Chance Discovery and Scenario Analysis

Peter McBURNEY

Department of Computer Science
University of Liverpool
Liverpool L69 7ZF
UK

Simon PARSONS

Center for Co-ordination Science
Sloan School of Management
Massachusetts Institute of Technology
Cambridge MA 02142
USA

`p.j.mcburney@csc.liv.ac.uk`

`sparsons@mit.edu`

Received 26 February 2002

Abstract

Scenario analysis is often used to identify possible chance events. However, no formal, computational theory yet exists for scenario analysis. In this paper, we commence development of such a theory by defining a scenario in an argumentation context, and by considering the question of when two scenarios are the same.

Keywords Argumentation, Chance Discovery, Ensembles, Forecasting, Scenarios.

§1 Introduction

The new discipline of chance discovery is concerned with the identification and management of rare, but significant, events, such as potential risks or opportunities, in some domain or system of application. [17, 10] In this paper, we focus attention on chance identification rather than on chance management. One method commonly used to identify chance events is by considering possible future values of some parameters of interest, allowing them to vary according to the actions of causal mechanisms assumed to operate. This approach yields a possible future state of the world, or a trajectory of such states, called a *scenario*; comparing two or more these is called scenario analysis.[19]

Scenario analysis has been applied extensively in business and in public policy, as well as in scientific domains, e.g., [8]. Perhaps the most important recent application of scenario analysis has been by the Intergovernmental Panel on Climate Change (IPCC), the United Nations agency tasked with assessing the current state and possible futures of the world's ecosystem, and attempting to devise appropriate regulatory policies to prevent or respond to global warming. [3] In this domain, scenario analysis has been used for scientific modeling and prediction, for the modeling of socio-economic variables, and for the comparison of environmental and other regulatory policies. Despite their widespread use, however, there appears to be no formal theory of scenarios. Without a formal theory of scenarios, many questions remain without rigorous answers, e.g., How should scenarios be constituted? How many scenarios should be considered? How should individual scenarios be analysed? How should aggregation of outcomes across scenarios be undertaken? How should the likelihood of occurrence of different scenarios be represented? How should such likelihoods be aggregated across scenarios? What is the relationship between scenarios and the domain of application? And without a formal theory of scenarios, there can be no computational theory, thus limiting the potential applicability of scenario analysis in intelligent systems.

The long-term aim of this research is a rigorous, formal, computational theory of scenarios. This paper takes one step towards this aim, by considering one type of scenario — those based on dialectical argumentation — applied to one problem — that of chance discovery. In earlier work [11], we showed how dialectical argumentation may be applied to the identification of chance events, and proposed a protocol for distributed communications between agents jointly engaged in chance discovery. Argumentation methods are appropriate when relevant knowledge is distributed between autonomous agents, or when the interests or values of such agents may diverge. In these circumstances, methods based on fusion of different knowledge bases or analysis of all the data held by the participating agents may be inappropriate, as participants may not wish to share all their information with each other. In Section 2, we review a basic model for dialectical argumentation about uncertain domains. Section 3 then defines our notion of scenario, while Section 4 considers the question of when two scenarios may be considered the same or not. We present a decision rule for deciding if two scenarios are distinct, based only on their initial premises and inference rules, and on the estimated likelihoods of new new information being presented in either scenario. Section 5 concludes with a brief discussion of applications to chance discovery, and discussion of related research.

§2 Dialectical Argumentation

In this section we briefly summarize the Agora framework for the qualitative representation of uncertainty which we presented in earlier work.[13] In this framework, arguments for and against claims are articulated by participants in an electronic space, called an *Agora*, with claims expressed as formulae in a propositional language. By means of defined locutions, participants in the Agora can variously posit, assert, contest, justify, rebut, undercut, qualify and retract claims, just as happens in real discourse. For example, a debate participant $\mathcal{P}_i$ could demonstrate her argument $\mathcal{A}(\rightarrow \theta)$ supporting a claim θ, an argument to which she was committed with strength D, by means of the locution:

$$\textbf{show\_arg(}\mathcal{P}_i : \mathcal{A}(\rightarrow \theta, D)\textbf{)}.$$

The rules governing the use of each permitted locution are expressed in terms of a formal dialogue-game between the participants. [6, 14] We assume that the Agora participants begin a debate with a set of agreed facts, or assumptions, and an agreed set of inference rules. Because we want to model many forms of reasoning, these rules need not be deductive and may themselves, in our Agora formulation, be the subject of argument.

We demonstrated the use of this framework for the representation of uncertainty by defining a set of uncertainty labels, which are assigned to claims on the basis of the arguments presented for and against them in the Agora. Essentially, one could say that claims have more credibility (and hence less uncertainty) the fewer and the weaker are the arguments against them. While any set of labels could be so defined, we drew on earlier work in argumentation [9] and used the set: {*Accepted, Probable, Plausible, Supported, Open*}, with the elements listed in decreasing order of certainty. For example, a claim was regarded as *Probable* at a particular time if at least one consistent argument had been presented for it in the Agora by that time, but no arguments for its negation (rebuttals) nor for the negation of any of its assumptions (undercuts) had been presented by then. We defined a claim as *well-defended* if there was an argument for it and any rebuttals or undercuts were themselves subject to counter-rebuttals or to undercuts. *Accepted* claims were defined as those which are well-defended.

We then defined the truth-valuation of a claim θ at time t, denoted $v_t(\theta)$, to be 1 if θ had the label *Accepted* at this time, otherwise it was 0. Such a valuation summarizes the knowledge of the community of participants at the particular time, since it incorporates, via the definitions of the labels, all the arguments for and against θ articulated to that time. Consequently, assessing the truth-status of a claim at a particular time can be viewed as taking a *snapshot* of an Agora debate. Of course, because these definitions are time-dependent, and arguments may be articulated in the Agora at any time, such an assignment of uncertainty labels and truth valuation must be defeasible. Claims accepted at one time may be overturned at another, in the light of new information learnt or arguments presented subsequently.

In using the Agora framework to represent uncertainty, attention will focus on the truth valuation function over the long-run.[*1] The sequence $(v_t(\theta) \mid t = 1, 2, \ldots)$ may or may not converge as $t \to \infty$. Suppose that it does converge, and denote its limit value by $v_\infty(\theta)$. What will the value of a snapshot taken at time t, namely $v_t(\theta)$, tell us about $v_\infty(\theta)$? Of course, any finite snapshot risks being overtaken by subsequent information or arguments, we cannot infer with complete accuracy from the finite snapshot to the infinite value. However, we have shown 13) that, under certain conditions, we can place a bound on the likelihood that such an inference is in error. The conditions essentially require that: (a) the snapshot is taken at a time after commencement sufficient for all the arguments using the initial information to have been presented, and (b) there is a bound on the probability that new information arises following the snapshot. This result is proved as Proposition 7 of 13), which we reproduce here. For this, we first need some definitions.

Definition 2.1
We write LE_θ for the statement: "*The function $v_t(\theta)$ converges to a finite limit as $t \to \infty$.*" We also write $\mathcal{X}_{t,\theta}$ for the statement: "*New information relevant to θ becomes known to an Agora participant after time t.*"

In general, at any time s, we do not know whether new evidence will become available to Agora participants at a later time t or not. Consequently, the variables $\mathcal{X}_{t,\theta}$, for t not in the past, represent uncertain events. Also uncertain for the same reason are statements concerning the future values of $v_t(\theta)$ for any θ. Because these events are uncertain, we assume the existence of a probability function over them, i.e., a real-valued measure function mapping such statements to $[0, 1]$ which satisfies the axioms of probability.

[*1] Strictly, we are assuming throughout that time in the Agora is discrete, and can be represented by a countably infinite set.

Definition 2.2

$Pr(.)$ is a probability function defined over statements of the form $\mathcal{X}_{t,\theta}$ and statements concerning the values of $v_t(\theta)$, for any formula θ.

Theorem 2.1

[Proposition 7 of [13]] Let θ be a formula and suppose that all arguments pertaining to θ and using the information available at commencement are articulated by participants by some time $s > 0$. Suppose further that $v_{t_m}(\theta) = 1$ for some $t_m \geq s$. Also, assume that $Pr(\mathcal{X}_{t_m,\theta}) \leq \epsilon$, for some $\epsilon \in [0,1]$. Then the following inequalities hold:

$$Pr(LE_\theta \ and \ v_\infty(\theta) \ = 1 \mid v_{t_m}(\theta) = 1) \geq 1 - \epsilon.$$

$$Pr(LE_\theta \ and \ v_\infty(\theta) \ = 0 \mid v_{t_m}(\theta) = 1) \leq \epsilon. \qquad \square$$

Like the standard (Neyman-Pearson) procedures for statistical hypothesis testing, this proposition provides us with some confidence in our use of finite snapshots to make inferences about the long-run truth-valuation function for a debate. While such inference is not deductively valid, at least its likelihood of error may be bounded.[*2] In the sections below, we will be comparing the results of debates in more than one Agora. We therefore assume that we have a single probability function Pr defined across all the relevant statements. We will also index symbols with superscripts (^{1},^{2}, etc) to denote the Agora to which they refer. We next define the concept of Scenario.

§3 Scenarios

The framework we have just outlined provides a means to represent the diverse arguments that may arise from a given set of assumptions, and using a given set of inference rules (deductive or otherwise). If we were to start with a different set of assumptions, and/or permit the use of a different set of inference rules, the arguments presented in the Agora may well be different. As a result, the uncertainty labels and truth values assigned to formulae may well also be different, both when taken at finite snapshots and in the limit. We define a scenario as follows:

Definition 3.1

A **Scenario** for a given domain consists of a set of assumptions and a set of inference rules, with which participants are equipped at the commencement of an Agora debate over formulae in that domain. We denote scenarios for a given domain by $\mathcal{S}^1, \mathcal{S}^2, \ldots$,

---

[*2] One may object that we can never know the value of ϵ. While this is true, participants in a debate are often quite willing to provide subjective estimates for such probability bounds. Scientists, for example, will often estimate the chance that new information will arise which overturns an established theory.

etc. For each scenario, $\mathcal{S}^i$, an Agora debate undertaken with the assumptions and inference rules of that scenario, is said to be the **associated Agora**, denoted $\mathcal{A}^i$. We assume only one Agora debate is conducted in association with any scenario.

In this paper, we will be assuming that all scenarios, and all the resulting Agora debates, relate to the same application domain. For this domain, suppose we are interested in a particular proposition θ. We imagine we have a number of scenarios in parallel, each with a different set of starting assumptions and possibly also different inference mechanisms. We now allow the associated Agora debates to proceed up to a certain time t, when we take a finite snapshot of each debate. It would be expected that the truth status of θ would be different under different scenarios. Not only are the assumptions and inference mechanisms different, but not all arguments may have been presented to each Agora debate at the time of the snapshot. For chance discovery of some claim θ (for example, a possible risk) we are interested in whether there are any scenarios in which $v_t(\theta)$ has been assigned the value 1. If so, there is a scenario in which the claim θ is well-defended. An immediate question would be how many such scenarios are there? To answer this question accurately, we need to ensure that each distinct scenario is only counted once, i.e., that no "double-counting" of identical scenarios takes place. In other words, we need a rule to determine whether two scenarios are the same or not. Such a decision rule is proposed in the next section.

An early use of scenario analysis was in nineteenth-century statistical mechanics, where physicists studied the extent to which properties of a physical system, such as its entropy at a given time, depended on the initial state of the system. Boltzmann [2] explored this question by comparing the given system to a collection of alternative, imaginary systems, each having different initial conditions; doing this, enabled an assessment of the extent to which the property of interest was independent of the initial system state. Gibbs [5] termed the collection of imaginary systems an *ensemble*, and we adopt this terminology also.

Definition 3.2
An **Ensemble** $\mathcal{S}$ is a finite collection of distinct Scenarios $\{\mathcal{S}^1, \ldots, \mathcal{S}^m\}$ relating to a common domain.

§4 Comparing Scenarios

4.1 Comparing two long-run debates

When are two scenarios the same? Obviously, we may consider them to be the same when they have identical sets of assumptions and inference rules. But two scenarios identical in this fashion may result in very different Agora debates, as different arguments may be presented in each, or the same arguments may be presented at different times. It is not clear, therefore, that identical scenarios will lead to identical assignments of truth-labels, even over the long-run; we show that, under certain conditions, they will do so. Throughout this section $\mathcal{S}^1$ and $\mathcal{S}^2$ will be two scenarios of interest, and $\mathcal{A}^1$ and $\mathcal{A}^2$ their associated Agora debates.

Theorem 4.1
Let θ be a claim. Suppose that $\mathcal{S}^1$ and $\mathcal{S}^2$ are identical scenarios, i.e., they have identical sets of assumptions and identical sets of inference rules. Suppose that in the corresponding Agora debates, $\mathcal{A}^1$ and $\mathcal{A}^2$, all possible arguments based on the initial assumptions and using the inference rules are eventually articulated in the Agora. Suppose further that no new information is presented to either debate following commencement. Then, the long-run truth status of θ in each debate is the same.

Proof **(Outline)** Given the premises, the only way the two debates will potentially differ will be in the order that arguments are articulated in the Agora. But if all arguments are eventually articulated, then after some finite time no further arguments will be presented in either debate. It therefore follows that the long-run truth status of a claim does not depend upon the order of presentation of the arguments for and against it. ■

If we relax the assumption that no new information arrives in either debate our conclusion acquires a probabilistic qualification. While this does not guarantee that two identical scenarios always lead to identical long-run truth assignments, it does bound the likelihood that this is not the case.

Theorem 4.2
Let θ, $\mathcal{S}^1$ and $\mathcal{S}^2$ be as before. Suppose there exist upper bounds $\epsilon^i \in [0,1]$ for the probability that new information arrives after commencement in debate i, i.e., that $Pr(\mathcal{X}^i_{0,\theta}) \leq \epsilon^i$, for $i = 1, 2$. Then we have:

$$Pr(v^1_\infty(\theta) = v^2_\infty(\theta)) \geq 1 - \epsilon^1 - \epsilon^2.$$

Proof **(Outline)** By the previous result, the two long-run assignments of truth to θ are only different if one or other debate receives new information. The probability that this occurs is less than or equal to the sum of the probabilities that either debate receives new information less the probability that they both do. This latter event has probability greater than or equal to zero, and the inequality follows by algebraic manipulation. ■

4.2 A decision rule for scenario comparison

We now provide a decision rule for determining if two scenarios $\mathcal{S}^1$ and $\mathcal{S}^2$ are the same. This decision rule classifies scenarios into two classes, labeled *distinct* and *non-distinct*. The rule proposed for determination of distinctness of scenarios uses two criteria (in order of application): (a) whether or not the two scenarios have identical assumptions and inference rules; (b) in the case where they do, whether or not either scenario is judged to have a high probability of receiving new information.

Case 1A: $\mathcal{S}^1 = \mathcal{S}^2$ **and** $Pr(\mathcal{X}^1_{0,\theta}), Pr(\mathcal{X}^1_{0,\theta})$ **both small.** In this case, the likelihood of new information arising in either scenario is small and Theorem 4.2 allows us to infer that $v^1_\infty(\theta) = v^2_\infty(\theta)$ with high probability. Conclude that the two scenarios are *non-distinct*.

Case 1B: $\mathcal{S}^1 = \mathcal{S}^2$ **and one or both of** $Pr(\mathcal{X}^1_{0,\theta}), Pr(\mathcal{X}^1_{0,\theta})$ **large.** In this case, the likelihood of new information arising in at least one scenario is large, and thus, Theorem 4.2, it is unlikely that $v^1_\infty(\theta) = v^2_\infty(\theta)$. Conclude that the two scenarios are *distinct*.

Case 2: $\mathcal{S}^1 \neq \mathcal{S}^2$**.** Conclude that the two scenarios are *distinct*.

In the first two cases (Cases 1A and 1B), where the underlying assumptions and inference rules are the same in the two scenarios, Theorem 4.1 says that the long run truth assignments for θ in the corresponding Agora debates, if they exist, will be identical, provided no new information is presented in either Agora debate following commencement. If new information is presented, then Theorem 4.2 provides a bound for the probability that the long-run truth assignments are the same, in terms of the probabilities of new information being received. In the case (Case 1A) when these probabilities are believed to be small, the two long-run truth assignments are most likely identical, and we can classify the two scenarios as being the same. In the other case (Case 1B), where one or both probabilities are large, we classify the two scenarios as not the same. In the final case (Case 2), where the two scenarios have different premises and/or inference rules, we also classify them as distinct. It may be, of course, that two such distinct scenarios may result in the same arguments being presented in both scenarios after some finite time.

Note that, although under Cases 1A and 1B we are making inferences about the long run truth assignments, $v^1_\infty(\theta)$ and $v^2_\infty(\theta)$, these inferences are based only on the premises and inference rules used and assessments of the probability of new information being received after commencement of the associated Agora debates. These inferences, and hence this classification, do not depend on the progress or status of the debates themselves. In other words, our classification of scenarios is not based on the output of the debates conducted under the scenarios.

§5 Discussion

In many domains chance events are identified by exploring possible scenarios which are consistent descriptions of possible futures in some domain. Despite their widespread use, there is as yet no formal, computational theory of scenarios and scenario analysis. In this paper, we have commenced work on such a theory for scenarios which describe debates over uncertain propositions, for example future states of some system. In our formalism a scenario is a debate in some domain with pre-specified and agreed premises and inference rules. The search for chance events becomes a matter of varying these pre-specifications and allowing a number of parallel debates to operate simultaneously. If a claim θ is assigned the truth-status of *true* in one of these debates, then there is a possible future world state in which θ is realized. To assess the likelihood that θ is realized in the actual world, we need to consider all the scenarios in which the proposition θ is assigned the value *true,* and determine their combined likelihood of occurrence. If this likelihood is small and θ refers to an important event, then we have identified a chance event.

The work presented here is novel. Although recent research in business strategy has considered the use of multiple scenarios to identify chance events, e.g. 16), that work has not been formalized. The closest research to ours is the *Ents* model of belief of Paris and Vencovska. 18) In that model, an agent's belief in a claim is determined by imagining possible worlds (analogous to our scenarios) in which the claim is decided, either true or false, and then setting the belief in the claim equal to the proportion of possible worlds in which it is true. Our scenarios may be viewed as argumentation analogs of these possible worlds, with the advantage that our argumentation system provides an operational mechanism for assigning truth-status labels to propositions, 13) a mechanism absent from the *Ents* model. Both these approaches, as was mentioned earlier, are conceptually similar to the Ensemble theory of Boltzmann 2) and Gibbs 5) in statistical mechanics.

In this paper, for simplicity, we have only considered the *likelihood* of events, and not their *significance*. Chance discovery is the identification of rare but important events. However, our framework could be readily modified to accommodate significance, either by the explicit incorporation of values in the search for rare events, as was done in our previous work 11), or by the prioritization of arguments according to preference-orderings, such as in 1). Moreover, in this paper we have only considered the identification of chance events, and not their management. We intend to pursue both these issues in our future work applying argumentation approaches to chance discovery and management, drawing on our earlier work developing argumentation formalisms for decision-making, such as 4, 7, 12). Finally, there could be connections between the approaches we have outlined here and the vast literature on default logics in non-monotonic reasoning, which may be interesting to explore.[*3]

References

1) L. Amgoud and S. Parsons. Agent dialogues with conflicting preferences. In J-J. Meyer and M. Tambe, editors, *Pre-Proceedings of the Eighth International Workshop on Agent Theories, Architectures, and Languages (ATAL 2001)*, pages 1–14, Seattle, WA, USA, 2001.

2) L. Boltzmann. Weitere studien über das wärmegleichgewicht unter Gasmolekülen. *Wissenschaftliche Abhandlungen*, 1:316–402, 1872.

3) T. R. Carter, E. L. La Rovere, R. N. Jones, R. Leemans, L. O. Mearns, N. Nakicenovic, A. B. Pittock, S. M. Semenov, and J. Skea. Chapter 3: Developing and Applying Scenarios. Pages 145–190 of 15).

4) J. Fox and S. Parsons. Arguing about beliefs and actions. In A. Hunter and S. Parsons, editors, *Applications of Uncertainty Formalisms*, LNAI 1455, pages 266–302. Springer, Berlin, Germany, 1998.

5) J. W. Gibbs. *Elementary Principles in Statistical Mechanics, Developed with Especial Reference to the Rational Foundation of Thermodynamics*. Yale University Press, New Haven, CT, USA, 1902. (Reprinted by Dover, New York, 1962.).

6) C. L. Hamblin. *Fallacies*. Methuen, London, UK, 1970.

7) D. Hitchcock, P. McBurney, and S. Parsons. A framework for deliberation dialogues. In H. V. Hansen, C. W. Tindale, J. A. Blair, and R. H. Johnson, editors, *Proceedings of the Fourth Biennial Conference of the Ontario Society for the Study of Argumentation (OSSA 2001)*, Windsor, Ontario, Canada, 2001.

8) J. W. Ironside. nvCJD: exploring the limits of our understanding. *The Biologist*, 46(4):172–176, 1999.

9) P. Krause, S. Ambler, M. Elvang-Gøransson, and J. Fox. A logic of argumentation for reasoning under uncertainty. *Computational Intelligence*, 11 (1):113–131, 1995.

10) P. McBurney. First international workshop on chance discovery. *Knowledge Engineering Review*, 16(2):215–218, 2001.

11) P. McBurney and S. Parsons. Chance discovery using dialectical argumentation. In T. Terano, T. Nishida, A. Namatame, S. Tsumoto, Y. Ohsawa, and T. Washio, editors,

---

[*3] We are grateful to the anonymous referees for their comments.

New Frontiers in Artificial Intelligence: Joint JSAI 2001 Workshop Post Proceedings, LNAI 2253, pages 414–424. Springer, Berlin, Germany, 2001.

12) P. McBurney and S. Parsons. Intelligent systems to support deliberative democracy in environmental regulation. *Information and Communications Technology Law*, 10(1):33–43, 2001.

13) P. McBurney and S. Parsons. Representing epistemic uncertainty by means of dialectical argumentation. *Annals of Mathematics and AI*, 32(1–4):125–169, 2001.

14) P. McBurney and S. Parsons. Games that agents play: A formal framework for dialogues between autonomous agents. *Journal of Logic, Language and Information*, 11(3):315–334, 2002.

15) J. McCarthy, O. F. Canziani, N. A. Leary, D. J. Dokken, and K. S. White, editors. *Climate Change 2001: Impacts, Adaptation and Vulnerability. Intergovernmental Panel on Climate Change*. Cambridge University Press, Cambridge, UK, 2001.

16) P. van Notten. Foresight in the face of scenario diversity. In H. Tsoukas and J. Shepherd, editors, *International Conference on Probing the Future: Developing Foresight in the Knowledge Economy*, Glasgow, Scotland, UK, July 2002. Graduate School of Business, University of Strathclyde.

17) Y. Ohsawa. Chance discoveries for making decisions in complex real world. *New Generation Computing*, 20(2), 2002.

18) J. Paris and A. Vencovská. A model of belief. *Artificial Intelligence*, 64:197–241, 1993.

19) P. Schwartz. *The Art of the Long View*. Doubleday, New York, NY, USA, 1991.