

REPORT DOCUMENTATION PAGE				<i>Form Approved</i> OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching data sources, gathering and maintaining the data needed, and completing and reviewing the 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 Washington Headquarters Service, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188) Washington, DC 20503.					
PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.					
1. REPORT DATE (DD-MM-YYYY) JUNE 2006		2. REPORT TYPE Conference Paper Postprint		3. DATES COVERED (From - To) 12 – 16 Apr 04	
4. TITLE AND SUBTITLE THEORETICAL FOUNDATIONS FOR RATIONAL AGENCY IN THIRD GENERATION WARGAMES				5a. CONTRACT NUMBER In-house	
				5b. GRANT NUMBER 	
				5c. PROGRAM ELEMENT NUMBER 62702F	
6. AUTHOR(S) Paul Bello				5d. PROJECT NUMBER 558S	
				5e. TASK NUMBER 3G	
				5f. WORK UNIT NUMBER WG	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Air Force Research Laboratory/IFSB 525 Brooks Rd Rome NY 13441-4505				8. PERFORMING ORGANIZATION REPORT NUMBER AFRL-IF-RS-TP-2006-2	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Air Force Research Laboratory/IFSB 525 Brooks Rd Rome NY 13441-4505				10. SPONSOR/MONITOR'S ACRONYM(S) 	
				11. SPONSORING/MONITORING AGENCY REPORT NUMBER AFRL-IF-RS-TP-2006-2	
12. DISTRIBUTION AVAILABILITY STATEMENT APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED. PA#04-163					
13. SUPPLEMENTARY NOTES This material is declared a work of the US Government and is not subject to copyright protection in the United States. Paper presented at the Enabling Technologies for Simulation Science VIII Conference, 12-16 Apr 2004.					
14. ABSTRACT Conflict between groups of armed men is as old as recorded history. Effective reasoning and decision-making are fundamental to the successful execution of military operations. These activities are of paramount importance, given the high-stakes nature of conflict; most especially in this modern era of asymmetric threats, and unconventionally armed rogue states. Yet as high as the stakes are, there does not exist a sufficiently formal military theory of reasoning and decision-making that instantiates modern war-fighting doctrine. Large bodies of knowledge of reasoning and decision-making exist, but they are not integrated, and they (to author knowledge) have not been cast effectively into a military context. Herein, I describe a new theory of military rationality which fully captures the reasoning and decision-making processes of military personnel. The goal of the third generation wargaming effort as the Air Force Research Laboratory's Information Directorate is to produce a high-fidelity simulation of conflict environments in order to facilitate a new brand of highly immersive training for warfighters. I shall briefly highlight the philosophical foundations for the construction of such entities, and the formal techniques by which they may be modeled and engineered.					
15. SUBJECT TERMS Third generation wargame; wargame; rational agency					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UL	18. NUMBER OF PAGES 11	19a. NAME OF RESPONSIBLE PERSON David O. Ross
a. REPORT U	b. ABSTRACT U	c. THIS PAGE U			19b. TELEPHONE NUMBER (Include area code)

Theoretical Foundations for Rational Agency in Third Generation Wargames

Paul Bello
Air Force Research Laboratory
Information Directorate
AFRL/IFSB - C4ISR Modelling and Simulation Branch
Rome NY 13441 USA
The Rensselaer AI and Reasoning Laboratory
Department of Cognitive Science
Rensselaer Polytechnic Institute (RPI)
Troy NY 12180 USA
Paul.Bello@rl.af.mil · bellop@rpi.edu

ABSTRACT

Conflict between groups of armed men is as old as recorded history. Effective reasoning and decision-making are fundamental to the successful execution of military operations. These activities are of paramount importance, given the high-stakes nature of conflict; most especially in this modern era of asymmetric threats, and unconventionally armed rogue states. Yet as high as the stakes are, there does not exist a sufficiently formal military theory of reasoning and decision-making that instantiates modern war-fighting doctrine. Large bodies of knowledge on reasoning and decision-making exist, but they are not integrated, and they (to the author's knowledge) have not been cast effectively into a military context. Herein, I describe a new theory of *military rationality* which fully captures the reasoning and decision-making processes of *homo militius*, military man. The goal of the third generation wargaming effort at the Air Force Research Laboratory's Information Directorate is to produce a high-fidelity simulation of conflict environments in order to facilitate a new brand of highly immersive training for our war-fighters and supporting personnel. This environment will be populated by a new breed of intelligent agents that we affectionately call ASC-ME's (Advanced Synthetic Characters for Military Environments). I shall briefly highlight the philosophical foundations for the construction of such entities, and the formal techniques by which they may be modelled and engineered.

1. INTRODUCTION

The theory of decisions has played a vital role in economic policy, the design of automated tools for medical decision-support, and in the practice of law. Yet there remain unexplored issues in the application of the decision sciences to military conflicts and organizations. In this modern age of terrorist threats, weapons of unimaginable destructive power, and information overload, there does not exist a computational model of decision making which will allow our leadership to make strategic decisions that are centered on the operational concepts that are defined by the Joint Chiefs of Staff in 1. We struggle as a community to understand, and to be able to model the motivations behind terrorist networks seeking the subversion and dismantling of democratic and free societies. In preparation to fight these enemies of freedom, we must develop more flexible notions of rationality, notions that may be inconsistent with, yet coexist with our own ideologies and traditions. Our only hope in waging war against an army of "martyrs" is to understand the motives, means, and methods that the enemy possesses so that we may preempt, deny, and ultimately win the war against terrorism, and its national sponsorship by rogue states.

The Third Generation Wargaming effort is meant to tie together some very important, but disparate concepts from the domain of military science by integrating them via computational models. These models will be the algorithmic basis for decision-support tools that are usable in the predictive modelling and simulation of conflict on a number of different levels. Most importantly, in support of the potential user of this system (military

leadership), this research effort seeks to define the concept of *homo militius*, the military man, and to give an axiomatic treatment of the principles that guide his decision-making process.

The U.S. military is an ever-evolving entity. In order to persist in its position of preeminence among the armed fighting forces of the world, it needs to continually predict, adapt to, and exploit the weaknesses of emerging threats. Information dominance plays a crucial role in the preemption of hostilities against the United States, its friends, and allies. Requisite to this dominance is decision-superiority, or the ability of leadership to make high-quality decisions in the presence of uncertainty, and under severe stress. Decision-superiority is defined in¹ as:

Information superiority provides the joint force a competitive advantage only when it is effectively translated into superior knowledge and decisions. The joint force must be able to take advantage of superior information converted to superior knowledge to achieve "decision superiority" - better decisions arrived at and implemented faster than an opponent can react, or in a noncombat situation, at a tempo that allows the force to shape the situation or react to changes and accomplish its mission. Decision superiority does not automatically result from information superiority. Organizational and doctrinal adaptation, relevant training and experience, and the proper command and control mechanisms and tools are equally necessary.

With these considerations in mind, we seek to revolutionize the training process for the men and women in our defense community by providing the merger of bleeding-edge applications of AI, and the doctrinally sound principles that have kept our nation safe for so many years. This paper is organized to address three issues, corresponding to the three requirements identified as being crucial to the construction of a third-generation wargame:

1. Rigorous theoretical foundations for describing the agents which populate the gaming environment. The theoretical framework in question will embody a system-on-system approach to the conduct of war, under the assumption that the enemy (at least in principle) will not be fighting a conventional conflict against the United States and coalition allies; trying instead to defeat us through breaking the spirit of our societal institutions and the organizations that protect them.
2. The inclusion of empirically justified models of human behavior into the formalisms developed for the wargaming agents.
3. Building a referentially opaque system, suitable for extension to a full-scale, post-game explanatory analysis (in natural language)

2. THIRD-GENERATION WARGAMING AGENTS: THEORY AND DOCTRINE

While names such as Von Neumann, Morgenstern, and Simon have made indelible marks on the formulation of the theory of decisions, one cannot help but think that their contributions lack a certain degree of realism. Military and political decision-making require a less idealized notion of "rationality" than these popular schools of thought can provide. One of my favorite quotes from Herbert Simon goes along the lines of "... there is no place for ethics in a science of decisions." Military man, whether the soldier on the ground, or the commander-in-chief, must be extremely concerned with ethics, and highly sensitive to the risks of each decision he makes. A bad decision on the part of a commander could not only impact the outcome of war, but could jeopardize the lives of the troops under his watch. We must not lose sight of the fact that implicit relationships exist between members of every type of organization, whether you classify the organization as a society, a nation, a military, a political party, a business, or even a family. A decision-maker wearing the shoes of a family man is constrained by the ethical standards associated with the context of being a family man. Certain actions are permissible, while others are forbidden. The life of military man is fraught with context-sensitive decision-making of this sort. Without it, chain-of-command would cease to exist, and the uniform code of military justice would have no usefulness. Military man must operate within well defined boundaries, sensitive to the wishes of their superiors,

their nation, and to their brothers-in-arms. Such complex mental accounting cannot be simply simulated. A new theory of decision is required, one that is founded on semantic mechanisms for dealing with such difficult concepts as knowledge, belief, obligation, permission, justification and truth.

2.1. Effects-Based Operations

The advent of information warfare, and other forms of low-intensity combat has heralded a new way of thinking among the leadership of the U.S. military. Engagement with the enemy can no longer be confined to force-on-force conventional engagements. The enemy of the future (and in some respects, the present) will rely on information operations, asymmetric tactics, and symbolic attacks to achieve their political agendas. Sadly, all of these things have begun to come to pass, from the ariel attack on the World Trade Center, to the brutal intimidation tactics employed by armed forces such as Iraq's Elite Republican Guard. Media is being used, now more than ever, as a means of delivering psychological attacks, and to boost morale. A new concept of operations is needed to address these alternative forms of conflict in a unified manner, and they are the *raison d'être* for the effects-based operations (EBO) philosophy. EBO is a process for obtaining a desired strategic outcome or *effect* on the enemy through the synergistic, multiplicative, and cumulative application of the full range of military and nonmilitary capabilities at the tactical, operational, and strategic levels. EBO consists of a set of processes, supported by tools and accomplished by people in organizational settings, that focuses on planning, executing, and assessing military activities for the effects produced, rather than merely attacking targets or simply dealing with objectives. EBO complements, rather than replaces, target-based or objectives-based approaches (such as strategy to tasks) and is very amenable to mission-type orders and strategy options that do not emphasize attrition-based approaches. Since EBO is not specifically battlefield oriented, we can begin to model any interconnected set of entities as a system, so that we may apply the principle of effects-based operations to its analysis.

2.1.1. Center of Gravity Analysis

Centers of Gravity (COG's) are considered the hub of the adversary's power. The will of the adversary to keep fighting often is inextricably linked to these important targets. Barlow² has identified seven National Elements of Value: Leadership, Industry, Armed Forces, Population, Transportation, Communication, and Alliances. Noting that some of these COG's are not feasible to attack directly (such as population and alliances), it is often possible to attack them through the application of nonlethal methods (political and psychological coercion). Barlow contends that attacking the right combination of these elements will induce a strategic paralysis that may win the day without excessive devotion of conventional forces to the effort.

2.1.2. Target System Analysis

John Warden has proposed a five-ring model for the analysis of target systems.³ Warden's model is composed of the following five "rings":

1. **Leadership:** The central command unit for the strategic entity in question. The leadership is the "brains", so to speak, of the system under consideration. This entity is usually the core decision-making entity for the system as a whole. Entities can survive without leadership, but the system cannot perform any strategic analysis.
2. **Organic Essentials:** These can be classified as the mechanisms whereby the leadership performs its function. Usually, these are material resources such as food, energy, money, power, etc. While the leadership performs the integration and direction of these resources towards the accomplishment of its objectives, both can survive independent of one another, but system effectiveness relies on the synergistic interplay of these elements.
3. **Infrastructure:** Infrastructure is how resources are disseminated for use by each part of the system. Abstractly, these are the connections which keep information, utilities, money, weapons, and people moving through the system.

4. **Population:** Population can be loosely thought of as the workhorse of the system. They are the key enablers for transport of resources, execution of decisions, and almost every other important system function. However, people are a redundant system for the most part. The elimination of a large majority of population may not be enough to paralyze the system as a whole.
5. **Fighting Mechanism:** The systems which we would like to consider have all got protective mechanisms of one sort or another. Whether it's debugging code for a complex computer program, or armed forces defending a nation-state, the system defense mechanisms comprise the fifth, and final part of our model.

Dimensions	Body	State	Cartel	Electric Grid
Leadership	Brain	Government, Security, Energy	Leader, Communication, Security	Central control
Organic Essentials	Food and oxygen	Energy (Electricity, oil, food, money)	Coca source plus conversion	Input (heat, hydro) and output (electricity)
Infrastructure	Vessels, bones, muscles	Roads, airfields, factories	Roads, airways, seaways	Transmission lines
Population	Cells	People	Growers, distributors, processors	Workers
Fighting Mechanism	Leukocytes	Military, police, fireman	Street soldiers	Repairmen

Table 1. An Example of the Five-Ring Model

As the reader can observe from the construction of Table 1, almost any set of entities with a non-trivial set of interactions can be viewed as a system. In military affairs, our objective is to make war on the *system*, and to understand how the dependencies in the system may be exploited in order to induce certain types of behavior in other parts of the system (for example, we'd like to force the capitulation of the leadership element, in the best case scenario). In any case, our view of warfare, and our attempts to model both the battlespace and the entities which exist within it must be a systemic view; where interaction among entities and environment plays the central role in defining the conduct of operations. A theory of decisions to capture the causal relationships that are at the heart of effects-based operations must put its emphasis on outcomes, and how to achieve them, rather than focusing solely on the analysis of choice. In the latter part of this paper, we shall see how the confluence of reasoning, decision-making, and organizational theory can synergistically interact to provide a substrate for effects-based operations; and thus a solid platform on which to construct a third-generation wargame.

We will first begin our discussion with some background on rational agency, and some ideas as to why the rationality debate rages at the heart of our problem.

2.2. Rational Agency for Individuals and Organizations

A major objective of our work is the comparison and contrasting of the most popular theories of rationality that have been postulated by the academic community over the past sixty or so years. In doing so, we seek to understand the benefits and shortcomings of applying these theories in a meaningful way to military decision problems, where the stakes are unusually high, and uncertainty reigns supreme.

2.3. The Rationality Debate and Four Requirements

The rationality debate is centered around four requirements which remain the nexus of an extremely volatile academic debate. *Homo Economicus*, or the Economical Man who embodies these four requirements, is an oft used (and abused) notion that at one time was the cornerstone of economic theory. Informally, these four requirements defining economic man amount to the following:

1. **Perfect Knowledge:** The decision-maker has a clear picture of the choice problem he faces. He is fully aware of all the alternatives from which he has available to him. He is also fully cognizant of the outcomes associated with each alternative.
2. **Unambiguous Preference:** the decision-maker has a full-scale preference ordering over his entire space of alternatives. There exists some function over the set of alternatives which takes two preferences, and returns at the least, a singular weakly preferred choice.
3. **Unlimited Calculative Capacity:** The decision-maker is capable of any and all calculations required to discover his optimal course of action.
4. **Indifference:** The decision-maker is required to remain ambivalent when faced with a set of logically equivalent (from the standpoint of utility) alternatives.

It has been widely accepted that *Homo Economicus* is an extraordinary idealization of the decision-making process. Certainly, these assumptions presuppose an almost God-like omniscience, and do not seem to manifest themselves as behavioral traits attributed to any human being in a real-world setting.

The Economical Man was born out of the theory of utility developed by Von Neumann and Morgenstern in their classic treatise on game theory.⁴ Decisions are made after a process of mentation concerning the feasibility and desirability of the set of alternatives and associated outcomes. *Homo Economicus* is a maximizer, who seeks to maximize his utility by choosing the alternative from the set of feasible choices which will result in the most desirable outcome *on the average*.

2.3.1. Utility Theory and Individual Decision-Making

Let us briefly examine the most widely accepted formalization of individual decision-making. Choices in this case are defined as *prospects*, or lotteries over a set of possible outcomes.

Let $C_A = \{c_0, c_1, \dots, c_n\}$ represent the set of *choices* available to agent *A*. Each choice $c_i \in C$ may have associated with it a set of *outcomes*, $O_i = \{o_0, o_1, \dots, o_m\}$, representing the potential uncertainty involved with choosing a particular c_i . Each outcome $o_{ij} \in O$ will have a certain probability of occurring, as well as having some numerical notion of "desirability", representing how "good" the outcome would be if it occurred. The rational decision-maker will perform the following calculation, in order to determine a ranking of choices via expected utility (EU), from highest to lowest, in an unambiguous way:

$$EU(c_i) = \sum_{c_i \in C} \sum_{o_{ij} \in O} p(o_{ij}) \cdot g(o_{ij}) \quad (1)$$

The Von Neumann and Morgenstern utility rests upon four assumptions, summarized nicely in 5. The existence of a utility function, or a mapping from relative desirability of a prospect to $\mathbb{R}$ is rigorously demonstrated in 4. Rational choice corresponds to the maximization of expected utility, given by the following:

$$MEU = \arg \max_{c_i \in C} \left(\sum_{c_i \in C} \sum_{o_{ij} \in O} p(o_{ij}) \cdot g(o_{ij}) \right) \quad (2)$$

2.3.2. Administrative Man and Simon's Theory of Organization

In his nobel-prize winning work, *Administrative Behavior*,⁶ Herbert Simon defined the *Administrative Man*, as a close cousin of economic man, and by doing so, reinvented the theory of organizations. Administrative man, much like his cousin, is a maximizer, but with a few caveats: he only has a subset of the total decision-problem for an organization to deal with (unless he is the CEO, or at the top of the organizational chain), and he may only be able to do a limited amount of calculation, based on the resources he has available to him. In this sense, Simon defined the abstract notion of management as being a fundamental part of organizational decision processes.

Most notably, Simon is known for raising the following three objections to the notion of economic man:




Figure 1. Traditional Levelled Rationality

1. Rationality requires complete knowledge and anticipation of the consequences that will follow on each choice. In fact, knowledge of consequences is always fragmentary
2. Since the consequences lie in the future, imagination must supply the lack of experienced feeling in attaching value to them. But values can be only imperfectly anticipated.
3. Rationality requires a choice among all possible alternative behaviors. In actual behavior, only a very few of all these possible alternatives ever come to mind.

Given these objections, Simon provides somewhat of a *lingua franca* for talking about administrative man, and his decision process:

- Intended: "Human behavior in organizations is, if not wholly rational, at least in good part intendedly so. human behavior is intendedly rational, but only limitedly so.." (Therefore there is room for a genuine theory of organization and administration)
- Subjective: "A decision is 'subjectively' rational if it maximizes attainment relative to actual knowledge of the subject."
- Conscious: A decision "is 'consciously' rational to the degree that the adjustment of means to ends is a conscious process
- Deliberate: A decision is "deliberately' rational to the degree that the adjustment of means to ends has been deliberately brought about (by the individual or by the organization)"
- Organizational (Personal): A decision is organizationally (personally) rational if it is oriented to the organization's (individual's) goals.
- Bounded: Simon's "bounded rationality" refers to the simplified and constrained model of intended rationality which decision makers construct for themselves to deal with complex decision situations. The important constraints are informational and computational limits on rationality by human beings.
- Procedural: "Behavior is procedurally rational when it is the outcome of appropriate deliberations. Its procedural rationality depends on the process that generated it. When psychologists use the term 'rational', it is usually procedural rationality they have in mind."

2.3.3. Organizational Linearity

Most importantly, Simon has provided good reason to approach decision-making, even of the complex variety that we are discussing, in a formal, computable way. Administrative rationality is an instantiation of *Means-Ends Analysis* (MEA), made popular by Simon and Newell's GPS system.⁷ MEA is a goal-directed search procedure that can be considered seminal in the development of classical artificial intelligence techniques such as automated planning, and backward-chaining rule-based systems. As the reader may notice in the figure above, Simon's notion of organization begins with an organizational goal at the top level, which (potentially) decomposes into a number of different subgoals (at the second level), and so on. Identifying desired ends, and the means Organizational design and streamlining can be thought of in this sense as assigning the right people to the right subgoals, since invariably, each subgoal has some number of decision problems associated with it. In essence, organizational planning is much the same as effects-based planning: an end-state is chosen, and then the decision problem becomes that of choosing mechanisms by which to apply resources to achieve the end-state under a specified set of constraints.

The goal, of course, is to maximize efficiency and effectiveness of organizational decision-making through the proper decomposition and assignment of tasks to lower-level decision-makers. Ideally, this is a perfect-information scenario, where each decision-maker is responsible for a niche decision-problem, and the results of his deliberation are trusted absolutely. These assumptions define a linearized model of decision-making within an organization, which to the author, doesn't seem nearly realistic enough to accurately account for complex phenomena generated by missing information, multiple-team-to-single-task assignments, and the non-homogeneity of the heuristics for option evaluation by each decision-maker.

2.4. Military Rationality as an Organizational Framework

Military Rationality is a multi-resolution framework describing the decision processes of military man. Within this framework, we can construe the military decision-making process (which is by nature, a group decision-making activity) as an individual decision-making process, consisting of dynamics between each level of organization. Intra-level decision processes may drive the outcome of inter-level processes, making the decision problem nonlinear, usually complex from the standpoint of modelling, and generally a good fit for the systems-of-systems architecture that we have identified as being desirable.




Figure 2. Ratifiable Decision Processes: Military Rationality

Decision-making in general can be broadly categorized as either analytical, or intuitive. Analytical decision-making usually involves the quantitative analysis of a number of strategies in order to make preference among them, and eventually to act on one (or more) of them. Intuitive decision-making is quite different in nature. This type of decision process involves the judgment, experience, and intuition of the decision-maker in order to recognize and act on emerging patterns of evidence. Both methodologies have their respective place in decision-making, and the two rarely appear as disjoint algorithms. Analytical decision-making supports the intuitive process through the provision of subjective probability distributions concerning outcome spaces, whereas the intuitive process is supported implicitly through the internalization of complex preference relations which are rooted in experiential knowledge. While these two paradigms may have been studied separately in the past, we shall consider them in unison under the assumption that decision processes, in general, are complex algorithms. These algorithms draw upon a large number of semantically diverse information structures, consisting of both explicit accessible knowledge (such as the experiential/episodic memory required to perform intuitive decision-making), and implicit information which may be responsible for ad-hoc generation/evaluation of utility functions and preference orderings. Bringing these disparate methodologies together may only require a singular concept: the ratifiable decision.

2.4.1. Ratifiable vs. Non-ratifiable Decision Processes

Before we launch into a discussion of ratifiability, it is useful to discuss one of the most commonly cited frameworks for decision-making: the OODA loop.⁸ OODA stands for *Observe, Orient, Decide, Act*. The OODA process is cyclical, where the decision-maker is constantly observing new information, re-orienting himself to his environment as it changes (making comparisons between the current state, and the end state he has envisioned), deciding between alternative strategies to achieve the end state, and acting on these strategic considerations. In many ways, it parallels Simon's theory of organization found in 6. OODA is driven by means-ends, goal-directed behavior. OODA has become an increasingly popular framework for comparative analysis of decision processes in the military, however, as many of the existing frameworks, there doesn't exist a sufficiently formal model to describe it, and implement it in intelligent agents.

Ratifiable decision-making forms a large part of the technical core of material at the heart of Military Rationality. For the sake of brevity, I will leave the formal description of ratifiable decisions in organizations to another document. The basic idea behind ratifiable decisions is in stark contrast to the normative framework of Bayesian decision theory. Normative theory tells us that decision is inextricably linked with action. Action must be taken in order for the deciding party to adapt to its environment. This is in fact the mantra of the *Rational Analysis* approach posited by Anderson.⁹ This type of approach relies on the statistical structure of the environment to provide description of behavior through the vehicle of adaptation. Action-oriented theories such as these, while within the general spectrum of cognitive theories, do not address the issue of representation and cognitive capacity (the innate ability to represent and compute). Ratifiable decisions, on the other hand, depend on the decision-maker's ability to "see into his future", so to speak, and to make decisions based on a (potentially hypothetical) desirable outcome. In short, the ratifiable decision-maker prefers to make decisions based on the person he'd like to be, when the time for acting is necessary, rather than accept the less desirable, uncertain reality associated with the present situation. Ratifiability can be used to explain a wide variety of phenomena, including the famous "retirement planning"¹⁰ problem in decision theory by providing a formal mechanism for establishing higher-order preference relations (preferences regarding existing preferences), which in essence, describes the conflict that the ordinary man faces in many decision-making scenarios. For a clear exposition of ratifiability and higher-order preference, see 11. It is the author's hypothesis that the development of higher-order preference relations is the result of some compelling reasoning processes involving doxastic, epistemic, and deontic considerations, motivating the integration of formal reasoning with decision-making.

2.5. Military Rationality as an Empirical Framework

While Military Rationality can be seen as a framework for organizational behavior, it also becomes a rich testing bed for the conduct of psychological experimentation. Military Rationality is, at its core, a theory of higher-order cognition, and as such, must be empirically motivated. Psychological models of representation, inference, and decision must be integrated to produce artificial decision-makers who behave as we expect real military men to behave.

2.5.1. Mental Models, Mental Logic, and Conditioned Utility

The mechanism for decision in the Military Rationality framework is a quasi-logic, closely related to Jeffrey's Logic of Decision.¹¹ This being the case, we commit to a certain form of representation. Semantically speaking, a number of "possible worlds" describe the space of outcomes, which when taken in proper subsets, implicitly define the notion of a choice, since all choices lead to a number of outcomes. However, the worlds themselves contain sentences built from the language of standard propositional logic. When choices are perceived in the environment by an agent, they are translated into a propositional form, suitable for representation as *mental models*.¹² Certainly, it seems silly to ratify a "choice", or an "outcome", but the ratification of *propositional attitudes* seems to provide a much more intuitive description of the revision process. Mental model theory provides a set of semantic principles for representation (such as the so-called *principle of truth*), and inference. For increasingly complex logical constructions (such as quantified compound statements), the use of *mental logic theory*¹³ can expedite the representation and reasoning processes. A thorough discussion of mental model theory and mental logic theory is beyond the scope of this paper, but the debate between the two paradigms still rages.¹⁴ These empirical theories come particularly in handy when computing the utility for a particular outcome. Since we are able to rank the relative difficulty of representing and reasoning about particular mental models (or schemas in the case of mental logic), we can assume that the agent will do his best to minimize cognitive effort by weighting low-effort inferences and simpler representations as being more desirable in his utility calculations.

2.6. Military Rationality as a Computational Framework

In concordance with the goal of the Third Generation Wargaming initiative, the theory of Military Rationality must be able to be instantiated in the form of intelligent agents, existing in a simulation environment. The enabling technologies to make militarily rational agents into a reality are currently under development, including tools for reasoning/decision-making across multiple modes (proof-based/model-based); and multiple levels of expressivity (classical first-order logic, various non-alethic modal systems including epistemic, doxastic, and deontic logics), and architecturally motivated by the empirical data found in the psychology of reasoning and the psychology of decision-making.¹⁵ Along with these fundamental advances in reasoning, algorithms for ratification and higher-order preference formation are under development. All of these tools are designed to produce agents whose process of mentation is completely inspectable, debuggable, and understandable from the standpoint of a human user of the system.

3. SUMMARY AND ACKNOWLEDGMENTS

The goal of the Third Generation Wargaming initiative at AFRL/IF is to research the technologies which will enable the development of future wargaming and simulation capabilities. Such a capability must be driven by a strong commitment to an underlying philosophy of rational agency. This paper has attempted to broadly define that philosophy in the light of similar normative accounts. The author would like to acknowledge contributions from Yingrui Yang, Selmer Bringsjord, Ron Sun and Kostas Arkoudas of the Rensselaer Cognitive Science Department; as well as continued support within the Air Force Research Laboratory's Information Directorate.

REFERENCES

1. H. Shelton. General, Chairman - Joint Chiefs of Staff, "Joint vision 2020." The Pentagon, 2000.
2. J. Barlow, "Strategic paralysis: An airpower theory for the present," (1), 1994.
3. J. Warden, "The enemy as a system," *Aerospace Power Journal* (1), pp. 40-55, 1995.
4. J. Von Neumann and O. Morgenstern, *Theory of Games and Economic Behavior*, John Wiley and Sons, New York, NY, 1944.
5. H. Chernoff and E. Moses, *Elementary Decision Theory*, Wiley, New York, NY, 1959.
6. H. Simon, *Administrative Behavior, Fourth Ed.*, The Free Press, 1997.
7. H. Simon and A. Newell, "GPS, a program that simulates human thought," in *Lernende Automaten*, H. Billing, ed., pp. 109-124, R. Oldenbourg, Munich, Germany, 1961.
8. J. Boyd, "A discourse on winning and losing." An unpublished set of slides available at Air University Library, Maxwell AFB, Alabama, 1987.

9. J. Anderson, *The Adaptive Character of Thought*, Lawrence Earlbaum Associates, Hillsdale, NJ, 1990.
10. A. Gupta and Y. Yang, "Ordinary rationality and turning points in retirement planning," in *Proceedings of the Joint International Conference on Cognitive Science*, 2003.
11. R. Jeffrey, *The Logic of Decision*, University of Chicago Press, Chicago, IL, 1965.
12. P. N. Johnson-Laird, *Mental Models*, Harvard University Press, Cambridge, MA, 1983.
13. M. D. S. Braine, "On the relation between the natural logic of reasoning and standard logic," *Psychological Review* **85**, pp. 1-21, 1990.
14. D. O'Brien, M. Braine, and Y. Yang, "Propositional reasoning by mental models? simple to refute in principle and in practice," *Psychological Review* **101**, pp. 711-724, 1994.
15. D. Kahneman, P. Slovic, and A. Tversky, *Judgement Under Uncertainty: Heuristics and Biases*, Cambridge University Press, Cambridge, MA, 1982.