Metacognition and Intuition in a Denominator Neglect Task: Converging Evidence from Individual Differences and Gaze Tracking Analyses

Valerie Thompson

It is a well-established finding that people with high IQ and who are motivated to think analytically perform better on a variety of reasoning tasks than their less able and less motivated counterparts. The usual interpretation of this relationship is that the former group have the capacity and motivation to inhibit appealing intuitive answers and engage the necessary resources to derive a correct one. In this talk, I will present some counter-intuitive evidence to suggest that high-IQ people are more accurate, even when they must respond under time pressure, suggesting that their intuitions, as well as their deliberations, are better than low-IQ reasoners. I will also provide an analysis of gaze patterns that help to elucidate strategies that lead to accurate and confident responses.

Individual differences in bias detection during thinking

Darren Frey & Wim De Neys

Decades of reasoning and decision-making research has established that human judgment is often biased by intuitive heuristics. Recent "error" or bias detection studies have focused on reasoners' abilities to detect whether their heuristic answer is erroneous and conflicts with logical or probabilistic principles. A key open question is whether there are individual differences in this bias detection efficiency and how this affects reasoning performance. Here we present three studies in which co-registration of different error detection measures allowed us to assess bias detection sensitivity at the individual participant level in a range of reasoning tasks. Results indicate that although most individuals show robust bias detection (as indexed by increased latencies and decreased confidence associated with erroneous answers), there are subgroups of reasoners who consistently fail to do so. We discuss implications for the debate on human rationality and popular dual process theories.

The development of the CRT-Long and what we have learnt in the process

Kinga Morsanyi

The CRT-Long (Primi et al., 2015) is an extension of the CRT (Frederick, 2005) that includes 3 new items. The new and the old items share the important property that, although the items are open-ended, more than 80% of respondents either give a correct or a typical incorrect (i.e., heuristic) response. A major limitation of the original CRT is the difficulty of the items, which can lead to floor effects in populations other than highly educated adults. The CRT-Long has been developed using item response theory (IRT) analyses. These analyses confirmed the good psychometric properties of both the original and the new items (i.e., that

they discriminate reliably between responders with lower and higher levels of the cognitive reflection trait). Additionally, the new scale measures cognitive reflection reliably among a wider range of abilities (i.e., it is also appropriate for participants with lower ability levels; for example, developmental samples). In this talk I will present some data regarding the validity of the CRT-Long, the effects of some experimental manipulations (i.e., applying working memory load, manipulating font fluency, and employing performance pressure by introducing the CRT as a measure that is highly diagnostic of reasoning abilities). I will also present the results of comparisons between adolescents and adults, and between participants with autism and a typically developing control group. I will discuss the implications of these findings for research on reasoning heuristics and cognitive reflection.

Measuring belief bias

Dries Trippas

Belief bias is the tendency for people to rely on their prior beliefs even when reasoning deductively. Recent work by Heit and Rotello (2014) suggests that way we have been measuring belief bias is flawed, questioning most of our theoretical conclusions about the phenomenon. The solution they propose is to apply signal detection theory (SDT) to ensure more appropriate measurement and theory. A potential issue with the routine application of SDT to reasoning is that the theory makes certain assumptions. One key assumption is that argument strength is normally distributed with an unequal variance ratio between valid and invalid arguments. A second assumption is that the link between argument strength and confidence is straightforward. In this talk I evaluate the plausibility of these assumptions. In light of this evaluation I propose a statistical (logistic regression) and a methodological (forced choice) alternative to SDT which can also address some of the issues raised by Heit and Rotello (2014) without relying on confidence ratings (Trippas, Verde, & Handley, 2014).

Jaydeep Singh and Mike Oaksford

Discounting Testimony with the Argument Ad Hominem and a Bayesian Congruent Prior Model

When directed to ignore evidence of a witness' previous bad character because of a violation of the rules of evidence, are jurors' beliefs still affected? The intuition is that they will be because in everyday argumentation, fallacies, like the ad hominem, are effective argumentative strategies. An ad hominem argument (against the person) undermines a conclusion by questioning the character of the proposer. This intuition divides current theories of argumentation. According to pragmadialectical theory (e.g., Van Eemeren & Grootendorst, 2004) procedural rules exactly like the rules of evidence are part of our cognitive resources for evaluating arguments. If one of these rules is violated, an argument should be treated as a fallacy and so it should not alter someone's belief in the conclusion. Some recent experiments investigating how reasonable these arguments are perceived to be seem to support this account (Van Eemeren, Garssen, & Meufells, 2009). These experiments are critiqued from the perspective of the relevance (Walton, 2008, 2009) and epistemic (Hahn & Oaksford, 2006, 2007; Oaksford & Hahn, 2004) approaches to argumentation. An experiment investigates the predictions of these approaches for a graded belief change

version of Van Eemeren et al.'s (2009) experiment and the results are modelled using a Bayesian congruent prior model. These results cannot be explained by the pragmadialectical approach and show that in everyday argument people are extremely sensitive to the epistemic relevance of evidence. Moreover, it seems highly unlikely that this can be switched off in more formal contexts such as the courtroom.

Simon Handley, Dries Trippas and Valerie Thompson

TBC

David Lobina

On System 2 – The linguist's input

According to dual-theories of reasoning, System 1 is fast, parallel, and automatic, possibly universally present in animal cognition, and perhaps similar in nature to Fodorian modules (Fodor 1983). System 2, on the other hand, is slow, sequential, allows abstract hypothetical thinking, and is definitely unique to humans (Evans 2003). The beliefs and reasonings involved in System 2, moreover, Frankish (2004) argues, are frequently language-involving (in the case of belief formation) and language-driven (in the case of reasoning). The latter point, in particular, connects with the work of Carruthers (1998, 2006), who argues that natural language may be the (main) medium of thought – and must certainly be so in the case of conscious types of thinking.

I here offer a critical appraisal of the connection between natural language and System 2. Firstly, I argue that the role of language in thought, be this conscious or unconscious, is very unlikely, either in terms of the formats (inner speech or logical forms, lately termed SEMs) or the forms (linear phonetic representations for inner speech, hierarchical structures in the case of SEMs) that the language faculty generates. Speech doesn't yield access to the underlying structure of our thoughts, as the relationship between internalised/externalised speech and SEMs (the underlying linguistic structure) is opaque, indeed inaccessible to speakers, and therefore constitutes an impossible conduit to the structure of, ultimately, thought. Moreover, the structure of linguistic expressions doesn't match up with the structure of thought representations, as evidenced in the many misalignments there are between syntactic structure and semantic/conceptual representations: lexicalisation, ambiguity, grammatical illusions, ungrammatical sentences for which a thought can be entertained, etc. Put together, these factors cast doubt on whether linguistic representations can really provide the representational medium for System 2 reasoning. Secondly, even if natural language were to somehow be involved in the operations of System 2, it must be as a type of linguistic behaviour (either in interior monologue or externalised speech), and it is a well-established feature of this type of behaviour that it is effectively stimulus *independent* (see Chomsky 1959, who emphasised this very point in his critique of behaviourist theories of language learning). That is, given a specific set of circumstances, there really is no way to determine what a person might say at any given moment, *if indeed anything at all*. A person may be incited to produce speech, or some particular piece of speech, but one is never actually compelled to do so. This feature of linguistic behaviour doesn't bode well for scholars who

claim that System 2 reasoning is language-involving/driven, for rationality studies, much as the rest of cognitive psychology, aim to elucidate *causal* mental processes, in order to in turn provide *causal* theories of cognition, and therefore the representations involved in System 2 cannot be linguistic. It is not only that we cannot be sure if the *right* linguistic material can be elicited in the rarefied setting of psychological experiments ("right" in the sense of pertaining to a causal, therefore scientific, account of behaviour); we actually can't be sure that *any* linguistic material is in fact being elicited.

I don't mean to suggest that dual-theories of reasoning, and in particular System 2 processes, are not well-founded; but I do mean to claim that language certainly plays no role therein. Instead, I suggest that System 2 requires an amodal, representational system; a *language of thought*, in the sense of Fodor (1975). An old story, to be sure, but nonetheless correct.

Eoin Travers

The Time Course of Conflict on the Cognitive Reflection Test

Reasoning that is deliberative and reflective often requires the inhibition of intuitive responses. The Cognitive Reflection Test (CRT) is designed to assess people's ability to suppress incorrect heuristic responses in favor of deliberation. Correct responding on the CRT predicts performance on a range of tasks in which intuitive processes lead to incorrect responses, suggesting indirectly that CRT performance is related to cognitive control. Yet little is known about the cognitive processes underlying performance on the CRT. In the current research, we employed a novel mouse tracking methodology to capture the time-course of reasoning on the CRT. Analysis of mouse cursor trajectories revealed that participants were initially drawn towards the incorrect (i.e., intuitive) option even when the correct (deliberative) option was ultimately chosen. Conversely, participants were not attracted to the correct option when they ultimately chose the incorrect intuitive one. We conclude that intuitive processes are activated automatically on the CRT and must be inhibited in order to respond correctly. When participants responded intuitively, there was no evidence that deliberative reasoning had become engaged.

Barnabas Szaszi, Balazs Aczel, Aba Szollosi

Studying Processes Underlying Cognitive Reflection Test

Performance on the Cognitive Reflection Test is supposed to be determined by one's capacity to suppress the wrong intuitive answer that first comes to mind and to substitute it with the correct one. The present study attempts to reveal whether this original suppression-based explanation can be hold or whether there are a reasoners who, due to some early selection process, manage to solve the tasks without the need to inhibit the intuitive incorrect

response. The study also aimed to identify how individual differences in numeracy and reflectivity influence whether the correct or incorrect response alternatives come first to mind. In order to trace the participants' decision processes, we used concurrent verbal reports and protocol analysis of the CRT along with several measures of reflectivity and numeracy (N= 210). The results suggest that there are a substantial part of reasoners who do not show this late correction pattern but start the thinking process with the right strategy. Our data also imply how the current literature sometimes underscores the importance of investigating subgroups of participants separately when trying to understand how reasoning works.

Balazs Aczel, Aba Szollosi, and Bence Bago

On the Determinants of Confidence in Conjunction Fallacy

The general assumption that people fail to notice discrepancy between their answer and the normative answer in the conjunction fallacy task has been challenged by the theory of Logical Intuition. This theory suggests that people can detect the conflict between the heuristic and normative answers even if they do not always manage to inhibit their intuitive choice. This theory gained support from the finding that people report lower levels of confidence in their choice after they commit the conjunction fallacy compared to when their answer is not in conflict with logic. In four experiments we asked the participants to give probability estimations to the options of the conflict and no-conflict versions of the tasks in the original setup of the experiment or in a three-option design. We found that participants perceive probabilities for the options of the conflict version less similar than for the no-conflict version. As people are less confident when choosing between more similar options, this similarity difference is proposed to serve as a mediator in the task in a way that the conflict and no-conflict conditions have their effects on confidence ratings through manipulating the similarity of the answer options.

Philipp Koralus

Systematic illusory inferences with disjunction and quantifiers

Human reasoners are subject to fallacious inferences from very simple premises that have been described as tantamount to cognitive illusions (Walsh & Johnson-Laird, 2004; Khemlani & Johnson-Laird, 2009). We present new experiments that show that these phenomena are much more general and systematic than has previously been thought, including inferences from disjunctive premises and premises involving quantifiers. The novel illusory inferences we present are predicted by the erotetic theory of reasoning (Koralus and Mascarenhas, 2013). The key idea is that, by default, we reason by interpreting successive premises as questions and maximally strong answers to those questions, which generates the observed fallacies.

Phil Johnson Laird and Geoff Goodwin

The truth of conditionals

Given a basic conditional of the form, *if A then C,* with a sensible content, individuals usually list three cases as possible:

>   A and C

>   not-A and not-C

>   not-A and C.

This result corroborates the theory of mental models. In contrast, individuals often judge that the conditional is true only in the case of *A and C,* and that cases of *not-A* are irrelevant to its truth or falsity. This result corroborates other theories of conditionals. To try to resolve the discrepancy, we devised two new tasks. In the first, a 'collective' truth task, participants judged whether sets of assertions, such as:

>   if A then C

>   not-A

>   C

could all be true at the same time. In the second task, participants judged the truth or falsity of conditional predictions of an elliptical sort, such as: 'John will fire Charlie, if not Annie' in the light of what happened. The results of both tasks matched the three possibilities, thereby corroborating the model theory. They led us to a significant reinterpretation of the theory, which explains why individuals so often judge cases of *not-A* as irrelevant to the truth value of a conditional.

Denis Hilton, Laetitia Charalambides, Bertrand Fauré, and Christophe Schmeltzer

A societal exchange model of deontic rule-giving and reasoning

We propose an analysis of rights and duties in terms of societal   utility transfers. Rights involve a transfer of benefits from a social group to the individual, and duties imply a transfer of   benefits from the individual to the group. Rights and duties are   typically expressed by identifying a class of people (e.g. "The unemployed") to whom an authority (e.g., a government) representing a group gives an obligation, permission, interdiction or dispensation (e.g., registering for unemployment allowance). In the first study, using 34 deontic rules, we demonstrate that rules that are typically understood as rights (e.g. "The unemployed have a right to claim unemployment benefit") are typically perceived as benefiting the target class, whereas rules that are typically perceived as duties (e.g. "Train passengers are required to punch their tickets") are perceived as benefiting the group. Infringements of the rules are typically seen to be cases where the benefit is taken without being a member of the target class in the case of rights, or who fail to perform the required behaviour in the case of duties. In the second experiment, we show that ambiguous rules

(e.g., "Nurses who do the day shift...attend professional training.") are interpreted as rights if they benefit the addressee and as duties if they benefit the institution represented by the person giving the rule. In the final experiment, we show that cases of force majeure where situational factors modify the normal utility structure (e.g., a strike prevents the train company from transporting passengers who have a ticket). In such cases, a rule that is normally expressed as a duty will be more naturally expressed as a right (e.g., train passengers may punch their tickets), if it is in the interest of the passengers to be able to punch their ticket (e.g., to prove to their employer that they attempted to make the journey despite the strike). The third experiment finds strong support for the societal exchange model, as participants' expressions of rights and duties (and perceptions of rule violations) track the changes in the underlying utility structure induced by the force majeure context.

Shira Elqayam, Meredith R. Wilkinson, Valerie A. Thompson, Jonathan St.B.T. Evans, David E. Over

Inference from *is* to *ought* mediates moral judgement

The moral judgement literature makes a qualitative distinction between utilitarian and deontological moral judgements. Utilitarian or consequentialist judgements are based on calculations of the costs and benefits of the consequences from each course of action. Deontological judgements are based on moral rules, often articulated in terms of rights and duties. We hypothesised that consequentialist but not deontological judgement is mediated by deontic introduction (inference from 'is' to 'ought'). Participants were presented with moral dilemmas which allowed for either consequentialist or deontological judgement; for each of these we measured separately deontic conclusions and overall moral judgement. Results show that utilitarian judgement, but not deontological judgement, was affected by the same manipulations that affected deontic introduction.

IgorDouven
Moral Bookkeeping

Abstract: There is widespread agreement among philosophers about the Mens Rea Asymmetry (MRA), according to which praise requires intent, whereas blame does not. However, there is evidence showing that MRA is descriptively inadequate. We hypothesize that the violations of MRA found in the experimental literature are due to what we call 'moral compositionality', by which we mean that people evaluate the component parts of a moral problem separately and then reach an overall verdict by aggregating the verdicts on the component parts. We have subjected this hypothesis to the test and here report the results of our experiment

Matthew E. Roser, Jonathan St. B. T. Evans, Lauren S. Carroll, Nicolas A. McNair, Giorgio Fuggetta and Anna Kharko.

Within-trial repetitive transcranial magnetic stimulation affects belief bias in conditional reasoning.

Human reasoning is affected by numerous biases. When asked to decide on the logical validity of conditional (if-then) statements, many people exhibit a bias towards endorsing inferences drawn from believable statements as logically valid. This is known as belief bias and it is hypothesized to involve neural systems for working memory and inhibition. To investigate the neural sources of belief bias we used brief within-trial repetitive transcranial magnetic stimulation (rTMS), guided by neuronavigation, to transiently disrupt activity in frontal brain regions identified as active in a large functional magnetic resonance imaging (fMRI) experiment of belief bias in conditional reasoning. Group fMRI results were warped into individual-subject space to identify cortical targets. The inferior frontal gyri (IFG, associated with inhibition) and dorsolateral prefrontal cortex (DLPFC, associated with working memory) bilaterally, and vertex, were targeted with 500ms of stimulation at 10Hz and 80% of active motor threshold at two separate stages (major premise and minor premise) of trials in which participants were instructed to conclude on the logical validity of conditional statements. Logical validity was crossed with the (independently-rated) believability of conditional statements in a design which has been shown to reliably produce belief bias. The experiment was run in six blocks of 24 trials with each block corresponding to stimulation of a different cortical site, the order of which was randomised. At the stage of Major-premise presentation, at which believability of the statement can be assessed, rTMS over the IFG increased belief bias relative to a sham condition. Stimulation of DLPFC at this timepoint did not affect belief bias. At the later stage of Minor-premise presentation, at which premises can be integrated and logical reasoning is possible, rTMS over IFG and DLPFC had no effect upon belief bias. These results suggest that belief bias arises as a result of insufficient inhibition of beliefs based on real-world knowledge at an early stage of the reasoning process.

Vinod Goel

Lesions to Polar/Orbital Prefrontal Cortex Selectively Impair Reasoning about Emotional Material

While it is widely accepted that lesions to orbital prefrontal cortex lead to emotion related disruptions and poor decision-making, there is very little patient data on this issue involving actual logical reasoning tasks. We tested patients with circumscribed, focal lesions largely confined to polar/orbital prefrontal cortex (BA 10 & 11) (N=17) on logical reasoning tasks involving neutral and emotional content, and compared their performance to that of an age and education-matched normal control group (N=22) and a posterior lesion control group (N=11). Our results revealed a significant group by content interaction driven by a selective impairment in the polar/orbital prefrontal cortex group compared to healthy normal controls, in the emotional content reasoning trials. Subsequent analyses of congruent and incongruent reasoning trials indicated that this impairment was driven by the poor performance of patients with polar/orbital lesions in the incongruent trials. We conclude that the polar/orbital prefrontal cortex plays a critical role in filtering emotionally charged content from the material before it is passed on to the reasoning system in lateral/dorsal regions of prefrontal cortex. Where unfiltered content is passed to the reasoning engine, either as a result of pathology (as in the case of our patients) or as a result of individual differences, reasoning performance suffers.

G. Politzer & J. Baratgin.

Deductive arguments with uncertain premises and nonnumerical probability expressions.

Six arguments (and-elimination, and-introduction, or-introduction, and to if, or to if-not, contraposition) were studied using two kinds of context ("people in cities", e.g., Nicolas is Lyon and Marie is in Marseille, and "random draw", e. g., Pierre draws a black square chip). Participants were presented with premises with one of the following five levels of probability: "the chances are /very low /low /average /high /very high that [premise]" and then asked: "In your opininon, what are the chances that [conclusion]?" using a response format allowing to select any number of options among: greater than [premise level], just [premise level], smaller than [premise level]. For the five informative arguments, after correction for chance responding, 60% to 80% of the evaluations were coherent in the sense of de Finetti. For contrapositon there was a bias toward a decrease in probability. The rate of conjunction errors was limited to 20%.

David Over, Nicole Cruz, and Jean Baratgin

Recent studies of ifs and ands and ors and their significance

Recent studies of traditionally neglected inferences containing conjunctions, disjunctions, and conditionals have started to reveal significant points about conditionals. A prime example is centering: inferring if p then q from p & q. Centering is valid for the material conditional, the probability conditional, and Stalnaker / Lewis possible-worlds conditionals. It is invalid for some other possible-world theories, such as Kratzer's, and invalid for inferentialist accounts of conditionals, such as Douven's. We will make relevant points about the "suppression" of inferences and discuss the importance of further experiments on the validity of centering for counterfactuals and missing-link conditionals.

Nicole Cruz, Jean Baratgin, Mike Oaksford, and David Over

Centering and the meaning of conditionals.

The new probabilistic approach to the study of deductive reasoning generalises binary consistency to coherence, and binary validity to probabilistic validity, or p-validity. We extend Cruz et al. (2015), using the new approach to study neglected inferences between conjunctions, disjunctions, and conditionals, particularly centering: inferring if p then q from p & q. We assess the extent to which people comply with coherence, and whether they do so more when the inferences are not just probabilistically informative, but also p-valid. The experiments compare three hypotheses concerning the meaning of natural language conditionals: the material conditional interpretation equivalent to not-p or q, the probability conditional interpretation implying the Equation $P(\text{if } p \text{ then } q) = P(q|p)$, and the inferentialist

interpretation postulating the requirement of an inferential connection between p and q for a conditional to be believable.

Robert Mackiewicz, Monica Bucciarelli, Sangeet S. Khemlani, and P.N. Johnson-Laird

Eye movements as signs of kinematic mental simulations

We report two experiments showing that participants' eye movements reflect their kinematic mental simulations as they create informal algorithms. In both experiments, the participating students looked at a static picture of a railway track with a siding, as they tried to describe algorithms for rearranging the order of the cars in trains by shunting them to and from the siding. An eye tracking system recorded their eye movements during the task. In Experiment 1, participants imagined separate moves from different parts of the tracks (e.g. from left track to right track). Two independent judges saw pictures of a scan path for each individual move and were able to infer what type of move the participants were imagining in 76% of cases ($p < .001$). In Experiment 2, participants described the sequences of moves required to make complete rearrangements of the cars in a train. The difficulty of the task (as shown in the times of correct solutions) depended, not on the number of moves in the algorithm, but on its computational complexity. The judges identified correctly 48% of the moves ($p < .001$) from the eye tracking videos alone.

Sangeet Kehmlani, Geoff Goodwin and Phil Johnson-Laird

Causal relations from kinematic simulations

"Reasoners distinguish between different types of causal relations, such as causes, enabling conditions, and preventions. Psychologists disagree about the representations that give rise to the different relations, but agree that mental simulations play a central role in inferring them. We explore how causal relations are extracted from mental simulations. The theory of mental models posits that people use a kinematic simulation to infer possibilities. It predicts that the time it takes to infer a causal relation should correlate with the number of moves in a mental simulation. To test these two predictions, we adapted a railway domain designed to elicit mental simulations, and we devised problems in which reasoners had to infer causal relations from simulations of the movements of cars in this domain. Two studies corroborated the model theory's predictions. We discuss the results in light of recent theories of causation and mental simulation."

Célia Rasga , Ana Cristina Quelhas and Ruth M.J. Byrne

Autistic children's reasoning: counterfactual and false-belief inferences about reasons for actions

We report one experiment that examined children's reasoning about intentions. Children, aged 8 years and 10 years, read scenarios in which an observer witnessed an actor carrying out an action, and the actor had an initial reason for the action that was subsequently replaced by a different reason for the same action. There is a consensus that theory of mind's deficit can explain the severe difficulties in social interaction and communication often reported in

Autistic children. The ability to interpret others' actions in terms of mental states is critical to evaluate others and especially to form moral judgment. However, theory of mind impairments in people with High Functioning Autism and Asperger Syndrome are not always present. Our experiment was designed to study autistic children's reasoning about other people's intentions. The first aim was to exam whether children aged 8 years make more errors than children aged 10, in false-belief and counterfactual inferences about intentions, and whether children make more errors in inferences about a Desire-Obligation condition, compared to Obligation-Desire condition. We found that children aged 8 years made more mistakes, and more correct counterfactual inferences than false belief inferences about intentions. We also found that autistic children made more errors in Desire-Obligation condition, compared to Obligation-Desire condition, because their answers focused on the obligation. We discuss the implications of the results for alternative theories of counterfactual and false-belief reasoning.

Vittorio Girotto and Stefania Pighin

Errors in diagnostic reasoning: Reliable and unreliable remedies

A common view is that many errors in diagnostic reasoning depend on the format in which information is provided: Respondents fail to evaluate the positive predictive value of a test result, when they reason about single-event probabilities, but they perform better when they reason about natural frequencies. Accordingly, many organizations have promoted the use of natural frequencies for communicating to patients. Here, we test whether this recommendation is well founded. Past studies have typically tested convenience samples of educated respondents. In five studies, we have tested respondents sampled from the general public, including samples of pregnant women interested in prenatal testing, on realistic clinical scenarios. The results show that these respondents do not benefit from natural frequency information. By contrast, they succeed when they reason about test results expressed as numbers of chances, and complete non standard probability tasks. We discuss the theoretical and clinical implications of these findings.

Stephanie Dornschneider

Deciding (Not) to Kill. A Computational Analysis of the Reasoning Processes Connected to Political Violence.

Why do some individuals pick up arms against their states as opposed to others who live under the same conditions but engage in nonviolent activism? In search for an answer, I adopt Axelrod's cognitive mapping approach (1976), and systematically compare the belief systems of individuals who have engaged in violence and nonviolent activism. Half of the individuals are Muslim and come from a formerly authoritarian state – Egypt -, whereas the other half is non-Muslim and from a democratic state – Germany. To gather information, I conducted in-depth interviews in Arabic and German; and to analyze the interviews, I coded the

individuals' direct speech for beliefs, belief connections, and decisions for actions. In this way, I constructed *cognitive maps*, which trace the reasoning processes underlying the individuals' decisions to pick up arms; or to refrain from doing so. Since these reasoning processes are very complex, I moreover developed a computer program that systematically analyzes the cognitive maps. The program systematically evaluates the beliefs underlying the individuals' decisions to pick up arms. It also examines alternative worlds in which the individuals would *not* have decided to pick up arms by systematically modifying their beliefs. My findings put into perspective existing explanations of violence by showing that decisions to pick up arms are primarily motivated by beliefs that the state is aggressive; and not by beliefs related to Islam, economic deprivation, or mental illness.

Andrew Stewart and Matthew Haigh

Reasoning as we read: A psycholinguistic perspective on the processing of conditionals

The rapid, apparently effortless, comprehension of conditional statements in everyday life poses a particularly interesting set of problems for models of conditional processing. In the context of comprehension, understanding the dynamics associated with how and, more importantly, when, different sources of information come to have an influence is key. In the last few years, there has been increasing interest in investigating the comprehension processes that are employed within the first few hundred milliseconds of conditionals being encountered. In this talk I will review a number of these findings, present some new data, and outline potential implications with respect to how the reasoning processes involved in the comprehension of conditionals might be construed. Briefly, there is emerging evidence of a rapid divergence in how indicative vs. counterfactual conditionals are treated by the language comprehension system. There is also evidence that utility information, conditional probability information, the type of speech acts communicated by conditionals, and pragmatic factors associated with the contexts in which conditionals are uttered all rapidly influence the processing of conditionals. I will suggest that a 'search for meaning' perspective of comprehension allows aspects of comprehension processes to be viewed as reasoning processes (and vice versa), and will outline ways in which there is potential for the application of psycholinguistic methods that focus on the temporal dynamics of comprehension (e.g., eye-tracking, and the examination of event-related brain potentials) to help us understand conditional processing.

Marta Couto , Ana Cristina Quelhas and Ruth M.J. Byrne

Counterexamples in Reasoning about Advice Conditionals: Tips and Warnings

We report two experiments that examine how people reason with advice conditionals that contain tips such as 'if you study more your grades will improve', and warnings, such as, 'if you stop exercising you will gain weight'. Experiment 1 showed that participants readily generated counterexamples to tips and warnings, and they generate more counterexamples to tips than warnings. Experiment 2 showed that participants made very few MP and MT

inferences from tips and warnings, and fewer MP and MT inferences than AC and DA inferences. The experiment also showed that generating disablers suppressed the MP and MT inferences, but generating alternatives had few effects. The experiments show that people readily retrieve alternative and disabling counterexamples to these mild inducements, and they make few inferences from them.

Henrik Singmann, Karl Christoph Klauer, David Over

Testing the empirical adequacy of coherence as a norm for conditional inferences.

There has been a major shift in research on human reasoning toward Bayesian and probabilistic approaches, which has been called a new paradigm. The new paradigm sees most everyday and scientific reasoning as taking place in a context of uncertainty, and inference is from uncertain beliefs and not from arbitrary assumptions. We present results from one published (Singmann, Klauer, & Over, 2014) and one new experiments testing whether or not participants responses to conditional inferences are coherent (in the sense of Pfeifer & Kleiter, 2005, 2010) using a novel fully probabilized conditional reasoning task. Results show that, while we find above chance coherence for MP and DA, participants are not generally coherent.