LOFT 2008

Logic and the Foundations of Game and Decision Theory

Institute for Logic, Language and Computation
Universiteit van Amsterdam

3 – 5 July 2008

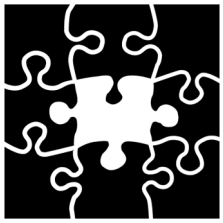

LOFT 2008
Logic and the Foundations of Game and Decision Theory
Institute for Logic, Language and Computation
Universiteit van Amsterdam
3 – 5 July 2008

The conference LOFT 2008 is the eighth in a series of conferences on the applications of logical methods to foundational issues in the theory of individual and interactive decision-making. Traditionally, preference is given to papers which bring together the work and problems of several fields, such as game and decision theory, logic, computer science and artificial intelligence, philosophy, cognitive psychology, mathematics and mind sciences. The previous LOFT conferences were held in Marseille (LOFT 1994), Torino (LOFT 1996, LOFT 1998, LOFT 2000, LOFT 2002), Leipzig (LOFT 2004), and Liverpool (LOFT 2006).

LOFT 2008 is hosted by the Institute for Logic, Language and Computation (ILLC) of the *Universiteit van Amsterdam*, and is funded by the *Nederlandse Organisatie voor Wetenschappelijk Onderzoek* (NWO, DN 612.080.000 CN 2008/05570/EW) and by the Marie Curie Research Training Site GLoRiClass ("**G**ames in **Lo**gic **R**eachi**n**g out to **Class**ical game theory"; MEST-CT-2005-020841).

General Chairs & the Programme Committee

The General Chairs of LOFT 2008 were Giacomo Bonanno (University of California at Davis, U.S.A.), Benedikt Löwe (Universiteit van Amsterdam, The Netherlands), and Wiebe van der Hoek (The University of Liverpool, United Kingdom). The General Chairs would like to thank the programme committee (listed below) for their input during the refereeing of the submissions. Special thanks go to Peter van Ormondt and Sanchit Saraf for their support during the organization of the conference.

Programme Committee.

- Krzysztof Apt, Centrum voor Wiskunde en Informatica, The Netherlands
- Alexandru Baltag, Oxford University, United Kingdom
- Johan van Benthem, Universiteit van Amsterdam, The Netherlands
- Adam Brandenburger, New York University, United States of America
- Hans van Ditmarsch, University of Otago, New Zealand
- Ulle Endriss, Universiteit van Amsterdam, The Netherlands

- Erich Grädel, Rheinisch-Westfälische Technische Hochschule Aachen, Germany
- Joe Halpern, Cornell University, United States of America
- Wojtek Jamroga, Technische Universität Clausthal, Germany
- Jérôme Lang, Université Paul Sabatier, France
- Hannes Leitgeb, University of Bristol, United Kingdom
- Larry Moss, Indiana University, United States of America
- Rohit Parikh, City University of New York, United States of America
- Marc Pauly, Stanford University, United States of America
- Andres Perea, Universiteit Maastricht, The Netherlands
- Ramanujam, Institute of Mathematical Sciences, India
- Robert van Rooij, Universiteit van Amsterdam, The Netherlands
- Hans Rott, Universität Regensburg, Germany
- Robert Stalnaker, Massachusetts Institute of Technology, United States of America
- Moshe Tennenholtz, Technion, Israel
- Rineke Verbrugge, Rijksuniversiteit Groningen, The Netherlands

Proceedings.

Many speakers of the conference have sent us PDF files of their papers. These can be found on the conference webpage `http://www.illc.uva.nl/LOFT2008/` (go to "Accepted Papers" and click on the PDF icons).

The proceedings of LOFT 2008 will be published in the new book series Texts in Logic and Games with Amsterdam University Press in which the LOFT 2006 proceedings have just appeared (as volume 3). For more information about the book series, visit its web site at `http://www.illc.uva.nl/TLG/`. The General Chairs will invite about twenty to twenty-five papers presented at the conference for submission in the proceedings volume. All papers will be thoroughly peer-reviewed according to the standards of a high-level research journal. The proceedings volume will appear within two years of the conference and will be sent to all registered participants of the conference who paid the fee. Make sure that you keep us informed about changes of address so that we will be able to send you your copy of the proceedings.

Schedule

Thursday, July 3

10:00-11:10: Registration

11:10-11:30: Opening (Room P.227)

11:30-12:30: **Invited talk** (Room P.227): Thomas Henzinger, *Games in System Design and Verification*

12:30-14:30: Lunch Break

14:30-15:30: **Session Plenary 1** (Room P.227)

14:30-14:50: Dov Samet, *Agreeing to Disagree: the Non-Probabilistic Case*

14:50-15:10: Joseph Halpern, Dov Samet, and Ella Segev, *Defining Knowledge in Terms of Belief*

15:10-15:30: Dov Samet, *S5 Knowledge without Partitions*

15:30-15:50: Coffee Break

15:50-16:50: **Session Plenary 2** (Room P.227)

15:50-16:10: Giacomo Bonanno, *Rational Choice and Belief Revision*

16:10-16:30: Jérôme Lang and Leon van der Torre, *From Belief Change to Preference Change*

16:10-16:50: Jan van Eijck and Floor Sietsma, *Multi-Agent Belief Revision with Linked Plausibilities*

16:50-17:10: Coffee Break

17:10-18:10: **Session Triple 1** (Rooms P.015A, P.016, P.017)

P.015A

17:10-17:30: Davide Grossi, *Proving Theorems of Judgment Aggregation as Corollaries of Theorems of Preference Aggregation*

17:30-17:50: Christian Klamler and Daniel Ecker, *A Geometric Approach to Judgment Aggregation*

17:50-18:10: Felix Brandt and Paul Harrenstein, *Dominance in Social Choice and Coalitional Game Theory*

P.016

17:10-17:30: Bernard Tchantcho and Lawrence Diffo Lambo, *A Note on the Stability of Voting Games*

17:30-17:50: Luc Bovens and Claus Beisbart, *Measuring Influence for Independent Voters: A Generalisation of the Banzhaf Measure*

17:50-18:10: Satoru Suzuki, *Preference Logic and its Measurement-Theoretic Semantics*

P.017

17:10-17:30: NO TALK

17:30-17:50: Till Grüne Yanoff, *Evolutionary Game Theory: Natural Selection Implies Interpersonal Comparisons*

17:50-18:10: Giacomo Sillari, *Quantified Logic of Awareness and Impossible Possible Worlds*

18:10-18:30: Coffee Break

18:30-19:30: **Session Plenary 3** (Room P.227)

18:30-18:50: Hans van Ditmarsch and Tim French, *Simulation and Information: Quantifying over Epistemic States*

18:50-19:10: Nils Bulling and Wojciech Jamroga, *What Agents can Probably Enforce*

19:10-19:30: Dietmar Berwanger, *Infinite Coordination Games*

Friday, July 4

9:50-10:50: **Invited talk** (Room P.227) Christian List, *Majority voting on restricted domains*

10:50-11:10: Coffee Break

11:10-12:30: **Session Duo 1** (Rooms P.017 and P.018)

P.017

11:10-11:30: Wojciech Jamroga, Thomas Ågotnes, and Wiebe van der Hoek, *A Simpler Semantics for Abilities under Uncertainty*

11:30-11:50: Claus Beisbart, *Joint Action Can Make a Difference: Measures of Voting Power Generalised*

11:50-12:10: Michael Franke, *Meaning and Inference in Case of Conflict*

12:10-12:30: Kris de Jaegher, Stephanie Rosenkranz, and Utz Weitzel, *Economic Laboratory Experiment on Horn's Rule*

P.018

11:10-11:30: Itai Sher, *Persuasion and Limited Communication*

11:30-11:50: Franz Dietrich, *Anti-terrorism Policies: the trade-off Between Provocation and Deterrence*

11:50-12:10: Konstantinos Georgatos, *Geodesic Conditioning*

12:10-12:30: Christian Fermüller, *Truth Value Intervals, Bets, and Dialogue Games*

12:30-14:30: Lunch Break

14:30-15:30: **Session Plenary 4** (Room P.227)

14:30-14:50: Johan van Benthem and Cédric Dégremont, *Building Bridges between Dynamic and Temporal Doxastic Logics*

14:50-15:10: Eric Pacuit and Jonathan Zvesper, *Assumption-incompleteness in Modal Logic*

15:10-15:30: Nicolas Troquard, Wiebe van der Hoek, and Michael Wooldridge, *Reasoning about Strategic Games with Hybrid Logic of Choice and Preferences*

15:30-15:50: Coffee Break

15:50-16:50: **Session Plenary 5** (Room P.227)

15:50-16:10: Lena Kurzen, *A Logic for Cooperation, Actions and Preference*

16:10-16:30: Sujata Ghosh, *Strategies Made Explicit in Dynamic Game Logic*

16:30-16:50: Robert van Rooij, *Revealed Preferences and Satisficing Behavior*

16:50-17:10: Coffee Break

17:10-18:10: **Session Plenary 6** (Room P.227)

17:10-17:30: Krzysztof Apt, Andreas Witzel, and Jonathan Zvesper, *Strategic Games with Interaction Structures*

17:30-17:50: Aviad Heifetz, Martin Meier, and Burkhard Schipper, *Dynamic Unawareness and Rationalizable Behavior*

17:50-18:10: Frederick Herzberg, *Hyperreal Expected Utilities and Pascal's Wager*

18:10-18:30: Coffee Break

18:30-19:30: **Invited talk** (Room P.227): Nicole Schweikardt, *Logic and Data Exchange: Which Solutions are "Good" Solutions?*

Saturday, July 5

9:50-10:50: **Invited talk** (Room P.227): Hervé Moulin, *Cost Sharing in the Capacity Synthesis Problem*

10:50-11:10: Coffee Break

11:10-12:30: **Session Duo 2** (Rooms P.016 and P.017)

P.016

11:10-11:30: Thomas Ågotnes, Valentin Goranko, and Wojciech Jamroga, *Strategic Commitment and Release in Logics for Multi-Agent Systems*

11:30-11:50: Gabriella Pigozzi, Marija Slavkovik, and Leon van der Torre, *Conclusion-Based Procedure for Judgment Aggregation Satisfying Premise Independence*

11:50-12:10: Guiseppe Primiero, *Aggregating Collective Judgments by Selecting Disagreements*

12:10-12:30: no talk

P.017

11:10-11:30: Tomohiro Hoshi, *Public Announcement Logics with Constrained Protocols*

11:30-11:50: Matthias Hanauske, Steffen Bernius, Wolfgang König, Berndt Dugall, *Experimental Validation of Quantum Game Theory*

11:50-12:10: Oliver Board and Kim-Sau Chung, *Object Based Unawareness II: Applications*

12:10-12:30: Olivier Roy, *Intentions and Transformations of Strategic Games*

12:30-14:30: Lunch Break

14:30-15:10: **Session Triple 2** (Rooms P.014, P.016, P.017)

P.014

14:30-14:50: Alexandru Baltag, Sonja Smets, and Jonathan Zvesper, *When All is Done and not (yet) Said: Dynamic Rationality in Extensive Games*

14:50-15:10: Sascia Pavan, *Disambiguation Games in Extended and Strategic Form*

P.016

14:30-14:50: Riccardo Pucella and Mehrnoosh Sadrzadeh, *Towards a Concrete Semantics for Announcements*

14:50-15:10: Davi Romero de Vasconcelos and Edward Hermann Haeusler, *Towards an Implementation Theory via a Game Logic Approach*

P.017

14:30-14:50: Conrad Heilmann, *A Representation of Time Discounting as Diminishing Intrapersonal Connectedness*

14:50-15:10: Brian Hill, *Dynamic Awareness*

15:10-15:30: Coffee Break

15:30-16:30: **Invited talk** (Room P.227): Wolfgang Spohn, *Dependency Equilibria:*
A New(!) and Reasonable(?) Equilibrium Concept

Invited Talks

Games in System Design and Verification

Tom Henzinger
École Polytechnique Fédérale de Lausanne, Switzerland
thenzinger@gmail.com

We consider graph games, where a token is moved along the edges of a directed graph in an infinite sequence of moves. Graph games come in many different flavors, depending for example, on the number of players and their relationship, the information that each player has available for choosing a move, the mechanism (probabilistic or not) by which the successor vertex is determined from a given choice of moves, and the type of objective for each player. Graph games can serve as models for discrete systems: vertices represent states; players represent components; moves represent transitions; and objectives represent specifications. The variety of possible graph games permits the modeling of many different interaction paradigms for component-based systems.

We present a convenient notation for conversing about graph games, which is based on Strategy Logic. We then survey several applications of graph games in system design and verification. The classical application is component synthesis, which goes back to a problem posed by Alonzo Church. More recent applications include the following: discrete-event control, multi-agent planning, refinement relations between systems, modeling fairness in open systems, early counterexample detection in model checking, the composition of behavioral interfaces, model-based testing, and assume-guarantee modular verification.

Majority rule revisited

Christian List
London School of Economics, United Kingdom
C.List@lse.ac.uk

Although majority rule is one of the most commonly used collective decision rules, it notoriously fails to generate consistent collective outputs for some combinations of individual inputs. For example, if one third of a group accepts "p", "if p then q" and "q", a second third accepts "not p", "if p then q" and "not q", and the final third accepts "p", "not (if p then q)" and "not q", there are majorities for "p", "if p then q" and "not q", an inconsistent set of propositions. Variants of this observation were the starting points for both classical social choice theory and the more recent theory of judgment aggregation. But how serious is the problem of majority inconsistency? Does it really undermine the viability of majority rule as a collective decision rule? In this talk, I will show that the level of pluralism in the relevant group or society is the key determinant for the viability of majority rule. Drawing on joint work with Franz Dietrich on judgment aggregation,

I will discuss several constraints on pluralism that are sufficient for consistent majority decisions. Our most general condition generalizes a classic condition introduced by Amartya Sen in standard social choice theory: the condition of value-restriction. I will also present a new characterization theorem: for a large class of domains of individual inputs, if there exists any decision rule satisfying some democratic conditions, then majority rule is the unique such decision rule. Taken together, these findings provide new support for the robustness of majority rule.

Cost sharing in the capacity synthesis problem

Hervé Moulin
Rice University, United States of America
moulin@rice.edu

We consider two resource allocation problems where one dimensional mechanisms achieve allocative efficiency in equilibrium under quasi-linear preferences, but are unable to balance the budget. We define the efficiency loss of a mechanism as its largest ratio of budget imbalance to efficient surplus, over all profiles of individual preferences. We look for mechanisms minimizing the efficiency loss. To assign one (desirable or undesirable) object among n agents, we find a subsidy-free VCG mechanism guaranteeing voluntary participation and an efficiency loss of $\frac{2n}{2^n}$. This mechanism is optimal in the VCG class, and among all strategy-proof mechanisms treating equals equally. The result generalizes partially to the assignment of multiple identical objects, where the optimal mechanism is no longer a VCG one, and the tradeoff between voluntary participation and minimizing the efficiency loss can be severe. To exploit a one input-one output convex technology, n users request a quantity of output and the mechanism assigns cost shares. In the canonical residual* mechanism, total surplus is maximized in any Nash equilibrium of a natural potential game. Total payment may differ from the actual cost and a user with a null demand may be subsidized. If the cost function is totally monotone (e.g., polynomial with positive coefficients, or exponential), participation is voluntary and total payment covers at least actual cost. The efficiency loss is at most $\min\{2/\log n, 1\}$. For power cost functions, $C(a) = ap$, $p < 1$, the above ratio converges to zero as $\frac{1}{n^{p-1}}$. Participation appears to be voluntary however a small budget deficit is possible. These properties are lost if the cost function is not smooth.

Logic and Data Exchange: Which solutions are "good" solutions?

Nicole Schweikardt
Johann Wolfgang Goethe-Universität Frankfurt, Germany
schweika@informatik.uni-frankfurt.de

A data exchange setting $D = (\sigma, \tau, \Sigma)$ consists of a source schema σ (i.e., a finite set of relation symbols), a target schema τ, and a set Σ of logical formulas that (a) specify the relationship between the source and the target and (b) specify particular properties that the target should have. Data exchange deals with the following problem: given a data exchange setting D and a source database S (i.e., a finite structure of vocabulary σ), construct a database T over the target schema τ that satisfies the relationships and properties listed in Σ. Such a target instance T is called a solution for S with respect to D. Preferably, in case that solutions exist at all, one would like to find particular solutions that reflect the given source data "as accurately as possible". In this talk I want to give an introduction to the area of data exchange, with a special emphasis on the question of which solutions can be considered to be "good" solutions. I will concentrate on notions of "good" solutions for query answering over target instances in the contexts of open world semantics, closed world semantics, and a combination of both semantics.

Dependency Equilibria: A New (!) and Reasonable (?) Equilibrium Concept

Wolfgang Spohn
Universität Konstanz, Germany
Wolfgang.Spohn@uni-konstanz.de

The notion of a Nash equilibrium is fundamental for game theory and its rich applications. It rests on an assumption that is is more or less explicit and taken for absolutely granted, namely the assumption of the causal independence of the decisions and the actions of the players. In my talk I want

1. to show how one can meaningfully undermine this assumption,
2. to explain how abandoning this assumption leads to a new equilibirum concept that I call dependency equilibrium,
3. to indicate some consequences of these considerations the most important of which is that mutual cooperation in the one-shot prisoners' dilemma turns out to be rational, i.e. to form such a dependency equilibrium.

Contributed Talks

Strategic Commitment and Release in Logics for Multi-Agent Systems

Thomas Ågotnes
Bergen University College, Norway
tag@hib.no

Valentin Goranko
University of the Witwatersrand, South Africa
Valentin.Goranko@wits.ac.za

Wojciech Jamroga
Clausthal University of Technology, Germany
wjamroga@in.tu-clausthal.de

We analyze the semantics of **ATL*** and related logics for strategic abilities. Our focus is on how these logics treat agents' commitments to strategies in the process of formula evaluation. We point out some questionable effects in that respect in the standard semantics, and discuss some alternatives leading to amendments of that semantics. We also propose and discuss various syntactic and semantics mechanisms for handling commitments to strategies and release from such commitments in the semantics of **ATL***, leading to more expressive and semantically refined versions of that logic.

Strategic games with interaction structures

Krzysztof Apt
Universiteit van Amsterdam, The Netherlands
K.R.Apt@cwi.nl

Andreas Witzel
Universiteit van Amsterdam, The Netherlands
awitzel@science.uva.nl

Jonathan Zvesper
Universiteit van Amsterdam, The Netherlands
jonathan@illc.uva.nl

A customary assumption in the study of strategic games is that the players have common knowledge of the underlying game. We weaken this assumption and investigate game-theoretic and epistemic consequences of limited knowledge of payoffs in strategic games.

To examine the dynamics of such situations, we introduce "interaction structures" which consist of hyperarcs across which players can communicate. A set of players A being connected in an interaction structure means that they can multicast (synchronously) across their common communication channel, creating

common knowledge amongst the group A. Each message multicast specifies a particular player's payoff for some particular strategy profile.

One can think of an interaction structure as the groups of coauthors who are involved in writing joint papers: information is shared amongst subgroups and one author might be in multiple groups. In the case we consider here, the information that can be shared is about players' payoffs. At each intermediate stage of the information sharing across an interaction structure it is possible (under the standard assumption of common knowledge of rationality) for players iteratively to eliminate (some) strictly dominated strategies. In the case when there is a single hyperarc that consists of all players, the communication of payoffs across this hyperarc yields the customary iterated elimination of strictly dominated strategies.

We also study the the epistemic foundations for the resulting solution concept, which is parametrised by the interaction structure.

When All is Done and not (Yet) Said: Dynamic Rationality in Extensive Games

Alexandru Baltag
Oxford University, United Kingdom
albaltag@yahoo.com

Sonja Smets
Vrije Universiteit Brussel, Belgium
sonsmets@vub.ac.be

Jonathan Zvesper
Universiteit van Amsterdam, The Netherlands
jonathan@illc.uva.nl

In this paper we provide another analysis of the epistemic foundations for backward induction in finite extensive games of perfect information.

We ground this analysis in 'Conditional Doxastic Logic', a dynamic logic for beliefs. In grounding our analysis in beliefs we depart from the tradition started by Aumann, who analyses backward induction in terms of common *knowledge*. For us, knowledge is a very strong notion: if a node is commonly known, before any play of the game, not to be reachable, then it should not appear in the game tree – a condition we call 'open future'. Another departure from the standard analysis is that, because we allow for *conditional* beliefs we do not require that a player has a fully determined strategy (though we certainly allow for that possibility). Our main theorem states that:

> Common knowledge of the game structure, of 'open future' and of stable belief in 'rational play' entails common belief in the backward induction outcome.

We extensively discuss connections with the literature.

Joint action can make a difference: Measures of voting power generalized

Claus Beisbart
Technische Universität Dortmund, Germany
Claus.Beisbart@udo.edu

The political influence that a voter has under a specific decision rule is often measured in terms of the probability that that voter's vote is critical. If it is calculated on the basis of the Bernoulli model, the popular Banzhaf measure of (a priori) voting power is obtained. But the probability of being critical can also be calculated for alternative probability models. If they are constrained by empirical data, measures of a posteriori voting power arise.

In the recent literature, it has been argued that the probability of criticality does not provide a suitable measure of a posteriori voting power. As an example due to G. Wilmers shows, this measure will assign zero voting power to every voter, if voting profiles under which at least one voter is critical have zero probability. It seems odd that nobody has any power whatsoever, though.

I therefore propose to go beyond the probability of being critical for measuring voting power. The probability of criticality quantifies the extent to which a voter has the opportunity to make a difference as to whether a bill passes or not. Likewise, one can calculate the extent to which a voter has the opportunity to find other voters in order to form a group that makes a difference as to whether a bill passes or not. Put differently, for each voter, we look at the opportunities for group actions that involve her.

I proceed in two steps. I first define criticality for a group. Roughly, given a specific coalition, a group is critical, iff the following is true: There is some way in which the group could have voted differently such that the outcome of the vote would have been different. The second step introduces criticality of higher ranks for individual voters. Roughly, a voter is critical of rank κ, iff there is a group of votes including a and with cardinality κ such that the group is critical. The proposed new measures are then the probabilities that a voter is critical of a fixed rank. In order to avoid double-counting, I introduce a differential counting. As a result, for each voter there is a hierarchy of measures with ranks ranging from 1 to the cardinality of the assembly. The powers of the different ranks add up to 1 for each voter. Roughly, you have overall more voting power than I have, if your measures start growing for smaller ranks than mine do. For rank 1 and the Bernoulli model, the new measure coincides with the Banzhaf measure. For higher ranks, additional information is provided.

Infinite Coordination Games

Dietmar Berwanger
Rheinisch-Westfälische Technische Hochschule Aachen, Germany
dwb@logic.rwth-aachen.de

We investigate the prescriptive power of sequential iterated admissibility in coordination games of the Gale-Stewart style, i.e., perfect-information games of infinite duration with only two payoffs. We show that on this kind of games the procedure of eliminating weakly dominated strategies is independent of the elimination order and that, under maximal simultaneous elimination, the procedure converges in countably many stages. We identify the class of games on which iterated admissibility ensures coordination and point out that the class where it cannot avoid coordination failure is unsatisfactorily large. For the particular case of games with ω-regular winning conditions, we propose a solution based on preplay agreement.

Object-Based Unawareness II: Applications

Oliver Board
University of Pittsburgh, United States of America
ojboard+@pitt.edu

Kim-Sau Chung
University of Minnesota, United States of America
sau@umn.edu

Board & Chung (henceforth BC) propose a class of structures, *OBU structures*, designed to model agents' awareness of their own unawareness. In particular, OBU structures are capable of modeling an agent who is "not sure whether or not she is aware of everything", without the implication that she does not know what he is aware of. In BC, OBU structures were used to provide semantics for a language of first-order modal logic; in this paper, we show that essentially the same structures can be interpreted without appeal to a formal language, creating a purely set-theoretic model of unawareness.

More formally, an *OBU structure for n agents* is a tuple

$$\langle W, O, \{O_w\}, \{\mathcal{I}_i\}, \{\mathcal{A}_i\}\rangle,$$

where W is a set of states; O is a set of objects; $O_w \subseteq O$ is the set of *real* objects at state w; $\mathcal{I}_i : W \to 2^W$ is an *information function* for agent i; and $\mathcal{A}_i : W \to 2^O$ is an *awareness function* for agent i. Intuitively, $\mathcal{I}_i(w)$ indicates the states that agent i considers possible when the true state is w, while $\mathcal{A}_i(w)$ indicates the objects she is aware of.

In the standard information partition model familiar to economists, events are subsets of the state space, corresponding (roughly) to the set of states in

which some given proposition is true. In the our model, an event is an ordered pair (R, S), where $R \subseteq 2^W$ is a set of states and $S \subseteq 2^O$ is a set of objects; we call R the *reference* of the event (denoted $\text{ref}(R, S)$), corresponding (as before) to the set of states in which some proposition is true; and S is the *sense* of the event (denoted $\text{sen}(R, S)$), listing the set of objects referred to in the description of the proposition. (To give an example, the events representing the propositions "the dog barked" and "the dog barked and the cat either did or did not meow" have the same reference but difference senses.)

Two operators on events, corresponding to "not", "and", are defined as follows.

$$\neg(R, S) = (W \setminus R, S) \qquad \text{and} \qquad \wedge_j (R_j, S_j) = (\cap_j R_j, \cup_j S_j)$$

The negation of an event holds at precisely those states at which the event does not hold, but it refers to the same set of objects. The conjunction of two events holds only at those states at which *both* events hold, and it refers to both sets of objects.

We also introduce three modal operators for each agent, representing awareness (A_i), implicit knowledge (L_i), and explicit knowledge (K_i):

$$\begin{aligned}
\mathsf{A}_i(R, S) &= (\{w \mid S \subseteq \mathcal{A}_i(w)\}, S) \text{ (awareness)} \\
\mathsf{L}_i(R, S) &= (\{w \mid \mathcal{I}_i(w) \subseteq R\}, S) \text{ (implicit knowledge)} \\
\mathsf{K}_i(R, S) &= \mathsf{A}_i(R, S) \wedge \mathsf{L}_i(R, S) \text{ (explicit knowledge)}
\end{aligned}$$

Intuitively, an agent is aware of an event if she if aware of the objects used to describe it, while she implicitly knows and event if the event holds at every state she considers possible; finally, she explicitly knows an event if she is aware of it and she implicitly knows it. In the full version of the paper (http://kimsauchung.googlepages.com/awareapp.pdf.), we prove characterization results which provide a precise description of the features of these structures.

In the full version of the paper, two applications are presented. The first examines the legal interpretive doctrine *verba fortius accipiuntur contra proferentem*, which instructs judges to resolve any contractual ambiguity against the party who drafted the contract. We show that ambiguity in human language alone does not explain why this doctrine is better than possible alternatives; assuming that there is persistent asymmetric unawareness between the contracting parties, however, as well as ambiguity in language, we can provide such an explanation.

The second application revisits a central result in Heifetz, Meier, and Schipper. HMS claim that "unawareness can be interpreted as a special form of delusion". However, surprisingly, they prove that the No-Trade Theorem continues to hold despite this existence of delusion. In our framework, unawareness and delusion are distinct, and unawareness and non-delusion can co-exist in general. Interpreted in these terms, the HMS model is a special case where they do not. We show

that it satisfies a property called terminal-non-delusion, and it is this property that drives their result.

Rational Choice and Belief Revision

Giacomo Bonanno
University of California at Davis, United States of America
gfbonanno@ucdavis.edu

We show a correspondence between the structures used in rational choice theory and the AGM theory of belief revision. A choice structure $\langle W, C, f\rangle$ consists of a set W of alternatives, a collection C of subsets of W, representing possible choice sets, and a function f from C into the set of subsets of W, representing choices made. If E belongs to C then $f(E)$ is the subset of E consisting of the chosen alternatives in E. It is assumed that if E is non-empty then so is $f(E)$. The main objective of rational choice theory is to investigate the conditions under which the function f can be rationalized by a total pre-order R on W (that is, a binary relation on W which is connected, reflexive and transitive) in the sense that $f(E)$ coincides with the best elements of E relative to R. We re-interpret choice structures in terms of belief revision. We interpret the set W as a set of states or possible worlds and we restrict attention to choice frames where the collection C contains W and does not contain the empty set. Given a propositional language L, a model based on a given choice structure is obtained by adding a valuation V that assigns to every atomic proposition p the set of state at which p is true. Truth of an arbitrary formula at a state is then defined as usual. Given a model $\langle W, C, f, V\rangle$ we define initial beliefs as the set of formulas X such that $f(W)$ is a subset of $\|X\|$, where $\|X\|$ denotes the truth set of X, that is, the set of states where X is true. If Y is a formula such that $\|Y\|$ belongs to C, then we define the revised beliefs upon learning that Y as the set of formulas X such that f($\|Y\|$) is a subset of $\|X\|$. Thus associated with every model $\langle W, C, f, V\rangle$ is a partial belief function B from the set of formulas Y such that $\|Y\|$ belongs to C into the set of subsets of formulas. We say that a choice structure $\langle W, C, f\rangle$ is AGM-consistent if, for every model based on it, the associated partial belief function can be extended to a full-domain belief function that satisfies the AGM axioms. We prove the following Proposition: a choice structure $\langle W, C, f\rangle$ with W finite is AGM-consistent if and only if it is rationalizable in terms of a total pre-order on W.

Measuring Influence for Dependent Voters: A Generalisation of the Banzhaf Measure

Luc Bovens
London School of Economics, United Kingdom
L.Bovens@lse.ac.uk

Claus Beisbart
Technische Universität Dortmund, Germany
Claus.Beisbart@udo.edu

We construct a measure of voting power that is a reasonable measure of influence even if the votes are not cast independently. The crucial building blocks of our measure are probabilities of counterfactuals, such as the probability that the outcome of the vote would have been yes, had a voter voted yes rather than no as she did in the real world. The probabilities of such counterfactuals are calculated on the base of causal information, following the approach by Balke and Pearl. Opinion leaders who have causal influence on other people's votes can have significantly more voting power under our measure than the latter. We provide several examples in which our measure yields intuitively plausible results and show that the measure reduces to Banzhaf voting power in the limiting case of independent and equiprobable votes.

Dominance in Social Choice and Coalitional Game Theory

Felix Brandt
Ludwig-Maximilians-Universität München, Germany
brandtf@tcs.ifi.lmu.de

Paul Harrenstein
Ludwig-Maximilians-Universität München, Germany
paul.harrenstein@ifi.lmu.de

McGarvey (1953) has shown that any irreflexive and anti-symmetric relation can be obtained as a relation induced by majority rule. We address the analogous issue for the dominance relations of cooperative games with non-transferable utility (coalitional NTU games). We find any irreflexive relation over a finite set can be obtained as the dominance relation of some ordinary and monotonic NTU coalitional game. We also show that any dominance relation can be induced by a non- cooperative game via beta-effectivity. For a dominance relation to be obtainable via alpha-effectivity, however, it has to comply with more restrictive conditions. Finally, we consider the formal interrelationships between dominating sets and undominated sets in finite NTU games.

Towards an Implementation Theory via a Game Logic Approach

Nils Bulling
Clausthal University of Technology, Germany
bulling@in.tu-clausthal.de

Wojciech Jamroga
Clausthal University of Technology, Germany
wjamroga@in.tu-clausthal.de

Alternating-time Temporal Logic (ATL) is probably the most influential logic of strategic ability that has emerged in recent years. The idea of ATL is centered around so-called cooperation modalities: $\ll A \gg \gamma$ is satisfied if the group of agents A has a collective strategy to enforce temporal property γ against *the worst* possible response from the other agents. So, the semantics of ATL shares the "all-or-nothing" attitude of many logical approaches to computation.

Economic Laboratory Experiment on Horn's Rule

Kris De Jaegher
Utrecht University, The Netherlands
K.deJaegher@econ.uu.nl

Stephanie Rosenkranz
Utrecht University, The Netherlands
S.Rosenkranz@econ.uu.nl

Utz Weitzel
Utrecht University, The Netherlands
U.Weitzel@econ.uu.nl

This paper presents an economic laboratory experiment testing the validity of different game-theoretic rationales for Horn's rule. The starting point of Horn's rule is that it is costly for speakers to disambiguate their utterances to the full extent. For this reason, speakers only disambiguate if they are referring to an unlikely event; for likely events, they do not disambiguate, expecting that rational listeners will then infer that they are referring to a likely event.

In the basic signalling game that is used to account for Horn's rule, there is an unlikely state and a likely state. The sender knows the state, the receiver does not; sender and receiver have common interests. By sending a signal without a preset meaning, the sender tries to induce the receiver to correctly guess the true state is. Sending the signal is costly. In the absence of communication, the receiver guesses the likely state. In case there is communication, Horn's rule applies if the sender only sends the signal in the unlikely state (note that an equilibrium also exists where the sender only sends a message in the likely state). We start by reviewing the different game-theoretic rationales that have been given for application of Horn's rule in this game. Focal point theory points

out that the equilibrium obeying Horn's rule is the best one for both players, and that it is therefore focal to both players. The intuitive criterion says that, if originally no signalling was taking place, the receiver of a signal should realise that only the sender in the unlikely state has an incentive to send a costly signal. The evolutionary argument also starts from the premise that at first, there is no communication between players. Receiving no signal then "means" to the listener that the likely state should be guessed; this will therefore also be the case in any signalling equilibrium that can evolve from this situation.

Unfortunately, when the different game-theoretic rationales for application of Horn's rule are applied to the basic game, they mostly make completely identical predictions. This is why we introduce variants on the basic signalling game such that the different theories do yield different predictions. In a first variant, in the absence of signalling, the receiver guesses the unlikely state. This is possible when it is risky to wrongly guess that the likely state occurs. In a second variant, there is no cost involved with sending a signal. In our laboratory experiment, we replicate both the basic signalling game, and the variants on this game, by giving participants financial incentives. We find that, as soon as the signal is costly, Horn's rule predominantly applies. This supports focal point theory, which predicts that players will focus on the efficient outcome. However, in the variant of the game where the signal is costless, we find support for the evolutionary argument. The way in which players played in the absence of signalling determines the type of signalling equilibrium that is obtained.

Anti-terrorism politics and the risk of provoking

Franz Dietrich
Maastricht University, The Netherlands
F.Dietrich@KE.unimaas.nl

The level of terrorism in a population depends on two factors: people's preferences (especially, whether they would like to create damage) and the constraints under which people act (especially, what damage they could create and what punishment the have to expect). Cause-related policies, such as raising the level of social stability or education, aim at appeasing preferences, thereby reducing terrorism. Symptom-related policies, such as military interventions or criminal legislation, change the constraints ('deterrence'), but they may also have side effects on preferences ('provocation'); whether terrorism overall decreases depends on whether deterrence overweighs provocation. I formally model the trade-off between deterrence and provocation. I argue that the provocation effect of tough policies is easy to overlook in practice, and I show that in this case a too tough policy is used, leading to more terrorism than in optimum.

Truth Value Intervals, Bets, and Dialogue Games

Christian Fermüller
Technische Universität Wien
chrisf@logic.at

Mainly due to rather recent developments, t-norm based fuzzy logics are well understood technically, i.e., algebraically and proof theoretically, by now. Fuzzy logics are usually motivated as logics for reasoning with vague and imprecise notions. However, these logics are hardly ever explicitly related to models of processing vague and imprecise information. We look into various problems that arise from insisting on truth functionality for a particular type of fuzzy logic intended to capture reasoning under 'imprecise knowledge'. To address these problems we follow an approach initiated already in the 1970s by Robin Giles, that combines a dialogue game with a scheme for betting on the truth of atomic assertions tested by 'dispersive experiments'. This analysis leads to an alternative characterization of an important interval based fuzzy logic, introduced by Godo, Garcia, and Esteva, but closely related to many similar logics that have been studied recently. Our result can be seen as a justification of truth functionality relative to a particular interpretation of truth value intervals. In any case, it relates concerns about properties of fuzzy logics to reflections on rationality qua playing optimally in a dialogue game intended to model a form of 'approximate reasoning'.

Meaning and Inference in Case of Conflict

Michael Franke
Universiteit van Amsterdam, The Netherlands
mfranke@science.uva.nl

This paper applies a model of boundedly rational "level-k thinking" (cf. Stahl & Wilson 1995, Crawford 2003, Camerer, et al. 2004) to a classical concern of game theory: when is information credible and what shall I do with it if it is not? The model presented here extends and generalizes recent work in game- theoretic pragmatics (Stalnaker 2006, Jäger 2007, Benz & van Rooij 2007). Pragmatic inference is modeled as a sequence of iterated best responses, defined here in terms of the interlocutors' epistemic states. Credibility considerations are a special case of a more general pragmatic inference procedure at each iteration step. The resulting analysis of message credibility improves on previous game-theoretic analyses, is more general and places credibility in the linguistic context where it, arguably, belongs.

Geodesic Conditioning

Konstantinos Georgatos
City University of New York, United States of America
kgeorgatos@jjay.cuny.edu

This paper presents an axiomatization of a class of set-theoretic conditional operators defined by minimization of the geodesic distance generated by the accessibility relation on a frame. The objective of this modeling is to define conditioning based on a notion of similarity generated by degrees of indistinguishability.

Strategies made explicit in Dynamic Game Logic

Sujata Ghosh
Indian Statistical Institute, India
sujata@science.uva.nl

We propose an explicit logic of strategies (SDGL) in the Dynamic Game Logic (DGL) framework and provide a complete axiomatization for this logic. Some discussions are put forward regarding SDGL and DGL, raising an interesting issue about their combination.

From preferences to judgements and back

Davide Grossi
University of Luxembourg, Luxembourg
davide.grossi@uni.lu

Recent results (Dietrich and List, 2007) have shown how Preference Aggregation theorems, such as Arrow's impossibility, can be obtained as corollaries of impossibility theorems concerning the aggregation of judgments. The idea behind this reduction consists in viewing preferences between issues as special kind of judgments, i.e., formulae to which a truth-value is attached. In (Dietrich and List, 2007) the formulae used for representing preferences are first-order formulae of the type xPy ("x is strictly preferred to y") where x, y are variables for the elements in the set of issues of the Preference Aggregation problem. Then, in order for the judgments concerning such formulae to behave according to our intuition of a strict preference relation the three axioms of asymmetry, transitivity and connectedness are added to the Judgment Aggregation framework.

The present paper proposes a different approach to obtain the same kind of reduction. More precisely, preferences will be studied as implicative statements $y \to x$ ("y is at most as preferred as x") in a many-value logic setting. The insights gained in this reduction of preferences to judgments are then used to obtain also an inverse reduction of judgments to special kind of preferences, namely those preferences definable in the Boolean algebra on the support $\{0, 1\}$. To the best of our knowledge, this is the first work advancing a proposal on how to reduce

Judgment Aggregation problems to Preference Aggregation ones. Finally, it is shown how to import an impossibility result from each framework to the other. This shows how the aggregation of preferences and judgments can be viewed as two faces of a same coin.

Evolutionary Game Theory: Natural Selection Implies Interpersonal Comparisons

Till Grüne-Yanoff
University of Helsinki, Finland
gruneyan@mappi.helsinki.fi

In recent years, evolutionary game theory has become an important tool for the social scientist. One apparent advantage that this theory offers is that it models natural selection in a social context. I argue that it indeed allows the modelling of natural selection, yet at a high price: namely, that payoffs are interpersonally comparable. This is an uncomfortable conclusion, as payoffs of evolutionary game models in a social context often must be interpreted as subjective preferences. I show that there is a way to avoid such interpersonal preference comparisons, but this way out also prevents evolutionary game theory from modelling natural selection.

Defining Knowledge in Terms of Belief

Joseph Halpern
Cornell University, United States of America
halpern@cs.cornell.edu
Dov Samet
Tel Aviv University, Israel
dovsamet@gmail.com
Ella Segev
Technion, Israel
esegev@ie.technion.ac.il

The question of whether knowledge is definable in terms of belief, which has played an important role in epistemology for the last fifty years, is studied here in the framework of epistemic and doxastic logics.

Three notions of definability in multimodal logic are considered. Two are analogous to the notions of explicit definability and implicit definability introduced by Beth in the context of first-order logic. However, while by Beth's theorem the two types of definability are equivalent for first-order logic, such an equivalence does not hold for multimodal logics. A third notion of definability, reducibility, is introduced; it is shown that in multimodal logics, explicit definability is equivalent to the combination of implicit definability and reducibility. The three notions

of definability are characterized semantically using (modal) algebras. The use of algebras, rather than frames, is shown to be necessary for these characterizations.

It is shown that if knowledge satisfies any set of axioms contained in S5, then it cannot be explicitly defined in terms of belief. S5 knowledge can be implicitly defined by belief, but not reduced to it. On the other hand, S4.4 knowledge and weaker notions of knowledge cannot be implicitly defined by belief, but can be reduced to it by defining knowledge as true belief. It is also shown that S5 knowledge cannot be reduced to belief and justification, provided that there are no axioms that involve both belief and justification.

Experimental Validation of Quantum Game Theory

Matthias Hanauske
Johann Wolfgang Goethe-Universität Frankfurt, Germany
hanauske@wiwi.uni-frankfurt.de
Steffen Bernius
Johann Wolfgang Goethe-Universität Frankfurt, Germany
Wolfgang König
Johann Wolfgang Goethe-Universität Frankfurt, Germany
Berndt Dugall
Johann Wolfgang Goethe-Universität Frankfurt, Germany

Mathematical concepts from quantum theory have been taken to formulate game theory in a new way. Quantum Game Theory (QGT) — a model merging two different scientific disciplines - has been used to understand and describe cooperative behavior pattern from another perception. The main aim of this article will be an essay to specify and reinterpret the used mathematical concepts of quantum theory. e.g. the non-observable complex wave function which characterizes the quantum state of a quantum system have been used in QGT to describe the players decision state including all possible ways of thinkable cerebration. The complex vector space of the players strategies are spanned with the eigenvectors (pure strategies) of the strategy-Hilbert space. A process of measurement in quantum mechanics can be described with a collapse of the wave function into one of its eigenstates. In QGT the process of decision making enforces the collapse of the player's complex mixed strategy in one of the real measurable pure strategy eigenstates. In contrast to classical game theory, QGT allows an entanglement of the player's strategies. Although no measurable accord is present between the player's strategy choices, the imaginary parts of their strategy wave functions might interact, if their individual states are entangled. This quantum phenomenon might possibly be interpreted as the player's ability to empathize into the other players thinking lanes that could be due to similar historical or cultural background (for example see Physica A 382, 650 (2007)). Players with

strongly entangled strategies appear to act more like a collective state. While a quantum game on a quantum computer has been experimental realized already (e.g. New Journal of Physics 9 (2007) 205) the experimental proof of a meaningful adaptability of QGT in real human games has only been rudimentary successful. The main part of the paper focuses on an experimental validation of QGT. The cooperation percentage of two experimental studies of two-person Prisoner's Dilemma games is compared with the theoretic predictions calculated within classical and QGT. The experimental findings of the cooperation percentage indicate a strong connectivity with the properties of a novel function (N), which depends on the payoff parameters of the game and the entanglement parameter. A classification scheme depending on four quantum cooperation indicators is developed to describe cooperation in real two-person games. The quantum indicators lead to results, which are at least as good as the cooperation predictions derived from classical game theory.

Dynamic Unawareness and Rationalizable Behavior

Aviad Heifetz
The Open University of Israel
heifetz@post.tau.ac.il

Martin Meier
Universitat Autònoma de Barcelona, Spain
martin.meier@iae.csic.es

Burkhard Schipper
University of California at Davis, United States of America
bcschipper@ucdavis.edu

We define generalized extensive-form games which allow for mutual unawareness of actions. Such games comprise of a partially ordered set of game trees together with information sets. A player's information set at a node of a game tree may consist of nodes in another game tree. We extend Pearce's (1984) notion of extensive-form correlated rationalizability to this setting, explore its properties and prove existence. We show by an example that unawareness of an action has different strategic implications from the unavailability of this action. A refinement of rationalizability, prudent extensive-form rationalizability is also defined and its effects are demonstrated. We define the normal form associated with a generalized extensive-form game, and characterize in it extensive-form rationalizability as well as full support rationalizability. While extensive-form rationalizability is equivalent to iterated conditional strict dominance in the normal form, prudent rationalizability corresponds to iterated conditional weak dominance.

A Representation of Time Discounting as Diminishing Intrapersonal Connectedness

Conrad Heilmann
London School of Economics, United Kingdom
c.heilmann@lse.ac.uk

The paper offers a representation of time discounting in intertemporal decisions. Drawing on a multiple-self interpretation of personal identity over time, it is shown that discounting in intertemporal decisions can be represented as diminishing connectedness between different temporal selves within a person. The representation facilitates a discussion of special cases of discounting, namely exponential and hyperbolic discounting. With additional assumptions, the representation can also be used as a normative justification of time discounting. The paper is structured as follows.

Section one discusses the normative status of time discounting in decision theory and economics and rejects several deflationist arguments against a representation of time discounting. It further introduces the concept of pure intertemporal decisions as intertemporal decisions with stable preferences and beliefs and fully compensated risk.

Section two suggests a framework for the representation of discounting by formalizing concepts of personal identity over time, introducing a model of the multiple-self. This account treats different temporal parts of a person as distinct. The model expresses personal identity as connectedness between the selves within the framework of a multiple-self. It contains a connectedness function that measures the strength of the relation between the selves within a person, characterizing the degree of persistence of personal identity. Analyzing intrapersonal connectedness between temporal selves in a multiple-self, two interpretations of connectedness are offered, namely psychological connectedness and empathic access. It is shown that under each interpretation, time discounting can be represented as diminishing connectedness between selves in a multiple-self.

Section three introduces conditions on the connectedness function to represent and discuss the special cases of exponential and hyperbolic discounting. Different sets of conditions on the connectedness function can mimic different methods of discounting: for exponential discounting, these are reflexivity, uniform sequence, reversibility and geometric transitivity. For hyperbolic discounting these are reflexivity, declining sequence, irreversibility and uniqueness.

Section four discusses further assumptions that are necessary and sufficient to use the representation as a normative justification of time discounting. It is concluded that taking the "disuniting metaphysics of personhood" (Broome) of the multiple-self model as base for rational agency permits a defence of discounting, depicting it as diminishing connectedness between selves within a person.

Hyperreal expected utilities and Pascal's Wager

Frederik Herzberg
Universität Bielefeld, Germany
fherzberg@uni-bielefeld.de

This paper proves a hyperreal version of the von Neumann-Morgenstern Theorem and from there deduces a theory of hyperreal expected utility, which subsequently is applied to Pascal's Wager. In particular, we determine additional conditions on the Wagerer's metaphysical views under which one can address both Hajek's dilemma and McClennen's objection through a hyperreal formalisation of the Wager.

Dynamic Awareness

Brian Hill
Université Paris 1, France
brian@brian-hill.org

In recent years, much work has been dedicated by logicians, computer scientists and economists to understanding awareness, as its importance for human behaviour has become more and more evident. Although several logics of awareness have been proposed, little work has been done on change in awareness. However, one of the most crucial aspects of awareness is the changes it undergoes, which have countless important consequences for decision, knowledge and interaction. The aim of this contribution is to make a start on the problem of awareness change. In the first part, a versatile model of awareness is proposed, and a model of awareness change is developed using it. In the second part, logics are given both for awareness and awareness change. The logic of awareness is close to that proposed by Halpern ('Alternative Semantics for Unawareness'. Games and Economic Behavior 37, 2001), which is sound and complete not only with respect to his propositionally determined awareness structures, but also with respect to the generalized standard models of Modica and Ristichini ('Unawareness and partitional information structures'. Games and Economic Behavior 27, 1999). Hence the model of awareness proposed here is comparable to many of the major models proposed in the literature. Two sorts of logic of awareness change are proposed, mimicking two paradigms in the belief revision literature. Postulates are presented in the AGM-style (G??rdenfors P., Knowledge in Flux. MIT Press, 1988), and axioms for dynamic change-of-awareness operators are presented in the DEL-style (Van Benthem J., 'Dynamic Logic for Belief Change'. Journal of Applied Non-Classical Logics 17, 2007).

Public Announcement Logics with Constrained Protocols

Tomohiro Hoshi
Stanford University, United States of America
thoshi@stanford.edu

Public announcement logic (PAL) is a paradigm case of dynamic epistemic logic, which models how agents' epistemic states change when pieces of information are communicated publicly. PAL extends epistemic logic with the operator $[A]$, where the intended reading of $[A]B$ is "After a public announcement that A, B holds." This logic has recently received two improvements. One improvement is to extend PAL with a generalized public announcement operator that allows quantification over public announcements. The other is a semantic setting to model "announcement protocols" to restrict the announcable sequences of formulas, while whatever is true is assumed to be announcable in PAL itself. The purpose of the present paper is to merge these two kinds of improvements. We consider the extension of public announcement logic with the generalized public announcement operator in the semantic setting of restricted announcement protocols.

A Simpler Semantics for Abilities under Uncertainty

Wojciech Jamroga
Clausthal University of Technology, Germany
wjamroga@in.tu-clausthal.de
Thomas Ågotnes
Bergen University College, Norway
tag@hib.no
Wiebe van der Hoek
University of Liverpool, United Kingdom
Wiebe.Van-Der-Hoek@liverpool.ac.uk

Modal logics of strategic ability form one of the fields where logic and game theory can successfully meet. However, defining a good semantics of abilities under imperfect information proved to be a difficult task. In our previous work, we have proposed a logic which –we believe– captures such abilities in an elegant and general way. However, the semantics of the logic is fairly non-standard, as it defines the truth of formulae for sets of states (rather than single states). In this paper, we present a standard, state-based semantics for a relevant subset of the logic. We also discuss explicit operators for complement, union, intersection, and transitive closure of epistemic relations. The resulting language is studied on a variant of the coordinated attack problem.

A geometric approach to judgement aggregation

Christian Klamler
University of Graz, Austria
christian.klamler@uni-graz.at
Daniel Eckert
University of Graz, Austria
daniel.eckert@uni-graz.at

he problem of judgement aggregation consists in aggregating individual judgments on an agenda of logically interconnected propositions into a collective set of judgments on these propositions. This relatively new literature is centred on problems like the discursive dilemma which are structurally similar to paradoxes and problems in social choice theory like the Condorcet paradox and Arrow's general possibility theorem. Saari (1995) has successfully introduced a geometric approach to the analysis of such paradoxes the extension of which to judgment aggregation seems promising.

A Logic for Cooperation, Actions and Preferences

Lena Kurzen
Universiteit van Amsterdam, The Netherlands
lkurzen@science.uva.nl

The concepts of cooperation actions and preferences play central roles in situations in which agents interact with each other. In many cases, agents cooperate and act together as groups because as a group they can achieve something better than they could individually.

If we want to make explicit how agents or groups of agents can achieve certain results, this leads us to the concept of actions, since it is by performing actions that agents interact. Preferences are an important concept in interactive situations because the agents make their decisions based on the preferences they have over possible alternatives.

In this paper, a modal logic for reasoning about cooperation, actions and preferences of agents is developed. It is shown to be sound and complete and the satisfiability problem of its fragment that does not contain strict preferences is shown to be NExpTime-complete.

From belief change to preference change

Jérôme Lang
Université Paul Sabatier
lang@irit.fr

Leon van der Torre
University of Luxembourg, Luxembourg
leon.vandertorre@uni.lu

Various tasks need to consider preferences in a dynamic way. We start by discussing several possible meanings of preference change, and then focus on the one we think is the most natural: preferences evolving after some new fact has been learned. We define a family of such preference change operators, parameterized by a revision function on epistemic states and a semantics for interpreting preferences over formulas. We list some natural properties that this kind of preference change should fulfill and give conditions on the revision function and the semantics of preference for each of these conditions to hold.

Assumption-Incompleteness in Modal Logic

Eric Pacuit
Stanford University, United States of America
epacuit@stanford.edu

Jonathan Zvesper
Universiteit van Amsterdam, The Netherlands
jonathan@illc.uva.nl

We study the notion of assumption-completeness of belief models, which was introduced by Brandenburger. Roughly speaking, a belief model M is assumption-complete for a language L if for every L-definable subset A of M there is a state in M where the player's knowledge is precisely A. Brandenburger and Keisler have argued that assumption-completeness of M with respect to L means that the players whose beliefs are being modeled "have access to" the language L. They have established the negative result that no belief model is assumption-complete for a natural first-order language.

As we show here however, it turns out that there are assumption-complete belief models for the basic modal language. Furthermore, we motivate a conjecture that there are assumption-complete belief models for the more expressive "bounded fragment".

We also define assumption-completeness for more standard epistemic models than the two-sorted belief models that Brandenburger and Keisler use, and briefly discuss the game-theoretical implications of the existence or non-existence of complete belief models.

Disambiguation games in extended and strategic form

Sascia Pavan
University of Eastern Piedmont, Italy
s.pavan@tele2.it

The substance of this work is a new game-theoretic analysis of the capacity humans have to communicate using ambiguous expressions. The background hypothesis is that humans can accomplish these tasks because they are rational creatures, and, when two people are involved in a conversation, they crucially capitalize on this fact, assuming that it is common knowledge. I build on ideas first developed by Prashant Parikh, raising some objections that lead me to modify his models.

I build a game of imperfect information in extensive form, where a hearer and a speaker are the two players, the speaker has some private information, and his task is to convey this piece of information to the hearer. Here lies the main difference between my analysis and Parikh's, since, in his model, the relevant private information of the speaker is the intended meaning of his speech act, while I propose that this private information be identified with some piece of knowledge possessed by the speaker. I argue that my reform renders the theory more natural and conceptually simpler.

The examples I choose as sample cases are simpler to analyze than more general cases, because of the structural features of the resulting model. In the end I retain Parikh's conclusion that speakers tend to focus on efficient equilibria, but I also propose a solution to a problem that has been left open, namely, the strategy adopted by the players when there is not a unique efficient equilibrium. I argue that, in this case, the speaker goes for the unambiguous expression, which is costly, but safe. The argument I use to back both of these tenets hinges on the idea that the players are able to guess the joint strategy they would agree on, were they allowed some preplay communication before the beginning of the game. This kind of argument is not new. It is crucial that the players do not really need to entertain this kind of communication in order to know what would ensue from it.

I also show that the relevant equilibria are plausible even if we consider that a conversation is sequential in nature, proving that they are trembling hand perfect.

Conclusion-Based Procedure for Judgment Aggregation Satisfying Premise Independence

Gabriella Pigozzi
University of Luxembourg, Luxembourg
gabriella.pigozzi@uni.lu

Marija Slavkovik
University of Luxembourg, Luxembourg
marija.slavkovik@uni.lu

Leon van der Torre
University of Luxembourg, Luxembourg
leon.vandertorre@uni.lu

In the doctrinal paradox, a choice has to be made between a premise and conclusion-based procedure. In general frameworks for judgment aggregation, however, the distinction between premise and conclusion does not play a role. We show how the distinction between premise and conclusion can be introduced in a weakened independence assumption for the doctrinal paradox, and we show by an operator that the weakened independence property does not lead to the standard impossibility results. Generalization of the properties for arbitrary judgment aggregation problems is left for further research.

Aggregating collective judgements by selecting disagreements

Giuseppe Primiero
Ghent University, Belgium
Giuseppe.Primiero@UGent.be

The discursive dilemma (or paradox) is well-known from recent literature in social choice, political and economical theories. It raises the problem of the application of majority voting procedures on a set of propositions and the consequent result of aggregating individual judgments into collective decisions. It shows that when majority is applied to a set of premises and to a logically equivalent conclusion, the respective outcomes might be different. In the context of decision-making procedures the dilemma presents a very relevant problem and the usual solutions can be easily criticized: the premise-based approach (Pettit, 2001; Bovens, Rabinowicz, 2006) provides justifications for the collective outcome, but it can be manipulated by taking into account compound formulas rather than literals; the conclusion-based solution (Kornhauser, Sager, 1986, 1993) does not provide explicit justifications for the collective decision. In Pigozzi (2006) an argument-based approach is proposed which uses the technique of belief merging operators from artificial intelligence: it provides a solution by setting integrity constraints that rule out paradoxical outcomes. A different approach to the discursive dilemma is here given based on the Adaptive Logic $\mathbf{ADM}^{c}$ (Adaptive

Doxastic Merging by Counting). By applying the adaptive strategy of $\mathbf{ADM}^{c}$ to a premise set mimicking the case of the dilemma, one formulates a resulting consequence set by selecting the disagreements arising from the agents' belief bases. It will be shown how this strategy is successful in avoiding the inconsistency result of the dilemma. In order to formulate a selection procedure whose result is not only non-paradoxical but which can also support a decision, additional information can be provided in the premise set, by focusing on individual preferences towards models of the collectively admissible judgments.

Towards a Concrete Semantics for Announcements

Riccardo Pucella
Northeastern University, United States of America
riccardo@ccs.neu.edu

Mehrnoosh Sadrzadeh
Université Paris 7, France
mehrs@comlab.ox.ac.uk

The story of our paper is inspired by the existence of two schools of thought in modeling and reasoning about information flow in interactive multi-agent systems where agents acquire new information after communicating with each other. One of these schools is the interpreted systems of Fagin et al. and the other is the dynamic epistemic logic (DEL) approach of van Benthem et al. The former is more algorithmic and inclined towards applications to distributed computing and artificial intelligence, whereas the latter is more logical and focuses on formalizing concepts such as rational agents. One key difference between the two schools is the semantic models used to represent a situation of interest, about which we may want to reason. Fagin et al. use a very concrete semantics based on sequences of local states, which represent the possible executions of the system under consideration and are referred to as runs. These can be obtained from an algorithmic description of the system in question. Whereas DEL uses a more abstract semantics in terms of Kripke structures of states and also actions involved in a scenario of interest. These structures are equipped with an update product that constructs new Kripke structures of states and is used to reason about the evolution of states, caused by public and private announcement actions and referred to as epistemic actions.

The aim of our paper, which is very much work in progress, is to bring together these two schools by developing an interpreted systems semantics for DEL. A full development would involve devising a model for public and private announcements, which can be both honest and dishonest. For the time being, however, we only present a system for logics of honest public and private announcements, leaving the dishonest announcements for future work. Our development consists of

1. A set of axioms describing runs of a system with public and private announcements. The most important of these are about registering the announcements in the states of the execution of the system, e.g. a private announcement is registered in the local state of the environment and the local states of the agents who receive the announcement, whereas the agents who do not receive the announcement get a neutral action (tick) registered in their the local state. A public announcement is treated as a private announcement to all the agents of the system.

2. An LTL-style logic enriched with epistemic modalities; this logic includes an atomic proposition that stands for 'an announcement is just being made to a sub-group of agents'. This proposition holds at a global state of the system at time m iff the action of announcement has been registered (according to our axioms) as the last action of the global state at time $m + 1$.

3. A translation from dynamic epistemic logic to our logic, showing that we can interpret the existing logics of announcements in our setting. For a system with only public announcements the knowledge modality of DEL is translated to our epistemic modality and the dynamic modality of DEL to an implication formula whose antecedent is the announcement proposition described above. For a system with private announcements, we introduce a new belief modality and use it to translate the (now possibly wrong) epistemic modality of DEL.

We show that the runs of a system with public and private announcements have an interesting property, namely that the indistinguishability of states to an agent at time $m + 1$ implies their indistinguishability at the previous time m. This property becomes handy in proving that our translation is sound, that is, its image indeed forms a DEL. It is not hard to see that our logic is also sound and complete with respect to the particular class of interpreted systems we consider.

We believe that our semantics, because of its more concrete nature, provides new insights into the nature of dynamic epistemic logic, and will allow us to study announcements in real systems, as well as more easily accommodate extensions (such as fact-changing actions and parallel announcements), explore design space for informative belief in the context of cheating and lying, and calculate the complexity of logics of public and private announcements. Our semantics in general, and the calculation of its complexity in particular, should go along recent interesting results of van Benthem and Pacuit ("The tree of knowledge in action"). Studying these connections is also left to future work. One nice feature of our approach is that we use standard primitives, such as time and knowledge, with an S5 interpretation for knowledge that comes from a completely natural accessibility relation for every agent, namely indistinguishability of the agent's local state. Using these relations, we define a new modality, a form of belief, that seems to be required to capture the state of uncertainty of agents as a result of

private announcements. This new modality helps explain the effect of private announcements on the knowledge of agents.

Finally, note that the relationship between our logic and DEL is similar to the relationship between temporal logic and dynamic logic: endogenous versus exogenous logics; difference in complexity; often interested in reasoning about a particular protocol, rather than all protocols. They indeed follow same ideas, but have different points of view.

Towards an Implementation Theory via a Game Logic Approach

Davi Romero de Vasconcelos
Pontifcia Universidade Católica do Rio de Janeiro, Brazil
daviromero@ufc.br

Edward Hermann Haeusler
Pontifcia Universidade Católica do Rio de Janeiro, Brazil
hermann@inf.puc-rio.br

Recently, the general approach of Game Theory, which is to find the set of outcomes (the equilibria) of a solution concept in a game, has been modeled by means of Game Logics, in which a game and a solution concept are related to models and formulas of a Game Logic. Thus, the general problem can be stated as a model-checking problem. Moreover, the set of solutions can be found automatically by using model-checkers of the respective Game Logic. The implementation problem in Game Theory focus on the inverse problem, i.e., finding a game that yields a given set of outcomes as equilibria of a solution concept. In this work, we relate the implementation problem to a model checking problem of an adequate Game Logic. Specifically, we illustrate the relationship between the implementation problem of extensive games with perfect information and the model checking problem of GAL (Game Analysis Logic). As a consequence, we can benefit of the use of model-checkers in order to solve the implementation problem as well. It is worth mentioning that the approach used in this article seems to be adequate for other Game Logics as well.

Intentions and transformations of strategic games

Olivier Roy
Rijksuniversiteit Groningen, The Netherlands
o.roy@rug.nl

In this paper we take a game-theoretic perspective to study the effects of previously adopted intentions in rational decision making. We investigate the question of how agents transform the decision problems they face in the light of what they intend, and provide conditions under which such transformations, when iterated, leave room for deliberation, i.e., do not exclude all the options of the decision maker.

Agreeing to disagree:The non-probabilistic case

Dov Samet
Tel Aviv University, Israel
dovsamet@gmail.com

A non-probabilistic generalization of Aumann's agreement theorem is proved. Early attempts at such a theorem were based on the sure thing principle which assumes intrapersonal-interstate comparison of knowledge. Unfortunately, such comparisons are impossible in partitional knowledge structures. The theorem proved here is based on a new version of the sure thing principle that makes interpersonal-intrastate comparison of knowledge.

S5 Knowledge without partitions

Dov Samet
Tel Aviv University, Israel
dovsamet@gmail.com

We study set algebras with an operator (SAO) that satisfy the axioms of S5 knowledge. A necessary and sufficient condition is given for such SAOs that the knowledge operator is defined by a partition of the state space. SAOs are constructed for which the condition fails to hold. We conclude that no logic singles out the partitional SAOs among all SAOs.

Persuasion and Limited Communication

Itai Sher
University of Minnesota, United States of America
isher@umn.edu

This paper studies persuasion as a game. A speaker must decide which arguments to present and a listener which arguments to accept. Glazer and Rubinstein (2006) showed that a listener's optimal persuasion rule is credible in the sense that commitment has no value for the listener. This paper provides an algorithm for finding an optimal persuasion rule in persuasion problems without time constraints along with the speaker's strategy in the credible implementation of the rule. The credibility result, as well as the algorithm, is extended to the case where the listener's decision is not binary under a concavity assumption. Qualitative properties of optimal rules are derived. In particular, in the absence of time constraints there exists an optimal rule that treats equivalent evidence equivalently whereas this may fail with time constraints. All of the results depend on a relation between the persuasion problem and the maximum flow problem, a well-known combinatorial optimization problem.

Quantified Logic of Awareness and Impossible Possible Worlds

Giacomo Sillari
University of Pennsylvania, United States of America
gsillari@sas.upenn.edu

Among the many possible approaches to the problem of logical omniscience, I consider here awareness and impossible worlds structures. The former approach, pioneered by Fagin and Halpern, distinguishes between implicit and explicit knowledge, and renders the agents immune from logical omniscience with respect to explicit knowledge. The latter, first developed by Rantala, allows for the existence of logically impossible worlds to which the agents are taken to have 'epistemological' access; since such worlds do not behave consistently, the agents's knowledge is fallible relative to logical omniscience. The two approaches are known to be equally expressive in the propositional case with Kripke semantics. This paper contains two contributions. First, I show that the two approaches are equally expressive in the propositional case with neighborhood semantics. Second, I provide predicate systems of both awareness and impossible worlds structures and prove them to be equally expressive.

Preference Logic and Its Measurement-Theoretic Semantics

Satoru Suzuki
Komazawa University, Japan
bxs05253@nifty.com

How can we explain why we should ascribe some logical properties to an agent's preferences, and why we should not ascribe others to them? Although this problem is very important, little attention has been given to it. We break this problem down into the three:

(Problem 1) When we attempt to ascribe some logical properties to the preferences of an agent, how can we explain this ascription?

(Problem 2) What logical properties should we ascribe to his preferences when and only when our solution to Problem 1 can be adopted?

(Problem 3) What logical properties should we ascribe to his preferences and what logical properties should not we ascribe to them?

The aim of this talk is to propose a new version of sound and complete preference logic (PL) that can furnish a solution to all of these problems. We solve all of them by providing PL with a Domotor-type semantics that is a kind of measurement-theoretic and decision-theoretic one. If an agent's belief state can be represented by a probability function, and his desire state can be represented by a utility

function, and his preferences should correspond to expected utility maximisation, then in terms of his beliefs and desires, we can explain ascribing some logical properties to his preferences, which can furnish a solution to Problem 1.

In this Domotor-type semantics, by virtue of a representation theorem, when and only when our solution to Problem 1 can be adopted, we should ascribe (connectedness) and (projectivity) to an agent's preferences, which can furnish a solution to Problem 2. We define a model of PL and propose a proof system of PL. We then show that (reflexivity), (transitivity) and (impartiality) are all provable in PL, and that neither (contraposition), (conjunctive expansion), (disjunctive distribution) nor (conjunctive distribution) is provable in PL, but that (restricted contraposition), (restricted conjunctive expansion), (restricted disjunctive distribution) and (restricted conjunctive distribution) are all provable in PL, which can furnish a solution to Problem 3.

A Note on the Stability of Voting Games

Bertrand Tchantcho
University of Yaoundé I, Cameroon
btchantcho@yahoo.fr
Lawrence Diffo Lambo
University of Yaoundé I, Cameroon
yldiffo@yahoo.fr

The models of vote in this work are devoted to bill adoption contexts. In such a context, a bill is considered by the community. Assume for instance that the issue is a referendum on whether or not to amend the constitution. The constitution in force (say y) is called the status quo. A projected constitution (say x) is a candidate to the replacement of y. Enforcing the bill consists of adopting the replacement, while rejecting the bill consists of discarding the candidate x. So, there are two options for the collective decisions : bill adopted, or bill rejected.

Simple games, also called yes-no system, introduced by von Neumann (1944) assume that a coalition of voters is either decisive (otherwise called winning, in which case it holds the entire power of decision), or losing (in which case it is absolutely powerless). In such model, there are no provision for abstentions, and abstentions (if any occurs) have effects that are equivalent to votes against the bill.

In the well known relative majority voting, the collective decision enforces the bill if the number of voters who vote for the bill is greater than the number of those who vote against. Social decision systems (Sds), owed to Rubinstein (1980), are a generalization of the relative majority voting. In addition to the two voting options (for or against), a voter can also abstain. The casting of all individual votes thus realizes a partition of the set of all voters into three subsets. Those who abstain, those who vote for and those who vote against the bill. Decisiveness is therefore described using couples of disjoints coalitions.

We study the core of Sds, with emphasis on the characterization of its non-emptiness for every preference profile of voters. Such a characterization for simple games was obtained by Nakamura (1979). Herein, we obtain a necessary and sufficient condition for which the core of an Sds is non-empty regardless of the profile of individual preferences, and we show that our result is a generalization of the Nakamura's theorem. As a corollary, we show that the relative majority voting is not stable, provided that there are at least three voters.

Reasoning about strategic games with hybrid logic of choice and preferences

Nicolas Troquard
University of Liverpool, United Kingdom
Nicolas.Troquard@liverpool.ac.uk

Wiebe van der Hoek
University of Liverpool, United Kingdom
Wiebe.Van-Der-Hoek@liv.ac.uk

Michael Wooldridge
University of Liverpool, United Kingdom
mjw@csc.liv.ac.uk

Solutions concepts are a fundamental tool for the analysis of game-like systems, and as a consequence, much effort has been devoted to the problem of characterising them using logic. However, one problem is that, to characterise solution concepts such as Nash equilibrium, it seems necessary to refer to strategies in the object language, which tends to complicate the object language. We propose a logic in which we can formulate important properties of games (and in particular pure-strategy solution concepts) without recourse to naming strategies in the object language. The idea is that instead of using predicates which state that a particular collection of strategies forms a solution, we define formulae of the logic that are true at a state if and only if this state constitutes a particular equilibrium outcome. We demonstrate the logic by model checking equilibria of strategic games.

Building bridges between dynamic and temporal doxastic logics

Johan van Benthem
Universiteit van Amsterdam, The Netherlands
johan@science.uva.nl
Cédric Dégremont
Universiteit van Amsterdam, The Netherlands
cdegremo@science.uva.nl

In this paper, we compare two modal frameworks for multi-agent belief revision: dynamic doxastic logics computing stepwise updates and temporal doxastic logics describing global system evolutions, both based on plausibility pre-orders. We prove representation theorems showing under which conditions a doxastic temporal model can be represented as the stepwise evolution of a doxastic model under successive 'priority updates'. We define these properties in a suitable doxastic-temporal language, discuss their meaning, and raise some related definability issues.

Simulation and information: quantifying over epistemic events

Hans van Ditmarsch
University of Otago, New Zealand
hans@cs.otago.ac.nz
Tim French
University of Western Australia, Australia
tim@csse.uwa.edu.au

We introduce a multi-agent logic of knowledge with time where $F\varphi$ stands for 'there is an informative event after which φ'. Formula $F\varphi$ is true in a model iff it is true in all its refinements (i.e., 'atoms' and 'back' are satisfied; the dual of simulation). The logic is 'almost' normal, and positive knowledge is preserved. The meaning of $F\varphi$ is also "after the agents become aware of new factual information, φ is true", and on finite models it is also "there is an event model (U, e) after which φ". The former provides a correspondence with bisimulation quantifiers in a setting with epistemic operators.

Multi-Agent Belief Revision with Linked Plausibilities

Jan van Eijck
Centrum voor Wiskunde en Informatica, The Netherlands
jve@cwi.nl

Floor Sietsma
Universiteit van Amsterdam, The Netherlands
floorsietsma@gmail.com

Van Eijck and Wang (unpublished) show how propositional dynamic logic (PDL) can be interpreted as a logic of belief revision that extends the logic of communication and change (LCC) given by van Benthem, van Eijck, and Kooi (Information and Computation 2006). This new version of epistemic/doxastic PDL does not impose any constraints on the basic relations and because of this it does not suffer from the drawback of LCC that these constraints may get lost under updates that are admitted by the system. Here, we will impose one constraint, namely that the agent's plausibility relations are linked. Linkedness is a natural extension of local connectedness to the multi-agent case and it assures that we know the agent's preferences between all relevant alternatives.

Since the belief updates that are used by van Eijck and Wang may not preserve linkedness, we limit ourselves to a particular kind of belief change that does preserve it.

Our framework has obvious connections to coalition logic and social choice theory. We show how we can use it to model consensus seeking in plenary Dutch meetings. In Dutch meetings, a belief update is done for all agents in the meeting if a majority beliefs the proposition that is under discussion. A special case of these meetings is judgement aggregation, and we apply our framework to the discursive dilemma in this field.

Revealed preferences and satisficing behavior

Robert van Rooij
Universiteit van Amsterdam, The Netherlands
R.A.M.vanRooij@uva.nl

A much discussed topic in the theory of choice is how a preference order among options can be derived from the assumption that the notion of 'choice' is primitive. Assuming a choice function that selects elements from each finite set of options, Arrow (1959) already showed how we can generate a weak ordering by putting constraints on the behavior of such a function such that it behaves as a utility maximizer. Arrow proposed that rational agents can be modeled by such choice functions. Arrow's standard model of rationality has been criticized in economics and gave rise to approaches of 'bounded rationality'. Two standard assumptions of rationality will be given up in this paper. First, the idea that agents are utility 'optimizers' (Simon). Second, the idea that the relation of 'indifference' gives rise to an equivalence relation. To account for the latter, Luce (1956) introduced semi-orders. Extending some ideas of Van Benthem (1982), we will show how to derive semi-orders (and so-called interval orders) based on the idea that agents are utility 'satisficers' rather than utility optimizers.

Lunch Suggestions

1. **Euclides Cafeteria**, P Building, Plantage Muidergracht 24, On the ground floor of the conference building. Small cafeteria offering simple sandwiches and drinks. **Not open on Saturday**.

2. Mensa **"Restaurant Agora"**
   H Building, Roetersstraat 11
   Typical Dutch student mensa offering sandwiches and a few warm dishes. **Not open on Saturday**.

3. **Sapori del Mondo**
   Plantage Middenlaan 30a
   Somewhat more upmarket sandwiches. During lunchtime queues can be expected and service rather slow.

4. **Plantage Lunch**
   Plantage Middenlaan 30b
   Turkish take-away selling sandwiches and various warm slices.

5. Bakery **Ed Rooth**
   Plantage Kerklaan 14

6. **Plancius**
   Plantage Kerklaan 61
   Nice place to sit down opposite the Artis zoo.

7. **Café Koosje**
   Plantage Middenlaan 37
   From sandwiches to more complete meals. Open for lunch and dinner. Reasonably priced.

8. **Meneer Nilsson**
   Plantage Kerklaan 41
   Mediterranean café open for lunch and dinner.

9. **Mediacafé Plantage**
   Plantage Kerklaan 36
   Right next to the entrance of the zoo. This large café offers sandwiches and full meals. Open for lunch and dinner. No take-away.

10. **De Groene Olifant**
    Sarphatistraat 510
    Quite a nice pub. Serves pub meals.

11. **Coffee to walk**
    Plantage Middenlaan 44
    Coffee to go.

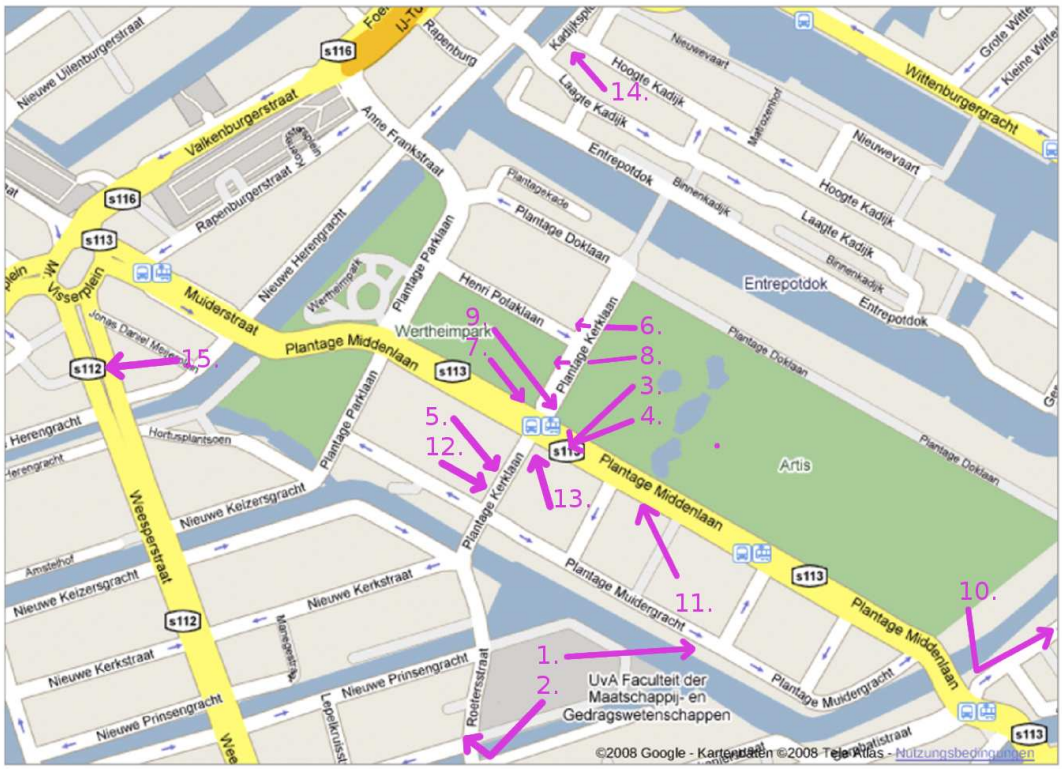

Dinner Suggestions

The following website has dinner suggestions and reviews:

`http://www.iens.nl/english/restaurantsIn/Amsterdam/`

In the following, we give a number of recommendations. If you want to stay close to the conference site, some of the places listed for lunch, also serve dinner (Plancius, Meneer Nilsson, to name a few). The restaurants listed as (12) to (15) are also in the Plantage neighbourhood (see lunch map) and can be reached on foot in less than ten minutes from the conference site. The other restaurants are roughly ordered by distance from the conference location; all of them are in walking distance, numbers (16) to (22) reachable in less then 30 minutes. The restaurant *De Bolhoed* (23) is not on the map; it is on the northwest part of *Prinsengracht* (about 40 minutes walk, or take tram number 14 from *Plantage Badlaan* to *Westermarkt*).

12. **Kerklaan** (Indian / Surinamese), Plantage Muidergracht 69, Phone (020) 421 49 39

    Mostly a take-out and delivery place, but has a number of tables and decent food.

13. **Taman Sari** (Indonesian), Plantage Kerklaan 32, Phone (020) 623 71 30

    Small and very unpretentious Indonesian restaurant.

14. **Koffiehuis van de Volksbond** (International)
    Kadijksplein 4, Phone (020) 622 12 09, kitchen open 18.00-22.00
    Simple and good restaurant without pretence. Warm atmosphere. Reasonably priced.

15. **Asmara** (African), Jonas Daniël Meijerplein 8, Phone (020) 627 10 02

    East African cuisine without cutlery.

16. **Bird** (Thai), Zeedijk 77, Amsterdam, Phone (020) 420 62 89, kitchen open 3.00pm-10.00pm

    Good Thai food on the street parallel to the Red Light District.

17. **'t Tuinfeest**, Geldersekade 109hs, Amsterdam, Phone (020) 620 88 64 / Fax (020) 620 88 64, kitchen open 5.00pm-10.30pm

    Bar food (Great place to eat when the weather is nice and you can sit outside)

18. **Oriental City** (Chinese), Oudezijds Voorburgwal 177-179, Phone (020) 626 83 52 / Fax (020) 626 92 95, kitchen open 11.30am-10.30pm

    Don't let the touristy exterior scare you away. Good for larger groups. If feeling adventurous, ask the waiter/waitress to recommend dishes from the Chinese menu

19. **Bubbles and Wine** (Champagne and Wine Bar), Nes 37, Phone (020) 422 33 18, Kitchen open 3:30-"onwards"

    50 wines by glass, 180 bottles–good place to drink and eat lightly and take in the scene

20. **Krua Thai Classic** (Thai), Staalstraat 22, Phone (020) 622 95 33

21. **Koh-I-Noor** (Indian), Rokin 18, Phone (020) 627 21 18, kitchen open 5.00pm-11.30pm

    Very good Indian food.

22. **Fromagerie Crignon Culinair** (French/Fondue), Gravenstraat 28, Phone (020) 624 64 28 / Fax (020) 646 10 34 kitchen open 6.00pm-9.30pm, cash only

    This is a small place, so reserve in advance if you are in the mood for a good cheese fondue. Their specialty is an eight cheese fondue

23. **De Bolhoed** (Vegetarian), Prinsengracht 60, Phone (020) 626 18 03, kitchen open noon-10.00pm

    Good vegetarian food. In a beautiful neighborhood

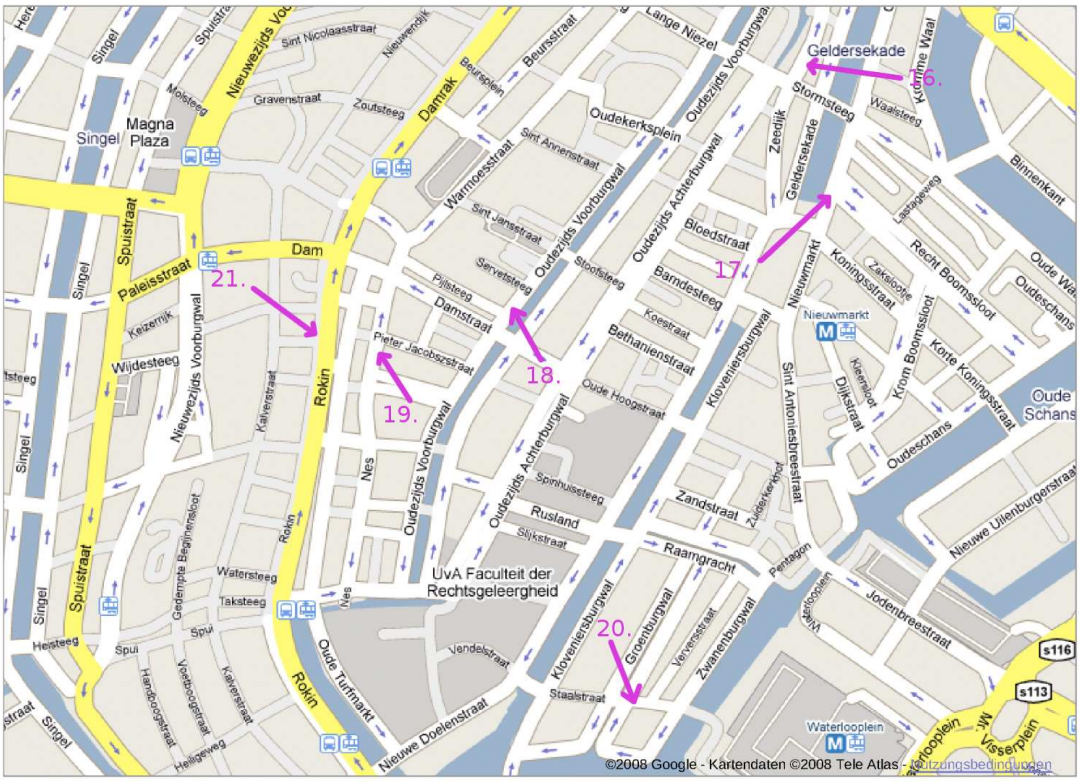
Geldersekade
6.
Magna Plaza
Dam
21.
17.
18.
19.
20.
UvA Faculteit der Rechtsgeleergheid
Oudezijds Voorburgwal
Oudezijds Achterburgwal
Damrak
Rokin
Nieuwezijds Voorburgwal
Spuistraat
Paleisstraat
Nieuwmarkt
Oude Schans
Waterlooplein
Nieuwe Doelenstraat
Oude Turfmarkt
Raamgracht
Zandstraat
Rusland
Oudekerksplein
Nieuwe Uilenburgerstraat
Jodenbreestraat
s116
s113
©2008 Google - Kartendaten ©2008 Tele Atlas - Nutzungsbedingungen