Internal Use Only

TR-I-0342

The Need for Second-order Probability Distributions Under Repeated Trials with Nonlinear Utilities or Catastrophic Outcomes

John K. Myers

March 12, 1993

Abstract

Previous researchers have indicated that it is not necessary to use full second-order probability distributions when making decisions: the first-order equivalent probability is sufficient. This paper shows that, under certain conditions, this is incorrect.

© ATR Interpreting Telephony Research Laboratories
© ATR 自動翻訳電話研究所

The Need for Second-order Probability Distributions Under Repeated Trials with Nonlinear Utilities or Catastrophic Outcomes

John K. Myers

ATR Interpreting Telephony Research Laboratories
Sanpeidani, Inuidani, Seika-cho, Soraku-gun, Kyoto 619-02, Japan
myers@atr-la.atr.co.jp

Abstract

Previous researchers have indicated that it is not necessary to use full second-order probability distributions when making decisions: the first-order equivalent probability is sufficient. This paper shows that, under certain conditions, this is incorrect.

First, the philosophical meaning of a "second-order probability" is discussed, and two interpretations are offered. Under the *Nondeterministic Probability (NDP)* interpretation, the "actual probability" of an event is a random variable that keeps changing; this invalidates the theory of a constant first-order probability and is thus judged unacceptable. Under the *Unknown Deterministic Probability (UDP)* interpretation, the "actual probability" is a constant with an uncertainly known value; this interpretation is judged to be realistic. Utility is typically a nonlinear function of value. It is demonstrated that when using a nonlinear utility function, under the UDP interpretation, with repeated trials, and maximizing expected utility, the use of a second-order probability distribution gives results that are significantly different from those obtained using the equivalent first-order probability.

A significant subclass of nonlinear functions is the catastrophe surface. Problems with *catastrophic outcomes* can be represented by catastrophe-surface utility functions of value. Catastrophic outcomes are useful for modeling emotional problems such as user acceptance, common-sense reasoning problems such as naive physics, and limited-resource problems such as spending-allocation tasks. Under two-way catastrophic outcomes, advantaged agents will tend to be conservative, while disadvantaged agents will tend to choose highly uncertain actions.

The results indicate that full second-order probability distributions should be used for making decisions regarding situations with nonlinear utility or catastrophic outcomes, when an action may have to be repeated.

© 1991, 1993 ATR Interpreting Telephony Research Laboratories

1 Introduction

In a landmark paper, Cheeseman states that when using a second-order probability distribution to make a decision, "*exactly the same decision is reached whether a point value or a density function is used*"(italics original) [Che85]. Kyburg also states that "so-called second order probabilities have nothing to contribute conceptually to the analysis and representation of uncertainty." [Kyb89].

Although these statements are accurate under some conditions, they are not true in a general sense. Second-order probabilities are indeed useful when dealing with repeated trials in cases in which the utility curve is a nonlinear function of value. A special case of a nonlinear curve occurs when the utility curve forms a catastrophe surface. This case will be defined further in the paper, and will be referred to as a *catastrophic outcome*. This paper will demonstrate that a representation of uncertain outcomes employing second-order probabilities obtains different, more realistic results from a first-order probability representation in these cases, and thus that second-order probability distributions are in fact required to make accurate decisions.

2 Basic Assumptions

Throughout this paper, objective statistical examples will be used, based on concrete repeatable trials. We make no claims as to whether probability is "actually" an objective statistical ratio or a subjective normative opinion (see [Che85]). Obviously, if our results hold for objective statistical cases, then subjective opinions can be formed based on projections of what could happen. Similarly, our results can be extrapolated from concrete events to subjective opinions of hypothetical, imaginary, or nonrepeatable events.

It will be assumed that the acting agent makes decisions based on *maximizing expected utility* (MEU). Note that utility may be a highly nonlinear function of *value*, the prima facie outcome of an action; this paper departs sharply from maximizing expected value (MEV)[1]. The definition of "utility" and MEU seems to be broad enough to cover almost all practical objective or normative decision-making situations. Note, however, that empirical results in human decision-making can contradict the MEU theory [Rac89].

3 Statement of the Problem

A concrete problem will be used for illustration purposes. Suppose there is a natural-language understanding program that has the task of comprehending the implications of an input utterance. The program has a body of inference rules (the *rule set*) which it uses to construct new facts. A single *trial* consists of choosing a rule from the set and applying it to the current knowledge-base. Either the rule will match successfully and generate a new fact, or the rule application will fail, with a certain probability of success. The success or failure outcome is randomly and impartially determined by an entity known as *the Universe*. The program receives a cumulative *score* of +1 for each successful new fact, and -1 for each failure (due to temporal cost). The trial is repeated 100 times for each utterance, constituting a *run*. Possible scores thus range from +100 to -100. The program is not permitted to stop in the middle of a run; the number of trials is predetermined.

The program has the option to choose between different rule sets. Each rule set has a different characteristic probability of success, corresponding to how well it is tuned to the types of utterances received. Statistics for each rule set are displayed in this paper by simulating

[1] Also called maximizing "expected monetary value" EMV in [Rai68].

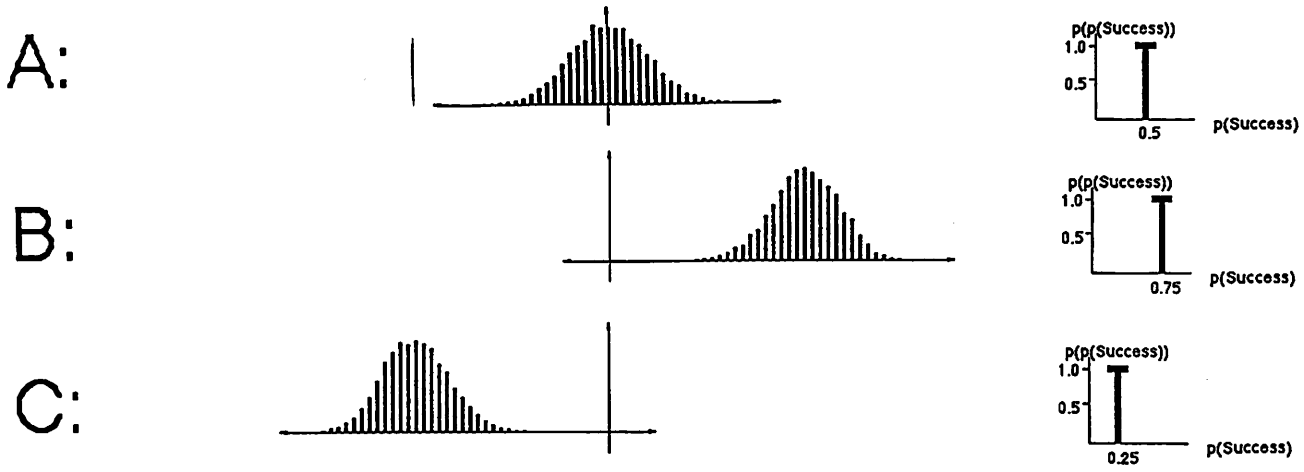

Figure 1: 10,000 Runs with First-order Probabilities of 0.5, 0.75, and 0.25

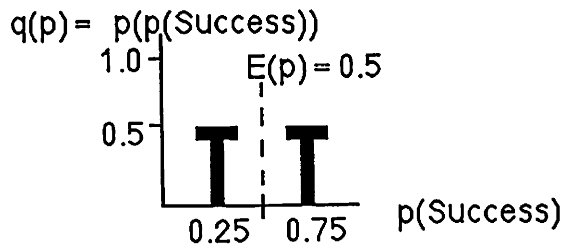

Figure 2: An Uncertain Second-order Probability

10,000 different runs for that set, representing 10,000 different *possible worlds*, and plotting the histogram of the resulting distribution. (The rule-sets discussed in this paper are hypothetical constructs given for illustrative purposes, represented by simulations only.)

For example, assume that rule set A is known with certainty to be successful with a normal (first-order) probability of 0.5. The resulting familiar distribution is displayed in the first part of Figure 1.[2] Results of first-order probabilities of 0.75 and 0.25 for sets B and C are displayed in the second and third parts, respectively.

Now assume that the correct probability of matching is not known with certainty. There is a 50% chance that it could be 0.75, while there is a 50% chance that it could be 0.25. The resulting *second-order probability* is shown in Figure 2. Note that the *effective probability E(p)*, equal to the weighted mean of the distribution,[3] is 0.5. There are at least two possible philosophical interpretations of what this representation could mean, which result in widely differing outcomes.

4 Second-order Probability as Nondeterministic Probability (NDP)

One interpretation of second-order probability is that the "actual" probability of an event is a nondeterministic **random variable** which varies according to the second-order distribution. There is no such thing as *the* (single) probability of an event; rather, probability is a *variable* which *keeps changing every time an experiment is performed.* Sometimes the actual probability is high, while sometimes it is low, for the *same type of experiment.* In effect, the Universe probabilistically keeps switching the actual world between a set of multiple possible worlds each time a trial occurs. The nondeterministic current value of this probability then determines whether the event happens or not. This interpretation is known as the *Nondeterministic*

[2] An aspect ratio of 1:40 is consistently used throughout the paper to give clean graphs. Actual graphs are 40 times taller. The zebra stripes are due to no 0 incremental score.

[3] $E(p) = \int_0^1 p \cdot q(p)\, dp$, where $q(\mathrm{p}) \equiv p(\mathrm{p} = p_i)$ is defined as the second-order distribution.

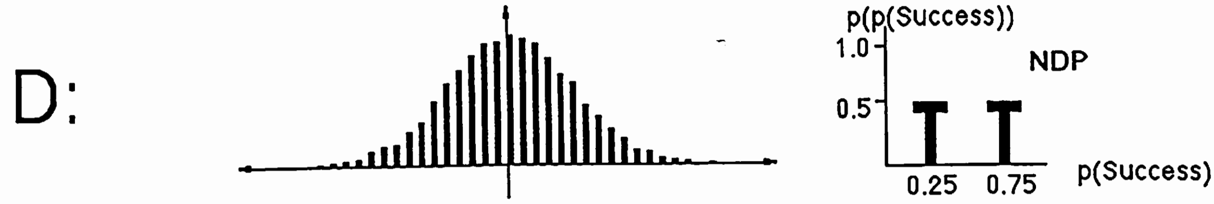

Figure 3: Nondeterministic Probability (NDP): Repeated Single Trials

Probability (NDP) assumption.

To represent this, imagine that you are requesting the Universe for balls from an urn. Imagine that the Universe has a cave full of urns. Every time you request a ball, the Universe randomly pulls an urn out of the cave, draws a ball from the urn, reads you the color, and replaces the ball and the urn. The urn, representing "real" probabilities in the "real world", keeps changing.

A run using this interpretation is known as a **repeated single trial** experiment. In our simulation of rule-set D, for each *trial*, the probability of success is determined randomly by the Universe (in this case, going with 0.25 or 0.75, each for half of the time). The Universe then uses this probability to determine the success or failure of the trial. The results for the repeated single trial are shown in Figure 3. Note that the results are identical with the distribution from the first-order probability corresponding to the effective probability of the distribution for rule-set D, i.e. rule-set A. Note also that the expected value of the results is zero. An agent making decisions under the NDP interpretation could always use the first-order effective probability $E(p)$ (in this case, rule-set A) with identical results.

This interpretation is quite radical. It rejects the very existance of unique first-order probabilities in all cases where the probability is not completely certain (i.e., all practical cases). The universe keeps changing probabilities behind our backs. It is possible that this interpretation could be useful in quantum mechanics[4]; however, this seems difficult to reconcile with macro-scale behavior. Probability and statistics are based on the assumption that the probability of an event is unique, and does not keep changing back and forth. Thus, this interpretation is judged to be unacceptable.

5 Second-order Probability as Unknown Deterministic Probability (UDP)

The second interpretation is that an objective probability of success actually does exist in the universe. It is logically consistent to discuss *the* probability of a trial, and perhaps hypothetically possible to be able to measure it. However, there may be several different conflicting sources of evidence as to what this probability could be. Thus, the actual probability is *deterministic* but *unknown*. In effect, the Universe probabilistically picks a single world from a set of possible worlds, and then uses only that world thereafter. This interpretation is known as the *Unknown Deterministic Probability* (UDP) assumption.

This interpretation corresponds to the Universe randomly picking a single urn from the cave; you are not sure which one, however. When you request a ball, the Universe draws a ball from the urn, reads you its color, and then replaces the ball. However, the urn is not replaced;

[4]For instance, the Everett interpretation states that the actual world keeps changing from an almost infinite number of parallel possibilities [Gri84].

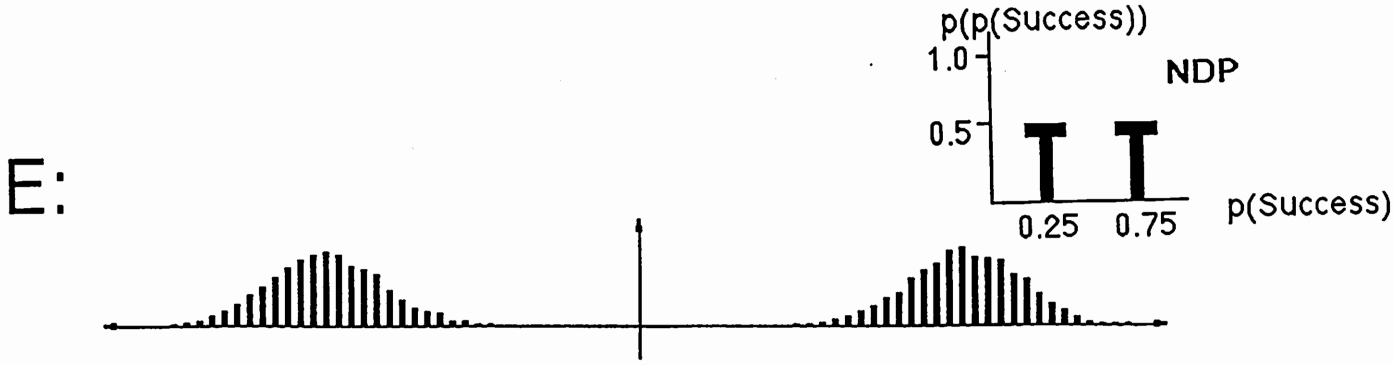

Figure 4: Unknown Deterministic Probability: Repeated Trials

the same urn is always used in the same situation.

A run using this interpretation is known as a **repeated trial** experiment. In the simulation for rule-set E, for each *run*, (each *possible world*,) the probability of success is determined randomly by the Universe (again going with 0.25 or 0.75, each for half of the time). The Universe then uses this probability to determine the success or failure of each trial in the run. The results for the repeated trial are shown in Figure 4. Note that the results are identical with the weighted average of the results from the corresponding first-order probabilities, i.e. 50% of rule-set B plus 50% of C. Note also that the expected value of the results is again zero. An agent making decisions based only on expected value could use $E(p)$ (rule-set A) and come up with the same results.

The UDP interpretation of second-order probability is not the same as using marginal probabilities, unless the second-order distribution is trivially defined as a representation of marginal first-order probabilities conditioned on a (second-order) random selection of which world is actual. Marginal probabilities can be collapsed to an equivalent first-order probability, while this paper shows that the second-order distribution cannot be collapsed.

The UDP interpretation intuitively seems to be correct. First-order probabilities do in fact exist in the universe; they simply are unknown to us. It is thus necessary to work with probabilities of probabilities. The universe is consistent and does not change around on us; it simply is important to learn which world we are in. It seems that this interpretation will be the most useful in ordinary experiments. This interpretation is thus judged to be a realistic model of the everyday world.

6 Repeated Trials

A *repeated trial* represents a situation in which the agent either chooses or is forced to execute the same type of action again with the "same setup" or the "same situation". An example is a gambler who must decide whether to spend his afternoon at Casino A, at Casino B, or at home. Casino A has a reputation for scrupulously fair odds. Casino B has either a reputation for cheating for the customer (to draw business) or for cheating against the customer, but the gambler forgets which. The casinos are too far apart to switch in midafternoon. The gambler must make a decision based on making *repeated trials* in the same situation.

It is significant that the agent is forced to "live with its decision". That is, the agent must make a decision as to which action to choose *before* a set of repeated trials starts. The agent suffers a penalty for changing situations. In such cases, it is insufficient to collapse the second-order probability down into a single world by integrating over marginals–it is necessary to maintain a full distribution of the different possible worlds, represented by a full second-order distribution.

The use of second-order probabilities provides a different value distribution than the use of equivalent first-order probabilities (Compare rule-set E in Figure 4 against rule-set A in

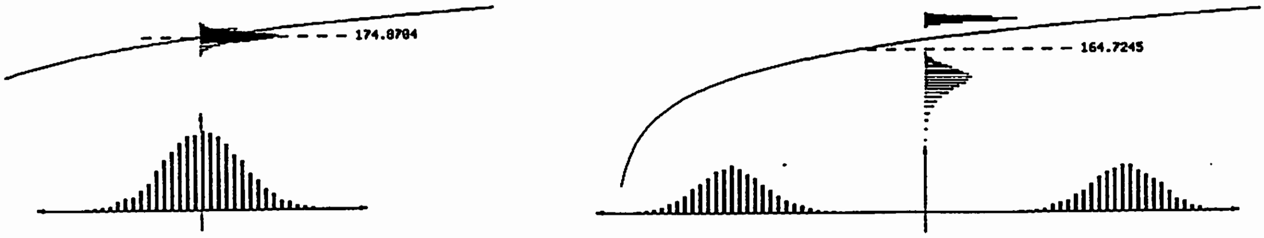

Figure 5: NDP and UDP Expected Utilities using a Nonlinear Curve

Figure 1). However, such results are only significant if the UDP interpretation is taken, and if more that one trial is performed. Note importantly that, assuming only a single trial, the NDP interpretation is equivalent to the UDP interpretation. It can be shown (by integrating possible worlds) that both are equivalent to using the equivalent first-order probability *under a single trial.* Thus, a deciding agent that is quite certain it will only be executing an uncertain action exactly once may use first-order probabilities–however, if this is not the case, second-order distributions must be used.

It is also significant to note that the question of whether an agent will make a single trial of an action or multiple trials of an action may not be completely determinable ahead of time. In such cases, the agent might think that it *might* have to repeat the trial, with the same unknown probability, at some unknown time in the future. Even this consideration of a hypothetical repetition could be enough to make an agent choose a situation with a particular degree of uncertainty over another situation–especially if the choice is irrevocable.

We note in passing that, under *fallible execution*, an agent may make repeated trials even though the agent is performing the action only once. The philosophy of fallible execution states roughly that "some things don't work right the first time"; this is significant to *intentional action theory.* If the agent *intends* to produce a particular effect, and *endeavors* to produce this effect by repeatedly executing an action (which nondeterministically results in this effect) until the effect is achieved, then the agent is in fact making repeated trials of the action–even though the effect is produced only once. (An example is trying to shoot a basketball through a hoop once, or trying to find the right parse of a sentence.) Such actions fall in the domain of this paper.

7 Nonlinear Utility Curves

The expected value of second-order probability runs, under either the NDP or the UDP interpretations, will always be the same as the expected value computed using first-order probabilities. Thus, as long as utility is a straight-line function of value, the expected utility will also be the same, whether NDP, UDP, or equivalent first-order probabilities are used.

However, in real life, utility is almost never a straight-line function of value. The actual shape of utility curves is normally concave [How70], and may in fact typically be logarithmic [Bar88, p. 300]. Under the NDP assumption, the value distribution and thus the utility distribution come out the same as cases using first-order probabilities. However, under the UDP assumption with repeated trials, the value distribution is different. Thus, with a nonlinear utility curve, the expected utility will differ as well. See Figure 5.

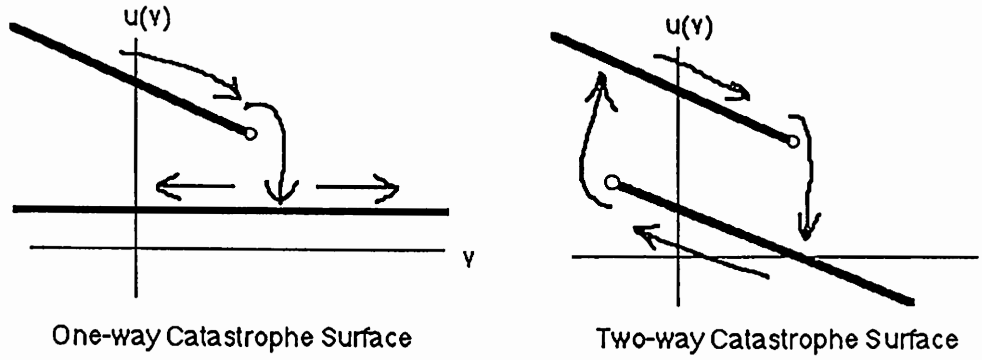

Figure 6: Catastrophe Surfaces

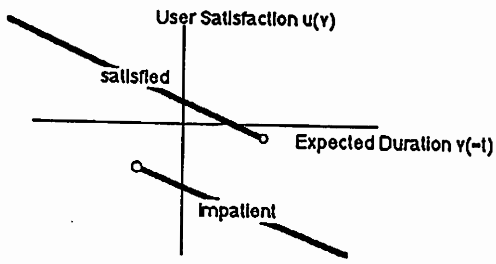

Figure 7: User Satisfaction as a Function of Expected Program Duration

8 Catastrophic Outcomes

One significant class of nonlinear curves is the *catastrophe surface* [Tho75] (see Figure 6).[5] The catastrophe surface represents a function that depends on the history of previous results. A series of results follows the upper surface until it goes past the "catastrophe" point, whereupon it follows the lower surface. "One-way" catastrophes are forced to stay on the lower surface; "two-way" catastrophes can jump back up to the upper surface. The catastrophe surface is important in representing emotions [Zee76]. Since utility is determined to a large extent by how happy or satisfied an organism is with a particular state of affairs, and happiness is an emotion, it is obvious that there is much to be gained by representing utility functions with catastrophe surfaces. Not only emotions, but also perception [SP83] and many physical phenomena can be represented by catastrophe surfaces [Tho75], which use concepts in common with chaos theory [Gle87]. Catastrophe functions depend on the history of the results; thus, repetition is again significant.

A situation that contains an outcome with a utility that can be modeled by a catastrophe surface is defined as a situation with a *catastrophic outcome.* A catastrophic outcome will typically involve an unrecoverable change of state. Examples are given in the next section.

9 The Uses of Catastrophic Outcome Theory

9.1 User Satisfaction

The theory predicts that user satisfaction with the repetitive results of a computer program should be capable of being modeled by a catastrophe surface, where the abscissa is the amoung of time expected to be taken by a single task, and the ordinate is the degree of the user's satisfaction. See Figure 7. For instance, suppose a user is interacting with an interpreting telephone, and listening to the amount of time the telephone system takes to translate each

[5]This paper will actually deal with two-dimensional cross-sections of higher-dimensional catastrophe surfaces. The distinction is unimportant for our purposes.

utterance. If the actual duration is approximately what the user expects, the user will have a certain amount of satisfaction with the system, and exist in the upper part of the catastrophe surface, labeled "satisfied". If the duration is significantly longer than the user expects, the user will reach a point where he or she suddenly loses patience, and satisfaction drops significantly ("impatient"). Even if the following utterances are then translated in a timely fashion, the user will still be dissatisfied with the system. However, if the system then interprets an utterance significantly faster than the user expects, the theory predicts that the user will then suddenly increase his opinion of the system again, and regain satisfaction. Psychological experiments must be performed to confirm these predictions, and empirically establish the parameters of the function.

We note that such a theory of user satisfaction is a necessary input to a scheduling system that has to decide whether to attempt a fast, rough interpretation, or a slow but polished one. It is dangerous to take too long; this theory permits defining the degree of risk.

9.2 Naive Physics

Catastrophic outcomes are significant in naive physics. For instance, a program simulating the results of cooling the solar cells on a space station will have to take catastrophic outcomes into effect. A utility function might represent the efficiency of the solar cells as a function of the actual amount of cooling probabilistically obtained from a particular amount of cooling effort. If the solar cells get too hot, performance degrades until suddenly the solar cells melt and efficiency drops dramatically. Thereafter, even if cooling is restored, the solar cells will still operate at the new low efficiency. This is a one-way catastrophe that does not allow jumping back up to the upper surface again.

9.3 Limited Non-negative Semi-renewable Resources

Catastrophe surfaces are especially useful in representing semi-renewable limited resource situations, where one has to trade a small amount of resource to uncertainly receive more of the resource. One example is an automatic telephone interpretation system, that has hesitation or delaying utterances (English: "uhhhh..."; Japanese: "ehhh touuuu") or can ask the user for more time if it is necessary to explore or polish a difficult translation. However, planning the delaying tactic costs time, and the tactic might not be successful–it might increase the impatience of the user. At some point the system runs out of time, and the utility of further computation suddenly drops significantly.

A second example is a Cadillac driving from gas station to gas station in the middle of the desert. The decision as to whether to use the air conditioner (which consumes fuel) or not depends a lot on how certainly the fuel consumption rate of the car is known. Running out of fuel on the highway stops the action and catastrophically results in highly negative utility.

Other examples of catastrophic outcomes include a gambler running out of money to gamble with; a plant running out of water; a pet running out of food; or a hospital life-support system temporarily running out of control resources.

10 Computing with Catastrophic Outcomes

Perhaps the best method for working with catastrophic outcomes is to perform multiple Monte Carlo simulations of the problem (as is shown in Section 11). It is possible, however, to solve for a closed-form expression of the value distribution for the top and bottom parts of the catastrophe surface, given the surface, the starting value, a second-order probability expression for incremental value, and the number of trials; this can then be used to find the utility of the problem. A full exposition of the mathematics is regrettably beyond the scope of this

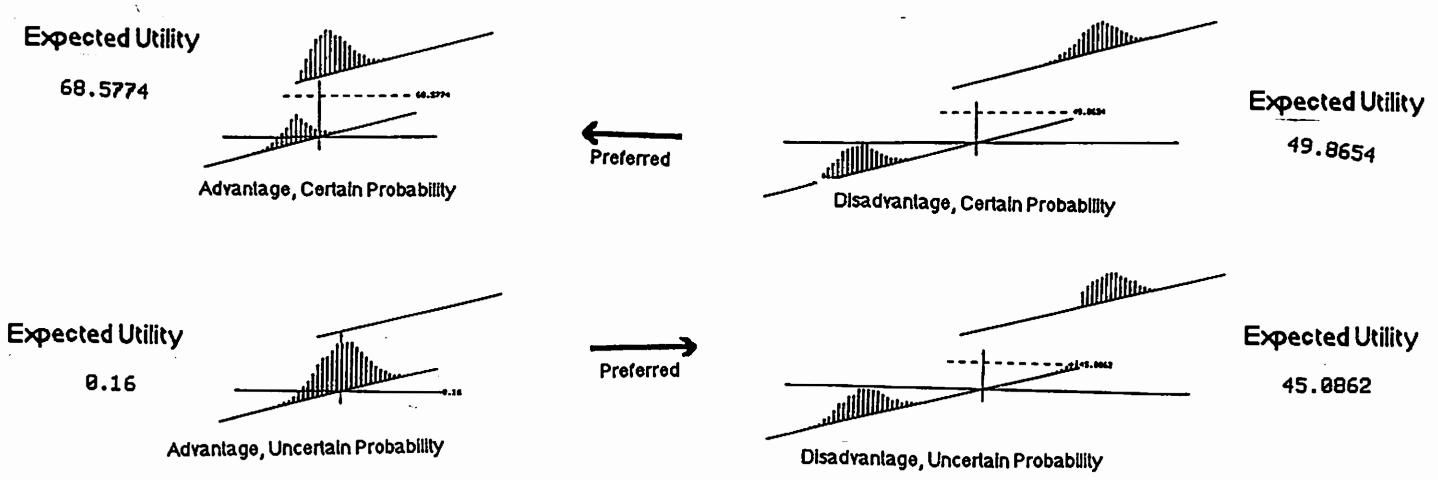

Figure 8: Uncertainty levels with Catastrophic Outcomes

paper. The method of solution involves defining a random walk with an absorbing state at the catastrophe point. See the author for further details.

11 The Preference for Uncertainty

The *degree of uncertainty* is defined as the variance of the second-order probability distribution for a particular event. When the degree of uncertainty is zero, the "actual probability" is completely certain (in the objective case) or is believed to be known with certainty (in the subjective case). Correspondingly, when the degree of uncertainty is high, the "actual probability" can vary widely. The degree of uncertainty for a repeated trial is correspondingly defined as the variance of the resulting value distribution.

When working with nonlinear utility curves, it is important to investigate situations with different degrees of uncertainty. Under repeated trials, different situations with various degrees of uncertainty can result in widely different expected utilities. These must be computed and compared against each other.

If the utility curve is *concave* (from the bottom), the acting agent is called *risk adverse* [KR76]. With a concave curve, a value distribution with a high degree of uncertainty will, in general, tend to have a lower expected utility than a value distribution with the same expected value and a lower degree of uncertainty. Conversely, if the utility curve is *convex* (from the bottom), the acting agent is called *risk prone* [KR76]. With a convex curve, a value distribution with the same expected value but a higher degree of uncertainty will, in general, tend to have a *higher* expected utility.

In the case of a catastrophic outcome, if the agent's initial value position is on the upper side of the catastrophe (at an *advantage*), repeated trials with a high degree of uncertainty in the second-order distribution will, in general, tend to send some of the value distribution over the edge and result in a lower expected utility than repeated trials with a lower degree of uncertainty but the same effective probability and expected value. Conversely, if the agent starts on the lower side of a two-way catastrophe (at a *disadvantage*), repeated trials with a high degree of uncertainty may result in some cases "jumping up" to the upper level, which would tend to increase the expected utility of the outcome when compared to a similar case with a lower degree of uncertainty. See Figure 8, which compares cases from rule-set A against cases from rule-set E. These interesting results indicate that agents in good situations will tend to be conservative, while agents in bad situations will tend to seek high-uncertainty actions—they have "nothing to lose".

It is important to note that these trends are heuristic tendencies only. The expected value plus the degree of uncertainty does *not* uniquely determine an expected utility [KR76, pp.135-

136], and specific cases may differ from these guidelines. It is thus important to use the whole distribution for computations, and not just parameters.

12 Acknowledgements

Peter Davis supplied the field problem formulation removed from Section 10. Wuyi Yue significantly clarified the mathematics used to obtain a closed-form solution for the same section. Philippe Quinio provided a useful critique on marginal probabilities.

13 Conclusion

This paper has shown that there are two philosophically consistent interpretations of the concept of a second-order probability: the Nondeterministic Probability (NDP) interpretation, and the Unknown Deterministic Probability (UDP) interpretation. The Nondeterministic Probability interpretation requires the universe to constantly change "actual probabilities" and thus invalidates the concept of a constant first-order probability; it is therefore judged unacceptable. The UDP interpretation allows the universe to have constant first-order probabilities which may however be unknown; it is judged to be realistic. Under UDP and repeated trials of experiments with nonlinear value/utility curves, second-order probability distributions result in significantly different expected utilities when compared against experiments made with equivalent first-order probabilities. An important subset of the class of nonlinear utility curves is the set of those represented by a catastrophe surface; these represent events from experiments having *catastrophic outcomes*. In such situations, advantaged agents will generally tend to choose actions with certain probabilities, while disadvantaged agents will generally tend to choose actions with highly uncertain probabilities. The results indicate that, unlike the results of other researchers, the full distribution of a second-order probability, and not simply its first-order equivalent probability, is required when making decisions maximizing expected utility.

References

[Bar88] Jonathan Baron. *Thinking and Deciding.* Cambridge University Press, Cambridge, Mass., 1988.

[Che85] Peter Cheeseman. In defense of probability. In *IJCAI'85: The Ninth International Joint Conference on Artificial Intelligence*, pages 1002–1009, Los Angeles, CA, 1985.

[Gai86] Haim Gaifman. A theory of higher order probabilities. In Joseph Y. Halpern, editor, *Reasoning About Knowledge*, pages 275–292. Morgan Kaufmann Publishers, Inc., Los Altos, CA, 1986.

[Gle87] James Gleick. *Chaos: Making a New Science.* Viking Penguin Inc., 1987.

[Gri84] John Gribbin. *In Search of Schrödinger's Cat.* Bantam Books, 1984.

[How70] Ronald A. Howard. Risk preference. In Ronald A. Howard and James E. Matheson, editors, *The Principles and Applications of Decision Analysis*, pages 629–663. Strategic Decisions Group, 1970.

[KR76] Ralph L. Keeney and Howard Raiffa. *Decisions with Multiple Objectives: Preferences and Value Tradeoffs.* John Wiley & Sons, New York, NY, 1976.

[Kyb89] H. E. Kyburg. Higher order probabilities. In Laveen N. Kanal, Tod S. Levitt, , and John F. Lemmer, editors, *Uncertainty in Artificial Intelligence 3*, pages 15–22. North-Holland, Amsterdam, The Netherlands, 1989.

[Pea88] Judea Pearl. *Probabilistic Reasoning in Intelligent Systems.* Morgan Kaufmann Publishers, Inc., Los Altos, CA, 1988.

[Rac89] Howard Rachlin. *Judgment, Decision, and Choice.* W. H. Freeman and Company, New York, NY, 1989.

[Rai68] Howard Raiffa. *Decision Analysis.* Addison-Wesley Publishing Co., Menlo Park, CA, 1968.

[SP83] I. N. Stewart and P. L. Peregoy. Catastrophe theory modeling in psychology. *Psychological Bulletin*, 94(2):336–362, 1983.

[Tho75] R. Thom. *Structural Stability and Morphogenesis.* Addison-Wesley Publishing Co., Menlo Park, CA, 1975. Translated by D. H. Fowler.

[Zee76] E. C. Zeeman. Catastrophe theory. *Scientific American*, 234(4):65–83, 1976.