Contextually-Based Utility: An Appraisal-Based Approach at Modeling Framing and Decisions

Jonathan Ito and Stacy Marsella

Institute for Creative Technologies
University of Southern California
12015 Waterfront Drive
Playa Vista, CA 90094-2536 USA
{ito,marsella}@ict.usc.edu

Abstract. Creating accurate computational models of human decision making is a vital step towards the realization of socially intelligent systems capable of both predicting and simulating human behavior. In modeling human decision making, a key factor is the psychological phenomenon known as "framing", in which the preferences of a decision maker change in response to contextual changes in decision problems. Existing approaches treat framing as a one-dimensional contextual influence based on the perception of outcomes as either gains or losses. However, empirical studies have shown that framing effects are much more multifaceted than one-dimensional views of framing suggest. To address this limitation, we propose an approach which combines the psychological principles of cognitive appraisal and decision-theoretic notions of utility and probability. We show that this approach allows for both the identification and computation of the salient contextual factors in a decision as well as modeling how they ultimately affect the decision process. Furthermore, we present preliminary fitting analyses which suggest that our multi-dimensional, appraisal-based approach is more descriptively accurate for realistically-framed decision scenarios than competing decision models.

1 Introduction

Creating accurate computational models of human decision making is a vital step towards the realization of socially intelligent systems capable of both predicting and simulating human behavior. This is critical across a diverse range of applications such as those found in social simulation [5], virtual-human based training [22], interactive health intervention [18, 3], and embodied conversational agents [11]. Additionally, accurate decision models are a necessary component in *prescriptive* decision systems, which assist people in making rational and goal-oriented decisions.

One key factor in the descriptive modeling process is the "framing" of a decision. Framing describes the manner in which a decision scenario is presented, or embedded, in some context. Research has shown that the context in which a decision is framed has a profound impact on the manner in which it is *perceived*, *interpreted*, and ultimately *evaluated*. In a seminal study involving what is now referred to as the Asian Disease Scenario (ADS), Kahneman and Tversky [33] demonstrate the impact that framing can have in even relatively simple decision scenarios. The scenario, shown in Listing 1, consists of two distinct frames in which the same underlying outcomes are described as either gains (positive frame) or losses (negative frame). For instance, the outcome of Program A in the positive frame is equivalent to the outcome in Program C of the negative frame; both describe a situation in which 200 people live while 400 die. In the original study, a majority of respondents (72%) in the positive frame preferred the risk-averse choice of Program A. However, when presented with the negative frame, a clear majority of respondents (78%) preferred the risk-seeking choice of Program D. This dramatic reversal of choice, observed when the same underlying decision problem is presented in differing frames, is commonly referred to as the "framing effect" and underscores the impact that framing can have on decisions.

Imagine that the U.S. is preparing for the outbreak of an unusual Asian disease, which is expected to kill 600 people. Two alternative programs to combat the disease have been proposed. Assume that the exact scientific estimate of the consequences of the programs are as follows:
Survival Frame:

- Program A: 200 people will be saved.
- Program B: 1/3 probability that 600 people will be saved and 2/3 probability that none will be saved.

Mortality Frame:

- Program C: 400 people will die.
- Program D: 1/3 probability that nobody will die and 2/3 probability that 600 people will die.

Listing 1: Asian Disease Scenario

Subsequent studies involving domains as diverse as financial planning [28], taxes [12] and Acquired Immune Deficiency Syndrome [15] have also demonstrated framing effects to varying degrees.

Existing approaches for modeling framing include Cumulative Prospect Theory (CPT) [34] and the Security-Potential/Aspiration Model (SP/A) [16]. However, an inherent limitation of both approaches is that framing is treated as a one-dimensional contextual influence predicated solely on the perception of outcomes as gains or losses. While both models still account for the classic framing effect, i.e., risk aversion in gains and risk seeking in losses, numerous empirical studies within the framing literature have shown a great deal of variance and inconsistency in the classic framing effect which suggests that a one-dimensional view of framing is descriptively incomplete [27, 4, 8, 7].

In this work, we propose a framework for decision making under risk which is descriptively flexible enough to model a wide range of framing effects while remaining intuitive and easily operationalizable. Specifically, we propose that by employing psychological principles of cognitive appraisal we can identify, extract, and compute the salient contextual aspects of a given decision scenario in a highly principled manner. This appraisal information can then be used to inform a weighted utility function resulting in a decision process which is descriptively more accurate than existing approaches. Additionally, we present preliminary fitting analyses which suggest that for more realistically-framed decision scenarios, a Contextually-Based Utility model is more descriptively accurate than both CPT and SP/A models.

2 Contextually-Based Utility (CBU)

Our approach in creating a contextually-based, descriptive, decision framework consists of two primary components: the computational appraisal of decision outcomes and a decision evaluation function incorporating the appraisal information [13].

2.1 Cognitive Appraisal of Outcomes

The process of cognitive appraisal, as defined in Appraisal Theory, is the evaluation of personal significance over several dimensions regarding an individual's relationship with herself, the environment, and others by which emotions are differentiated and elicited for a given stimulus [1, 14]. It is precisely this assessment of personal significance, with respect to the environment of the decision maker, which provides us with a rich discretization of the key contextual factors, or situational variables, involved in a decision.

In this work, computational appraisal is defined as a function of diminishing sensitivity evaluated with respect to some reference point. This definition represents both the tendency of emotional habituation, in which continually increasing affect loses poignancy, and the notion that emotions and appraisals are elicited not by the outcomes themselves but by the *changes* accompanying them [10]. This suggests an S-shaped function in which the reference point serves as the point of inflection. Therefore, a logistic function is employed to model diminishing sensitivity, which forms the computational basis of appraisal, shown in (1), in which the impact of t diminishes as it moves further from the reference point r and where the constant k, such that $k \geq 1$, controls the degree of diminishing sensitivity. The result is a real value between -1 and 1 such that positive values correspond to positive appraisals, e.g., appraisals of pleasantness, and negative values correspond to negative appraisals, e.g., appraisals of unpleasantness.

$$\mathrm{dimsens}\left(t, r, k\right) = \frac{2}{1.0 + \exp\left(-k\left(t - r\right)\right)} - 1 \tag{1}$$

In this work, we present three concrete appraisal functions defined over outcomes: pleasantness, goal congruence, and control. These dimensions are chosen both because they are widely represented in the appraisal literature and are also linked to notions of probability and value, both highly pertinent in decision-theoretic domains.

It is important to note that our goal in this work is *not* to implement a full computational model of appraisal, as many promising implementations already exist [11, 2, 6, 29]. Rather, we seek to leverage the principles underlying cognitive appraisal to model the effect of framing on decision processes.

In what follows, a decision scenario is represented using a standard decision-theoretic representation. A scenario consists of n distinct alternatives, in which the ith alternative is labeled x_i. Each alternative results in one of n_i possible outcomes, in which the jth outcome is denoted as outcome $x_{i,j}$. The probability of obtaining outcome $x_{i,j}$, when alternative x_i is chosen, is given by the probability function $\mathrm{p}\left(x_{i,j}\right)$. Additionally, the associated value of outcome $x_{i,j}$ is $\mathrm{v}\left(x_{i,j}\right)$ and is generally encoded directly from the numerical description of an outcome, e.g., amount of money won, number of lives lost.

Pleasantness. Pleasantness is the intrinsic attractiveness or unattractiveness of an outcome. Appraisals of pleasantness play a central role in many predominant structural theories of cognitive appraisal [23, 26, 31, 20]. Pleasantness is an evaluation of value made with respect to the *status quo* (sq) as in (2).

$$\mathrm{pleas}\left(x_{i,j}\right) = \mathrm{dimsens}\left(\mathrm{v}\left(x_{i,j}\right), sq, \mathrm{k}\right) \tag{2}$$

The status quo serves as the reference point differentiating pleasant from unpleasant outcomes. In other words, outcomes more preferred than the status quo are appraised as pleasant, whereas less-preferred outcomes are perceived as unpleasant. The status quo itself is the state resulting in no change of current state for the decision-maker. In many decision scenarios, an explicit status quo is not provided. Oftentimes, the most natural location for it, and which is usually implied by convention, is the state associated with the 0-value, e.g., 0 lives lost, 0 dollars gained.

Goal Congruence. Goal congruence is the degree to which an outcome fulfills the adopted goals of the decision maker which are heavily influenced by standards, expectations, and responsibilities. Goal congruence, like pleasantness, is well represented in the appraisal literature [26, 23, 32]. We represent goal congruence as an evaluation of value made with respect to the *aspiration outcome* (ao) as in (3).

$$\mathrm{gc}\left(x_{i,j}\right) = \mathrm{dimsens}\left(\mathrm{v}\left(x_{i,j}\right), ao, \mathrm{k}\right) \tag{3}$$

The aspiration outcome is the outcome which the decision maker aspires to and represents her adopted expectations, standards, morals, and obligations.

Outcomes more preferred than the aspiration outcome are said to be *congruent* with the goals of the decision maker whereas less-preferred outcomes are *incongruent*. It can be argued that a rational location for the aspiration outcome is the maximum expected value of the scenario. However, as noted earlier, the imposition of external standards and expectations, even when self-imposed, may influence the aspiration outcome.

Control. Control is a measure of the degree of agency and influence that a decision maker has in a scenario. It is conceptualized as both a form of agency [23, 31] and as a critical component in coping potential [32, 26]. Control is an evaluation of decumulative probability made with respect to some *expectation of control* (ec) as seen in (4). The decumulative aspect of control, shown in (5), is the total probability of achieving an outcome as good as $x_{i,j}$, in which outcomes are ordered from worst to best by value. This corresponds to the intuition that even a highly uncertain outcome, i.e., an outcome with a very low probability of occurrence, can be perceived as highly controllable if the other alternate outcomes are even more preferred. In other words, control is implemented not as a function of likelihood over the occurrence of a *singular* outcome but rather as a function of likelihood over the achievement of a certain value threshold.

$$\mathrm{ctrl}\left(x_{i,j}\right) = \mathrm{dimsens}\left(\mathrm{d}\left(x_{i,j}\right), ec, \mathrm{k}\right) \tag{4}$$

$$\mathrm{d}\left(x_{i,j}\right) = \sum_{k=j}^{n_i} \mathrm{p}\left(x_{i,k}\right) \tag{5}$$

The expectation of control establishes a reference point by which it is judged whether outcomes are controllable or uncontrollable. While the expectation of control is certainly influenced by situational variables such as the manner in which the uncertainty in a scenario is described, it is also subject to personal and individual factors such as locus of control [25]. The concept of locus of control distinguishes between two sources of control: internal and external. Internals tend to believe that they control their own destiny and that outcomes are largely determined by their efforts which is suggestive of a lower expectation of control. That is, outcomes need only exceed a fairly low threshold of likelihood in order to be perceived as controllable. Externals, on the other hand, tend to attribute successes or failures to external forces such as nature, destiny, and other agents which is suggestive of a higher expectation of control.

2.2 Contextually-Based Utility (CBU) Decision Evaluation

The second component in our approach to modeling contextually-based decision making is the implementation of a decision function over the alternatives available in a decision.

The CBU function evaluates an alternative, x_i, and maps it to a real-value such that x_i is preferred at least as much as alternative x_j if and only if

$\mathrm{CBU}(x_i) \geq \mathrm{CBU}(x_j)$. The CBU function takes the form of a standard decumulative weighted utility function, also commonly referred to as a rank-dependent utility model [21]; it is comprised of both a weighting component, w, and a utility function, u, as in (6) in which outcomes are ordered from worst to best. The decumulative component, $\mathrm{d}(x_{i,j})$, is the total probability that an outcome *as good as* $x_{i,j}$ will be obtained and is defined identically to (5).

$$\mathrm{CBU}(x_i) = \sum_{j=1}^{n_i} \mathrm{w}(\mathrm{d}(x_{i,j}))(\mathrm{u}(x_{i,j}) - \mathrm{u}(x_{i,j-1})) \tag{6}$$

The decumulative form is employed to allow for the variable weighting of outcomes, such as either an optimistic over-weighting of desirable outcomes or a pessimistic over-weighting of undesirable ones, while preserving stochastic dominance. Stochastic dominance, stated simply, is a highly desirable decision-theoretic principle asserting that an alternative x_i should be preferred at least as much as alternative x_j if all outcomes of x_i are preferred at least as much as all outcomes of x_j.

In what follows, we show how the various dimensions of our previously defined appraisal process relate to the weighting and utility components of the CBU function. Specifically, we show that appraisals of control inform the weighting component while appraisals of pleasantness and goal congruence form a multi-attribute utility function.

Control as a Decision Weight. The decision-weighting function, a function of probability and likelihood, represents the intensity and relative importance of a particular outcome. Note that both our implementation of control, shown in (4), and the weighting function $\mathrm{w}(x_{i,j})$ are defined as functions of decumulative probability. Therefore, we use the previously defined appraisal of control to inform the weighting function.

The full control-based decumulative-weighting function, w, is given in (7) in which the weighting function is the appraisal of control normalized such that its range is between 0 and 1.

$$\mathrm{w}(x_{i,j}) = \frac{1}{2}\,\mathrm{ctrl}(x_{i,j}) + \frac{1}{2} \tag{7}$$

Pleasantness and Goal Congruence as a Multi-Attribute Utility Function. The utility function, which maps an outcome to a real value, is representative of the "goodness" of an outcome and is independent of considerations of weight or probability. We therefore employ the appraisals of pleasantness and goal congruence, both based on evaluations of value, to inform the utility function.

We represent utility as a multi-attribute utility function which models the trade-off between considerations of pleasantness and goal congruence as in (8), where $0 \leq \alpha \leq 1$ and in which α and $1-\alpha$ are the relative weights given to considerations of pleasantness and goal congruence respectively.

$$\mathrm{u}\left(x_{i,j}\right) = \alpha\,\mathrm{pleas}\left(x_{i,j}\right) + \left(1-\alpha\right)\mathrm{gc}\left(x_{i,j}\right) \tag{8}$$

3 Descriptive Accuracy and Fitting

In what follows, we replicate existing work originally conducted by Lopes and Oden [17], comparing the relative accuracy, in terms of parametric fitting and the root mean squared deviation (RMSD), of both Cumulative Prospect Theory (CPT) [34] and the Security-Potential/Aspiration Model (SP/A) [16]. Additionally, we present preliminary work suggesting that Contextually-Based Utility (CBU) is descriptively more accurate than both CPT and SP/A for more realistically-framed decision scenarios.

In their analysis of CPT and SP/A, Lopes and Oden collected and aggregated subject preference data over various pairs of lotteries. Lotteries in each pairing are roughly equivalent in expected utility and differ only in the value of each outcome and their overall probability distribution. In total, 6 distinct lotteries containing strictly positive outcomes were created resulting in 15 distinct lottery pairings. Additionally, the outcomes of the 6 distinct lotteries were also scaled by a positive constant or alternatively shifted by the addition of a positive constant to create 2 additional groups of 6 lotteries for an additional 30 pairings. Moreover, the outcomes of all lotteries are also negated resulting in 3 *negative* lottery groups for an additional 45 distinct *negative* pairings.

In addition to replicating the original findings of Lopes and Oden, which show that SP/A converges to a more accurate fit than CPT given the original lottery data, we also add a fitting analysis of both Expected Utility and our previously described Contextually-Based Utility (CBU) process. Furthermore, a representative sampling of empirical framing data is taken from a range of domains and empirical studies shown in Table 1 and separately analyzed with respect to fitting accuracy.

The results of both the fitting for the lottery data and the framing data is shown in Figure 3 where the y-axis is indicative of the percentage of *improvement*, in terms of the root mean square deviation (RMSD), over an expected utility optimization. Our analysis suggests that while the descriptive flexibility afforded by appraisal theory may not be fully utilized in relatively simple, inherently one-dimensional decision scenarios, such as choosing between probabilistic lotteries, it becomes increasingly important as the contextual complexity and realism of scenarios increase.

4 Conclusion

Computational models of human decision making are becoming more important in a wider range of applications. Therefore, it is vital that these models can

Table 1. Empirical Framing Studies

Author	**Domain**
Tversky & Kahneman [33]	Asian Disease
Miller [19]	Asian Disease
Fagley [9]	Asian Disease, Cancer, School Dropout Prevention, Job Layoffs
Sieck [30]	Asian Disease
Roszkowski [24]	Financial Investment
Fagley [8]	Cancer
Fagley [7]	School Dropout Prevention
Levin [15]	AIDS Treatment

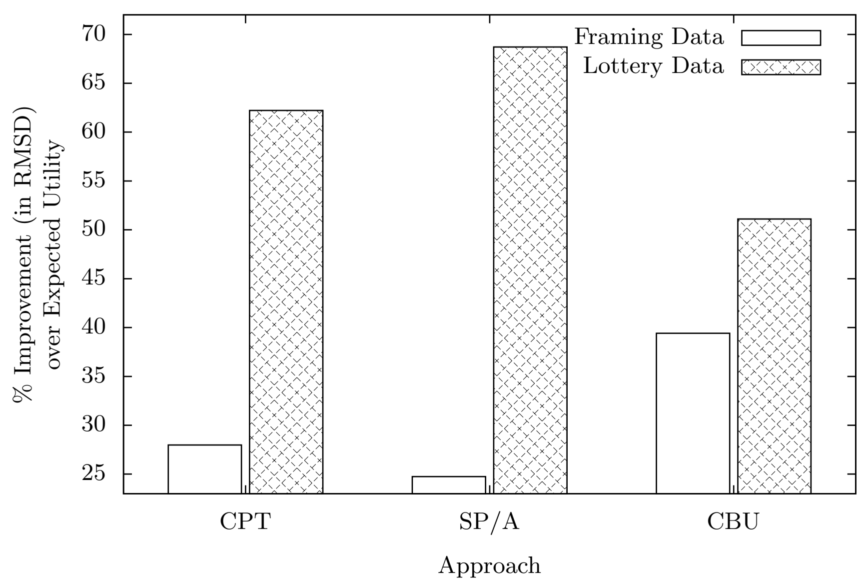

Fig. 1. Fitting Accuracy of Lottery and Framed Scenario Preferences

account for a broad range of human behavior while remaining readily implementable. One well-established hallmark of human decision behavior, typically problematic for theories of choice, is the phenomenon of framing, by which the same underlying problem described differently produces different results.

Existing approaches at modeling framing treat it as a one-dimensional contextual influence and therefore lack the descriptive flexibility to account for a broad range of behavior. Therefore, we have proposed a contextually-sensitive decision framework which integrates the psychological principles of cognitive appraisal with traditional decision-theoretic approaches for modeling preferences. In particular, we have shown that appraisals of pleasantness and goal congruence can be used in a multi-attribute utility function while appraisals of control inform the weighting function.

Furthermore, we have shown through preliminary analyses that the descriptive flexibility offered by cognitive appraisal and decumulative weighted utility in the CBU framework may be significantly more accurate than competing decision models in more realistically-framed and contextually-rich decision scenarios.

References

1. Arnold, M.: Emotion and personality. Columbia University Press New York (1960)
2. Aylett, R., Dias, J., Paiva, A.: An affectively-driven planner for synthetic characters. Proceedings of ICAPS 2006 (2006)
3. Bickmore, T., Gruber, A., Picard, R.: Establishing the computer–patient working alliance in automated health behavior change interventions. Patient Education and Counseling 59(1), 21–30 (2005)
4. Bohm, P., Lind, H.: A note on the robustness of a classical framing result* 1. Journal of Economic Psychology 13(2), 355–361 (1992)
5. Davidsson, P.: Multi agent based simulation: beyond social simulation. Multi-Agent-Based Simulation pp. 141–155 (2001)
6. Dias, J., Paiva, A.: Feeling and reasoning: A computational model for emotional characters. Progress in Artificial Intelligence pp. 127–140 (2005)
7. Fagley, N., Kruger, L.: Framing effects on the program choices of school psychologists. In: Annual Meeting of the American Psychological Association, Los Angeles, CA (1986)
8. Fagley, N., Miller, P.: The effects of decision framing on choice of risky vs certain options. Organizational Behavior and Human Decision Processes 39(2), 264–277 (1987)
9. Fagley, N., Miller, P.: The effect of framing on choice: Interactions with risk-taking propensity, cognitive style, and sex. Personality and Social Psychology Bulletin 16(3), 496 (1990)
10. Frijda, N.: The laws of emotion. Lawrence Erlbaum Associates (2007)
11. Gratch, J., Marsella, S.: A domain-independent framework for modeling emotion. Cognitive Systems Research 5(4), 269–306 (2004)
12. Highhouse, S., Paese, P.: Problem domain and prospect frame: Choice under opportunity versus threat. Personality and Social Psychology Bulletin 22(2), 124 (1996)
13. Ito, J.Y., Marsella, S.C.: Contextually-based utility: An appraisal-based approach at modeling framing and decisions. In: Proceedings of the 25th Conference by the Association for the Advancement of Artificial Intelligence AAAI (2011), to appear

14. Lazarus, R.: Emotion and adaptation. Oxford University Press, USA (1991)
15. Levin, I., Chapman, D.: Risk taking, frame of reference, and characterization of victim groups in aids treatment decisions. Journal of Experimental Social Psychology 26(5), 421–434 (1990)
16. Lopes, L.: Between hope and fear: The psychology of risk. Advances in experimental social psychology 20(3), 255–295 (1987)
17. Lopes, L., Oden, G.: The role of aspiration level in risky choice: A comparison of cumulative prospect theory and sp/a theory. Journal of mathematical psychology 43(2), 286–313 (1999)
18. Marsella, S., Johnson, W., LaBore, C.: Interactive pedagogical drama for health interventions. In: Conference on Artificial Intelligence in Education, Sydney, Australia. Citeseer (2003)
19. Miller, P., Fagley, N.: The effects of framing, problem variations, and providing rationale on choice. Personality and Social Psychology Bulletin 17(5), 517 (1991)
20. Ortony, A., Clore, G., Collins, A.: The cognitive structure of emotions. Cambridge Univ Pr (1990)
21. Quiggin, J.: A theory of anticipated utility. Journal of Economic Behavior & Organization 3(4), 323–343 (1982)
22. Rickel, J., Marsella, S., Gratch, J., Hill, R., Traum, D., Swartout, W.: Toward a new generation of virtual humans for interactive experiences. IEEE Intelligent Systems, pp. 32–38 (2002)
23. Roseman, I.: Cognitive determinants of emotion: A structural theory. Review of Personality & Social Psychology 5, 11–36 (1984)
24. Roszkowski, M., Snelbecker, G.: Effects of "framingon measures of risk tolerance: Financial planners are not immune. Journal of Behavioral Economics 19(3), 237–46 (1990)
25. Rotter, J.: Generalized expectancies for internal versus external control of reinforcement (1966)
26. Scherer, K.: Emotion as a process: Function, origin and regulation. Social Science Information 21(4-5), 555 (1982)
27. Schneider, S.: Framing and conflict: Aspiration level contingency, the status quo, and current theories of risky choice. Journal of Experimental Psychology: Learning, Memory, and Cognition 18(5), 1040–1057 (1992)
28. Schoorman, F., Mayer, R., Douglas, C., Hetrick, C.: Escalation of commitment and the framing effect: An empirical investigation. Journal of Applied Social Psychology 24(6), 509–528 (1994)
29. Si, M., Marsella, S., Pynadath, D.: Modeling appraisal in theory of mind reasoning. Autonomous Agents and Multi-Agent Systems 20(1), 14–31 (2010)
30. Sieck, W., Yates, J.: Exposition effects on decision making: Choice and confidence in choice. Organizational Behavior and Human Decision Processes 70(3), 207–219 (1997)
31. Smith, C., Ellsworth, P.: Patterns of cognitive appraisal in emotion. Journal of personality and social psychology 48(4), 813–838 (1985)
32. Smith, C., Lazarus, R.: Emotion and adaptation. Handbook of personality: Theory and research pp. 609–637 (1990)
33. Tversky, A., Kahneman, D.: The framing of decisions and the psychology of choice. Science 211(4481), 453 (1981)
34. Tversky, A., Kahneman, D.: Advances in prospect theory: Cumulative representation of uncertainty. Journal of Risk and uncertainty 5(4), 297–323 (1992)