

Advice and Trust in Decision Making (ATDM)

June 29-30 2006

A two day international meeting on theoretical and applied research
into factors affecting advice taking and trust in advisors.

University College London, Department of Psychology.

Ground Floor Lecture Theatre (GO3), 26 Bedford Way, London WC1E 6BT.

With thanks to:

The Economic and Social Research Council (ESRC)
The ESRC Centre for Economic Learning and Social Evolution (ELSE)
The European Association for Decision Making (EADM)
UCL Department of Psychology

Thursday June 29th

8.00 – 9.00 *Conference registration*

Self, other, and similarity

9.00 – 9.30 Nigel Harvey
Using advice to judge risk acceptability for self and others.

9.30 – 10.00 Ilan Yaniv
Receiving advice on matters of taste: How relevant are other peoples' preferences?

10.00 – 10.30 Barbara Fasolo & Annika Wallin
Help! I need somebody, Help! Not just anybody: A field study of agent selection.

10.30 – 11.00 Ilan Fischer
On the rationality of irrational trust.

11.00 – 11.30 *Coffee break*

Context and presentation of advice

11.30 – 12.00 Joachim Meyer
Task difficulty and reliance on potentially costly advisory information.

12.00 – 12.30 Clare Harries
How does risk format affect the way people take advice?

12.30 – 2.00 *Lunch & Poster session*

Social judgment and emotions

2.00 – 2.30 J. Richard Eiser
Trust, decision-making and the social judgement of risk.

2.30 – 3.00 Mathew P. White
Trust in the police following (suspected) terrorist incidents: A social judgement approach.

3.00 – 3.30 Liat Hadar
Decision making and social emotions.

3.30 – 4.00 Peter Taylor-Goodby
The efficiency/trust dilemma in public policy reform: How the new public policy can deliver better services but undermine trust at the same time.

4.00 – 4.30 *Coffee break*

Credibility, calibration, and competence

- 4.30 – 5.00 Michael Birnbaum
Decision-making based on advice from sources who vary in expertise and bias.
- 5.00 – 5.30 Thomas S. Wallsten
The effect of advice from sources differing in calibration.
- 5.30 – 6.00 Helmut Jungermann
The effect of advisors' competence and self-interest in a dyadic decision situation.
- 7.00 pm London Eye & Conference Dinner

Friday June 30th

Statements of trust and use of advice

- 9.00 – 9.30 Lyn Van Swol
The effects of advisor motives on confidence and advice utilization.
- 9.30 – 10.00 Gideon Keren
Incompatible messages implied from positive and negative frames regarding trust and choice decisions.
- 10.00 – 10.30 Matt Twyman
Stated and revealed trust in advice.
- 10.30 – 11.00 *Coffee break*

Taking advice from machines

- 11.00 – 11.30 Peter Ayton
Diverse effects of computerised advice on expert decision making.
- 11.30 – 12.00 Paul Goodwin
Accepting advice from a machine: the role of statistical advice in forecasting support systems.
- 12.00 – 12.30 Philip Bonhard
"Thanks for the advice, but do I know you?": Improving decision-making support of recommender systems through social networking.
- 12.30 – 2.00 *Lunch*

Weighting advice

- 2.00 – 2.30 Elise Weaver
A new measure of relative change of importance weights when given advice.
- 2.30 – 3.00 Stian Reimers
How do we know if people really use advisors' track records: Regression weights and optimal performance.

3.00 – 3.30 Derek J. Koehler
Illusion of confirmation from exposure to another's hypothesis.

3.30 – 4.00 *Coffee break*

Models and analyses

4.00 – 4.30 Robert D. Sorkin
Signal detection analysis of advice behavior.

4.30 – 5.00 David V. Budescu
To Bayes or not to Bayes? A comparison of two classes of models of information aggregation.

5.00 – 5.30 Timothy Earle
Advice and trust in risk management contexts: The TCC model perspective.

PAPER ABSTRACTS:

Self, other, and similarity

Nigel Harvey, Matt Twyman and Clare Harries – University College London

Using advice to judge risk acceptability for self and others.

People used advice from two sources to assess the risk associated with various activities and then estimated the likelihood that both they and a friend similar to or different from them would engage in these activities. For activities that people tended to avoid, they judged likelihood of avoidance to be greater for themselves than for others different from them but not for others similar to them. For activities that people tended to engage in, no self–other differences were found. When people rated the activities for social acceptability and attractiveness, self–other differences were again greater for activities that people tended to avoid. However, no such effect was found for ratings of controllability. Results can be interpreted in terms of motivated reasoning, via purely cognitive mechanisms, or by arguing that empathy gaps depend on self–other differences.

Ilan Yaniv – The Hebrew University of Jerusalem

Receiving Advice on Matters of Taste: How Relevant Are Other Peoples' Preferences?

We distinguish between advice on objective matters of belief (e.g., forecasts) and on subjective matters of taste (e.g., preferences). Two fundamental questions of theoretical and practical importance are: How relevant is one person's preference to another person's decision? Under what conditions might people benefit from such advice? I will review two important strategies, similarity and aggregation, and a trade-off between them. We conducted several experiments, in which participants made realistic choices based on advice. They could follow either similarity-based or majority-based recommendations. Decision makers relied heavily on similarity rather than majority of opinion, although this preference varied as a function of discriminability. We also conducted simulation analyses testing the validity of the similarity and the aggregation strategy. The results of these simulations demonstrate that, as discriminability of tastes decreases, the aggregate advice increases in accuracy. I will discuss our theoretical framework and results in relation to social comparison theories and corrective procedures in decision making.

Barbara Fasolo¹ & T. Annika Wallin²

¹London School of Economics

²Lund University

Help! I need somebody, Help! Not just anybody: A field study of agent selection.

There are many real world situations in which we need to appoint somebody else to choose on our behalf. Because we are the ultimate consumers of this surrogate choice, it is crucial that this person is 'not just anybody,' but can predict our preferences well. For instance, it is common to appoint a real estate agent to scan the housing market for us. Choice of a good agent is important, as we only get to buy or rent one of the houses the agent predicts we like.

In this talk we review the scant available research on consumers' ability to select decision agents and describe a field study conducted in a Berlin cafe' where we asked real customers to appoint one of their friends as 'decision agent'. This study allowed us to address questions pertaining to the process of agent selection [How do we choose an agent? What are the different cognitive processes that underlie agent selection?] as well as the outcome of agent selection [How well do we choose an agent? Are we able to select as agent the person that is best able to predict our preferences?].

Our preliminary results point to a general ability in selecting the best out of two prospective agents. We describe three possible mechanisms that concur to explain how our participants chose their agents: 1) exposure assessment (people choose as agent the friend who had a greater chance to observe him or her ordering a drink in the past), 2) perceived reciprocity (people choose as agent the friend towards whom s/he expects to make more accurate predictions), and 3) similarity assessment (people choose as agent the friend whom they consider more similar to themselves).

Ilan Fischer- University of Haifa

On the rationality of irrational trust.

For over half a decade the Prisoner's Dilemma (PD) has been used as a primary model for the study of competitive behaviors and the evolution of cooperation. However, no theory has yet been capable of predicting or explaining why many participants cooperate and trust their opponents, even when engaging in single (non-repeated) PD interactions.

The current study introduces Subjective Expected Relative Similarity (SERS), which provides both a normative solution and a descriptive theory. SERS explains under what conditions trust and cooperation actually increase expected payoffs in single step PD interactions. SERS shows that cooperation is directly related to opponent's expected similarity and inversely related to the games' critical similarity, an index that is calculated from the game's payoff matrix.

SERS is experimentally tested in three experiments that manipulate the perception of opponent's similarity prior to the participation in a single Prisoner's Dilemma game. The findings are also extended to the iterated game, showing that the same principle may be developed into a strategy that, under some conditions, outperforms strategies like Tit For Tat and Win-stay, lose-shift.

SERS provides a powerful tool for predicting competitive and cooperative behaviors. It further points to the necessity of revising social and economic thought by integrating the perception of similarity into rational decision making models.

Context and presentation of advice

Joachim Meyer- Ben Gurion University of the Negev

Task difficulty and reliance on potentially costly advisory information.

Predicting how users will respond to information from advisory systems remains one of the major challenges in the design of such systems. Two experiments dealt with the question how the effort required to obtain task relevant information and the diagnostic value of an alerting system affect performance in a signal detection task that is aided by binary alerts. In a simulated production task participants had to decide whether to produce or not, based on the length of a rectangle that represented the quality of raw material and on the output of a binary alerting system. The experimental conditions differed in the contrast of the rectangle from the background (high vs. low contrast) and in the diagnostic value of the warning (non-valid or valid alerts). Results showed non-optimal weighting of the warning information. When the perceptual task was easy (i.e., a high contrast rectangle appeared) information from valid alerts was used less, and information from non-valid alerts was used more, than when the task was difficult. This led to farther-from-optimal performance in the easier perceptual task than in the more difficult task. The results demonstrate that task characteristics and the context in which advisory information is received may have intricate effects on the responses to this information.

Clare Harries, Nigel Harvey and Matt Twyman - University College London

How does risk format affect the way people take advice?

Presenting risks in different formats can make the information easier or harder to understand. People are better at processing relative frequencies over probabilities, and direct comparison of two risks is improved when they have the same denominator, rather than the same numerator. Feedback about the accuracy of advisors' estimates of risk in some situations (e.g. probability of death from certain activities) will come as summary statistics, and can be compared directly to the statistical advice from advisors. In other situations (e.g. probability of winning football matches), whilst advisors may estimate risk in summary formats, the decision maker will experience the probabilities as cumulative experiences. In two experiments we investigated the effect of risk format on peoples' ability to combine advice and learn about the accuracy of advisors. In the first experiment advice about risk of death from various activities was presented and estimated as percentages or as relative frequencies with common denominators or common numerators. Errors in risk estimates were significantly lower for participants seeing information in percentage formats. Actual trust in advisors was only higher in better advisors when risk was presented as percentages. In a second experiment participants took advice about the probability of a football team winning, either a game at a time, or as a summary statistic. They received feedback that either matched the format of advice or did not.

Social judgment and emotions

J. Richard Eiser- University of Sheffield

Trust, Decision-Making and the Social Judgement of Risk.

Trust is important when decision-makers rely on information from others under conditions of uncertainty, and uncertainty matters where it involves risk. Whereas risk is commonly seen as a property of natural or human-made physical systems, I shall argue that risk arises from an interaction between such systems and human decision-making. Individuals' decisions can affect the level of risk for both themselves and others, so risk is partly a social product. Hence, judgement of risk is (or should be) a social judgement, involving consideration of the quality of our own and others' decisions and success or failure at predicting and controlling a hazard and mitigating its consequences. A framework will be presented for conceptualising perceived decision-making quality, based on an extension of Signal Detection Theory, with implications for factors that can increase or decrease trust. However, a constraint on such

judgement processes is that feedback on the consequences of decisions can be unevenly distributed in individuals' personal experience. On the one hand, unsafe decisions may not always be punished; on the other hand, risk aversion (false alarms) may appear to be reinforced by avoidance of feared consequences, even if these consequences would not actually have occurred. Trust in, and acceptance of advice from, others is thus likely to be constrained by its consistency with personal experience.

Mathew P. White¹ & J. Christopher Cohrs²

¹ *University of Sheffield*

² *Friedrich-Schiller University*

Trust in the police following (suspected) terrorist incidents: A social judgement approach.

The "Police Officer's Dilemma" - whether or not to shoot a potential suspect - has increasingly serious implications with the terrorist risk posed by suicide bombers. Although two tragic cases of innocent suspects being shot have occurred, we know little about how these events affect people's perceptions of the police, in particular the extent to which they are trusted to deal with the terrorist threat. We investigated these issues using two scenario experiments where outcomes were manipulated according to Signal Detection Theory. Study 1 (N = 164) found higher trust following correct (Hits & All Clears) than incorrect judgements (False Alarms & Misses) and following decisions to Shoot (Hits and False Alarms) than Not Shoot (All Clears and Misses). Trust in the police did not fall significantly following a False Alarm, suggesting a general tolerance of such errors. Study 2 (N = 156) replicated these patterns but found that this tolerance was moderated by Right Wing Authoritarianism (RWA). Those high in RWA were more tolerant of False Alarms than those low on RWA. The findings help explain the heterogeneity of public reactions to such events and aid our understanding of how trust in key decision makers is built and lost.

Liat Hadar - UCLA

Decision making and social emotions.

A growing body of empirical evidence suggests that emotions play an important role in decision making. For example, it has been demonstrated that fearful individuals tend to be more risk averse whereas angry people tend to be more risk seeking (e.g., Lerner & Keltner, 2001). In my talk I will review these findings and suggest a framework for understanding the interaction between emotions and prospect theory's value- and weighting-functions. I will focus on social emotions (e.g., accountability and empathy) and present empirical data that demonstrate their impact on probability judgments and on the weighting function. Finally, I will discuss the implications for advice judge-advisor systems.

Peter Taylor-Goodby - University of Kent

The Efficiency/Trust Dilemma in Public Policy Reform: How the new public policy can deliver better services but undermine trust at the same time.

Government is spending more on key social services and performance is improving, yet there is widespread concern about declining trust in the state sector. This is important because new developments throw increased emphasis on the role of organisational trust in enabling co-ordination of increasingly complex and decentralised organisations and on public trust in the acceptance of controversial and far-reaching reform programmes as legitimate.

New research in psychology and sociology provides an explanation, with strong implications for future public sector reform: trust includes both deliberative and affective dimensions. Recent reforms are based on an organisation theory derived exclusively from a rational actor model. The new public policy operates through careful management of the framework of incentives presented to service providers and users, so that their deliberations will produce the behaviour desired by policy makers. It fails to engage the affective component in trust and may actually undermine it. The argument is illustrated through analysis of British Social Attitudes survey data to demonstrate the importance of the

affective component in public trust. Current reforms may generate improved services, at the same time as they undermine the pattern of relationships that sustains affective trust. The result will be better welfare at the price of an increasingly unstable welfare state settlement. Nothing is easy.

Credibility, calibration, and competence

Michael Birnbaum – California State University, Fullerton

Decision-making based on advice from sources who vary in expertise and bias.

This talk will review a theory of source credibility that has held up quite well over the last thirty years in a variety of decision-making tasks. People can be asked to judge the buying or selling price of a used car or an investment, the probability that an event will occur, to decide which medical procedure they would prefer, or which gamble they would rather play. The model assumes that value of an alternative is a configurally weighted average of the opinions from advisors, which are adjusted by perceived bias of the source and weighted by the perceived expertise of the source. The configural weight (weight transferred from high to low opinions) depends on the judge's point of view, which is what differs between buyers and sellers. This model is much more accurate than the theory later introduced by Tversky and Kahneman that buying and selling prices differ because of "loss aversion".

Yaron Shlomi & Thomas S. Wallsten – University of Maryland

The effect of advice from sources differing in calibration.

Probability judgments range from over-confident through calibrated to under-confident. Concomitantly, they are weakly through moderately to very diagnostic. In other words, there is a trade-off between extent of calibration and diagnosticity in the quality of probability judgments. This research is aimed at understanding how judges use advice from advisors who vary on this continuum.

In two experiments, research participants experienced various advisors by seeing their probability estimates and the associated outcomes. Following this experience, they were asked to make judgments in the presence of the advisor's probability estimates. This procedure allowed us to assess the impact of experiencing the advisor's calibration on the weighting of his recommendations.

Overall, responses to an under-confident advisor were more extreme than to a calibrated advisor, which in turn, were more extreme than to an overconfident advisor. This suggests that judges give more weight to diagnostic than to calibrated advisors. In addition, response variability was greatest in the presence of an overconfident advisor and smallest in the presence of a diagnostic advisor. We present a formal model for understanding how judges weight recommendations from differentially diagnostic-calibrated advisors for purposes of updating their own opinions.

Helmut Jungermann – Technical University of Berlin

The effect of advisors' competence and self-interest in a dyadic decision situation.

In the kind of dyadic decision situations we study, there exists an informational asymmetry between the decision maker and another person, an advisor. We assume that in such situations four factors determine whether the decision maker accepts or rejects the advisors recommendation: the judgment of the decision maker regarding the recommended option, the judgment of the advisor regarding the recommended option, the advisors credibility, and the clients confidence. When there is a conflict between the decision makers preference and the advisors recommendation, the decision makers trust in the advisor can become significant. In two internet-based experiments we manipulated two

important components of trust, the advisors' alleged competence and their self-interest in the outcome, and analysed the effect of these variables on decision makers' choices.

Statements of trust and use of advice

Lyn Van Swol- Northwestern University

The Effects of Advisor Motives on Confidence and Advice Utilization.

The talk examines research using a Judge-Advisor System model (JAS) to examine the effects of different advisor motives on the advisors' behavior and on the judge's utilization of the advisors' advice. In 2 experiments, 2 advisors gave a decision-maker (judge) advice. One advisor (persuasive motive) was instructed to try to get the judge to accept his or her advice, regardless of its quality. The other advisor (quality motive) was instructed to help the judge make the best possible decision. In Experiment 1, the judge was more likely to match advice from the advisor with the ulterior, persuasive motive and rated this advisor as more influential. In Experiment 2, half the judges were warned that some advisors may have a persuasive motive. Whether warned or not, judges were more likely to accept advice from the persuasive advisor, but judges were less likely to accept advice overall from both advisors when warned about the advisors' possible motives. Judges were unable to distinguish the advisors' differing motives. The persuasive advisors were more confident in their advice, and their confidence was correlated with their influence upon the judge. The research is discussed in light of the fact that much previous research examining advice-giving has assumed that the advisors have no ulterior motives. Further, the role of confidence as a tool of influence and manipulation in the advice process is discussed.

Gideon Keren - Eindhoven University of Technology

Incompatible messages implied from positive and negative frames regarding trust and choice decisions.

Following common wisdom, if one trusts agent A more than B, then one should prefer to conduct transactions with the former rather than with the latter agent. Several experiments are presented that are incompatible with this conjecture. For example, when faced with a choice between two butchers, whose ground beef is advertised as containing 25% fat (negative frame) or 75% lean (positive frame), respectively (Levin, 1987; Levin and Gaeth, 1988), most people have more trust in the former yet most indicate they would buy their meat from the latter butcher. This phenomenon, in which negative framing weighs more in trust assessments, and positive framing weighs more in choice, is labeled *trust – choice incompatibility*. The robustness of the phenomenon is further demonstrated in several experiments, and possible explanations for its occurrence are discussed.

Matt Twyman, Nigel Harvey and Clare Harries - University College London

Stated and revealed trust in advice.

People often learn about the levels of risk associated with different activities through advice, and their use and assessment of such advice may depend on factors such as the identity of the advisor, and the perceived quality of that advice. Earle & Cvetkovich (1999) demonstrated that explicit verbal estimates of trust in advisors correlate with perceived shared values between advisor and advisee. Here we apply that finding to a risk communication paradigm. Earle & Cvetkovich's findings were replicated in two experiments, in which participants were given advice about a range of risky activities. However, declared trust in advice sources did not correlate with how much those sources were used in making risk judgments. Relative measures of use and assessment of advisors were also found to bear different relationships to the accuracy of advice. Use of advisors was not reflected in explicit verbal estimates of trust in those advisors.

Taking advice from machines

Peter Ayton, Eugenio Alberdi, Andrey Povyakalo and Lorenzo Strigin – City University

Diverse Effects Of Computerised Advice On Expert Decision Making.

Because X-ray reading requires scarce, highly qualified staff and is a difficult and attention-intensive job, computer-based advisory systems have been designed to help readers not to overlook details that they ought to examine to reach their decisions (Computer Aided Detection: CAD). Indeed, developments in computing now offer experts in many fields specialised support for decision making under uncertainty. However, the impact of these technologies is somewhat controversial; in particular, it is not clear how advice of variable quality from a computer may affect human decision makers. Here we discuss research designed to investigate how the graphic “prompts” that CAD provides on regions of interest on mammogram images affect breast cancer screening decisions.

Ideas from previous research on the reliability of software systems that use redundant, diverse components (i.e. that perform the same task but in different ways) suggested that we should study the dependability of the overall “system” (CAD + human) as a function of variation and co-variation in the “difficulty” of input cases for the two different components of the CAD system (computer and human). By contrast, in clinical trials, decision support technologies are often assessed in terms of their simple average effects. However this methodology overlooks the possibility of differential effects on decisions of varying difficulty, on decision makers of varying competence, of computer advice of varying accuracy and of possible interactions among these variables.

Re-analyses of a clinical trial reporting no effects of CAD that teased apart the aggregated data revealed strikingly diverse effects of computer support on expert decision-making. Decision support can both systematically improve or damage the performance of decision makers in subtle ways depending on the decision maker's skills, variation in the difficulty of individual decisions and the reliability of advice from the support tool. Computer support was less useful for - and sometimes hindered - professional experts who were relatively good at difficult decisions without support; at the same time the same computer support tool helped those experts who were less good at relatively easy decisions without support. Our experiments on expert readers further indicated that when CAD failed to prompt cancers the readers using CAD were biased by the incorrect computer output. Readers with computer support were more likely to miss unprompted cancers than readers in a control group who saw the same cases without CAD.

Paul Goodwin - University of Bath

Accepting advice from a machine: the role of statistical advice in forecasting support systems.

Managers in supply-chain companies usually have access to expensive forecasting software which provides advice in the form of statistical forecasts of future demand for their products. Although these forecasts are usually based on extrapolations of past patterns, and take no account of forthcoming special events like promotion campaigns, they can often provide reliable estimates of future demand. However, the forecasts produced by these software packages are frequently ignored by forecasters or the available time series data is manipulated in order to obtain the statistical advice that the forecaster is comfortable with. This talk will use the results of a study in a pharmaceutical company to examine the reasons for this. It will show that simply accepting the advice from the software would have saved the forecasters considerable effort and usually have resulted in more accurate forecasts. It will then draw upon the results of a number of experimental studies to examine how statistical advice might be made more acceptable to forecasters.

Philip Bonhard – University College London

“Thanks for the advice, but do I know you?”: Improving Decision-Making Support of Recommender Systems through Social Networking.

In today's information age, choice abundance has become a burden. Whether we are going out to a restaurant, choosing a DVD to rent or buying a CD online, the amount of available choices far outstrips what we can possibly rationally consider. Seeking advice and recommendations from friends, peers and other trusted sources is normally an efficient solution to solve this problem. In an online context, recommender systems are meant to emulate this advice seeking process through filtering available options based on user taste matching via item ratings. The resulting options however only carry data value for the advice seeker as they usually lack explanations as to how the recommendations were generated. Whereas in face-to-face interactions we usually have some information about the recommender and can question their advice directly, in online interactions this social element of advice seeking is missing. This talk presents a new view of how the functionality of online recommender systems and social networking applications can be integrated to leverage the benefits of both. The aim is to reintroduce the social element of advice seeking and decision making into this interaction and thus help the user judge the appropriateness of given recommendations.

Weighting advice

Elise Weaver – Worcester Polytechnic Institute

A new measure of relative change of importance weights when given advice.

I present a multidimensional measure of message effectiveness where the message is a vector of importance weights for weighing the attributes of a case in judgment. In other words, an advisor tells someone how to weigh information for all cases rather than how to respond to a specific case. The person, who starts with a certain policy for weighing information in judgment, moves to a new policy after advice. The measure is a relative change measure in that perfect influence by the advisor is represented by a 1, partial influence is represented by a fraction, no influence by 0, but an individual can also react to advice in the opposite direction scoring a negative value.

Stian Reimers - University College London

How do we know if people really use advisors' track records: Regression weights and optimal performance.

Many experiments on advice taking investigate whether participants can learn, through repeated-play feedback, to give more weight to good advisors over bad advisors. Multiple linear regression is often used: Higher beta weights suggest that an advisor's predictions are being used more in a participant's estimates. Using simulations, I argue that this can be artefactual: Unweighted strategies such as taking the median or the trimmed mean can lead to a large differential in beta weights, favouring better advisors. In other situations, researchers compare participants' performance with that predicted by unweighted strategies, like taking the median or mean: If performance is better, people must be using track-record information. However, the optimal unweighted strategy is context dependent, so finding an appropriate baseline may be impossible. As an alternative, I suggest we should measure unweighted performance experimentally, with incentivised participants.

Derek J. Koehler¹ & T. Alexandra Beauregard²

¹ *University of Waterloo*

² *London School of Economics*

Illusion of Confirmation from Exposure to Another's Hypothesis.

We examine the influence of exposure to an advisor's hypothesis, in the form of a point estimate of an uncertain quantity, on subsequent point estimates and confidence judgments made by advisees. In three experiments, a group of unexposed advisees produced their own estimates before being presented with that of the advisor, while a group of

exposed advisees were presented with the advisor's estimate before making their own. Not surprisingly, exposed advisees deliberately incorporated the information conveyed by the advisor's estimate in producing their own estimates. But the exposure manipulation also had a contaminating influence that shifted what the advisees viewed as their own, independent estimates toward those of the advisor. Seemingly unaware of this influence, exposed advisees were subject to an illusion of confirmation in which they expressed greater confidence in the accuracy of the advisor's estimate than did unexposed advisees.

Models and analyses

Robert D. Sorkin¹ & Shenghua Luan²

¹*Air Force Research Laboratory & University of Florida*

²*Max Planck Institute for Human Development*

Signal Detection Analysis of Advice Behavior.

Suppose that you have a chance to purchase an antique British auto, a 1960 Triumph TR3B. Your estimate of the car's value is moderately high and you have a chance to get advice from two friends who are knowledgeable about old cars. You have time to get advice from only one of your friends—they are equally expert but one friend is very conservative (usually says "don't buy it") and the other is very liberal (usually says "buy it"). Which expert should you ask? Our talk will discuss optimal detection analyses of advice acquisition and use: A person must make a decision about the occurrence of a noisy event. Advice about the event's occurrence is available, but the person must decide whether or not to purchase advice, which advice source to choose, and how to integrate the obtained advice into a final decision. We use a detection analysis to define these issues and prescribe optimal strategies for their resolution. We describe the results of some experiments with trained and paid human decision makers in similar situations. The results of these experiments indicate that human advice behavior is generally consistent with the optimal strategy. For example, in the TR3B example, people generally make the optimal choice—they choose the liberal advisor: the so-called "confirmation bias" strategy.

David V. Budescu & Hsiu-Ting Yu - University of Illinois at Urbana-Champaign

To Bayes or not to Bayes? A Comparison of Two Classes of Models of Information Aggregation.

We model the aggregation process used by individual Decision-Makers (DMs) who obtain probabilistic information from multiple, possibly non-independent, sources. We distinguish between two qualitatively different aggregation approaches: "compromising" by averaging the advisors' opinions and combining the forecasts according to Bayes rule. The DMs in our studies received forecasts from 2 or 3 advisors who had access to multiple diagnostic cues, and made a large number of decisions. Our data are unusually rich in many respects since the studies involve natural sampling of cues with various levels of dependence and various patterns of information overlap. This provides an excellent opportunity to compare the quality and fit of these models.

Overall, the DMs' judgments were closest to the averaging model but, clearly, they do not use this model exclusively and uniformly. It appears that DMs rely on a variety of external cues in adopting a particular aggregation rule. Interestingly, the most relevant cues for this purpose are those that affect the DM's confidence (and trust) in the advice. Their aggregates were more in line with Bayes rule when the advisors provided extreme forecasts, and were highly consistent with each other.

Timothy Earle - Western Washington University

Advice and Trust in Risk Management Contexts: The TCC Model Perspective.

Advice, in risk management contexts, refers, in large part, to the nature of risks and how they should be dealt with. Typically, advice is provided by individuals or institutions charged with the management of a risk to those exposed to it. The trust that the latter have in the former is a major factor affecting the acceptance of the advice. The TCC model of cooperation based on trust and confidence is framework for understanding the antecedents and consequences of trust as well as the effects of trust on the interpretation of risk information or advice. In the TCC model, trust is social and relational; confidence is instrumental and calculative. We define trust as the willingness, in the expectation of beneficial outcomes, to make oneself vulnerable to another based on a judgment of similarity of intentions or values. Confidence is the belief, based on experience or evidence (e.g., past performance), that certain future events will occur as expected. The basis for trust is fundamentally a judgment of similarity between one person and another, that the person to be trusted (or entity treated as a person) would act as the trusting person would. Thus, trust is based on social relations, on shared values. The basis for confidence is past performance, or institutions/procedures designed to constrain future performance. Trust is based on morality-relevant information; confidence is based on performance-relevant information. According to the TCC model, trust should affect the interpretation of risk information most forcefully when strong morality information is present, e.g. when individuals believe that a risk management issue is morally important to them. Individuals' knowledge of the issue—which is often positively related to moral importance—should also affect the impact of trust. Data from several studies are used to explore these hypotheses, and the implications of these findings for risk communication and risk management practice are discussed.

POSTER ABSTRACTS:

Buczek, Katarzyna - Max Planck Institute for Human Development

How biases enable better decisions: The influence of player heterogeneity on formation of informational cascades.

Imagine you have to decide whether to adapt some new technological solution. You don't know for sure whether it is beneficial to do so, you only have some imperfect information. However, you can ask for advice from others who also have imperfect information, possibly different than your own. Say their advice is public so that each subsequent actor knows what the previous actors said. Assume furthermore that some have preferences simply for good technology, but others are eager to apply almost all new solutions, while still others are afraid of any changes. From whom would you ask advice? When individuals make decisions sequentially an informational cascade may arise. This happens when every subsequent actor, based on the observations of others, makes the same choice independently of his/her private information. As a result, even if there is overwhelming information in favour of the correct choice, everybody could be mistaken. The work to be presented is a formal investigation of the influence of heterogeneous preferences on the formation of such cascades. It is shown that the way information is aggregated depends on the quality of private information, information concerning others' biases, and the order in a queue. There are situations in which some individuals do cascade and some do not. A single individual with a preference against the cascade may stop it. At the same time biased individuals tend to cascade sooner, if it is in favour of their bias or simply choose the less risky option. In general heterogeneity of a group may fulfil a role similar to the role of overconfidence and increase the overall accuracy of choices. In some instances paying attention to the advice of a biased player is more beneficial than relying on choices of neutral individuals.

Davies, Greg - UCL & Decision Technology

The sources of trust: Consumer trust in large retail organizations.

Consumer decision making requires that consumers have trust in the organisations from which they purchase. However, it is unclear how overall consumer perceptions of “Trustworthiness” are formed. Trust could arise from a) perceived short and long-term alignment of corporate values with those of the customer, b) from perceptions of competence, and c) from the perceived altruism of the organisation. Furthermore, perceptions on each of these dimensions could differ for each low level issue on which customers expect the organisations to deliver (e.g., delivering low prices, safe food, or treating staff and suppliers fairly), each of which could contribute differentially to overall trust levels depending on importance to consumers. Finally, individuals may form perceptions of trust in different ways.

Working together with a major British retailer we develop a model of trust formation incorporating each of these elements and then employ data from both large scale online surveys and controlled experiments to validate and test the model. We show that perceptions of corporate trustworthiness are distinct from perceptions of corporate self-interest. The data allow us to test the importance of the individual components of trust in forming overall trust perceptions, as well as providing some insight into how trust levels for individual issues combine to form global trust perceptions for an organisation. Finally, we show that individuals do indeed differ in both the issues that drive their trust perceptions, and in the manner in which they build overall perceptions of corporate trustworthiness and self-interest.

Mata, J., Scheibehenne, B. & Todd, P.M.

Parents’ advice giving on children’s lunch choices: How successful can they be?

Food choice is one area where it is important for parents to give good advice to their children: Children often make poor decisions (e.g., eat greasy foods with low nutritional value) when parents are not present. Giving appropriate advice e.g., to eat food that is both nutritious AND tasty to the child, could help children make better choices. In this domain, successful advice giving thus requires parents to know their child’s food preferences. But do they? To find out, we investigated how well parents predict their child’s meal preferences and whether parents use projection of their own preferences as a strategy to do this. In addition, we tested parents’ calibration whether their perceived accuracy in meal prediction matches actual accuracy, to find out if they are overconfident of their ability to predict their child’s preferences.

Our study assessed 59 parent-child dyads. We asked each child to rate the degree to which he/she would like to eat a number of individual lunch meals. Simultaneously, each parent was asked to (1) predict how much their child would like the meals, (2) estimate the accuracy of their predictions, and (3) rate how much the parent himself/herself would like to eat each lunch meal.

Overall, parents were able to predict 47% of children’s choices correctly, which matches the results found in a previous study (Mata, Scheibehenne, & Todd, 2006). Concerning calibration, only a minority of parents correctly estimated their prediction accuracy. On average, parents overestimated their accuracy by about 20%. Parents used projection as a strategy in only a small percentage of the cases, and otherwise used additional cues to achieve higher accuracy. In sum, parents’ knowledge of their children’s meal preferences is not perfect. We discuss the potential impact of this for successful advice giving in the food domain.

McColl, A. - Durham Business School, University of Durham

The benefit of an additional opinion.

Seeking advice from others seems very much in the realm of the commonplace, yet studies exploring the influence of advice on decision-making are few. Of these studies, most are concentrated in the Judge Advisor Systems (JAS) stream of research. Here, inputs are provided by *advisors*, but the *judge* is the final arbiter over the weight and use of advice. One robust finding from this work is that people generally prefer their own opinion, over that of an advisor - people *egocentrically discount* advice. Arguably, this is because where advice is numerical a judge has unrestricted

access to their own cognitions, but constrained access to the reasoning of an advisor (Yaniv & Kleinberger, 2000; Yaniv, 2004). We tested this *task specific* explanation in two phases by tasking participant_s to estimate the date of a notable historical event. In the second phase participants could revise their original estimate in the light of either no advice, purely numerical advice, or numerical advice with additional reasons. We hypothesised that egocentric discounting should disappear when reasons were provided in addition to numerical advice. The persistence of egocentric discounting would argue for a cognitive account of sub-optimal advice use. Further, we hypothesised that participant_s in receipt of numerical advice and additional reasons, would estimate more accurately, than participants in receipt of purely numerical advice. Findings indicate that there is no difference between numerical advice with reasons and purely numerical advice, in terms of the weight judges place upon advice, and the ultimate accuracy of their estimates. For a cognitive activation account of egocentric discounting to be accepted, participants would have to beneficially revise their estimates in the face of inaccurate advice. This was found not to be the case, and we discuss our data in terms of an *anchor and adjust* conceptualisation of egocentric discounting.

Marianne Promberger, M, and Baron, J. - University of Pennsylvania

Do patients trust computers?

In two studies, we inquired whether patients accept medical recommendations that come from a computer program rather than from a physician. In study 1, we found that subjects, when deciding whether to have an operation or not in different medical scenarios, were more likely to follow a recommendation that came from a physician than one that came from a computer program. Subjects stated that they would feel less responsible when following a recommendation than when deciding against it. Following a physician's recommendation reduced feeling of responsibility more than following that of a computer program. The reduction of responsibility when following partly mediated subjects' inclination to follow the physician more. In our second study, we found that subjects were more decision seeking when they received a recommendation or decision from a computer program, and they were more decision seeking when they had to accept a decision than when they received a recommendation. Subjects also trusted the physician more than the computer program to make a good recommendation or decision.

Pulford, B.D., & Evans, L. - University of Leicester

Influences on advisors' confidence

How people's confidence in their advice is influenced by other people's advice was experimentally examined in this study. In a pencil and paper task 80 participants gave written advice to a friend about the truth/falsity of 16 general knowledge statements. A high/low financial reward was supposedly at stake for a friend who received the advice. Participants saw the statement of general knowledge and also the advice that another person was also giving to the friend. The status of the other person (high/ low) was crossed with their gender (male/female advisor), and their level of certainty in their advice (high/low confidence). This resulted in 8 statements, such as "This is definitely true : Professor Jane Harteau" which were used twice, once when the statement was judged to be true and once when it was judged to be false. The written advice that participants then gave was scored by two raters for the level of confidence/certainty it conveyed, from 1 (low) to 7 (high). The average score across the raters was the DV entered into a 2 X 2 X 2 ANOVA. Status interacted with speaker confidence when there was a low reward, but not when there were high-stakes rewards involved. With low rewards, advisor confidence was unimportant when the speaker was of low status, but it significantly influenced how confident the participants felt when the advisor had high status. A high confidence, high status speaker inspired the participants to express higher confidence in their own advice in the low reward condition. Speaker confidence and speaker gender interacted slightly differently according to reward level but, most noticeably, the advice participants gave had the highest certainty after hearing advice from highly confident female advisors.

Qian, Jing - Max Planck Institute for Human Development

How the distribution of information and the payoff function influence the communication between adviser and addressee.

In risk communication, it is commonly found that information-provider (advisor) prefer to use qualitative expressions whereas information-receiver (advise-taker) prefer to receive quantitative information (e.g., Erev & Cohen, 1990; Wallsten, Budescu, Zwick & Kemp, 1993). Although the two parties in communication have different intentions for precision of expression, they share the common goal of minimizing error in communication. In verbal categorization of psychophysical magnitudes (such as sizes of squares), Parducci (1965, 1995) noted that the distribution of magnitudes influenced their categorization. Specifically, people tend to follow two rules in verbal categorization: They tend to assign equal number of stimuli under each verbal label (the frequency principle), and they also tend to divide the total range of the all stimuli into equal bins, and assign stimuli to them according to their relative size (the range principle). The relative weight people adhere to these rules is usually a balanced compromise, empirically derived in psychophysical judgements. However, in risk communication, this weight can be determined by the pay-off function employed by the two parties. In this study, a communication game is simulated, where the advisor assign verbal labels to quantitative risk magnitudes and the advise-taker recover this verbal information back into quantitative terms. The goal of the game is to minimize error in information transmission and recovery. Two payoff structures were employed, one of which penalize only the direction of error, and the other penalize the size of error. It is found that when error is penalized only by direction, adherence to the frequency principle is more fruitful than to the range principle, whereas when the size of error is penalized, compromise with the range principle minimizes error more than either the range-only strategy or the rank only strategy. These results can be modeled using Range Frequency Theory.

Scheibehenne, B., Mata, J., and Todd, P.M.. - Max Planck Institute for Human Development

Are parents able to predict the lunch choices of their children?

Knowing individual preferences is an important prerequisite for advice giving and subcontracted decision making. In a study on school lunch choices with 99 school children (age 8 to 12) and 30 of their parents we investigated how well parents were able to predict the lunch choices of their children. As parents usually are not certain about the preferences of their children they have to infer them based on proximal information cues. If family members were similar to each other they could effectively use their own likes and dislikes as a cue for prediction of each others preferences. However, in the food domain, past research has shown that food preferences within a family are surprisingly dissimilar ("family paradox").

Our data confirms that paradox. Nevertheless, parents in our study use their own preferences as a cue to predict their children's food choices. As a consequence, the more similar parents are to their children the more accurate their predictions. Our empirical data indicates that children's lunch choices can also not be predicted (by parents or others) based on health knowledge. However, as many children tend to choose the same dish, the preferences of other children potentially provides a valid cue. Overall, parents' predictive accuracy is rather low. To some extent this can be explained by the large variability of children's food preferences (assessed based on a retest 3 month later). In a follow up study with 69 pairs of parents and their children we further investigated the ability of parents to predict their children's food preferences as well as their confidence of doing so.

Ungemach, Christoph - University of Warwick

Decisions From Experience And Their Implications For Advice On Risk Communication

In everyday life people very often learn about risky behaviours and their consequences only through exploration and experience (e.g. health risks) though advice on risk communication is mainly focused on the communication of objective risk assessments. However, experiments on decision from experience have been shown to result in a different choice pattern than the presentation of descriptions of lotteries, that is, decision from description. According to Hertwig, Barron, Weber and Erev (2004), this difference is based on the combination of a reliance on small samples and overweighting of recently sampled information which results in a choice behaviour that appears to be driven by

underweighting the probabilities of rare events. In scenarios where information about outcomes and their probabilities is collected through repeated experience decision makers would therefore be advised to invest more time and effort into the information search by considering a bigger sample in order to minimise the bias. This research investigates the processes underlying decision from experience and seeks to challenge the current lines of explanation thereby questioning the type of advice appropriate in such scenarios. In a study (N = 50) investigating decisions from experience under two conditions – free sampling and forced sampling (39 samples per option) – the choice pattern could be replicated even when participants were forced to sample more exhaustively, reducing the probability of experiencing the rare event less frequently than expected. Across all prospects used, the majority of the sets of samples drawn in this condition actually overestimated the real probability. In addition, no significant recency effect was found for the sampling process indicating that the reason/s for the different choice pattern remain unclear. These results suggest that different choice behaviour may derive from reasons other than underestimation as a result of reliance on small samples and recency weighting, pointing to a need for further investigation into the links with probability judgments and their applications for advice on risk communication.

Van Dongen, K., & Van Maanen, Peter-Paul - TNO Human Factors

Under-trust in decision support systems.

It is often assumed that two heads are better than one, but reliance on decision aids is often inappropriate. Decisions to rely on an aid are thought to be based on a comparison between the perceived reliability of own performance and that of the decision aid (Dzindolet, et al., 2003). Unfortunately, perceived reliabilities are unlikely to be perfectly calibrated. This can result in inappropriate decisions to rely on advice. In a laboratory experiment with 40 participants, we studied whether calibration improves after practice, whether calibration of own reliability differs from calibration of the aid's reliability and whether unreliability of the aid is attributed differently. Under-trust in own reliability disappears after practice but under-trust in the aid's reliability does not. Unreliability of the decision aid is less likely to be attributed to temporary external factors. The asymmetry in attribution and calibration may explain under-reliance on decision aids.

Wesson, C and Pulford, B - University of Wolverhampton

Abstract: Knowing what other people know: Implications for the confidence heuristic.

When we are uncertain we may turn to other people for their advice. The level of confidence they have in their advice can influence the choices we opt in and out of. Making decisions in this way indicates that confidence is used heuristically, whereby a speaker's confidence is taken as being a cue to their knowledge, competence or accuracy. This study considers whether people apply the confidence heuristic when provided with information relating to a speaker's performance, and hence their judgmental biases. 86 participants observed three speakers answering general knowledge questions, with each speaker demonstrating a different level of confidence and bias. Half the participants also received feedback about the speakers' performance. Participants then completed a general knowledge questionnaire, choosing from answers given by the three speakers. A speaker's level of confidence had a greater effect on choice when no feedback was given than when it was. The addition of feedback led to a reduction in the influence of a highly confident, but overconfident, speaker and increases in the influence of the medium and low confidence speakers. Feedback did not have any significant effects on participants' confidence in their answers, but it did result in a slight increase in accuracy. A further experiment considered how speakers are perceived when listeners are provided with information relating to the speaker's performance. The results indicate that people do, to a certain extent, continue to apply the confidence heuristic in their judgments of others even when there is evidence to suggest that the level of expressed confidence used by the speaker is misleading.

A note regarding the conference dinner on June 29th:

The conference is scheduled to finish for the day at 6.00pm on Thursday June 29th, and all those attending the conference dinner are asked to meet at the London Eye at 7.00pm. Your flight on the eye begins at 7.30pm, but we may not be issued tickets for those who arrive after 7.00pm. Dinner will be at Chez Gerard, a 5 minute walk from the Eye, at 8.15pm. We will walk to the restaurant as a group after the Eye Flight.

Below are some directions to the London Eye, which cover every eventuality except Black Cabs. You can hail a cab in the street, ask your hotel to call one, or ask Matt Twyman for help. You are strongly advised to not take an illegal minicab.

By tube:

The Eye is approximately five minutes walking distance from Waterloo tube (follow signs for the South Bank) and Westminster tube station (exit one, follow signs for Westminster pier).

Unless you intend to take a Black Cab, the conference organisers recommend taking the tube. Get on the tube at Russell Square (very close to the conference venue) and travel west four stops to Leicester Square. Change trains there to get on the Northern Line, and travel a further three stops south to Waterloo station.

By rail:

The Eye is five minutes from Waterloo station. Take exit six for the South Bank and follow the signs.
The Eye is 15 minutes from Charing Cross station, which is accessed via Hungerford pedestrian bridge.

By bus:

Buses to the London Eye include the 211, 24 and 11. We are on most London sightseeing bus tours and on the new RV1 route that connects the London Eye to the Tate Modern and Covent Garden.

By car:

We advise against driving to the London Eye. However if you do drive, there are three car parks within walking distance located around the South Bank area.