Auton Agent Multi-Agent Syst (2010) 20:3–13
DOI 10.1007/s10458-009-9096-7

Modeling self-deception within a decision-theoretic framework

Jonathan Y. Ito · David V. Pynadath · Stacy C. Marsella

Published online: 10 May 2009

Abstract Computational modeling of human belief maintenance and decision-making processes has become increasingly important for a wide range of applications. In this paper, we present a framework for modeling the human capacity for self-deception from a decision-theoretic perspective in which we describe an integrated process of wishful thinking which includes the determination of a desired belief state, the biasing of internal beliefs towards or away from this desired belief state, and the final decision-making process. Finally, we show that in certain situations self-deception can be beneficial.

Keywords Self-deception · Wishful thinking · Decision theory · Belief maintenance

1 Introduction

A mother has been shown seemingly incontrovertible evidence of her son's guilt. Although the information is provided by reliable sources, the mother continues to proclaim her son's innocence. This illustrates an important characteristic of human belief maintenance: that our beliefs are not formed merely by the evidence at hand. Rather, desires and intentions interfere with the processes that access, form and maintain beliefs and thereby bias our reasoning.

Research on human behavior has identified a range of rational as well as seemingly irrational tendencies in how people manage their beliefs [10]. Research in human emotion has detailed a range of coping strategies such as denial and wishful thinking whereby people will

J. Y. Ito (✉) · D. V. Pynadath · S. C. Marsella
Institute for Creative Technologies, University of Southern California,
13274 Fiji Way, Marina del Rey, CA 90292, USA
e-mail: ito@ict.usc.edu

D. V. Pynadath
e-mail: pynadath@ict.usc.edu

S. C. Marsella
e-mail: marsella@ict.usc.edu

be biased to reject stressful beliefs and hold on to comforting ones [11]. Research on cognitive dissonance [8] has demonstrated that people often seek to achieve consistency between their beliefs and behavior. Specifically, cognitive dissonance research has especially focused on how we alter beliefs in order to resolve inconsistencies between a desired positive self-image and our behavior [1], much like Aesop's fable of the fox and the grapes in which after repeatedly failing to reach a bunch of grapes the fox gives up and concludes that the grapes did not look so delicious after all. Similarly, research has also shown a tendency for what is called motivated inference, the tendency to draw inferences and therefore beliefs, based on consistency with one's motivations as opposed to just the facts. Research on how people influence each other has also identified a range of influence tactics that are not simply based on providing factual evidence. However, these are not unconstrained; people do not, cannot, simply believe whatever they choose.

Computational modeling of these human belief maintenance mechanisms has become important for a wide range of applications. Work on virtual humans and Embodied Conversational Agents increasingly has relied on research in modeling human emotions and coping strategies to create more life-like agents [9]. Work in agent-based modeling of social interaction has investigated how persuasion and influence tactics [4] can be computationally modeled [15] for a variety of applications such as health interventions designed to alter user behavior [3].

In this work, we approach the issue of human belief maintenance from the perspective of decision-theoretic reasoning of agents in a multi-agent setting. Specifically, we argue that a range of self-deceptive phenomena can be cast into a singular framework based upon Subjective Expected Utility (SEU) Theory. To cast the seemingly irrational process of wishful thinking and self-deception into a decision-theoretic framework may in itself seem irrational. However, we argue that seemingly irrational behavior such as wishful thinking, motivated inference, and self-deception can be grounded and integrated within an agent's expected utility calculations in a principled fashion such that the desired beliefs of the agent, which are crucial to the self-deceptive process, can be derived form the preferences of the agent itself. Furthermore, we will also show that not only is our self-deceptive framework a principled descriptive theory for human decision making but it also serves as a normative theory under certain circumstances.

2 Self-deception framework

Psychological literature on self-deception commonly refers to the *act* of self-deception as the internal biasing processes involved in adopting a desired belief in the face of possibly contradictory evidence [5,17]. By focusing on these biasing processes, oftentimes the definition and specification of the desired belief state itself remains very abstract. However, by employing Subjective Expected Utility Theory, we are provided a means by which to not only bias beliefs towards a desired belief state and thus influence the subsequent decision-making process but also to designate the desired belief state itself given the decision-maker's own preferences.

Our model of self-deception is a psychologically inspired model of decision making under risk and uncertainty. We loosely define risk as the environment in which objective probabilistic information regarding the occurrence of the possible states in the world is known. Alternatively, an environment of uncertainty is one in which objective probabilities are wholly or partially unknown requiring the formulation of subjective probability estimates reflecting the beliefs of the decision maker.[1] We define a general decision problem as the selection of an action from among a set, A, of actions in an environment consisting of a set, S, of possible states in which exactly one will prevail. A specific outcome may be represented by the pair (a, s), where $a \in A$ and $s \in S$ and which is obtained by performing action a when the prevailing state of nature is s.

2.1 Subjective expected utility maximization

According to SEU Theory, we may define a utility function $\mu(a, s)$ over all possible outcomes which is representative of the decision maker's preferences over outcomes. Furthermore, SEU Theory states that by selecting action a_{u} among the available set of actions such that expected utility is maximized as in (1) in which $pr(s)$ is the subjective probability estimate that state s prevails, it can be shown that the decision maker is acting in accord with their preferences.

$$a_{\mathrm{u}} = \operatorname*{argmax}_{a \in A} SEU\,(a) = \operatorname*{argmax}_{a \in A} \sum_{s \in S} pr\,(s) \cdot \mu\,(a, s) \tag{1}$$

The standard SEU paradigm presents a decision process as depicted in Fig. 1a in which the beliefs of an agent are represented by its probability function pr over the possible states of nature and its preferences over outcomes is represented by the utility function $\mu(a, s)$.

2.2 Self-deceptive decision making

Our generalized formulation for the act of self-deceptive decision making within a decision-theoretic framework is the biasing of the rational belief state towards some alternative belief state. The formal definition of the process is the choice of some action a_{sd} such that the self-deceptive expected utility is maximized as seen in (2) in which α is a constant controlling the magnitude of self-deception evinced by a decision maker and $pr_{\mathrm{sd}}(s)$ is a probability mass function describing an alternative probability distribution over the set of possible states. Notice that when $\alpha = 0$ a decision maker reverts to maximizing expected utility and when $\alpha = 1$ as in (3) a decision maker is fully self-deceptive and bases his actions entirely on the maximization of his self-deceptive expected utility.

$$a_{\mathrm{sd}} = \operatorname*{argmax}_{a \in A} \sum_{s \in S} ((1-\alpha)pr(s) + \alpha pr_{\mathrm{sd}}(s)) \cdot \mu(a, s) \tag{2}$$

$$a_{sd} = \begin{cases} \operatorname{argmax}_{a \in A} \sum_{s \in S} pr(s) \cdot \mu(a, s) & \text{if } \alpha = 0 \\ \operatorname{argmax}_{a \in A} \sum_{s \in S} pr_{\mathrm{sd}}(s) \cdot \mu(a, s) & \text{if } \alpha = 1 \end{cases} \tag{3}$$

In contrast to the traditional decision-making process of SEU maximization, the self-deceptive process is depicted by Fig. 1b.

2.3 Wishful thinking

Wishful thinking is a special case of self-deception, in which optimistic (pessimistic) decision makers bias themselves toward believing positive (negative) outcomes to be more likely than reality would suggest. EU-Theory provides the basis for our formulation of both optimistic and pessimistic self-deceptive wishful-thinking. With it, we not only are able to define the

[1] It is important to note that subjective probability estimates are not the only way to deal with an uncertain decision environment. For a more complete definition and discussion of certainty, risk, uncertainty, and ignorance with respect to decision-theory we refer the reader to the work of Luce and Raiffa [13].

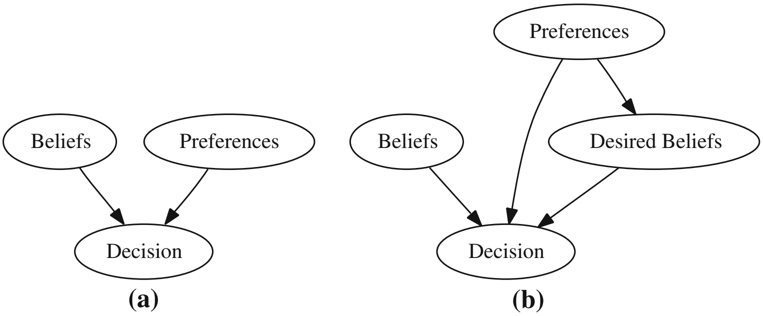

Fig. 1 SEU maximization and self-deceptive decision processes. (**a**) SEU maximization process, (**b**) SD decision process

final decision-making process, but also derive the desired belief state of an agent engaged in the self-deceptive process of wishful thinking.

Under our wishful-thinking formulation of self-deception, we specify the alternative belief state to be certainty that a particularly favorable state $s_\mathrm{k} \in S$ will occur, or more precisely $pr_\mathrm{w}(s_\mathrm{k}) = 1.0$.

Because EU Theory provides a robust axiomatic treatment of the relationship between utility and the preference ordering over outcomes for a decision maker, we can utilize the decision-maker's utility function, μ, in the determination of the most preferred state. Equation 4 presents the formal determination of s_k as the state which would maximize utility if it were to occur with certainty (i.e., $pr(s_\mathrm{k}) = 1$) and if the decision maker has full knowledge of its certainty. Note that this definition is not concerned with the probability estimates of the decision maker. This omission is intentional and reflects the desire for the best possible outcome irrespective of reality. The alternative probability mass function pr_w, is then defined as (5).

$$s_k = \underset{s \in S}{\operatorname{argmax}} \left(\max_{a \in A} \mu(a, s) \right) \tag{4}$$

$$pr_\mathrm{w}(s) = \begin{cases} 1 & \text{if} \quad s = s_\mathrm{k} \\ 0 & \text{if} \quad s \neq s_\mathrm{k} \end{cases} \tag{5}$$

We now illustrate the desired belief formulation process with a simple example:

Example 1 Let us revisit the example of a mother proclaiming her son's innocence despite iron-clad evidence to the contrary. We can represent the mother's dilemma as a simple decision problem consisting of 2 states as shown in Table 1. Furthermore we make the assumption that the best possible outcome, with respect to the mother's preferences, is a steadfast belief of her son's innocence coinciding with actual innocence. We also assume that the worst possible outcome is a belief in her son's guilt when in actuality he is innocent. After making these assumptions, only 2 other possible preference orderings remain: $a \succ b \succeq c \succ d$ or $a \succ c \succeq b \succ d$. The former preference ordering is one in which the mother will always choose to proclaim her son's innocence regardless of the evidence presented. And since this behavior is coincidental with the mother's desired belief that her son is innocent, let us instead consider the preference ordering of $a \succ c \succeq b \succ d$. To illustrate the process of desired belief formulation we assign numerical utilities to the various outcomes in accordance with our preference ordering as seen in Table 2. Let us now consider the following two candidate belief distributions: b_0, in which the son is certainly innocent, and b_1 where the

Table 1 Mother's dilemma

	Son innocent	Son guilty
Proclaim innocence	a	b
Proclaim guilt	d	c

Table 2 Sample utility table for mother

	Son innocent	Son guilty
Proclaim innocence	3	1
Proclaim guilt	0	2

Table 3 Candidate belief table

	$pr_{\mathrm{sd}}(\text{innocent})$	$pr_{\mathrm{sd}}(\text{guilty})$	$max_{a \in A} \sum_{s \in S} pr_{\mathrm{sd}}(s) \cdot \mu(a, s)$
b_0	1.0	0.0	3.0 when proclaiming innocence
b_1	0.0	1.0	2.0 when proclaiming guilt

son is certainly guilty as shown in Table 3. For each candidate belief distribution we calculate the best-case scenario under the assumption that the candidate belief distribution is both true and fully known by the decision maker. For instance, with belief b_0 in which the son is certainly innocent we see that the expected utility of proclaiming innocence is greater than that of proclaiming guilt. And finally, by maximizing over the best-case scenarios of each of the candidate belief distributions we arrive at the *desired belief state* of the decision-maker for use in our self-deceptive biasing process.

2.4 Belief integration and decision-making

The purpose of the belief integration and decision-making phase is to choose an action while considering both the rational belief $pr(s)$ and the desired belief specified via wishful thinking $pr_{\mathrm{w}}(s)$. The manner in which this final decision is reached depends on both the type and magnitude of self-deception employed.

Mele distinguishes between two distinct forms of self-deception [16]:

1. Being self-deceived into believing something that we desire to be true
2. Being self-deceived into believing something we desire to be false

We call the former *optimistic* self-deception and the latter *pessimistic* self-deception.

2.4.1 Optimistic wishful thinking

The decision rule employing optimistic wishful thinking, thereby selecting action a_{wo}, in which the agent's belief is biased *towards* the desired belief state is shown in (6). Note that s_{k} is the state specified using our wishful thinking formulation specified in (4).

$$a_{\mathrm{wo}} = \underset{a \in A}{\operatorname{argmax}} \left[(1 - \alpha) \cdot \left(\sum_{s \in S} pr(s) \cdot \mu(a, s) \right) + \alpha \cdot \mu(a, s_{\mathrm{k}}) \right] \tag{6}$$

2.4.2 Pessimistic wishful thinking

The decision rule employing pessimistic wishful thinking, thereby selecting action a_{wp}, in which the agent's belief is biased *away* from the desired belief state is shown in (7).

$$a_{\mathrm{wp}} = \underset{a \in A}{\operatorname{argmax}} \left[(1-\alpha) \cdot \left(\sum_{s \in S} pr\,(s) \cdot \mu\,(a, s) \right) - \alpha \cdot \mu\,(a, s_{\mathrm{k}}) \right] \tag{7}$$

2.4.3 Magnitude of self-deception

Both optimistic and pessimistic definitions of self-deception utilize the constant α as a representation of the magnitude or strength of the self-deceptive tendencies evinced by a decision-maker. More formally, $0 \leq \alpha \leq 1$ and is defined such that when $\alpha = 0$ the decision-maker behaves in a purely rational manner as ascribed by SEU-Theory and when $\alpha = 1$ the decision-maker behaves in a purely self-deceptive manner in which all rational evidence is rejected and the desired belief is wholly adopted in either an optimistic or pessimistic fashion.

3 Simulation

Here we present our self-deceptive framework within the context of a game commonly referred to as the "Battle of the Sexes". With these experiments we seek to illustrate the behavior of both rational and self-deceptive agents as well as explore the interaction between the two.

The "Battle of the Sexes" traditionally represents a couple attempting to coordinate their actions for the evening without the benefit of communication. Their two choices are attending either an opera or a football game. Each partner has different preferences as to which event they would like to attend. However, each partner would also rather attend their non-preferred event if it results in coordinating with their partner. At its core, the "Battle of the Sexes" is about coordination and synchronization since participants choosing to synchronize actions have higher utility both individually and collectively than they would alternatively. An illustrative utility matrix for the "Battle of the Sexes" is depicted in Table 4 in which the *row* player prefers attending a football game and the *column* player prefers the opera. Each entry in the table contains two utility values in which the first value refers to the utility received by the row player and the second value is the utility received by the column player.

3.1 Scenario setup

In order to cast the "Battle of the Sexes" into a form amenable to analysis within our framework, we must probabilistically represent beliefs. Most traditional game-theoretic analyses focus on equilibrium strategies in which the utilities for both participants is common knowledge rather than forming probabilistic beliefs regarding the action of an opponent. However,

Table 4 Example "Battle of the Sexes" payoff matrix

	Opera	Football
Opera	2,3	0,0
Football	1,1	3,2

Table 5 Utility of outcomes for Terry

	Pat attends football game	Pat attends opera
Go to opera	1	3
Go to football game	2	0

Table 6 Belief distributions for Terry

	Pat attends football game p (football)	Pat attends opera p (opera)
Desired belief	0	1
Rational belief	.5	.5

a probabilistic treatment of the game is appropriate in situations in which little or no information is available regarding a partner's preferences, strategies, or knowledge and when the only available information is probabilistic in nature, e.g., a relative frequency of past observations.

Consider the following scenario:

Example 2 Terry and Pat are players in the "Battle of the Sexes" in which Terry prefers attending the opera and Pat prefers football. We represent Terry's outcome preferences using the utilities shown in Table 5 which capture both Terry's primary goal of coordinating activities with Pat and a more general preference for opera.

Let us assume that the initial beliefs for both players indicate an equally likely chance of attending either event. Irrespective of this rational belief, Terry's *desired* belief is one in which the possibility of Pat attending the opera is certain since this allows Terry to both coordinate events with Pat and attend the opera. Table 6 shows both Terry's rational and desired belief distributions. Figure 2 is a graph of Terry's decision thresholds with respect to the various decision-making processes described in this paper. A point on the graph is designated on the

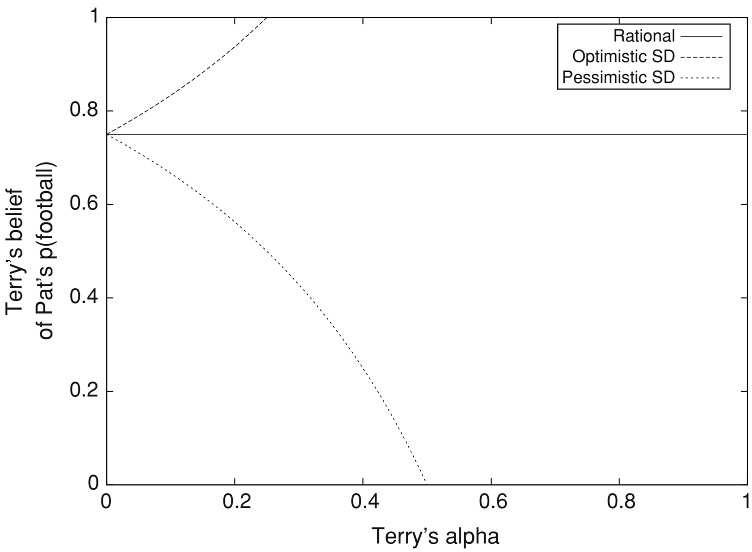

Fig. 2 Terry's decision threshold

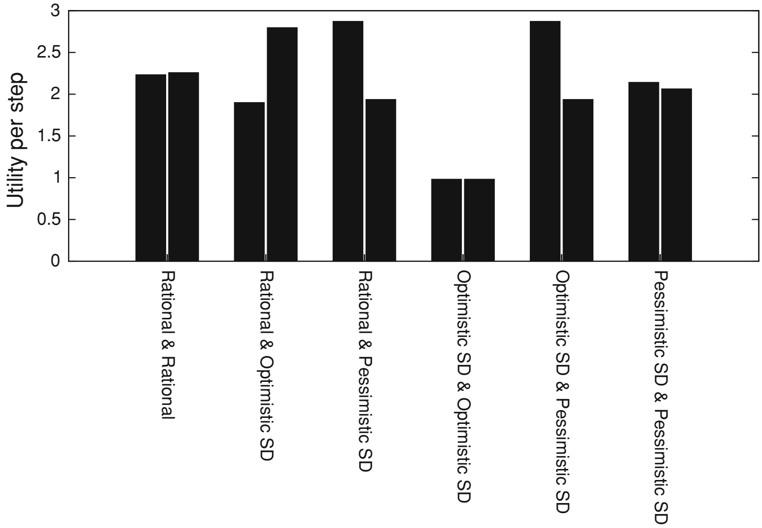

Fig. 3 Average step utility for "Battle of the Sexes"

x-coordinate by α, representing the magnitude of self-deception, and on the y-coordinate by Terry's probabilistic estimate that Pat attends the football game. If the indicated point lies above the threshold curve of Terry's decision process Terry will choose to attend the football game. If it lies below the threshold curve Terry will attend the opera. For instance, when employing an optimistic self-deceptive decision-making process with $\alpha = .2$ and a rational belief that Pat's likelihood of attending the football game is .8, Terry will choose to attend the opera. However, given the same parameters utilizing a rational decision-making process, Terry will choose to attend the football game.

3.2 Simulation results and discussion

We now present the experimental results for the scenario of Terry and Pat. In each of the six possible decision-making match-ups we average the results over 500 runs in which each game is played iteratively for 200 rounds. Figure 3 depicts a graph of the mean utility per step for each agent in all of the six possible match-up combinations. The graph in Fig. 4 shows the mean utility received in any given step for a particular match-up. Table 7 shows the approximate number of steps required in any given match-up to converge upon a stable solution of either coordination or miscoordination.

Our experimental results show that situations involving participants employing dissimilar decision styles converge more quickly to a coordination of actions than do situations in which the participants employ identical decision styles. One situation in particular consisting of two agents employing a pessimistic self-deceptive style never attains a state of coordination while the other two combinations of identical decision processes take roughly 80 steps to reach coordination in contrast to the approximately 20 steps required for the combinations of dissimilar decision processes to converge on coordination.

Another interesting aspect of our experimental results is that in situations of eventual coordination, the agent that is most optimistic has a higher individual utility, i.e., always goes to its preferred event, than its partner. In situations where both participants utilize the same decision-making strategy, each partner is equally likely to emerge as the one attend-

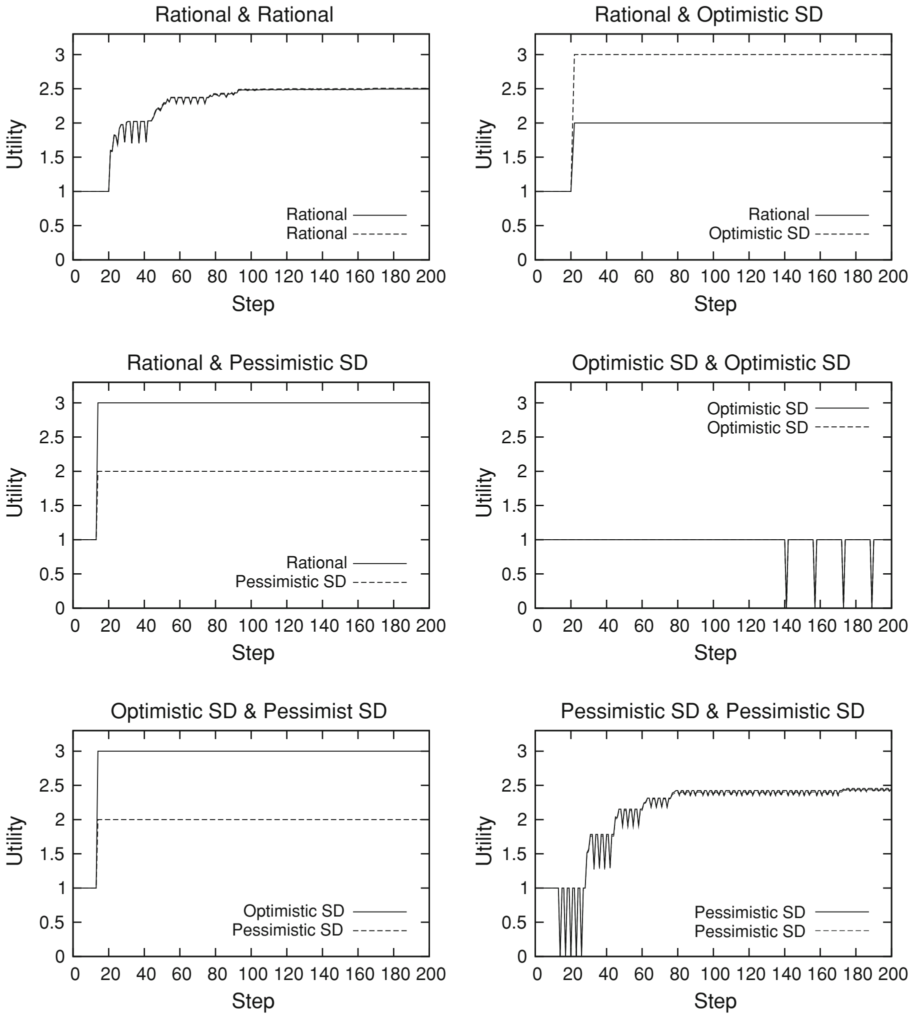

Fig. 4 Utility by step for "Battle of the Sexes"

Table 7 Convergence for "Battle of the Sexes"

Scenario	Number of steps until convergence	Convergence type
Rational and rational	80	Coordination
Rational and optimistic SD	20	Coordination
Rational and pessimistic SD	15	Coordination
Optimistic SD and optimistic SD	0	Miscoordination
Optimistic SD and pessimistic SD	15	Coordination
Pessimistic SD and pessimistic SD	80	Coordination

ing its preferred event. Here we should note that in all cases of eventual coordination, once the first coordinating event is established, agents will continue coordinating on the same event for the duration of the game. For instance, the mean step utility of roughly 2.5 for two rational participants is an average of 500 runs and indicates the equally likely possibility of the coordinating event being the preferred event of a given participant.

4 Related work

In this paper we have attempted to operationalize the psychological concept of self-deception within a decision-theoretic framework. The notion of formalizing a psychological construct within the decision-theoretic domain is not without precedent. In fact, much of the work in decision theory since the ground-breaking efforts by Von Neumann, Morgenstern and Savage have concentrated on introducing a psychological dimension into the formal decision process in order to provide for more robust and descriptive approaches. Regret Theory [2,12] models the tendency to avoid decisions that could lead to an excessive feeling of regret. Prospect Theory [21] is a purely descriptive framework that employs a series of heuristics in order to approximate the mental shortcuts that people seem to employ when making decisions. Ellsberg's Index [7] and the Ambiguity Model of Einhorn and Hogarth [6] both model the perceived aversion to uncertainty that decision-makers sometimes express.

Within the realm of self-deception, Talbott presents a model based on the desirability of adopting some preferred belief. Talbott's notion of desirability is utility-based and is a weighting of the possibility that the belief is accurate against the chance that it is not [20]. Based on this assessment, Talbott's model then calculates the expected utility for both behaving rationally and attempting to bias one's cognitive processes towards the desired belief. The primary difference between Talbott's work and the work presented in this paper is that Talbott defines the desirability of the possible belief outcomes externally while we derive that desirability using an agent's internal preference structures and then integrate the desired belief into the decision-making process using an externally defined constant representing the magnitude of deception thus further operationalizing and formalizing the notion of a desired belief with respect to a decision-theoretic framework.

In addition to decision-theoretic formulations of psychological phenomena, there exist a number of computational models of emotion and bias. The Affective Belief Revision system of Pimentel and Gravo [18] describes a logic-based system of maintaining the consistency of a propositional knowledge base in which the belief revision activities are influenced by the affective state of the individual. Other computational models of emotion [9,14] utilize self-deception as a coping mechanism to ease an agent's emotional stress. These computational models however, do not provide a manner in which to model the repercussions and trade-offs of possibly adopting a false belief. The PsychSim modeling framework [15,19], allows decision-theoretic agents to possibly influence the belief state of other agents by sending messages containing a hypothetical belief state. One factor that is assessed when evaluating these messages is self-interest. In other words, an agent will be more likely to accept a change in belief if the proposed belief is more amenable to the agent's desires and preferences. Since self-interest is evaluated entirely outside the reality of the current situation, it is in principle similar to the notion of self-deception. The key difference between these computational models of emotion and our work is that we present both the determination of the desired belief state and the subsequent process of self-deceptive belief revision all within a decision-theoretic framework.

5 Conclusion

In this work we have detailed a descriptive framework for modeling the psychological act of self-deception within a decision-theoretic environment based on the tenets of SEU-Theory. Our self-deceptive theory utilizes SEU-Theory for not only the desired belief integration and decision-making process but also for the formation of the desired belief state that is central to the biasing processes of self-deception. Through a series of experimental simulations using the "Battle of the Sexes" game formulation we've shown that our framework operationalizes both optimistic and pessimistic self-deception processes and that within certain situations, a healthy dose of self-deception is beneficial. In future work we may explore the possibility of a slightly altered and more relaxed definition of the desired belief. Specifically, rather than ignoring the reality of the situation in the formulation of the desired belief, we can choose a desired belief state *given* the current belief state. So once a course of action is chosen under the current belief state, we can then determine the outcome that is *desired*.

Acknowledgments This work was sponsored by the U.S. Army Research, Development, and Engineering Command (RDECOM), and the content does not necessarily reflect the position or the policy of the Government, and no official endorsement should be inferred.

References

1. Aronson, E. (1968). Dissonance theory: Progress and problems. *Contemporary Issues in Social Psychology, 2*, 310–323.
2. Bell, D. E. (1982). Regret in decision making under uncertainty. *Operations Research, 30*(5), 961–981.
3. Bickmore, T., Gruber, A., & Picard R. (2005). Establishing the computer–patient working alliance in automated health behavior change interventions. *Patient Education and Counseling, 59*(1), 21–30.
4. Cialdini, R. B. (2001). *Influence: Science and practice*. Allyn and Bacon.
5. Demos, R. (1960). Lying to oneself. *The Journal of Philosophy, 57*(18), 588–595.
6. Einhorn, H. J., & Hogarth, R. M. (1986). Decision making under ambiguity. *The Journal of Business, 59*(4), 225–250.
7. Ellsberg, D. (1961). Risk, ambiguity, and the savage axioms. *The Quarterly Journal of Economics*, 643–669.
8. Festinger, L. (1957). *A theory of cognitive dissonance* (Vol. 1, pp. 65–86). Evanston, IL: Row, Peterson.
9. Gratch, J., & Marsella, S. C. (2004). A domain-independent framework for modeling emotion. *Cognitive Systems Research, 5*(4), 269–306.
10. Kunda, Z. (1990). The case for motivated reasoning. *Psychological Bulletin, 108*(3), 480–498.
11. Lazarus, R. S. (1991). *Emotion and adaptation*. Oxford: Oxford University Press.
12. Loomes, G., & Sugden, R. (1982). Regret theory: An alternative theory of rational choice under uncertainty. *The Economic Journal, 92*(368), 805–824.
13. Luce, R. D., & Raiffa, H. (1957). *Games and decisions*. New York: Wiley.
14. Marsella, S. C., & Gratch, J. (2003). Modeling coping behavior in virtual humans: Don't worry, be happy. In *Proceedings of the second international joint conference on autonomous agents and multiagent systems* (pp. 313–320).
15. Marsella, S. C., Pynadath, D. V., & Read, S. J. (2004). PsychSim: Agent-based modeling of social interactions and influence. In *Proceedings of the international conference on cognitive modeling* (pp. 243–248).
16. Mele, A. R. (1997). Understanding and explaining real self-deception. *Behavioral and Brain Sciences, 20*(01), 127–134.
17. Mele, A. R. (2000). *Self-deception unmasked*. Princeton University Press.
18. Pimentel, C. F., & Gravo, M. R. (2005, December 5–8). Affective revision. In *Progress in artificial intelligence: 12th Portuguese conference on artificial intelligence, EPIA 2005, Covilhã, Portugal, Proceedings*.
19. Pynadath, D. V., & Marsella, S. C. (2005). PsychSim: Modeling theory of mind with decision-theoretic agents. In *Proceedings of the international joint conference on artificial intelligence*, (pp. 1181–1186).
20. Talbott, W. J. (1995). Intentional self-deception in a single coherent self. *Philosophy and Phenomenological Research, 55*(1), 27–74.
21. Tversky, A., & Kahneman, D. (1979). Prospect theory: An analysis of decision under risk. *Econometrica, 47*(2), 263–292.