Edited by

Anton Benz, Gerhard Jäger and Robert van Rooij

Game Theory and Pragmatics

Palgrave Studies in Pragmatics, Language and Cognition

Game Theory and Pragmatics

Edited by

Anton Benz

Gerhard Jäger

Robert van Rooij

palgrave
macmillan

First published 2006 by
PALGRAVE MACMILLAN
Houndmills, Basingstoke, Hampshire RG21 6XS and
175 Fifth Avenue, New York, N. Y. 10010
Companies and representatives throughout the world

PALGRAVE MACMILLAN is the global academic imprint of the Palgrave Macmillan division of St. Martin's Press, LLC and of Palgrave Macmillan Ltd. Macmillan® is a registered trademark in the United States, United Kingdom and other countries. Palgrave is a registered trademark in the European Union and other countries.

ISBN-13: 978-1-4039-4572-3
ISBN-10: 1-4039-4572-1

This book is printed on paper suitable for recycling and made from fully managed and sustained forest sources.

A catalogue record for this book is available from the British Library.

Library of Congress Cataloging-in-Publication Data
Game theory and pragmatics / edited by Anton Benz, Gerhard Jäger, Robert van Rooij.
p. cm. — (Palgrave studies in pragmatics, language and cognition)
Includes bibliographical references and index.
ISBN 1-4039-4572-1 (cloth)
1. Mathematical linguistics. 2. Game theory. 3. Pragmatics. 4. Language and logic. I. Benz, Anton, 1965– II. Jäger, Gerhard, 1967– III. Van Rooij, Robert, 1966– IV. Series.

P138.G36 2005
410'.1'51—dc22

2005044356

10 9 8 7 6 5 4 3 2 1
15 14 13 12 11 10 09 08 07 06

Transferred to Digital Printing 2006

Game Theory and Pragmatics

Palgrave Studies in Pragmatics, Language and Cognition

Series Editors: **Noël Burton-Roberts** and **Robyn Carston**

Series Advisors: **Kent Bach, Anne Bezuidenhout, Richard Breheny, Sam Glucksberg, Francesca Happé, François Recanati, Deirdre Wilson**

Palgrave Studies in Pragmatics, Language and Cognition is a new series of high quality research monographs and edited collections of essays focusing on the human pragmatic capacity and its interaction with natural language semantics and other faculties of mind. A central interest is the interface of pragmatics with the linguistic system(s), with the 'theory of mind' capacity and with other mental reasoning and general problem-solving capacities. Work of a social or cultural anthropological kind will be included if firmly embedded in a cognitive framework. Given the interdisciplinarity of the focal issues, relevant research will come from linguistics, philosophy of language, theoretical and experimental pragmatics, psychology and child development. The series will aim to reflect all kinds of research in the relevant fields – conceptual, analytical and experimental.

Titles include:

Anton Benz, Gerhard Jäger and Robert van Rooij (*editors*)
GAME THEORY AND PRAGMATICS

Reinhard Blutner and Henk Zeevat (*editors*)
OPTIMALITY THEORY AND PRAGMATICS

Corinne Iten
LINGUISTIC MEANING, TRUTH CONDITIONS AND RELEVANCE
The Case of Concessives

Ira Noveck and Dan Sperber (*editors*)
EXPERIMENTAL PRAGMATICS

Forthcoming titles:

María J. Frápolli (*editor*)
SAYING, MEANING AND REFERRING
Essays on François Recanati's Philosophy of Language

Christoph Unger
ON THE COGNITIVE ROLE OF GENRE

Palgrave Studies in Pragmatics, Language and Cognition Series
Series Standing Order ISBN 0–333–99010–2
(*outside North America only*)

You can receive future titles in this series as they are published by placing a standing order. Please contact your bookseller or, in case of difficulty, write to us at the address below with your name and address, the title of the series and the ISBN quoted above.

Customer Services Department, Macmillan Distribution Ltd, Houndmills, Basingstoke, Hampshire RG21 6XS, England

Acknowledgements

The roots of this book lie in a small but very successful workshop on *Games and Decisions in Pragmatics* held at the Centre for General Linguistics (ZAS) in Berlin in October 2003. This event was organized by Anton Benz as part of the ZAS project on Bidirectional Optimality Theory, which was then directed by Gerhard Jäger and Manfred Krifka. We have to thank the members of this project who helped with reviewing and organising this workshop: Jason Mattausch, Yukiko Morimoto, and Tue Trinh. Thanks go also to Klaus Robering, who helped with the reviewing, and to the director of the ZAS, Manfred Krifka, for the encouragement and general support we received. Jason was also very helpful in the final stage of the book project when he corrected the English of the introductory chapter.

A crucial role in making this book project happen was played by Richard Breheny. After the first Call for Papers, we received e-mails from Richard encouraging us to use the conference contributions as a basis for a book on this newly developing field of pragmatics. He suggested contacting Palgrave Macmillan and the editors of the series *"Palgrave Studies in Pragmatics, Language and Cognition"*. He receives our special thanks. We would also like to thank the series editors Robyn Carston and Noël Burton-Roberts for their encouragement for and confidence in our project and to Jill Lake and Melanie Blair for their advice.

Last but not least, special thanks go to Mirco Hilbert, who did a terrific job in typesetting the book and preparing the index.

Anton Benz, Gerhard Jäger and Robert van Rooij

Notes on the Contributors

Nicholas Allott is a Ph.D. student at University College London, working on relevance and rationality in communication.

Nicholas Asher is a Professor of Philosophy at the University of Texas at Austin. His interests include natural language semantics, pragmatics, and philosophy of language. He is author of *Reference to Abstract Objects in Discourse* and, together with Alex Lascarides, of *The Logics of Conversation*.

Anton Benz is Assistant Professor of Humanistic Information Science at the University of Southern Denmark, Kolding. His main research interests lie in the field of formal pragmatics and dialogue modelling. He received his Ph.D. in philosophy, logic, and theory of science at Ludwig Maximilians University Munich.

Jacob Glazer is a Professor of Economics at the Faculty of Management, Tel Aviv University and the Department of Economics, Boston University. He received his Ph.D. from Northwestern University in 1986 His research areas include economic theory and health economics.

Kris de Jaegher is a postdoctoral Research Fellow of the Fund for Scientific Research, Flanders (Belgium), stationed at the Vrije Universiteit Brussel (Belgium). He teaches economics at the Roosevelt Academy in Middelburg (Netherlands). He received his Ph.D. at the Department of Economics of the Vrije Universiteit Brussel in 2002. His research interests lie in the fields of microeconomics and game theory. His research in the field of game theory includes applications in economics, biological signalling theory, and the philosophy of language.

Gerhard Jäger is Professor of Linguistics at the University of Bielefeld in Germany. He received his Ph.D. from Humboldt University in Berlin. He has published on a variety of topics relating to semantics, pragmatics and the mathematics of language. His current focus of interest is the application of evolutionary models to linguistics.

Tom Lenaerts works as a postdoctoral Researcher at IRIDIA, l'Institut de Recherches Interdisciplinaires et de Développements en Intelligence Artificielle, at the Université Libre de Bruxelles and guest professor at the computer science department situated at the Vrije Universiteit Brussel. He received his Ph.D. from the latter institution in May 2003. His research interests include game dynamics, evolutionary computation, cultural evolution, complex networks and the biochemical origin of life.

Prashant Parikh is a Senior Research Scholar at the Institute for Research in Cognitive Science at the University of Pennsylvania. He pioneered the application of game theory to semantics/pragmatics in the mid-nineteen-eighties and his interests include all aspects of the connections between game theory and language.

Robert van Rooij is a senior staff member of the Institute for Logic, Language and Computation (ILLC), stationed at the Department of Philosophy of the University of Amsterdam. He received his Ph.D. at the Institut für Maschinelle Sprachverarbeitung at the University of Stuttgart in 1997. His working areas include formal semantics and pragmatics, philosophy of language, and the evolution of language.

Ariel Rubinstein is a Professor of Economics in Tel Aviv University and New York University. He received his Ph.D. at the Hebrew University in 1979. His major working areas are economic theory, game theory and models of bounded rationality.

Merlijn Sevenster is currently employed as a Ph.D. student at the Institute for Logic, Language and Information in Amsterdam, working on a project on imperfect information games. His approach combines a linguistic, computational, and logical viewpoint.

Robert Stalnaker is the Laurance S. Rockefeller Professor of Philosophy at the Massachusetts Institute of Technology. He is the author of two recent collections of papers, *Context and Content* and *Ways a World Might Be*, both published by Oxford University Press. His areas of interest include metaphysics, the philosophy of language, and the foundations of game theory.

Bart de Vylder is a researcher at the Artificial Intelligence Laboratory at the Free University of Brussels. His research comprises robotics, computer vision and the modelling of language evolution.

Madison Williams is a graduate student who is finishing his Ph.D. at the University of Texas at Austin. He is writing his thesis on rational inference and decision making.

Contents

Acknowledgements vi

Notes on the Contributors vii

1 An Introduction to Game Theory for Linguists 1
Anton Benz, Gerhard Jäger and Robert van Rooij

2 Saying and Meaning, Cheap Talk and Credibility 83
Robert Stalnaker

3 Pragmatics and Games of Partial Information 101
Prashant Parikh

4 Game Theory and Communication 123
Nicholas Allott

5 Different Faces of Risky Speech 153
Robert van Rooij and Merlijn Sevenster

6 Pragmatic Reasoning, Defaults and Discourse Structure 176
Nicholas Asher and Madison Williams

7 Utility and Relevance of Answers 195
Anton Benz

8 Game-Theoretic Grounding 220
Kris de Jaegher

9 A Game Theoretic Approach to the Pragmatics of Debate 248
Jacob Glazer and Ariel Rubinstein

10 On the Evolutionary Dynamics of Meaning-Word Associations 263
Tom Lenaerts and Bart de Vylder

Author Index 285

Subject Index 287

1
An Introduction to Game Theory for Linguists

Anton Benz, Gerhard Jäger and Robert van Rooij

1 Classical game theory

In a very general sense we can say that we play a game together with other people whenever we have to decide between several actions such that the decision depends on the choice of actions by others and on our preferences over the ultimate results. Obvious examples are card games, chess, or soccer. If I am to play a card to a trick, then it depends on the cards played by my playing partners whether or not I win the trick. Whether my move in chess leads to a win usually depends on the subsequent moves of my opponent. Whether I should pass the ball to this or that team member depends not in the least on my expectations about whether or not he will pass it on to a player in an even more favourable position. Whether or not my utterance is successful depends on how it is taken up by its addressee and the overall purpose of the current conversation. This provides the basis for applications of game theory in pragmatics.

Game theory has a prescriptive and a descriptive aspect. It can tell us how we should behave in a game in order to produce optimal results, or it can be seen as a theory that describes how agents actually behave in a game. In this book, the latter interpretation of game theory is of interest. The authors of this volume will explore game theory as a framework for describing the use of language.

1.1 Decisions

At the heart of every game theoretic problem there lies a decision problem: one or more players have to choose between several actions. Their choice is governed by their preferences over expected outcomes. If someone is offered a cherry and a strawberry but can only take one of them, then if he prefers the strawberry over the cherry, he will take the strawberry. This is not a prescription. It is an explication of the semantics of the word *preference*. If I can choose between actions a_1 and a_2, and prefer the outcome s_1 of a_1 over s_2 of a_2, then it is the very meaning of the word *preference* that I choose action a_1. In general, one can distinguish between decision making under *certainty*, *risk* and *uncertainty*. A decision is made under *certainty* if the decision maker knows for each action which, outcome it will lead to. The cherry and strawberry example is such a case. A decision is made under *risk* if each action leads to a set of possible outcomes, where each outcome occurs with a certain probability. The decision maker knows these probabilities, or behaves as if he knew them. A decision is made under *uncertainty* if no probabilities for the outcomes are known to the decision maker, and where not even reasonable assumptions can be made about such probabilities. We consider here only decision making under certainty or risk, as does the majority of literature on decision theory.

Decision under risk

Before we enter into game theory proper we want to say more about decision under risk. Decision theory found interesting applications in pragmatics, and its ideas and concepts are fundamental for game theory. The decision maker may be uncertain about the outcomes of his actions because he has only limited information about the true state of the world. If Adam has to decide in the morning whether or not to take an umbrella with him, this depends on whether or not he believes that it will rain that day. He will not know this but will have some expectations about it. These expectations can be represented by probabilities, and Adam's information state by a *probability space*.

We identify a proposition A with sets of possible worlds. In probability theory they are called *events*; but we will stick here to the more familiar terminology from possible worlds semantics. If a person is convinced that A is true, then we assign probability 1 to it, and 0 if he thinks that it can not be true. If there are two propositions A and B that cannot be true at the same time, e.g. that the sky is sunny and that the sky is cloudy, then the probability of A *or* B is just the sum of the probability of A and the probability of B. The latter property is generalised in the following definition to arbitrary countable sequences of pairwise incompatible propositions.

Let Ω be a countable set that collects all possible states of the world. P is a *probability distribution* over Ω if P maps all subsets of Ω to the interval $[0, 1]$ such that:

1 $P(\Omega) = 1$;

2 $P(\sum_{j\in J} A_j) = \sum_{j\in J} P(A_j)$ for each family $(A_j)_{j\in J}$ of countably many pairwise disjoint sets. The sum $\sum_{j\in J} A_j$ here denotes the (disjoint) union of the sets A_j.

We call (Ω, P) a (countable) *probability space*. The restriction to countable Ω's simplifies the mathematics a lot. It follows e.g. that there is a subset $S \subseteq \Omega$ such that $P(\{v\}) > 0$ for each $v \in S$ and $P(A) = \sum_{v \in A \cap S} P(\{v\})$ for all $A \subseteq \Omega$, i.e it follows that P is a *count* measure. For $P(\{v\})$ we write simply $P(v)$.

If (Ω, P) describes the information state of a decision maker, what does his new information state look like if he learns a fact E? Adam may look out of the window and see that the sky is cloudy, or he may consult a barometer and see that it is rising. E would collect all worlds where the sky is cloudy, or, in the second scenario, where the barometer rises. If neither fact contradicts what Adam previously believed, then his probabilities for both sets must be greater than zero. Whatever proposition E represents, how does *learning* E affect (Ω, P)? In probability theory this is modelled by *conditional probabilities*. In learning theory, these are known as *Bayesian updates*. Let H be any proposition, e.g. the proposition that it will rain, i.e. H collects all possible worlds in Ω where it rains at some time of the day. The probability of H *given* E, written $P(H|E)$, is defined by:

$$P(H|E) := P(H \cap E)/P(E) \text{ for } P(E) \neq 0. \tag{1.1}$$

In particular, it is $P(v|A) = P(v)/P(A)$ for $v \in A \neq \emptyset$. For example, before Adam looked out of the window he may have assigned to the proposition $(E \cap H)$ that it is cloudy *and* that it rains a probability of $\frac{1}{3}$ and to the proposition (E) that it is cloudy a probability of $\frac{1}{2}$. Then (1.1) tells us that, *after* observing that the sky is cloudy, Adam assigns probability $\frac{1}{3} : \frac{1}{2} = \frac{2}{3}$ to the proposition that it will rain. Bayesian updates are widely used as a model for learning. P is often said to represent the *prior* beliefs, and P^+ defined by $P^+(A) = P(A|E)$ the *posterior* beliefs.

As an illustration we want to show how this learning model can be applied in Gricean pragmatics for explicating the notion of *relevance*. We discuss two approaches. The first one measures relevance in terms of the amount of information carried by an utterance and is due to Arthur Merin (Merin 1999b). The second approach introduces a measure that is based on expected utilities and is used by Prashant Parikh (Parikh 1992, Parikh 2001), Rohit Parikh (Parikh 1994) and Robert van Rooij (van Rooij 2003b).

The fact that the barometer is rising (E) provides evidence that the weather is becoming sunny. We can see the situation as a competition between two hypotheses: (H) *The weather will be sunny*, and $(\overline{H})$ *The weather will be rainy*. For simplicity we may assume that H and $\overline{H}$ are mutually exclusive and cover all possibilities. E, the rising of the barometer, does not necessarily imply that H, but our expectations that the weather will be sunny are much higher after learning E than before. Let P represent the given expectations before learning E, i.e. P is a probability distribution over possible states of the world. Let P^+ represent the expectations obtained from epistemic context (Ω, P) when E, and nothing but E, is learned. Modeling learning by conditional probabilities as above, we find that $P^+(H) = P(H|E)$, where we have to assume that $P(E) \neq 0$, i.e. we can only learn something that doesn't contradict our previous beliefs.

Our next goal is to introduce a measure for the *relevance* of E for answering the question whether H or $\overline{H}$ is true. Measures of relevance have been extensively studied in statistical decision theory (Pratt et al. 1995). There exist many different explications of the notion of *relevance* which are not equivalent with each other. We choose here Good's notion of relevance (Good 1950). It was first used by Arthur Merin (Merin 1999b), one of the pioneers of game theoretic pragmatics, in order to get a precise formulation of Grice's Maxim of Relevance.[1]

If we know $P(H|E)$, then we can calculate the reverse, the probability of E given H, $P(E|H)$, by *Bayes' rule*:

$$P(E|H) = P(H|E) \times P(E)/P(H). \tag{1.2}$$

With this rule we get:

$$P^+(H) = P(H|E) = P(H) \times (P(E|H)/P(E)). \tag{1.3}$$

$\overline{H}$ denotes the complement of H. Learning E influences our beliefs about $\overline{H}$ in the same way as it influences our beliefs about H: $P^+(\overline{H}) = P(\overline{H}|E)$. This leads us to:

$$\frac{P^+(H)}{P^+(\overline{H})} = \frac{P(H|E)}{P(\overline{H}|E)} = \frac{P(H)}{P(\overline{H})} \times \frac{P(E|H)}{P(E|\overline{H})}. \tag{1.4}$$

Probabilities are non-negative by definition. In addition we assume that all probabilities in this equation are positive, i.e., strictly greater than 0. This allows us to apply a mathematical trick and build the log of both sides of this equation. As the logarithm is strictly monotone it follows that (1.4) is true exactly iff

$$\log(P^+(H)/P^+(\overline{H})) = \log(P(H)/P(\overline{H})) + \log(P(E|H)/P(E|\overline{H})). \tag{1.5}$$

We used here the fact that $\log(x \times y) = \log x + \log y$. Furthermore we know that $\log x = 0$ iff $x = 1$. This means that we can use the term $r_H(E) := \log(P(E|H)/P(E|\overline{H}))$ as a measure for the ability of E to make us believe H. If it is positive, E favors H, if it is negative, then E favors $\overline{H}$. In a competitive situation where a speaker wants to convince his addressee of some proposition H it is reasonable to call a fact E more relevant the more evidence it provides for H. Merin calls $r_H(E)$ also the *argumentative force* of E.[2]

Whether or not this is a good measure of relevance in general depends on the overall character of communication. Merin sees the aim to convince our communication partner of something as the primary purpose of conversation. If Adam has an interview for a job he wants to get, then his goal is to convince the interviewer that he is the right person for it (H). Whatever he says is more *relevant* the more it favors H and disfavors the opposite proposition. We could see this situation as a battle between two agents, H and $\overline{H}$, where assertions E are the possible moves, and where $\log(P(E|H)/P(E|\overline{H}))$ measures the win for H and the loss for $\overline{H}$. Using the terminology that we will introduce in subsubection 1.2.1, we can say that this is a *zero-sum* game between H and $\overline{H}$.

We want to elaborate a little more on this example. The basis for Merin's proposal lies in the assumption that the main purpose of communication is to provide evidence that helps one decide whether a proposition H or its opposite is true. Hence, it works fine as long as we concentrate on yes-no questions or situations where one person tries to convince an addressee of the truth of some hypothesis. In general decision problems, the decision maker has to decide between different actions. Hence, the preferences over outcomes of these actions become important. It is not surprising that we find examples where a measure of relevance based on pure information becomes inadequate. Imagine that Ω consists of four worlds $\{v_1, \ldots, v_4\}$ of equal probability and that the decision maker has to decide between two actions a_1 and a_2. Suppose that she prefers a_1 in v_1 and v_2 and a_2 in v_3 and v_4 but that the value she assigns to a_1 in v_1 is very large compared to the other cases. If the decision maker learns $E = \{v_2, v_3\}$, then, using Merin's measure, this turns out to be irrelevant for deciding whether it is true that it is better to perform a_1 (i.e. $H = \{v_1, v_2\}$), or a_2 (i.e. $\overline{H} = \{v_3, v_4\}$) because $\log(P(E|H)/P(E|\overline{H})) = 0$. But, intuitively, it is relevant for the decision maker if she learns that the most favoured situation v_1 cannot be the case.

Let us return to the job interview example, and turn from Adam the interviewee to the interviewer. Let's call her Eve. From Eve's perspective the situation can be seen as a decision problem. She has to decide between two actions, *employ Adam* (a_1) or *not employ Adam* (a_2). Depending on the abilities of Adam these actions will be differently successful. The abilities are part of the various possible worlds in Ω. We can represent the success of the actions as seen by Eve by her preferences over their outcomes. We assume here that we can represent these preferences by a (von Neumann-Morgenstern) *utility measure*, or *payoff* function U that maps pairs of worlds and actions to real numbers. How does U have to be interpreted? If v is a world in Ω, then an equation like $U(v, a_1) < U(v, a_2)$ says that the decision maker prefers the outcome of action a_2 in v over the outcome of a_1 in v. $U(v, a_1)$ and $U(v, a_2)$ are real numbers, hence their difference and sum are defined. In utility theory, it is generally assumed that utility measures are unique up to *linear rescaling*, i.e. if $U(v, a) = r \times U'(v, a) + t$ for some real numbers $r > 0$ and t and all v, a, then U and U' represent the same preferences. If Eve values employing an experienced specialist twice as much as employing a trained and able novice, and she values employing an able novice as positively as she values employing an inexperienced university graduate negatively, then this can be modeled by assigning value 2 in the first case, value 1 in the second and value -1 in the third. But it could equally well be modeled by assigning 5 in the first, 3 in the second and -1 in the third case. Putting these parts together we find that we can represent Eve's decision problem by a structure $((\Omega, P), \mathcal{A}, U)$ where:

1 (Ω, P) is a probability space representing Eve's information about the world;

2 $\mathcal{A}$ is a set of actions;

3 $U : \Omega \times \mathcal{A} \longrightarrow \mathbf{R}$ is a utility measure.

In decision theory it is further assumed that decision makers optimize *expected utilities*. Let $a \in \mathcal{A}$ be an action. The *expected utility* of a is defined by:

$$EU(a) = \sum_{v \in \Omega} P(v) \times U(v, a) \tag{1.6}$$

Optimizing expected utilities means that a decision maker will choose an action a only if $EU(a) = \max_{b \in \mathcal{A}} EU(b)$. Let's assume that Eve assigns a probability of $p = \frac{3}{4}$ to the proposition that Adam is an inexperienced novice, but gives a probability of $1 - p = \frac{1}{4}$ to the proposition that he has some training. We further assume that she assigns value 1 to employing him in the first case, and value -1 to employing him in the second case. Furthermore, we assume that she values the state where she employs no candidate with 0. Then her expected utilities for employing and not employing him are $EU(a_1) = \frac{3}{4} \times (-1) + \frac{1}{4} \times 1 = -\frac{1}{2}$ and $EU(a_2) = 0$ respectively. Hence she should not employ Adam.

This may represent the situation before the interview starts. Now Adam tells Eve that he did an internship in a company X specialized in a similar field. This will change Eve's expectations about Adam's experience, and thereby her expected utilities for employing or not employing him. Using the ideas presented before, we can calculate the expected utility of an action a *after learning* A by:

$$EU(a|A) = \sum_{v \in \Omega} P(v|A) \times U(v, a); \tag{1.7}$$

where $P(v|A)$ denotes again the conditional probability of v given A. If Eve thinks that the probability that Adam is experienced increases to $\frac{3}{4}$ if he did an internship (A), then the expected utility of employing him now rises to $EU(a_1|A) = \frac{1}{2}$. Hence, Adam was convincing and will be employed. A number of people (P. Parikh, R. Parikh, R. van Rooij) proposed measuring the *relevance* of a proposition A in terms of how it influences a decision problem that underlies the current communication. Several possible ways to measure this influence have been proposed. One heuristic is to say that information A is relevant if and only if it makes a decision maker choose a different action from before, and it is more relevant the more it increases the expected utility. This is captured by the following measure of *utility value* of A:

$$UV(A) = \max_{a \in \mathcal{A}} EU(a|A) - EU(a^*|A). \tag{1.8}$$

a^* denotes here the action the decision maker had chosen before learning A — in our example this would have been a_2, not employing Adam. The expected utility value can only be positive in this case. If Eve had already a preference to employ Adam, then this measure would tell us that there is no relevant information that Adam could bring forward. So, another heuristic says that information is more relevant the more it increases expectations. This is captured by the following measure:

$$UV'(A) = \max_{a \in \mathcal{A}} EU(a|A) - \max_{a \in \mathcal{A}} EU(a). \tag{1.9}$$

If we put the right side in absolutes, then it means that information is the more relevant the more it changes expectations. This would capture cases where Adam could only say things that diminish Eve's hopes.

$$UV''(A) = |\max_{a \in \mathcal{A}} EU(a|A) - \max_{a \in \mathcal{A}} EU(a)|. \tag{1.10}$$

Obviously, Adam should convince Eve that he is experienced. Following Merin we could say that arguments are more relevant for Adam if they favor this hypothesis and disfavor the opposite. If Adam uses the utility-based measure of relevance, then he should choose arguments that make Eve believe that the expected utility after employing him is higher than that after not employing him. Given our scenario, this is equivalent with choosing arguments that favor the thesis that he is experienced. Hence, we see that for special cases both measures of relevance may coincide.

We want to conclude this section about decision theory with a classical example of Grice's. In this example there is no obvious hypothesis for which the provider of information could argue. Nevertheless, we can explain the relevance of his statement by a criterion based on the maximization of expected utilities.

A and B are planning their summer holidays in France. A has an open map in front of him. They would like to visit C, an old friend of B. So A asks B: *Where does C live?* B answers: *Somewhere in the south of France.* We are not concerned here with the question of how the implicature '*B does not know where C lives*' arises but with the question why B's answer is relevant. In Merin's model, there must be an hypothesis H that B argues for. But it is not immediately clear what this hypothesis H should be. We can model the situation as a decision problem where Ω contains a world for each sentence *C lives in* x, where x ranges over cities in France and where each of these worlds is equally possible. $\mathcal{A}$ contains all actions a_x of *going to* x, and U measures the respective utilities with $U(v, a) = 1$ if a leads to success in v and $U(v, a) = 0$ if not. Let E be the set of all worlds where C lives in the south of France. Calculating the expected utilities $EU(a_x|E)$ and $EU(a_x)$ for an arbitrary city x in the south of France would show that E increases the expected payoff of performing a_x. Hence, if B has no more specific information about where C lives, then a criterion that measures *relevance* according to whether or not it increases expected utilities would predict that E is the most relevant answer that B could give.

1.2 Games

What distinguishes game theory from decision theory proper is the fact that decisions have to be made with respect to the decisions of other players. We start this section with some fundamental classifications of games and introduce the normal form. We look then at one example in more detail, the *prisoners' dilemma*. In section 1.2.3 we present the most fundamental solution concepts of game theory, especially the concept of a *Nash equilibrium*. Finally, we introduce the extensive form. The latter is more suitable for sequential games, a type of game in terms of which communication is studied a lot.

1.2.1 Strategic games and the normal form

There exist several important different classifications of games which are widely referred to in game theoretic literature. We provide here a short overview.

A first elementary distinction concerns that between *static* and *dynamic games*. In a static game, every player performs only one action, and all