilities to generate process-level models of both emotion generation via cognitive appraisal, and emotion effects on cognition, supported the construction of alternative mechanisms for several observed decision biases.

MAMID is a symbolic architecture of high-level cognition, which implements a see-think-do model of sequential, recognition-primed decision-making (with some limited parallelism). MAMID uses Bayesian belief nets as its primary knowledge-representational formalism for the long-term memory (LTM). MAMID dynamically generates emotions via a dedicated Affect Appraiser module, and thus in effect implements a see-[think / feel]-do sequence.

MAMID was built for the explicit purpose of modeling the effects of multiple, interacting affective factors, both traits and states (Hudlicka, 2002; 2003), and is thus well-suited for exploring the mechanisms of the associated decision biases. MAMID implements the sequential the 'see-think/feel-do' decision process in terms of several modules, each corresponding to a distinct stage of decision-making (see figures 3-1 and 3-2). The modules progressively map the incoming stimuli (cues) onto the outgoing behavior (actions), via a series of intermediate internal representational structures (situations, expectations, and goals). The MAMID modules are as follows: *Sensory Pre-processing*, translating the incoming raw data into high-level task-relevant perceptual cues; *Attention*, selecting a subset of cues for further processing; *Situation Assessment*, integrating individual cues into an integrated situation assessment; *Expectation Generation*, projecting the current situation onto possible future states; *Affect Appraiser*, dynamically deriving the affective state from a combination of external and internal stimuli; *Goal Manager*, selecting the most relevant goal for achievement; and *Action Selection*, selecting the most suitable action for achieving the highest-priority goal within the current context. Each module has an associated long-term memory (LTM), consisting of either belief nets or rules, which represent the knowledge necessary to transform the incoming mental construct (e.g., cues for the "Situation Assessment" module) into the outgoing construct (e.g., situations for the "Situation Assessment" module).

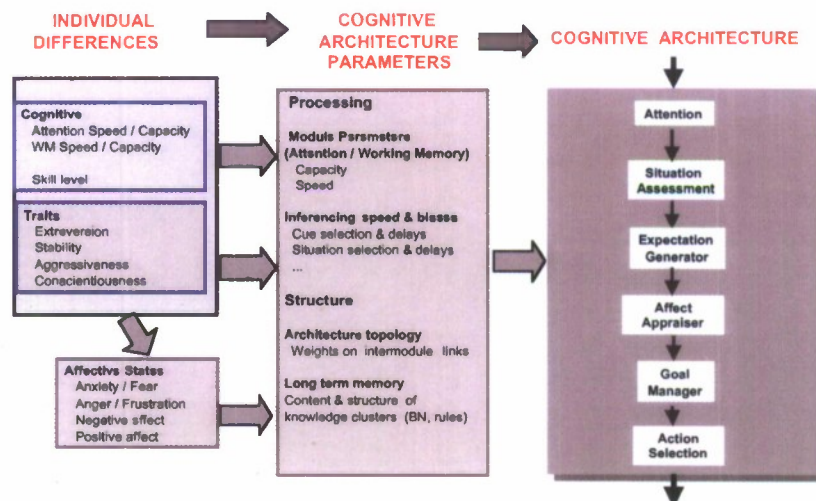


Figure 3-1: Schematic Illustration of MAMID Modeling Methodology




Figure 3-2: Diagram of the MAMID Architecture, Showing the Modules & Mental Constructs

The underlying assumption of the MAMID approach to modeling state and trait effects on decision-making is that the combined effects of a broad range of factors can be modeled *by varying the fundamental properties of the processes and structures* mediating decision-making (Hudlicka, 1997; 2002; 2003). Examples of these ‘fundamental properties’ are the *speed* of the individual modules (e.g., fast or slow attention), the capacities of the working memories associated with each module, and the *content and organization of LTM* (e.g., LTM for situation assessment has a predominance of self- and threat-related schemas for a high-neuroticism individual).

These ‘fundamental properties’ are controlled by a series of parameters, whose values are derived from the decision-maker’s state and trait profile. Modeling different types of decision-makers then requires only changing these individual profiles, rather than the architecture components. The parameters cause ‘micro’ variations in processing (e.g., number and types of cues processed by the Attention Module), which lead to ‘macro’ variations in observable behavior (e.g., high-anxious decision-maker misses a critical cue due to attentional narrowing and selects the wrong action).

The MAMID parameter space thus provides a means of encoding the effects of a variety of interacting individual differences, enabling the development of human decision-making models which provide a basis for modeling the detailed mechanisms of the affective factors’ influence on decision-making, including the role of these factors in risk assessment, uncertainty interpretation and particular decision heuristics and biases. The parameter space also supports accommodation of high-level differences such as those characterizing cultures (e.g., uncertainty avoidance), and effects physiological factors on cognition (e.g., fatigue).




Figure 3-3: MAMID Models of 'Normal' and 'Anxious' Commanders' Decision-Making, Showing Processing Differences Within Each Module in Reaction to Encountering a 'Hostile Crowd', in the "Peacekeeper Scenario" Implementation of MAMID

An initial evaluation in the peacekeeping context established MAMID's ability to model a broad range of interacting individual differences and their effects on individual behavior and task outcome (Hudlicka, 2003). Figures 3-3 and 3-4 illustrate in detail the internal processing of two instances of MAMID architecture, representing a 'normal' and a 'high-anxious' commander encountering a particular problematic situation (hostile crowd) during a peacekeeping mission, and provide a summary of the distinct behaviors produced by the 'normal', 'anxious', and 'aggressive' commanders.

MAMID has recently been transitioned to a different task domain (a collaborative, multi-player search-and-rescue task), where it is used to explore the effects of individual team players' traits and states on both individual performance and overall team effectiveness (Hudlicka, 2006b), for purposes of risk-reduction and safe human-system design. Instances of the MAMID architecture were used to model individual team members with distinct trait/state profiles (e.g., task-focused vs. process-focused leader, high-neuroticism vs. low-neuroticism player), based on empirical studies at NASA-Ames (Orasanu et al., 2003). Experiments demonstrated significant differences in team interactions and task outcome for the different types of individual players. MAMID can thus provide insights into the likely effects of particular personality configurations on team behavior, and thereby contribute to the identification of team configurations best suited for particular task contexts.




Figure 3-4: Summary of Behavior by 'Normal', 'Anxious', and 'Aggressive' Commanders in the "Peacekeeper Scenario" Implementation of MAMID

4.0 Task Context: Search-and-Rescue Synthetic Task

Below we discuss the rationale for selecting the synthetic search-and-rescue task, highlighting its characteristics that make it suitable for exploring the nature of affective biases. The specific task vignettes used for the empirical studies are described in section 5, and the vignettes used for the computational modeling studies, focusing on identification of affective bias mechanisms, are described in section 6.

The choice of an appropriate domain is critical for investigation of affective biases in tactical and strategic decision-making. A major limitation of current research is the historical focus on short-term, tactical decisions, with well-defined options and outcomes. These contexts typically do not provide environments that are sufficiently rich in stimuli, interpretive ambiguities, competing goals, and course-of-action alternatives to provide opportunities for realistic, complex tradeoffs and demonstrate robust affective biases. The selected task must therefore meet several requirements. *First*, it must provide a rich task environment affording detection of dynamic information, and tradeoffs among multiple, competing goals; situations likely to induce affective reactions; decisions involving both information and outcome uncertainties; opportunities for both tactical and strategic decision-making; and opportunities for both individual and coordinated team decision-making. *Second*, since the key aspect of the proposed research program is a systematic comparison of human decision-making with a computational model of these processes, the task must serve the dual role of being a basis for a computational model (i.e., the model must be able to perform the task), and providing the context for the empirical study (i.e., human subjects must be able to perform the task). These criteria dictate that the task provide sufficient complexity to require the range of decision-making outlined above, and yet be amenable to computational modeling, and that the task be sufficiently compelling to support cognitive and affective engagement with human subjects. *Third*, the task simulation must be sufficiently flexible to support the construction of a broad range of specific scenarios, varying in uncertainty, complexity and workload. Together, these characteristics enable the exploration of decision-biases across a range of situations that more closely resemble real-world decision-making contexts, where decision options and outcomes are constrained, but somewhat open ended, to investigate the interplay between affect and decision-making in decision types ranging in time frames, risk, uncertainty, complexity, and associated subjective stress levels.

The use of a synthetic, interactive game-like task has various advantages for this purpose. Computer games and synthetic tasks are a recognized tool for investigating human decision-making, offering greater complexity, realism and participant motivation than standard laboratory tasks (Washburn, 2003; Warren et al., 2004; Parasuraman et al., 2005). Galster et al. (2005) have argued in favor of the use of synthetic task environments in conducting performance-based research to enhance air battle manager capabilities and situation awareness while decreasing workload. Furthermore, manipulations of game events have been shown to induce congruent changes in emotion, appraisal and psychophysiological response (Scheirer et al., 2002; Van Reekum et al., 2004).

The research described here used a synthetic search-and-rescue task (S&R task) that met requirements set out by Galster et al. (2005). These include its applicability to theory-driven research relevant to C2 environments, metrics for rapid evaluations of theory driven constructs, and high degree of experimental control. The task thus afforded study of defining features of C2 contexts, including decision-making in complex, dynamic environments; experimental control

over key constructs of uncertainty, risk and workload; and a focus on both the individual and team interactions. We briefly describe the relevant aspects of the task below.

The S&R task was embedded in an interactive game-like environment, involving one or more simulated players. Multiple game configurations are possible, varying in the nature of specific events and objectives to be achieved by the players, as well as sources and types of data available to accomplish these objectives. For empirical evaluations with human subjects, the task was set up as an individual game-play, focusing on single-frame, single-decision scenarios.

Two geographical contexts were available: an Antarctic scenario involving snowcats, and a Mars scenario involving Mars rovers. The players navigate their vehicles over the inhospitable terrain, and attempt to reach missing members of a previous expedition. The players need resources (e.g., fuel, range of task-specific resources, such as medical, communication, and repair equipment). During the course of the search, the players encounter 'surprise situations', represented as 'tasks', each requiring the expenditure of specific resources (e.g., mechanical breakdown requires a specific number of repair kits), which may need to be replenished at supply stations distributed throughout the terrain. Upon the successful completion of a task, the player is awarded a certain number of points. The players can encounter bad weather and terrain problems, which hinder or prevent travel over a particular terrain segment. Figure 4-1 shows a graphical depiction of a bird's-eye view of the task, along with displays showing additional information about the task status and possible routes. The game format is loosely based on the DDD game developed by Aptima, Inc. (Orasanu et al., 2003).

The interactive task environment, supported by the MAMID testbed, allows the modeler/experimenter to manipulate a range of task variables, including: number of players; total game time available to find the lost party; location of the lost party; characteristics of the tasks such as location, resource requirements, time constraints, points awarded; characteristics of the supply stations, such as location, resources available, availability; channels available for communication among players; and 'broadcast type' messages providing additional game relevant information. The accomplishment of some of the tasks may require collaboration among players (e.g., if one player lacks some of the resources necessary for a task, s/he may ask another player for help). The players interact with the game environment, and each other, via a graphical user interface, which provides necessary information about the surprise events (e.g., location, resource requirements, status).

The S&R task thus provides opportunities for decisions that vary across a range of complexity levels, uncertainty of information and outcome, risk type and magnitude, time frames, and the number and type of tradeoffs required. A number of task features make it especially suitable as a testbed for investigating affective biases, including:

(1) *freedom of action*, allowing choice among differing strategies (e.g., cooperate with other players vs. 'go it alone'), and numerous tactical decisions (e.g., clear vs. bypass blocked terrain; risk running out of fuel vs. delay progress by backtracking to refuel at a supply station). This open-endedness contributes to making decision-making more sensitive to affect (Forgas, 2001), helps maintain motivation, and increase the emotional impact of success and failure outcomes;

(2) *multiple, possibly-conflicting, goals, operating over different timescales, requiring complex tradeoffs* (e.g., select a safe but longer route vs. a faster route with a high-risk of terrain obstruction; delay mission to obtain more resources vs. risk running out of supplies to save time);

(3) multiple, stress-inducing factors including time pressure, task complexity, risk and uncertainty, and social pressures in team contexts;

(4) longer time-frames providing opportunities for multiple, related decisions, and the need for, and opportunities to manifest, *longer-term strategies* and associated affective biases.

Specific configurations of the task, defined in terms of targeted scenarios for assessing or modeling of particular biases, are described below, in section 5.1 (for the empirical study scenarios), and section 6.1 (for the modeling scenarios).




Figure 4-1: Graphical Depiction of the Antarctic Task Environment and Examples of Possible Displays Depicting Additional Task Information for the Empirical Studies

The figure shows the player-controlled snowcats, Tasks ("E", "T", "+"), supply stations (S), and "blocked terrain" (avalanche icon).

5.0 Empirical Studies

This section describes the empirical study component of this effort, which focused on assessing the effects of anxiety and anger on tactical decision-making (goal #5). Emotion was manipulated in each study using inductions previously validated in published research (see below). The objectives of the studies included:

- (1) determining the effects of anxiety and anger on affective biases in tactical decision contexts;
- (2) determining moderating effects of scenario properties (e.g., complexity, risk and uncertainty) and selected personality traits (e.g., neuroticism) as moderators of the biases and processes above;
- (3) investigating the role of affect in dynamic gameplay where emotional response to performance feedback may perpetuate or amplify bias; and
- (4) investigating the mental structures and processes that may mediate emotion effects on the ultimate decisional choices (i.e., identifying possible mechanisms of the biases).

We first describe the tasks developed within the search-and-rescue context that were used to assess these biases (goal #1) (section 5.1), and briefly describe the software developed for administering the experiments (goal #2) (section 5.2). We then discuss the design and administration of the empirical studies in more detail (section 5.3), and the results (section 5.4).

5.1 Search-and-Rescue Task Vignettes for Assessing Effects of Affective Biases on Tactical Decision-Making

The empirical studies used the existing search-and-rescue task context, modified and augmented as necessary to define a series of dynamic situations and scenarios of varying complexity, risk, uncertainty, and tradeoff types and magnitudes (and associated stress levels).

The 'vignettes', representing individual stimuli in the study, consisted of single-frame situation 'snapshots', representing a choice point in the search-and-rescue task, where the players had to select one of several routes to reach a 'lost party' (see example in figure 5-1). Players were required to evaluate the costs and benefits of each route in making their decisions.

This task can be configured with quantitative information to test whether decisions are optimized. The task can also serve to investigate qualitative style of decision-making during extended game-play. Follow-up questions probing the participants' situation assessment, expectancies, goals etc. can then provide information regarding the possible cognitive mechanisms to support detailed process modeling.

This task configuration allows the experimenter to manipulate a broad range of task variables, such as:

- Level of threat and task difficulty
- Degree of certainty associated with incoming information and decision-outcome
- Probability structure of costs and benefits
- Number, type, timing and difficulty of en-route tasks (and associated decisions)

- Cooperative vs. competitive team environment





Figure 5-1: Example of a Stimulus Representing a Single-Frame Decision Point, Presented to the Study Participants

5.2 Software Developed for Administration of the Empirical Studies

A stand-alone experiment administration application was designed and developed, to support the administration of the experimental studies. This consisted of developing the following components: the overall user interface, the map displays consisting of distinct decision vignettes, the capabilities that allowed the participants to obtain additional information about the distinct routes, and the data collection capabilities. Screenshots illustrating the MAMID Experiment software GUI are shown in figure 5-2.

The software enables the experimenter to flexibly specify and display a range of stimuli, varying in degree of risk vs. benefits, uncertainty, and threat level. These variations are accomplished by varying the location and type of routes through the game terrain, initial experimental description and route descriptions, location and type of 'cost' and 'benefit' icons along these routes, and detailed verbal descriptions associated with these routes and icons.

Following the presentation of the stimulus, the participants are presented with a series of follow-up questions in multiple formats. The participants' behavior is tracked by the system, allowing precise control of the amount of time the subject views 'costs' vs. 'benefits' (e.g., total amount of time spent viewing 'cost' icons vs. 'benefit' icons for each route). The responses to the questions are also timed, and the system provides facilities to specify different amounts of time available for each question. Once the experiment is complete, the software calculates mean times for the 'cost' and 'benefit' viewing, and provides a summary output file of these results, along with all other experimental data, for each subject.



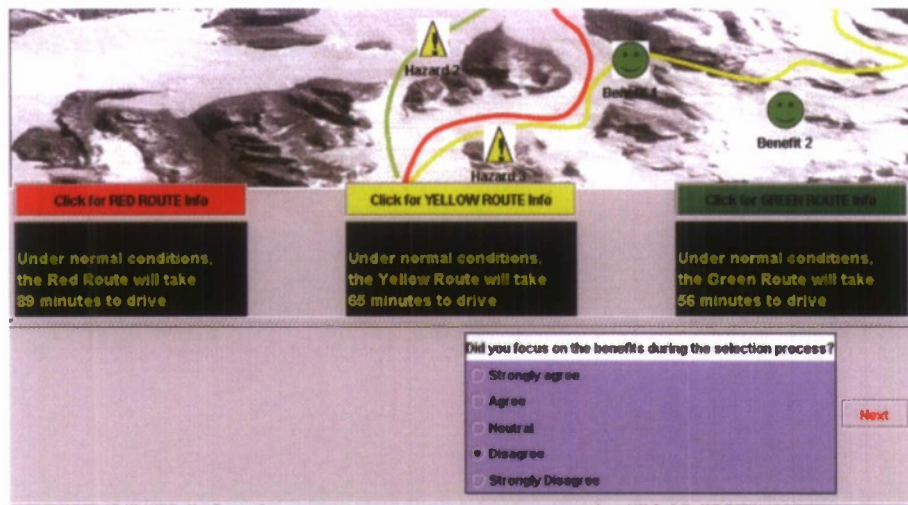


Figure 5-2: Screenshots of the MAMID Experiment Administration Software Illustrating a Single Stimulus (Descriptions of the Routes and Their ‘Costs’ and ‘Hazards’) (top) and an Example of a Follow-Up Question (bottom)

5.3 Design and Administration of the Empirical Studies

Below we first provide an overview of the empirical study tasks, and then describe in detail both the pilot study and the full-scope tactical decision-making study.

The participants must choose between different routes to find the lost party. The aim is to minimize expected travel time to find a 'lost party'. The choices vary in expected travel time, with some choices being clearly superior to others. Choices also vary qualitatively; e.g., the player might be required to choose between a slow but safe route, and a fast but hazardous route.

Participant views a map-like display, from which all information relevant to decision must be obtained (see figure 5-2). Potential costs and benefits of each route are shown as icons (a "smiley" symbol for a benefit, and a hazard symbol used for a hazard). The participant can obtain additional information about the hazard or benefit by using the mouse to 'hover' over the icon. Participant is also informed about baseline travel time for each route, and likelihood of success relative to a specific target time. The specific instructions to the participant are as follows:

- Your objective is to save the lost party.
- You must decide which route will be the best path to take.
- You have a short time to view
 - Benefits
 - Hazards
- Select route descriptions to view the amount of time a particular route will take.

Participant responds by choosing a specific route (from a multiple-choice display, showing the different routes) and then answers questions relating to the situation assessment and his/her emotional state (via multiple choice questions, using the mouse).

The participants' stress level can be manipulated by increasing the time pressure, by not allowing sufficient time to compute expected travel time for each route. This in effect forces the participant to use heuristics & 'intuition' (experiential processing) to make their choice




Figure 5-2: Example of a Map Display Used in the Empirical Studies

5.3.1 Pilot Study

Prior to the tactical decision-making studies, we conducted a pilot study. Its objectives were as follows. *First*, to validate methods for mood induction used in the full-scale studies, and to verify sensitivity of the search-and-rescue scenario to affective bias. (Three moods were induced (happiness, anger, fear, and neutral), using methods validated in existing published studies (Mayer, Allen & Beauregard, 1995), and consisting of guided imagery and music. *Second*, to check usability and difficulty of decision-making task (task difficulty should be moderate so that participants are challenged to distinguish optimal and suboptimal routes), and to ensure that the experiment administration user interface is understandable and easy to use. *Third*, to check sensitivity of task to biases; that is, to determine whether the decisions are sensitive to uncertainty, risk and threat; whether there are any trends towards emotion effects; and whether participants appear to be using 'intuitive' experiential processing.

To accomplish these aims, participants were randomly assigned to one of the mood induction conditions. The mood induction materials consisted of eight vignettes for each mood, which were used as a focus for eliciting the specific mood. Examples of vignettes include enjoying ice cream with friends on a beautiful day (happiness), someone damaging one's car (anger), hearing someone breaking into one's apartment (fear) and doing a week's shopping at the supermarket (neutral). Participants were asked to imagine themselves in the situations described by these guided imagery vignettes. Moods were also enhanced by use of emotional music.

Emotion was assessed at various time-points using the sets of adjectives for basic emotions employed by Mayer et al. (1995). The speed and accuracy of their performance on the decision-making task was measured.

The experiment consisted of a between groups comparison of mood and decision-making in four different conditions: neutral, happy, anxious and angry. Participants were randomly assigned to one of the four conditions. All participants completed a short personality questionnaire, followed by a baseline mood assessment. They then practiced the decision-making task. They were then exposed to the appropriate mood-induction manipulation, and performed the decision-making task. During the period of task performance, participants were exposed to one further mood induction, to maintain the mood induced initially, and they also completed a mood assessment to track mood changes. Participants completed a final mood assessment, after which they were debriefed.

The decision-making task used in the pilot study was as follows. Each participant was tasked with finding a lost party in the Antarctic, by driving a snowcat to their location. The task was made up of a series of discrete items. Each item presented the participant with a map of the terrain, and symbols indicating the positions of the participant and the lost party. Four alternative, color-coded routes were shown. The participant's task was to find the optimal route for reaching the lost party rapidly. Each route carried risks and potential benefits. By use of the mouse, the participant was able to examine the potential costs and gains of each route. Costs related to obstruction of progress, due to terrain and mechanical breakdown. Each cost had a probability and a fixed increase in journey time. For example, there may be a 10% probability of damage to the snowcat due to rough terrain, leading to a time increase of 20 minutes. Conversely, benefits related to enhanced performance of the snowcat, and decreases in journey time. For example, there may be a 20% probability of finding a short cut to reduce the journey, leading to a time decrease of 10 minutes. After assessing the costs and benefits of each route, the participant was asked to choose one of the four, using the mouse to register the choice of route. Following the choice of route, the participant was asked to rate key features of the decision-making problem including its level of risk and uncertainty.

In essence, the task was to choose between alternate routes across an Antarctic landscape in order to minimize travel time. The pilot study aimed primarily to verify that the difficulty and workload of the task was appropriate. Several sets of items were evaluated that presented the participant with qualitatively different choices, such as whether to choose a fast but risky route, or to choose between routes with high and low uncertainty of outcome. The frequency with which the participant chose an optimal over suboptimal routes was assessed, together with qualitative preferences, e.g., for 'risky' or 'safe' routes. Performance data were analyzed to test whether the participant has picked the optimal solution, and for biases in being more strongly influenced by the costs and benefits of each route, depending on the mood.

A detailed description of the study is provided in Appendix A.

Results of the Pilot Study

Results from 40 participants of the pilot study are presented below.

Emotion ratings

The mean rating was calculated for each of the three sets of emotion descriptors, to provide indices of happiness, anxiety and anger. Effects of time of administration of the emotion measure and of mood induction were analyzed using three 4 x 4 (time x induction) mixed-model ANOVAs, with repeated-measures on the time factor. Box's correction was applied in calculating significance levels, because of violations of sphericity; uncorrected dfs are reported here. The critical test is for the time x induction interaction; a significant interaction effect indicates that the time course of emotion was influenced by the manipulation.

The time x induction interaction was significant for happiness ($F(9,108) = 5.82$, partial $\eta^2 = .327$, $P < .01$), anxiety ($F(9,108) = 5.59$, partial $\eta^2 = .318$, $P < .01$), and anger ($F(9,108) = 2.88$, partial $\eta^2 = .193$, $P < .05$). In addition, the main effect of induction was significant for happiness ($F(1,3) = 8.19$, partial $\eta^2 = .406$, $P < .01$), anxiety ($F(1,3) = 2.87$, partial $\eta^2 = .193$, $P < .05$), and anger ($F(1,3) = 2.95$, partial $\eta^2 = .197$, $P < .05$). Main effects of time were significant for happiness ($F(3,108) = 22.53$, partial $\eta^2 = .385$, $P < .01$) and anger ($F(3,108) = 3.74$, partial $\eta^2 = .094$, $P < .05$), but not for anxiety.

Effects of the induction on happiness are shown in Figure 5-3. All groups were similar in happiness initially. Participants tended to experience moderate happiness at baseline; the mean level of endorsement for happiness items on the 1-4 scale was around 2.5. Following the first emotion induction (Time 2), happiness increased in the happiness-induction condition, remained similar in the neutral-induction condition, and declined in the two negative emotion conditions. At time 3 – following the first phase of task performance – some attenuation of happiness in the happiness-induction and neutral-induction conditions was evident. However, group differences were maintained at this time. A similar pattern of group differences was found following the second task performance phase, at Time 4.




Figure 5-3. Effects of emotion induction on happiness at four time points. (Time 1 = Baseline, Time 2 = After first induction, Time 3 = After first task phase, Time 4 = After second task phase).

Figure 5-4 shows effects of emotion-induction on anxiety. Initial levels of endorsement of anxiety items were low in all conditions. Anxiety levels showed little change in the happiness-, neutral- and anger-induction conditions. In the anxiety-induction condition, there was an increase in anxiety at time 2 (following the first induction). The level of anxiety dropped in this condition following the two task phases (times 3 and 4), but anxiety levels remained elevated relative to baseline and relative to the other groups.




Figure 5-4. Effects of emotion induction on anxiety at four time points. (Time 1 = Baseline, Time 2 = After first induction, Time 3 = After first task phase, Time 4 = After second task phase).




Figure 5-5. Effects of emotion induction on anger at four time points. (Time 1 = Baseline, Time 2 = After first induction, Time 3 = After first task phase, Time 4 = After second task phase).

Effects of the emotion-inductions on anger are shown in Figure 5-5. Initial levels of endorsement of anger items were low in all groups. They remained low during performance in the happiness- and neutral-induction groups. Both the anger- and the anxiety-inductions raised anger (time 2). There was some loss of anger during the performance phases (times 3 and 4), but anger remained elevated in both groups relative to baseline and to the other groups.

Workload (NASA-TLX)

Figure 5-6 shows the NASA-TLX profile defined by the mean ratings for the six sources of workload assessed. For comparison, the profile for a vigilance task requiring sustained monitoring of a display is shown. The decision-making task imposes especially high mental and temporal demands. However, it also elicits higher levels of effort than vigilance, and participants attribute less workload to maintaining performance and frustration. Although the task is demanding, it appears to maintain engagement with task demands.




Figure 5-6. Workload profile for the decision-making task, in comparison with a visual vigilance task (N=187; Reinerman et al., 2006). Error bars are standard errors. (MD = Mental demands, PD

= Physical demands, TD = Temporal demands, Perf. = (poor) Performance, Eff. = Effort, Frust. = Frustration)

Effects of the emotion-induction on overall workload were analyzed using a 1-way ANOVA with four levels. The main effect of the induction was significant ($F(3,36) = 3.26$, partial $\eta^2 = .214$, $P < .05$). The highest workload was reported for the anxiety condition (mean = 36.5), followed by anger (33.3), neutral (30.9) and happiness (29.4). Negative emotion inductions appear to elevate workload.

Performance data

Examples of findings are provided here. Three types of performance measure are available.

Performance efficiency. Participants may choose either an optimal or suboptimal route, as defined by expected travel time. Performance efficiency may be indexed as the percentage of route choices that are optimal.

Bias in response. Each of the three item sets (see pp. 3-4) for description) was configured to require the subject to make a qualitative choice, e.g., between a 'safe' or a 'risky' option. Bias

may then be expressed as the percentage choices made for each of the two options, e.g., the percentage of choices that are 'safe'.

Viewing frequencies. The program records the frequencies with which the participant accesses the information on benefits and hazards offered by each route, by moving the mouse over the appropriate icon. On each trial, bias towards viewing hazard information may then be expressed as the percentage: $100 \times \text{frequency of viewing hazards} / (\text{frequency of viewing hazards} + \text{frequency of viewing benefits})$. Note that dwell times for each viewing are also available but have not yet been analyzed.

Mean values for these various indices were calculated for each of the six item sets defined in the previous description of item types (i.e., three item types $\times$ 2 levels of threat; see pp. 3-4).

Performance efficiency

Effects of item type and threat were analyzed using a 3×2 repeated-measures ANOVA. The effect of item type showed a trend towards an effect ($F(2,78)=2.43$, $P=.10$), but there was no trend towards any threat effect. Mean performance levels (% optimal routes) were 58% (*Risky vs. safe options*), 67% (*High vs. low outcome probabilities*), and 59% (*High vs. low uncertainty*). The major feature of these data is that performance levels were not higher for the *Risky vs. safe options* items. These items are simpler than the remaining two item sets, in that there is only a single cost or single benefit to process, by contrast with other items sets, in which the participant must balance a cost against a benefit in evaluating each route. The *Risky vs. safe options* items feature different initial, baseline travel times for the different routes, and this feature may have added to task difficulty. By contrast, baseline times are the same for each route in the other two conditions.

Bias in response

The qualitative biases for each condition were analyzed separately. Table 5-1 shows qualitative biases for each one, in terms of the appropriate percentage choice measures. *t*-tests did not show any significant effects of threat. It is noteworthy that, in two out of three cases, a substantial qualitative bias is evident. For the first item type (*Risky vs. safe options*), there appears to be a strong bias towards choosing safe route (i.e., longer baseline time; chance of a benefit) over a riskier route (shorter baseline time; chance of a loss). The mean number of choices that were risky was 27% in the low threat condition, and 33% in the high threat condition. This may be a 'loss aversion' effect. For the second item type (*high vs. low outcome probabilities*), the bias is in favor of routes that offer a high chance of a small benefit with a low chance of a high cost, as opposed to routes that provide a probable small cost with a high chance of large benefit. This finding may suggest that participants are more sensitive to probabilities of costs and benefits, than to their quantitative values. For the third item type, bias is less marked, with a small trend towards avoidance of high uncertainty options.

Table 5-1. Indices of qualitative bias in response for item sets defined by response alternatives and threat level.

Item Set		Mean	Std. Error Mean
<i>Risky vs. safe options</i> - % risky options chosen	Low threat	.27	.05
	High threat	.33	.03
<i>High vs. low outcome probabilities</i> - % 'probable small loss' options chosen	Low threat	.40	.06
	High threat	.35	.04
<i>High vs. low uncertainty</i> - % high uncertainty options chosen	Low threat	.48	.04
	High threat	.46	.04

*Bias in viewing frequencies***Table 5-2. Indices of bias in viewing hazard and benefit information for item sets defined by response alternatives and threat level.**

Item Type		Mean	Std. Error Mean
<i>Risky vs. safe options</i> - % viewing of hazards	Low threat	.41	.02
	High threat	.44	.02
<i>High vs. low outcome probabilities</i> - % viewing of hazards	Low threat	.56	.02
	High threat	.54	.02
<i>High vs. low uncertainty</i> - % viewing of hazards	Low threat	.56	.02
	High threat	.57	.02

Table 5-2 shows the percentages of viewing responses that were directed towards hazard information, as opposed to benefits, for each of the three item types. *t*-tests did not show any significant effects of threat. Biases are evident, but they vary according to item type. For the first type (*risky vs. safe options*), participants view hazards *less* frequently than benefits, but for the remaining two types, the bias is towards viewing of hazard information. In the first case, it may be that participants are risk-averse, as suggested by the response bias data just described, and so they are less motivated to process risk information. Possibly, the need to balance risks against costs, as required for the other types, increases attention to hazards.

Analysis of emotion effects

In view of the small N, and small numbers of items (4) representing each task condition, strong emotion effects were not expected, but the data have been examined for promising trends to investigate in future research. Effects of both induced emotion and trait measures of anxiety and anger on the Spielberger scale are currently being examined. In general, the performance indices analyzed thus far do not seem to be highly sensitive to induced emotion. One index that may be sensitive to induced emotion is bias towards viewing hazard information in the '*High vs. low outcome probabilities*' item set. Figure 5-7 shows that all three induced emotions appear to produce a bias towards focusing on hazard information, whereas little bias is evident in the neutral condition. A similar, but weaker, effect is also apparent in the *high vs. low uncertainty* item set.




Figure 5-7. Effects of emotion-induction on bias towards viewing hazard information, at two levels of threat (*high vs. low outcome probabilities* item set).

Other analyses have suggested that trait anxiety may relate to biases that are threat-dependent. Trait anxiety relates to preference for risky choices on the *risky vs. safe option* items in the low-threat condition ($r = .41$, $P < .01$), but not in the high-threat condition ($r = .09$). Trait anxiety also

correlates with preference for high-uncertainty choices in the high-threat condition ($r = .49$, $P < .01$), but not in the low-threat condition to preference for high-uncertainty choices in high-threat condition ($r = -.04$), on the *high vs. low uncertainty* items. It is difficult to make much of these findings, given the small numbers of participants and item sets, but they suggest that further investigation of interactions between trait anxiety and threat may be productive.

Conclusions of the Pilot Study

Subjective data

- The emotion-induction manipulation was effective in changing emotion, and effect sizes (partial η^2) were substantial. Emotion change was generally as expected, except that the anger induction also elevated anxiety.
- Induced emotions were attenuated somewhat during task performance, but group differences persisted through both phases of performance
- Data from the 'neutral' condition suggest that the task itself produces loss of initial happiness, but has little effect on negative emotion.
- NASA-TLX data confirmed that the task imposes high demands, but participants maintain effort and performance.
- Negative emotion inductions (especially anxiety) appeared to elevate workload.

5.3.2 Tactical Decision-Making Study

Participants and procedure

120 participants (46 men and 74 women) were recruited from a pool of University of Cincinnati undergraduate psychology students. Ages ranged from 18 to 44 years, with a mean of 25 years. They were randomly allocated to either an anxiety or neutral mood-induction. After completing baseline state and trait anxiety scales, they were exposed to the mood-induction. Next, they practiced and then performed a first version of the decision-making task. Then, they were exposed to a second mood-induction, and performed a second version of the task. The two versions varied in threat, defined as likelihood of completing the mission successfully. Order of low and high threat versions was counterbalanced. The state anxiety questionnaire was repeated following each of the two mood inductions, and at the end of the study. Thus, manipulated independent variables were the mood-induction and threat.

Decision-making task

The participant was tasked with finding a 'lost party' in the Antarctic by choosing the optimal route for a snowcat to follow to their location. The task was made up of a series of 24 discrete items. Each item presented the participant with a map of the terrain, symbols indicating the positions of the participant and the lost party, and a target travel time necessary to save the lost party. Four alternative, color-coded routes were shown. Each route carried risks and potential benefits, each marked by an icon on the map. The participant was able to use the mouse to bring up separate windows showing further information on the potential costs and gains of each route. Only a single window could be viewed at one time, necessitating multiple views of the different

information sources. Costs related to obstruction of progress, due to terrain and mechanical breakdown. Each cost had a probability and a fixed increase in journey time. For example, there might be a 10% probability of damage to the snowcat due to rough terrain, leading to a time increase of 20 minutes. Conversely, benefits related to enhanced performance of the snowcat, and decreases in journey time. Participants could calculate an expected value for each route. On this basis, two routes on each trial were 'optimal' (faster expected travel time), and two were suboptimal.

In addition, the task was designed to require participants to evaluate costs and benefits against one another. Half the options presented small but probable benefits and large but improbable costs, whereas the other half involved likely small losses and unlikely large benefits. In the low threat condition, expected travel times were less than the target time, signaling expected success, whereas in the high threat condition, failure was expected. Dependent variables were (1) accuracy (percentage of optimal routes chosen), (2) preference for routes with small but probable benefits, and (3) frequencies of sampling the cost and benefit information using the mouse to access the relevant icons.

Questionnaires

Trait and state anxiety were assessed using scales from the State-Trait Personality Inventory (STPI: Spielberger & Reheiser, 2004). Workload was measured by calculating the unweighted mean of the six 0-10 rating scales of the NASA Task Load Index (Hart & Staveland, 1988).

Mood-induction

Use of the guided imagery with music mood induction (Mayer et al., 1995) began with subjects listening to a piece of music for one minute. As they continued listening, they next imagined themselves in situations described by guided imagery vignettes presented via Powerpoint on a computer screen at 30 s intervals. Situations for the anxious induction were threatening, whereas those for the neutral induction were mundane. Mayer et al. (1995) and other authors have reported data validating the technique.

Study Results

Three sets of findings are reported:

- (1) effects of the mood manipulation on subjective state,
- (2) effects of threat and anxiety on choice of route, and
- (3) effects of threat and anxiety on attention to costs and benefits.

Mood-Induction on Subjective State

Effects of the mood induction were assessed using analysis of covariance (ANCOVA), using a 2 x 4 (induction x phase) design, with trait anxiety as a covariate. 'Phase' is a within-subjects factor referring to the four time points at which the state anxiety measure was administered. The main effect of the induction and the induction x phase interaction were significant at $p < .01$. The effect of trait anxiety on state anxiety was also significant ($p < .01$), but trait anxiety did not interact with phase. The effects of the mood induction are shown graphically

in Figure 5-8. The neutral and anxiety induction groups were matched on state anxiety initially. State anxiety was elevated by the first induction, and remained about 1 SD above that for the neutral group throughout the study. In addition, a t-test showed that mean workload as assessed by the NASA-TLX was higher for the anxious-induction group than for the neutral-induction group (means: 6.16 vs. 5.78).




Figure 5-8. Significant effects of mood induction on subjective responses to state anxiety items. (standard error bars shown).

Route Choice

Effects of the experimental manipulations on performance accuracy and on route preference were analyzed with two 2 x 2 (induction x threat) repeated-measures ANOVAS. No effects on accuracy were found. For route preference, there was a significant effect of threat ($p < .01$), and of the induction x threat interaction ($p = .05$). There was a general bias towards selection of the route offering a probable small gain (and improbable large loss), but bias was stronger in the low threat (mean = 68.9%) than in the high threat (mean = 59.9%) condition. In the low threat condition, there was a greater bias in the low anxiety (mean = 72.8%) than in the high anxiety (mean = 65.9%) group, but anxiety had no effect under high threat. Further analyses showed no associations between trait and state anxiety and these performance indices.

Attention to costs and benefits

The frequencies with which the cost and benefit icons were accessed were analyzed with a 2 x 2 x 3 (induction x threat x icon) ANOVA, with repeated measures for the threat and icon (cost vs. benefit) factors. The main effect of threat and the threat x icon interaction reached significance ($p < .01$), but there were no significant effects of the mood induction. In general, the hazard icon was sampled more frequently than the benefit icon, but the effect was stronger under low threat (means: 7.38 vs. 6.04) than under high threat (means: 7.13 vs. 6.64).

Table 5-3 shows the associations between trait and state anxiety and the frequencies of viewing benefit and cost items in low and high threat conditions. Data are shown for each mood induction condition separately, and for the whole sample. Three features of the data are noteworthy. First, the influence of anxiety is stronger with the neutral mood. Second, anxiety tends to relate to a higher frequency of sampling in general, especially for trait anxiety. Third, in the neutral condition, trait anxiety relates more strongly to sampling costs rather than benefits.

TABLE 5-3: Correlations of Trait and State Anxiety with Icon Viewing Frequency

	Low Threat		High Threat	
	Benefits	Cost	Benefits	Costs
Neutral mood induction (N=60)				
Trait anxiety	.185	.437**	.232	.399**
State anxiety	.258*	.306*	.243	.329**
Anxious mood induction (N=60)				
Trait anxiety	.242	.027	.185	-.002
State anxiety	.075	.040	.042	-.030
Whole sample (N=120)				
Trait anxiety	.211*	.225*	.212*	.211*
State anxiety	.130	.146	.110	.143

* $p < .05$, ** $p < .01$

The data confirm that mood-induction methods may be used in the human factors context of affective bias in decision-making. The anxiety induction was effective in maintaining elevated anxiety during a complex, high-workload task. Findings also suggest that, within this paradigm, affective biases are subtle, and rather different depending on whether 'affect' is defined by task threat, induced mood, or individual differences in trait and state anxiety. Contrary to expectation, trait anxiety effects were stronger in the neutral rather than the anxious mood-induction.

Data showed some biases that were typical of most participants. They preferred routes that offer a high chance of a small benefit with a low chance of a high cost, as opposed to routes that provide a probable small cost with a high chance of large benefit. Participants may be more sensitive to probabilities of costs and benefits, than to their quantitative values. Both threat and anxious mood induction (under low threat) appeared to increase sensitivity to loss. Preference for small high-probability gains was not simply a function of differential attention to gains and benefits, as participants typically accessed the cost icon more frequently than the benefit icon.

Trait anxiety did not relate to choice of route but it did influence attention to the benefits and cost icons. In the sample as a whole, trait anxiety was associated with a tendency to sample both benefits and costs more frequently. This finding is consistent with Eysenck's (1997) processing efficiency theory, which states that anxious individuals typically compensate for loss of efficiency through increased effort. The effect of anxiety on workload is also consistent with this hypothesis.

In the neutral condition, anxiety was also more strongly related to sampling information on costs than on benefits. This finding resembles the classic bias in selective attention to threat stimuli shown by anxious individuals (Bar-Haim et al., 2007). Perhaps surprisingly, the effect disappears with the anxious mood induction. One explanation is that the affective context provided by the induction influences framing and strategy. In the neutral condition, anxious subjects may frame decisions as requiring vigilance to threat (i.e., elevated attention and analysis), whereas in the anxious condition, the frame is one of escape (requiring less analysis).

At an applied level, the data suggest that decision-making in complex, uncertain environments, such as search-and-rescue, may be sensitive to a variety of emotional biases, consistent with computational models (Hudlicka, 2004). Individual differences in anxiety appear to influence active search for information on potential threats. Although anxiety did not affect the overall quality of decision-making, under other circumstances, a focus on threat might have either beneficial or harmful consequences. In operational settings it may be important to monitor for such biases to ensure that decision-makers are sampling different information sources optimally.

6.0 Computational Modeling

This section describes the computational modeling component of the research conducted under this effort, which focused modeling the mechanisms of selected affective biases (anxiety and anger), focusing primarily on goal #7 (see section 1): identifying candidate hypotheses for observed affective biases. We first describe the search-and-rescue task ‘vignettes’ that were used for the computational modeling studies (section 6.1). We then describe the simulation experiments conducted and the results (section 6.2).

6.1 Search-and-Rescue Task ‘Vignettes’ Used for the Simulation Studies

The modeling simulation studies used several simplified scenarios within the tactical decision-making task context, analogous, but not identical, to the scenarios used for the empirical studies. For the anxiety-bias modeling studies, an anxiety-targeted scenario was defined to demonstrate MAMID’s ability to model several of the anxiety-linked biases on cognitive processes; specifically: prioritizing of self- and high-threat cues and other mental constructs (situations, expectations), in conditions of high state and trait anxiety, where trait anxiety was defined as a condition of low extraversion and high neuroticism. The scenario (shown in figure 6-2) provides opportunities for selecting higher-threat cues and situations related to the emergency task (“E” in the figure), or the lower-threat cues and situations related to the supply station task (resupplying resources) (“S” in the figure). The scenario also provides opportunities for generation of intense emotional states, which are represented by self (vs. task) mental constructs, and as such provides opportunities for demonstrating self-biased processing.




Figure 6-1: Map Illustrating the Simple Scenario Focused on Demonstrating Trait- and State-Anxiety Biases: “E” (Emergency Task) is Inherently More Threatening Than “S” (Supply Station Task)

This task configuration was then further modified to model anger-linked biases. Since anger is associated with increased risk tolerance and impulsive action, the scenario needed to provide opportunities for these types of biases, and their consequences. To this end, we introduced a modification in the point accumulation algorithm, where points could be not only gained, upon successful completion of a specific task (e.g., finding the lost party, completing the emergency task), but also lost, if the agent attempted to process a task without adequate resources. The task could be configured with different configurations of available resources, providing opportunities for different levels of point loss or gain, as a function of the agent's affective profile, both state and trait.

Descriptions of the simulation studies using these scenarios are provided in section 6.2 below.

6.2 Simulation Studies Aimed at Identifying Candidate Hypotheses Regarding Affective Bias Mechanisms

Below we describe several simulation studies where MAMID was used to model selected affective biases, and generate possible alternative hypotheses regarding their mechanisms. The studies focus on affective biases associated with anxiety and anger. The anxiety biases focus on trait- and state-anxiety related attentional and situation assessment biases for self (vs. task or other) and high-threat cues and situations. These capture the documented effects of anxiety on threat-biased interpretation of ambiguous stimuli and predictions of negative future outcomes.

The anger biases focused on reduced risk tolerance.

Data from existing literature were used for this modeling phase. Future modeling studies will use data from the empirical studies.


Two agent stereotypes were defined for these studies: aggressive-angry (trait-state) agent, and anxious (trait-state) agents. The behavior of these agents was modeled within the scenarios outlined above, which provided opportunities for both rewards (point gain) and punishment (point loss). Points were gained by successfully completing a task. Points were lost if an agent attempted to process a task without adequate resources. The scenarios were run with different configurations of available resources, and the agent's ability to re-supply. The modeling studies demonstrated differences in internal processing, performance and task outcome across distinct agent types, as a function of the task configuration. Section 6.2.1 describes the anxiety-focused simulation studies. Section 6.2.2 describes the anger-focused studies. Section 6.2.3 describes the studies focusing on modeling alternative mechanisms of selected biases.

6.2.1 Modeling Anxiety-Associated Biases on Decision-Making


Below we first describe how MAMID models the specific anxiety-linked effects (threat and self bias), and then describe a simulation study where these biases resulted in differences in task outcomes, as a function of specific task configurations.

Figure 6-2 shows the output of the Attention Module buffer for the "Normal" player and the "Anxious" player. The buffers illustrate the anxious player's preference for self cues (related to its own affective state) and for high-threat cues. Figure 6-3 shows the output of the Situation Assessment buffer for these players. Again, the bias toward self and high-threat cues is evident in the ordering of the situations. In both cases notice also the generally higher threat assessments


for the threatening cues and situations. Buffers for cycles 1 and 2 are shown, indicating the differences in the cues and situations as the players approach the tasks.




Attention Module output: CUES: Cycle 1: Normal Player



Attention Module Output: CUES: Cycle 1: Anxious Player




Attention Module output: CUES: Cycle 2: Normal Player




Attention Module output: CUES: Cycle 2: Anxious Player


Figure 6-2: Output of the Attention Module, Showing the Ordering of Cues for the Normal (upper) and Anxious (lower) Players as They Approach the Emergency Task (E) and the Supply Station Task (S). Note the difference in the ranking of the self cues and the threatening cues. A trait and state anxiety-related bias towards self and high-threat cues is demonstrated in the anxious player. Output buffers are shown for cycles 1 and 2.




Situation Module Output: SITUATIONS: Cycle 1: Normal Player



Situation Module Output: SITUATIONS: Cycle 1: Anxious Player



Situation Assessment Module output: SITUATIONS: Cycle 2: Normal Player



Situation Assessment Module output: SITUATIONS: Cycle 2: Anxious Player

Figure 6-3: Output of the Situation Assessment Module, Showing the Ordering of Situations for the Normal (upper) and Anxious (lower) Players as They Approach the Emergency Task (E) and the Supply Station Task (S). Note the difference in the ranking of the self situations and the threatening situations. A trait and state anxiety-related bias towards self and high-threat situations is demonstrated in the anxious player. Output buffers are shown for cycles 1 and 2.

To conduct the simulation studies of anxiety effects, we defined a personality profile for a high-trait anxious agent, characterized by high degree of neuroticism and low extraversion. The affective dynamics for this agent, during the performance of the task described above (see figure 6-1), are shown in figure 6-4. The task contingencies were configured such that the agent had no resources, was unable to resupply due to closed supply stations, and was therefore unable to process the emergency tasks encountered. Note the high-intensities of negative emotions (anxiety, negative affect) that persist throughout the period when the agent is within range of the emergency tasks and closed supply stations, as well as the persisting high intensities once the "Lost Party" task is reached.

The labels above the figure refer to the supply station status (SS open or closed), the distance from the emergency task (E-WVR = emergency task within visible range; E-WPR = emergency task within processing range) and the agent's ability to process the emergency task (E = able / unable to process). Resources Adequate (ResAdeq for E (emergency) and LP (Lost Party)).




Figure 6-4: Differences in Affective Profiles for a High Trait-Anxious Agent, As a Function of Task Context.

Figure 6-5 shows the affective profile of the same agent but with the task contingencies, configured for adequate resources and ability to resupply, thus enabling the agent to successfully process the emergency tasks and the "Lost Party" task. Note the transient peaks in negative affect as the emergency tasks (and the Lost Party task) are approached, but which immediately disappear as the tasks are successfully processed.

The labels above the figure refer to the supply station status (SS open or closed), the distance from the emergency task (E-WVR = emergency task within visible range; E-WPR = emergency task within processing range) and the agent's ability to process the emergency task (E = able / unable to process). Resources Adequate (ResAdeq for E (emergency) and LP (Lost Party)).




Figure 6-5: Differences in Affective Profiles for a High Trait-Anxious Agent, As a Function of Task Context

6.2.2 *Modeling Anger-Associated Biases on Decision-Making*

Anger is associated with increased tolerance of risk and reduced loss aversion. In MAMID, these biases are modeled within multiple modules, by including factors in the construct rank calculation that reflect a preference for, or disregard of, high-risk constructs, and a reduced sensitivity to loss. Figure 6-6 shows examples of these biases in the Behavior Selection module. Specifically, the rankings of behaviors produced by the Behavior Selection module are shown for two agent stereotypes: aggressive-angry (top), and an anxious (bottom). The two behaviors are associated with different levels of risk: a high-risk behavior: "Process Emergency Task", and a low risk behavior: "Re-supply from Supply Station".

The contents of the Behavior Selection buffer show the rank of these behaviors, and some of factors that contribute to this rank, most notably the risk level (both are circled in red above). For the angry agent stereotype (top), the higher-risk behavior ("Emergency Task") is ranked more highly, and thus executed first. In contrast to this, the anxious agent's buffer shows that this high-risk behavior is ranked lower, indicating a risk-avoidance tendency associated with anxiety, and resulting in the Re-supply behavior being executed first.

These small changes within the individual modules eventually result in distinct task outcomes. In this case, the angry agent's tendency to choose the high-risk behavior leads to an attempt to process the emergency task without sufficient supplies, and results in loss of points and task delay. In contrast, the anxious agent's preference for low-risk behavior causes it to re-supply first, which leads to adequate supplies and successful processing of the emergency task. The anxious agent thus not only finishes the task in less time, but also with more points, because it does not have points deducted for attempting to process a task without adequate resources.

Figure 6-7 illustrates the differences in the affective dynamics during the evolving scenario, for the "angry" agent profile, as a function of different task configurations.

Object	State	Value	Recipient	Rank	Desirab...	Confidence	Type	Source	Risk
Red Snow Cat Driver	Process task - Emergency	Yes	task - Emergency	6.70	0.30	0.40	Task	External	5.00
Red Snow Cat Driver	Resupply - All kits	Yes	SupplyStation 1	3.70	0.30	0.40	Task	External	2.00
Red Snow Cat Driver	Move Forward Normal	Yes		3.70	0.30	0.40	Task	External	2.00

“ANGRY” AGENT Behavior Buffer


Angry Agent ranks the Processing of the Emergency task higher than Re-supplying, because of its preference for high-risk behaviors, causing it to attempt processing with inadequate supplies. This results in a loss of points and an overall task delay.

Object	State	Value	Originator	Recipient	Comperatr	Rank	Desirability	Risk	Confidence	BN Probability
Red Snow Cat Driver	Resupply - All kits	Yes		SupplyStation 2		0.90	0.80	2.10	0.10	1.00
Red Snow Cat Driver	Process task - Emergency	Yes		task - Emergency		0.00	0.00	5.10	0.10	1.00


ANXIOUS AGENT Behavior Buffer

Anxious Agent ranks the processing of the Re-supply task higher than Emergency task, because of its preference for low-risk behavior, causing it to re-supply first, and having the available resources to successfully process the Emergency Task. This results in a point gain and reduced overall task time.


Figure 6-6: Differences in Anger and Anxiety Biases in Behavior Selection, Resulting in Different Task Outcomes



Aggressive-Angry Agent: Has supplies and is able to re-supply once when needed (cycle 10, to process emergency task).



Aggressive-Angry Agent: Does not have supplies but is able re-supply as needed (in cycles 3, 10, and 18, to process 1st and 2nd emergency task and the Lost Party task).



Aggressive-Angry Agent: Does not have supplies and is unable re-supply, causing it to fail on all three tasks (in cycles 3, 10, and 18 - 1st and 2nd emergency task and the Lost Party task).

Figure 6-7: Differences in anger intensities for an aggressive-angry agent, as a function of resource availability, ability to re-supply, and ability to successfully process the tasks (2 Emergency tasks in cycles 3 and 10, and a Lost Party task in cycle 18).

6.2.3 Modeling Anxiety-Associated Biases on Decision-Making

Below we describe how MAMID can model multiple mechanisms of anxiety effects, for anxiety intensities ranging from low to extreme states, such as a panic attack. Anxiety was selected because of its direct relevance for decision-making and behavior, because robust empirical data regarding its effects are available, and because anxiety emerged as the most significant effect in the empirical validation studies (Panganiban, Matthews & Hudlicka, 2009).

Panic attack is an interesting state to explore because its extreme nature provides a useful context in which to model the effects of anxiety on cognition, and cognition-emotion interaction in general. Panic attack is a state where the confluence of multiple anxiety effects produces a type of a 'perfect storm', frequently inducing behavioral paralysis. Three anxiety-linked effects are involved: *threat processing bias*, *self processing bias*, and *capacity reductions in both attention and working memory*. MAMID models all three of these effects, and provides parameters that control their relative contributions to the overall effect on processing.

Recall that a given parameter value is the results of a linear combination of the weighted factors influencing the parameter. The same global effect (e.g., reduced module capacity) can thus be obtained from multiple combinations of factor values and weights. These alternative configurations then provide the means of defining alternative mechanisms mediating specific

effects. MAMID provides facilities that support the rapid construction of these alternative mechanisms, via interactive manipulation of the factors and weights, which allow the modeler to control the magnitude and contribution of each influencing factor.

Threat bias is modeled by first calculating the threat level of each cue, situation and expectation, from factors that include an a priori 'fixed' threat level (e.g., low level of resources is inherently more threatening than adequate resources), state and trait anxiety factors, and individual history. The threat level is then used as a weighted factor in the function calculating the overall construct rank, which determines the likelihood of its processing. In states of high-anxiety, high-threat constructs have a higher ranking, and are thus processed preferentially (cues attended, situations derived) (refer to figure 6-8).

Self bias is modeled by including a weighted factor reflecting the self vs. non-self origin of each construct in its rank calculating function. High levels of state or trait anxiety then induce a higher ranking for self-related constructs, contributing to their preferred processing.

The *capacity reduction* effects on attention and working memory are modeled by dynamically calculating the capacity values of all modules during each execution cycle, from weighted factors representing the emotion intensities, the four traits represented in MAMID, baseline capacity limits, and skill level.




Figure 6-8: Modeling Threat Bias Within MAMID




Figure 6-9: Affective Dynamics and Module Capacity Parameters (top) Associated with a “Panic Attack” State (Cycle 8), and Subsequent Reduction in the Number of Constructs Processed (bottom)

MAMID’s ability to model alternative mechanisms of anxiety effects was demonstrated in the context of the search-and-rescue vignette shown in figure 6-1. Briefly, the agent’s task is to find a “lost party” in an inhospitable terrain, where “emergency situations” arise unexpectedly. The agent may need to obtain supplies from available “supply stations”, to maintain adequate resources (fuel, first aid kits). In the experiment described below, the agent approaches a difficult “emergency situation”, and lacks the required resources. The agent’s state of anxiety, dynamically calculated by the Affect Appraiser module, is high; in part because of a trait-induced tendency towards higher anxiety, and in part because of the difficult task ahead and lack of adequate resources.

Within this context, MAMID models a panic attack state as follows. Stimuli, both external and internal, arrive at the Attention Module, whose capacity is reduced. Because of the threat- and self-bias, self-related high-threat cues are processed preferentially, in this case resulting in the agent’s focus on a self-related anxiety cue (see figure 6-9, lower left). This cue, reflecting the agent’s anxious state, consumes the limited module capacity, leading to the neglect of external and non-threatening cues (e.g., proximity of a supply station). This results in a continued self- and threat-focus in the downstream modules (Situation Assessment and Expectation Generation). No useful goals or behaviors can be derived from these constructs, and the agent enters a positive feedback-induced vicious cycle (an endless self-reflection), where the reduced-capacity and biased processing excludes cues that could lower the anxiety level and trigger adaptive behavior. Figure 6-9 shows a diagram of the emotion intensities and module capacities, and representative contents of the cue and situation buffers, providing input to Attention and Situation Assessment modules, respectively.

The model parameters are then modified to increase attentional and processing capacities, thereby enabling the processing of additional cues. This allows the agent to begin processing a larger set of incoming cues, which eventually result in a decreased state of anxiety, and trigger task-related goals and associated task-relevant behavior. Refer to figure 6-10.

A number of factors can be modified to induce the effects described above, simultaneously or sequentially, reflecting multiple, alternative mechanisms mediating the anxiety biasing effects. In the case of the capacity parameters, alternative mechanisms can be defined from the agent's overall sensitivity to anxiety (reflected in the weights associated with trait and state anxiety intensity factors), the baseline, 'innate' capacity limits (reflected in the factors representing the minimum and maximum attention and working memory capacities), and the anxiety intensity itself. This factor can be further manipulated via the set of parameters influencing the affect appraisal processes, including the nature of the affective dynamics (e.g., maximum intensity, and the intensity ramp-up and decay functions).




Figure 6-10: Changing the weights of the anxiety factor in the capacity-calculating functions for attention and working memory from -1000 (in figure 6-9) to -100, led to a sufficient increase in the module capacities to enable the agent to process information other than its own affective state, enabling the agent to derive subsequent goals and behaviors (e.g., "Move forward"). The upper part of the figure shows the cues, situations, goals and behaviors generated at the beginning of the task. The lower part of the figure shows the affective profile and the capacity parameter values generated during each cycle of the task.

MAMID's abilities to model alternative mechanisms of emotion effects are continuing to be evaluated, with a focus on anxiety and anger. The on-going empirical study with human subjects is demonstrating anxiety biases in both information seeking and behavior selection, within the search-and-rescue task context. The results are being used to tune the MAMID parameters, as outlined above. Since the parameters correspond to specific psychological variables or functions, we hope that this parallel empirical-computational approach will provide a useful means for validating MAMID models of emotion effects on cognition.

7.0 Summary and Conclusions

7.1 Summary of Empirical Findings

Key results from the empirical study were as follows:

- An experimental induction of anxiety using a combination of guided imagery and music was highly effective in producing a sustained elevation in state anxiety, measured with the Spielberger State-Trait Anxiety Inventory (STAI), during performance of a demanding tactical decision-making task
- Data confirmed a finding from pilot data that participants were more sensitive to probabilities of costs and benefits, than to their quantitative values.
- The impact of affect is subtle, and depends on whether 'affect' is defined by task threat, induced mood, or individual differences in trait and state anxiety.
- Both threat and anxious mood induction (under low threat) appeared to increase sensitivity to loss.
- With a neutral emotion-induction, trait anxiety was associated with a classic selective attention bias. Anxious individuals sampled information on potential costs more frequently than information on potential gains. This bias was eliminated in the anxious emotion-induction condition. In the neutral condition, anxious subjects may frame decisions as requiring vigilance to threat (i.e., elevated attention and analysis), whereas in the anxious condition, the frame is one of escape (requiring less analysis).

Implications of the study are as follows:

- Findings make a methodological contribution in demonstrating how experimental emotion-induction can be successfully employed in a task that is longer, more complex and more demanding than those typically used in affective bias research
- The data support the validity of the empirical-computational approach of this project. The biasing effects of anxiety cannot be characterized as a global bias towards prioritizing processing of threat. Instead, anxious emotion has several independent effects, tentatively assigned to selective attention, framing and weighting of probabilistic information, that requiring modeling within a cognitive architecture comprised of multiple processing modules.
- The biases revealed in the study suggest that decision-makers may be vulnerable to a variety of potentially damaging biases in conditions characterized by uncertainty and threat, including neglect of the magnitudes of outcome values, and over-attention to costs over benefits.

- A study of bias in an Air Force relevant synthetic task environment is in progress, in collaboration with Dr. Benjamin Knott of the Air Force Research Laboratory at Wright Patterson AFB. The study is examining biases in attention and decision-making during an Air Battle Management (ABM) task of more direct operational relevance. Participants are required to defend assets from incoming enemy aircraft that present various degrees of threat and uncertainty. Participants are allocated to high or low anxious two-person teams on the basis of STAI trait anxiety scores. Data are not yet available from this study, but it is expected that anxious teams will show impairments in decision-making linked to over-attention to threat.

7.2 Summary of Modeling Results

The modeling simulation studies demonstrated MAMID's capability to model effects of anxiety- and anger-linked biases (e.g., self and threat bias in processing for anxiety, increased risk tolerance for anger) at the micro-level; that is, on the low-level processes comprising the cognitive-affective architecture processing sequence. These micro-effects then resulted in observable differences in task outcomes, in terms of time required for task completion, as well as the total points accumulated, including whether or not the task was completed successfully.

The simulation studies also demonstrated MAMID's ability to model alternative mechanisms for the same observed phenomenon, by manipulating the parameter-calculating functions of the architecture parameters defining the modules' capacities, and the ranking of the mental constructs processed by each module.

Together, these capabilities demonstrate MAMID's effectiveness in modeling both the overall effects of emotions on cognitive processes comprising decision-making, as well as its ability to generate detailed hypotheses regarding possible mechanisms of these effects.

Future plans include further validation studies of the MAMID model, using emerging data from empirical studies, as well as its use to help define the nature of the mechanisms mediating affective biases.

7.3 Future Work

Work performed under this contract has provided the basis for an ongoing collaboration with Drs. Ben Knott and Gregory Funke at the Air Force Research Laboratory, Wright-Patterson AFB. Anxiety may bias judgment and decision in military decision-making. A study is in progress using a simulated air defense task, in which two-person teams are tasked with defending various assets against attack from enemy aircraft. Some aircraft are initially of unknown status so that effective selective attention is necessary for monitoring both actual and potential threats. Teams have been selected to be high and low in trait anxiety. In addition, anxiety is manipulated experimentally by using the technique validated by the grant-funded research. The study will demonstrate whether anxiety is associated with suboptimal patterns of attentional deployment. The impact of the research will be to demonstrate whether operator anxiety is a factor to accommodate within air defense displays, and will indicate possible countermeasures such as adaptive aids to attention that might be explored in further studies.

A further collaboration is in progress with the Dr. Richard Roberts of the Educational Testing Service. Funding is being sought for a study that would test whether emotional intelligence, measured by a situational judgment test, relates to vulnerability to emotional basis on the search-and-rescue task developed during the research.

8.0 References

- Bar-Haim, Y., Lamy, D., Pergamin, L., Bakermans-Kranenburg, M. J., & van Ijzendoorn, M. H. (2007). Threat-related attentional bias in anxious and nonanxious individuals: A meta-analytic study. *Psychological Bulletin*, 133, 1-24.
- Eysenck, M. W. (1997). *Anxiety and cognition: A unified theory*. Hove, England: Psychology Press.
- Hart, S.G., & Staveland, L.E. (1988). Development of a multidimensional workload rating scale: Results of empirical and theoretical research. In P.A. Hancock & N Meshkati (Eds.), *Human mental workload* (pp. 139-183). Amsterdam: Elsevier.
- Hudlicka, E. (2004). Beyond cognition: Modeling emotion in cognitive architectures. In M. Lovett, C. Schunn, C. Lebiere, & P. Munro (Eds.), *Proceedings of the Sixth International Conference on Cognitive Modeling: ICCCM 2004: Integrating Models*, pp. 118-123. Mahwah, NJ: Lawrence Erlbaum.
- Loewenstein, G.F., Weber, E.U., Hsee, C.K & Welch, E. (2001). Risk as feelings. *Psychological Bulletin* 127, 267-286.
- Matthews, G., Deary, I.J., & Whiteman, M.C. (2003). *Personality traits* (2nd ed.). Cambridge: Cambridge University Press.
- Mayer, J.D., Allen, J.P., & Beauregard, K. (1995). Moodinductions for four specific moods: A procedure employing guided imagery vignettes with music. *Journal of Mental Imagery*, 19, 151-159.
- Spielberger, C.D., & Reheiser, E.C. (2003). Measuring anxiety, anger, depression, and curiosity as emotional states and personality traits with the STAI, STAXI and STPI. In M.J. Hilsenroth and D.L. Segal (Eds.). *Comprehensive handbook of psychological assessment, Vol. 2: Personality assessment*, pp.70-86. Hoboken, NJ: John Wiley. .

Appendix A: Detailed Description of the Empirical Study Procedures

Participants

Participants were drawn from the introductory psychology pool at the University of Cincinnati. They received course credit for participation. The data presented here are based on a sub-sample of 40 (33 females and 7 males; age range 18-30). Note that the gender ratio is typical of the subject pool; 1-2 males were allocated to each emotion condition.

Emotion-induction

The technique used is a combination of guided imagery and music, initially developed and validated by Mayer et al. (1995), and currently a leading method for inducing emotions or basic moods (Marzillier & Davey, 2005). The participant listens to emotive music, while imagining themselves in situations described by guided imagery vignettes. The advantages of this technique for the present research are that (1) Mayer et al. (1995) validated the induction against behavioral data, (2) music may be continued into the task performance session to maintain the emotion, and (3) the qualities of music that confer emotion appear to be well-understood (Kallinen, 2005). Mayer et al. (1995) list the specific pieces of music and vignettes to be used to induce anger, fear and happiness. Marzillier and Davey (2005) developed comparable materials for a neutral emotion induction. Mayer et al. (1995) report three studies that demonstrated the efficacy of the technique for each of the three emotions of interest.

Use of the guided imagery with music emotion induction (Mayer et al., 1995; Marzillier & Davey, 2005) begins with subjects listening to a piece of music for one minute. As they continue listening, they next imagine themselves in situations described by guided imagery vignettes presented via Powerpoint on a computer screen at 30 s intervals. Vignettes for happy, angry, and anxious emotions were identical to those used by Mayer et al. (1995); vignettes for the neutral emotion were taken from Marzillier and Davey (2005).

Instructions were similar to those of Marzillier and Davey (2005), as follows:

"I am going to ask you to enter a (name of target) mood. I will ask you to listen to some music for one minute, after which I will start the computer program that involves reading some sentences. The music will continue to play throughout the experiment. Please listen to the music, read the sentences and try to get into the mood as much as possible. Please try and stay in that mood for the duration of the experiment. After the screen goes blank, you will be asked to rate your mood. Although I would like you to try and get into the mood, please answer these questions honestly and don't just say what you think would like you to say. You may ask the experimenter to terminate the procedure at any time if you do not feel comfortable with it."

Decision-making task

The task puts the participant in an Antarctic exploration scenario. The participant is tasked with finding a 'lost party' in the Antarctic by driving a snowcat to their location. The task is made up of a series of 24 discrete items (a sample display is shown below), presented successively. Each item presents the participant with a map of the terrain, and symbols indicating the positions of the participant and the lost party. The participant's task is to find the optimal route for reaching the lost party rapidly. The participant is also provided with a target time to attain to rescue the lost party, and an expectation of success that corresponds to the probability of success on the item,

where success is defined as finding a route with an expected travel time that is less than the target time.

Four alternative, color-coded routes are shown on the display. Each route carries risks and potential benefits. By use of the mouse, the participant is able to examine the potential costs and benefits of each route. Costs relate to obstruction of progress, due to terrain and mechanical breakdown. Each cost has a probability and a fixed increase in journey time. For example, there may be a 10% probability of damage to the snowcat due to rough terrain, leading to a time increase of 20 minutes. Conversely, benefits relate to enhanced performance of the snowcat, and decreases in journey time. For example, there may be a 20% probability of finding a short cut to reduce the journey, leading to a time decrease of 10 minutes. A verbal description that may enhance its affective impact is also presented. For example, "5% chance of a hazard causing an 80 minute delay" might be expressed as "there is a small chance that you will fall into a large crevasse, causing a major delay". In each case the numbers are chosen so that the calculation is simple to perform mentally, and results in a whole number (e.g., $5\% \times 80 = 16$).

After assessing the costs and benefits of each route, the participant is asked to choose one of the four routes, using the mouse to register the choice of route. After 30 s have elapsed, the participant is prompted to respond rapidly, and the cost and benefit information is rendered inaccessible. Following the decision, the participant is asked to rate key features of the decision-making problem including its level of risk and uncertainty.

Item types. The pilot study makes use of three different types of item, involving qualitatively different choices between routes (e.g., high or low uncertainty, risky or safe). The factors manipulated may relate to some classic biases demonstrated in the decision-making literature. For example, the 'loss aversion' principle suggests that potential losses may have a disproportionate effect on utility. If possibility of loss outweighs the actual probability of loss, major but improbable hazards may be over-weighted in the decision process. These item types are as follows.

1. *Risky vs. safe options, e.g.,:*

Route 1. Baseline (normal) driving time is 82 minutes. There is a 5% chance of a hazard causing an 80 minute delay. Expected time is 86 min (range: 82-162).

Route 2. Baseline time is 94 minutes. There is a 20% chance of a benefit causing a 40 minute gain. Expected time is 86 min (range: 54-94).

Route 1 is fast, but there is a small risk of a major delay. Route 2 is slower, but there is no risk of delay, and a chance that progress will be faster. Are people excessively influenced by the possibility of a 'disaster'?

2. *High vs. low outcome probabilities, e.g.,*

Route 1. Baseline (normal) driving time is 90 minutes. There is a 67% chance of a benefit causing a 12 minute gain. There is a 5% chance of a hazard causing an 80 minute delay. Expected time is 86 min (range: 78-170).

Route 2. Baseline time is 90 minutes. There is a 10% chance of a benefit causing an 80 minute gain. There is a 67% chance of a hazard causing a 6 minute delay. Expected time is 86 min (range: 10-96).

The previous item set referred to single, low-probability events. This item set, in effect, asks what happens when the person must balance likely small but probable benefits against large but improbable costs (and vice versa)? Here, Route 1 offers the prospect of a probable small gain and an improbable large delay. Route 2, conversely, offers a probable small delay and an unlikely large gain. Will decision be more sensitive to the probabilities than to outcome values (favoring Route 1)? Or will participants be sensitive to the possible disaster scenario (favoring Route 2)?

3. *High and low uncertainty in outcome, e.g.,:*

Route 1. Baseline (normal) driving time is 90 minutes. There is a 33% chance of a benefit causing a 9 minute gain. There is a 50% chance of a hazard causing a 2 minute delay. Expected time is 88 min (range: 81-92).

Route 2. Baseline time is 90 minutes. There is a 50% chance of a benefit causing a 16 minute gain. There is a 67% chance of a hazard causing a 9 minute delay. Expected time is 88 min (range: 74-99).

In both these cases, benefits and costs are both moderately probable. Route 1 offers a narrow range of outcomes; Route 2 offers a wider range. Do subjects prefer the more or less certain option?

Routes are configured to be *optimal* or *suboptimal*, as defined by expected travel time. For example, the choices for an item might be:


- Route 1. 'Risky': expected travel time = 85 mins.
- Route 2. 'Safe': expected travel time = 85 mins
- Route 3. 'Risky': expected travel time = 90 mins.
- Route 4. 'Safe': expected travel time = 90 mins

Routes 1 and 2 are thus optimal, and the participant's ability to choose optimal routes, independent of their qualitative preferences, may then be assessed.

In addition, items can be configured as *high* or *low threat* simply by changing the target time in relation to expected times. In high threat conditions, expected times typically exceed target times. In low threat conditions, the opposite applies. Thus, high and low threat items are designed to have the same structure in terms of probabilities of loss and gain.

The version of the task used items as follows (presented in a pseudo-random sequence), for a total of 24 items in the entire test session:

	Risky vs. safe	High vs. low outcome probabilities (for costs and benefits)	High vs. low uncertainty
Low threat	4 items	4 items	4 items
High threat	4 items	4 items	4 items



A sample task display is shown above.

Subjective measures

Subjective measures are used as a check on the effectiveness of the emotion-induction, and also to gauge the impact of the task in relation to participant workload and stress. The use of subjective measures may be justified on the following grounds:

- Psychological theories of emotion emphasize the centrality of appraisal (Scherer, 2001) and assignment of personal meaning (Lazarus, 1999), suggesting that language-based assessments are appropriate.
- Studies of effects of emotion-inductions on cognitive bias have typically shown comparable effects on subjective scales, which provide the primary manipulation check for such studies (Mätt, Vazquez & Campbell, 1992; Westermann et al., 1996). Emotion-inductions have also been shown to produce concurrent subjective and psychophysiological responses (Kreibig et al., 2007).
- Subjective scales for emotion and related constructs are typically extensively validated in studies that link them to objective factors. For example, subjective scales are sensitive to objective psychophysical parameters of stimuli (Warm, Finomore & Matthews, 2008). They also correlate with psychophysiological indices (e.g., Fairclough & Venables, 2006), and, when tasks are appropriately selected, they predict performance on a subsequent task (Warm et al, 2008).
- Demand characteristics are a concern but there are reasons to believe that they have no impact or minor impact. Evidence adduced by Westermann et al. (1996) suggests that (1) typical participants have little motivation to follow demand characteristics, (2) the demand characteristic hypothesis does not predict the detailed patterns of state change observed experimentally, (3) mood-inductions influence behavioral as well as subjective indices of emotion.

The primary subjective measures used in the pilot study were the adjectival rating scales used by Mayer et al. (1995). The respondent is asked to rate their mood in relation to a series of adjectives covering the three emotions most relevant to the study. These are *happiness* (cheerful, lively, happy, joyful), *anger* (angry, furious, mad, hostile), and *fear* or *anxiety* (scared, fearful, afraid, nervous). Each adjective is rated on a 4-point Likert scale, where 1 = definitely not felt, and 4 = definitely felt. Instructions emphasized the need to “record how you feel right now, at the present time.” Scales for the three emotions were constructed by computing the mean rating for each adjective set.

Participants also completed a short 30-item version of the Dundee Stress State Questionnaire (DSSQ: Matthews et al., 2002), a questionnaire that measures more broadly-defined subjective state factors labeled as task engagement, distress and worry. The DSSQ includes an embedded workload measure, a modification of the NASA Task Load Index (TLX: Hart & Staveland, 1988). It comprises six 0-10 rating scales for key components of workload. In this study, an unweighted sum of ratings was used as an index of overall workload.

The final scale used was a short personality inventory, comprising the trait anxiety and trait anger scales of the State-Trait Personality Inventory (STPI: Spielberger & Reheiser, 2004). Each trait construct is measured by 10 items answered on 1-4 response scales.

Procedure

Upon reporting for the experiment, participants completed an informed consent form. They then completed a sequence of assessments as follows:

1. Short personality questionnaire: STPI

2. Baseline subjective state: Mayer adjectival mood ratings + short Dundee Stress State Questionnaire (pre-task) + NASA-TLX workload
3. Description and practice of task
4. Music – 1 minute [and then continues into next phases]
5. Emotion-induction: 8 vignettes for guided imagery
6. Adjectival ratings of mood
7. First task phase: 12 decision-making items [music continues]
8. Adjectival ratings of mood
9. Emotion-induction 8 vignettes for guided imagery: total of 4 minutes
10. Second task phase: 12 items [music continues]
11. Adjectival ratings of mood + short DSSQ (post-task) + NASA-TLX workload

At the end of the study, participants were debriefed; an optional positive mood induction was offered to participants in negative emotion conditions.