

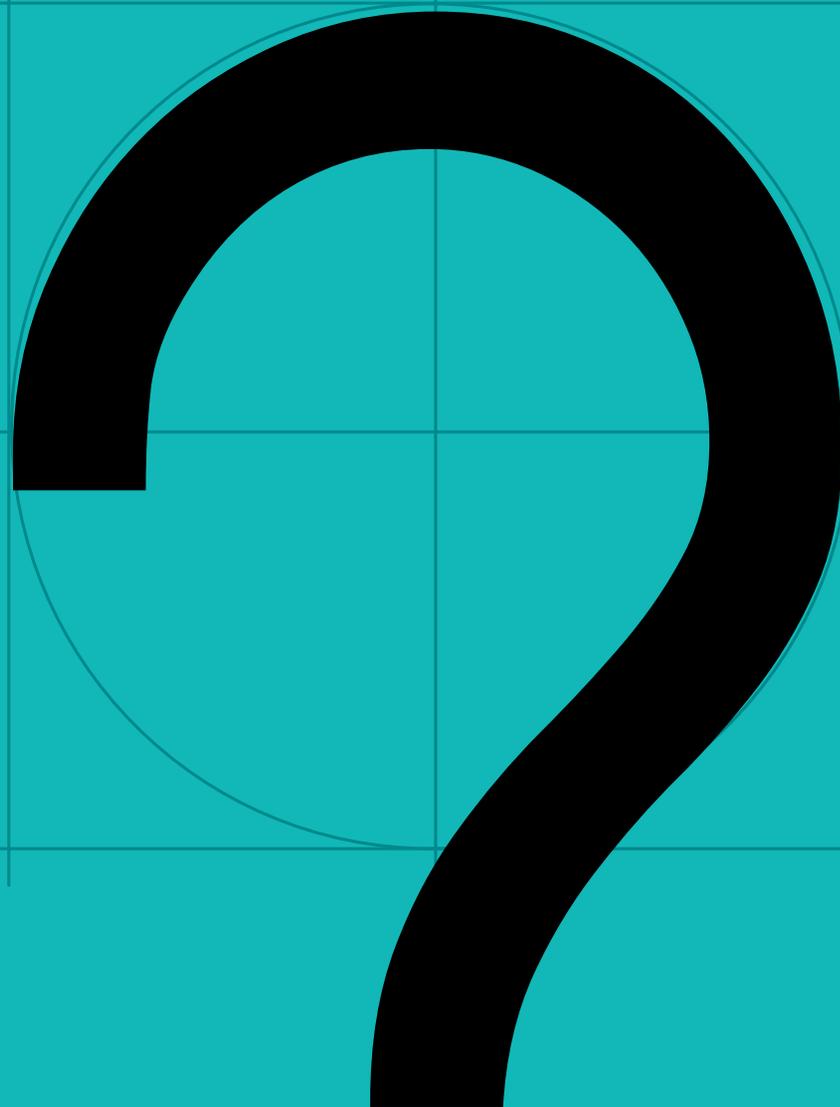
**EPSRC**

Engineering and Physical Sciences  
Research Council



University of  
**BRISTOL**

If found, please return to the person  
above or to the reception desk.

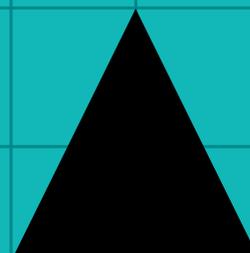
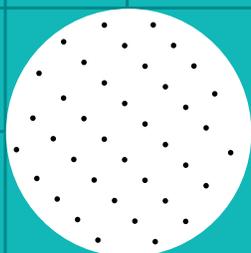


**KEYNOTES**

Gerd Gigerenzer, Jeff Schall,  
Tom Griffiths, Antonio Rangel

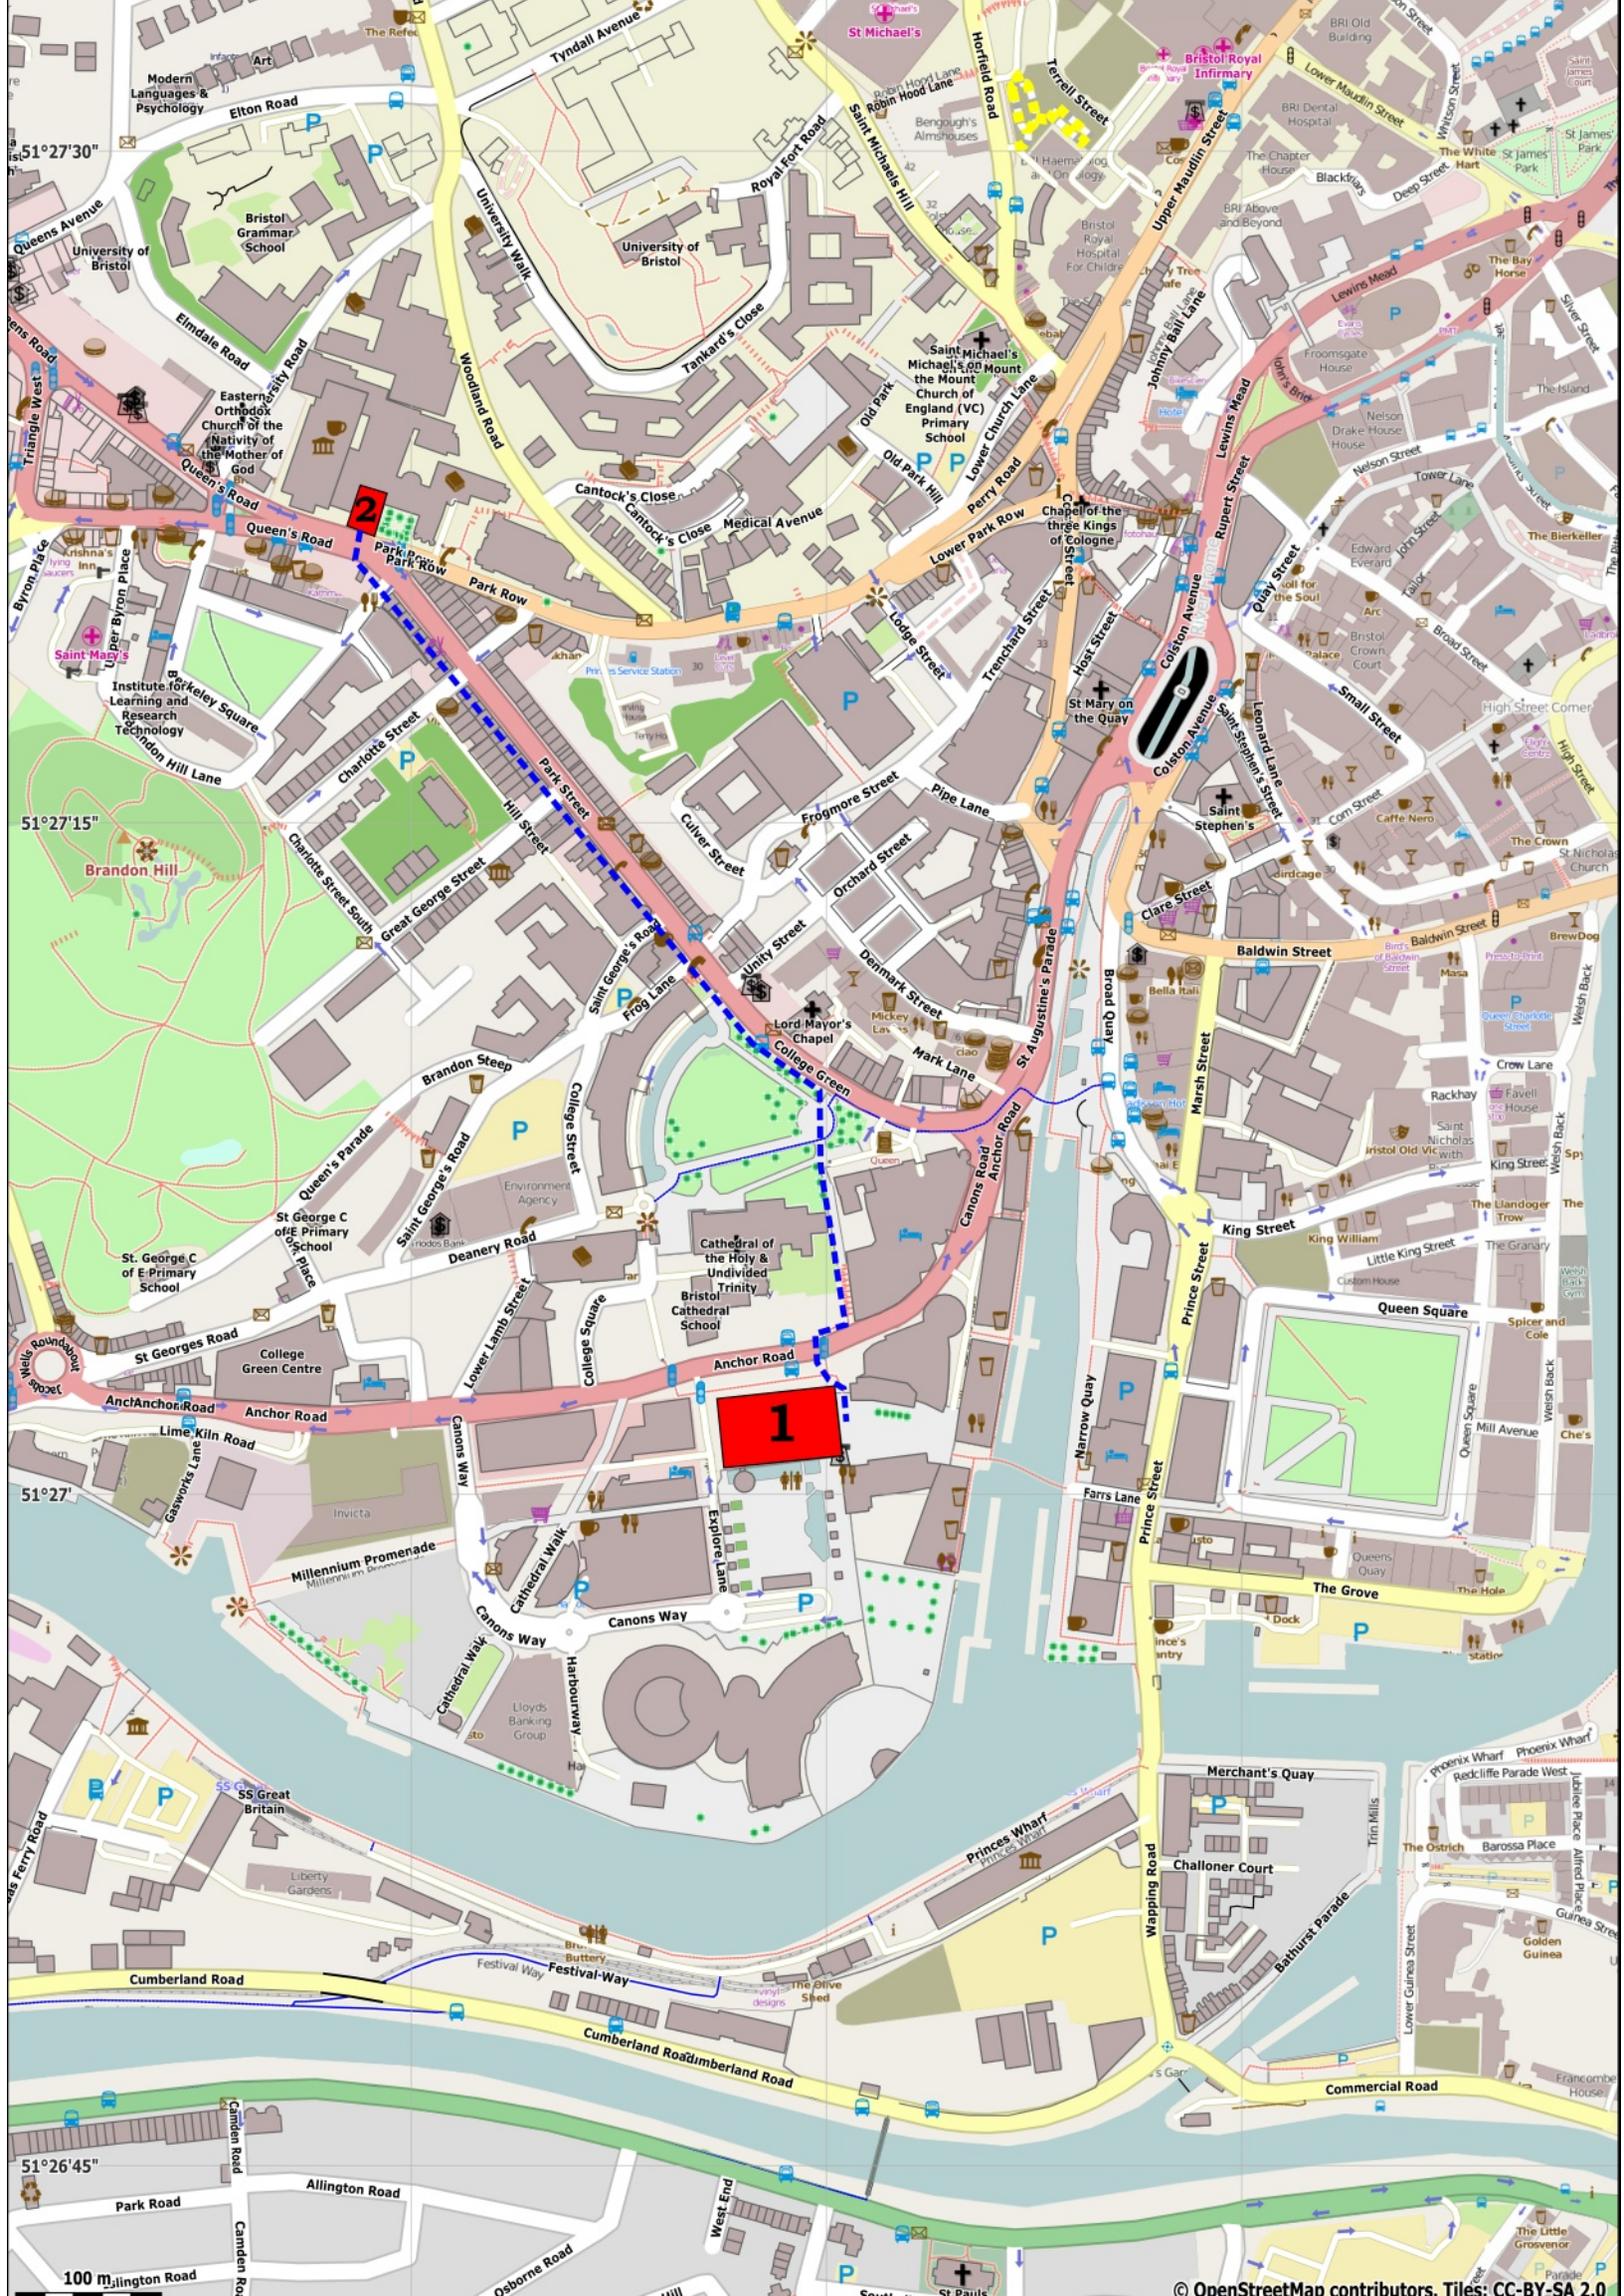
**NEW PERSPECTIVES**

Eric-Jan Wagenmakers,  
Nick Chater, Daniel Wolpert



**DECISION MAKING BRISTOL 2014**

THEORETICAL AND EMPIRICAL RESEARCH IN DECISION-MAKING



51°27'30"

51°27'15"

51°27'

51°26'45"

1

2

This conference brings together researchers from a wide range of disciplines who have an interest in decision making. We think this is a particularly exciting and important time for research in this area: a time where there is real potential for a convergence between theory, experimental work, applied research and policy to affect everyday decisions that impact on health and well being. We hope that you will take this opportunity to hear about new discoveries and perspectives from across the wide range of disciplines represented and hope that this will inspire your own research.

The conference is part funded by a grant from the Engineering and Physical Sciences Research Council (UK) who fund the Decision-making group here in Bristol University (<http://www.bris.ac.uk/decisions-research/>).

We really hope that you enjoy the conference and find a little time to explore our beautiful city. Our organising team and volunteers are on hand to help you - you can spot us in the red t-shirts.

On behalf of the organising team and the members of the Decision-making group here in Bristol - welcome!

**Iain Gilchrist** and **David Leslie**  
Co-Directors Bristol Decision-Making Group

Organising Team:  
Roland Baddeley, Rafal Bogacz, Rosie Clark, Adnane Ez-Zizi,  
Simon Farrell, Rebecca Floyd, Andreas Jarvstad, Casimir Ludwig,  
Gaurav Malhotra, Alice Mason, John McNamara and Sally-Ann Parry

# LIST OF PARTICIPANTS

MARLÈNE	ABADIE	RONG	GUO	TAKAO	NOGUCHI
LUIGI	ACERBI	CLAIRE	HALES	LUANA	NUNES
SILVIO	ALDROVANDI	STEPHEN	HEAP	KARSTEN	OLSEN
FLORIAN	ARTINGER	CRAIG	HEDGE	PHILLIP	PÄRNAMETS
ROLAND	BADDELEY	ANDREW	HIGGINSON	PAULA	PARPART
TIMOTHY	BALLARD	THOMAS	HILLS	SALLY-ANN	PARRY
DAN	BANG	JANINE	HOFFART	ALVIN	PASTORE
ROBERT	BEIGLER	JANINA	HOFFMANN	GILES	PAVEY
ULRIK	BEIERHOLM	AUDREY	HOUILLON	ELIZABETH	PAUL
SHIRLY	BERENDS	ANDREW	HOWES	REBECCA	PIKE
SUNDEEP	BHATIA	ALASDAIR	HOUSTON	ANGELO	PIRRONE
ROBERT	BIEGLER	LAURENCE	HUNT	GERIT	PFUHL
VALERIO	BISCIONE	TOM	JAHANS-PRICE	GANNA	POGREBNA
WOUTER	BOEKEL	ANDREA	ISONI	LIAM	POLLOCK
RAFAL	BOGACZ	PERKE	JACOBS	ANDREA	POLONIOLI
ALINE	BOMPASS	ANDREAS	JARVSTAD	DIANE	POMEROY
JEFFREY	BOWERS	HENRY	JORDAN	ROSIE	PRIOR
NICOLA	BOWN	HANNAH	JULLIENNE	MIKAEL	PUURTINEN
NEIL	BRAMLEY	DAVID	KELSEY	TIM	RAKOW
DANIEL	BRAUN	MAX C.	KEUKEN	ANTONIO	RANGEL
JULIA	BRETTSCHNEIDER	METTE	KJER KALTOFT	PETER	RIEFER
GORDON	BROWN	TOMAŽ	KOLAR	JOHANNES	RITTER
JOANNA	BURCH-BROWN	TATIANA	KORNIENKO	DAVID	RONAYNE
EGLE	BUTT	NATHAN	LEPORA	ADAM	SANBORN
NICK	CHATER	DAVID	LESLIE	JEFF	SCHALL
ROSIE	CLARK	STEPHAN	LEWANDOWSKY	MICHELLE	SCHALL
MATTHEW	COLLETT	KEVIN	LLOYD	ANYA	SKATOVA
JOHN	COOK	BRADLEY	LOVE	DYLAN	SOUBIEN
ANTHONY	CRUICKSHANK	SHENGHUA	LUAN	SARAH	SMITH
GILLES	DE HOLLANDER	ELLIOT	LUDVIG	TOM	STAFFORD
BARNEY	DOBSON	CASIMIR	LUDWIG	SANDRA	STARKE
CAROLINA	DORAN	JIE-YU	LV	JACK	STECHER
JACK	DOWIE	JASMIN	MAHMOODI	HRVOJE	STOJIC
GILLES	DUTILH	NISHANTHA	MALALASEKERA	NEIL	STEWART
ULLRICH	ECKER	GAURAV	MALHOTRA	BRIONY	SWIRE-THOMPSON
TOBIAS	ENDRESS	DAVE	MALLPRESS	PETROC	SUMNER
ADNANE	EZ-ZIZI	ENCARNI	MARCOS	FLORIAN	TANDONNET
RICHARD	FAIRCHILD	ALICE	MASON	PETE	TRIMMER
A. ALDO	FAISAL	ANJALI	MAZUMDER	JAMES	TRIPP
GEORGE	FARMER	JOHN	MCMAMARA	KONSTANTINOS	TSETSOS
TIM	FAWCETT	EUGENE	MCSORLEY	HENK	VAN DEN BERG
CAROLINA	FEHER DA SILVA	LIANNE	MEAH	RONALD	VAN DEN BERG
REBECCA	FLOYD	SOPHIE	MEAKIN	ERIC-JAN	WAGENMAKERS
KERSTIN	GIDLÖF	GEOFFREY	MÉGARDON	LUKASZ	WALASEK
GERD	GIGERENZER	MICHAEL	MENDL	LEONARDO	WEISS-COHEN
IAIN	GILCHRIST	JOLYON	MILES-WILSON	JAN K.	WOIKE
RUSSELL	GOLMAN	RANI	MORAN	DANIEL	WOLPERT
JOHN	GOUNTAS	TIMOTHY	MULLETT	JIAXIANG	ZHANG
TOM	GRIFFITHS	JERYL	MUMPOWER		
JOHN	GROGAN	PHILLIP	NEWALL		

- 6 Information
- 8 Conference Schedule
- 13 **Gerd Gigerenzer**  
Simple Heuristics for a Complex World
- 14 **Contributed Talks: Day 1**
- 33 **EJ Wagenmakers**  
Reinforcement Learning  
and the Iowa Gambling Task
- 35 **Jeff Schall**  
Decisions, accumulators and neurons:  
How secure a bridge?
- 36 **Contributed Talks: Day 2**
- 57 **Nick Chater**  
Virtual bargaining: A micro-foundation  
for social decision making and interaction
- 59 **Tom Griffiths**  
Priors and processes  
in Bayesian models of cognition
- 60 **Contributed Talks: Day 3**
- 81 **Daniel Wolpert**  
Probabilistic models of sensorimotor  
control and decision making
- 83 **Antonio Rangel**  
The neuroeconomics of simple choice
- 84 **Contributed Talks: Day 4**
- 98 **Poster Session 1**
- 107 **Poster Session 2**

## Conference Venue

The majority of DMB 2014 activities take place in the Bristol science museum, At-Bristol, located on Anchor Rd, Harbourside, Bristol BS1 5DB (see map reference 1). The exception to this is the conference dinner, which will take place in the Great Hall of the beautiful Wills Memorial Building, Queen's Road (map reference 2).

### Practical Information relating to The Venue

*Parking:* The nearest car park is Millennium Square Car Park, a safe and secure car park adjacent to At-Bristol. It is an underground car park, open 24 hours a day with on site security and CCTV cameras. The car park is situated underneath Millennium Square just a few minutes walk from At-Bristol. There is parking for 550 cars. There is a 2m height restriction. Using Sat Nav Enter BS1 5LL as your destination.

*Toilets:* As you enter the main Rosalind Franklin Room the toilets are through the double doors at the far end of the room. There are two accessible toilets, one at each end of the room.

*Fire Alarm:* If the fire alarm is activated a loud siren will sound. Please leave via the pink or purple staircase at each end of the room, following the green fire exit signs. At-Bristol staff will guide guests out of the building to the assembly point. The assembly point is on Anchor Square (the cobbled square in front of the At-Bristol building). At-Bristol staff will inform guests when it is safe to re-enter the building.

*Terraces:* There is a smoking area on the North Terrace. Smoking is not allowed anywhere else within At-Bristol. Please note terraces become slippery when wet. If in use please take care.

### Registration Desk

The registration desk can be found in the entrance to the Rosalind Franklin Room, situated on the first floor of At-Bristol. Please enter At-Bristol through the main, ground floor entrance. You will be directed to the stairs or lift to the first floor. The registration desk will open at 9am on Tuesday 9 September and will remain open for late registrations and queries throughout the conference. A cloakroom facility will also be available.

## Conference Name Badge

Participants are kindly asked to wear the conference badge at all times during the conference. It entitles them to participate in all activities and is required for entry to the Great Hall for the conference dinner.

## WiFi Internet Access

Free wireless internet access is available in the conference venue. You can connect by entering

WiFi Code: **At-Bristol Events**

Password: **Max-hertz!WE11**

If you have any problems connecting, please go to the registration desk.

## Coffee Breaks and Lunches

Coffee breaks and lunches are served in the Maurice Wilkins Room. The timings of lunches vary during the conference so please check the schedule carefully.

## Conference Dinner

The conference dinner will be held on Tuesday, 9th September in the beautiful Wills Memorial building (map reference 2). Delegates are invited to arrive at 19h15 for a welcome drink. Guests will be seated at 19h45; dinner will be served at 20h00.

If you are walking, the map shows how to get there on foot from the conference venue. There is no parking at the venue itself; limited parking can be found in the streets surrounding the venue and at Berkeley Square.

There are several bus services which stop in the Park Street / Triangle / Whiteladies Road area including the 1, 2, 8, 9, 40, 40A and 41. Leave the bus at the Berkeley pub on Queen's Road which is opposite Wills Memorial Building.

## Drinks Reception

The drinks reception will be held on Thursday, 11th September, directly following the last talk of the day. Delegates are asked to make their way down to the ground floor where they will also be able to view the exhibits at At-Bristol.

## Instructions to Presenters

### Talk Presentations

Talks last no longer than 22 minutes, followed by 8 minutes of discussion. These time limits will be strictly enforced so that participants can switch sessions to hear specific talks. Each session will have a dedicated laptop that speakers can use to connect to the projector. Speakers will also be able to connect their own machines to the projector using a VGA cable. If their laptops require any other connector (HDMI, Display port, etc), the speakers are required to bring their own adapter to a VGA connector. Please bring your presentation on a USB stick a minimum of 10 minutes before the beginning of the session. If using your own laptop, please bring it and test your presentation 15 minutes before the session starts.

## Poster Presentations

Poster sessions will be held on Wednesday and Thursday 15h30 – 16h45. Poster boards will be set up in the Rosalind Franklin Room. Each poster board will be numbered and you are asked to put your poster on the allocated board. Material to attach your poster will be provided. On the day of your presentation please put up your poster by 10h30 and remove it after the last session of the day.

## Practical Information

### How to get to the conference

*By Air:* The nearest airport to the venue is Bristol International Airport. From the airport, you can either take the bus (Bristol Flyer Express LInk) to Bristol Coach Station or a taxi (Checker Cars) directly to At-Bristol. If you're flying to London Heathrow, you can either get a National Express bus directly to Bristol Coach Station or train (via Paddington) to Bristol Temple Meads.

*By Train:* The city of Bristol is served by two major train stations - Bristol Temple Meads and Bristol Parkway. The nearest station to the venue is Bristol Temple Meads, Bristol BS1 6QF. From the train station you can either get a taxi to At-Bristol or any of the First Buses that come to the City Centre and walk from the Centre to At-Bristol (2-3 minutes on foot).

*By Bus:* The closest bus station is Bristol Coach Station on Marlborough Street (top right corner of map on the inside front cover). You can walk from the coach station to At-Bristol (10-15 minutes on foot) or take a taxi available at the exit of the coach station.

### Getting Around

*Walking:* We find the best way to get around town is to walk. At-Bristol is at the centre of town and a walking distance from public transport, cafés and hotels.

*Bus:* Several bus companies run buses around Bristol. On all of them, you pay the driver when you get on. Route details and timetable can be found using <http://www.travelwest.info/bus>.

*Taxi:* Blue taxi cabs can be taken without pre-booking. These can be found at various points around the city (including the train station).

Private hire: V-Cars: 0117 925 2626

### Emergency

*Emergency Phone Number:* 999. This number goes through to the police, the fire brigade and ambulance services. The nearest accident and emergency department is situated at the Bristol Royal Infirmary, Marlborough Street, Bristol, Avon, BS2 8HW.

Non urgent medical advice can be sought from the NHS by calling 111

## About Bristol

### Attractions

We hope you will have time to explore the beautiful city of Bristol whilst you are here. To find out what Bristol has to offer, take a look at <http://visitbristol.co.uk>.

*Tourist information:* Bristol Visitor Information Centre is located on the waterfront at the E Shed, 1 Canons Road, Bristol, BS1 5TX (open: Mon-Sat 10am – 5pm and Sun 11am – 4pm).

### Eating Out

There are some amazing restaurants in Bristol - we have award-winning restaurants, restaurants on boats, gastro pubs, cafés and bistros of all varieties. You will find a wealth of eateries within walking distance of the venue in Millennium Square and along the water front as well as on Park street and Whiteladies Road.

### Our Favourite Restaurants

Bordeaux Quay: *V-Shed, Cannons Way, BS1 5UH*

Cherry Duck: *3 Queen Quay, Welshback, BS1 4SL*

Glass Boat: *Welsh Back, BS1 4SB*

Hotel du Vin: *Narrow Lewins Mead, BS1 2NU*

Sergio's: *1-3 Frogmore Street, BS1 5NA*

Spyglass: *Welsh Back, BS1 4SB*

The Olive Shed: *Princes Wharf, BS1 4RN*

The Stable (*good pizza*): *Harbourside, BS1 5UH*

### Our Favourite Cafés

Arnolfini: *16 Narrow Quay, BS1 4QA*

Boston Tea Party: *75 Park St, BS1 5PF*

Falafel King (*van*): *Harbourside, City Centre*

Mud Dock: *40 The Grove, BS1 4RB*

The Folk House: *40A Park St, BS1 5JC*

Watershed: *1 Canons Road, BS1 5TX*

### Our Favourite Bars/Pubs

Grain Barge: *Mardyke Wharf, Hotwell Rd, BS8 4RU*

Harveys Cellars: *12 Denmark St, BS1 5DQ*

Llandoger Trow: *King St, BS1 4ER*

Milk Thistle (*book ahead*): *Colston Avenue, BS1 1EB*

The Cornubia (*live music*): *142 Temple St, BS1 6EN*

The Cottage Inn (*by boat*): *Cumberland Rd, BS1 6XC*

The Hope & Anchor: *38 Jacob's Wells Rd, BS8 1DR*

The Old Duke (*for Jazz*): *45 King St, BS1 4ER*

The Ostrich Inn: *Lower Guinea Street, BS1 6TJ*

The White Lion: *Sion Hill, BS8 4LD*

### Live Music

Fleece: <http://www.thefleece.co.uk/>

Louisiana: <http://www.thelouisiana.net/>

# DAY 1: TUESDAY, 9TH SEPTEMBER

	Rosalind Franklin Room	Francis Crick Room	James Watson Room
09h00 - 10h30	Registration Opens		
10h30 - 10h45	Welcome Remarks		
	Chair: John McNamara		
10h45 - 11h45	Keynote: <b>Gerd Gigerenzer</b>		
11h45 - 12h00	Coffee Break		
	Chair: Rafal Bogacz	Chair: Casimir Ludwig	Chair: Joanna Burch-Brown
12h00 - 12h30	<b>Adam Sanborn</b> Decision making with ordered discrete responses	<b>Rong Guo</b> Amygdala activity correlates with the transition from saliency-based to reward-based choices	<b>John Cook</b> Rational irrationality: modeling climate change belief polarization using Bayesian networks
12h30 - 13h00	<b>Gilles Dutilh</b> The adaptive shifts of decision strategies	<b>Nishantha Malalasekera</b> Neuronal correlates of value based decision making differ with information gathering strategies	<b>Joanna M. Burch-Brown</b> Diversification, risk and climate change: Towards a principled guide to portfolio optimisation in climate investment
13h00 - 14h15	Lunch		
	Chair: Rafal Bogacz	Chair: Andreas Jarvstad	Chair: David Leslie
14h15 - 14h45	<b>Luana Nunes</b> Ethological decision making with non-stationary inputs using MSPRT-based mechanisms	<b>Leonardo Weiss-Cohen</b> Incorporating conflicting descriptions into decisions from experience	<b>Jan K. Woike</b> Pyrrhic victories in public-goods games: when relative comparisons matter more than absolute payoffs
14h45 - 15h15	<b>Gaurav Malhotra</b> Changing decision criteria in rapid & slow decisions: do people behave optimally?	<b>Jack Stecher</b> Description and Experience Based Decision Making: An Experimental and Structural Estimation Approach to the Description-Experience Gap	<b>Andrea Isoni</b> Preference and Belief Imprecision in Games
15h15 - 15h45	<b>Rani Moran</b> The Collapsing Confidence Boundary Model: A Unified Theory of Decision, Confidence, and Response-Latencies	<b>Thomas Hills</b> Choice-Rich Environments Induce Risk-Taking and Reduce Choice Deferral	<b>Liam Pollock</b> Reduced uncertainty improves children's coordination, but not cooperation
15h45 - 16h15	<b>Petroc Sumner</b> Competitive accumulator models explain why there is often no correlation between error and RT measures of inhibitory control	<b>Janine Christin Hoffart</b> The Influence of Sample Size on Decisions from Experience	<b>Philip Newall</b> Failing to Invest by the Numbers
16h15 - 16h45	Coffee Break		
16h45 - 17h30	New Perspectives: <b>Eric-Jan Wagenmakers</b>		
19h15 - 23h00	Conference Dinner, Wills Memorial Building		

## DAY 2: WEDNESDAY, 10TH SEPTEMBER

	Rosalind Franklin Room	Francis Crick Room	James Watson Room
	Chair: Iain Gilchrist		
09h00 - 10h00	Keynote: <b>Jeff Schall</b>		
10h00 - 10h30	Coffee Break		
	Chair: Iain Gilchrist	Chair: Jeff Bowers	Chair: David Leslie
10h30 - 11h00	<b>Eugene McSorley</b> Dissociation between the impact of evidence on eye movement target choice and confidence judgements	<b>Neil Bramley</b> Acting informatively - How people learn causal structure from interventions	<b>Florian Artinger</b> Risk preferences in context
11h00 - 11h30	<b>Andreas Jarvstad</b> The neural substrate of eye-movement control: Inhibition, selection and choice	<b>Hrvoje Stojic</b> Function learning while making decisions	<b>Andrea Polonioli</b> Stanovich's challenge to the "adaptive rationality" project: an assessment
11h30 - 12h00	<b>Aline Bompas</b> A unified framework for speeded saccadic and manual decisions	<b>Alice Mason</b> Chance based uncertainty of reward improves long-term memory	<b>Ganna Pogrebna</b> Ambiguity attitude in coordination problems
12h00 - 12h30	<b>Philip Pärnamets</b> Eye gaze reflects and causes moral choice	<b>Tim Fawcett</b> 'Irrational' behaviour can arise from decision rules adapted to complex environments	<b>David Kelsey</b> Ambiguity in Coordination Games
12h30 - 14h00	Lunch		
	Chair: Iain Gilchrist	Chair: Roland Baddeley	Chair: Steve Lewandowsky
14h00 - 14h30	<b>Jiaxiang Zhang</b> Decision-making deficits in neurodegenerative diseases	<b>Gordon D. A. Brown</b> Relative Rank Theory	<b>Ullrich Ecker</b> The Effects of Misinformation on Reasoning and Decision Making
14h30 - 15h00	<b>Encarni Marcos</b> The variance of neuronal activity in dorsal premotor area predicts behavioral performance: The modulation of behavior and neural activity by trial history	<b>Tatiana Kornienko</b> Nature's Measuring Tape: A Cognitive Basis for Adaptive Utility	<b>Sudeep Bhatia</b> Naturalistic Multi-Attribute Choice
15h00 - 15h30	<b>Gilles de Hollander</b> Large Cortical Networks in a Small Nucleus: a 7T fMRI Study on Limbic, Cognitive and Associative Subparts in The Subthalamic Nucleus During Decision-making	<b>Sarah Smith</b> Is financial risk attitude entirely relative?	<b>Julia Brettschneider</b> Decision-making in adjuvant cancer treatment choices based on complex information, including genomic recurrence risk scores
15h30 - 16h45	Poster Session 1 (Coffee Available)		
16h45 - 17h30	New Perspectives: <b>Nick Chater</b>		

## DAY 3: THURSDAY, 11TH SEPTEMBER

	Rosalind Franklin Room	Francis Crick Room	James Watson Room
	Chair: David Leslie		
09h00 - 10h00	Keynote: <b>Tom Griffiths</b>		
10h00 - 10h30	Coffee Break		
	Chair: Kevin Lloyd	Chair: Tim Fawcett	Chair: Gaurav Malhotra
10h30 - 11h00	<b>Bradley C. Love</b> Optimal Teaching with Limited-Capacity Decision Makers	<b>Shenghua Luan</b> Threshold-based Inference	<b>Christophe Tandonnet</b> Decision making and response implementation: the missing link to constrain models
11h00 - 11h30	<b>Adnane Ez-zizi</b> Learning to solve working memory tasks	<b>Konstantinos Tsetsos</b> The value of economic irrationality: why and when being irrational is advantageous	<b>Valerio Biscione</b> Rate Domain Distributions in Simple Reaction Time Task
11h30 - 12h00	<b>Carolina Feher da Silva</b> A Simple Method to Characterize the Exploration-Exploitation Trade-Off in Binary Decision Making	<b>Perke Jacobs</b> A Competitive Test of Satisficing Heuristics in Choice From Serially Dependent Sequences	<b>Angelo Pirrone</b> Evolutionary and computational considerations on the Drift Diffusion Model
12h00 - 12h30	<b>John Grogan</b> Dopamine affects encoding and retrieval of positive and negative reinforcement learning	<b>Paula Parpart</b> Heuristics as a special case of Bayesian inference	<b>Takao Noguchi</b> The Dynamics of Evidence Accumulation and Choice
12h30 - 14h00	Lunch		
	Chair: Casimir Ludwig	Chair: Gaurav Malhotra	Chair: Andrew Higginson
14h00 - 14h30	<b>Timothy Mullett</b> Visual attention in choice: Do individuals pay more attention to more important information?	<b>Elliot Ludvig</b> Big, fast, and memorable: Increased gambling in risky decisions from experience	<b>Claire Hales</b> Validation of the diffusion model for investigating the decision making processes underlying cognitive affective bias in rodents
14h30 - 15h00	<b>Rosie Clark</b> The effect of spatial probability and reward on response time	<b>Lukasz Walasek</b> Loss aversion is a property of the experimental design, not the participant	<b>Pete Trimmer</b> Optimal moods in behavioural models
15h00 - 15h30	<b>Luigi Acerbi</b> A matter of uncertainty: Optimality and sub-optimality in sensorimotor learning	<b>David Ronayne</b> Multi-Attribute Decision-by-Sampling	<b>Carolina Doran</b> Economical decision making by <i>Temnothorax albipennis</i> ant colonies
15h30 - 16h45	Poster Session 2 (Coffee Available)		
16h45 - 17h30	New Perspectives: <b>Daniel Wolpert</b>		
17h30 - 19h30	Drinks reception and @-Bristol exhibits		

## DAY 4: FRIDAY, 12TH SEPTEMBER

---

	Rosalind Franklin Room	Francis Crick Room	James Watson Room
	Chair: Casimir Ludwig		
09h00 - 10h00	Keynote: <b>Antonio Rangel</b>		
10h00 - 10h15	Coffee Break		
	Chair: Roland Baddeley		Chair: Pete Trimmer
10h15 - 10h45	<b>Daniel A. Braun</b> Ellsberg's paradox in sensorimotor learning		<b>Russell Golman</b> Polya's Bees: A Model of Decentralized Decision Making
10h45 - 11h15	<b>A. Aldo Faisal</b> A learning rule for optimal feedback control that predicts continuous action selection in complex motor tasks		<b>Peter S. Riefer</b> Coordination in large groups: What would we do if zombies attack?
11h15 - 11h30	Coffee Break		
	Chair: Nathan Lepora	Chair: Joanna Burch-Brown	Chair: John McNamara
11h30 - 12h00	<b>Rafal Bogacz</b> To act or not to act: Learning the value of not acting	<b>Andrew Howes</b> Computationally Rational Decision Making	<b>Dan Bang</b> Confidence matching in social interaction
12h00 - 12h30	<b>M.C. Keuken</b> The role of the subthalamic nucleus in multiple alternative perceptual decision-making revealed by 7T structural and functional MRI	<b>Laurence Hunt</b> Sub-optimal information search strategies during multi-attribute choice	<b>Rebecca Floyd</b> Better Together - Understanding Collaborative Decision Making
12h30 - 13h00	<b>Nathan Lepora</b> Perceptual decision making in robotics and neuroscience	<b>George Farmer</b> The attraction effect is rational given uncertainty in expected value calculation	<b>Karsten Olsen</b> Keeping track of others' choices: Predicting changes in the perceptual decisions and decision confidence of other individuals
13h00 - 13h15	Closing remarks		

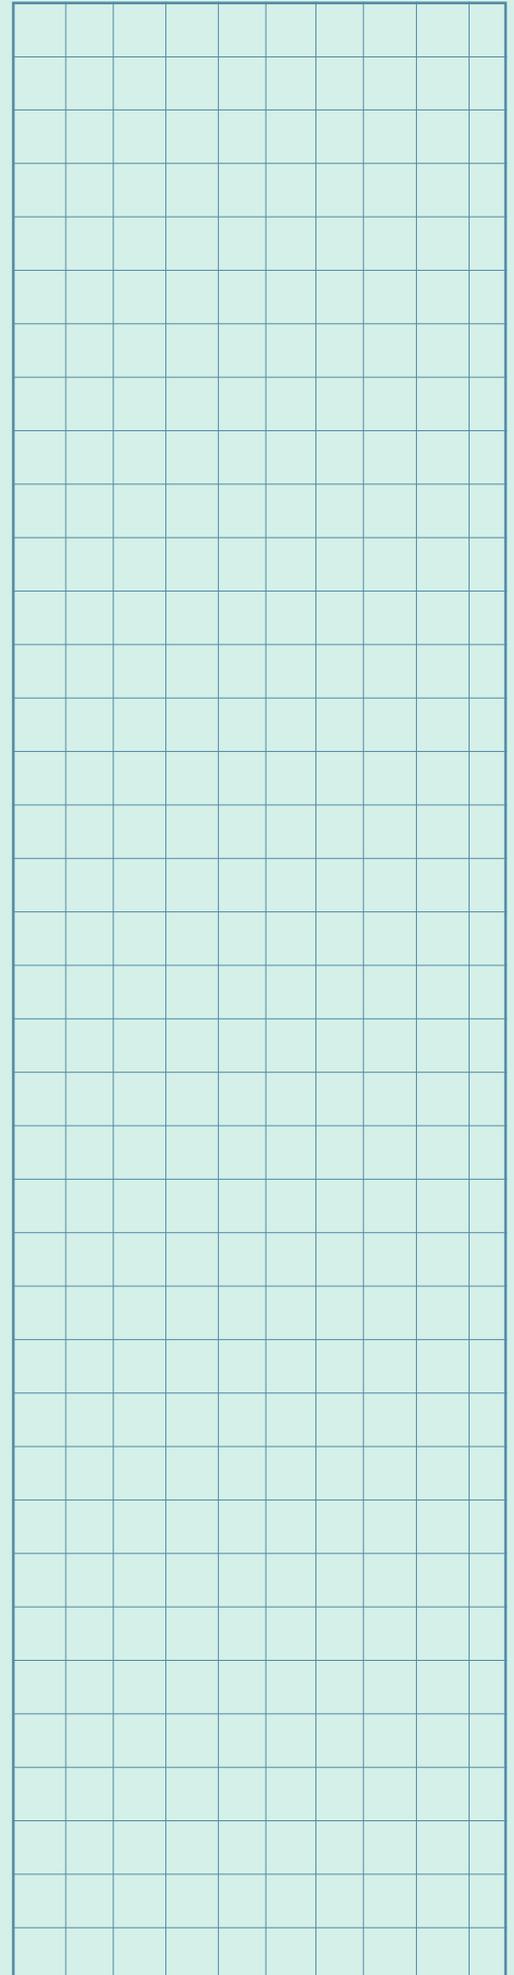


# GERD GIGERENZER

## SIMPLE HEURISTICS FOR A COMPLEX WORLD

GERD GIGERENZER *Max Planck Institute for Human Development;  
Director of the Harding Center for Risk Literacy in Berlin*

How to invest? Whom to trust? Complex problems require complex solutions – so we might think. And if the solution doesn't work, we make it more complex. That recipe is perfect for a world of known risks, but not for an uncertain world, as the failure of the complex forecasting methods leading to the 2008 financial crisis illustrates. In order to reduce estimation error, good inferences under uncertainty counter-intuitively require ignoring part of the available information. Less can be more. Yet although we face high degrees of uncertainty on a daily basis, most of economics and cognitive science deals exclusively with lotteries and similar situations in which all risks are perfectly known or can be easily estimated. In this talk, I invite you to explore the land of uncertainty, where mathematical probability is of limited value and people rely instead on simple heuristics, that is, on rules of thumb. We meet Homo heuristicus, who has been disparaged by many psychologists as irrational for ignoring information—unlike the more diligent Homo economicus. In an uncertain world, however, simple heuristics can be a smart tool and lead to even better decisions than with what are considered rational strategies. The study of heuristics has three goals. The first is descriptive: to analyze the heuristics in the “adaptive toolbox” of an individual or an institution. The second goal is normative: to identify the ecological rationality of a given heuristic, that is, the structures of environments in which it succeeds and fails. The third goal is engineering: to design intuitive heuristics such as fast-and-frugal trees that help physicians make better decisions.



## DECISION MAKING WITH ORDERED DISCRETE RESPONSES

ADAM SANBORN *University of Warwick*

ULRIK BEIERHOLM *University of Birmingham*

Analyses of decision-making behavior have often compared people against a normative standard, assuming that people are attempting to maximize their expected rewards. The Bayesian formulation of this problem prescribes what should be done: people's prior beliefs are combined with the likelihood of the observed stimulus given the response via Bayes rule to produce posterior beliefs about which response is correct. The posterior beliefs are then integrated with the loss function, which describes the gain in value for a given choice for each possible true answer, to determine which response has the highest expected value.

When people are asked to choose from among a small set of discrete responses, the most straightforward approach is to treat the responses as unrelated, allowing the use of a multinomial prior distribution. The lack of order to the responses calls for an all-or-none loss function that prescribes picking the response with the highest posterior probability. However, when asked to make a continuous response, order is very important. People's responses are often modeled with normal distributions for the prior and likelihood, which leads to a normal posterior. Most loss functions prescribe choosing the mean of the posterior.

Lying between these two well-explored cases is decision making with ordered discrete responses. This kind of decision-making is prevalent outside the laboratory. People are often asked to make discrete responses about variables that are experienced as continuous: On what day this week will you finish the project? Ordered discrete responses particularly arise when it comes to counting numbers of objects. Here the ground truth, such as the number of horses in a paddock, is irreducibly discrete and a discrete response is necessary.

We investigated how people make decisions with a small number of ordered discrete responses, using a numerosity task in which participants are asked to count the number of dots that were shown in a very briefly presented display. We characterized the kinds of prior distributions, likelihood distributions, and loss functions that people used in this task. People's choices were not well described by either common discrete or common continuous models of normative decision-making. The likelihoods and loss functions reflected the ordered nature of the responses, but people learned prior distributions that reflected the discrete nature of the task. Hence the best explanation of people's decision making with ordered discrete responses involved aspects of both approaches.

**Keywords:** Bayesian decision making, ordered discrete responses, numerosity

Francis Crick Room

## AMYGDALA ACTIVITY CORRELATES WITH THE TRANSITION FROM SALIENCY-BASED TO REWARD-BASED CHOICES

**RONG GUO** Technische Universität, Berlin

**KLAUS OBERMAYER** University Medical Center, Hamburg-Eppendorf

**JAN GLÄSCHER** Bernstein Center for Computational Neuroscience, Berlin

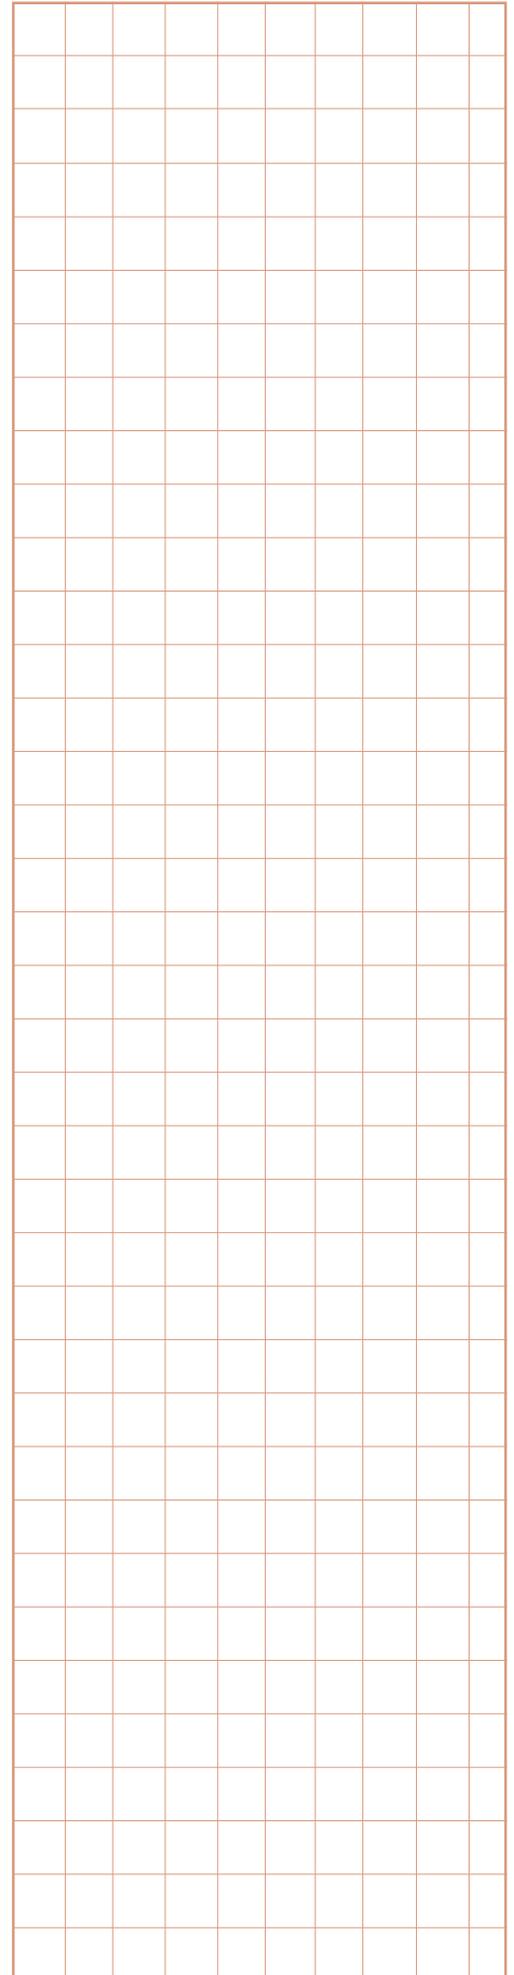
*Introduction:* Ten years ago a finding that the ventral striatum (vStr) responds to salient, but non-rewarding stimuli (Zink et al. (2003)) challenged the hitherto accepted theory of the vStr as one of the critical hubs in the decision-making network encoding rewards and reward prediction errors (O'Doherty (2004)). Although the co-existence both reward- and saliency-related signals in vStr has ever since been confirmed the precise interaction of the two signals has not been fully resolved. Here, we approach this question computationally using reinforcement learning theory in the context of model-based fMRI in a task that requires subject to make probabilistic choices for salient visual stimuli that may or may not lead to a subsequent reward.

*Methods:* 27 participants were instructed to predict the occurrence of a visual stimulus by choosing between left and right button presses in the scanner. Stimuli appeared probabilistically on the left and right screen side in two conditions: (a) equal stimulus probability of 0.5 or (b) biased stimulus probabilities of 0.7 and 0.3 counterbalanced between runs. Conditional reward occurred probabilistically with 0.2 and 0.8 for correct predictions on the left and on the right side. Critically, in the biased stimulus condition, higher reward probability was assigned to the lower stimulus probability, thus creating a situation of dis-information, in which the salient visual stimulus was not predictive of higher reward probabilities. We modeled choice behavior with two independent reinforcement learning algorithms (Rescorla-Wagner models for stimulus and reward occurrences) that were negotiated by a decaying trade-off parameter  $\eta$ . We extracted model-derived signals for  $\eta$ , stimulus and reward prediction error. These signals are used in a model-based fMRI analysis.

*Results:* Model-free behavioral analyses confirmed that the (mis-leading) side with higher stimulus probability in the biased condition was chosen more frequently than in the equal condition. This was also confirmed through model-based analysis, in which the trade-off parameter  $\eta$  decayed more quickly in the equal condition suggesting a faster transition to purely reward-based choices than in the stimulus condition. At the neural level we found a co-existence of stimulus and reward prediction errors in the ventral striatum suggesting that this region responds to surprising perceptual events as well as unexpected reward delivery or omission. Furthermore, the amygdala correlated with the trade-off parameter  $\eta$ .

*Conclusions:* Our findings suggest that vStr is responsive to expectancy violations both at the pure perceptual and at the reward level. This is consistent with previous studies that posit a dual role for vStr in reward learning and saliency detection. Furthermore, the amygdala appears to be negotiating between an initial emphasis on choosing the salient stimulus and pure reward-based choices later. This may reflect its prominent role in Pavlovian stimulus-outcome association, which in our study have to be overcome, in order to maximize the payoff. In summary, our studies highlights the roles of key parts of the decision-making network in learning stimulus- and reward-based choices and trading off both associations against each other.

**Keywords:** Reinforcement learning, model-based fMRI, reward

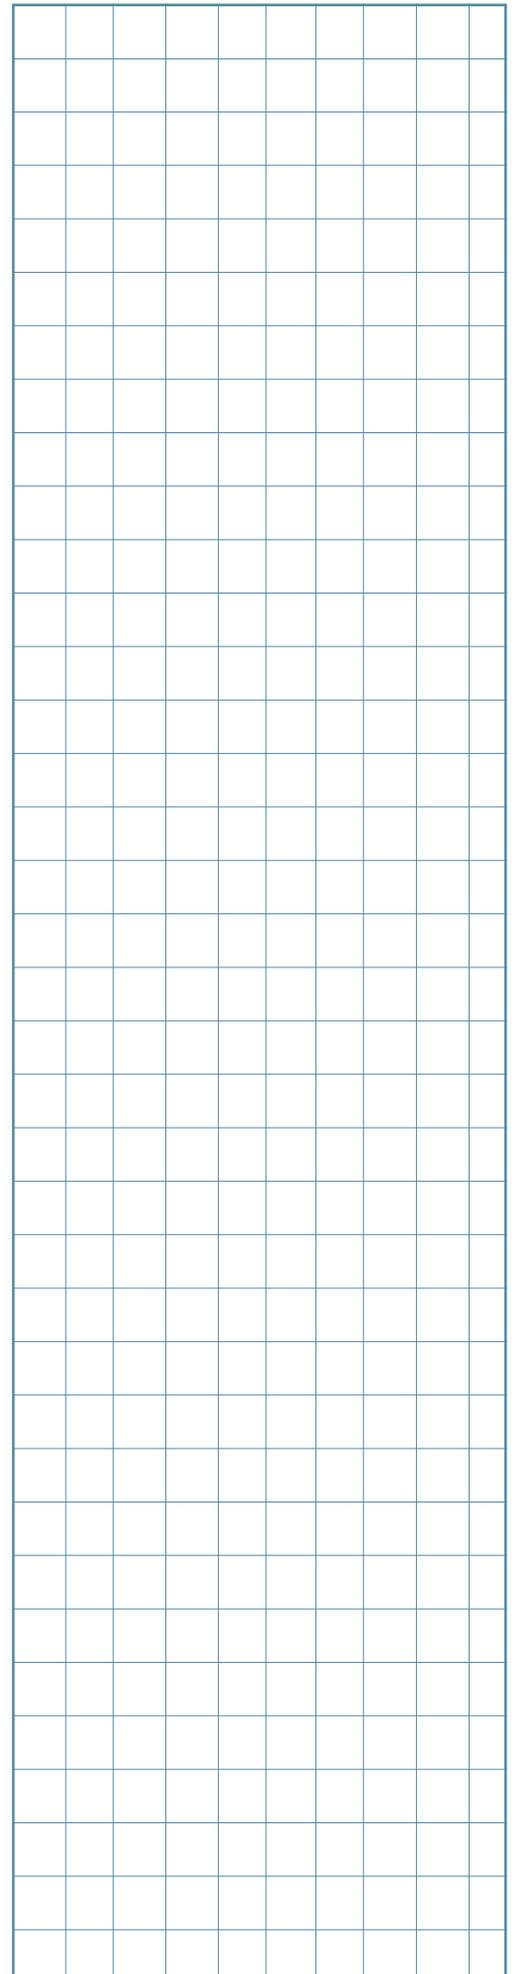


## RATIONAL IRRATIONALITY: MODELING CLIMATE CHANGE BELIEF POLARIZATION USING BAYESIAN NETWORKS

JOHN COOK *University of Queensland; University of Western Australia*  
STEPHAN LEWANDOWSKY *University of Bristol; University of Western Australia*

Belief polarization is said to occur when two people respond to the same evidence by updating their beliefs in opposite directions. This response is considered irrational because it appears to violate normatively optimal belief updating, as for example dictated by Bayes' theorem. In light of much evidence that people are capable of normatively-optimal decision making, belief polarization seemingly presents a puzzling exception. We show using Bayesian Networks that belief polarization can occur even if people are making decisions rationally. We develop a computational cognitive model that uses Bayes Nets to simulate Bayesian belief polarization. The Bayes Net model is fit to experimental data involving representative samples of Australian and US participants. Among Australians, intervention text about the scientific consensus on climate change partially neutralised the influence of ideology, with conservatives showing a greater increase in acceptance of human-caused global warming relative to liberals. In contrast, consensus information had a polarizing effect among the U.S. participants, with a reduction in acceptance of human-caused global warming among conservatives. Fitting Bayes Net models to the survey data indicates that a heightened expectation of fraudulent behaviour by climate scientists drives the backfire effect among U.S. conservatives.

**Keywords:** Belief polarization, Bayesian Networks, climate change, scientific consensus



## THE ADAPTIVE SHIFTS OF DECISION STRATEGIES

**GILLES DUTILH** *University of Basel*

**BENJAMIN SCHEIBEHENNE** *University of Basel*

**HAN L. J. VAN DER MAAS** *University of Amsterdam*

**JÖRG RIESKAMP** *University of Basel*

In many domains of cognitive psychology, it is proposed that people are equipped with a repertoire of strategies to solve the problems they face. For instance, it is proposed that when people have to make a probabilistic inference based on the information of multiple cues, they sometimes rely on simple heuristics (like the Take The Best rule) and sometimes apply strategies that require more elaborate processing. Indeed, many studies suggest that people's inferences can be described by different strategies in different situations. However, critics of this view suggest that people do not apply different strategies in each situation but instead adjust one single strategy to the characteristics of each situation. In this study we examine the strategies that individuals use when available resources change. Specifically, we continuously changed the relative importance of speed and accuracy in a multiple-cue inference task. By doing so, participants were incentivized to gradually shift from very fast guessing behavior to slow but highly accurate behavior (and back and forth). Thorough investigation of participants' behavior at the transitions between fast and accurate behavior offers insight into whether people select different strategies or rather continuously adjust one single strategy to meet varying task requirements. Results show that participants are able to adjust behavior to some extent, but also display qualitative shifts in behavior.

**Keywords:** Multi-Attribute Decision Making, Strategy Selection, Speed-Accuracy Trade-Off.

## NEURONAL CORRELATES OF VALUE BASED DECISION MAKING DIFFER WITH INFORMATION GATHERING STRATEGIES

**NISHANTHA MALALASEKERA** *UCL Institute of Neurology, London*

**LAURENCE T HUNT** *UCL Institute of Neurology, London*

**BRUNO MIRANDA** *UCL Institute of Neurology, London*

**TIMOTHY E BEHRENS** *UCL Institute of Neurology, London*

**STEVE W KENNERLEY** *UCL Institute of Neurology, London*

Optimal decision-making depends on gathering sufficient information to determine the best outcome. However, multi-attribute decisions present the additional challenge of deciding what information to use to compare the alternatives. Patients with damage to the ventromedial prefrontal cortex (vmPFC) not only exhibit deficits in value-based decision-making, they also use different information to make choices; while healthy humans tend to analyse choices using a within-attribute comparison strategy (e.g., comparing the price of each house), patients with vmPFC damage tend to use a within-option strategy (e.g., determine the size, cost, etc. of individual houses). These findings suggest a critical function of vmPFC may be to regulate the process by which choices are compared. Other animal and human lesion studies have described a double dissociation between anterior cingulate cortex (ACC) and orbitofrontal cortex (OFC) damage where ACC lesions cause specific deficits in action based decision-making whereas OFC lesions profoundly affect stimulus based decisions. These results suggest that different parts of prefrontal cortex may represent relevant decision variables in different frames of reference.

To study multi-attribute value based decision-making, monkeys were first taught a set of picture value associations for each of two attributes (reward probability, reward size; 5 pictures per attribute). Subjects were then presented with two pictures (one from each attribute) on each side of a fixation point (4 pictures in total) and made choices between left and right options. Importantly, subjects were free to saccade to the different pictures to gather information about attributes/options. Subjects were also free to indicate their choice (by moving a joystick to the left/right) without necessarily viewing all the picture information. Eye movements therefore provided a proxy for the information gathering strategies influencing decision-making. Each cue consisted of two important properties other than its value; firstly its position on the screen (left or right) which was tied to the action required to obtain the option. Secondly, the cue can denote either a probability or magnitude attribute.

Single neurons were recorded from ACC, OFC, dorsolateral PFC (DLPFC) and vmPFC while subjects performed the task. When subjects were presented with the first picture cue in the trial, many of neurons throughout all four brain areas encoded its value. This encoding was substantially stronger and earlier in ACC and OFC than vmPFC and DLPFC. However, a significant subpopulation of ACC and DLPFC neurons differentially encoded the value of the cue when presented on the left compared to the right suggesting an action value frame of reference for these neurons. In contrast, a significant subpopulation of OFC neurons differentially encoded the value of the cue depending on whether it was of probability or magnitude type. Analysis of neuronal responses to subsequent cues have found that OFC and ACC maintain this dissociation of value reference frames throughout the trial but value signals evolve from encoding the value of what is currently presented to encoding the value of what will be eventually chosen.

**Keywords:** Decision Making, Information Gathering, Prefrontal Cortex, Anterior Cingulate Cortex, Orbitofrontal Cortex, Ventromedial Prefrontal Cortex

# DIVERSIFICATION, RISK AND CLIMATE CHANGE: TOWARDS A PRINCIPLED GUIDE TO PORTFOLIO OPTIMISATION IN CLIMATE INVESTMENT

**JOANNA M. BURCH-BROWN** *University of Bristol*

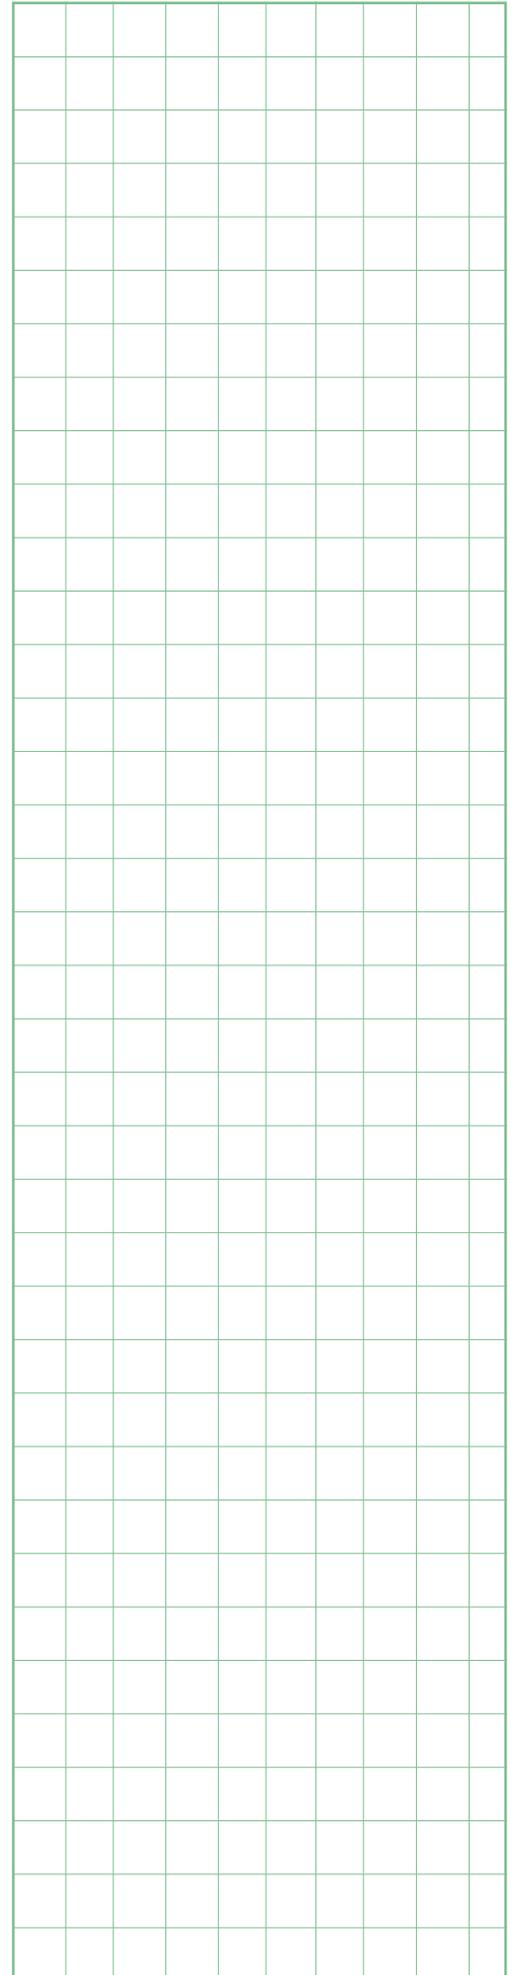
Climate policies are likely to have important side-effects that policy-makers and investors will fail to anticipate in advance. For example, the World Bank estimates that rising food prices in 2010 drove an additional 44 million people below the poverty line in Africa, and research suggests that increased demand for biofuels was an important factor contributing to the rise in food poverty. Demand for biofuels can also lead to destruction of important habitats. In Indonesia, over nine million acres of rainforest have been cleared in recent decades to support palm oil plantations. Most of this production has historically been used for food, but palm oil is also one of the main biofuels, and projections suggest that increased demand for biofuel in the next decade is likely to lead to further clearing of large areas of rainforest in South East Asia and elsewhere. Although biofuels are often advocated as an important technology for emissions reduction, they have become more controversial in recent years as a result of these adverse side-effects.

The possibility that climate investment may have adverse side effects is important for two reasons. One is that adverse side-effects risk undermining the goals that justify the policies in the first place (eradication of poverty, social and economic development, protection of biodiversity and habitat, etc). The second is that perceptions of adverse side-effects can lead to withdrawal of political will for climate mitigation. This in turn increases 'policy risk' - instability in policy support for green investment. Potential investors typically identify policy risk as one of the major blocks to current investment in green technologies.

What strategies can policy-makers and investors adopt to mitigate the risk of adverse side-effects, given that they have good reason to believe that they must make their decisions without full understanding of the systematic effects of their actions?

In this paper, I argue that we can substantially reduce risk of adverse consequences from climate investment by drawing on key principles of risk management from the financial theory of portfolio analysis. I begin by explaining three central principles of portfolio theory. I then show how we might apply these principles to manage risk in relation to climate policy and investment. I argue in particular that we should systematically use diversification to reduce risk in our responses to climate change. I explain some of the challenges for figuring out whether we have successfully diversified risk in this context; and suggest some simple heuristics that we might apply to make the theory useful in practice for climate policy makers and investors.

**Keywords:** Climate change, risk, uncertainty, portfolio theory, risk management, climate ethics

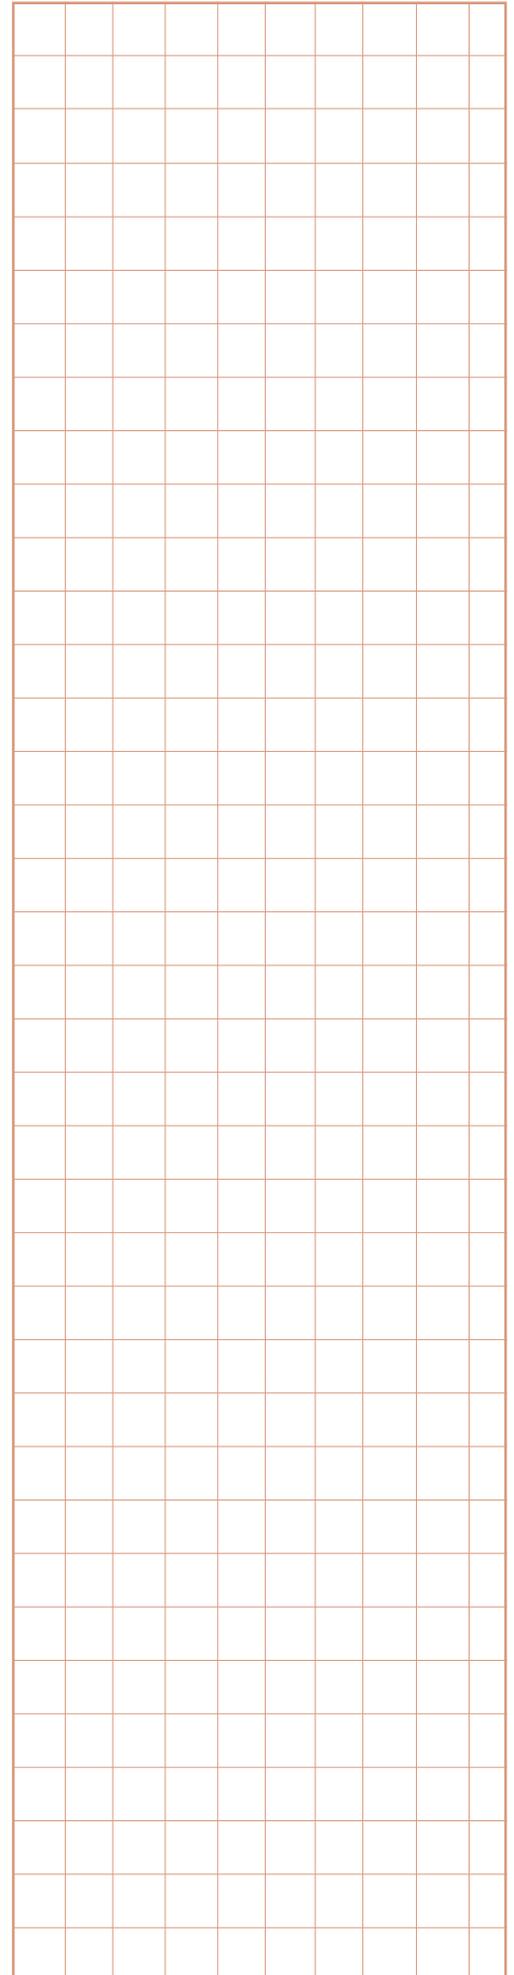


## ETHOLOGICAL DECISION MAKING WITH NON-STATIONARY INPUTS USING MSPRT-BASED MECHANISMS

**LUANA NUNES** *University of Sheffield*  
**KEVIN GURNEY** *University of Sheffield*

The Multihypothesis Sequential Probability Ratio Test (MSPRT) is an asymptotically optimal, sequential test for decision making between several alternatives. It works by integrating 'evidence' on each of its 'channels' and makes a decision when the first channel output crosses a threshold. The overwhelming majority of its applications focus on making single, discrete decisions with stochastically stationary inputs supplying evidence, and in which the evidence is initialized to zero prior to the decision process. Here, the decision making process is terminated upon crossing the threshold, which therefore constitutes an absorbing boundary. In contrast, ethologically plausible decision making in animals (and situated agents) should proceed in a continuous fashion, and will occur in environments where the information supporting each alternative changes dynamically throughout the integration process; that is the inputs will be non-stationary. Models of decision making and action selection in animals are often based on the hypothesis that the basal ganglia - a set of sub-cortical brain nuclei - act as the central mechanism mediating decisions. These models work with continuously varying inputs, and select actions (make decisions) via a threshold crossing which does not terminate the decision making process. Intriguingly, (Bogacz & Gurney, 2007) showed there was close correspondence between the functional anatomy of the basal ganglia and the algebraic manipulations of the MSPRT, but considered only discrete decision trials in reporting their results. However, viewed as an algorithmic abstraction of an anatomical architecture, we ask the question: can the MSPRT mechanism be recruited for decision making with non-stationary inputs, and without use of an ad-hoc 'reset' mechanism? We present a preliminary series of simulation experiments exploring this possibility, and comparing the results with the full basal ganglia model (Gurney, Prescott, & Redgrave, 2001) working in continuous selection mode. In order to do this, we introduced what we call a transparent boundary, which records choices made but does not stop the integration. Further, as with the basal ganglia model, threshold crossing from above comprises a selection, whereas crossing from below indicates its de-selection. Using inputs in which two out of three channels had very similar means, the MSPRT mechanism showed behavior consistent with that obtained from the basal ganglia model (Gurney, Prescott, & Redgrave, 2001); thus, closely matched competitors can induce a 'selection limiting' effect on each other and take more time to get selected than a channel with the same input but no close competitors. Continuous decision making of the kind we envisage here introduces a problem: continuous evidence accumulation, results in unfeasibly large inputs, but also in the inability to 'forget' past evidence and focus on more recent input. We therefore, also introduced an integration window so that only the N most recent samples of information are used in the decision process. With this device, time-to-select after introduction of a new input signal contrast, was considerably decreased because the "baggage" of the previous evidence history could be discarded more quickly. Taken together, our results indicate how the MSPRT mechanism may be recruited for continuous, ethologically plausible decision making with non-stationary inputs.

**Keywords:** decision making, multihypothesis, MSPRT, sequential test, continuous algorithm, non-stationary inputs



## INCORPORATING CONFLICTING DESCRIPTIONS INTO DECISIONS FROM EXPERIENCE

**LEONARDO WEISS-COHEN** *University College London*  
**EMMANOUIL KONSTANTINIDIS** *University College London*  
**MAARTEN SPEEKENBRINK** *University College London*  
**NIGEL HARVEY** *University College London*

Most of the research on the "description-experience-gap" so far has focused on presenting different participants with either description-only or experience-only tasks separately. However, decisions in everyday life are also commonly made on a combination of descriptive and experiential information, and these two sources of information often contradict each other. Our aim was to shed additional light on how individuals combine information acquired by description and by experience when making decisions. Our experiments have shown that individuals exposed to both description and experience can be influenced by the two sources of information, in particular in situations in which the description was in conflict with the experience. We looked at how individuals might integrate experience with prior beliefs about the outcomes of their choices, with different weights given to each source of information, depending on their relevance. Cognitive models that included the descriptive information fitted the human data more accurately than models that did not include descriptions. We discuss the wider implications for understanding how these two commonly available sources of information can be combined for daily decision-making, and the impact on areas such as warning labels.

**Keywords:** description, decisions from experience, description-experience gap, reinforcement learning, repeated decisions

## PYRRHIC VICTORIES IN PUBLIC-GOODS GAMES: WHEN RELATIVE COMPARISONS MATTER MORE THAN ABSOLUTE PAYOFFS

SEBASTIAN HAFENBRÄDL *University of Lausanne*

**JAN K. WOIKE** *Max Planck Institute for Human Development*

Participants in repeated public-goods games do not fully contribute to the public good, surprisingly even when the typical dilemma is absent and contributions maximize both individual and group outcome. Here, we demonstrate that destructive competition based on relative comparisons is the driving factor for suboptimal (and spiteful) decisions in games where full cooperation is optimal for the individual player. Adding social information to the individual outcome feedback can destroy future self-beneficial cooperation when focused on relative gains and rankings and help to maintain cooperation levels in true dilemmas when focused on absolute group performance and efficiency. Ranking information triggered relative comparisons that led to an erosion of cooperation, whereas spiteful behavior is severely limited in the absence of relative information. Explanations centering on mere confusion, unconditional social preferences or a universal cooperative motivation cannot account for these results. Our findings demonstrate that feedback-structures can influence cooperation-levels and trigger social comparisons that lead to value-destroying competition. Designers of information environments -- inside and outside the lab -- need to be aware of these consequences. Furthermore, interpretations of behavior shown in economic games need to reflect the demonstrated dependency of cooperative actions on highly variable features of social environments.

**Keywords:** social comparisons, cooperation, economic games, confusion

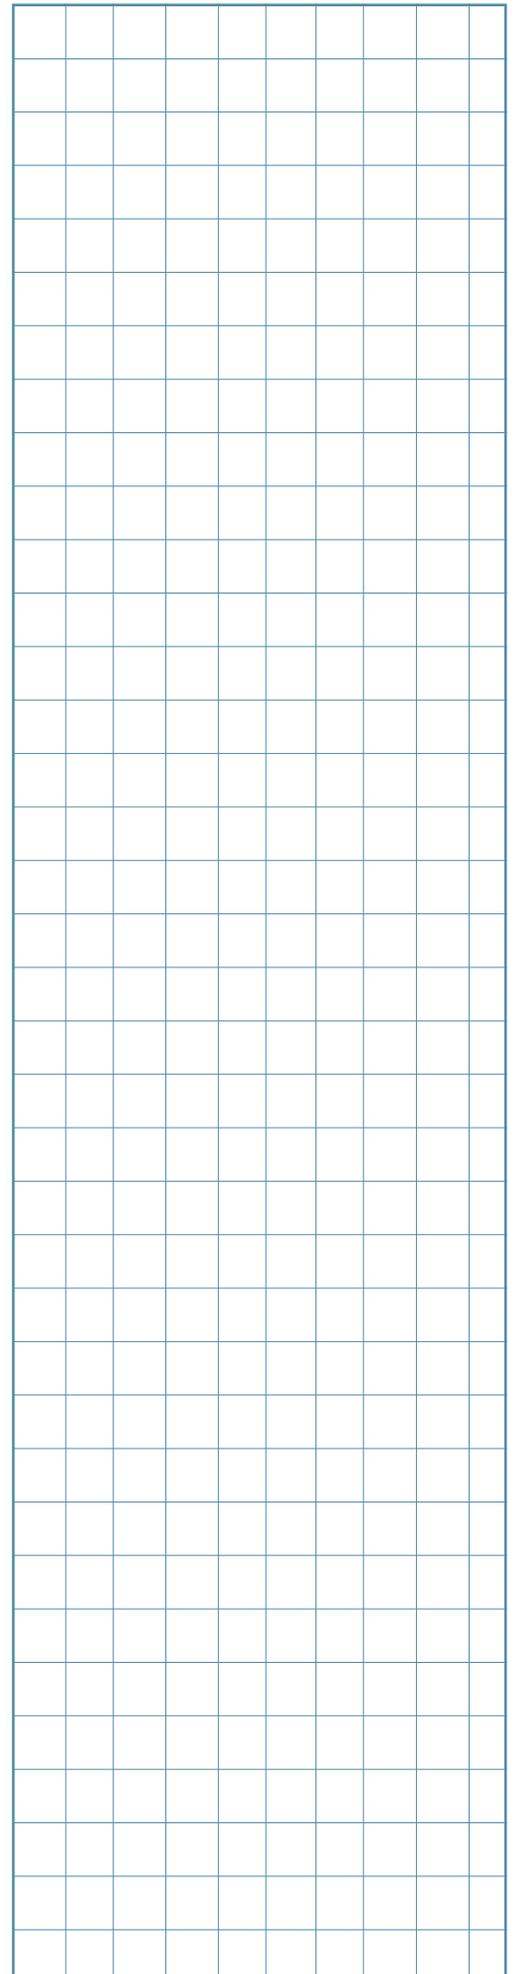
## CHANGING DECISION CRITERIA IN RAPID AND SLOW DECISIONS: DO PEOPLE BEHAVE OPTIMALLY?

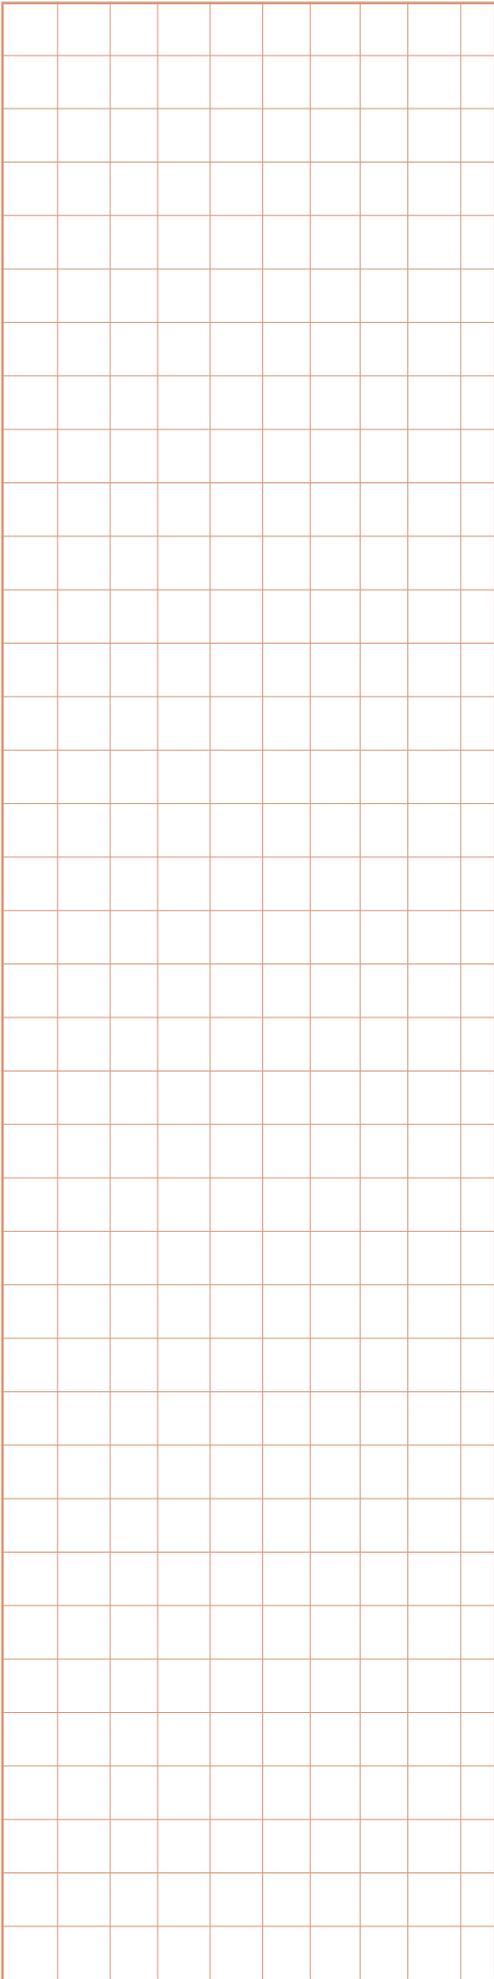
**GAURAV MALHOTRA** *University of Bristol*  
**DAVID LESLIE** *University of Bristol*  
**CASIMIR LUDWIG** *University of Bristol*  
**RAFAL BOGACZ** *University of Bristol; University of Oxford*

Research on decision making has indicated that the trade-off between the speed and accuracy with which people make decisions is determined by a decision criterion; a threshold on the evidence that needs to be gathered before the decision-maker commits to a choice. A fundamental question that has attracted a lot of interest is whether humans decrease the decision threshold within the course of a single decision. We investigated this issue from both a normative and an empirical perspective. Using a simple theoretical framework, we determined the minimal conditions under which a decision-maker should vary their decision criterion to maximise reward rate. We found that the shape of the optimal threshold depends crucially on how noise in signals varies across trials. Optimal thresholds stay constant within a trial when noise in sensory input is consistent across trials. In contrast, when noise across trials is variable, time spent within a trial becomes informative and the optimal decision criterion may not only decrease, but also increase or remain constant, depending on the mixture of noise levels across trials.

We compared the performance of this optimal decision-maker to behaviour using a series of experiments. An expanded judgement paradigm that precisely matched the theoretical framework showed that, in agreement with the optimal model, participants had on average constant thresholds in blocks of trials with constant difficulty, and decreasing thresholds in blocks of trials with mixed difficulties. We then tested whether results from this expanded judgement paradigm generalise to faster decisions, like perceptual decisions, with mean RTs below 1000ms. We found that the shape of decision boundaries critically depends on the specific mixture of difficulties - only when participants believe that some trials in a block are extremely difficult, do they decrease their decision boundaries significantly over the course of a decision.

**Keywords:** decreasing bounds, optimal decisions, reward rate, drift-diffusion model





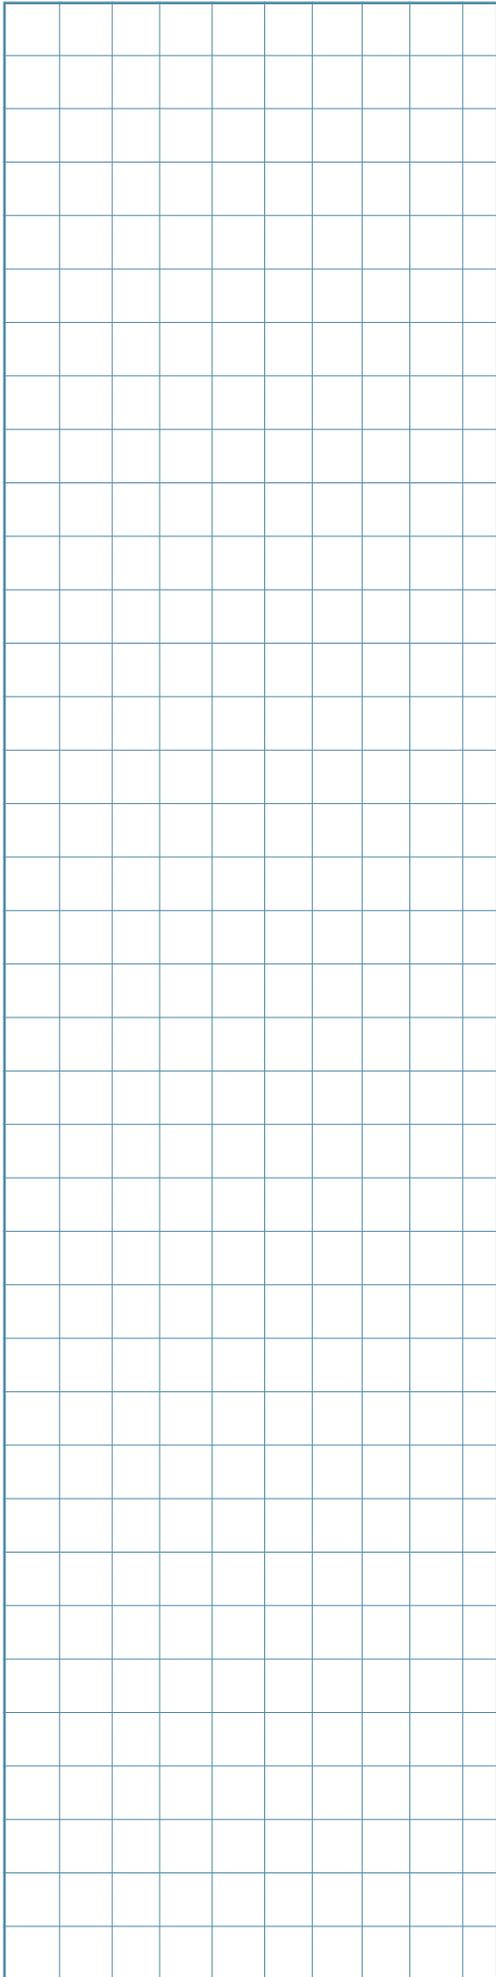
## DESCRIPTION AND EXPERIENCE BASED DECISION MAKING: AN EXPERIMENTAL AND STRUCTURAL ESTIMATION APPROACH TO THE DESCRIPTION-EXPERIENCE GAP

**JACK STECHER** *Carnegie Mellon University*  
**LINDA MOYA** *Carnegie Mellon University*

In experiments on decision making under risk, participants who learn probabilities by sampling behave differently from those to whom the experimenter describes probability distributions. This difference is known as the description-experience gap. We argue the chief reason for the description-experience gap is that participants who learn by sampling become familiar with a decision problem, prior to making a choice. Participants who learn from description do not have this luxury. We run laboratory experiments, in which participants learn by sampling, and use their behavior to estimate the distribution of risk aversion coefficients. A different group of participants receives equal task exposure, but samples from a nonstationary, non-ergodic distribution, after which they are told the distributions in a choice problem. After controlling for task exposure, the gap disappears, and participants' behavior is shockingly consistent with expected utility theory. We re-analyze older data, and find that our results are robust across contexts.

**Keywords:** decisions from experience; description-experience gap; rare events; risk; probability weighting; utility; prospect theory





## THE COLLAPSING CONFIDENCE BOUNDARY MODEL: A UNIFIED THEORY OF DECISION, CONFIDENCE, AND RESPONSE-LATENCIES

**RANI MORAN** *Tel Aviv University*

**ANDREI R. TEODORESCU** *Tel Aviv University*

**MARIUS USHER** *Tel Aviv University*

While recent research has made impressive progress in understanding the mechanism underlying decision-making, decision-confidence – an important regulatory decision process – is still raising considerable challenges. The choice followed by confidence paradigm produces rich data sets consisting of four important measures of cognitive performance: Choice, confidence and their corresponding latencies. These variables combine and interact to form an extensive manifold of patterns that challenges confidence modelers. The resolution of confidence, the positive correlation between choice-accuracy and choice-confidence, is one such pattern, central to our investigation. In two perceptual choice experiments, we manipulated the perceptual availability of the stimulus following the decision and found systematic influences on the resolution of the confidence judgments. This indicates that post-choice evidence exerts a causal effect on confidence and that the post-choice time is pivotal in modeling confidence. We present a novel dual-stage model that provides a unifying account for all four dependent variables, namely the collapsing confidence boundary (CCB) model. According to CCB, decision is determined by a standard diffusion model. Choice, however, is followed by a second evidence-integration stage towards a stochastic collapsing confidence boundary

**Keywords:** Confidence, Diffusion model, Collapsing threshold/boundaries, Time pressure, Resolution of confidence, Choice, Perception.

Francis Crick Room

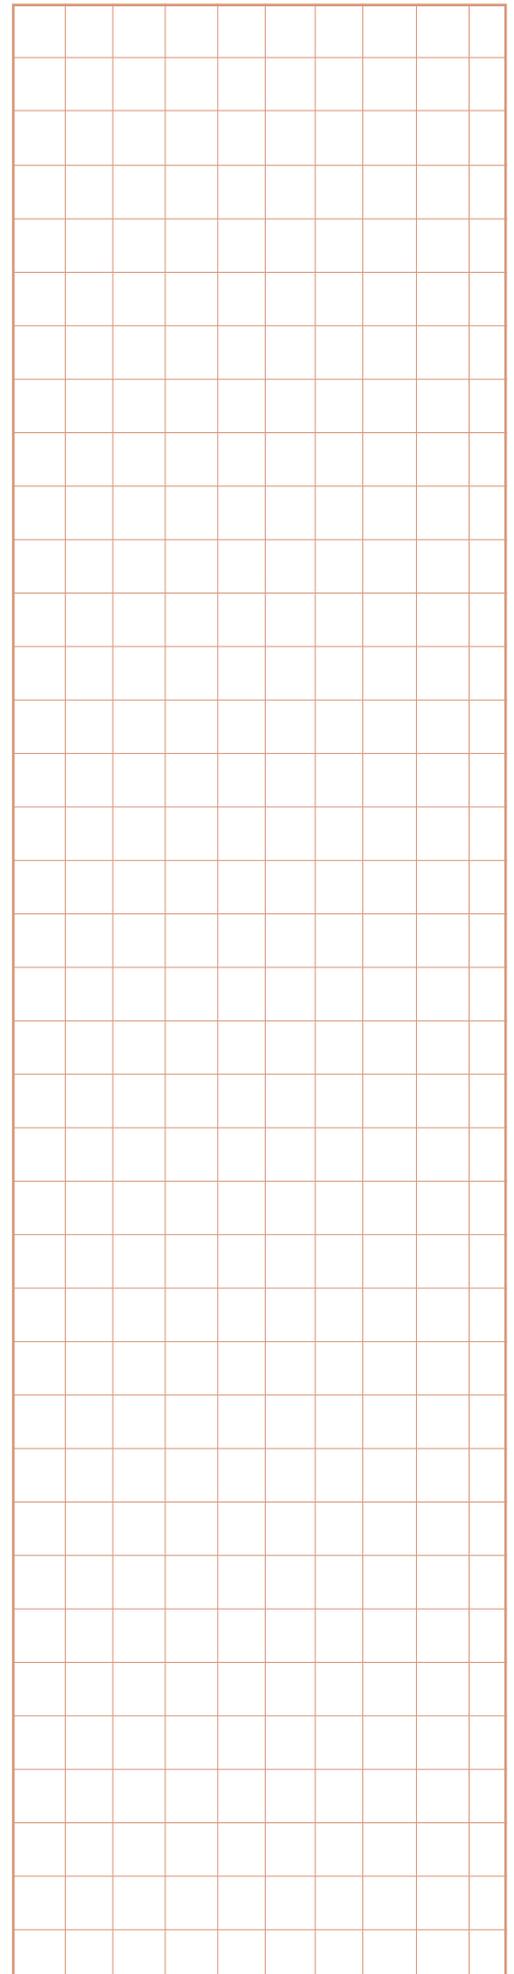
## CHOICE-RICH ENVIRONMENTS INDUCE RISK-TAKING AND REDUCE CHOICE DEFERRAL

**THOMAS HILLS** *University of Warwick*

**TAKAO NOGUCHI** *University of Warwick*

Choices can be risky. Unfortunately, more choices can be riskier. We present research showing that in choices between monetary gambles, choice enriched environments increase risk taking and lead to less choice deferral. Specifically, we compared decisions from experience with decisions from description. In decisions from experience, when people sample a gamble they experience individual outcomes; people must repeatedly sample gambles to learn about the gamble's outcomes and probabilities. In decisions from description, information about the probabilities and outcomes are provided at the same time and repeated sampling is unnecessary. A typical finding is that people underweight rare events in decisions from experience and overweight rare events in decisions from description: a phenomenon called the description-experience gap. To investigate the influence of set size on choice, we increased the number of alternatives (to more than 30 options) in two standard choice paradigms: The sampling-paradigm (decision from experience), where people can sample freely over alternatives before making a final decision, and the decision from description paradigm, where people can deliberate over complete information about outcomes and probabilities at the same time before choosing. In the first experiment, gambles were either very risky or fairly safe and payoffs were sampled from similar distributions. Increasing set size increased the sample size in both conditions, but led to fewer samples per gamble. In the sampling paradigm this increased the likelihood of experiencing rare, risky events. When choice was over gains, this led to increased choice for riskier options; people tended to choose options associated with large (rare) gains. Over losses, set size had no influence on risky choice; encountering rare events led participants to favor other risky gambles where rare events were not encountered, as they do for small set sizes. Importantly, the observation that decisions from experience leads to underweighting of rare events was reversed for large set sizes in the gains domain; instead, people chose as if they overweighted rare events. In a second experiment, we gave participants £1 for showing up and then offered them the opportunity to buy gambles with positive payoffs. We added a button to allow participants to, after free sampling of alternatives, either keep their money (defer choice) or to choose one of the alternatives. In both decisions from description and decisions from experience choice deferral went down as set sizes increased. Moreover, in both experiments, choices were not chosen at random but were predicted by what the participants saw when they sampled alternatives. Overall, our results indicate that context effects associated with increasing set sizes are not always readily apparent from two-alternative decisions. However, the increase in risky choice associated with experienced gains is predicted from standard choice models. These effects also extend to questions of choice deferral and information overload, for which we found no evidence of either.

**Keywords:** Risk, decisions from experience, too much choice, information overload, decisions from description



## REDUCED UNCERTAINTY IMPROVES CHILDREN'S COORDINATION, BUT NOT COOPERATION

**LIAM POLLOCK** *Queen Mary University of London*

**DR KEITH JENSEN** *University of Manchester*

**DR MAGDA OSMAN** *Queen Mary University of London*

In game theory the battle of the sexes (BOS) is a well-known set-up designed so that two agents (husband, wife) need to coordinate their actions (go to the opera or the football) in order to maximise their payoffs. The game has two Nash equilibria (both act the same), each of which favours only one agent, or a mixed strategy (both act differently).

Typically, in experimental applications of the BOS, without signalling their actions, adults follow inefficient strategies (poor coordination). Is this true of children, particularly given how differently they conceptualise fairness? To address this children aged between 5-13 were presented variations of the BOS game in which payoffs and signalling were manipulated, to see whether reduced risk of opponent's outcome (condition 2) or reduced uncertainty of opponent's outcome (condition 3) would foster greater mutual cooperation as compared to the standard (condition 1).

304 children aged from 5-13 years (134 males, 170 females; mean age = 7.7 years, SD = 2.01) took part in a BOS game with 30 iterations. Children were randomly allocated to one of three conditions: 1) standardBOS, 2) BOSincreased risk, and 3) BOSincreased risk + reduced uncertainty. Payoff matrices for each of the 3 conditions are shown in Table 1.

In all conditions, children were tested in dyads face-to-face across a table. Each player was given a stack of identical playing cards half of which had lions on the face, the rest with foxes. Players were randomly allocated an animal that was "theirs" to indicate their preference. The mutual playing of "their" animal would maximise both the relative and absolute payoffs for the animal's owner in the manner shown in table 1. It was clearly emphasised to all players during the habituation period that it was their absolute and not their relative payoff that they needed to maximise.

Our findings revealed that children showed the highest rates of coordination when uncertainty around the opponent's actions was reduced, whereas for reduced risk and the standard conditions coordination rates were comparably lower. Internal checks based on performance in condition 1 and 2 suggest that our participants understood the experimental set-up. In addition, we found that cooperation levels between the three conditions were identical, meaning that our experimental manipulations impacted children's decisions regarding cooperation, which occurred at a relatively low and, from a game theoretical perspective, is a highly inefficient level. Implications for the ontogeny of decision-making related to cooperation and fairness are discussed.

**Keywords:** battle of the sexes, game theory, children, cooperation, coordination

## COMPETITIVE ACCUMULATOR MODELS EXPLAIN WHY THERE IS OFTEN NO CORRELATION BETWEEN ERROR AND RT MEASURES OF INHIBITORY CONTROL

**PETROC SUMNER** *Cardiff University*

**CRAIG HEDGE** *Cardiff University*

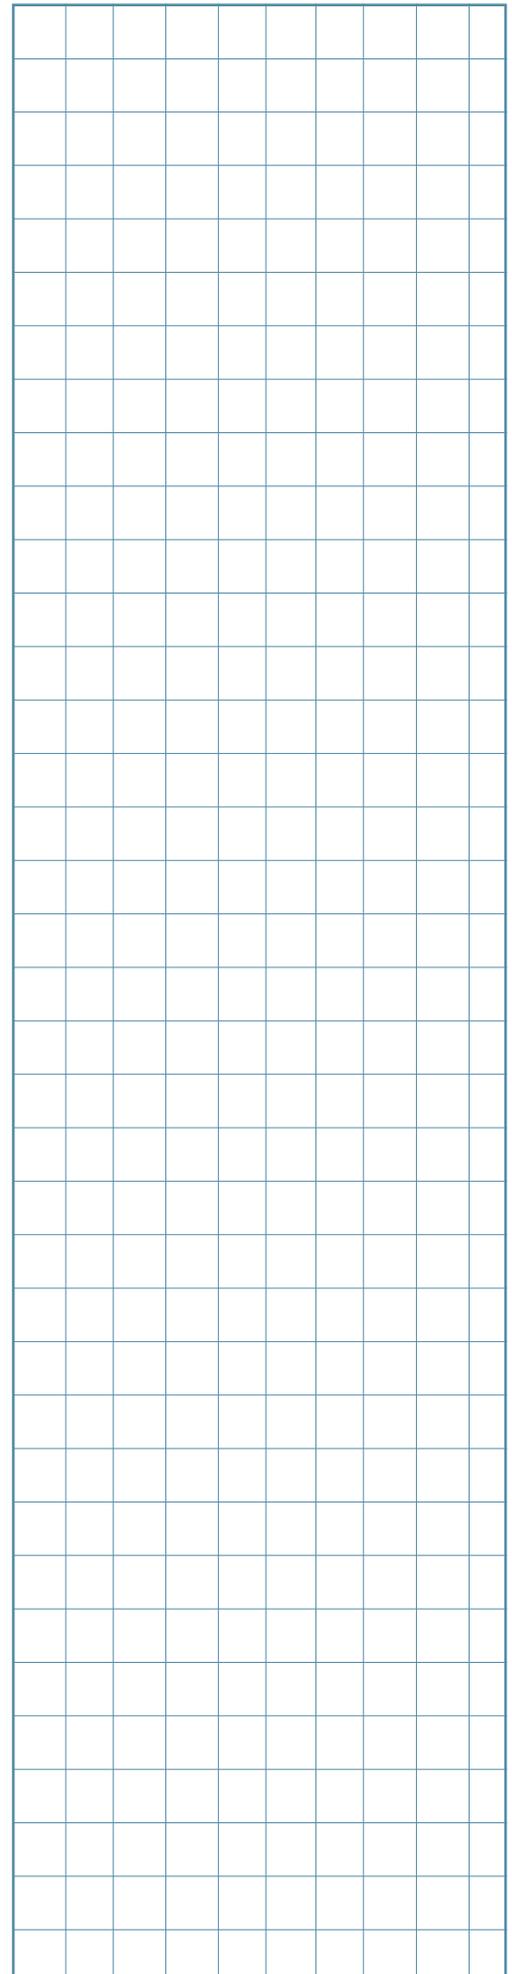
**ALINE BOMPAS** *INSERM; CNRS; Lyon Neuroscience Research Center*

Individual or group differences in rapid action decisions measured by tasks such as anti-saccades and Stroop are generally assumed to reflect ability to inhibit unwanted responses. However, there is little theoretical understanding of what such a conclusion actually means and how to generalize from it to real-world behaviour [13]. Progress is difficult because different response control tasks have typically shown surprisingly little correlation (or little repeatability of correlation) with each other or with self-reported impulsivity. There is even little correlation between interchangeably used measures from the same task – error rate (response control failure) and reaction time cost – which are universally assumed to tap the same process. This leads to confusion and apparent conflict in the literature.

Here, we replicate the lack of correlation between error and RT costs across several response conflict tasks. We then show that this situation is straightforwardly accounted for by competitive accumulator models in which impulsivity/reduced top-down control can be captured in (at least) two distinct ways: either as lower caution or lower selection ability (e.g., less good inhibition). Both of these increase errors, but they have opposite effects on RT costs. This means that if people vary in both caution and selection ability, there is expected to be no correlation between RT cost and errors.

The reason caution and selection oppositely affect RT cost is that increasing caution trades fewer errors for increased RT costs (as well as overall longer RT – the speed-accuracy trade-off), while increasing selection bias reduces errors by limiting how much the ‘wrong choice’ gets activated, and this reduces RT cost too. Therefore, although caution and inhibition/selection are both subsumed within the general concept of impulsive action, they are predicted to have opposite effects on some measures of response control. In this way the competitive accumulation framework makes an apparently confusing fact entirely understandable.

**Keywords:** Impulsivity; cognitive control; inhibition; speed accuracy trade-off.



## THE INFLUENCE OF SAMPLE SIZE ON DECISIONS FROM EXPERIENCE

**JANINE CHRISTIN HOFFART** *University of Basel*

**GILLES DUTILH** *University of Basel*

**JÖRG RIESKAMP** *University of Basel*

In recent years, there has been a growing interest in the so-called description-experience gap (D-E gap) in risky decision making. The D-E gap refers to the common finding that people choose differently in an experience-based choice paradigm and a descriptive choice paradigm. In the experience-paradigm people are allowed to draw samples to learn about a gamble's outcome distribution whereas in the description paradigm they are presented with plain numerical descriptions of the gamble's outcome distribution. The D-E gap holds that in the description paradigm, people behave as if they overweighted relatively small probabilities, whereas in the sampling paradigm, people behave as if they underweighted small probabilities.

Many studies on the D-E gap show that people draw relatively few samples before making a decision. Although this undersampling is often raised as an explanation for the D-E gap, little is known about how precisely sample size influences people's choice behavior. A reason for this lack of understanding follows from the fact that successful models for decisions from description, including prospect theory, make no predictions about sample size. Models that assume Bayesian updating, and to some extent memory based learning models, such as Gonzalez' instance based learning model (2003), do make predictions about the effect of sample size. However, until now these effects have not been studied experimentally yet.

Here, we apply a within-subjects design to experimentally test the effect of sample size on choice behavior. For this purpose, we systematically manipulated the number of samples (5, 10, 20 or 40) that participants drew before making a choice. For each pair of gambles, the outcomes a participant saw matched the according gamble's underlying outcome distribution. In addition, the exact same gambles were presented in a descriptive format to the participants.

On an aggregate level, our data reveals that the D-E gap is more pronounced for trials in which participants drew relatively small numbers of samples (i.e., 5 or 10 samples). So, it appears that overall, sample size does influence people's choices, yet it is not entirely clear to what extent. One reason, for which we cannot draw strong conclusions about this influence, is that binary choices contain little information about people's degree of preference. Therefore, we conducted a second experiment, in which we administered a gradual measure of people's valuation of gambles. Such a gradual measure allows us to apply more powerful statistical methods and thus to better understand differences in cognitive processes that might lead to the D-E gap.

**Keywords:** Decisions from experience, sample size, D-E gap

## FAILING TO INVEST BY THE NUMBERS

**PHILIP NEWALL** *University of Stirling*

**BRADLEY LOVE** *University College London*

A large body of empirical work in finance has concluded that mutual fund investors significantly reduce their welfare by selecting mutual funds with high past returns rather than funds with low fees. Future performance is on average reduced one-for-one with increases in fees. The large size of the mutual fund industry, and the dispersion of mutual fund fees, means that this is an area where behavior change interventions might increase investor welfare. Potential interventions must, however, be evaluated for not only their aggregate implications but potential perverse effects amongst sub-populations.

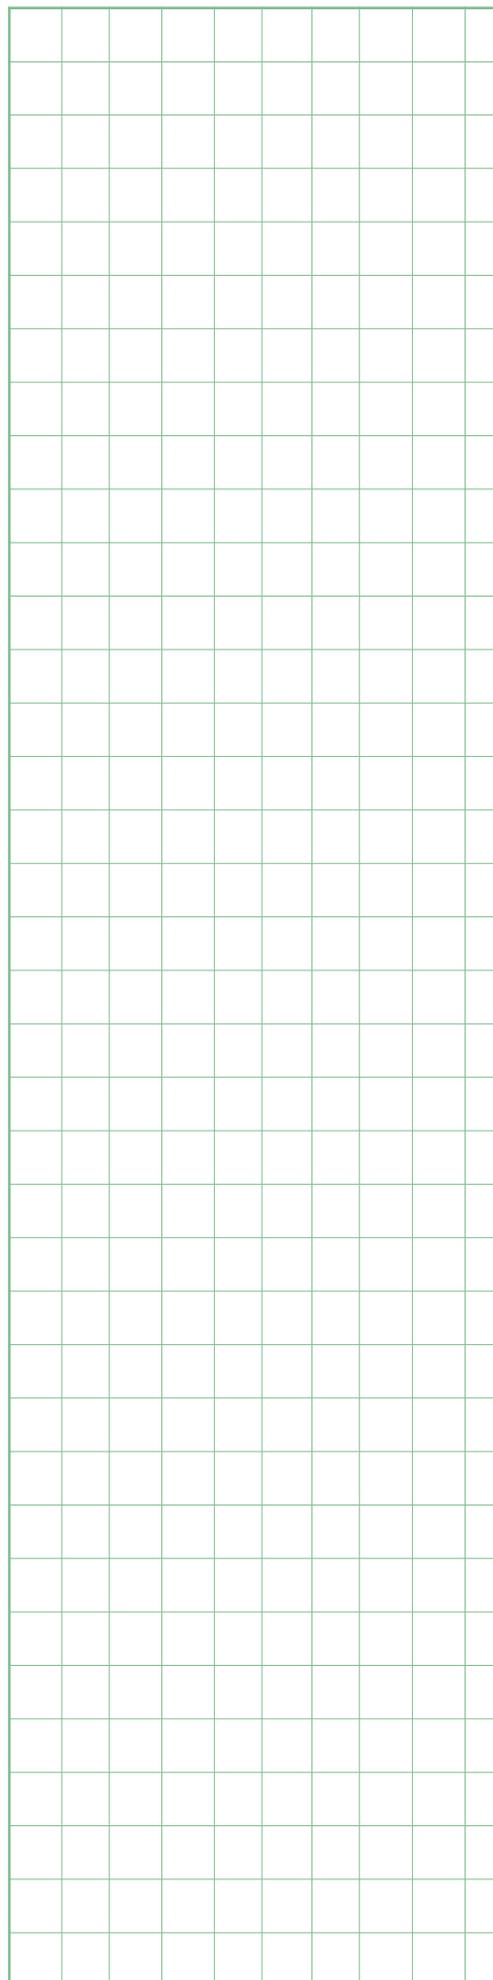
Previous work has attempted to debias mutual fund investors by reframing percentage fees into actual money amounts, where a 1% fee on a \$10,000 portfolio would be reframed as \$100 a year (Choi et al., 2010; Hastings & Tejada-Ashton, 2008). Although this effect has led to a statistically significant increase in fee-sensitivity, economically significant levels of bias remain. Given the low cost of framing manipulations, this might still be a welfare-increasing intervention if investors respond homogeneously.

In a simple 2x2 experiment with fees (actual money, percentage) and portfolio size (\$1,000, or \$1,000,000) manipulated between-participants, we show perverse effects of actual money framing of fees: this manipulation leads to a large decrease in fee-sensitivity when the fee translates into a low amount of \$10-\$15 a year, with the proportion of fee-minimizing participants dropping from 40.6% to 27.4% (compared to a percentage control). Furthermore, our results fail to replicate those of Choi et al. (2010) and Hastings and Tejada-Ashton (2008): representing the same percentage fees as \$10,000 or \$15,000 led to a statistically insignificant improvement over percentage framing (from 37.6% to 41.6%), despite per-cell averages of 250 participants.

Experiment two, which mirrored experiment one except that framing in percentages or actual money was now applied to past returns, shows that this effect can instead induce greater fee-sensitivity. 54.9% of participants who saw past returns of \$10-\$15 a year minimized fees, a significant improvement on 41.5% in the corresponding percentage group.

Our conclusions are twofold. Aggregate positive effects of an intervention can mask perverse responses amongst subpopulations. Interventions that improve behaviour on average may make certain groups systematically worse. Second, we view our results as a simple proof of concept. Previous research has attempted to debias mutual fund investors by increasing the salience of fees. Our results show that reducing the salience of past returns may be a more effective method of increasing the decision weight given to fees. Further work in experimental economics should aim to uncover efficient techniques of reducing the salience of mutual fund past returns, as it shows promise as a method of debiasing individuals in this economically significant market.

**Keywords:** Investing, Behavioural finance



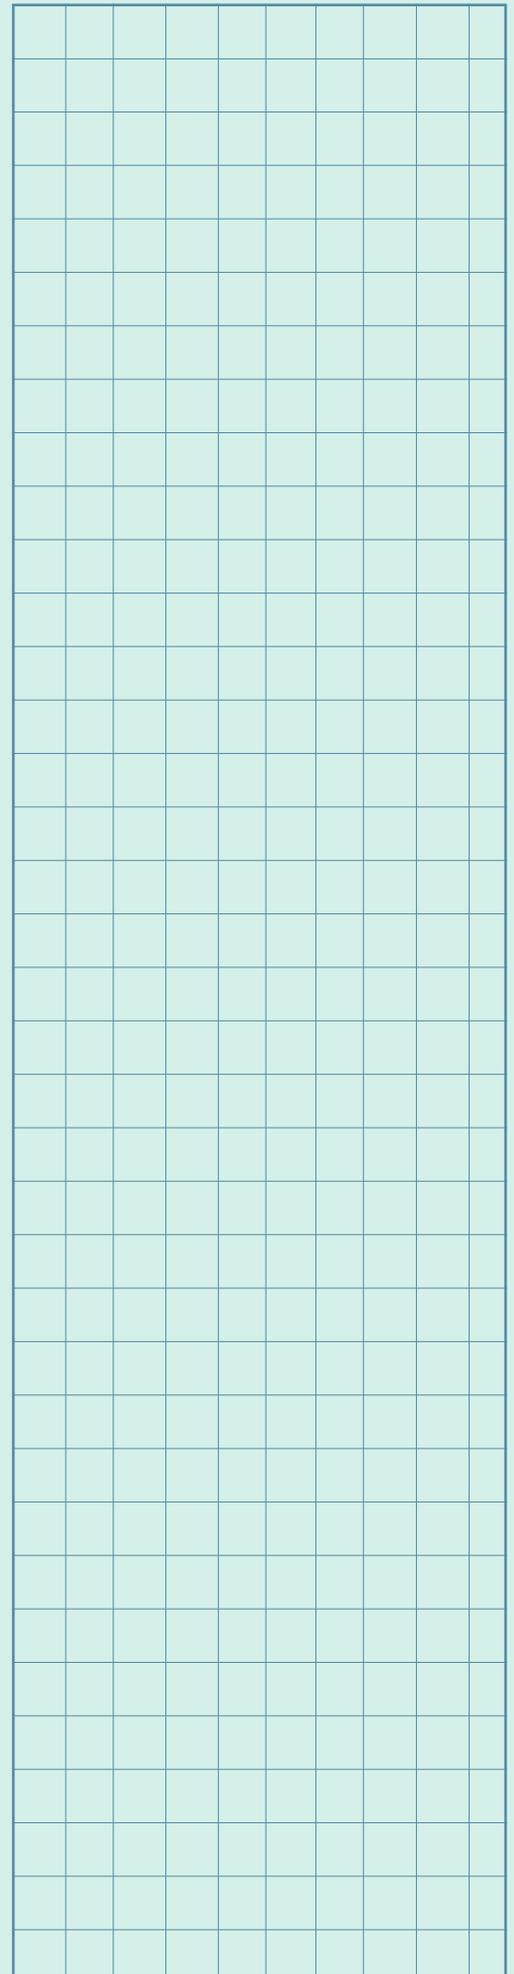


# ERIC-JAN WAGENMAKERS

## REINFORCEMENT LEARNING AND THE IOWA GAMBLING TASK

**ERIC-JAN WAGENMAKERS** *University of Amsterdam*

People can learn through trial and error, a process that can be quantified using models of reinforcement learning. In psychology, these models may be used for at least two purposes. First, researchers may be interested in what model best accounts for people's choice behavior. This goal is usually referred to as model comparison or model selection. Second, researchers may be interested in using a particular model to decompose observed choice behavior into distinct cognitive processes. This goal is usually referred to as parameter estimation. Here I will discuss a research line that has provided a suite of tools for model comparison and parameter estimation in the Iowa gambling task, one of the most popular reinforcement learning tasks to date.



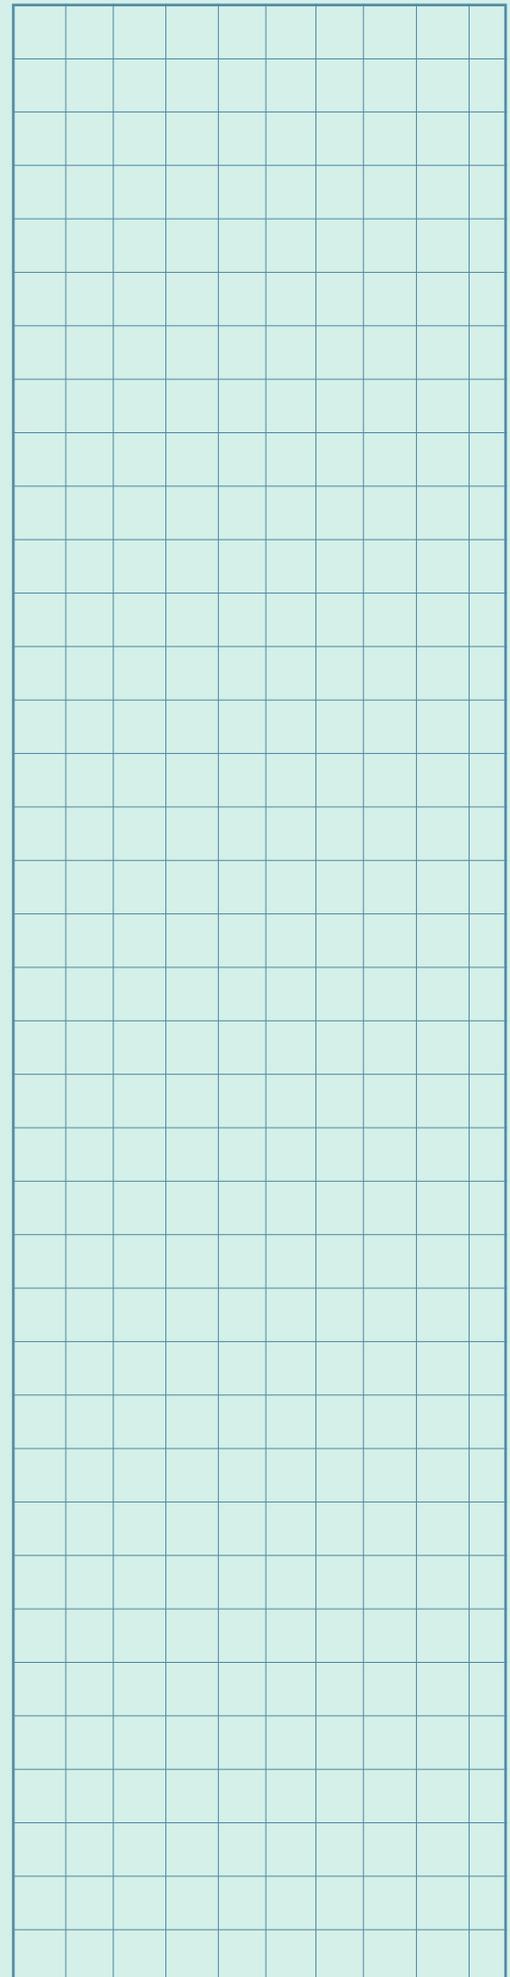


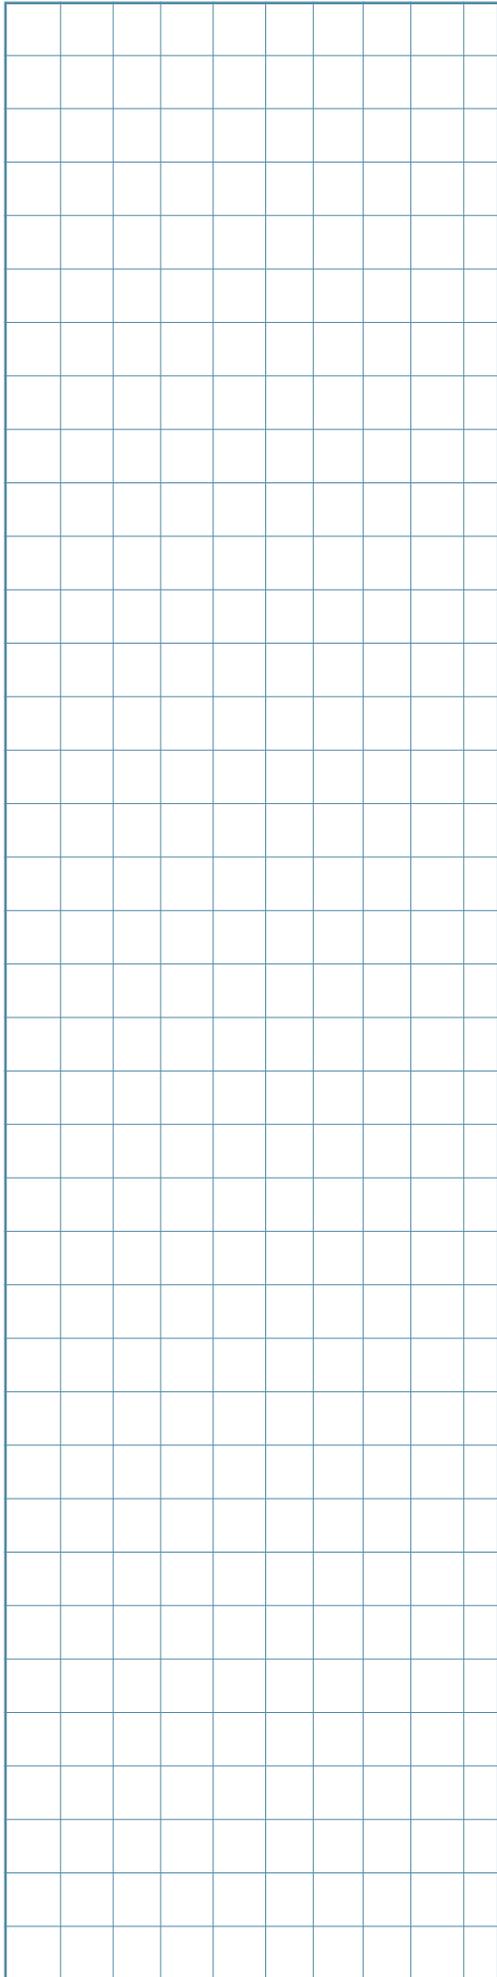
JEFF  
SCHALL

## DECISIONS, ACCUMULATORS AND NEURONS: HOW SECURE A BRIDGE?

**JEFFREY D. SCHALL** *E. Bronson Ingram Prof of Neuroscience, Vanderbilt University*

The conjecture identifying spiking activity of certain neurons with a stochastic accumulator decision process has inspired a vigorous and productive research effort for 20 years. This effort has described decision accumulation processes in multiple brain regions even extending to noninvasive ERP and fMRI measurements. It has been buttressed by models of perceptual categorization, response inhibition and visual search. Lately, though, several new findings have raised questions about the coherence and clarity of the mapping between neurons and decision accumulators. These include the first data showing how neurons identified with decision accumulators accomplish executive control and speed-accuracy tradeoffs and the first model of response time from ensembles of accumulators. The new results indicate that mapping between parameters of accumulator models and measurements of neural activity is not as transparent as originally presumed. The accumulator model framework will no doubt remain an effective means of quantifying performance and instantiating computations in various tasks. However, the construction of a more secure bridge between model and neural levels of description will require more assiduity in (1) accounting for multiple stages of processing each adding time and potential errors, (2) incorporating distinct neural processes from heterogenous neurons in diverse neural structures, (3) articulating the transformations between spikes, ERPs, and BOLD, (4) specifying converging constraints to limit parameters in more complex models and (5) appreciating the logical and rhetorical scope of the mapping – true identity, quantitative analogy, or interesting metaphor.





## DISSOCIATION BETWEEN THE IMPACT OF EVIDENCE ON EYE MOVEMENT TARGET CHOICE AND CONFIDENCE JUDGEMENTS

**EUGENE MCSORLEY** *University of Reading*

**CLARE LYNE** *University of Reading*

**RACHEL MCCLOY** *University of Reading*

It has been suggested that the evidence used to support a decision to move our eyes and the confidence we have in that decision are derived from a common source. Alternatively, confidence may be based on further post-decisional processes. In three experiments we examined this. In Experiment 1, participants chose between two targets on the basis of varying levels of evidence (i.e., the direction of motion coherence in a Random-Dot-Kinematogram). They indicated this choice by making a saccade to one of two targets and then indicated their confidence. Saccade trajectory deviation was taken as a measure of the inhibition of the non-selected target. We found that as evidence increased so did confidence and deviations of saccade trajectory away from the non-selected target. However, a correlational analysis suggested they were not related. In Experiment 2 an option to opt-out of the choice was offered on some trials if choice proved too difficult. In this way we isolated trials on which confidence in target selection was high (i.e., when the option to opt-out was available but not taken). Again saccade trajectory deviations were found not to differ in relation to confidence. In Experiment 3 we directly manipulated confidence, such that participants had high or low task confidence. They showed no differences in saccade trajectory deviations. These results support post-decisional accounts of confidence: evidence supporting the decision to move the eyes is reflected in saccade control, but the confidence that we have in that choice is subject to further post-decisional processes.

**Keywords:** Saccades Decision-Making Confidence

Francis Crick Room

# ACTING INFORMATIVELY: HOW PEOPLE LEARN CAUSAL STRUCTURE FROM INTERVENTIONS

NEIL BRAMLEY *University College London*  
DAVID LAGNADO *University College London*  
MAARTEN SPEEKENBRINK *University College London*

Most decision making research has focused on how people choose actions in static contexts, where the structure of the task environment is assumed to be known and fixed. However, in real world decision making, people are typically uncertain about what model is true or appropriate for the situation, meaning that part of the value of an action comes from the information its outcome can provide. The current research focuses on how people select sequences of actions (tests, manipulations or interventions on parts of a system) when their goal is to learn about the causal structure underlying that system. Steyvers et al (2003) showed that people can often select the most informative single variable to act on to distinguish effectively between a set of alternative causal models in a partially probabilistic scenario. We extend on this research over three experiments.

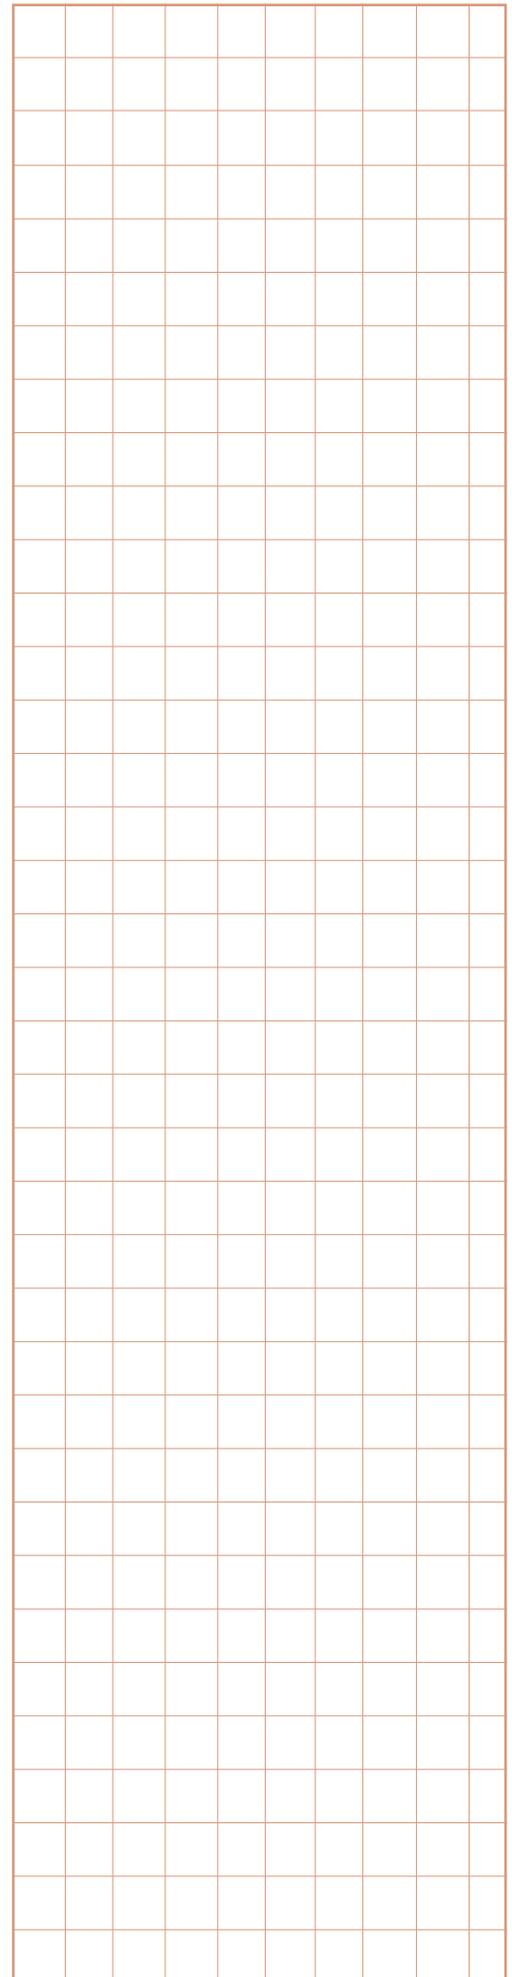
In experiment 1, we explored how people select and learn from sequences of interventions on several fully-probabilistic three-variable causal systems, where they are allowed manipulate more than one variable at a time. We developed and tested models of participants' intervention choices and causal structure judgments, using three measures of the usefulness of interventions: expected utility gain, probability gain and information gain. We also introduced plausible memory and processing constraints. We found that successful participants were best described by a model that acts to maximize information (rather than expected score, or probability of being correct); that discounts much of the evidence received in earlier trials; but that mitigates this by being conservative, preferring structures consistent with earlier stated beliefs.

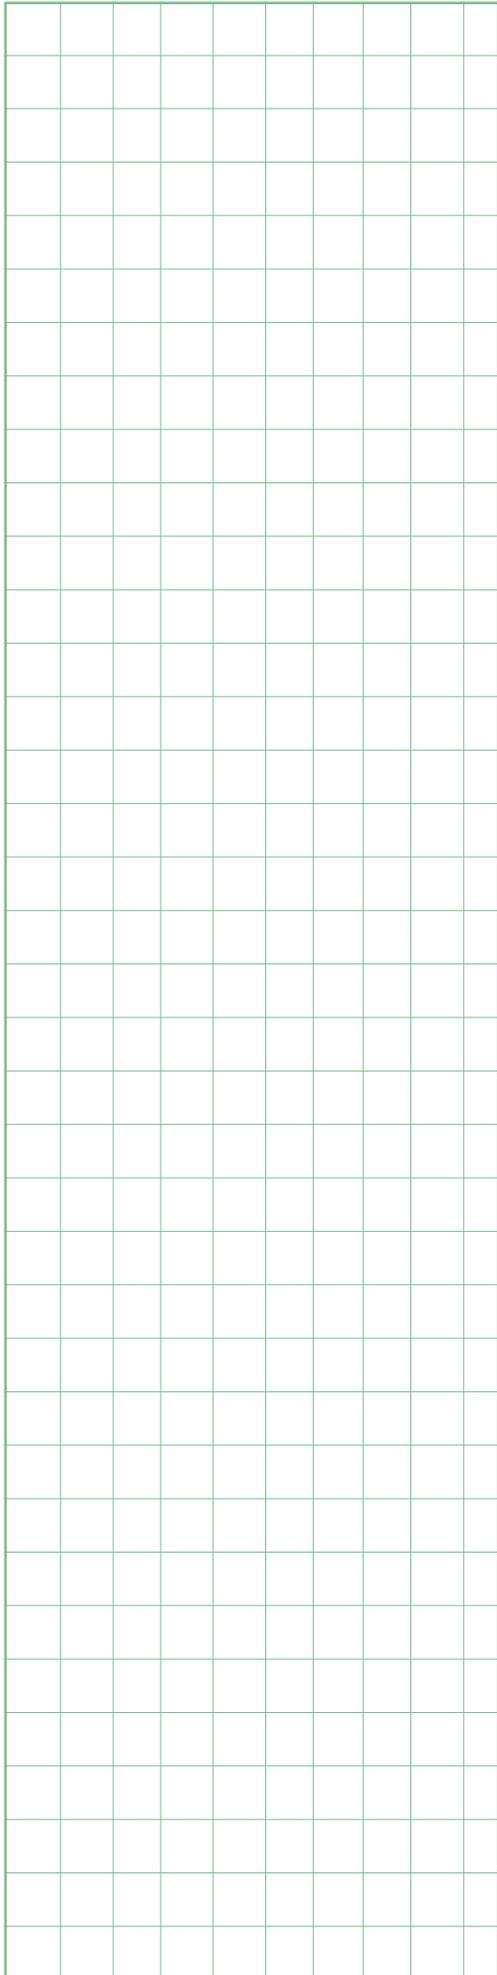
In experiment 2, we replicated our main findings from experiment 1 and established that providing summary information did not significantly reduce participants' accuracy. This suggested that limitations in memory for previous interventions was not a key determiner of learning.

In experiment 3, we focused on identifying boundary conditions for effective active causal learning and exploring active learning heuristics. We independently varied two aspects of the target causal systems: 1. How reliable the causal connections were and 2. How frequently the components of the system turned on by themselves. Within subjects, we also varied the number of causal components and number of connections between these components. We found that peoples' active causal learning abilities were robust and flexible across a wide range of environments. Participants were able to identify causal connections even when they were weak and unreliable but their learning broke down when spontaneous activations were very frequent. Accuracy was unaffected by the number of variables or links, but participants struggled with chain structures in particular.

Overall, the picture emerges that human active learning uses a mixture of heuristics. People exhibit periods in which they focus on "connecting" - endorsing connections between actions and any apparent "effects", ignoring dependencies such as the existence of competing explanations. At other times they focus on "pruning" - disambiguating causal paths and removing unnecessary links from their existing model. We demonstrate that a heuristic active learning model following these principles can largely explain participants' performances.

Keywords: causal, learning, information theory, heuristics, active learning





## RISK PREFERENCES IN CONTEXT

**FLORIAN ARTINGER** *Max Planck Institute for Human Development*

**GORDON D. A. BROWN** *University of Warwick*

**GRAHAM C. LOOMES** *University of Warwick*

Individual preferences for risk are commonly elicited using Multiple Price Lists (Holt & Laury, 2002). How stable are individuals' preferences for risk as assessed by this method? Participants responded to each of four different Multiple Price Lists on two separate occasions. Analysis of responses to identical lists showed that participants exhibit imprecision in their revealed preferences. Analysis of different Multiple Price Lists found that individuals' preferences for risk systematically vary beyond the effect of imprecision. The differences in CRRA obtained between different Multiple Price Lists are large and significant, with 34% of participants even switching from risk averse to risk seeking. The study highlights the context dependent nature of risky choices in one prominent elicitation method and questions the reliability of stable revealed preferences for risk.

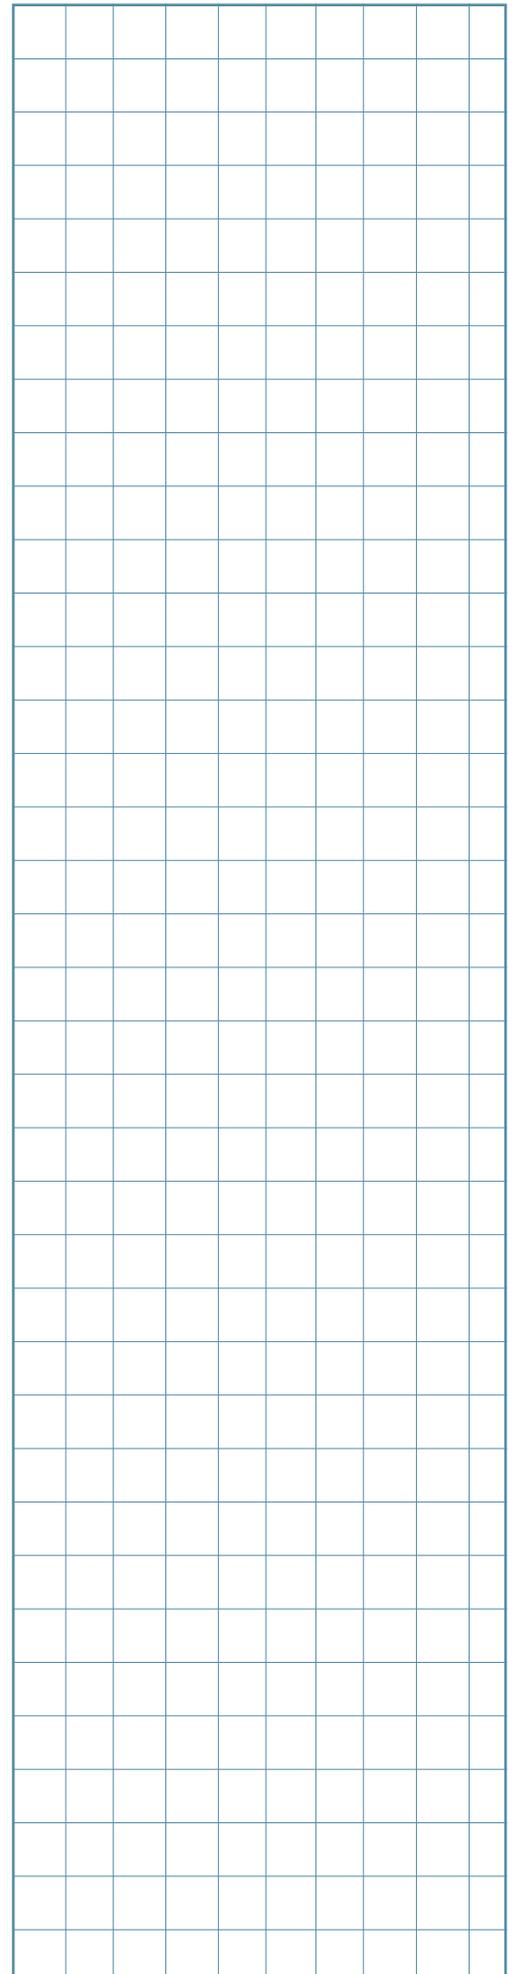
**Keywords:** Risk preferences, context

# THE NEURAL SUBSTRATE OF EYE MOVEMENT CONTROL: INHIBITION, SELECTION AND CHOICE

**ANDREAS JARVSTAD** *University of Bristol*  
**IAIN D. GILCHRIST** *University of Bristol*

The ability to inhibit task-irrelevant but salient stimuli, in favour of task-relevant but less salient stimuli, is crucial for the cognitive control of action. Here, we explore this ability and two other key aspects of cognitive control (selection and choice), with a novel eye movement fMRI paradigm. The paradigm avoids many of the potentially problematic confounds of existing tasks. It is continuous and naturally paced, involves looking at actual objects, deciding which objects to look at and utilises saccade-contingent displays to equate visual stimulation across tasks. In contrast to previous studies, we find that inhibition of task-irrelevant salient stimuli is associated with activity within the core saccadic network (posterior parietal cortex, frontal eye fields), seemingly without a strong reliance on areas external to it (e.g. dorsolateral prefrontal cortex, pre-supplementary motor area, inferior frontal gyrus). Our choice task, which attempts to expose the neural signature of choice, but contains many of the confounds present in existing tasks (e.g., task switching, task set complexity), activates the very same extra-saccadic regions that have previously been linked to inhibition. Our results highlight the importance and challenges of cleanly isolating cognitive functions and suggest limitations to the way cognitive functions are mapped onto neural substrate.

**Keywords:** neuroimaging, motor control, eye movement, cognitive control, inhibition



## FUNCTION LEARNING WHILE MAKING DECISIONS

**HRVOJE STOJIC** *Universitat Pompeu Fabra*

**PANTELIS P. ANALYTIS** *Max Planck Institute for Human Development*

**HENRIK OLSSON** *University of Warwick*

Predictions are essential to decision making. In making these predictions we rely on knowledge of functional relations between attributes we can observe and the outcome we seek to predict. A variety of tasks and models have been developed to investigate and explain how people learn functional relationships, but not in a decision making context. In a function learning task, participants learn to predict the continuous criterion of a single alternative described by some attributes, while in a multi-attribute decision making task, participants decide between two or more alternatives described by some attributes. Differences between these two tasks might be substantial enough that function learning operates differently in decision making situations. The possibility of making comparative judgments in a decision making task, for example, is likely to change attention dynamics. In addition, because in a decision making task people receive feedback only about the chosen alternative while at the same time unlabeled alternatives are present, the learning process might be better described by a semi-supervised learning model.

Given the prevalence of decision making in the real world, much of our conceptual knowledge is acquired while making decisions. Hence, our main contribution of our results lies in establishing relations between two rich literatures - function learning and decision making. We examine how learning is affected by the decision making context and whether qualitatively different models, rather than some of the classical function learning models, are needed to explain the learning process. We designed an experiment where a function learning task is yoked to a multi-attribute decision making condition where participants also provide continuous predictions. We have two function learning conditions where participants are either yoked to see only alternatives that decision makers have chosen, or to see all the alternatives and receive feedback only on the chosen alternatives, but presented in separate trials.

Firstly, we find that decision makers learn functional relationships faster and reach higher levels of accuracy than participants in the function learning tasks. Further, when we compare the decision making condition where participants provide predictions with another decision making condition where participants do not make predictions, we find that both the decision and learning performance deteriorates. To further investigate the source of the differences between the conditions, we use a Kalman smoother method to estimate the importance participants place on each of the attributes and fit formal learning models to the data. We find that a semi-supervised form of learning is unlikely to be the source of the advantage in decision making and further experiments are required to locate the sources of the differences in learning.

**Keywords:** multi-attribute decision making, function learning, learning, semi-supervised learning

James Watson Room

## STANOVICH'S CHALLENGE TO THE "ADAPTIVE RATIONALITY" PROJECT: AN ASSESSMENT

**ANDREA POLONIOLI** *University of Edinburgh*

Gigerenzer, Hertwig and their co-workers in the Centre for Adaptive Behaviour and Cognition (ABC) and Centre for Adaptive Rationality (ARC) have recently articulated an innovative perspective on rational behavior and cognition. I refer to this innovative framework as "adaptive rationality" (henceforth, AR) for simplicity's sake. This framework encompasses a descriptive and a normative element. First, evolution has endowed organisms with an "adaptive toolbox" of simple heuristics, which are basic rules that are easy to apply. Second, although these heuristics are "frugal" and may lead to violations of norms of the so-called "standard picture of rationality" (Stein 1996), they often lead to successful behavior and are therefore "adaptively rational".

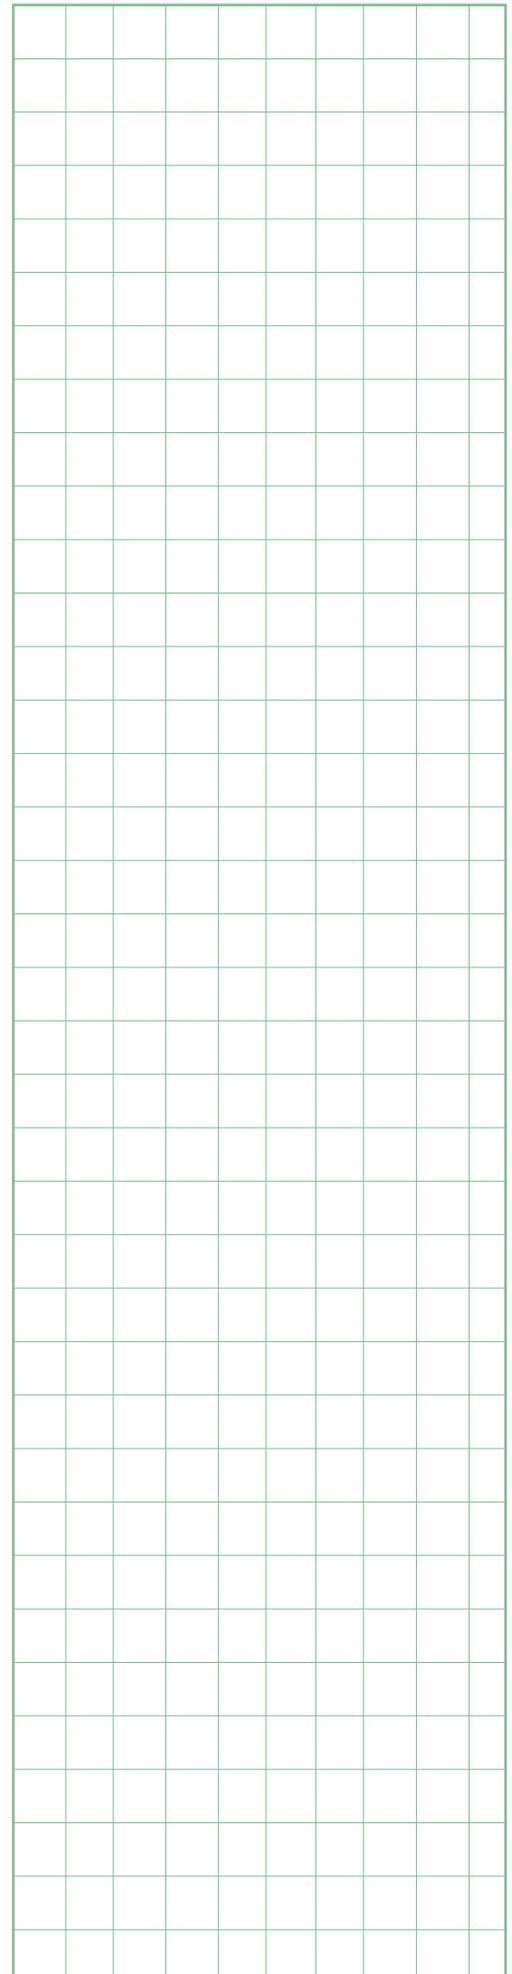
This paper focuses on a challenge to this framework mounted by Stanovich and his colleagues. Stanovich et al. (e.g. Stanovich and West 2000) have conducted an important stream of individual differences studies involving reasoning and decision-making. An important result is that not everyone in reasoning and decision-making tasks follows "simple heuristics". Overall, there seems to be heterogeneity in the use of heuristics. Moreover, Stanovich et al. have found correlations between measures of cognitive ability and tendency to follow standard norms of rationality.

I assess two arguments based on Stanovich's research. The first argument is supposed to challenge the adaptationist background of AR. Heterogeneity in the use of heuristics appears incompatible with the idea that adaptationist pressures led to the use of heuristics: one would expect the use of heuristics to be far closer to universal if they did. The second argument seeks to question the normative claims made by AR theorists: according to Stanovich, the fact that people with higher cognitive abilities follow standard norms of rationality suggests that the "standard picture of rationality" has normative force.

At least as things stand now, none of the arguments can be seen as fully compelling. At least in principle, AR theorists can accommodate findings on heterogeneity, and I discuss several moves open to them. Moreover, even the most plausible version of the second argument turns out to be unconvincing. It is true that there is evidence that people with higher mental abilities achieve better life outcomes (e.g. Deary et al. 2011), and one might be tempted to conclude that they do so because they follow standard norms of rationality. This would suggest that, pace AR theorists, following standard norms of rationality does lead to successful behavior. Yet, the argument remains speculative, as we are unable to draw causal conclusions from the available data.

Whilst both of Stanovich's arguments miss their target, his research might still push forward the "rationality debate" by putting forth two important research questions and encouraging greater experimental and theoretical work. Following recent and seminal studies (e.g., McNamara et al. 2009; Mohlin 2012), AR theorists should offer concrete models and explain how evolutionary processes led to the heterogeneity in the use of strategies that is found in the lab. In addition, causally informative studies are needed to establish whether people with higher cognitive abilities achieve better outcomes because they follow standard norms of rationality.

**Keywords:** Heuristics, Evolutionary Psychology, Adaptive Individual Differences, Cognitive Ability, Cognitive Epidemiology



## A UNIFIED FRAMEWORK FOR SPEEDED SACCADIC AND MANUAL DECISIONS

**ALINE BOMPAS** *Cardiff University; INSERM; CNRS; Lyon Neuroscience Research Center*

**PETROC SUMNER** *Cardiff University*

Action decisions are considered an emergent property of competitive response activations. As such, decision mechanisms are embedded in, and therefore may differ between, different response modalities. Despite this, the saccadic eye movement system is often promoted as a model for all decisions, especially in the fields of electrophysiology and modelling, while other research traditions predominantly use manual button presses, which have very different response distributions and are governed by different brain areas. Here we question and eventually validate the generalisability of core concepts of action decision models from saccadic to manual domains, using two diagnostic phenomena: the remote distractor effect (RDE) and 'saccadic inhibition'. We find that, despite apparent differences, manual responses are also sensitive to the interference of visual distractors and in fact the temporal profile of this interference indicates that the extra delay and variance in manual responses is best accounted for by non-decisional delays, not a difference in the decision process itself. Moreover, a biologically-inspired model (DINASAUR, based on non-linear input dynamics) can account for both saccadic and manual response distributions by simply adjusting the balance of fast transient and slow sustained input.

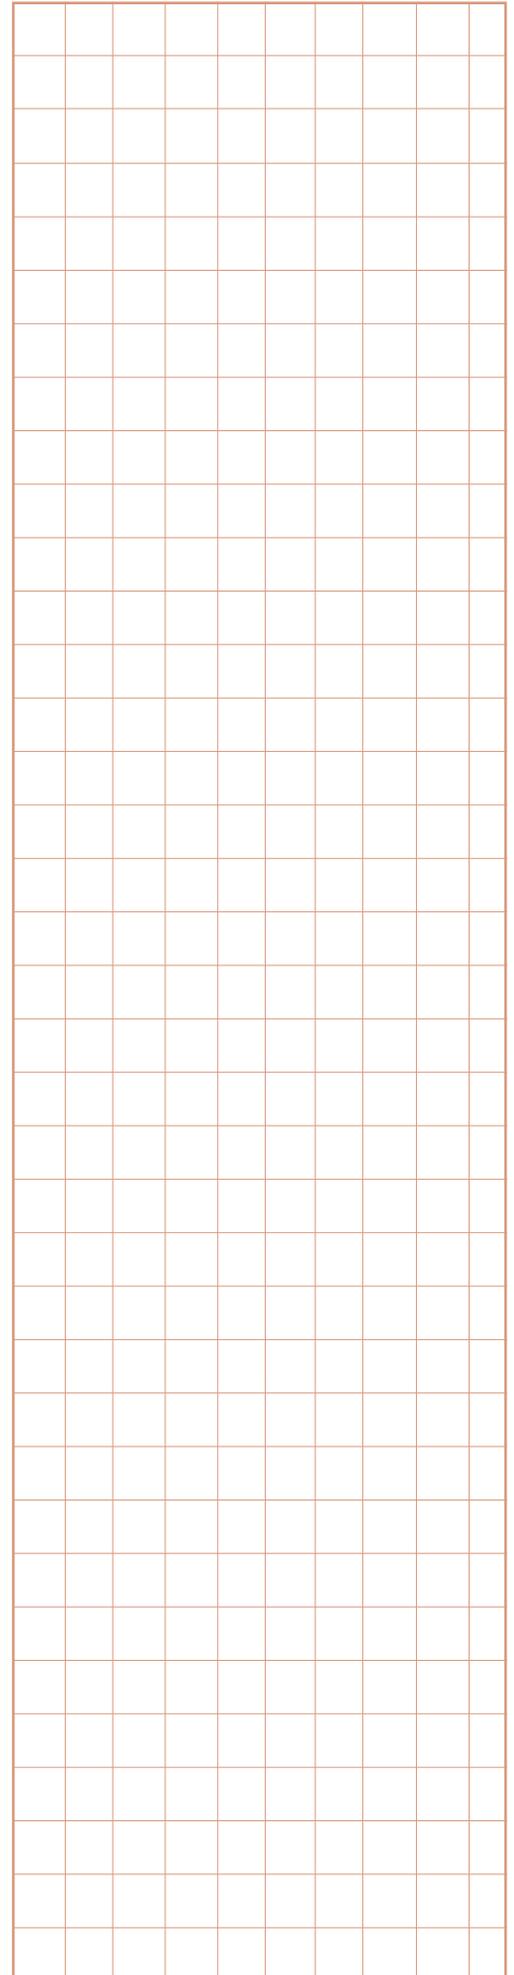
**Keywords:** Action selection; response modalities; competition; reaction time

## CHANCE BASED UNCERTAINTY OF REWARD IMPROVES LONG-TERM MEMORY

**Alice Mason** *University of Bristol*  
**Casimir Ludwig** *University of Bristol*  
**Paul Howard-Jones** *University of Bristol*  
**Simon Farrell** *University of Bristol*

The release of dopamine, in response to rewarding and motivationally salient events, is thought to increase plasticity in the hippocampus. This results in enhanced memory encoding for rewarding events (Lisman & Grace, 2005). How best to promote learning is a central question for both researchers and educators. In a series of studies we asked whether uncertainty of reward outcome has an effect on rewarded memory. We used a simple verbal memory task experiment in which participants were told they could earn small monetary incentives for each word they successfully remembered. Some of the rewards were fixed and others were delivered with a 50:50 probability of reward or no reward. Twenty four hours after the learning phase participants were invited back to lab to perform a recognition memory test. In Experiment 1, we tested whether uncertain rewards promote learning to a greater extent than certain rewards. Based on observations of a sustained ramping response in striatal dopamine between the cue signally reward and reward delivery (Fiorillo, Tobler & Schultz, 2003), we predicted that learning would be greatest during uncertain trials, when there is maximum dopamine for hippocampal dependent consolidation. We found uncertain reward produced a memory enhancement similar in magnitude to fixed reward. In Experiment 2, we used uncertain and certain reward cues, but we also revealed the reward outcome for each trial during the learning phase, so that the design was analogous to the neurophysiological studies. Mixed effects modelling of the data revealed that uncertainty and prediction error are significant drivers of recognition memory. Lastly, we wanted to investigate if uncertain rewards lead to better memory encoding than certain rewards regardless of value. We designed an experiment with uncertain and certain reward values and found superior recognition memory under uncertainty. In our last experiment we have been looking at the differences immediate and delayed memory recall. If these effects are dependent upon increased hippocampal consolidation we would expect them to only be evident or significantly increased after a twenty four hour delay. These studies are the first to demonstrate an influence of uncertainty of reward on motivated learning. Chance based uncertainty is seen as a key component of games and is of core interest to educators in understanding how best to maximise reward based learning (Howard-Jones et al, 2011).

**Keywords:** Reward Uncertainty, Incentivised learning, Episodic memory



## AMBIGUITY ATTITUDE IN COORDINATION PROBLEMS

GRAHAM LOOMES *Warwick Business School*

GANNA POGREBNA *Warwick Manufacturing Group*

Subjective Expected Utility (Savage, 1954) provides a normative account of decision making under risk (when probabilities are known) and under ambiguity (when probabilities are not objectively known). SEUT postulates that individuals have neutral attitude towards ambiguity and subjective probabilities of ambiguous events could be extracted from preferences (Machina and Schmeidler, 1992). Yet, the Ellsberg paradox (Ellsberg, 1961) questions the descriptive validity of SEUT by providing an example where individuals exhibit non-neutral attitude towards ambiguity.

To date, the validity of SEUT and other theories of decision making under ambiguity have been primarily tested in games of nature where an individual faces ambiguity which is constructed using a non-interactive environment. However, in the majority of the real world situations ambiguity arises from the difficulty of predicting the actions of others. This paper examines the implications of SEUT and other theories of decision making under ambiguity in games of strategy where individuals face ambiguity which stems from an interactive environment.

We design a simple experiment where participants are faced with a closed set of three items. Participants are informed that in order to receive a monetary payoff of £20, each of them needs to coordinate with another randomly selected participant by choosing the same item from the closed set. Failure to coordinate results in a zero payoff. Then participants are asked to think about the likelihood of the different choices the other participant could have made. Such events are ambiguous because they depend on the actions of other people and the probabilities are not objectively known.

We elicited certainty equivalents and probability equivalents for each ambiguous event separately as well as for pair of events within the closed set. We compare certainty and probability equivalents obtained from different decision problems and analyse whether and to what extent individuals make consistent decisions throughout the experiment and whether these decisions can be accounted for using SEUT and other deterministic theories. We find that individuals violate SEUT in this strategic environment. In particular, in one subset of experimental tasks many participants exhibit ambiguity averse behaviour, while in another subset their behaviour is often consistent with ambiguity seeking. We provide an explanation of observed behaviour which is based on a model of stochastic choice.

## EYE GAZE REFLECTS AND CAUSES MORAL CHOICE

**PHILIP PÄRNAMETS** *Lund University Cognitive Science*

Eye gaze reflects online processing in many cognitive tasks, such as spatial memory, and linguistic processing (Spivey, 2007). Choices, as they unfold over time, are also reflected in concurrent activation of perceptual and motor systems (Gold & Shadlen, 2000; McKinstry, Dale, Spivey, 2008). This allows for both faster action implementation and for the cognitive system to receive direct feedback from the environment. However, effects of gaze distribution on higher-level, abstract choices, such as moral ones, remain unexplored. Previous studies have established how gaze can affect choice, but only for simple preferential choices, and only by directly controlling the information that is attended to (Shimojo et al. 2003; Armel et al. 2008). Here we investigated this aspect further; we show that knowledge of gaze dynamics can be leveraged to influence choices, even those concerning complex moral issues, by manipulating the timing of the choice alone.

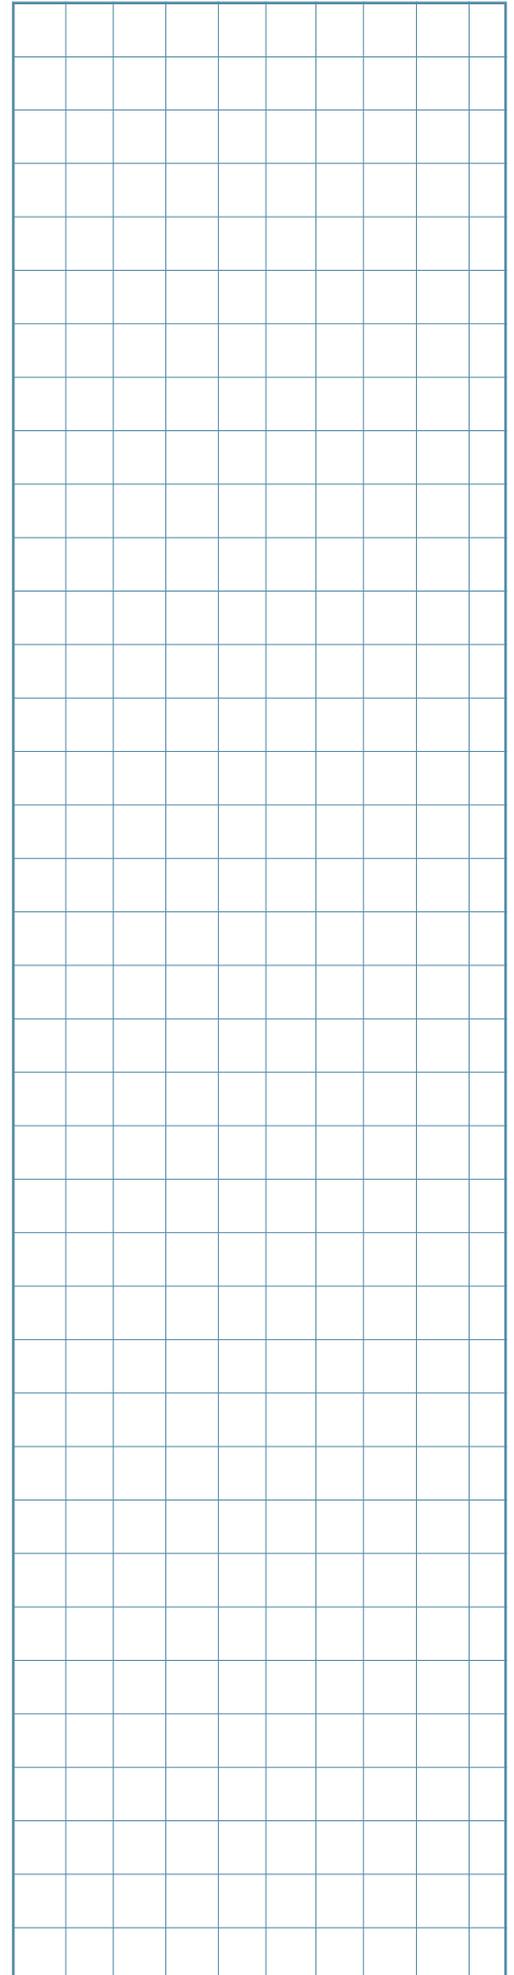
We monitored participants' eye movements during a two-alternative forced choice task with moral and non-moral factual questions. Participants would hear a statement read aloud and see two response alternatives. For example, they would hear a statement such as "Murder is sometimes justifiable" and see the alternatives "Sometimes justifiable" and "Never justifiable". One alternative was randomly predetermined as a target for the manipulation. At the moment participants had fixated the target alternative for at least 750ms and the non-target for at least 250ms, their deliberation was terminated and their choice was prompted. The timing and cause of the choice prompts was not known to the participants nor discovered by them during the course of the experiments.

We found that participants choose the target alternative in 58% of the trials for the moral statements and 55% of the trials for the factual non-moral statements. Response times from choice prompt to button press were significantly faster when choosing the target alternative as was participant confidence.

Having demonstrated this causal relationship between eye gaze and moral choice we also attempted to better characterise the underlying computational mechanisms. We fit diffusion models (Krajbich & Rangel, 2010) with and without dependence on eye gaze to data on moral choices using the same statements as before. We found that an eye gaze dependent model provides a better fit and can capture many features of the empirical data.

We conclude that by tracking the interplay between a decision maker, her perceptual-motor systems, and the environment, it is possible to influence the outcome of a choice without manipulating the information available to her. High-level, abstract cognition, such as moral choices, is partly caused by interactions with the immediate visual environment, and this can be described using similar models as for simple preferential and perceptual choices.

**Keywords:** Moral choice, visual attention, modelling



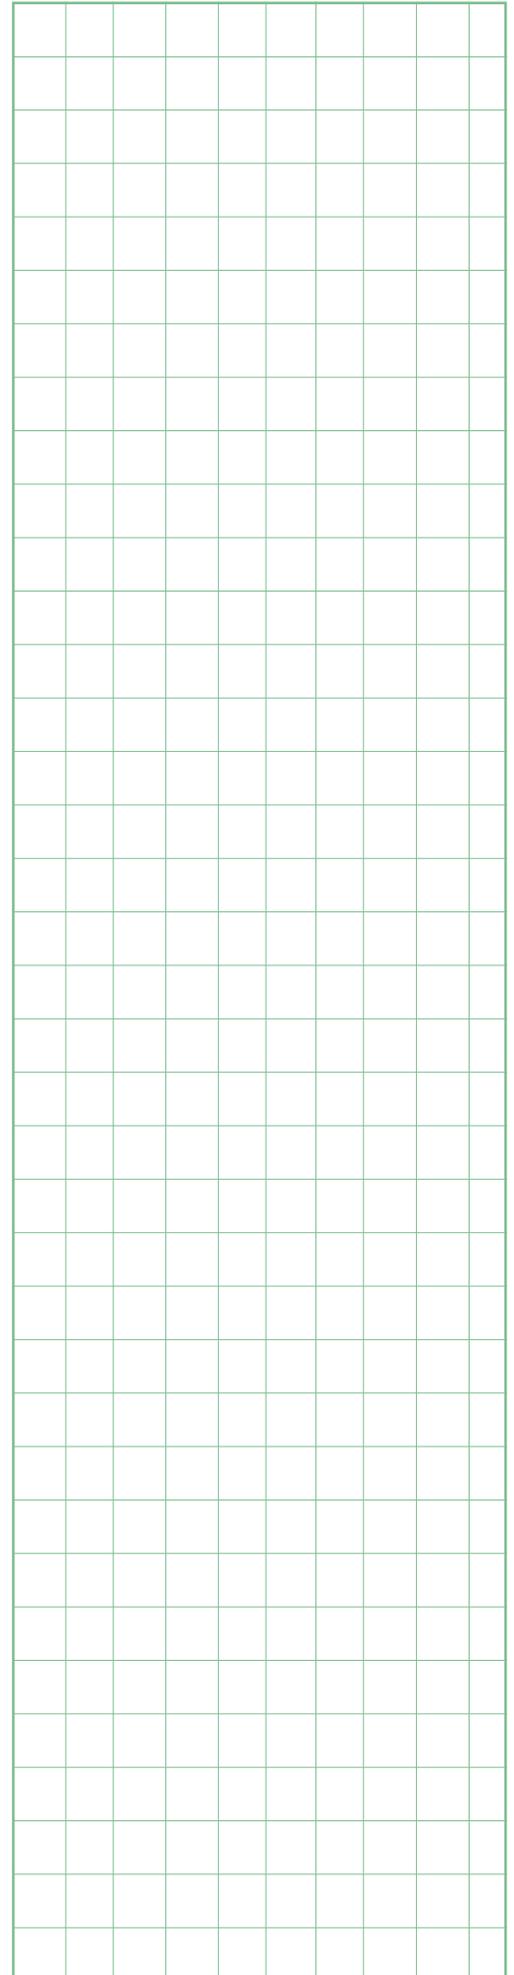
## ' IRRATIONAL ' BEHAVIOUR CAN ARISE FROM DECISION RULES ADAPTED TO COMPLEX ENVIRONMENTS

TIM W. FAWCETT *University of Bristol*  
BENJA FALLENSTEIN *University of Bristol*  
ANDREW D. HIGGINSON *University of Bristol*  
ALASDAIR I. HOUSTON *University of Bristol*  
DAVE E.W. MALLPRESS *University of Bristol*  
PETE C. TRIMMER *University of Bristol*  
JOHN M. MCNAMARA *University of Bristol*

Patterns of decision-making in humans reveal some striking deviations from economically rational expectations. These include distorted beliefs about external events, inconsistent preferences that are altered by past experience and current context, and apparent violations of the axioms of rational choice theory. There is mounting evidence that analogous biases exist in other organisms, which hints at an evolutionary explanation. Yet such cases of apparently irrational behaviour seem difficult to reconcile with the fundamental biological concept of natural selection as an optimising process. Here I summarise recent theoretical research from our group that suggests a possible solution to this conundrum. Our work demonstrates that several common biases in decision-making may in fact result from ecologically rational decision rules adapted to exploit temporally autocorrelated environments. Temporal autocorrelation is ubiquitous in nature and implies that conditions experienced now or in the recent past are predictive of future conditions, which can affect the consequences of current decisions. Many standard laboratory procedures used to study decision-making do not reflect this statistical structure, and in such settings ecologically rational decision rules can lead to biased or irrational behaviour. This evolutionary perspective can account for some well-known deviations from economic rationality, including intransitivity and irregularity, successive contrast effects and the fourfold pattern of prospect theory. We encourage other researchers to consider the richness of natural environments to understand better how evolution has shaped our cognitive systems. The real world can be complex, variable and autocorrelated, and we should expect cognitive and perceptual systems to have evolved to exploit its statistical structure.

**Keywords:** cognitive bias, irrational behaviour, evolution, economic rationality, ecological rationality, temporal autocorrelation, intransitivity, irregularity, successive contrast effect, prospect theory, fourfold pattern

James Watson Room



## AMBIGUITY IN COORDINATION GAMES

**DAVID KELSEY** *University of Exeter*  
**SARA LEROUX** *Oxford Brookes*

We report an experimental test of the influence of ambiguity on behaviour in a coordination game. We study the behaviour of subjects in the presence of ambiguity and attempt to determine whether they prefer to choose an ambiguity safe option. We find that this strategy, which is not played in either Nash equilibrium or iterated dominance equilibrium, is indeed chosen quite frequently. This provides evidence that ambiguity aversion influences behaviour in games. While the behaviour of the Row Player is consistent with randomising between her strategies, the Column Player shows a marked preference for avoiding ambiguity and choosing his ambiguity-safe strategy.

**Keywords:** Ambiguity; Choquet expected utility; coordination game; Ellsberg urn, experimental economics

## DECISION-MAKING DEFICITS IN NEURODEGENERATIVE DISEASES

JIA XIANG ZHANG *Medical Research Council, Cambridge, UK*

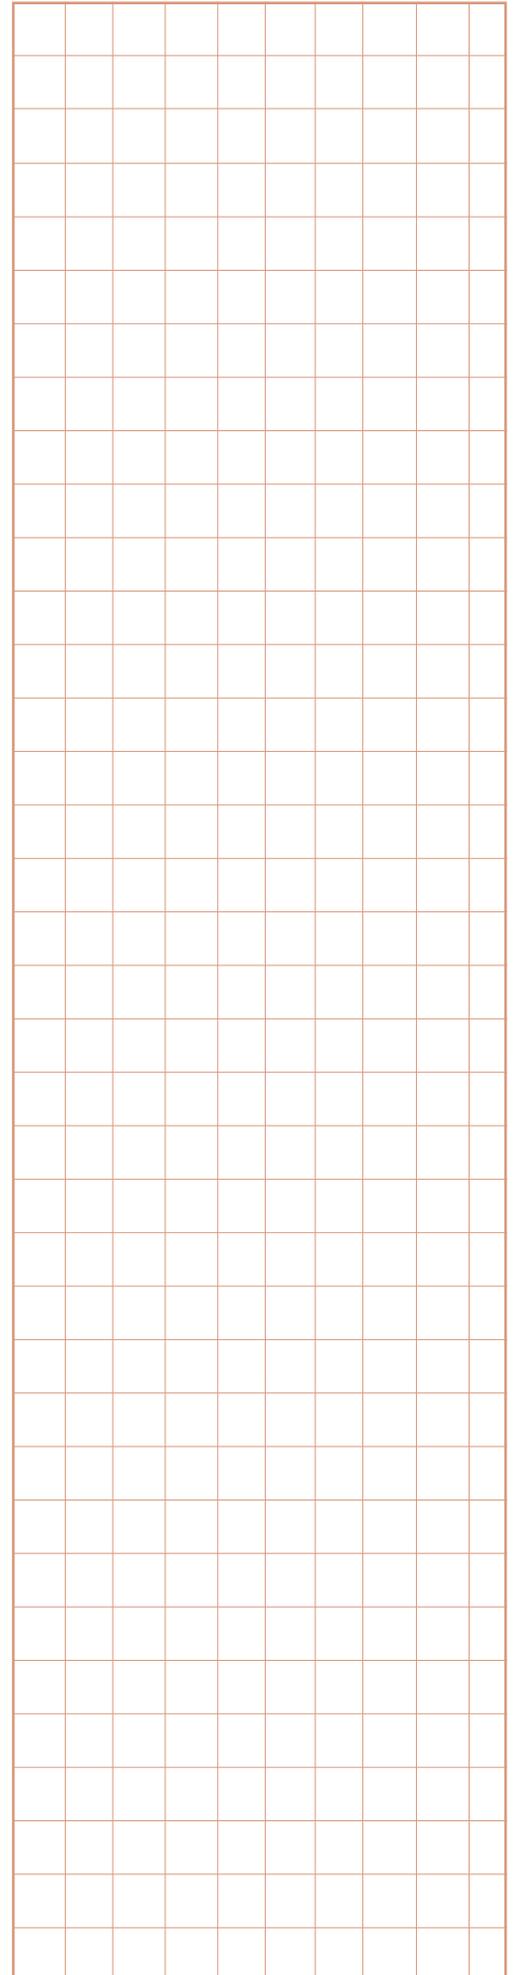
The search for symptomatic and disease-modifying therapies for neurodegenerative diseases urgently requires a mechanistic understanding of cognitive deficits in patients. Here we present two recent studies on decision-making impairments in different cognitive domains.

The first study investigated impairments of response inhibition in Progressive Supranuclear Palsy (PSP) and Parkinson's disease (PD). The two diseases have different neuropathology, but both affect cognitive function including impulsivity. We tested nineteen PSP, twenty-four PD and twenty-six healthy controls in a saccadic go/no-go task by using a head-mounted infrared saccadometer. Subjects were cued on each trial to make a pro-saccade to a horizontal target or withhold their saccadic responses. We modelled the paradigm as a diffusion process between a "go" and a "no-go" decision boundary and fitted a hierarchical drift-diffusion model to behavioural data by using Bayesian parameter estimation approach. Both PSP and PD patients were impulsive: they had more commission errors in the "no-go" condition, but the commission errors and mean saccade latencies were similar between PSP and PD patients. However, the posterior model parameters revealed significant disease-specific deficits in the mechanisms underlying go/no-go decisions. Both PSP and PD patients had slower information accumulation rate and faster non-decision time than controls, but PSP patients were severely biased towards the pro-saccade decision boundary than PDs and controls. The combination of a well-tolerated oculomotor paradigm and the sensitivity of the Bayesian modelling approach provide a useful biomarker for distinguishing different neurodegenerative diseases and provide a rational basis on which to develop and assess new therapeutic strategies.

In the second study, we examined the performance of sub-second interval discrimination in PD. Eighteen non-demented and medicated PD patients and nineteen healthy controls performed two duration discrimination tasks, during which the participants were required to categorize auditory durations into two (bisection) or three (trisection) interval categories. The patients with PD exhibited impaired interval sensitivity than the controls in both tasks and biased interval decisions in the bisection task. The PD patients also showed maladaptive decision strategies that failed to slow down in discriminating ambiguous intervals. Furthermore, the impaired timing performance in the PD patients correlated with their disease severity and doses of dopaminergic medication. The findings suggest that impaired interval discrimination in PD cannot simply be attributed to pathophysiological distortions in an internal clock, but also associated with impulsive decision-making processes which are biased towards premature responses.

**Keywords:** Decision-making, go/no-go, interval discrimination, PD, PSP, modelling

Francis Crick Room

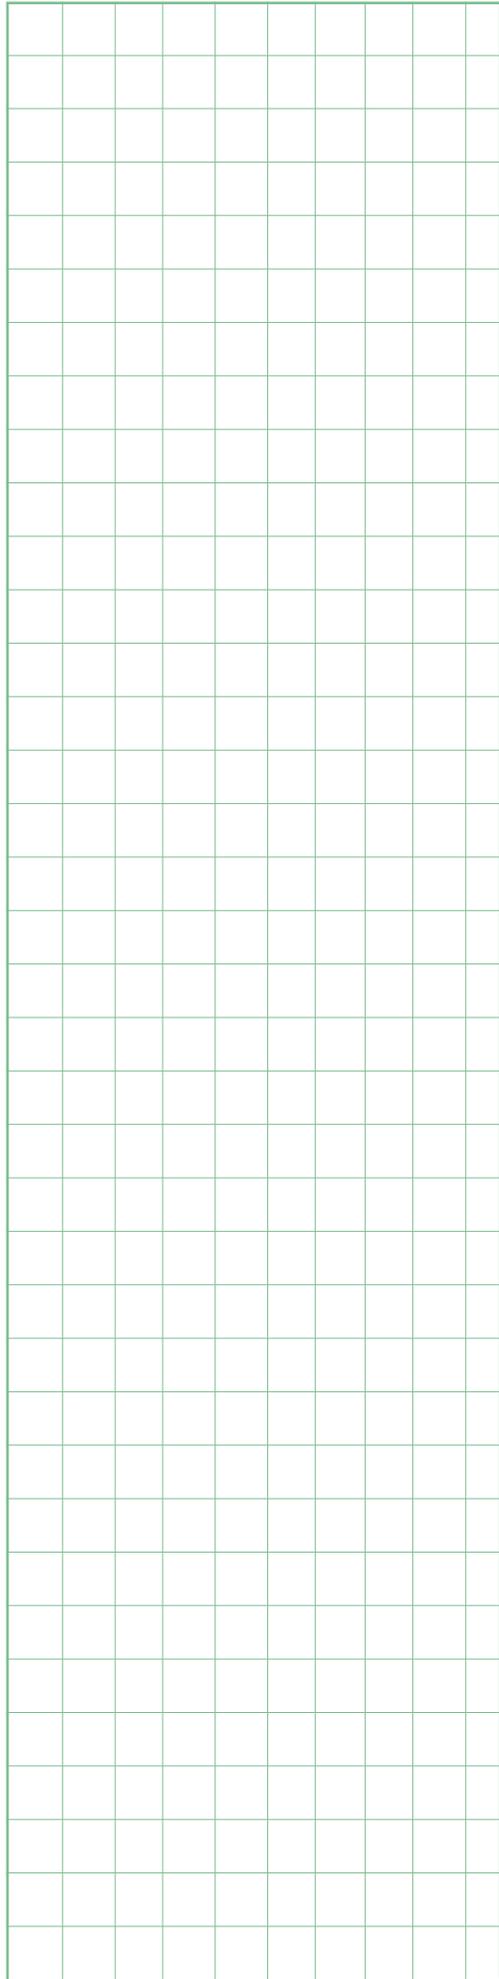


## RELATIVE RANK THEORY

GORDON D. A. BROWN *University of Warwick*

How do people value states of health, decide on a fair price for a product, or determine the appropriate amount of damages to award against a polluting company? Here I describe and discuss a process I call “relative rank matching”. The subjective magnitude of quantities such as prices, health states, or crimes are assumed to be determined by contextual comparison involving rank-based principles such as those embodied in Range Frequency Theory and Decision by Sampling. However, such models are often silent on the question of how comparisons are made across incommensurable dimensions. Subjective judgements are assumed to be entirely relative, yet we have no difficulty rejecting a “relatively good” bottle of wine in favour of a “relatively bad” house. Although relative judgements cannot themselves provide a common currency for comparing options, it is suggested that relativity-matching translation into a common distribution (such as market prices) is often possible and enables comparison across different dimensions. When a suitable matchable dimension such as a market price distribution is unavailable, however, our valuations are inconsistent and unreliable.

**Keywords:** Judgement; decision making; relative rank



## THE EFFECTS OF MISINFORMATION ON REASONING AND DECISION MAKING

**ULLRICH ECKER**, *University of Western Australia*  
**STEPHAN LEWANDOWSKY**, *University of Bristol*

The quality of decisions is determined largely by the quality of the information they are based on. Misinformation and inaccurate beliefs are therefore of major concern: If a majority believes in something that is factually incorrect, the misinformation may form the basis for political and societal decisions that run counter to a society's best interest; if individuals are misinformed, they may likewise make decisions for themselves and their families that are not in their best interest and can have serious consequences. For example, following the unsubstantiated claims of a vaccination-autism link, many parents decided not to immunize their children, which has had dire consequences for both individuals and societies, including a marked increase in vaccine-preventable disease and hence preventable hospitalizations and deaths, as well as unnecessary expenditure for follow-up research and public-information campaigns aimed at rectifying the situation.

Misinformation comes in many guises, including explicitly provided falsehoods, subtly misleading headlines, or urban myths. What makes misinformation particularly concerning is its resistance to correction. In this talk, we will illustrate how misinformation affects memory, reasoning, and decision making, and we will present experimental evidence for the ineffectiveness of corrections. We will discuss the role of people's attitudes towards the corrected misinformation, and will highlight strategies to increase the potency of myth rebuttals and refutations.

**Keywords:** Misinformation; Reasoning; Memory

THE VARIANCE OF NEURONAL ACTIVITY IN  
DORSAL PREMOTOR AREA PREDICTS  
BEHAVIORAL PERFORMANCE:  
THE MODULATION OF BEHAVIOR AND  
NEURAL ACTIVITY BY TRIAL HISTORY

**ENCARNI MARCOS** *Universitat Pompeu Fabra, Barcelona*

**PIERPAOLO PANI** *Sapienza University of Rome*

**EMILIANO BRUNAMONTI** *Sapienza University of Rome*

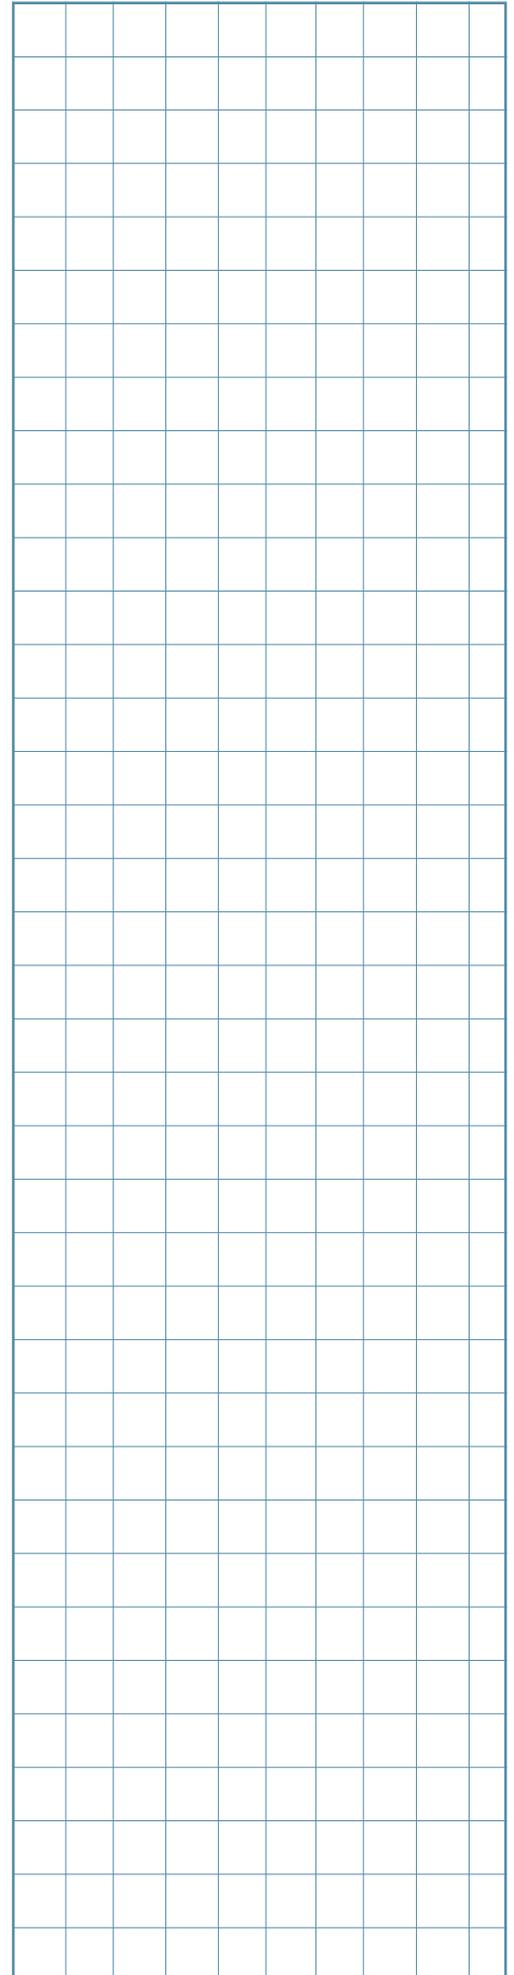
**GUSTAVO DECO** *Universitat Pompeu Fabra, Barcelona; ICREA Institució  
Catalana de Recerca i Estudis Avançats, Barcelona*

**STEFANO FERRAINA** *Sapienza University of Rome*

**PAUL VERSCHURE** *Universitat Pompeu Fabra, Barcelona; ICREA Institució  
Catalana de Recerca i Estudis Avançats, Barcelona*

In decision making information stemming from both memory and perception is integrated in the service of the goal oriented action of the agent. In a previous study, we have proposed that the behavior of the agent is biased by the sequentiality of external events and that this biasing is achieved by a proper chaining of the memory space (Verschure et al., 2003). However, the underlying neural mechanisms causing this bias in behavior are not well understood yet. Here, we investigate this issue by looking at the activity of neurons in the dorsal premotor area (PMd) of two macaque monkeys while they performed a countermanding reaching task. The countermanding task has been extensively used to study the ability to cancel a planned movement when an infrequent delayed stop signal appears. In this task, the probability of success mainly depends on the temporal distance between a go-signal (a visual cue that indicates that a movement should be initiated) and the so-called stop-signal (a second visual cue that in a minority of trials indicates that the previous planned movement has to be cancelled). Recent results have shown that the behavior of subjects in a single trial is also influenced by the history of that trial, i. e. reaction time (RT) of subjects is longer when one or many Stop trials have been recently experienced (Emeric et al., 2007). In the present study, we investigate the neural mechanisms causing this modulation in behavior due to the task history. We found no relationship between mean firing rate of PMd neurons and the changes of RT due to previous experience, i. e. whether a Go trial was preceded by a Stop or a Go trial. In contrast, we observed that the across-trial variability of the neural responses (measured by the variance of the conditional expectation, VarCE, Churchland et al., 2011) showed the same modulation with trial history than the mean and STD of the RT and that this variability could be used as a predictor of RTs in these cases. An additional test with a theoretical model suggested that a system that continuously monitors the recent history of a trial is necessary and sufficient to explain the observed behavioral and neural modulation with trial history. These results raise two important questions: whether this monitoring system actually exists and whether it lies inside or outside PMd.

**Keywords:** decision-making, variance, trial history, PMd, motoring system



## NATURE'S MEASURING TAPE: A COGNITIVE BASIS FOR ADAPTIVE UTILITY

**TATIANA KORNIENKO** *University of Edinburgh*

How happy would Robinson Crusoe be with a particular coconut on his desert island? Undoubtedly, that would depend on the coconut's size. Yet, measuring tapes and scales are evolutionarily advanced technology, inaccessible until very recently. How would Nature enable individuals to judge, or measure, magnitudes, and what is the relationship of such magnitude judgments to a utility function?

I explore the process and tools which can be used to judge, or measure, a given magnitude. Following Robson [2001, 2002], I assume that, rather than endowing living organisms with all necessary judgment scales, Nature only provides "tools" that enable one to extract information from one's environment and experience. To deal with more evolutionarily recent tasks, the candidate cognitive tools have to be independent of the nature of the measurement.

I identify two well-documented domain-independent cognitive tools and, using a parsimonious mathematical model, show that one can judge, or measure, a magnitude of an item entirely using ordinal comparisons which, by means of a frequency (proportion) tool, are keyed into a universal cardinal scale. By calculating how frequently a given object "wins" a pairwise ordinal "tournament" against all other objects in the reference set, one can map an arbitrary set of modalities (such as quantities, sizes, weights, durations, luminosities, and so on) onto the interval [0,1]. Such process of evaluation by ordinal rank is a basis of "decision by sampling" in Stewart, Chater and Brown [2006], which been used to explain "preference anomalies".

I argue that the resulting magnitude judgment is a candidate for a utility function representing "more is better" preferences. That is, a utility function may simply be an artefact of magnitude judgment. Yet such magnitude judgment is reference-dependent, leading to a possibility that one's marginal utility (and thus apparent attitudes towards risk) could be shaped entirely by one's experiences, memory, and cognitive imperfections. I show that a (neoclassical) context-independent utility function may arise as a special case of this magnitude judgment procedure when one judges a magnitude relatively to a remembered sample and has a long memory. In contrast, if one's memory is bounded, the well documented context effects tend to arise.

The model presented here allows to accommodate well known perceptual errors in the two cognitive tools and to explore the effects of such cognitive limitations on magnitude judgment. One notable result is that a certain class of cognitive distortions of the mental line would lead to apparent risk aversion in an environment where risk neutrality would have been optimal in the sense of Robson [2001] and Netzer [2009].

Despite the asocial nature of cognitive processes, the procedure involves built-in social comparisons. I show that welfare evaluation of redistributive policies and economic growth depend on the bias in ordinal imperfections. Furthermore, as judgment of own skill relatively to the others involves interpersonal comparisons, the proposed model allows one to relate individual difficulties in making upward ordinal comparisons with the observed patterns of overconfidence in relative skill judgment.

**Keywords:** Reference-dependent preferences, random utility, procedural approach, memory limitations, perceptual imperfections, Weber-Fechner law, decision-by-sampling, overconfidence, welfare.

## NATURALISTIC MULTI-ATTRIBUTE CHOICE

**SUDEEP BHATIA** *University of Warwick*

**NEIL STEWART** *University of Warwick*

Multi-attribute choice is an important area of enquiry in decision making research. While existing work on multi-attribute choice has greatly improved our understanding of behavior, it suffers from an important limitation. Multi-attribute choice theories need observable and quantifiable attributes in order to be tested. However, most everyday choices --from deciding which movie to watch, to what food to eat-- involve objects defined on latent attribute dimensions. How can we study everyday multi-attribute choice without quantified, observable information about the attributes that these choices involve?

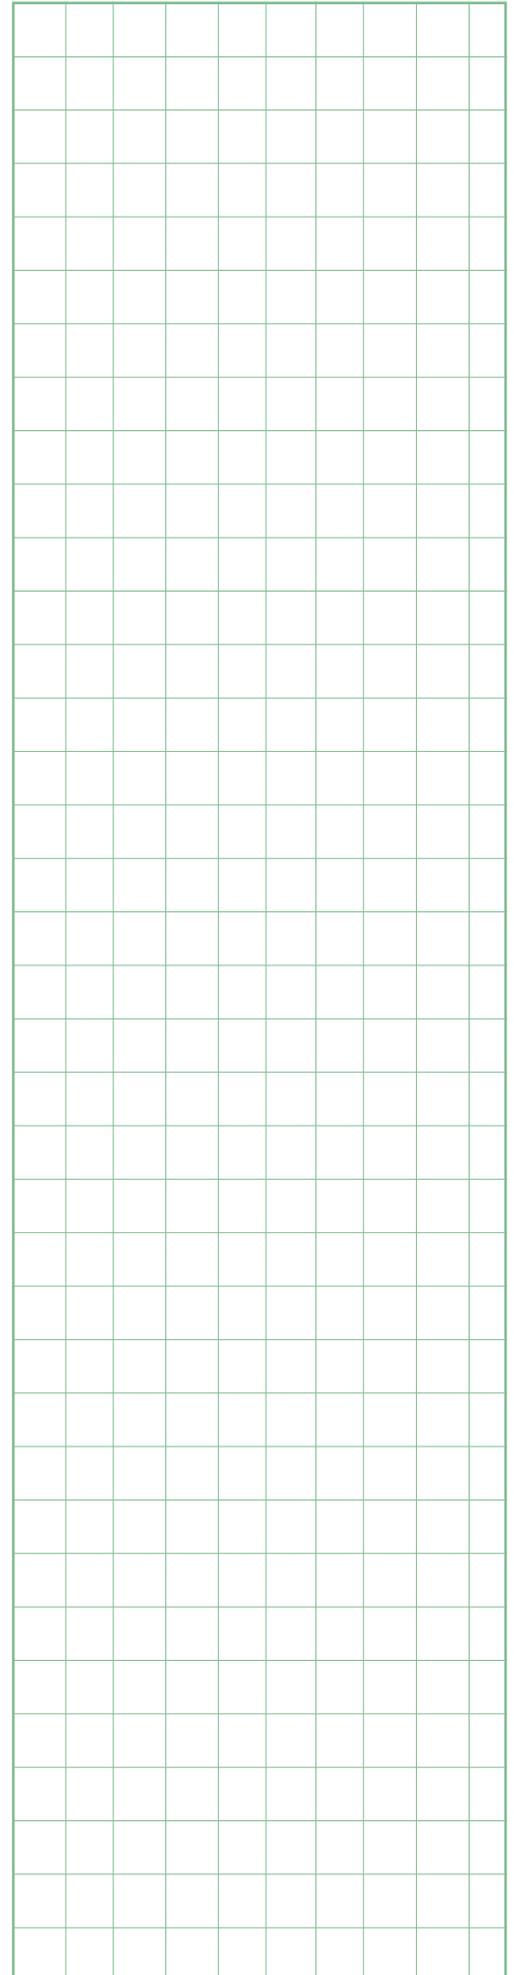
In this paper we offer a solution to this problem. We argue that large datasets with detailed text descriptions of everyday choice objects can be used to uncover the latent attributes that these objects are defined on. These latent attributes can then be used to study multi-attribute choice in its naturalistic setting.

We applied our approach to predicting movie choices, using text data from the Internet Movie Data Base (IMDB). Particularly, we obtained the plot synopses of the 500 most voted movies on IMDB. These synopses had an average length of 1,025 words, and offered detailed descriptions of the movie plots of the most popular movies in the world. We processed these synopses with Latent Semantic Analysis (LSA), a natural language processing technique, used to analyze the conceptual structure of text corpuses. By applying LSA to our dataset, we were able to discover the latent attribute dimensions that characterize a decision maker's movie universe.

Having obtained quantified attribute values for different movies, we tested whether these attributes did in fact play a role in decision making. For this, we asked 75 participants to make 200 hypothetical choices between different movie triplets, selected from the 100 most voted movies on Netflix. We assumed that decision makers utilized a probabilistic version of weighted-additive decision rule, applied to the top five latent attributes obtained from LSA. We fit both group-level and individual-level attribute weights to the movie choices obtained in our experiment, and found that our model provided a good fit to the data. Particularly, 55 out of the 75 participants had a statistically significant fit (average log-likelihood = -207.99,  $p < 0.05$  using the likelihood-ratio test) with the weighted additive model on the five latent attributes. Our model also had a statistically significant fit on the pooled choice data from our experiment (log-likelihood = -16,322.96,  $p < 0.01$  using the likelihood-ratio test).

Overall the result of our study show that latent attributes obtained from large text corpuses can be used to predict everyday choice. Decision making research does not need to limit itself to artificial decision environments in which attribute information is explicitly presented to participants as part of an experimental choice task. Rather, online datasets combined with powerful text processing techniques, can extend the insights of multi-attribute decision research to the naturalistic settings in which most of our choices are made. Future work will extend our analysis to other naturalistic domains, and additionally use these domains to compare the descriptive power of different behavioral theories of multi-attribute choice.

**Keywords:** decision making, latent semantic analysis, multi-attribute choice



# LARGE CORTICAL NETWORKS IN A SMALL NUCLEUS: A 7T fMRI STUDY ON LIMBIC, COGNITIVE AND ASSOCIATIVE SUBPARTS IN THE SUBTHALAMIC NUCLEUS DURING DECISION-MAKING

**GILLES DE HOLLANDER** *University of Amsterdam, Amsterdam Brain and Cognition*

**ANDREAS SCHÄFER** *University of Amsterdam, Amsterdam Brain and Cognition*

**ROBERT TURNER** *University of Amsterdam, Amsterdam Brain and Cognition*

**BIRTE FORSTMANN** *University of Amsterdam, Amsterdam Brain and Cognition*

The subthalamic nucleus (STN) is a nucleus in the basal ganglia of approximately 100mm<sup>3</sup> in size and a crucial node in the cortico-basal ganglia decision-making network. It is theorized to act as a global brake on the decision process, heightening the threshold during difficult decisions, as well as implementing stopping behavior.

Parkinson's disease (PD) is characterized by the loss of dopamenergic cell in the substantia nigra and, consequently, a hyperactivated STN. PD is now often treated using deep brain stimulation (DBS) of the STN, alleviating its 'motor' symptoms such as tremor. This treatment is known to induce severe side-effects such as cognitive decline, compulsive gambling, depression, and even suicide. In the clinical literature it is often proposed that the STN can be subdivided in discrete cognitive, limbic and motoric parts, differentially connected to functional counterparts in the cortex (Temel et al. 2005). The side-effects of DBS are then explained by the stimulation of non-motoric parts as opposed to the motor part of the STN. A recent review of the empirical literature (Keuken et al, 2012), however, shows inconsistent results in the number and location of these subdivisions across studies.

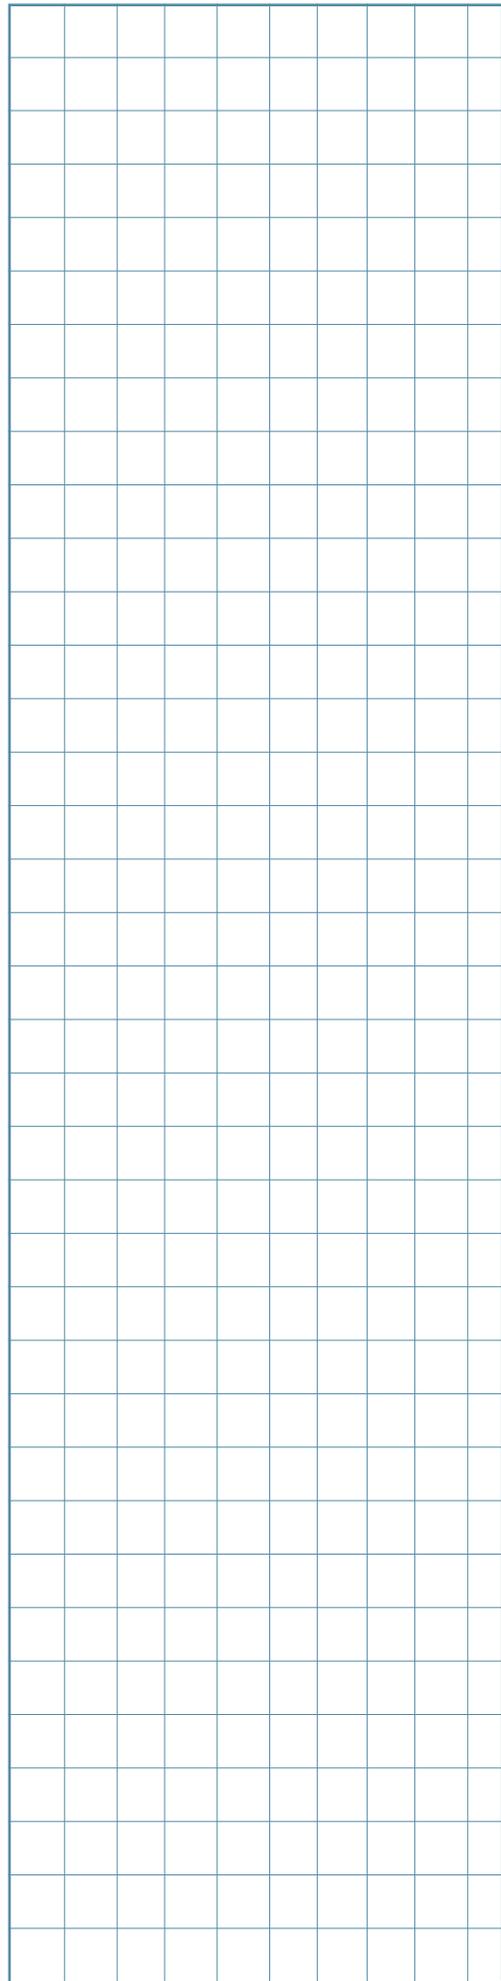
In this study we aimed to investigate these putative subdivisions using ultra-high resolution 7T functional magnetic resonance imaging (fMRI), with a voxel resolution of 0.83x0.83.1.0 mm and individually segmented STN masks, which is much more detailed approach than the often-used resolution at 3T imaging of approximately 3mm isotropic with standard population masks/coordinates. We used a newly developed decision-making paradigm, which contained 3 experimental manipulations, targeted at inducing differential activation of these putative subdivisions and their cortical counterparts.

The subjects were shown pictures from the IAPS-dataset, a standardized dataset of emotionally valenced pictures. One half of these pictures were rated as of neutral emotional valence, the other half as of negative emotional valence. The pictures were colorized using a filter, making them slightly green or slightly red. One half of the pictures was more colorized than the other half. The goal of the subject was to indicate whether the picture they saw was red or green. Additionally, on 25% of the trials an auditory stop-signal was presented and the subject was instructed to then withhold their response.

There were four experimental conditions that are related to fMRI BOLD signal in both cortex and STN: 1. Emotional Valence (limbic network) 2. Difficulty (cognitive network) 3. Stop Trial (cognitive/motor network) 4. Response direction (motor network). The stop signal linear ballistic accumulator model was used to control for any interactions between emotional valence, difficulty and stopping behavior.

Preliminary, unpublished results, both behaviorally and from the neuroimaging will be presented and discussed.

**Keywords:** 7T fMRI, Basal Ganglia, Subthalamic Nucleus, Subdivisions, Topography



Francis Crick Room

## IS FINANCIAL RISK ATTITUDE ENTIRELY RELATIVE?

**SARAH SMITH** *University College London*

**ADAM J. L. HARRIS** *University College London*

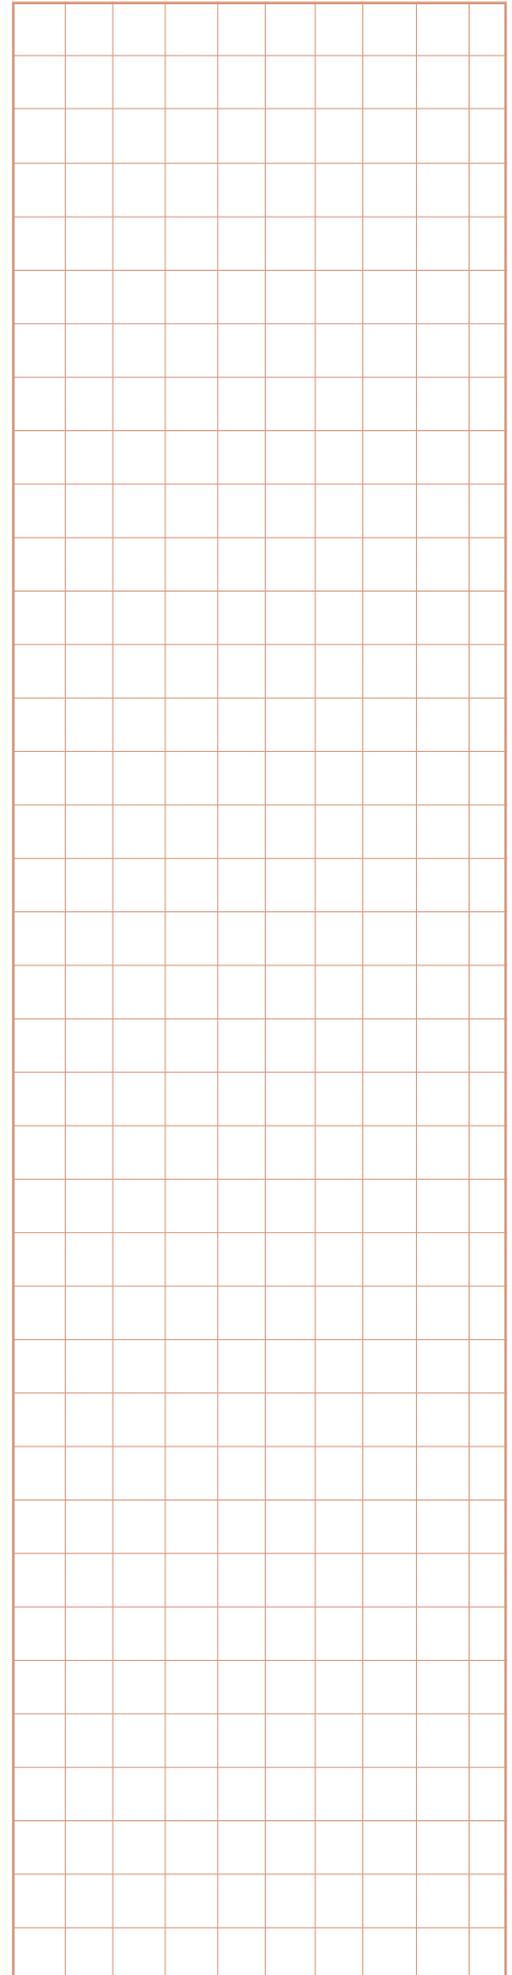
Evidence that preferences may be relative not absolute raises concerns about the practice of using risk attitude measures to guide investment choice. For example, Stewart et al. (2003) found that valuations of gambles almost entirely depended on the set of choices presented, and Vlaev et al. (2007) found similar effects in the selection of pension funds. In the current studies, we investigated the relativity of risk attitude by manipulating the prior experience of participants, rather than the options available to them in the immediate choice context. This addresses the potential that previous findings arose because the constrained set of choices did not include the decision maker's (absolutely) preferred option, and findings reflect a forced change in choice strategy rather than a lack of absolute preference. Additionally we aimed to test the degree to which the effects are the result of a judgment system that is fundamentally relative, rather than an easily overridden heuristic.

We conducted two studies in which the experimental manipulation was a prior experience of making pairwise choices between pension funds, and the dependent variable was the final choice of a fund from a full set. Thus, if a decision-maker had an internal, stable, absolute risk preference she could make her final choice in line with it. Study 1 had four conditions: prior experience constrained to high risk, to low risk, or to balanced prospects, and no prior experience. We found condition significantly affected final fund choice. Participants whose prior experience was of high risk prospects made a higher risk final choice from the full range. However, there was also evidence of some sensitivity to absolute values.

Study 2 examined context effects within-participants and tested a bias warning to see if context effects result from a heuristic process that can be overridden. All participants experienced either the high or low risk pair-wise choices and then made a choice from the full range. Fourteen-days later they did the opposite test. Half of the participants were warned about the potential influence on them of the choice set. We found significant context effects. Participants selected a lower risk fund after experiencing low risk prospects than they did after experiencing high risk prospects. Effects were not influenced by order, or attenuated by the bias warning. Again there was evidence of absolute preference: at the participant level 29% of participants chose the same investment fund after both high and low risk experiences.

These prospect relativity effects violate economic principles of stable preference and are similar to those found in psychophysical judgments. They are consistent with Range Frequency Theory (Parducci, 1965), and with a Decision by Sampling account of the choice process (Stewart et al. 2006). However, absolute preference appears to also have a role. How absolute preference and relative judgments might work simultaneously to influence risk attitude needs consideration. From an applied perspective, customers' risk preferences are susceptible to manipulation in the immediate decision making environment, and in any preceding preparation tasks or online journeys.

**Keywords:** judgment and decision making, prospect relativity, decision by sampling, financial risk attitude



## DECISION-MAKING IN ADJUVANT CANCER TREATMENT CHOICES BASED ON COMPLEX INFORMATION, INCLUDING GENOMIC RECURRENCE RISK SCORES

**JULIA BRETTSCHNEIDER** *University of Warwick*

**ANJALI MAZUMDER** *University of Warwick*

One of the recent developments in medical genomics are refined prognostic tests providing so-called recurrence scores for common cancers such as breast, prostate and colon to be used for individualized decision making about adjuvant treatment. They add an additional piece of information to an already complex decision process involving clinical evidence such as cancer stage, eligibility for treatment, health, life style and personal preferences. Such decisions involve a delicate trade-off between the short-term health risks and decreased quality of life on the one hand and potential decrease of total years of survival on the other hand.

In this project we characterize and define adjuvant treatment decisions, a class of decision tasks that includes the adjuvant treatment examples described above, model the decision making process with a particular emphasis on deviations from normative rules of probability and propose a prototype for a support tool.

Such genomic recurrence tests are covered by some US health insurance policies and by the Canadian health care system. The recurrence scores are numerical predictions for the recurrence risk based on the expression values of a panel of genes in the patient's tissue sample. (For example, for breast cancer, the American Oncotype DX test is based on a 21-gene signature measured by RT-PCR, and the Dutch MammaPrint test is based on a 70-gene signature assessed by a custom microarray. The predictive value of these tests has been established in initial studies, and long-term clinical trials are currently under way.)

Particular characteristics of adjuvant treatment decision problems are uncertainty and ambiguity on a number of levels, including reliability of the test results, validity of the genomic recurrence score and contradictions with traditional clinical information. In addition, they are not transparent, but build on recent developments in genomics, which not only remain a black box for patients and health care professionals, but are not yet fully understood scientifically. Furthermore, attitudes towards genomics as a new technology play into the decision process. In addition, mainly for patients, strong emotions such as fear, shame or anger can interfere with their decision process. Furthermore, the patient's social support system can play a central role.

We develop a rational model using Bayesian networks for decision tasks of our adjuvant treatment decision class and develop scenarios to gauge the impact of the modifying factors described above. We put a particular emphasis on the processing of risk information, such as an investigation of how fallacies known from descriptive studies may potentially occur in our decision class. Using a prospect theory framework, we analyse the effect of different risk attitudes. Different preferences are also studied and can vary between physicians and patients. In a context of shared decision-making this can be used to detect and overcome hidden obstacles. Finally, we propose a prototype for a decision-support tool that supports patients and physicians while they go through a complex decision task, with the aim of making this complex process transparent and manageable, while avoiding unintended fallacies.

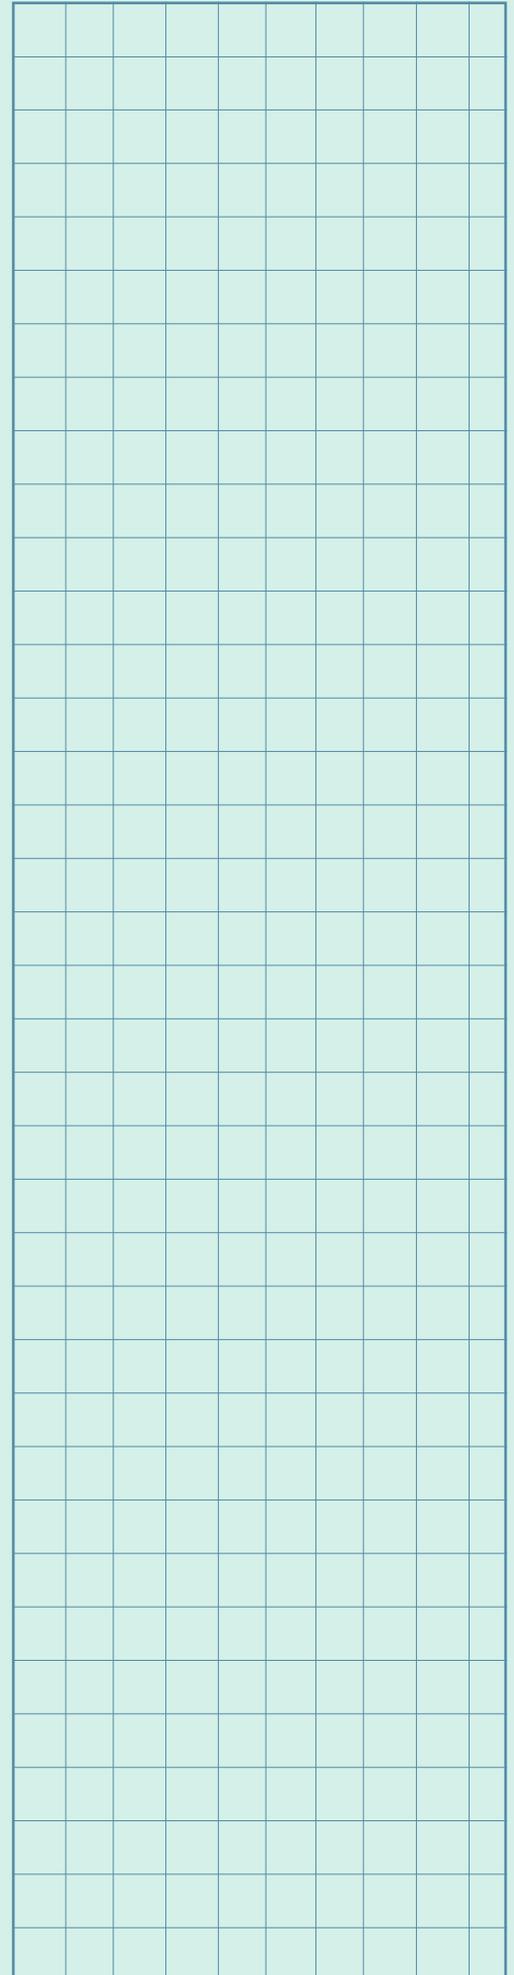
**Keywords:** cancer recurrence risk, adjuvant treatment decision, decision bias fallacies, decision support

# NICK CHATER

## VIRTUAL BARGAINING: A MICRO-FOUNDATION FOR SOCIAL DECISION MAKING AND INTERACTION

**NICK CHATER** *Professor of Behavioural Science, Warwick Business School*

Many important decisions involve interacting with other decision makers. What I should do depends on what other people will do. But this can lead to some apparently awkward problems of circularity. So person A can only decide what to do on the basis of what person B will do; but B person B can only decide to do on the basis of what person A will do! One way out of this deadlock is the notion of Nash equilibrium, as widely applied in economics. Here I explore an alternative way of breaking the deadlock, virtual bargaining, developed with Jennifer Misyak, Tigran Melkonyan and Hossam Zeitoun. Virtual bargaining proposes that decision makers implicitly compute what they would agree to do, if they were to bargain explicitly; and then carry out that agreement. I argue that virtual bargaining may underpin social behaviour and communication, and perhaps be distinctively human.



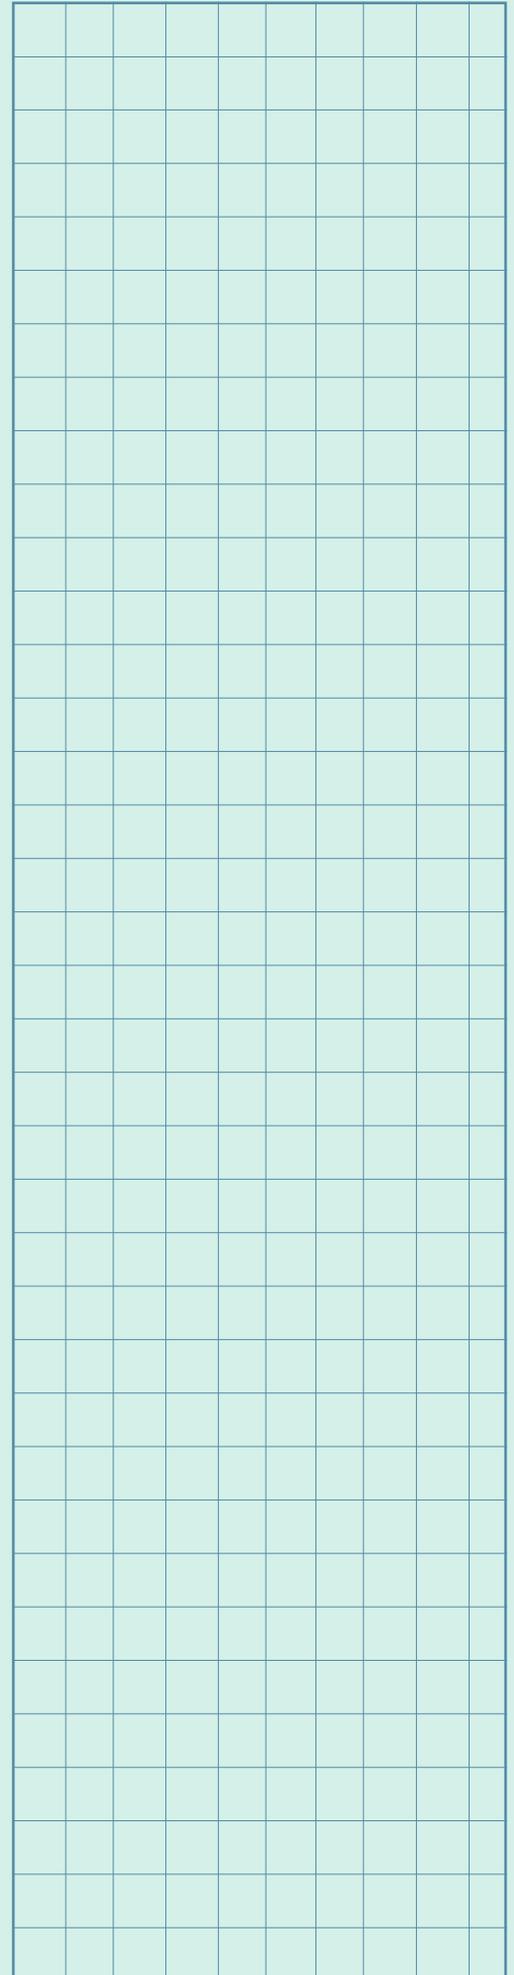


# TOM GRIFFITHS

## PRIORS AND PROCESSES IN BAYESIAN MODELS OF COGNITION

**TOM GRIFFITHS** *Associate Professor of Psychology and Cognitive Science,  
University of California, Berkeley*

Bayesian models of cognition explain human inductive leaps as probabilistic inferences, combining prior knowledge with available data. This raises two important questions: how can we identify the prior knowledge that people use, and what cognitive processes make it possible for people to act consistently with Bayesian inference despite its computational intractability? I will present possible answers to both questions, describing an experimental framework for estimating prior distributions that is more effective than traditional elicitation methods and a theoretical framework for connecting Bayesian models of cognition to psychological process models and heuristics. This theoretical framework - which we call "resource rationality" - assumes that people make effective use of their limited cognitive resources. I will argue that this approach gives us formal tools for defining and deriving heuristics, and show that classic heuristics from the decision-making literature naturally fall out of this kind of account.



## OPTIMAL TEACHING WITH LIMITED-CAPACITY DECISION MAKERS

**BRADLEY C. LOVE** *University College London*

**XIAOJIN ZHU** *University of Wisconsin*

**KAUSTUBH PATIL** *University College London*

**LUKASZ KOPEC** *University College London*

This talk explores the interconnections between capacity constraints in memory retrieval, the benefits of idealising training examples, and optimal teaching procedures. Optimal teaching procedures provide both a normative basis for idealisation and a means to evaluate competing cognitive models. The goal of this research is to understand why human decision making goes wrong and to suggest methods to improve performance.

Some decisions, such as predicting the winner of a baseball game, are challenging in part because outcomes are probabilistic. When making such decisions, one view is that humans stochastically and selectively retrieve a small set of relevant memories that provides evidence for competing options. We show that optimal performance at test is impossible when retrieving information in this fashion, no matter how extensive training is, because limited retrieval introduces noise into the decision process that cannot be overcome. One implication is that people should be more accurate in predicting future events when trained on idealised rather than on the actual distributions of items. In other words, we predict the best way to convey information to people is to present it in a distorted, idealised form. Idealisation of training distributions is predicted to reduce the harmful noise induced by immutable bottlenecks in people's memory retrieval processes. In contrast, machine learning systems that selectively weight (i.e., retrieve) all training examples at test should not benefit from idealisation. These conjectures are strongly supported by several studies and supporting analyses. Unlike machine systems, people's test performance on a target distribution is higher when they are trained on an idealised version of the distribution rather than on the actual target distribution. Optimal machine classifiers modified to selectively and stochastically sample from memory match the pattern of human performance.

Optimal teaching procedures provide a means to test and extend this basic theory of capacity constraints and the benefits of idealisation. Given a test environment and agent, the optimal teacher determines the set of training examples that should maximise future performance. Critically, the optimal set of training examples depends on the model of the agent. We find that limited-capacity agents yield sets of training examples that are idealised, which provides a normative basis for idealisation manipulations given the conclusion that humans are capacity limited. However, these training sets have some characteristics that diverge in informative ways from our initial intuitions about how to best idealise training examples to benefit human decision makers. Finally, optimal teaching provides a novel method to test competing cognitive models -- models should be favoured to the extent that they yield optimal training sets that maximise human performance. For example, we predict that optimal training sets derived from limited-capacity models should yield better human performance than those derived from unlimited-capacity models. Overall, the reported results suggest firm limits on human rationality and have broad implications for how to train humans tasked with important classification decisions, such as radiologists (one test domain is mammography), baggage screeners, intelligence analysts, and gamblers (one test domain is sport prediction).

**Keywords:** Optimal Teaching, Memory Retrieval, Model Selection, Idealisation

## THRESHOLD-BASED INFERENCE

**SHENGHUA LUAN** *Max Planck Institute for Human Development*

**LAEL J. SCHOOLER** *Max Planck Institute for Human Development*

**GERD GIGERENZER** *Max Planck Institute for Human Development*

In a lexicographic semiorders model for preference, a decision-maker is assumed to search cues in a subjective order and choose an alternative if its value on a cue exceeds those of other alternatives by a threshold  $\Delta$ . We generalized this model from preference to inference and refer to it as  $\Delta$ -inference. Unlike with preference, which is a matter of taste and for which accuracy is difficult to define, the problem a mind faces when making an inference is to select a  $\Delta$  that can lead to accurate judgments.

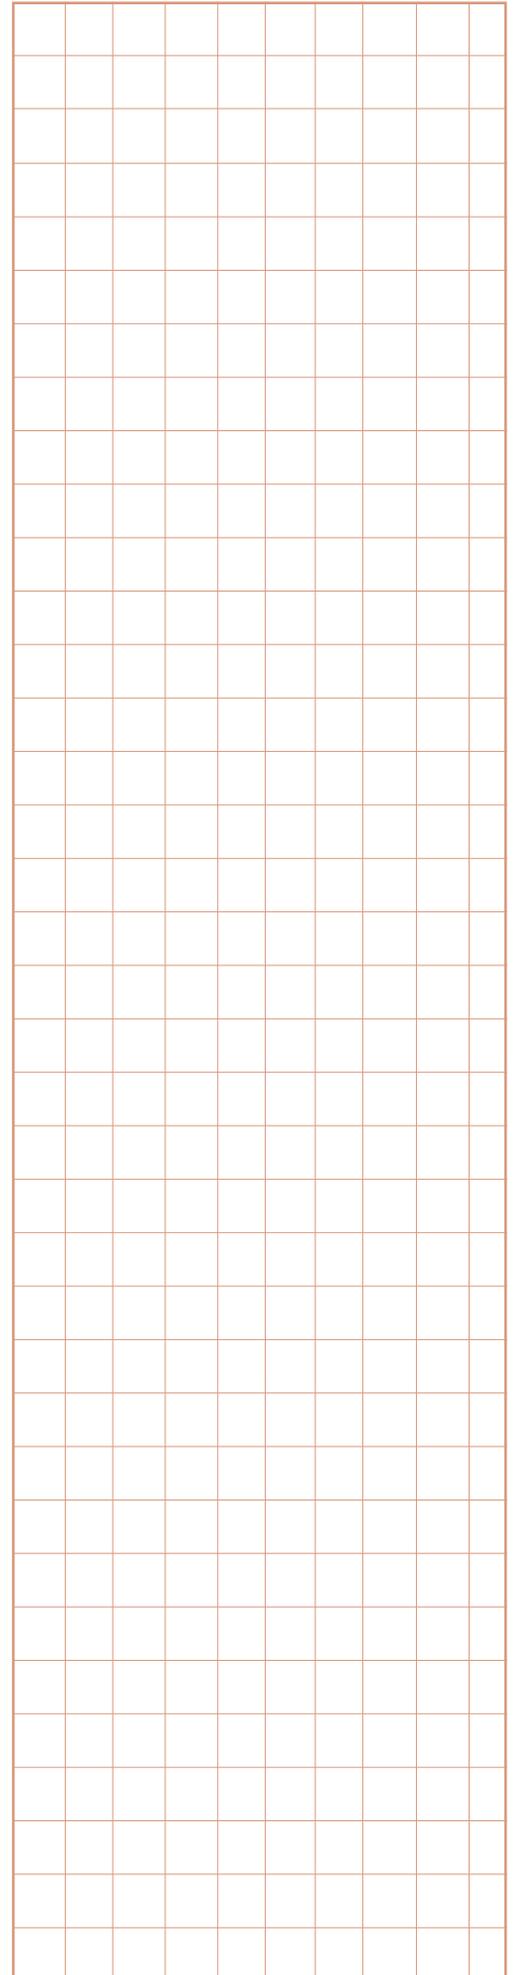
To find a solution to this problem, we applied Clyde Coombs's theory of single-peaked preference functions. We show that the accuracy of  $\Delta$ -inference can be understood as an approach-avoidance conflict between the decreasing usefulness of the first cue and the increasing usefulness of subsequent cues as  $\Delta$  becomes larger, resulting in a single-peaked function between accuracy and  $\Delta$ . We refer to the peak of this function as  $\text{Peak}\Delta$  and found it varied with the properties of the task environment: the more redundant the cues and the larger the differences in their information quality, the smaller the  $\text{Peak}\Delta$ .

Learning the  $\text{Peak}\Delta$  in each environment can be costly. An alternative way is to find a fixed  $\Delta$  that works the best across a variety of environments and apply that  $\Delta$  in a new environment. To gain an understanding of what that  $\Delta$  might be, we started an investigation that involved 746 simulated three-cue environments. In each environment, we examined both the fitting and prediction accuracy of  $\Delta$ -inference under a series of fixed  $\Delta$  values in three sample-size conditions ( $N=20$ , 100, and 2,000). We found that for prediction and averaged across all simulated environments, the best fixed  $\Delta$  is not large (i.e., 0.5 z-score) and it is the same for each sample-size condition.

To check the generalizability of the finding, we collected 39 real-world environments and replicated the analysis there. To our surprise, the best fixed  $\Delta$  turned out to be zero across those environments! Using  $\Delta$ -inference with  $\Delta=0$  means that a decision is made as soon as there is any difference between the alternatives on the first cue. In terms of prediction accuracy, this extremely simple model not only outperformed  $\Delta$ -inference with any other fixed  $\Delta$  values, but also on average performed as well as  $\Delta$ -inference with  $\text{Peak}\Delta$  for each environment. We further compared  $\Delta$ -inference with  $\Delta=0$  with multiple linear regression (ordinary and Bayesian) and the general monotone model (Dougherty & Thomas, 2012), and found that it performed either better than or at similar levels as the other models when the training set was less than 80% of the entire sample.

Overall, our study demonstrates the potential of integrating and extending existing concepts, models, and theories from preference to improve our understanding of how the mind makes inferences. It also adds to the growing literature that shows how and why the mind can make fast as well as accurate decisions using simple models that do not aim to optimize.

**Keywords:** Threshold,  $\Delta$ -inference, lexicographic semiorders, single-peaked function, approach-avoidance conflict, ecological rationality



## DECISION MAKING AND RESPONSE IMPLEMENTATION: THE MISSING LINK TO CONSTRAIN MODELS

CHRISTOPHE TANDONNET *Université de Genève, Genève, Switzerland*

Sequential-sampling models of decision-making differ according to two major distinctions: (i) how the evidence in favour of the different alternatives is accumulated, and (ii) what reflects the accumulation. First, accumulation of evidence may be independent for the different alternatives or interact, possibly through inhibition of the non-selected alternative. Second, when considering decisions between actions based on sensory information, accumulation of evidence may reflect either a purely perceptual decision that leads thereafter to an action or a decision involving the selection of an action. In the former case, the perceptual code has to be transmitted into a motor code with either a sequential or continuous transmission. However, a fully continuous transmission can be viewed as functionally equivalent as considering that decision involves action selection. Thus, at least for this type of models, response implementation directly reflects the accumulation of evidence, and hence can be used to assess the principles of decision making.

In previous studies, the state of the motor system was assessed during the reaction time of a visual choice task involving manual responses. Event-related potential data are consistent with an inhibition of the non-selected response: response-related activity showed a positive component symmetric to the negative component involved in the motor command of the response that develops over the sensorimotor cortices. Transcranial magnetic stimulation data showed that the motor activity related to the non-selected response is reduced after an initial facilitation for both responses. Taken together, these data are consistent with an initial activation of both responses followed by an inhibition of the non-selected response.

Such a pattern of response implementation provide constraints to the models. Indeed, these data are consistent without further assumptions only with the race-like models (i) either assuming continuous transmission from decision to motor implementation or assuming that decision involves action selection, and (ii) assuming inhibition between the different alternatives. In contrast, the data are not compatible with the race-like models involving independent accumulation. The data constrain also the dual-process models assuming sequential transmission from decision to motor implementation. Although, by definition, response implementation cannot constrain the decision in this case, at least two further assumptions are required to account for the activation and inhibition pattern: a non-specific activation independent of the decision process, and a specific feedforward motor inhibition between decision and motor implementation.

In conclusion, measures of response implementation can provide constraints to the models. One may speculate that this putative selective inhibition reflects an intrinsic principle of decision making. However, this selective inhibition may not be generalized to the oculomotor system, although other form of inhibition like non-selective lateral inhibition may play a key role in generation of saccadic eye movements. Future research assessing whether this selective inhibition can be evidenced in decisions involving ocular responses would provide more constraints to dissociate between race-like models with inhibition, assuming either that decision involves action selection or assuming continuous transmission from decision to motor implementation, and models assuming sequential transmission.

**Keywords:** sequential-sampling models; vision; action; neural inhibition; sensorimotor; oculomotor

## LEARNING TO SOLVE WORKING MEMORY TASKS

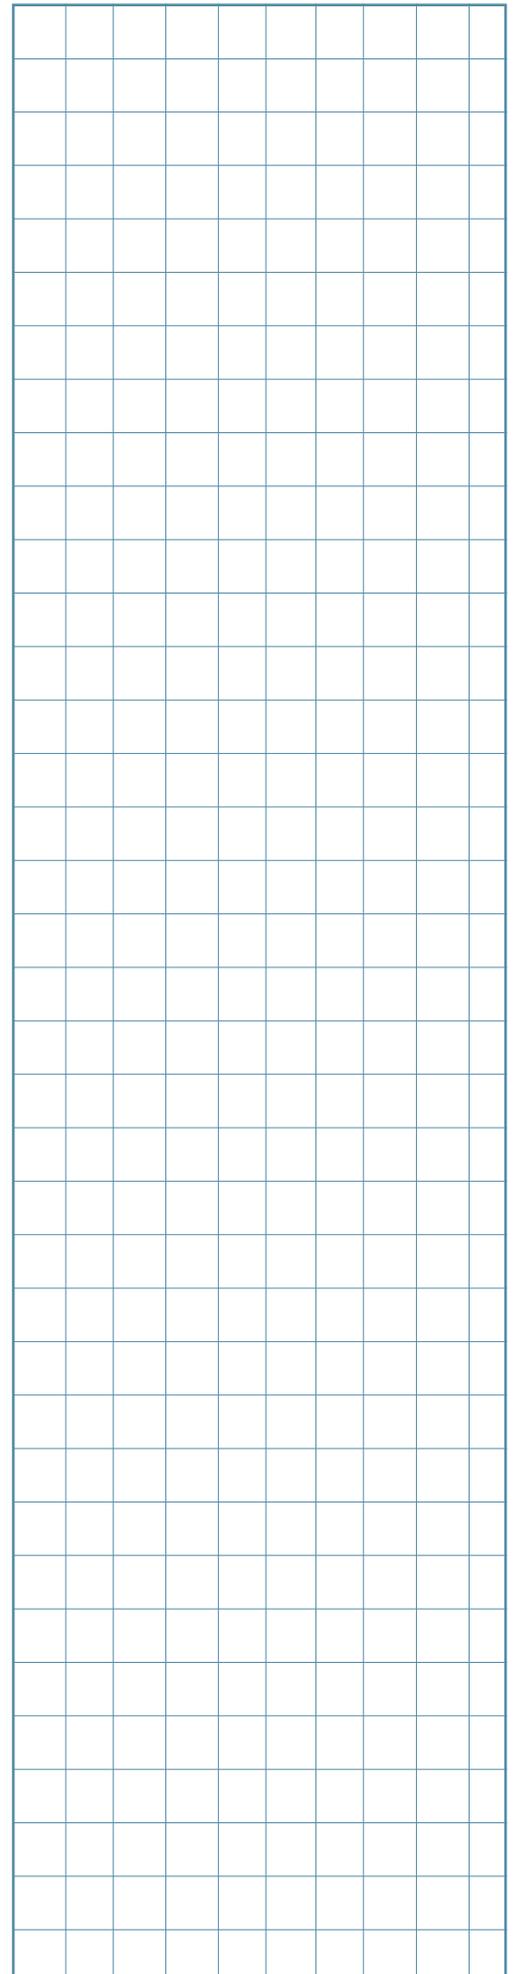
**ADNANE EZ-ZIZI** *University of Bristol*  
**DAVID LESLIE** *University of Bristol*  
**SIMON FARRELL** *University of Bristol*

Working memory (WM) tasks can be considered as dynamic decision making tasks, where the agent needs to decide which action to take for each new observation. The decision maker not only needs to use the current observation, but also should take into account a sequence of past observations. The environment created by these types of tasks can be seen as partially observable Markov decision process (POMDP). This means that it does have the Markov property—rewards depend only on the current state—but assumes that the underlying state of the environment cannot be fully observed.

One approach to solve POMDPs is to support the learner with a memory device to store past states and/or actions in order to disambiguate the underlying state (McCallum, 1995). This method fits nicely with WM tasks, as the memory device can simply be the agent’s WM. Based on this method, several reinforcement learning models (RL) of WM has been proposed and shown to solve the learning problem of many WM tasks such as the n-back and the 12-AX task (O’Reilly & Frank, 2006; Todd et al., 2009).

However, learning in these models is very slow, especially when the task requires more than one WM state, raising questions about whether WM could reasonably be learned purely from experience. Using simulations with different WM capacities, on a simplified version of the 12-AX task, I will show that there exists a trade-off between learning speed and learning performance, which might explain why our WM capacity is limited. I will also propose an approach for how an organism could deal with this trade-off. Finally, I will propose several ways to increase the performance and learning speed of these models. Hence, we can start testing and comparing learning performance given by these models with human data.

**Keywords:** Working memory tasks; working memory capacity; Reinforcement learning



## THE VALUE OF ECONOMIC IRRATIONALITY: WHY AND WHEN BEING IRRATIONAL IS ADVANTAGEOUS

KONSTANTINOS TSETOS *University of Oxford*

Years of research within the cognitive and decision sciences have converged to the idea that humans are systematically at odds with rational choice theory. What is the cost of being irrational? Rationality is a necessary condition for optimality, the ability of an agent to choose correctly the course of action that maximizes the expected reward. Although irrational agents fail by definition to achieve the ideal observer benchmark, this exact performance difference cannot be experimentally measured due to the subjective nature of "correctness" in preferential decisions. To circumvent this problem and inspired by research in perception and visual psychophysics, I will present a novel paradigm (termed "value psychophysics") that abstracts value-based decisions into a simple information-integration task in which correctness is objectively defined. Apart from offering precise control on the information preceding each decision, this task permits the orthogonal assessment of choice consistency (rationality) and choice accuracy (optimality). I will show that human behaviour violated the rational principles of transitivity and regularity but, surprisingly, the degree of these violations was not negatively correlated with choice accuracy.

From a theoretical perspective, choice accuracy is limited by two factors a) reliance on non-normative choice mechanisms and b) uncorrelated noise that distorts processing, offering an upper boundary on choice accuracy. Why did the more irrational participants did not also have lower accuracy? This empirical result suggests that noise is a much stronger predictor of poor performance compared to reliance on irrational decision mechanisms. To theorise this, I will describe a new sequential sampling model that explains rationality violations in decisions under uncertainty using a single selective mechanism, which prioritises the integration of the momentarily most valuable samples of incoming information. Surprisingly, under this computational scheme, a higher tendency towards irrationality (realised as stronger selectivity in the model) increases robustness to uncorrelated perturbations of the variable that drives the decision (e.g. integrated evidence) or in other words to "late" noise (corresponding for example to cortical noise). As a result, for moderate levels of late noise the irrational selective integration model outperforms the choice accuracy of the statistically optimal sequential probability ratio test (SPRT). I will present data from new experiments that aim to test the hypothesis that stronger reliance on irrational mechanisms helps (and does not deter) choice performance in the face of increasing late noise. This interpretation of non-normative behaviour downplays the severity of acting irrationally and offers a new explanatory framework of why humans evolved irrational mechanisms.

**Keywords:**rationality, optimality, noise

## RATE DOMAIN DISTRIBUTIONS IN SIMPLE REACTION TIME TASK

**VALERIO BISCIONE** Plymouth University  
**CHRISTOPHER HARRIS** Plymouth University

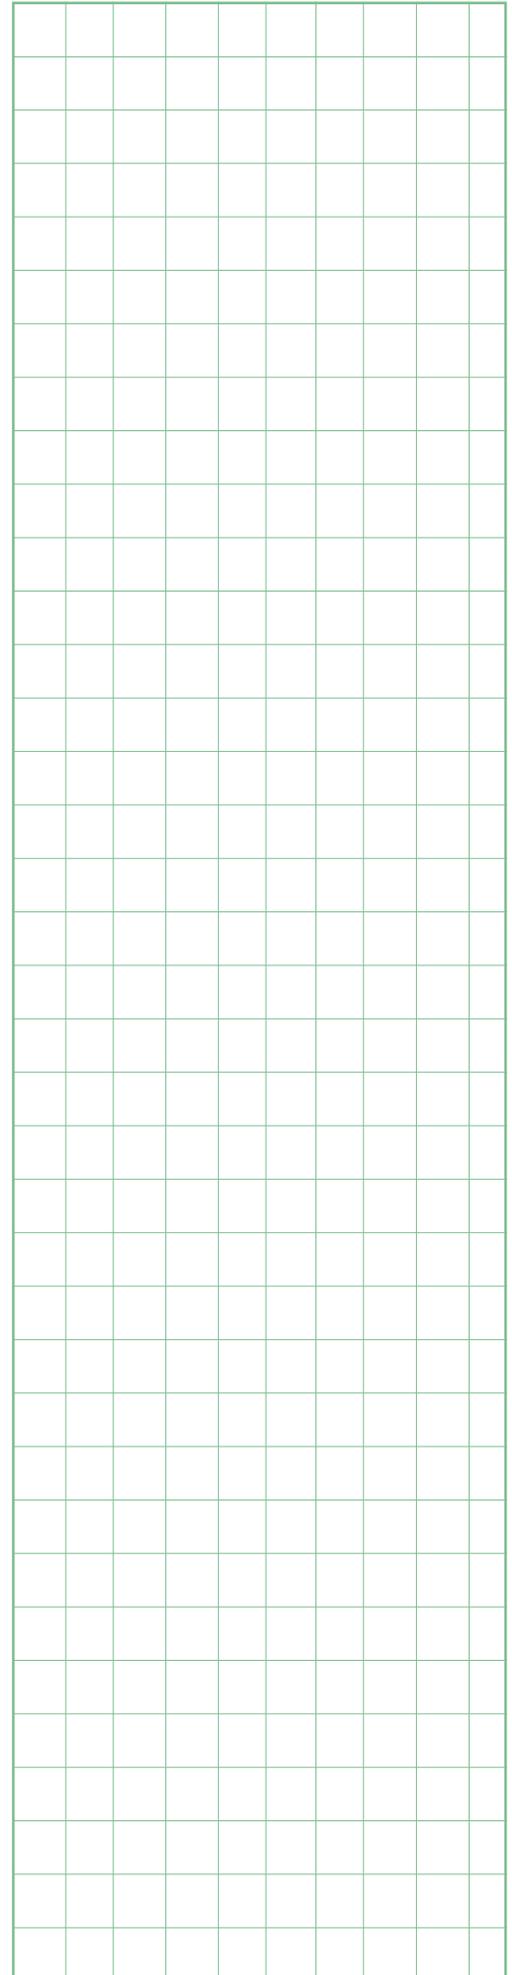
*Background.* It is widely accepted that human reaction times (RT) can reveal important information about decision strategies for “all-or-none” type of responses. During the past decades several models have been proposed to describe and predict mean RT and RT distributions and, amongst them, sequential sampling models are the most successful in accounting for data for simple and two-choice tasks. RT distributions are usually positively skewed which are often modelled as an Inverse Gaussian (the first passage time of a stochastic Wiener process) or an Ex-Gaussian. However it has been found that latency distribution of saccades are very close to the Reciprocal Normal, meaning that they are normally distributed in the rate domain. This represents a challenge for most stochastic rise-to-threshold models, since these models can only achieve normality in the rate domain with implausible parameters. We have already observed, in a previous study, the normality of the distribution in the rate domain. In that experiment we used a two-forced choice paradigm with an easy/difficult condition and an accurate/urgent instruction sets. We now investigate the possibility to find similar results with a simple RT paradigm, exploring the relation between the foreperiod (FP) time and the Piéron’s Law in the rate domain.

*Methods.* In this experiment (12 subjects, 3 blocks of 250 trials) we varied the FP time and the luminance of the stimuli. The participants were asked to press a button as soon as they saw the stimulus. We used 3 FP conditions (0.6, 1, and 2.4 seconds) and 5 luminance levels (0.42, 0.71, 1.21, 2.06, 3.50 cd/m<sup>2</sup>).

*Results.* As expected, the relationship between RT and luminance followed Piéron’s Law, and the mean RT increased with mean FP. The rate distributions approached normality with the increasing of the FP, which may indicate that with shorter FP a second process was contaminating the data. We fitted different distributions to the data and we found that the Reciprocal Normal provided a good fit. We investigated how the two independent parameters of the Normal Distribution varied across the FP/luminance conditions.

*Discussion.* We discuss several problems of the rise-to-threshold models and how a Reciprocal Normal Distribution is not consistent with some stochastic rise to threshold models. We propose a simple optimality model in which reward is maximized to yield to an optimal rate and therefore an optimal time to respond, exploring the connection between this model and the speed-accuracy trade-off. We also show how Piéron’s Law naturally arises from this model. Our key claim is that the main goal of the human decision process in simple decision tasks is to maximize the rate of reward.

**Keywords:** Reaction times, latency, Reciprocal Normal, Rate Domain, Piéron’s Law, Foreperiod



## A SIMPLE METHOD TO CHARACTERIZE THE EXPLORATION-EXPLOITATION TRADE-OFF IN BINARY DECISION MAKING

MARCUS VINÍCIUS CHRYSÓSTOMO BALDO *University of São Paulo, Brazil*

CAROLINA FEHER DA SILVA *University of São Paulo, Brazil*

CAMILA GOMES VICTORINO *University of Surrey, UK*

EDUARDO SANGIORGIO DOBAY *University of São Paulo, Brazil*

NESTOR CATICHA *University of São Paulo, Brazil*

In repeated binary choice experiments, when choosing between two possible alternatives that differ in their probability of occurrence, agents belonging to several different species, including humans, usually match their choice frequencies to the corresponding outcome probabilities, a behaviour known as probability matching. This strategy is suboptimal in comparison to always choosing the outcome with the higher probability (known as maximizing). We have previously shown that probability matching may be a likely by-product of adaptive cognitive strategies that are crucial in sequence learning, but may lead to suboptimal performances in less structured environments (Feher-da-Silva & Baldo, PLoS ONE 7(5): e34371). Here, we report on a simple but powerful analysis based on the computing of the conditional probabilities arising from the coupling between predicted and actual binary sequences. By means of this method we were able to show that, despite the common observed matching between the alternatives' probabilities and their respective choice frequencies, the decisional behaviour of human agents clearly deviates from what would be expected by their simply establishing probabilities of choice that match the underlying probability structure of the two alternatives (as learned from a previous sample of outcomes). In this sense, probability matching is a label that should better describe the empirical fact of roughly coincident frequencies, rather than an underlying decisional behaviour. Moreover, by applying the same analytical method to human participants belonging to different age groups (from 3.9 to 71.2-year-old volunteers), we were able to characterize the exploration-exploitation trade-off across a wide segment of the human life span. Our results suggest that ageing tends to shift the decisional behaviour from a predominantly explorative procedure, as observed in children, towards a more exploitative strategy, which was adopted by the elderly volunteers.

**Keywords:** decision making, probability matching, exploration-exploitation trade-off

## A COMPETITIVE TEST OF SATISFICING HEURISTICS IN CHOICE FROM SERIALY DEPENDENT SEQUENCES

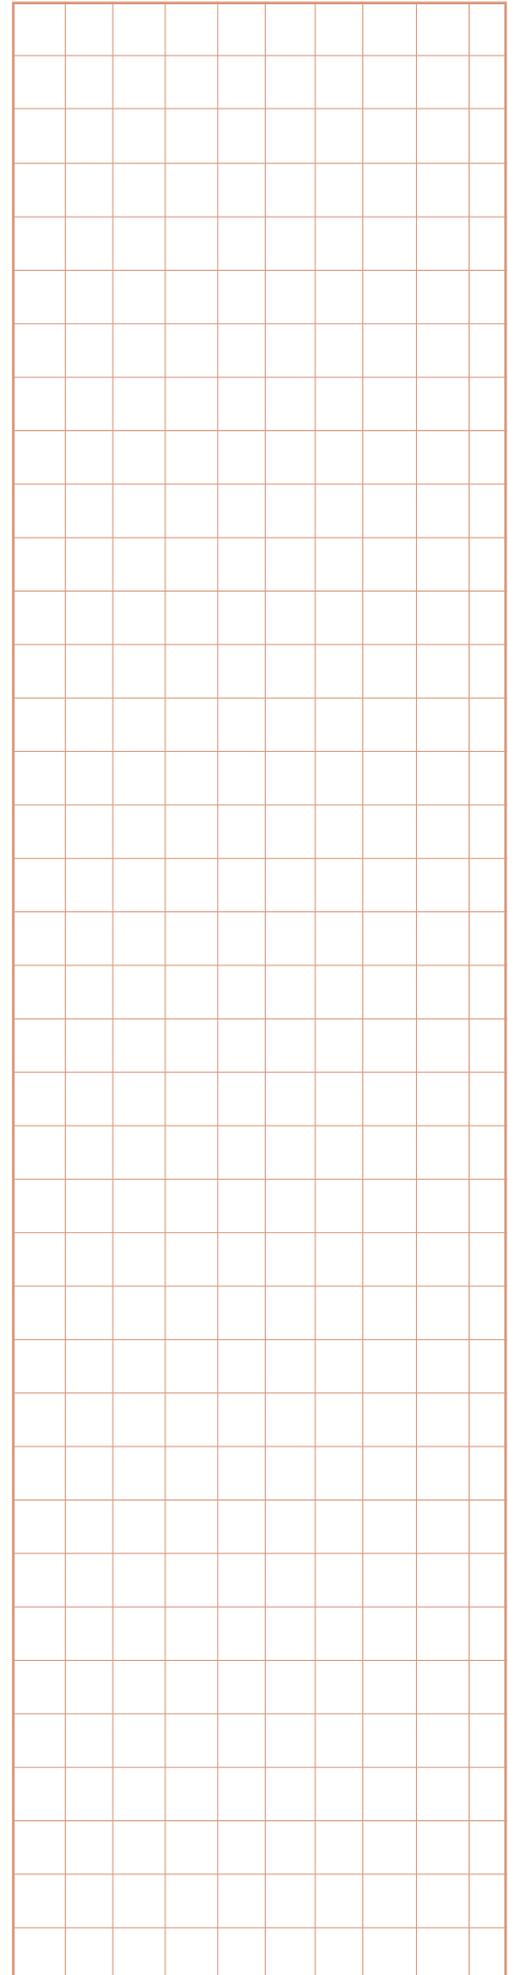
**PERKE JACOBS** *Max Planck Institute for Human Development*

There is an abundance of situations requiring decision makers to choose from sequences of options rather than option arrays. This study examines the performance of three classes of satisficing heuristics identified earlier by Seale and Rapoport (1997). Structurally, all three classes of heuristics are instances of the more general class of satisficing heuristics (Simon, 1955): Using different criteria, they formulate aspiration levels based on the first realizations of each sequence and then select the first option meeting the respective aspiration level. Seale and Rapoport (1997) demonstrated that experimental subjects' choices were consistent with all three classes of heuristics.

Later, Stein, Seale and Rapoport (2003) evaluated the decision quality of the three classes of heuristics in the context of the secretary problem, i.e., a fully random sequence with no option to recall previously encountered options. Given the fact that real-life sequences may not be fully random, the present study examines their performance A) when sequences are serially dependent; and B) when search costs are considered. Specifically, sequences were modeled as a multinomial process with fixed probabilities that the sequence does not change, changes upwards, or changes downwards by a fixed amount. The parametrization of the sequence was based on price data from a German online market for used automobiles (Artinger, 2012). This simulation study followed a full factorial design in which probabilities of price increase and decrease, the magnitude of price changes, and search costs were varied. Gilbert and Mosteller (1966) provided an optimal solution for secretary-problem-type of situations, which was used as a performance benchmark by Stein and colleagues. In contrast, the present study employed a competitive testing strategy: Performance of the heuristics was assessed relative to the performance of a Bayesian decision algorithm that is aware of the statistical structure of the sequence and inferred parameter values from observed realizations of the sequence.

The results indicated that the single best heuristic from each class differed depending on the trend of the sequence, with improving sequences requiring longer and deteriorating sequences requiring shorter sampling periods. Given the best performing heuristic for each class, performance was compared to that of the Bayesian algorithm across several conditions. The Bayesian algorithm offered the best performance across several conditions — perhaps not surprisingly, given its relative sophistication. Interestingly, all classes of satisficing heuristics offered considerably better performance than the Bayesian algorithm when 1) sequences exhibited an improving trend; and 2) search costs were zero so that performance differences stemmed only from finding the best option. Decision environments with these characteristics therefore allow decision makers to employ a simple, psychologically plausible strategy whose performance exceeds that of a computationally more intensive Bayesian algorithm. These results re-affirm Simon's (1956) conjecture that heuristics can be effective decision strategies when adapted to the specific decision environment.

**Keywords:** Satisficing, Heuristics, Sequential Choice, Competitive Testing



## EVOLUTIONARY AND COMPUTATIONAL CONSIDERATIONS ON THE DRIFT DIFFUSION MODEL

**ANGELO PIRRONE** *University of Sheffield*

**JAMES A.R. MARSHALL** *University of Sheffield*

**TOM STAFFORD** *University of Sheffield*

The DDM has been shown to be quantitatively accurate in describing binary choices and value-based choices. According to the DDM the decision-maker integrates the difference in evidence supporting two alternatives. In laboratory settings subjects are usually rewarded only on making a correct choice, so optimisation of a zero-one loss function is appropriate, and this is achieved by implementing a statistically-optimal decision procedure that gives the best compromise between speed and accuracy of decision-making; this trade-off can be implemented and optimised by the DDM. However many decisions, such as selecting food items of potentially different value, appear to be different since the animal is rewarded by the value of the item it chooses regardless of whether it was the best. We argue that most naturalistic decisions, which animals' brains should have evolved to optimise, are value-based rather than accuracy-based. While value-based decisions can be optimised using mechanisms for managing speed-accuracy trade-offs, this requires additional information on the decision problem at hand in order to allow optimal parameterisation since subjects should learn the values of correct and incorrect choices over time on a case by case basis. Recent theory by Pais et al. (2013) has presented mechanisms that can manage value-sensitive decision problems adaptively without direct parameterisation regarding each alternatives' values. In Pais et al.'s model when equal-but-low-value alternatives are presented, a decision deadlock is maintained that can be broken should a third, higher-value alternative, be made available while when equal-but-high-value alternatives are presented, or sufficient time passes, deadlock is spontaneously and randomly broken; when differences between alternative values are sufficiently large, the value-sensitive mechanism becomes closer to a classic DDM, allowing speed-accuracy trade-offs to be managed. In contrast, the DDM is insensitive to the absolute magnitude of evidence for alternatives, with its behaviour determined solely by their difference and by sensory noise. To discriminate these two decision-making models we conducted two studies, one a perceptual-judgement accuracy-based decision and one a value-based decision. The perceptual experiment involved a numerosity discrimination paradigm, with subjects instructed that only correct identifying the largest dot-cluster would be rewarded; in the value-based setting we adjusted the paradigm by instructing participants that they would receive a monetary reward in direct proportion to the sum of the dots in the clusters they selected over all trials. The main results of the two studies are (i) in line with the DDM, reaction times differ according to the difference in evidence of the two alternatives, (ii) in contrast with the DDM and in line with Pais' model predictions, reaction times differ according to the overall magnitude of the two alternatives for each level of 'difference in evidence'. Our theoretical considerations and the initial results from the study, that cannot be fitted by the DDM, raise issues with the idea that the DDM can be a comprehensive computational framework for decision making and suggest that human decision making is indeed value sensitive, even in simple perceptual tasks.

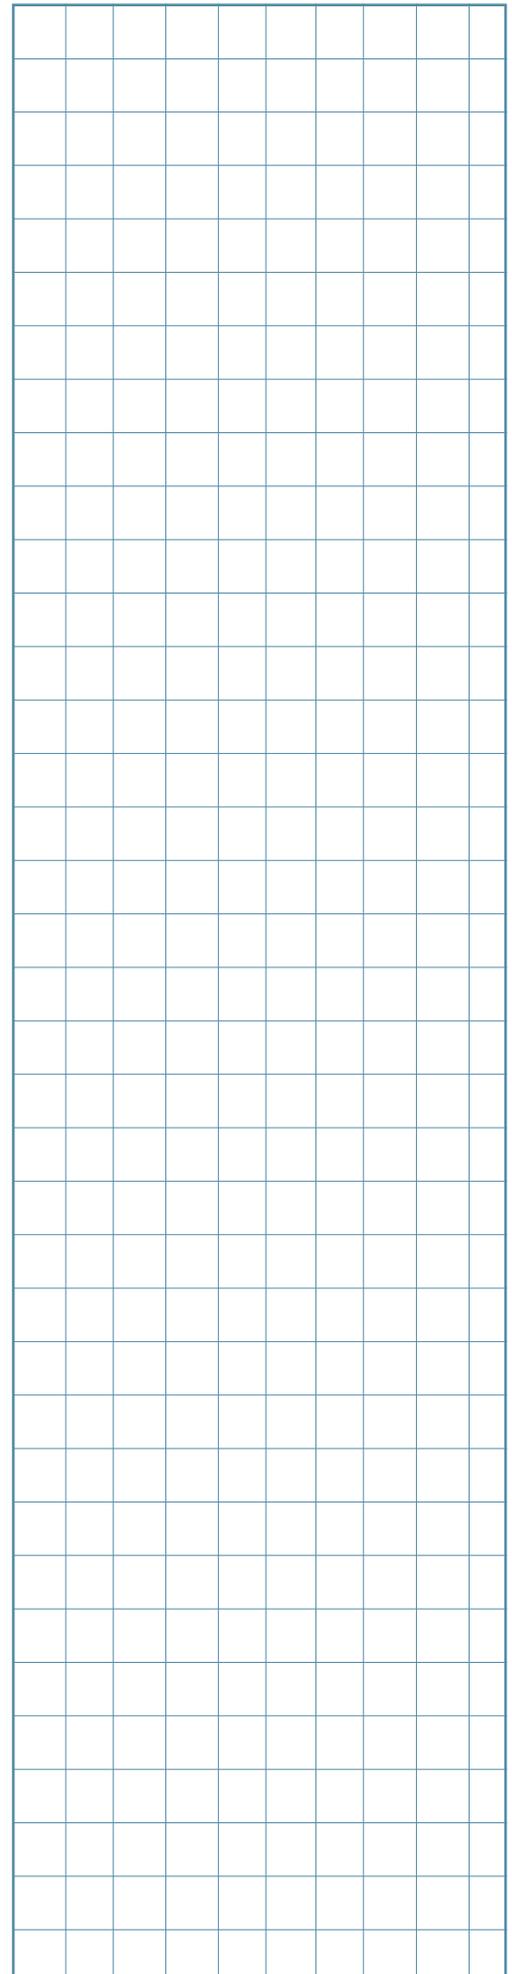
**Keywords:** decision-making, value, reward, drift-diffusion, mechanism, evolution, optimality

## DOPAMINE AFFECTS ENCODING AND RETRIEVAL OF POSITIVE AND NEGATIVE REINFORCEMENT LEARNING

**JOHN GROGAN** *University of Bristol*  
**RAFAL BOGACZ** *University of Oxford*  
**ELIZABETH COULTHARD** *University of Bristol; North Bristol Trust NHS*

Dopamine has been found to improve learning from positive reinforcement at the cost of negative reinforcement (e.g. Frank et al., 2004), but has also been suggested to work mainly through effects on testing and retrieval rather than learning (e.g. Shiner et al., 2012). We aimed to establish whether the dopamine levels affect encoding or retrieval of stimulus-action associations learnt from positive (or negative) feedback. Our paradigm included a 24 hour delay between learning and testing, that allowed Parkinson's Disease (PD) patients to be tested on or off their dopaminergic medication on both days (giving 4 conditions: on-on, on-off, off-on, off-off). This within subjects design was tested on 15 PD patients and 13 healthy age-matched controls. On day 1 participants were given learning trials of the probabilistic selection task in which the feedback generated after selecting one of two cards is probabilistic. For example, card A receives positive feedback on 80% trials while card B receives it on 20% (65% & 35% for C & D, respectively). After learning, memory for these cards is tested (without feedback given) both immediately and after a 30 minute delay. On day 2 participants were given another memory test of the cards, and then shown the new combinations they hadn't seen previously (AC, AD, BC, BD). The amount of times they chose A (the most rewarded card) or avoided B (the most punished card) reflected learning from positive and negative reinforcement, respectively. We also fit a variety of reinforcement learning models to the patient data, both to see which model best explained the data and to see any effects of dopamine on the learning rates of the patients. We found a significant interaction of day 1 and day 2 medication states on the amount of choose-A behaviour. When patients were off medication on day 1, day 2 medication increased choose-A and decreased choose-B as expected (off-on), relative to day 2 off medication (off-off). However, when patients were on medication on day 1, day 2 medication had no effect. Intriguingly, patients on medication for both days (on-on) showed the highest amount of avoid-B and lowest amount of choose-A behaviour of any of the conditions, suggesting that the effects of dopamine are very different when tested 24 hours later compared to immediately or only 1 hour after learning. The modelling revealed that the best fitting model (as determined by the BIC) was a dual learning rate SARSA model, but there were no significant effects of medication on the learning rates. These findings suggest that dopamine has a different pattern of effects when tested 24 hours after learning, and that dopamine state during learning determines whether later dopamine will have an effect.

**Keywords:** dopamine, reinforcement learning, parkinson's disease



## HEURISTICS AS A SPECIAL CASE OF BAYESIAN INFERENCE

**PAULA PARPART** *University College London*

**MATT JONES** *University College London*

**BRADLEY C. LOVE** *University College London*

In the judgement and decision making literature, rational inference models (e.g. Bayesian models) have usually been depicted as opposing simple heuristics (e.g., take-the-best or tallying). The current research demonstrates a novel finding wherein prominent decision heuristics are a special case of Bayesian inference.

We developed two Bayesian learning models based on two well-known regularization techniques in machine learning, namely lasso regression and ridge regression. The penalty terms of these Bayesian models each incorporate a prior that reflects the necessary environmental structures for tallying or take-the-best to succeed. For example, tallying performs best in an environment where cues are equally weighted, which is sometimes referred to as a compensatory weighting structure. In contrast, take-the-best thrives when the most heavily weighted cue outweighs all other cues such that no combination of theirs can compensate for the strongest weight, often referred to as a noncompensatory structure. We demonstrate that these Bayesian inference models become equivalent to the heuristics when the priors become very extreme (with an infinitely large penalty parameter), thereby making them a special case of the model, next to standard linear regression.

In a re-analysis of popular heuristic datasets which span a wide range of domains we show that our modified Bayesian ridge regression outperforms both tallying and standard linear regression. Similarly, our Bayesian extension of lasso regression could outperform both take-the-best and linear regression. A large simulation study furthermore illustrates the practical convergence of the Bayesian learning models and the heuristics with increasingly extreme priors: convergence is reached when the original cue weights have developed into perfectly compensatory (tallying) or noncompensatory (take-the-best) weighting structures. Results depict that the Bayesian learning models declare both heuristics and standard linear regression as a special case of the Bayesian inference model. This implies that heuristics can be adaptive to certain environments, while being continuously contained within a rational inference model. Thereby we create a formal relationship between two traditionally opposing theories of decision making.

In addition, these findings lead us to think that sometimes the appropriate psychological process and its matching environmental structure may lie somewhere in between a frugal heuristic and more complex, integrative regression approach. These new developments have far-reaching implications with respect to the judgement and decision making literature on heuristics and rational inference models, as well as other fields that make use of novel regularization algorithms.

**Keywords:** heuristics, Bayesian inference, take-the-best, regularized regressions, ridge regression

James Watson Room

## THE DYNAMICS OF EVIDENCE ACCUMULATION AND CHOICE

**TAKAO NOGUCHI** *University of Warwick*  
**BRADLEY C. LOVE** *University College London*

Sequential sampling models provide a way to understand both speed and accuracy of judgement. In these models, an individual is assumed to accumulate evidence supporting a judgement, and when the accumulated evidence reaches a response criterion, the individual makes a judgement. For example, in exemplar-based random-walk (EBRW) model (Nosofsky & Palmeri, 1997), items are retrieved from memory sequentially and their category membership provides evidence for the respective response. The probability of retrieving an item depends on its similarity to the stimulus.

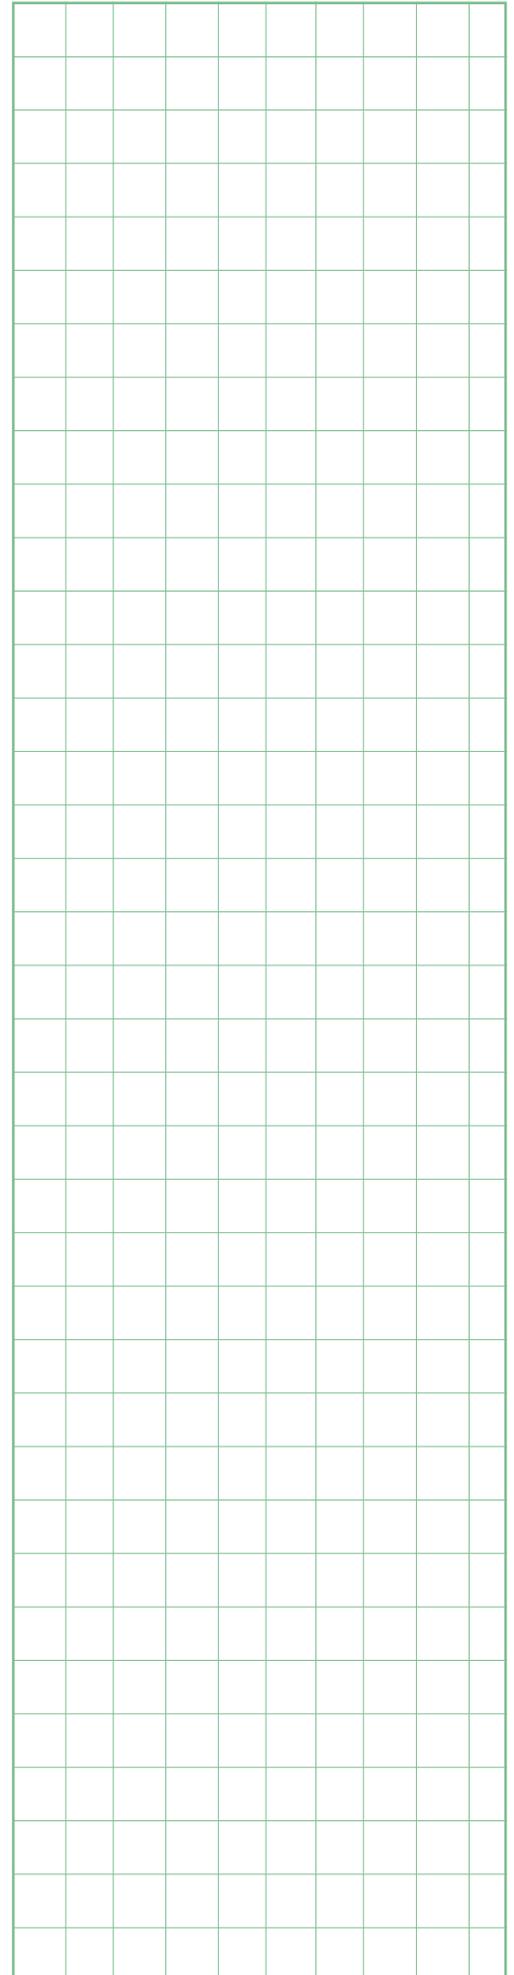
Accumulation rate in sequential sampling models is typically stationary over time. In EBRW model in particular, the probability of recalling a particular item is fixed until an accumulation reaches a response criterion. As an individual learns more about structure of a stimulus, however, the way this individual processes information may be altered over time. Previous studies report that especially when under time pressure, an individual is likely to evaluate an alternative as attending attribute dimensions (e.g., Lamberts, 1995). These findings indicate that accumulation rate may be based on a subset of dimensions at first, and gradually, more dimensions become incorporated. In other words, the similarity of a stimulus and items in memory may change within a trial.

To implement the process where a dimension is evaluated as attended, we propose a dynamic construction of similarity perception. Here, an individual recalls an item similar to the stimulus at hand based only on the dimensions the individual has attended. On the first fixation, similarity is a function of only the one dimension, but with time, similarity becomes based on more dimensions. This dynamic construction predicts that when an individual first attends to a stimulus dimension that provides misleading information about category membership, this individual accumulates evidence for an incorrect judgement at first. As a result, 1) the individual should be slower to make a correct judgement, as the individual has to override initially accumulated evidence, and also 2) the individual should be less likely to make a correct judgement.

These predictions concerning the order in which dimensions are sampled were empirically tested with eye-movement recordings. In the experiments, participant learned to classify an amoeba-like organism with three dimensions. In classifying these stimuli, all three dimensions are relevant but are imperfect predictors (Type IV structure; Shepard, Hovland, & Jenkins, 1961). Even when a stimulus belongs to Category A, for example, its shape at the top may be more readily shared with the stimuli in Category B. When such unrepresentative dimension is fixated at first, compared to when the same dimension is fixated later, participant takes longer to correctly classify the stimulus, and is less likely to make a correct classification judgement.

Thus, our results support the dynamic construction of similarity perception. We computationally formulate this construction and propose an extension to the EBRW model. This proposed model is compared against existing models in their predictive accuracy. We conclude with a discussion on implementation of dynamic processes to other sequential sampling models.

**Keywords:** Sequential sampling; process tracing; exemplar



**VISUAL ATTENTION IN CHOICE:  
DO INDIVIDUALS PAY MORE ATTENTION TO  
MORE IMPORTANT INFORMATION?**

**TIMOTHY MULLETT** *University of Warwick*  
**RICHARD TUNNEY** *University of Nottingham*  
**NEIL STEWART** *University of Warwick*

Attention is an important, yet often overlooked topic in decision making research. Many models make specific predictions about which items and attributes an individual is likely to attend to and at what point in the decision process. There are also other models which are not intentionally formulated to predict attention, but still imply that specific information will be deliberately ignored or that the decision process will terminate after a specific piece of information is made available to the model's deliberative core. We use eye-tracking to investigate a prediction which applies to several models: that individuals spend more time attending to information which they weight more strongly in their choices. Furthermore, we examine the prediction that this bias will increase over time as individuals move from an initial exploratory, information encoding phase to a second, more deliberative strategy. Previous research has demonstrated that in choices between simple, single attribute items, individuals have a bias to attend more to the item they eventually choose and that this bias grows over time. This is a phenomena referred to as the "gaze cascade". Our study is the first to examine these effects of attention bias within multi-attribute choice.

The results show that the gaze cascade effect is remarkably robust. The effect size is not significantly different to previous demonstrations, despite the significant additional complexity of items with 5-attributes each. However, the results also show that there are no attribute-wise effects: Individuals do not attend more to the information that they subsequently weight more highly in their choices. In fact we find no attribute-wise attention bias, with attention proportions not differing significantly from chance even when their influence on individuals' choices differs by an order of magnitude.

We also present a series of simulations which demonstrate that the existence of the gaze cascade is only compatible with models which posit that evidence is accumulated over time and that a decision is only made when the relative difference in accumulated evidence between the best and second best items becomes sufficiently large. The results are robust to very significant changes in the core model, so long as these two assumptions are retained. The findings from our computational simulations and empirical research provide significant constraints for existing and future models.

**Keywords:** Attention, value, choice, eye-tracking, modelling

# BIG, FAST, AND MEMORABLE: INCREASED GAMBLING IN RISKY DECISIONS FROM EXPERIENCE

**ELLIOT A. LUDVIG** *University of Warwick*  
**CHRISTOPHER R. MADAN** *University of Alberta*  
**MARCIA L. SPETCH** *University of Alberta*

When making decisions from experience, people often exhibit different patterns of risky choices than when making analogous decisions from description. For example, when asked whether they would prefer a guaranteed \$20 or a 50/50 chance of winning \$40, people are typically risk averse. In contrast, when the same risks and rewards are learned from experiences, people are often more risk seeking.

Here, we present data from a series of experiments that identify three factors that further increase risk seeking in decisions from experience: big wins, fast responses, and reminders for past wins. These factors may prove key to understanding why people engage in real-world gambling.

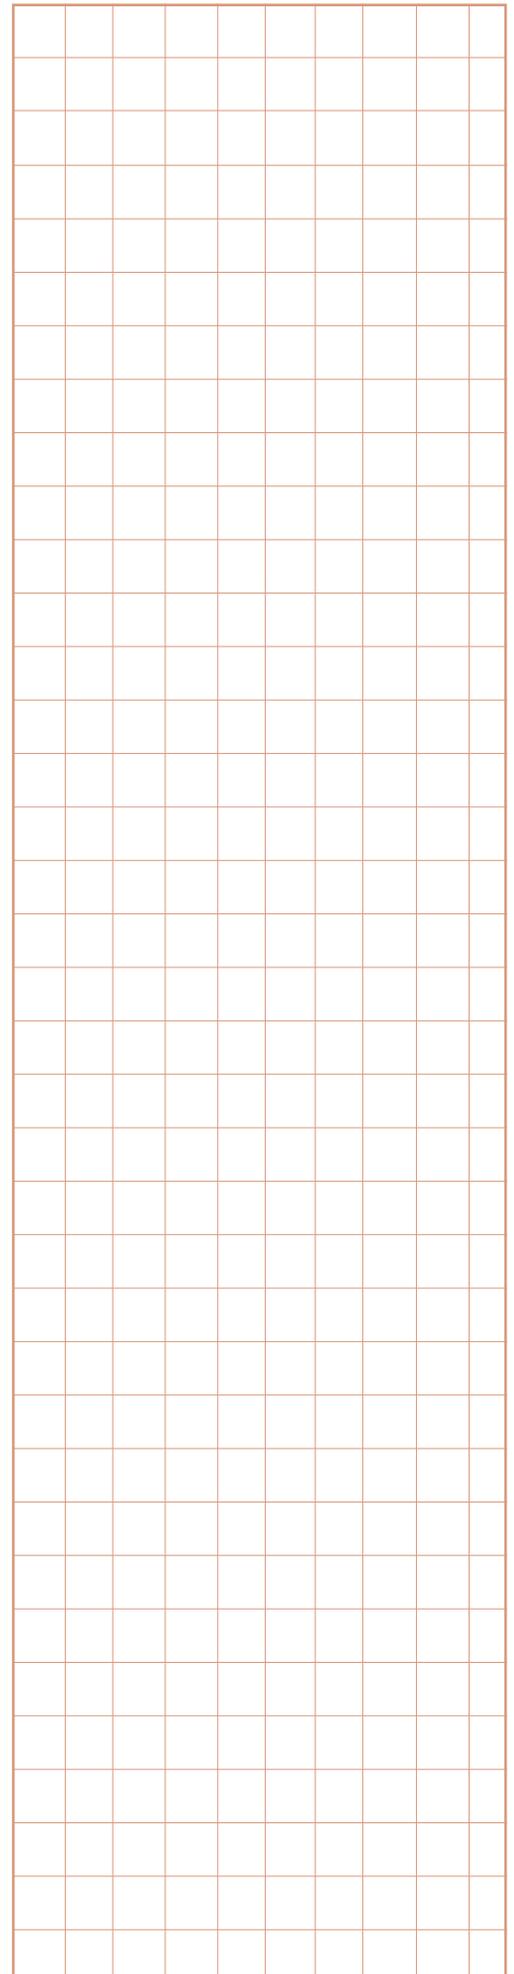
In Exp 1, people were tested on a pair of decisions. On high-value trials, they chose between a fixed 60 points and a 50/50 chance of 40 or 80 points. On low-value trials, they chose between a fixed 20 points and a 50/50 chance of 0 or 40 points. People were more risk seeking on the high-value trials than the low-value trials, reflecting an extra sensitivity to the biggest win (80) and the smallest win (0).

Exp 2 used an identical protocol as Exp 1, except there were two separate groups: fast and slow. People in the fast group had a short deadline for responding (1.5 s) and short inter-trial intervals (0.5 s). The slow group had a longer deadline (5 s) and longer inter-trial intervals (4 s). For both low- and high-value decisions, people were more risk seeking in the fast group, when they had to make their decisions rapidly.

In Exp 3, people were repeatedly tested on a decision between a fixed 40 points, and a 50/50 chance at 20 or 60 points. Each of the outcomes was also directly associated with a unique image. Before some of the decisions, people were primed with the unique image, thereby reminding them of the past outcome. Reminders for past wins increased risk seeking by nearly 20%, while reminders for the other two outcomes did not alter risky choice.

We interpret these results as reflecting a sampling-based decision process, along the lines of the decision-by-sampling (DbS) framework or the Dyna algorithm from reinforcement learning. On this view, people make decisions by sampling from the history of past outcomes and comparing the sampled outcomes for each option. These samples can be biased toward particular outcomes, either implicitly (e.g., saliency of big wins in memory) or explicitly (e.g., through reminders), allowing people's risky choice to shift based on the context.

**Keywords:** Risky choice, decisions from experience, gambling, extreme outcomes, time pressure, memory biases



## VALIDATION OF THE DIFFUSION MODEL FOR INVESTIGATING THE DECISION MAKING PROCESSES UNDERLYING COGNITIVE AFFECTIVE BIAS IN RODENTS

**CLAIRE HALES** *University of Bristol*

**EMMA ROBINSON** *University of Bristol*

**CONOR HOUGHTON** *University of Bristol*

Cognitive biases influence many aspects of human decision making. Some biases have also been shown in animals, such as judgement biases in interpretation of ambiguous information. This bias can also be altered by affective state. However, the cognitive processes underlying this cognitive affective bias (CAB) have not been explored. In rodents, a two-choice reaction time ambiguous cue interpretation task has been used to study CAB. Modelling data from two-choice reaction time tasks using the diffusion model can provide insight into cognitive aspects of the decision process, each of which is represented by a different parameter. We aimed to empirically validate use of the diffusion model with this ambiguous cue interpretation task by using specific manipulations that should affect corresponding model parameters of interest.

Male lister hooded rats (n=14) were trained to discriminate between two distinct auditory tones and make a response on the appropriate lever to either obtain high value (four sugar pellets) or low value (one sugar pellet) reward. Following successful discrimination, CAB was measured in test sessions where ambiguous tones and reference tones were presented. Tones with varying ambiguity were used to test the capacity of the diffusion model to represent differences in stimulus discriminability, which should affect the drift rate parameter. To test if the diffusion model is sensitive to changes in starting bias, reward value was altered so that a correct lever press after either tone resulted in a single sugar pellet reward. This should eliminate any differences in response bias between the two tones.

Experimental manipulations caused expected changes in corresponding diffusion model parameters that matched with behavioural data. The drift rate for the midpoint ambiguous tone was closest to zero, indicating it was the most difficult to discriminate. Near-midpoint ambiguous tones had drift rates that were not significantly different to the corresponding reference tone ( $p's > 0.05$ ), matching behavioural data that indicated comparable discrimination for these tones as for reference tones. Tones associated with the high value reward had significantly positively biased starting point parameters ( $p's < 0.001$ ). Interestingly, despite apparent neutral responding to the ambiguous midpoint tone from behavioural measures, the diffusion model indicated a significant positive bias in starting point for this tone ( $p = 0.027$ ). There were no significant differences between reference tone starting point parameters ( $p = 0.472$ ), or in the magnitude of the drift rate parameters ( $p = 0.948$ ) when both tones predicted the same outcome.

This study provides initial validation that the diffusion model can be applied to data from a rodent ambiguous cue interpretation task used to measure CAB. Experimental manipulations that should specifically affect stimulus discriminability and response bias did alter corresponding parameters in the diffusion model. Furthermore, application of the diffusion model to this CAB task can expose subtle changes in behaviour that are masked through traditional data analyses of behavioural measures. This supports further use of the diffusion model with this task to investigate how these aspects of the decision making process are altered by affective state manipulations and relate to changes in cognitive bias.

**Keywords:** cognitive affective bias, diffusion model, rodent

## THE EFFECT OF SPATIAL PROBABILITY AND REWARD ON RESPONSE TIME

**ROSIE CLARK** *University of Bristol*  
**RAFAL BOGACZ** *University of Oxford*  
**IAIN D. GILCHRIST** *University of Bristol*

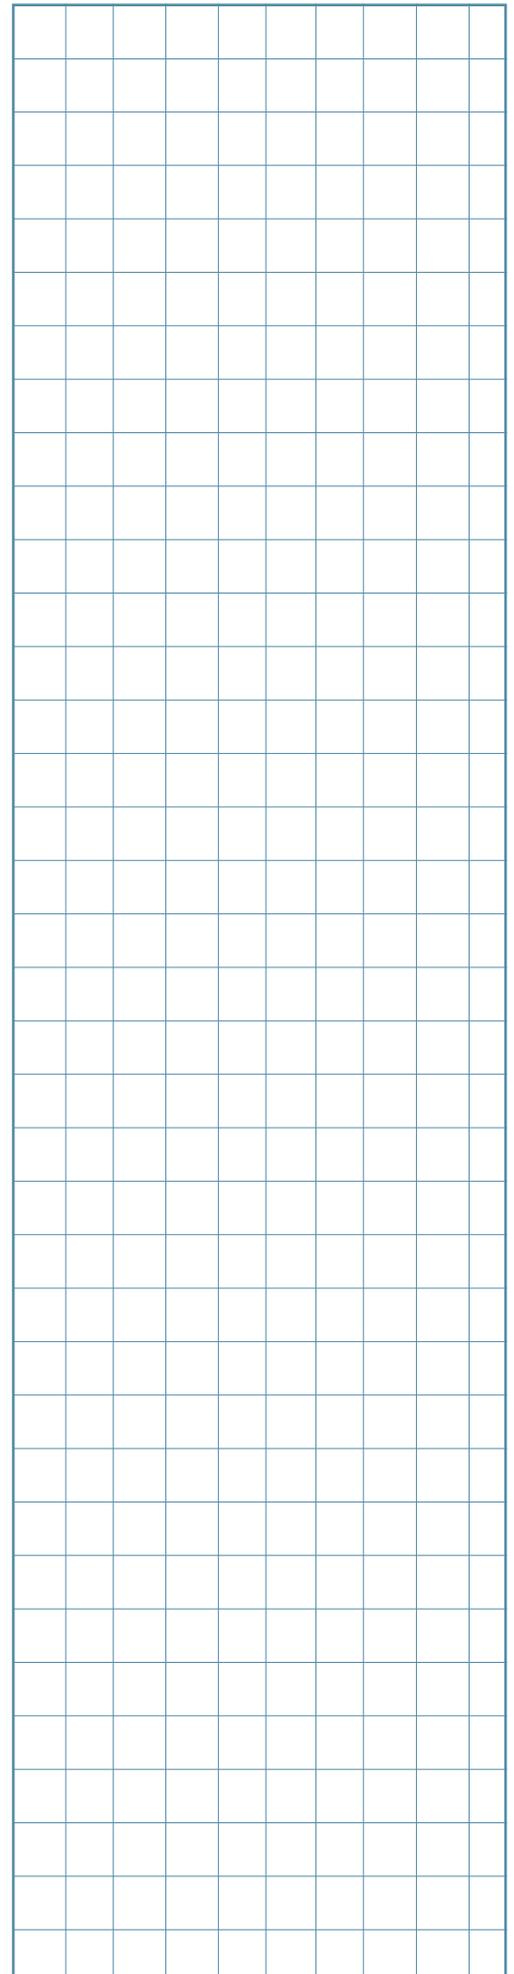
Prior information about our environment influences our decisions about where to look, and how urgent those decisions are. The relative probability of stimuli and the reward associated with them are two highly influencing forms of prior knowledge. Saccadic reaction times are faster to locations that occur more frequently (Carpenter and Williams, 1995), and some (largely primate) studies have shown mean saccadic latency was shorter in rewarded trials than in non-rewarded (Takikawa et al 2002). However, there is a lack of spatial reward contingencies in human studies (Bendiksby & Platt 2006). Within models of saccade latency distributions (such as LATER) both effects could be accounted for by a modulation in the starting point of the accumulation process.

In the first experiment we investigated the effect of spatial probability on manual response time and saccade latency concurrently. We found a significant difference between the saccadic response times across three probability levels, but no effect in the manual response time.

In a second experiment we investigated the effect of reward in a similar paradigm, we manipulated both the spatial distribution of reward within a block of trials (high reward and low reward side) and the overall reward level within a block (high reward block and low reward block). We found a significant effect of target side on the manual responses, and a significant interaction between block type and target, suggesting that the context of the reward values has an effect. For saccade latency there was also a significant interaction.

These results suggest that there are subtle differences in the way that reward and spatial probability have an impact on saccadic and manual responses respectively.

**Keywords:** saccades, reward, spatial probability, response times



## LOSS AVERSION IS A PROPERTY OF THE EXPERIMENTAL DESIGN, NOT THE PARTICIPANT

LUKASZ WALASEK *University of Warwick*

NEIL STEWART *University of Warwick*

One of the key assumptions in prospect theory is loss aversion. Loss aversion is the property of having a steeper value function for losses than gains. Surprisingly few attempts were made to elicit the loss aversion parameter  $\lambda$  in decision under risk. These estimates show great variability, with values of  $\lambda$  varying from 1.00 to 2.63.  $\lambda$  is the ratio of the slopes for losses and gains, with a value of 1 indicating no loss aversion, and values greater than one indicating loss aversion.

In four experiments we demonstrate that loss aversion results from the evaluation of a given gain or loss against the distribution of gains and losses within the experiment. We focus on one parametric methods for eliciting loss aversion, which requires an individual to make a series of accept/reject decisions about 50-50 gambles involving a gain and a loss of money (Tom et al., 2007). We experimentally manipulate the ranges for the distributions of gains and losses in the experiment.

In decision by sampling theory (Stewart, Chater &, Brown, 2006) the subjective value of a given gain/loss is given by the number of favourable comparisons to gains/losses in memory. For example a gain of £10 will have a subjective value of 1/4 in an experiment with a uniform distribution of gains in the range £0-£40 because 1/4 of gains in this range are smaller than £10 and make £10 look good. So decision by sampling predicts, in advance, that we should see more loss aversion in an experiment with a small range of losses and a large range of gains. For example, the magnitude £10 will have a subjective value of 1/4 when compared against gains in the range 0-40 and a subjective value of 1/2 when compared against losses in the range 0-20. This is loss aversion--the subjective value for a fixed magnitude is larger when it is a loss. If we were to reverse the ranges of gains and losses, decision by sampling predicts that we should observe the opposite of loss aversion.

In all four of the experiments--three online and one lab experiment with true incentives--we find that the loss aversion parameter is indeed largely determined by our choice of the distribution of gains and losses. We replicate loss aversion when the range of the gains is twice as large as the range of losses. When the ranges of gains and losses are equal we find no loss aversion. When the range of losses is twice as large as the range of gains we find the opposite of loss aversion. We conclude that while the aversion to losses may be a robust empirical phenomenon, loss aversion results, at least in part, from differences in the distributions of gains and losses in the environment. We offer a novel explanation of the origin of loss aversion, proposing that loss aversion should not be interpreted as a stable parameter of the prospect theory but instead as a property of the experimental design.

**Keywords:** loss aversion, Decision by Sampling, relative judgments, context, decision under risk, prospect theory

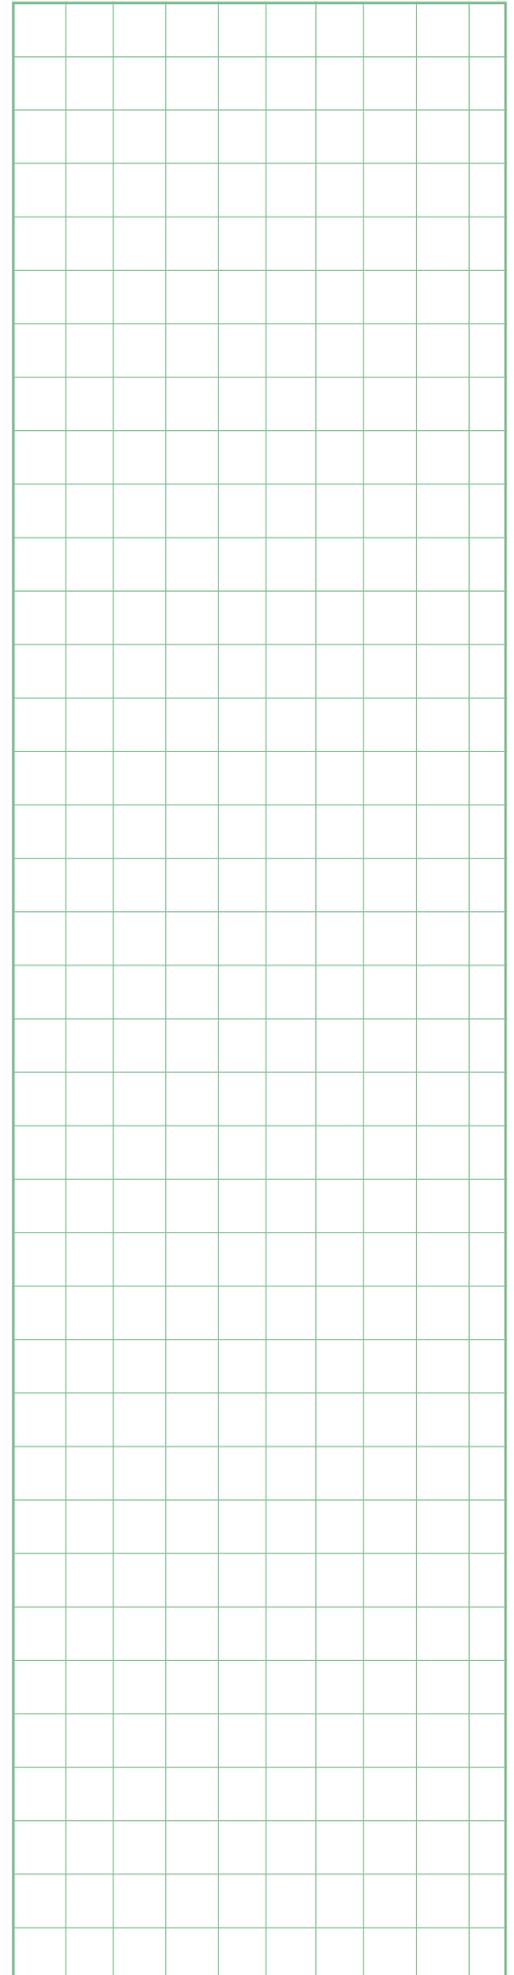
James Watson Room

## OPTIMAL MOODS IN BEHAVIOURAL MODELS

**PETE TRIMMER** *University of Bristol*  
**ELIZABETH PAUL** *University of Bristol*  
**MIKE MENDEL** *University of Bristol*  
**JOHN MCNAMARA** *University of Bristol*  
**ALASDAIR HOUSTON** *University of Bristol*

Moods can be regarded as fluctuating dispositions to make positive and negative evaluations. Developing an evolutionary approach to mood as an adaptive process, we consider the structure and function of such states in guiding behavioural decisions regarding the acquisition of resources and the avoidance of harm in different circumstances. We use a drift diffusion model of decision making to consider the information required by individuals to optimise decisions between two alternatives, such as whether to approach or withdraw from a stimulus that may be life enhancing or life threatening. We show that two dimensions of variation (expectation and preparedness) are sufficient for such optimal decisions to be made. These two dispositional dimensions enable individuals to maximize the overall benefits of behavioural decisions by modulating both the choice made (e.g., approach/withdraw) and decision speed. Such a structure is compatible with circumplex models of subjectively experienced mood and core affect, and provides hypotheses concerning the relationships that occur between valence and arousal components of mood in differing ecological niches. The approach is therefore a useful step toward being able to predict moods (and the effect of moods) using an optimality approach.

**Keywords:** Mood, emotion, evolution, core affect, circumplex, modelling.



## A MATTER OF UNCERTAINTY: OPTIMALITY AND SUB-OPTIMALITY IN SENSORIMOTOR LEARNING

**LUIGI ACERBI** *University of Edinburgh*

**SETHU VIJAYAKUMAR** *University of Edinburgh*

**DANIEL M. WOLPERT** *University of Cambridge*

Uncertainty is a ubiquitous feature in shaping human sensorimotor behaviour (Orban & Wolpert, 2011). Notably, humans are able to integrate prior information with sensory data to build 'posterior' estimates of perceptual and sensorimotor variables that take uncertainty into account (Kording & Wolpert, 2004).

In the first part of the talk, we show how combining probabilistic beliefs of varying complexity with a cost function leads to 'near-optimal biases' modulated by the detailed form of the posterior, in a study of sensorimotor timing (Acerbi et al., 2012). Optimal performance, though, is not always achievable when the task at hand requires the integration of complex probabilistic information, and sub-optimal biases may emerge. Crucially, such biases appear to be affected by estimation uncertainty too, as quantified by the spread of the observer's posterior distribution (Acerbi et al., 2012).

In the second part of the talk, we identify a possible source of sub-optimal performance in the failure of the motor system to correct for errors in conditions of increased estimation uncertainty. To test our hypothesis, unbeknownst to participants we randomly shifted visual feedback of their finger position during reaching in a centre of mass estimation task. Even though they were given enough time to compensate for this perturbation at the end of their movement (adjustment phase), participants only fully corrected for the induced error on trials with low uncertainty about target location; instead, correction was partial in conditions involving more uncertainty.

Albeit sub-optimal with respect to task demands, this lack of correction can be explained by considering an additional cost of adjusting one's response in conditions of uncertainty, a term in the cost function that can be interpreted as 'effort' (whether energy, time or computational). Even for simple, naturalistic tasks, such as centre of mass estimation, the effect of this additional cost can be significant and is strongly modulated by trial uncertainty. Our findings suggest that subjects' decision uncertainty, as reflected in the width of the posterior, is a major factor in determining how their sensorimotor system responds to errors, supporting theoretical models in which the decision making and control processes are fully integrated.

**Keywords:** Sensorimotor learning, Bayesian Decision Theory, Human performance, Probabilistic inference, Sub-optimality, Motor planning, Error correction.

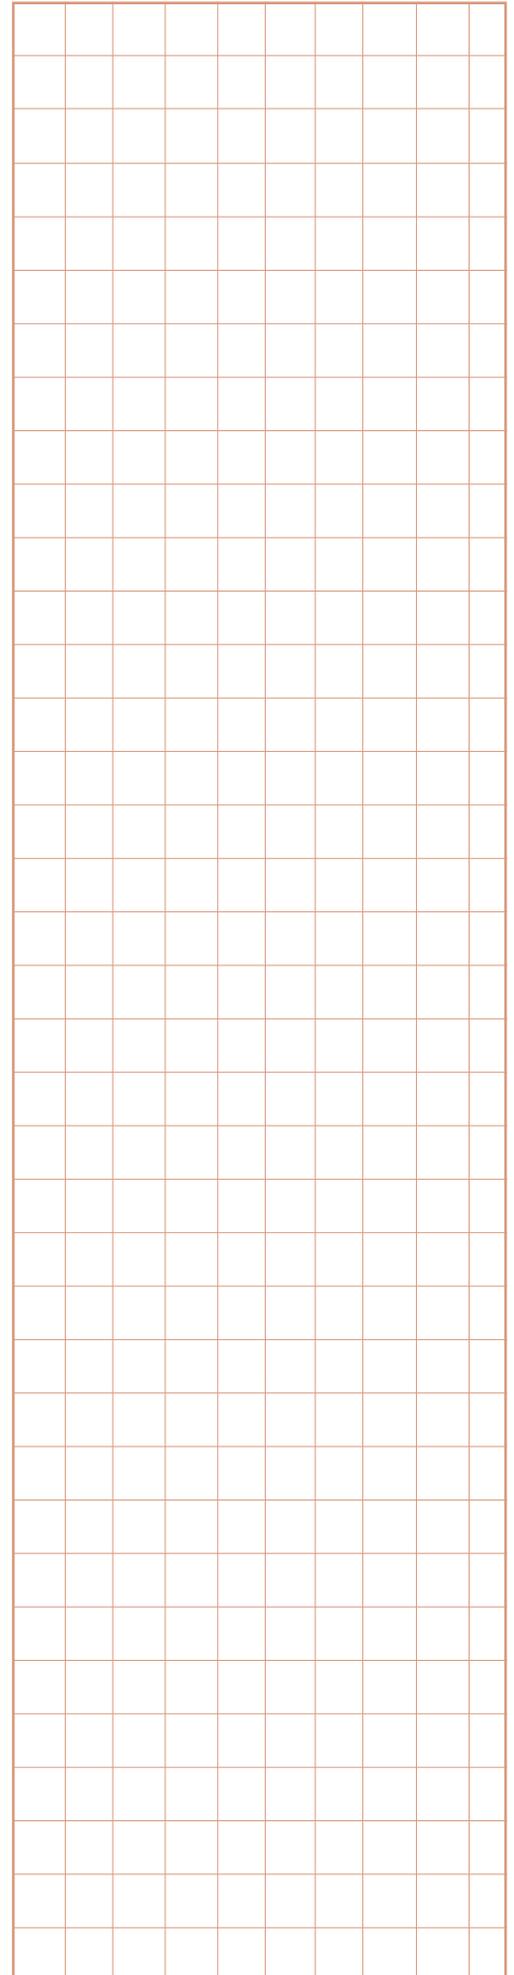
Francis Crick Room

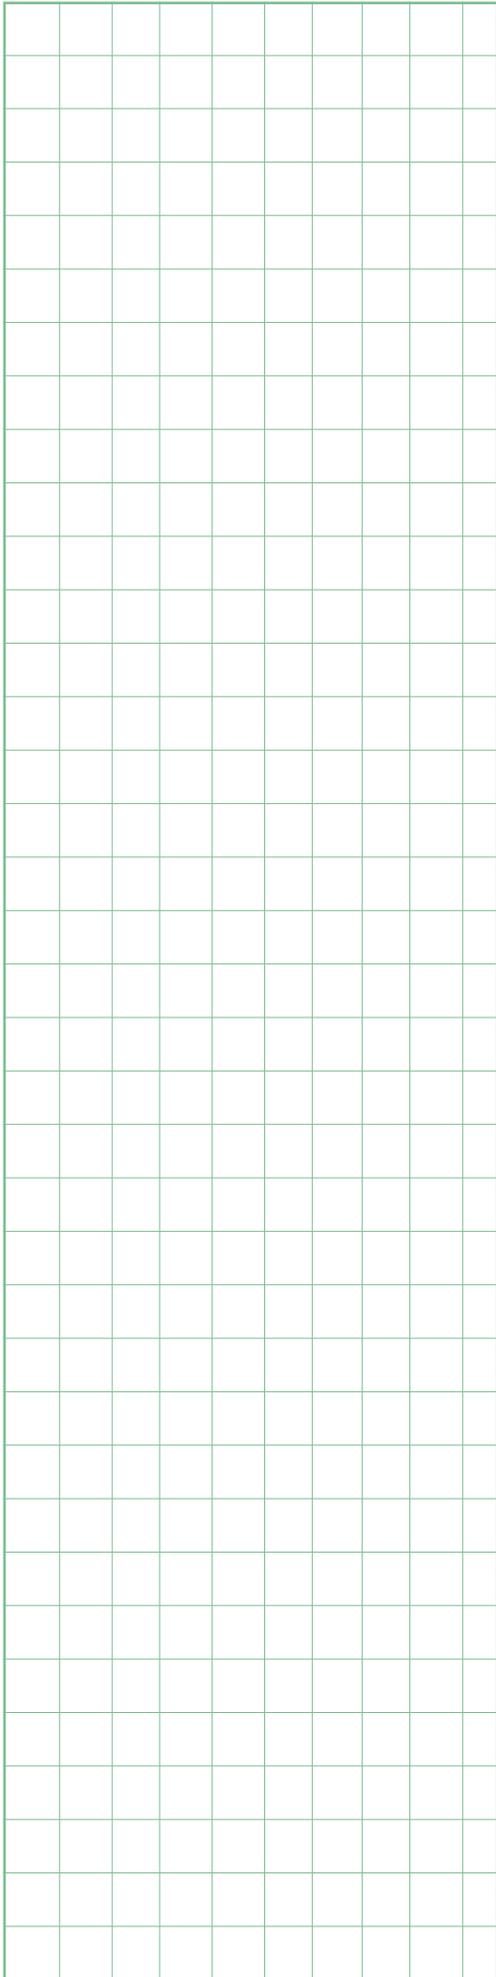
## MULTI-ATTRIBUTE DECISION-BY-SAMPLING

**DAVID RONAYNE** *University of Warwick*

I offer a theoretical account of how individuals evaluate multi-attribute alternatives along with empirical evidence. The model extends and develops the single-attribute decision-by-sampling (DbS) model of Stewart et al. (2006). The multi-attribute DbS model of how individuals assign value to multi-attribute alternatives assumes the following two-stage process: 1) The choice set presented to an individual is used as information to update the individual's belief over the total market offerings of that alternative (or product); 2) an alternative within the choice set is then evaluated by a finite series of dominance comparisons between it, the remaining alternatives in the choice set and others drawn from the posterior. The model captures all three of the most well-evidenced multi-attribute context effects: the attraction, compromise and similarity effects. Explaining all three of these effects within the same framework has been a challenge in the decision-making literature. I show how each of these "big three" context effects are produced theoretically by the model. I then provide experimental evidence, from both laboratory and online environments, that tests the model's predictions.

**Keywords:** Context Effects, Decision-by-Sampling, Amazon Mechanical Turk, Consumer Choice





## ECONOMICAL DECISION MAKING BY TEMNOTHORAX ALBIPENNIS ANT COLONIES

**CAROLINA DORAN** *University of Bristol*  
**NIGEL R. FRANKS** *University of Bristol*

Social insect colonies provide some of the richest examples of complex systems in nature. They are an excellent model for experimental investigation into questions of how group decisions are made as they allow direct manipulation of their components and observation of the collective behaviour. *Temnothorax albipennis* colonies are able to allocate the appropriate effort into gathering information regarding new homes in accordance with the quality of the nest they currently inhabit. Furthermore when faced with a risky choice they seem to be risk takers and to 'gamble' if the expected payoff is positive, i.e. it represents a gain. However, the time it takes colonies to reach a consensus gradually increases when the gain is smaller. Our latest research shows evidence that this species of ants exhibits economic rationality both at the individual and collective level. Both individual workers and colonies seem to base their decisions on the final expected payoff, namely the value of the target nest minus the value of the current nest minus emigration costs. We so concluded that when they are not totally isolated from other colony members and when they are not under severe time pressure, even the so called - simple components: the individual workers - are able to behave rationally. It is clear that animal collectives can solve more complicated problems than their isolated individual members. Nevertheless, research on collective decision-making should not underestimate the sophistication of the individual. We suggest that one of the clichés of complexity science - the stupid component and the sophisticated whole - needs to be over-turned in certain cases, especially in biology, by a recognition of sophistication at all levels.

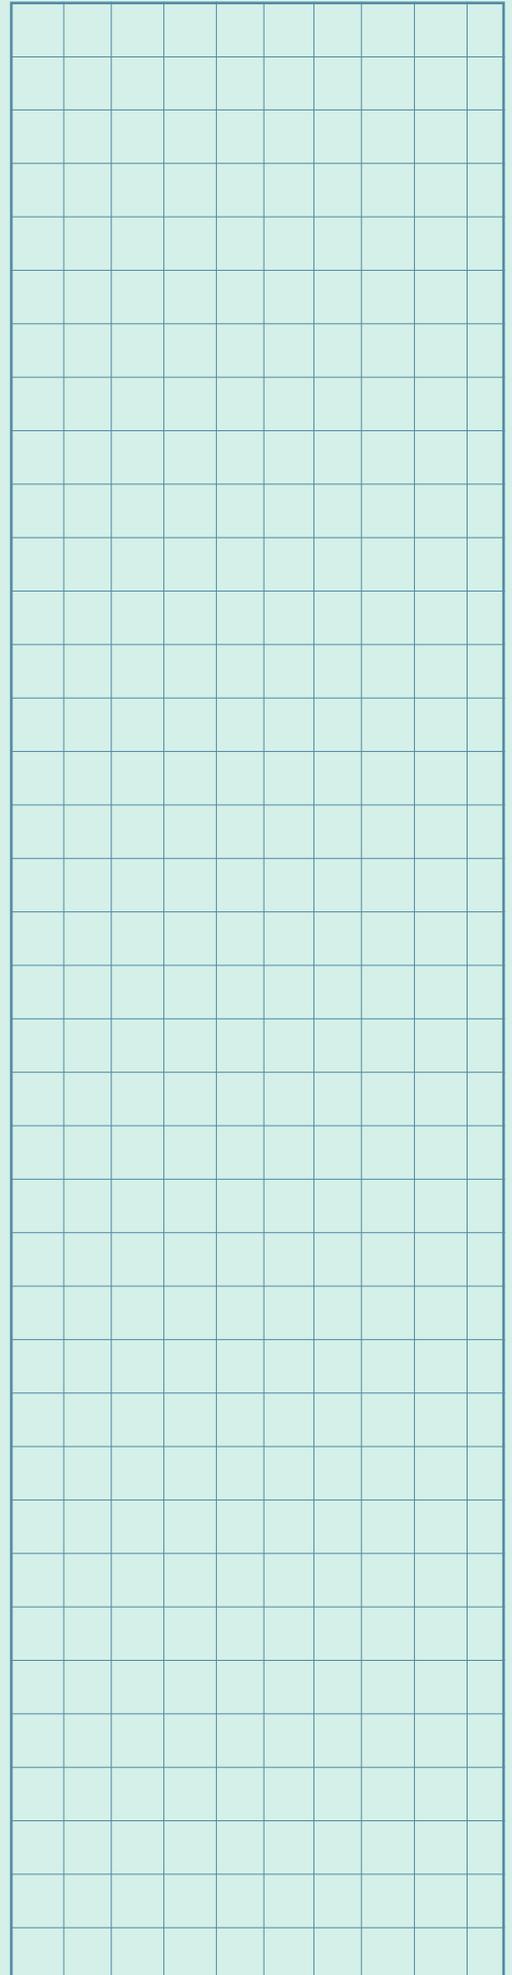
**Keywords:** risk, collective behaviour, decision-making

# DANIEL WOLPERT

## PROBABILISTIC MODELS OF SENSORIMOTOR CONTROL AND DECISION MAKING

DANIEL WOLPERT *Royal Society Noreen Murray Research Professor & Professor of Engineering, University of Cambridge*

The fields of decision making and sensorimotor control have developed in parallel although both require acting in real time on streams of noisy evidence. I will review our recent work covering both probabilistic models of decision making and sensorimotor control and the interactions between these processes. This will include the development of a cognitive tomography technique that can extract complex priors from simple decision making measurements. I will also review our work on the relation between vacillation and changes of mind in decision making, the bidirectional flow of information between elements of decision formations (such as accumulated evidence) and motor processes (such as reflex gains) and how active sensing through eye movements is used for visual texture categorisation.





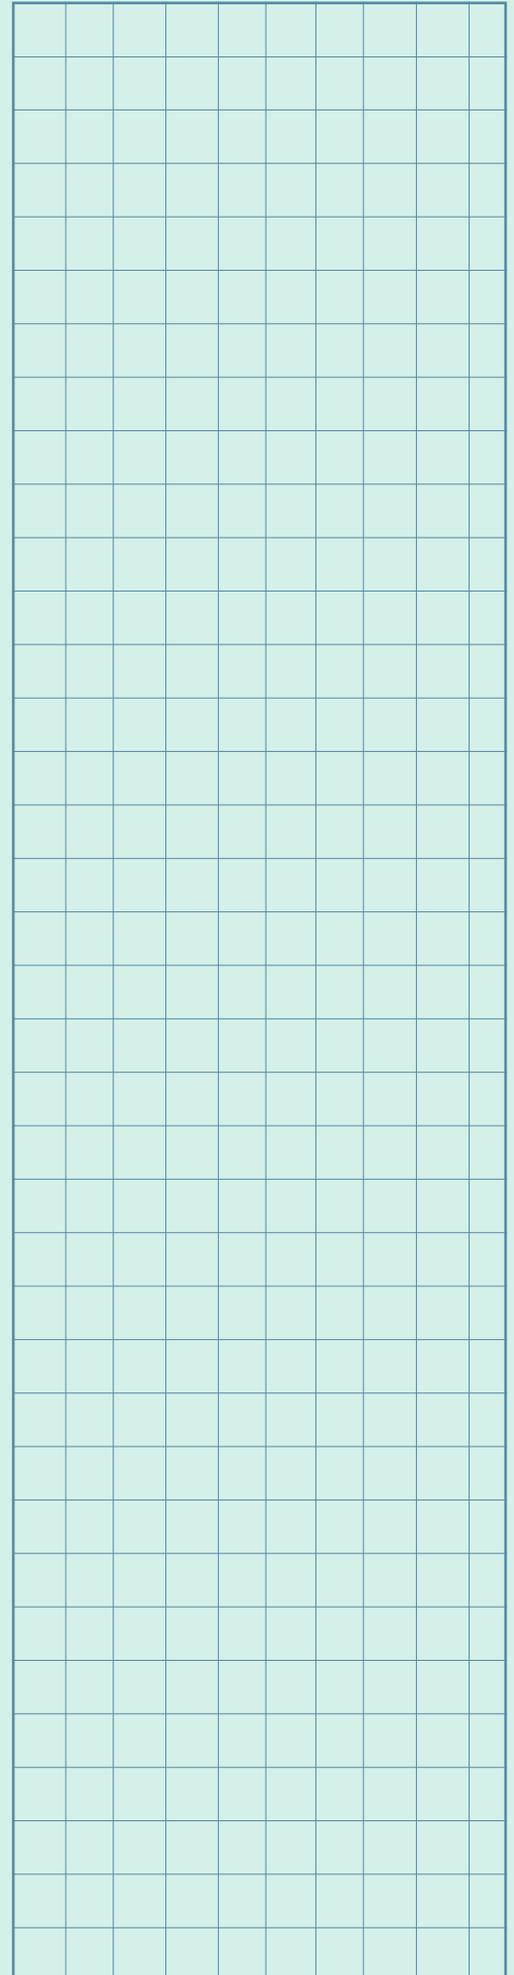
Rosalind Franklin Room

# ANTONIO RANGEL

## THE NEUROECONOMICS OF SIMPLE CHOICE

**ANTONIO RANGEL** *Bing Professor of Neuroscience, Behavioral Biology and Economics, California Institute of Technology*

Neuroeconomics studies what are the computations made by the brain in different decision situations, and how are these computations implemented and constrained by the underlying neurobiology. This talk describes recent fMRI, EEG and eye-tracking experiments designed to understand how the brain computes and compares values during simple decisions, like choosing between an apple and an orange.



## ELLSBERG'S PARADOX IN SENSORIMOTOR LEARNING

**DANIEL A. BRAUN** *Max Plank Institute for Biological Cybernetics and Max Plank Institute for Intelligent Systems*

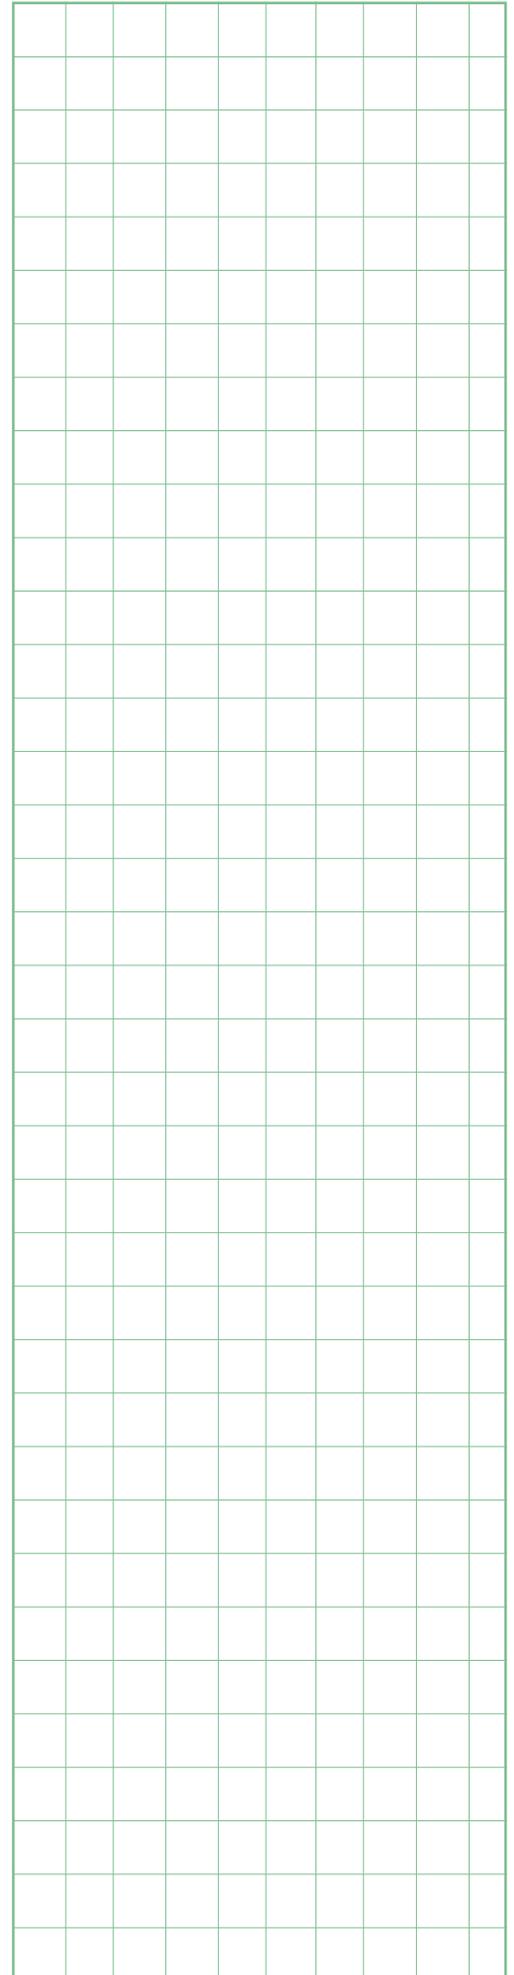
**JORDI GRAU-MOYA** *Max Plank Institute for Biological Cybernetics and Max Plank Institute for Intelligent Systems*

**PEDRO A. ORTEGA** *Max Plank Institute for Biological Cybernetics and Max Plank Institute for Intelligent Systems*

Both sensorimotor and economic behavior in humans can be understood as optimal decision-making under uncertainty specified by probabilistic models. In many important everyday situations, however, such models might not be available or be ambiguous due to lack of familiarity with the environment. Deviations from optimal decision-making in the face of ambiguity have first been reported by Ellsberg in economic choices between urns of known and unknown composition. Here we designed an urn task similar to Ellsberg's task and an equivalent motor task, where subjects choose between hitting partially occluded targets with differing degree of ambiguity. In both experiments subjects had to choose between a risky and an ambiguous option in every trial. The risky option provided full information about the probabilities of the possible outcomes. The ambiguous option was always characterized by a lack of information with respect to the probabilities. We could manipulate the degree of ambiguity by varying the amount of information revealed about the ambiguous option. In the motor task, we manipulated the extent to which an ambiguous target was occluded that subjects aimed to hit, whereas in the urn task we varied the number of balls drawn from the ambiguous urn before subjects made their decision. This way, we could test the more general hypothesis that decision-makers gradually switch from ambiguity to risk when more information becomes available. Ellsberg's paradox then arises in the limit case in which the ambiguous option gives away no information. We found that subjects tended to avoid ambiguous urns in line with Ellsberg's results, however, the same subjects tended to be ambiguity-loving or -neutral in the motor task. One of the most important points of Ellsberg's original experiment was to show that expected utility models—that is models that only care about maximizing expected success—cannot explain subjects' choice behavior under ambiguity. Since then a number of models for decision-making under ambiguity have been proposed. However, few of them are able to dynamically change the degree of ambiguity as new information arrives. Here we employ a multiplier preference model, that is a type of variational preference model for decision-making under ambiguity, and use it under a Bayesian update procedure to integrate novel information. We show that the deviations from optimal decision-making can be explained by such a robust Bayesian decision-making model. Our results suggest that ambiguity is a ubiquitous phenomenon, not only to understand economic choice behavior, but also sensorimotor learning and control.

**Keywords:** ambiguity, risk, sensorimotor learning

James Watson Room



# POLYA'S BEES: A MODEL OF DECENTRALIZED DECISION MAKING

**RUSSELL GOLMAN** *Carnegie Mellon University*  
**DAVID HAGMANN** *Carnegie Mellon University*  
**JOHN H. MILLER** *Carnegie Mellon University*

Decentralized systems of agents must often make critical choices, as, for example, when a swarm of bees (or a colony of ants) chooses a new hive (nest) location. Such natural decentralized systems seem to share a common decision-making process involving information search with positive feedback and consensus choice through quorum sensing. We model this process with an urn scheme that runs until hitting a threshold. We show that there is an inherent tradeoff between the speed and the accuracy of a decision. The proposed mechanism provides a robust and effective means by which a decentralized system can make reasonably good, quick choices. We find that this decision mechanism naturally produces risk aversion. Along with providing a deeper understanding of the behavior of many natural systems, this model provides a blueprint that could be applied in the creation of novel and productive decentralized decision mechanisms for use in human and artificial systems.

**Keywords:** consensus decision making, decentralized systems, Polya urn process, quorum, ants, bees

# A LEARNING RULE FOR OPTIMAL FEEDBACK CONTROL THAT PREDICTS CONTINUOUS ACTION SELECTION IN COMPLEX MOTOR TASKS

A. ALDO FAISAL *Imperial College London*  
ANASTASIA SYLAIDI *Imperial College London*

A rapidly growing body of evidence suggests that human decision making in the context of motor behavior is consistent with Optimal Feedback Control (OFC) (Landy & Wolpert, 2012, Todorov & Jordan, 2002, Nagengast et al, 2009). However the representations and mechanisms that underlie learning optimal controllers remain elusive, particularly in complex motor tasks that involve the manipulation of objects. We propose here that in such tasks the brain makes continuous decisions for the generation of complex trajectories by learning OFC based on the identification of unknown parameters of both body and object dynamics (Sylaidi & Faisal, 2012). We present experiments and a novel theoretical framework that can capture the temporal dynamics of motor learning on a trial-by-trial basis.

Human subjects (N=15) were instructed to move a virtual object of unknown dynamics from start to target locations in an unintuitive task that translated hand velocity into control forces on the object state. Subjects received a performance feedback after the completion of each trial in the form of a cost function capturing the end point error, a control penalty and an end state penalty.

While learning an optimal controller directly is a sophisticated non-linear dynamic programming problem, we test here the hypothesis that for the considered task context the brain only needs to learn the unknown task parameters, composed by arm and object dynamics, in a locally linear system identification process. This identification process of task dynamics enables in turn a decision on the form of the OFC policy. Our approach describes motor learning as gradient descent steps in the space of unknown task dynamics parameters. This mechanism is driven by the error between predicted and actually produced object movements in each trial and can be implemented at the neuronal level by the modification of synaptic weights via Hebbian learning rules. The aspect of motor adaptation in our approach involves both the update of task dynamics and the applied OFC strategy. It thus proposes a novel framework that expands and merges studies which on one hand use predictive models of OFC to fit human performance, assuming that the system dynamics are already known by humans at the initiation of a task (Nagengast et al, 2009) and on the other hand previous work which limits the investigation of motor learning to the assumption of updated task dynamics while considering a fixed control strategy throughout the experiment (Berniker & Kording, 2008).

Crucially, our adaptation model predicts accurately the gradual progression of human learning from trial to trial and performs better in capturing the behavior of subjects at the end of the experiment than an ideal observer model, which assumes complete knowledge of task dynamics. Our results suggest that the brain employs simple learning rules to support decisions implemented by near optimal control in complex object manipulation tasks. Our proposed framework provides thereby an algorithmic formalization, which can guide further experimental investigations on the neural foundation of cortical action selection and motor learning rules.

**Keywords:** learning, optimal feedback control, action selection

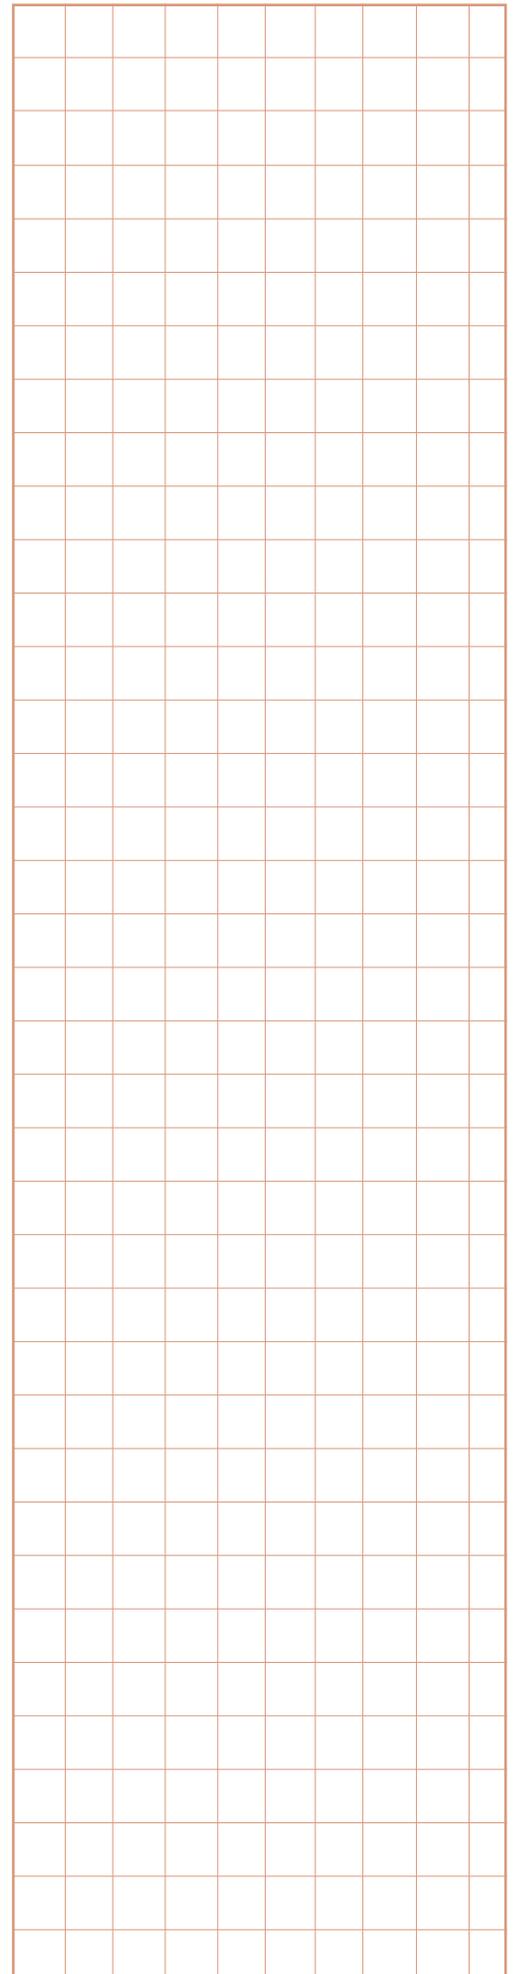
**RISKY CHOICES  
WITH TEMPORAL OR MONETARY OUTCOMES:  
DOES TIME EQUAL MONEY?**

**TIM RAKOW** *University of Essex*

**NATHANIEL J. S. ASHBY** *University of Essex; Carnegie Mellon University*

Many choices have consequences for the length of time that it takes to complete a task, activity or journey; yet laboratory research on risky choice is dominated by investigations of the choice between monetary lotteries. Previous research using hypothetical choice scenarios suggests that people's decisions are less sensitive to temporal expenditures and outcomes than to monetary ones; suggesting that money is not a good proxy for time if one wishes to understand decisions about time. In two studies, we conducted direct tests of decision makers' sensitivity to monetary and temporal outcomes in risky choice. For instance, the participant might face a choice between a risky alternative that imposes possible delays of 60 or 2 seconds, and a safer alternative with 31- or 25-second delays possible; and his/her preferences would be compared to those for an equivalent choice involving monetary losses. Choices were not hypothetical; the payoff distributions of each alternative were either presented via descriptions or learned through experience; and outcomes were framed as either losses or gains. Experiential choices over losses and gains were highly similar for time and money, with preferences developing in similar ways, and at similar rates, as "good" and "bad" outcomes were experienced. The only notable difference between time and money for these experience-based choices was a small reversed reflection effect in choices involving time, with slightly more risk seeking for gains than for losses. Described choices also differed little between the two domains, though choices involving time showed slightly lower sensitivity to changes in framing and expected value. Together, these studies suggest that when the outcomes of risky choices are consequential then time really is like money. These data challenge several theories which propose that time matters less than money when decisions are made; and speaks to the generalizability of laboratory research using monetary gambles.

**Keywords:** Risky choice; time delays; time saving decisions; decisions from experience; incentivization; loss-gain framing



## COORDINATION IN LARGE GROUPS: WHAT WOULD WE DO IF ZOMBIES ATTACK?

**PETER S. RIEFER** *University College London*

**DANIEL C. RICHARDSON** *University College London*

**BRADLEY C. LOVE** *University College London*

Research in game theory has investigated coordination in several scenarios, such as the battle of the sexes or market-entry games, but neglected the fact that we often face situations where we coordinate as part of a large group of people rather than dyads or small parties. For example, if we choose which of two routes to take to work, we will consider our experiences and try to anticipate which of the two options will be less busy and therefore faster to drive on. Similar to the routes scenario, our study examines the coordination behavior of large groups with about 50 people in a zombie game where players have to choose in which of two hiding places to hide from attacking zombies. Points were distributed according to how many people were in a certain hiding place, giving higher rewards to players in less crowded places. In order to investigate how people use experiences and build anticipations about other people's behavior, subjects played the game for several rounds.

We conducted two experiments using this zombie game. In experiment 2, we added volatility information for subjects in one condition who could then see how many people switched from one hiding place to the other compared to the previous round. Using this approach, we tried to examine how information about other people's switching influences individual beliefs and decisions in coordination.

Both experiments show that groups generally struggle to coordinate and reach equilibrium, which is at a fifty-fifty split between hiding places and describes the situation, where no one can do better by switching to the other hiding place. When groups come close to equilibrium or even reach it, they will almost certainly drift away from it, due to some individuals who keep on switching choices. Regarding these findings, it seems unlikely for large groups to find and retain stable equilibria, since there will most likely be some people who nudge the group right back into volatility. On closer look, we found that switching choices after unsuccessful rounds is mainly responsible for lower individual scores, whereas switching after successful choices showed no significant impact on performance. We furthermore found that while switching after unsuccessful choices relates to more myopic thinking, such as choosing the currently better option in the upcoming round, switching after successful choices is anticipatory and related to people's strategy of integrating expectations about others.

Results of experiment 2 replicated the findings of experiment 1. Although performance wasn't significantly different in the conditions, we observed some behavioral differences. Subjects with volatility information switched more frequently after unsuccessful rounds than successful ones, while those without volatility information didn't. When examining people's strategies individually with a probabilistic choice model, we found that volatility information seems to make people rely on more myopic strategies. Overall, these findings suggest that seeing the switching behavior of others causes people to fall back on simple strategies where they immediately change from an unsuccessful option to the currently better one.

**Keywords:** Group behavior, Decision-making, Coordination

Rosalind Franklin Room

## TO ACT OR NOT TO ACT: LEARNING THE VALUE OF NOT ACTING

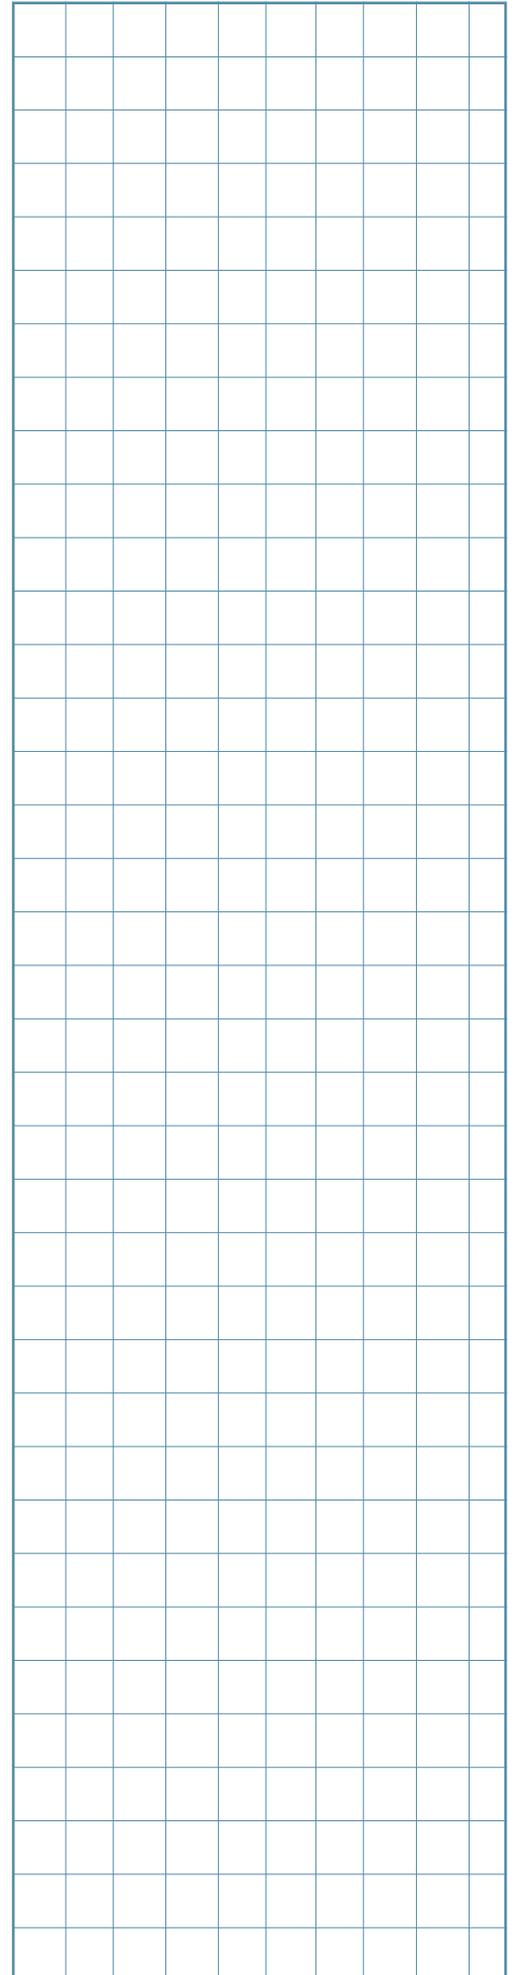
**RAFAL BOGACZ** *University of Bristol; University of Oxford*

**PAUL HOWARD-JONES** *University of Bristol*

**KEVIN LLOYD** *University of Bristol*

Sometimes the best course of action is to do nothing (non-action). However, it is unclear how the brain learns the value of non-action and how it is compared to the value of action during decision making. To address these questions, we repeatedly confronted human participants with the choice of selecting one alternative with a button press or the other by doing nothing. Analysis of the behavioural data with reinforcement learning models did not reveal differences between the rate at which participants learned the values of action and non-action. At feedback, in addition to brain activity correlated with the updated value of action and action prediction error, we found brain activities correlated with the updated value of non-action and non-action prediction error. The representations of non-action prediction error and non-action value were found in right and left inferior frontal gyri, respectively, which are regions previously implicated in studies of response inhibition. Additionally, we found value comparison activity at feedback in dorsomedial frontal cortex and frontopolar cortex, the latter predicting individual differences in exploitative behaviour and performance across participants. Our results suggest that the processes of learning the value of non-action resemble those for learning the value of action, engaging response inhibition regions in a manner analogous to the engagement of regions involved in the planning and implementing of action when learning the value of action.

**Keywords:** decision making, value, response inhibition, fMRI



## COMPUTATIONALLY RATIONAL DECISION MAKING

**ANDREW HOWES** *University of Birmingham*

**RICHARD L. LEWIS** *University of Michigan*

**PAUL WARREN** *University of Manchester*

**GEORGE FARMER** *University of Manchester*

Computational rationality is an approach to explaining behavior in which theories are specified as optimisation problems defined by a bounded machine, a utility function and the experience of an adaptation environment (Lewis, Howes, & Singh, 2014; Howes, Lewis, & Vera, 2009). Computational rationality holds out the promise of enhancing the contribution of rational analysis to testing theories of psychological mechanism and thereby offers a means of bringing together mechanistic and rationalistic approaches to psychology (Howes, Lewis, & Singh, 2014). It also promises explanations of behaviors that have been widely considered to be irrational but may be seen as rational given bounded neural architectures (Dayan, 2014; Holmes & Cohen, 2014) and bounded experience (Hahn & Warren, 2009; Le Mens & Denrell, 2011). In this talk I will introduce computational rationality and illustrate it with analyses of two key decision making phenomena: (1) preference reversals, and (2) risk preferences. I will show these phenomena are computationally rational consequences of uncertainty in expected value calculation.

Preference reversals occur when a preference for one choice over another is reversed by the addition of further choices. It has been argued that the occurrence of preference reversals in humans suggests that they violate the axioms of rationality and cannot be explained with utility maximization theories. In a recent article we used numerical simulation to demonstrate that for a range of types of contextual preference reversal, including the attraction, compromise and similarity effects, the rational, expected value maximizing, choice between existing options is one that is influenced by the addition of new choices (Howes, Warren, Farmer, El-Derredy, & Lewis, 2014). The analysis assumes that people rationally integrate two sources of information, one based on an estimate of expected value and one based on observation of the rank order of attribute values. I will also show that the same assumptions explain some observed risk preference effects, particularly risk aversion in the domain of gains and risk seeking in the domain of losses.

**Keywords:** Bounded optimality, computational rationality, decision making, preference reversals, risk preference.

James Watson Room

## CONFIDENCE MATCHING IN SOCIAL INTERACTION

**DAN BANG** *University of Oxford; Aarhus University*  
**LAURENCE AITCHISON** *University College London*  
**PETER E. LATHAM** *University College London*  
**BAHADOR BAHRAMI** *Aarhus University; University College London*  
**CHRISTOPHER SUMMERFIELD** *University of Oxford*

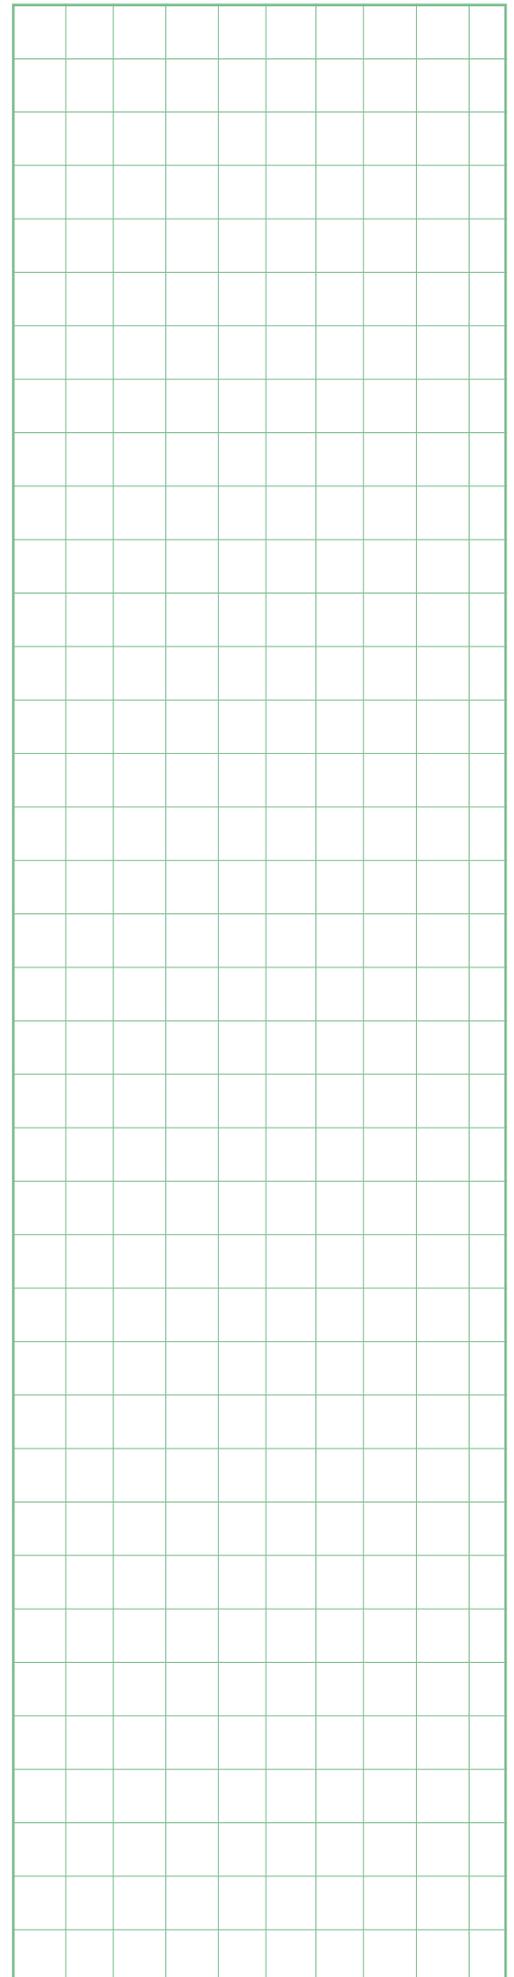
Humans and other animals use estimates about the reliability of sampled evidence to arbitrate between competing choice alternatives [Kepecs et al., 2008, Nature]. However, unlike other animals, humans can report on these estimates, saying "I'm sure it's this one" when the evidence in favour of a specific alternative is deemed highly reliable. Groups of individuals use such confidence reports to resolve disagreement about which choice alternative to select, usually opting for the one supported by higher confidence [Bahrami et al., 2010, Science]. To do so optimally, individuals must solve a 'mapping problem': they must map their 'internal' estimates onto 'shareable' confidence reports so as to maximise the probability that the group selects the better choice alternative [Bang et al., 2014, Concious Cogn]. For instance, if A on average samples more reliable evidence than B, then A should on average be more confident than B. In this talk, I report on two experiments in which we addressed the mapping problem in the context of a visual perceptual task.

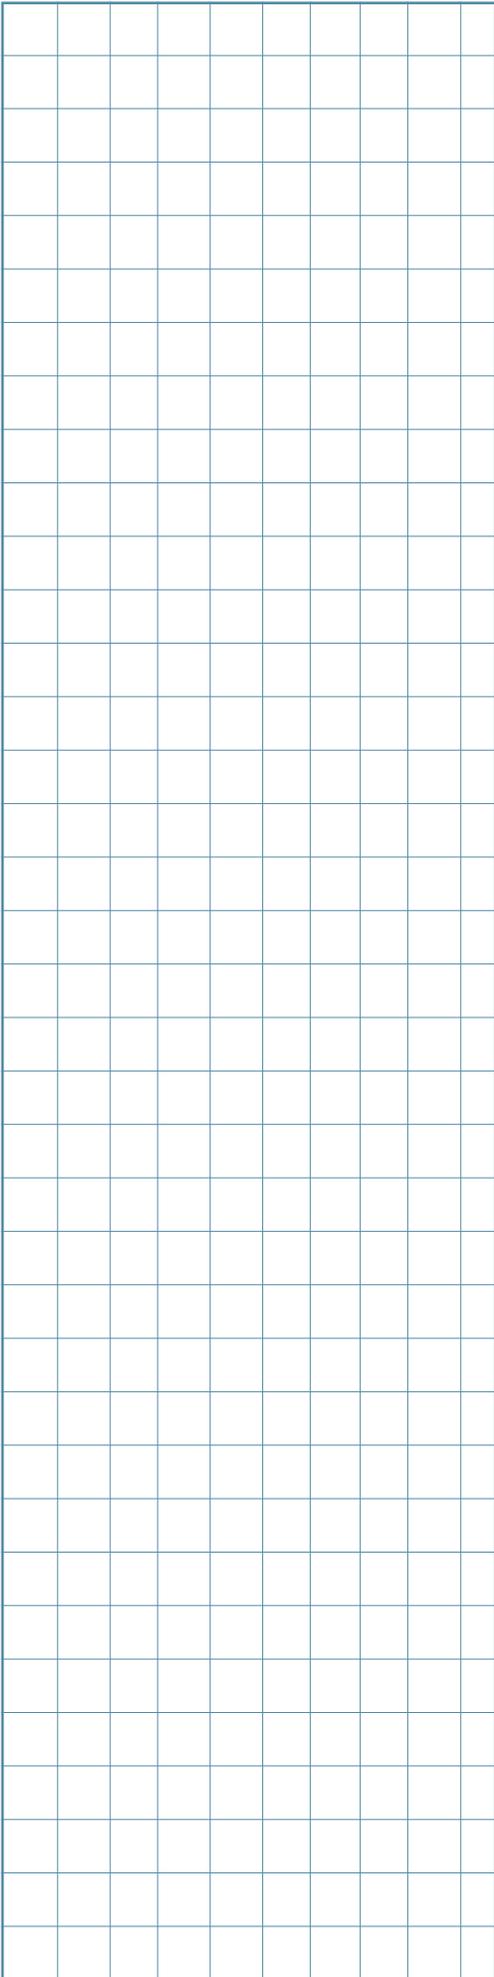
In the first experiment, we tested how pairs of individuals solve the mapping problem. Two participants (15 pairs in total) privately made a binary decision about a brief stimulus and indicated how confident they felt about this decision on a scale from 1 to 6. The responses were then shared, and the private decision made with higher confidence was selected as the group decision. Feedback about the accuracy of each decision was displayed before the initiation of the next trial. A comparison of models indicated that participants' strategy for maximising group accuracy was to match their distributions over the confidence scale. For participants with similar levels of individual accuracy, confidence matching yielded a relative increase in group accuracy (cf. they should on average be equally confident). In contrast, for participants with dissimilar levels of individual accuracy, confidence matching led to a relative decrease in group accuracy (cf. the more accurate participant should on average be more confident than the less accurate participant).

In the second experiment, we tested the adaptive/maladaptive effects of confidence matching more directly. One naive participant (38 participants in total) performed the visual perceptual task together with four simulated partners in a two (accuracy) x two (confidence) within-subject design. The simulated partners were less or more accurate and less or more confident than the participant. In line with model predictions, confidence matching yielded a relative increase in group accuracy when participants interacted with "poorly calibrated" partners ("less accurate but more confident" and "more accurate but less confident"), but led to a relative decrease in group accuracy when participants interacted with "well-calibrated" partners ("less accurate and less confident" and "more accurate and more confident").

We suggest that confidence matching arose as a default strategy for cooperation in a world where individuals are rarely perfectly calibrated [Fleming et al., 2010, Science]. Crucially, confidence matching is computationally inexpensive, works for different tasks and is not dependent on the presence of feedback about individual accuracy.

**Keywords:** social interaction; confidence; decision-making; computational model; perception





## THE ROLE OF THE SUBTHALAMIC NUCLEUS IN MULTIPLE ALTERNATIVE PERCEPTUAL DECISION-MAKING REVEALED BY 7T STRUCTURAL AND FUNCTIONAL MRI

**M.C. KEUKEN** *Max Planck Institute for Human Cognitive and Brain Sciences;  
Cognitive Science Center Amsterdam*

**L. VAN MAANEN** *Cognitive Science Center Amsterdam*

**R. BOGACZ** *University of Oxford*

**A. SCHÄFER** *Max Planck Institute for Human Cognitive and Brain Sciences*

**J. NEUMANN** *Max Planck Institute for Human Cognitive and Brain Sciences*

**R. TURNER** *Max Planck Institute for Human Cognitive and Brain Sciences*

**B.U. FORSTMANN** *Max Planck Institute for Human Cognitive and Brain Sciences;  
Cognitive Science Center Amsterdam*

In our everyday life we constantly have to make decisions between many different choice options. Recently, quantitative mathematical and neurocomputational models have been developed that make predictions about brain structures involved in decision-making with multiple alternatives. One such model is the Multiple Sequential Probability Ratio Test (MSPRT), which predicts that activity in the subthalamic nucleus (STh), a small structure in the basal ganglia, becomes more active with an increasing number of choice alternatives.

The present study set out to test this hypothesis using ultra-high 7T structural and functional magnetic resonance imaging in healthy human subjects. By simulating the MSPRT model, we generated concrete predictions that were then compared to the observed behavior as well as the activation pattern in the STh. Preliminary results indicate the involvement of the STh in decision-making.

**Keywords:** Multiple-choice decision-making, Subthalamic nucleus

Francis Crick Room

## SUB-OPTIMAL INFORMATION SEARCH STRATEGIES DURING MULTI-ATTRIBUTE CHOICE

**LAURENCE HUNT** *University College London*

**NISHANTHA MALALASKERA** *University College London*

**DANIEL GRIGAT** *University College London*

**RAY DOLAN** *University College London*

**STEVE KENNERLEY** *University College London*

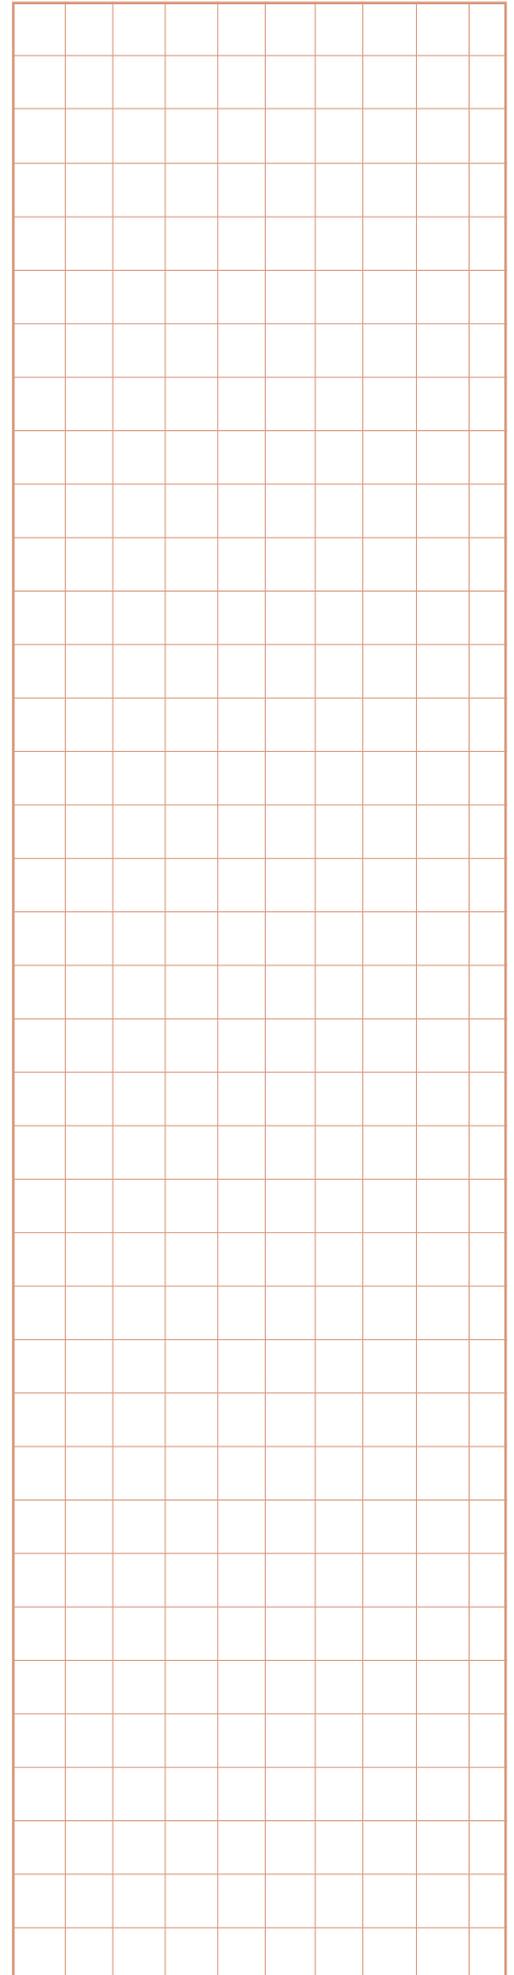
A central feature of many real-world decisions is that each alternative consists of several attributes. Such multi-attribute decisions may be realised in many different ways - ranging from attribute-based strategies (such as elimination-by-aspects) to alternative-based strategies (such as weighted adding of attributes). These strategies make opposing predictions as to how information will be acquired during decision formation, and, importantly, to how neural circuits can implement the decision as it is being made. Information acquisition may also frequently involve costs to the decision-maker.

In this study, we consider what normative principles might govern information search strategies in a multi-attribute choice task, and whether these match with empirical observations in humans and macaque monkeys. Two choice alternatives, consisting of two attributes, were presented. Subjects sequentially selected which feature of each alternative they wished to reveal; they were also able to terminate information sampling early in order to make a choice. Potential information that might be uncovered at each turn was drawn from a flat probability distribution, with which the subjects were familiarised with before commencing the task. This meant that a normative dynamic programming approach could be adopted, to derive the optimal strategy for calculating the value of gathering information at each turn. The dynamic programming model provides normative predictions of both when information sampling should be terminated, and also which information is most valuable to sample next, if reward is to be maximised.

We probed information gathering behaviour of both human and macaque subjects on analogous versions of the decision task. Human data was collected from a large subject pool (>8,000 participants) via a smartphone app, and compared to laboratory data from a smaller subject pool (21 participants, collected whilst undergoing magneto-encephalography). Choice data from two macaque monkeys was collected whilst undergoing neurophysiological recording from prefrontal cortex. Whereas human subjects paid explicit costs for sampling information, macaque subjects underwent the opportunity cost of time.

Some key aspects of human and macaque behaviour matched well with predictions from the normative model. For example, subjects would terminate information sampling early if informative cues had been received that made one alternative much more likely to be rewarding. However, there were also intriguing and unambiguous violations of the normative model. For example, human subjects were found to be particularly biased toward reducing uncertainty about the value of the currently preferred alternative, even if other information searches would prove more valuable. Macaque subjects showed a similar bias, in that they would often be unwilling to terminate a decision when confronted with one very poor alternative, without first ascertaining information about the other alternative. I will present potential considerations for the origins of these non-normative behaviours, and what they may imply for the implementation of such decisions in neural circuits.

**Keywords:** Information search, Multi-attribute decision making



## BETTER TOGETHER – UNDERSTANDING COLLABORATIVE DECISION MAKING

**REBECCA FLOYD** *University of Bristol*

**DAVID LESLIE** *University of Bristol*

**SIMON FARRELL** *University of Western Australia*

**ROLAND BADDELEY** *University of Bristol*

The adage 'two heads are better than one' suggests that when two people discuss an ambiguous situation, the outcome of a joint decision will be closer to reality than each individual's evaluation. It has been suggested that this has contributed to the success of our species by enabling us to carve out a unique 'socio-cognitive niche' (Whiten and Erdal, 2012). However, research has led to varying opinions about whether there is a benefit to working together e.g. 'Wisdom of the Crowd' (Galton, 1907) vs Groupthink (Janis, 1971).

Our studies seek to understand the processes underlying group decision making and why they might cause such divergent findings. We use ambiguous stimuli presented to pairs of participants. In each trial, participants individually indicate their perception of the correct answer, before then discussing and agreeing a joint answer. Together with pre-testing and self-reported personality questionnaires, the joint decision process is examined in the context of relative individual ability at the task and interpersonal factors.

Our findings so far suggest that at a broad level, the adage holds. On average, the joint answers were closer to correct than the answers of individual participants. However, looking at pair-wise performance, the joint answers were often less accurate than those given by the more "sensitive" (i.e. more accurate) participant, especially when there was a large discrepancy in participant's abilities to detect dot direction. It appears that a weighting of individual estimates in the consensus decision reflects the precision of the initial estimates and we aim to interpret these results in the context of a Bayesian model of how the participants combine their personal information.

**Keywords:** Joint Decision Making, Collaboration

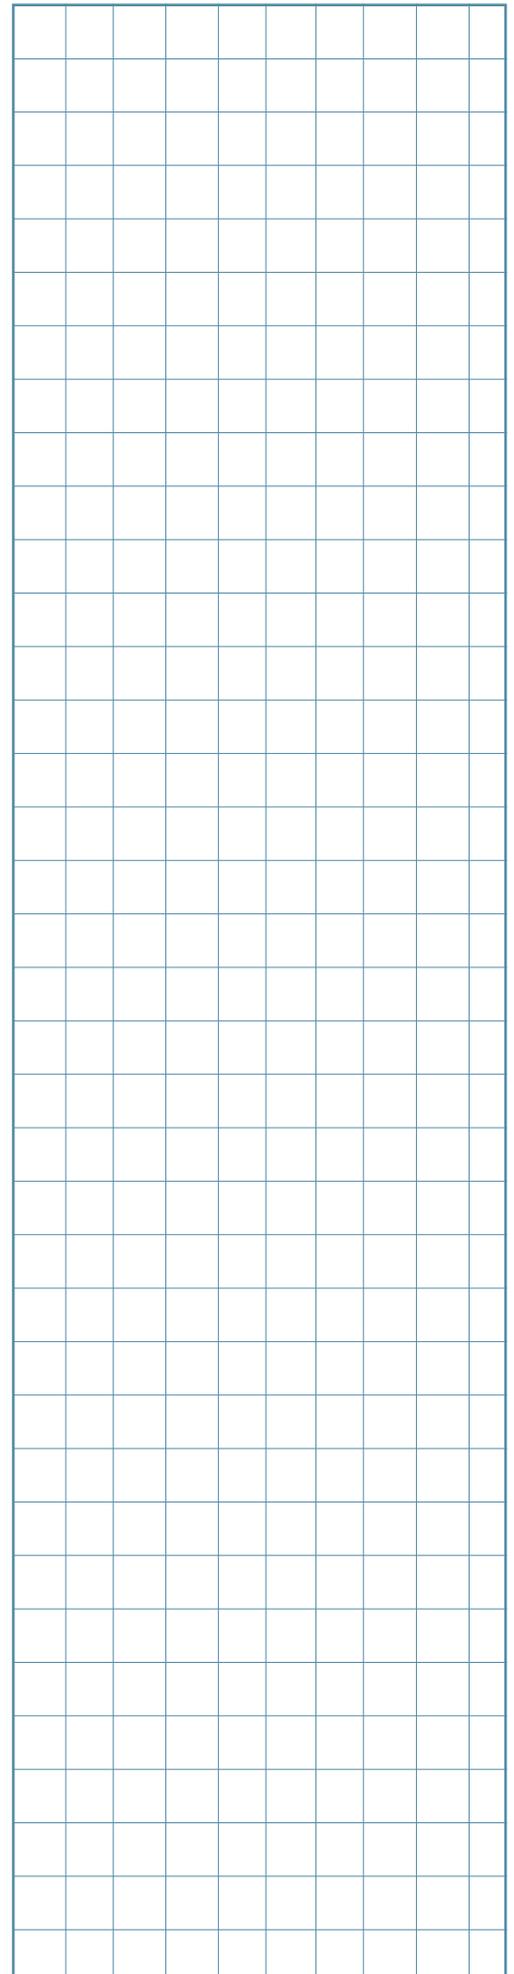
## PERCEPTUAL DECISION MAKING IN ROBOTICS AND NEUROSCIENCE

**NATHAN LEPORA** *University of Bristol and Bristol Robotics Laboratory*

A main principle underlying animal perception is the accumulation of evidence for multiple perceptual alternatives until reaching a preset belief threshold that triggers a decision, formally related to sequential analysis methods for optimal decision making. In a series of papers, we have formalized a Bayesian perception approach for robotics based on this understanding of animal perception. Our formalism extends naturally to active perception, by moving the sensor with a control policy based on evidence received during decision making. Benefits of active Bayesian perception include: (i) robust perception in unstructured environments; (ii) an order-of-magnitude improvement in acuity over passive methods; (iii) a general framework for Simultaneous Object Localization and Identification (SOLID), or perceiving 'where' and 'what'; and (iv) a formalism that naturally integrates with reinforcement learning so that both the active control policy and the appropriate belief threshold can be tuned appropriately to the contextual situation. A strength of the formalism is that it connects closely with leading work in neuroscience, allowing insights from animal perception to be transferred to robot perception. For example, these methods have enabled the first demonstration of hyperacuity in robot touch, giving perceptual acuity finer than the sensor resolution, as is common in animal perception. As discussed above, they also give robust perception in unstructured environments

in which there is uncertainty in both where and what objects are, as is also a central aspect of animal perception. In this talk, we describe this bio-inspired approach to robot perception and how it connects with computational neuroscience, in particular the macro-architecture of the cortex and basal ganglia.

**Keywords:** Decision making, robotics, active perception



## THE ATTRACTION EFFECT IS RATIONAL GIVEN UNCERTAINTY IN EXPECTED VALUE CALCULATION

**GEORGE FARMER** *University of Manchester*

**PAUL WARREN** *University of Manchester*

**Wael EL-DEREDY** *University of Manchester*

**ANDREW HOWES** *University of Birmingham*

The attraction effect (Huber, Payne & Puto, 1982) occurs when a choice between two alternatives is biased by the addition of a third irrelevant alternative. In a choice between a 70% chance of £20 (the competitor) and a 35% chance of £40 (the target), the addition of an alternative in which there is a 33% chance of £36 (the decoy) is likely to bias some people toward the target £40 prospect. This demonstration of context sensitivity in human decision-making is typically characterised as irrational. This is because rational value-maximising models of decision making show that the best alternative can only be guaranteed to be chosen if the available alternatives are evaluated independently of one another (Luce, 1959).

With a mathematical analysis of computationally rational choice we have shown that when there is uncertainty about the expected value of each alternative, the third decoy option is not, in fact, irrelevant. This is because the dominance relations provide information about the likelihood that either of the original (target or competitor) options will be best. A consideration of the rank order of choice probabilities and values leads to the attraction effect and higher expected value than would be achieved if the "irrelevant" alternative were ignored. The ordinal relations in a two alternative choice set where one alternative has higher probability, and the other has higher value imply approximately equal expected values. However, the ordinal relations necessary to elicit the attraction effect imply the target has greater expected value than the competitor.

The increase in expected value that results from selecting the target may be attenuated in two ways. First is simply to increase the difference in expected value between the target and competitor prospects. Second is to increase the accuracy which the decision maker perceives the expected value of each alternative. We report the results of an experiment in which we test this model using two different paradigms, participants either stated which prospect they preferred or which rectangle had the larger area. For prospects, we systematically manipulated the difference in expected value, while for rectangles, we manipulated the difference in area. As the differences increased so the attraction effect decreased. The increased ease of the area task relative to the prospect task resulted in a reduced attraction effect and a greater sensitivity to differences in area. The tendency for people to exhibit the attraction effect is strongest when expected value rank ordering is most uncertain, that is precisely when the effect is most likely to be rational.

**Keywords:** Attraction Effect; Preference Reversal; Expected Value

# KEEPING TRACK OF OTHERS' CHOICES: PREDICTING CHANGES IN THE PERCEPTUAL DECISIONS AND DECISION CONFIDENCE OF OTHER INDIVIDUALS

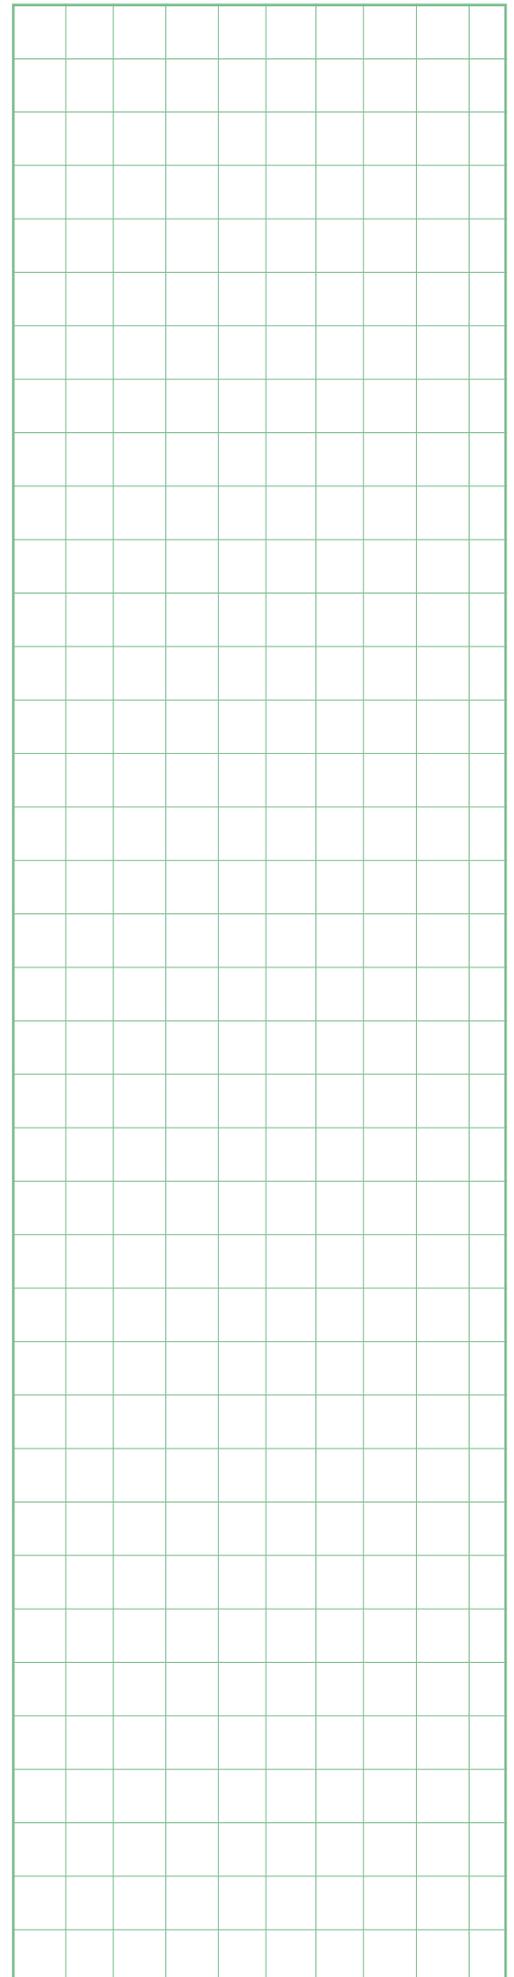
**KARSTEN OLSEN** Aarhus University; Aarhus University Hospital

Humans rarely make decisions in social isolation. Rather, we often make decisions in the context of other individuals, either via observation or in interaction with them. Previous research has indicated that through confidence sharing we can optimally integrate information from others into our decisions (Bahrami, Olsen et al, 2010), and, over time, build up a subjective estimate of other's performance even without access to objective feedback (Bahrami, Olsen et al., 2011; Olsen et al, in preparation). In the present study, we investigated how individuals learn about others' performance and their decision confidence in a two-person perceptual decision-making task.

We staircased one subject (the observer) and asked them to identify a Gabor patch target (and rate their decision confidence) in an orientation discrimination task. The other subject (the mentaliser) was shown the upcoming target and asked to guess on each trial the upcoming decision of the first subject and then rate their confidence in this guess. No interaction was permitted at this stage. Using a signal detection theoretic approach, we subtly manipulated difficulty and pay-off contingency (discriminability and bias) in the first subject's (the observer) task.

Results suggest that individuals can learn and track simple perceptual behaviour across changes in stimulus conditions. That is, individuals could reliably pick up on the underlying discriminability parameter behind their partner's perceptual decision and their decision confidence. In addition, we found that the 'mentaliser' systematically underestimated the 'observers' behaviour. These results, and inter-individual differences, are discussed in light of the role of social interaction and information sharing.

**Keywords:** observational learning, social interaction, collective decision-making, decision confidence



[ 1 ]

INFORMATION PRESENTATION FORMAT  
MODERATES THE UNCONSCIOUS-THOUGHT  
EFFECT: THE UNDERLYING MEMORY  
PROCESSES

**MARLÈNE ABADIE** *University of Toulouse*  
**LAURENT WAROQUIER** *University of Sussex*  
**PATRICE TERRIER** *University of Toulouse*

The unconscious-thought effect occurs when distraction improves complex decision making. In three experiments, we investigated the effect of a detailed or a global presentation format of decision information (i) on the quality of decisions made after distraction, conscious deliberation or immediately and (ii) on memory for decision-relevant information. We used the process-dissociation procedure to measure recollection and familiarity (Experiment 1 and 2) and the simplified conjoint-recognition paradigm to dissociate verbatim and gist representations (Experiment 3). Conscious deliberation resulted in better decisions when the format was detailed whereas distraction improved decision quality when the format was global. A detailed format allowed participants to retrieve precise memories as shown by an increase in recollection and verbatim memory. Gist memory increased after distraction when a global format was used. This suggests that conscious deliberation efficiency is dependent upon the availability of precise memories whereas the unconscious-thought effect is accompanied by enhanced gist memory.

**Keywords:** Conscious deliberation; Unconscious thought; Distraction; Presentation format; Decision making; Dual-memory processes

[ 2 ]

THE RELATIONSHIP BETWEEN MODELLING  
DISCRETE CHOICE BEHAVIOUR IN  
PSYCHOLOGY, ECONOMICS AND MACHINE  
LEARNING

**SOPHIE MEAKIN** *University of Bristol*

Decision-making is an intrinsic part of everyday life, and so being able to mathematically model such behaviour is of interest in a range of areas, including psychology, economics and machine learning. One approach to modelling discrete choice behaviour in psychology is based on Luce's Choice Axiom, from which we can derive the Independence from Irrelevant Alternatives property. The family of discrete choice models used in economics are known as random utility models. Thompson sampling and Boltzmann exploration are two heuristics for sequential decision making used in machine learning that trade-off between exploration and exploitation. Despite having dramatically different backgrounds, we can find links between each of the approaches, which demonstrate that they are in fact more similar than we first expect.

[ 3 ]

THE EFFECT OF OUTCOME SEVERITY AND  
BASE RATE ON THE INTERPRETATION OF  
LINGUISTIC PROBABILITY  
EXPRESSIONS

**EGLÉ BUTT** *Kingston University, London*  
**DR GAELLE VALLEE-TOURANGEAU** *Kingston University*  
**DR MARIE JUANCHICH** *Kingston University*  
**PROF FREDERIC VALLEE-TOURANGEAU** *Kingston University*

This study explored the effect of severity and base rate on the attributed meaning of probability in the light of politeness theory and examined how these factors may influence numerical estimates of a verbal probability phrase in a legal context where a lawyer communicates uncertainty about a client's imminent conviction. We expected that a frequent conviction rate and a severe court sentence would lead to increased preferences for face-management interpretations of the probability term possible uttered by the lawyer, whereas a low conviction rate and a mild court sentence would lead to a likelihood interpretation. We also expected that a frequent conviction rate and a severe court sentence would lead to higher probability estimates of the client's guilt. Data showed that base rate directly influenced the attributed meaning of possible as well as its numerical interpretation. Participants preferred the hearer-face-management interpretation and ascribed higher numerical estimates in high base rate conditions, whereas they preferred the likelihood communication interpretation and ascribed lower numerical estimates in low base rate conditions. Outcome severity led to an increased preference for likelihood communication and had no effect on numerical estimates of a verbal probability phrase. These findings suggest base rate is a strong determinant of both linguistic and numerical interpretations of the probability phrase 'possible' and might override a severity bias. We discuss methodological implications for future research on outcome severity.

**Keywords:** outcome severity, base rate, verbal probability, politeness theory, face-management devices.

[ 4 ]

EXPLORING PSYCHOLOGICAL BIASES OF  
INDIVIDUAL INVESTORS IN SLOVENIA

**TOMAŽ KOLAR** *University of Ljubljana*  
**PETRA DEBENEC** *University of Ljubljana*

Economy and finance heavily relies on assumption of rationality of economic subjects. Rich body of theoretical and empirical arguments however attests that at financial decisions rationality assumption is often and systematically breached. According to De Bondt (1998) and other authors individual investors are prone to the various decision-making and judgment errors; they "discover" naive patterns in past price movements, share popular models of value, investments are not properly diversified and trade in suboptimal ways.

[5]

## PRODUCING THE BEANS NEEDED FOR PERSON-CENTRED HEALTHCARE DECISION MAKING REQUIRES TRANSLATING THE WISDOM OF THE CLINICAL CROWD

**JACK DOWIE** *University of Southern Denmark*

**METTE KJER KALTOFT** *Norwegian Knowledge Centre for the Health Services*

**ØYSTEIN EIRING** *University of Southern Denmark*

**JESPER BO NIELSEN** *University of Sydney*

**GLENN SALKELD** *London School of Hygiene & Tropical Medicine*

Accordingly, one could expect that unexperienced investors from transitional markets like Slovenia are even less "rational" and more prone to various decision errors. Consistent with Prospect theory people in general regularly follow mental shortcuts (heuristics) and are prone to various biases. In order to examine some of these biases (i.e. ambiguity aversion, framing, anchoring and mental accountancy) an empirical study was carried out on a convenience sample of 140 individual, non-professional investors in Slovenia. Obtained results attest that with certain caution due to sample size and structure, we can conclude that Slovene investors are susceptible for three of the examined biases. We found that respondents are indicating aversion to ambiguity, since 73.6% of respondents prefer the "known ratio" choice. They are also prone to framing, since positively framed choice was chosen by 78.6% respondents. In addition they use different mental accountancies, as indicated by different structure of investments decisions for "easy-" vs. "regularly-earned" money. In regard to anchoring bias obtained results are inconclusive, since respondents evaded estimation of multiplication result (and calculated it). Such response suggests, that respondents prefer certain outcomes over uncertain and in fact tend to avoid judgement error (!). This finding is in line with the ambiguity avoidance bias and suggests that when investors have choice they are trying to operate with known, certain outcomes. In investment uncertainty is however not an option, but rather "certain sacrifice for uncertain benefit". The question is therefore how and to which extent investors can reduce and cope with market uncertainty (and when this tendency becomes biased), especially when it comes to private and inexperienced investors. The question of strength of bias effect is also relevant at the mental accountancy bias, which for instance was not found to be present in strong (extreme) form. Respondents namely did not decide to "lavishly spend" easy earned money entirely, but also invested it into real estate and other long term and stable investment options. The question of how investors can further reduce their mental accountancy bias and properly diversify their investments, however remains a challenging task. In this respect coping with the framing bias seems a bit simpler mission, as investors might be warned and trained in order to resist "framed" promotional descriptions of investment options (e.g. when only positive / negative attributes are emphasized). In any case biases reduce quality of investment decisions and proper education of investor is warranted, since various studies (e.g. Ginyard, 2001; Dorn & Huberman, 2005) attest that knowledge and experiences affect quality of investment decisions.

**Keywords:** Behavioral economics, Heuristics, Psychological biases, Investors, Slovenia

Person-centred care is the increasingly avowed aim of health services and professionals. To be meaningful such care requires a shared decision making process in which an individual's preferences over the multiple criteria that matter to them are synthesised with the Best Estimates Available Now (BEANs), at the point of decision, for how well each of the available options will perform on each criterion. Conventional evidence-based approaches can meet the latter requirement in relation to the required BEANs only by assuming professionals are able to make up the shortfalls remaining after the peer-reviewed published products of scientific research have been fully exploited.

Since the clinical judgement of individual professionals has never been subjected to scientific validation in this respect, we have a situation where demonstrated scientific rigour is simultaneously regarded as essential and irrelevant to clinical decision making. Attempts to increase the external validity of scientific studies (notably randomised controlled trials) are attractive to many, but can never get near to providing most of the BEANs needed in person-centred practice.

What is required is the translation of the wisdom of the clinical crowd through the systematic processing of the beliefs of expert professionals. Given the variations in clinician's beliefs and limited statistical competencies (as established by Gigerenzer and colleagues), these beliefs require analytically rigorous processing. Needed therefore is the reformulation of the research community's task as being the continuous ('living') production of BEANs, within a Bayesian framework and by a process that is ARAPAN - As Rigorous As Practical And Necessary. This will involve the systematic elicitation and analysis of expert beliefs, as well as the exploitation of observational studies and big data included in health records and elsewhere. We explore the use of Expertise-based Network Meta-Analyses in this context.

Analytic rigour should not be confounded with scientific rigour. A binary concept of 'evidence', as that which meets, or doesn't meet, some general, value judgement-based threshold (e.g.  $p < .05$ ), is inappropriate for individual decision makers, who have their own trade-offs and error loss functions. It creates much of the 'know-do' gap and perceived 'translation' problem at the final 'bedside' stage.

To give substance to this argument we inserted the results from a recent high-quality Mixed Treatment Comparison (Network Meta-Analysis) on medications for Generalised Anxiety Disorder into a Multi-Criteria Decision Analysis (<http://healthdecisions.org.au/respond/mental>). The criteria included (Response, Remission, and Tolerability) were

determined by data availability in trials, not by a survey of person-as-patient important outcomes. Minimally the latter would distinguish different types of side effects and adverse events, given the known existence of, and differential concern with, effects on sexual, and other functions. Working to the standards appropriate to practice as opposed to science, and simultaneously increasing the range of source inputs to include expert beliefs, is essential to give such modelling research practical relevance for person-centred care. Major benefits in establishing priorities for person-centred research will follow by way of 'backward translation' of the need for better BEANs for many person-as-patient important outcomes.

**Keywords:** person-centred care; decision making; expert beliefs; Network Meta-Analysis

[6]

### AN INVESTIGATION OF THE EFFECTS OF VISUAL ORGANISATION OF FLOOD-RISK COMMUNICATION MEDIUMS ON RISK INTERPRETATION

**BARNEY DOBSON** *University of Bristol*  
**JOLYON MILES-WILSON** *University of Bristol*  
**DAVID LESLIE** *University of Bristol*  
**THORSTEN WAGENER** *University of Bristol*  
**JONATHON ROUGIER** *University of Bristol*  
**IAIN GILCHRIST** *University of Bristol*

The need for effective public awareness of flood hazards is becoming increasingly evident, especially in light of the ongoing environmental consequences associated with climate change. Researchers and governmental bodies currently emphasise the necessity of publicly available flood risk information with which members of the public can base appropriate flood mitigation strategies, but the success of these communication efforts appears to vary. Much of this variation has been observed to be at least in part due to individual differences such as previous experience with flood hazards and socio-demographic variables such as homeownership and level of income. Relatively less attention has been directed toward how the visual organisation of communication mediums affects the viewer's interpretations and inferences, though there is some evidence that such factors are likely to be influential. Research in the cognitive sciences over the past 20 years has provided several theoretical models for how visual displays are perceived, analysed, and decoded by the human visual system, and it is highly likely that these processes influence interpretation in the context of risk communication. With this in mind, we investigated how the format of presentation flood-risk information influences viewers' decision-making and risk-taking behaviour. Participants were presented with pairs of houses and asked to decide which house they would prefer to buy. Information for each house - such as price, energy performance, etc. - was presented alongside flood-risk information to encourage participants to consider various aspects of the house. Within and between participants, we compared the flood hazard map form of communication currently used by the UK Environment Agency with two newly-devised formats; a table that presents a combination of

flood depth and flood frequency, and a graphic that depicts these depth-frequency combinations with reference to a cartoon house. The results indicated that participants were more tolerant of high probability flood-risk houses when this information was presented in map format, as compared to the table and graphic formats. Put another way, higher-risk houses were more frequently rejected when the flood-risk information was presented in the table and graphic formats. These results confirm the suggestion that the way in which flood-risk information is presented affects how acceptable the risk is, not only between participants but also for the same individual in making a series of decisions. The finding that flood hazard maps encourage greater risk acceptance than other types of risk presentation is potentially concerning in light of the fact that this method of communication is currently prescribed by the 2007 Flood Directive of European Parliament. The finding that this variation exists even within the individual is particularly important considering the emphasis placed on differences between individuals in the risk communication literature. It may therefore be useful for communicators to direct attention toward both the method of risk presentation in addition to the intended audience when developing flood-risk communication mediums. These early-stage findings highlight a need for greater consideration of presentation format in flood-risk communication, and further research in this area will hopefully contribute to more effective risk communication in the near future.

**Keywords:** flood, risk, presentation, communication, visual organisation, format, maps, environment, public perception, prevention, awareness, psychology,

[7]

### FOUR THINKING SYSTEMS, INSTEAD OF TWO: A NEW THEORY OF COGNITIVE PROCESSING SYSTEMS

**DR JOHN GOUNTAS** *Murdoch University*  
**DR JOSEPH CIORCIARI** *Swinburne University of Technology*

The dual cognitive processing model has a long tradition and is considered to be robust by a number of researchers (Evans and Over, 1996; Sloman, 1996, 2002; Stanovich, 1999; Kahneman, 2002; Evans, 2012). Stanovich and West (2000) summarised the two broad modes of thinking as system 1 (fast, holistic, heuristic, intuitive, emotional, experiential, implicit, automatic), and system 2 (slow, rational, analytical, reflective, sequential, explicit). However this conceptualisation has not been universally accepted as a reliable theory. A number of researchers have criticised the putative dual cognitive processing systems for its lack of conceptual coherence (Osman, 2004; Keren and Schul 2009). Kruglanski and Gigerenzer (2011) have expressed reservations regarding the validity and reliability of the current dual processing model and proposed that one system is possible, using a unified rule of ecological rationality for both intuitive and deliberate judgments. Sherman (2006) suggests that there are severe limitations with the dual system and postulates that a quad model is potentially more useful. Cools and Van Den Broeck (2007) suggest that there three cognitive styles which are an improvement on the dual system of analytic-intuitive model.

The main weakness of the dual cognitive systems theory is the diversity (lack of consistency) of definitions and disparate attributes to describe each system. Conceptual incoherence of what constitutes system 1 and 2 is evident in the literature because of the inclusion of incompatible attributes in systems 1 & 2, which are qualitatively different. The lack of valid and reliable measures to identify/measure empirically each system's attributes creates even more uncertainty about data interpretation and theoretical coherence. Simon's (1999) idea of invariance in physical phenomena and human attributes implies that there is probably invariance in human cognitive processing systems and decision making processes. The dual process systems model ignores the research literature on learning styles and personality differences. The research literature on personality and learning styles suggests that there is more variability in cognitive processing systems than the current dual model permits. This paper integrates personality theories and learning styles and proposes that there are four cognitive systems which use qualitatively different types of information processing systems. The four putative thinking systems are Logical-Ideational, Pragmatic-Physical, Emotion-Action, and Imaginative-Visualising. All four thinking systems are hypothesised to be representative of distinctive brain processing systems and therefore there are real physiological networks for each mode of information processing. The paper reports three separate studies. The first study reports the development of the psychometric constructs and scales for each of the four cognitive systems; the second study reports the findings of the EEG brain scans, which verify the different brain systems; and the third study report the fMRI findings which identify more accurately to BOLD activation for each of the four cognitive processing systems. All three studies support the hypothesis that people think and process information in distinctively different modes and there is neuroscientific evidence to support the existence of separate brain functions for each thinking system. The empirical survey, EEG and fMRI findings provide robust evidence of scale validity and reliability. Implications for decision preferences and behavioural applications are discussed.

**Keywords:** cognitive processing systems, thinking styles, dual processing, scale validation neuroscience testing.

[8]

## IMPROVING FOOD CHOICES THROUGH RANK-BASED SOCIAL NORMS NUDGING

**SILVIO ALDROVANDI** *Birmingham City University*  
**GORDON D.A. BROWN** *University of Warwick*  
**ALEX M. WOOD** *University of Stirling*

Despite extensive nutrition information campaigns and food labelling policies (e.g., Burton, Creyer, Kees, & Huggins, 2006), the prevalence of harmful effects arising from poor dietary choices is on the rise (e.g., Strazzullo, D'Elia, Kandala, & Cappuccio, 2009). The increase in the proportion of people who are either overweight or obese (e.g., Bassett & Perl, 2004) has been mostly associated with food overconsumption, over and above unhealthier life style (e.g., Chandon & Wansink, 2012). It has been often suggested that poor food choices and food overconsumption might be rooted in cognitive biases that undermine accurate

evaluations of foods and their healthiness (e.g., Wansink, Just, & Payne, 2009). In an effort to determine the cognitive mechanisms behind these biases, we assess the predictions of rank-based (e.g., Stewart, Chater, & Brown, 2006) and reference-level (e.g., Helson, 1964) models of judgment and decision making in the domain of food evaluation and choice. We also test whether the principles embodied in the models can be used to inform and improve consumer's dietary choices through social norm interventions that may reduce the preferences for relatively unhealthy foods. Recently behavioural 'nudging', a method for guiding and influencing consumer's behaviour by shaping the environment but without unduly restricting their freedom of choice (cf. Thaler & Sunstein, 2008), has been implemented within the social norms framework (e.g., Agostinelli, Brown, & Miller, 1995). Indeed, many behavioural 'nudges' provide social norm information and have been found to induce behaviour change in a variety of contexts such as energy consumption (e.g., Schultz, Nolan, Cialdini, Goldstein, & Griskevicius, 2007) and recycling (e.g., Goldstein, Cialdini, & Griskevicius, 2008). Generally, behavioural nudges based on changing perceptions of social norms often provide information about mean levels of behaviour (e.g., average energy consumption) in a social comparison group. However, recent cognitive models of judgment suggest that provision of rank-based information (e.g., percentage of other people consuming less energy) might tap more directly into natural ways of processing information. Studies 1 and 2 find that people's evaluations are highly sensitive to the options available in the decision-making context in ways consistent with rank-based models of judgment and decision-making. We also find that people's concern about their own consumption of unhealthy products is predicted not by the quantity of consumption, but by how this quantity is believed (inaccurately) to rank within a social comparison distribution (Study 3). In Study 4 we find that rank-based social norm feedback - telling people how they rank within a normative comparative sample - increases willingness to pay for a healthy food by over 30%, with greater effects for participants who most underestimate their own ranked position. Mean-based social norm information (i.e. telling people how they differ from the average person), in contrast, had no effect. We conclude that people's attitudes towards food consumption are determined by inaccurate beliefs about how much other people consume and that rank-based social norms feedback is most effective in addressing people's inaccurate beliefs and can stimulate healthy food purchasing.

**Keywords:** Food perception; Concern about food consumption; Social norms; Decision by Sampling; Context effects; Healthy eating

[9]

## GROUP FORECASTING AND THE QUALITY OF EQUITY PREDICTIONS IN ONLINE COMMUNITIES

**TOBIAS ENDRESS** *University of Gloucestershire*  
**TONY GEAR** *University of Gloucestershire*

This paper presents the results of a study of the quality of equity predictions in online communities. The study reviews secondary data from several special interest communities (Stockjaeger, Sharewise, and Spekonauten) that focus on stock market topics and create group predictions using group decision-making approaches. Additionally, primary data was generated through purposefully selected groups in a controlled environment. The experiments consist of a pilot and a main run. The groups were benchmarked with actual market prices as well as with each other. The data was collected over a period of 10 weeks. The objective of the research project as a whole is to assess the quality of the predictions in terms of accuracy and clarity. One of the main challenges in the online decision-making process seems to be that there are often just a few predictions (if any) for each stock, potentially outdated predictions, and limited activity of community members. There might be some moderation and/or user guidance needed to address these issues and to gain a critical mass. Another issue is due to the design of the communities. The occasional underperformance of online groups might be influenced by a conforming effect and community members' tendency to make overly optimistic recommendations and preferences for high beta shares (i.e., riskier shares). Such a speculation is consistent with the results from the study in which the recommendations of the existing communities turned out to be rather optimistic, as well. Through the application of well-known best practices from the literature, it might be possible to improve the accuracy of group predictions. Besides identifying possible areas for improvement, the study demonstrated that online groups might be able to deliver predictions of similar quality to financial experts. The experts represented in the Bloomberg consensus did not provide price targets with a higher accuracy than the online Sharewise group. While the study has provided some indications that in certain situations and with careful group design, stock price predictions can be superior to the predictions of experts, the main experiment indicates a more differentiated picture and provided valuable information about the underlying decision-making process.

**Keywords:** online community, group decision-making, collective intelligence, equity predictions, stock trading

[10]

## SHAME AND ANGER DRIVE REPAIR AND RESTORATION IN THE FACE OF RETALIATION

**ANYA SKATOVA** *University of Nottingham*  
**CAROLINE LEYGUE** *University of Nottingham*  
**ALEXA SPENCE** *University of Nottingham*  
**EAMONN FERGUSON** *University of Nottingham*

Cooperation is central to an effective and fair society. Common evolutionary mechanism like kinship and reciprocity cannot fully explain cooperation with strangers. Theoretically, emotions could play an important role in helping explain cooperation between strangers in everyday life. Anger has been suggested as a primary mechanism that sustains cooperation, specifically when it is possible to act on the angry impulse to punish free rider. Emotions of shame and guilt should lead to repair of damaged social interactions, and should turn the free rider "to be good to others". However, the precise consequences of emotions in cooperative situations when sanctioning mechanisms are not present are unknown.

In our experiment participants (N=118) took part in two blocks of inverted repeated-interactions Public Goods Games, with ten games in each block. After each game participants reported emotions they felt at that moment on Likert scales, including anger, guilt, shame. We analysed how decisions of other group members affected people's emotional reactions and behaviour. We demonstrated detrimental effects of anger towards free riding on the outcomes of interactions in social dilemmas. The consequences of anger were not counteracted by the effects of shame and guilt, experienced by individuals themselves. We found that strategies that showed the highest levels of free riding and caused the collapse of cooperation were not associated with shame or guilt after benefiting at the expense of others. Further, these profit-seeking strategies retaliated more frequently than others and their retaliation was fuelled by anger. We found, however, that such behaviour did not secure better profits. Instead, strategies that withheld the anger impulse and retaliation towards free riders made better profits through sustaining cooperation in the group.

Our results in the line with previous findings, demonstrating that in a repeated-interaction scenario retaliation and non-cooperation does not secure better payoffs. Further, we show that, at least in no sanctions scenarios, acting upon one's anger is detrimental to social interactions as it leads to further escalation of retaliation. The effects are enhanced by strategies that seem to play a zero-sum individual profit making game. They are angry when their profits suffer and subsequently they retaliate. In addition, they do not show concern for the public good or for being a free rider (no shame or guilt). On the other hand, those who signal their prosociality to the group from the start (by using little in the first round) and those who show concern for others (by feeling guilt and shame when they happened to benefit at expense of others), repair more often and use more lenient, forgiving strategies. However, our data is at odds with previous assertions that these individuals are not angry at free riders. On contrary, we demonstrated both strategies were associated with anger after being treated unfairly, however not everybody retaliated in response. The strategy

of not acting upon one's anger, non-retaliation and repair seem to be more profitable in the long run. We showed that during repeated interactions around public good it pays off to control one's anger and turn the other cheek.

[11]

## A "VEIL OF IGNORANCE" CAN ACHIEVE DISTRIBUTIVE JUSTICE IN MULTIPLAYER SHARING GAMES

**JAN K. WOIKE** *Max Planck Institute for Human Development*

**SEBASTIAN HAFENBRÄDL** *University of Lausanne*

Faced with a decision that influences yourself and others, it has been argued that your choice might heavily depend on the fact whether you know your own position as a recipient or not. While Rawls argues that decision makers behind a "veil of ignorance" will adhere to the maximin-principle, Harsanyi proposes the maximization of expected value (and total social welfare), instead.

We conducted a series of three-player money-sharing games, in which participants had to choose between different monetary distributions for three receiving positions. The distributions varied systematically in terms of achieved equality, total social welfare, and expected personal profit. In study 1 (N=154), one group was asked to choose the fairest distribution in hypothetical games, while other participants made choices for other players and incentivized choices under different degrees of ignorance concerning their receiver position.

When participants were disaffected by their choices or all three positions were equi-probable, choices were mostly aligned with choices in the hypothetical fairness condition: Distributions maximizing equality of outcomes were preferred over distributions maximizing average payoffs or the outcome for specific positions. When participants knew their receiver position for certain, the majority of participants chose to maximize their personal outcome. Finally, when participants were informed that they would be assigned one out of two of the three positions (with a chance of  $p=3/4$  for one of the two), choices were split into three clusters: some participants chose equal distributions, some chose to maximize the minimum payment for the two receivers and a third group maximized their expected outcome. Further analyses include reaction-time and qualitative data, and behavior in other economic games.

In Study 2 (N=111), we tested participants with sets of choices, where the equally distributed option was dominated by all other options. Indeed, some participants chose a guaranteed loss of CHF 1 for everyone over an average outcome of CHF 5 with an unequal split.

In summary, some variants of the veil of ignorance are able to achieve distributive outcomes that are considered fair by neutral observers. While social welfare maximization is not pursued by a large number of subjects, a substantial percentage chooses according to the maxi-min principle proposed by Rawls. Also, a non-negligible percentage of participants are willing to incur both personal and social-welfare costs to achieve equality.

**Keywords:** veil of ignorance, distributive fairness, group decision-making

[12]

## MODELING JUDGMENT AND DECISION MAKING PROCESSES AT THE SYSTEM LEVEL

**JERYL L. MUMPOWER** *Texas A&M University*

**GARY H. MCCLELLAND** *University of Colorado Boulder*

The distinction between idiographic and nomothetic approaches in psychology is widely attributed to Allport (1937). Idiographic research in JDM generally involves intensive study of a single individual to achieve an understanding of that particular individual. For example, Gonzalez-Vallejo et. al (1998) developed separate models for 32 individual physicians' diagnostic judgments and treatment decisions for acute otitis media in children. Such models permit predictions about how individuals will behave with respect to comparable cases in the future. Comparisons and contrasts can be made among individuals or clusters of individuals regarding their judgment processes or policies.

Nomothetic research in JDM generally involves investigations of large groups of people to find general laws of behavior that hold for people in general. An example is Tversky and Kahneman's (1981) studies of the effects of framing on preference reversal using the famous so-called "Asian disease" problem. The various cognitive heuristics and biases identified by Tversky, Kahneman, and others are generally presumed to hold to some degree for all persons.

Typically, idiographic research relies on a single individual-many cases paradigm, whereas nomothetic research relies on a many individuals-one case paradigm. These are not the only design possibilities for judgment and decision making research, however.

A third possibility is the single decision maker--single case, or case study, paradigm. A well-known example of the case study approach in JDM is Allison's (1971) study of decision making within the U.S. government during the Cuban Missile Crisis. Case studies are similar to nomothetic approaches in so far as they focus on particular decision makers, but the validity of generalizing to future cases is unclear since these studies are based on an n of one.

A fourth variant – and the focus of this paper – is the many individuals-many cases paradigm which yields descriptions or models of the judgment and decision making behaviors of systems. For instance, Mumpower and McClelland (2014) analyzed decision making processes during the referral and substantiation stages of the child welfare system. The data for the study represented the aggregation of literally millions of decisions by individuals across literally millions of cases. Analyzing those data with JDM techniques yields "as-if" models of system behavior that do not apply either to any individual decision maker or to "people in general". In Mumpower and McClelland (2014), the systemic, many individuals-many cases, approach uncovered quite different patterns of JDM behavior with respect to the treatment of different racial and ethnic groups in the referral and substantiation portions of the child welfare services system.

In this paper, we review and discuss the Mumpower and McClelland study, as well as others among the relatively small number of JDM studies that employ the systemic, many individuals-many cases paradigm. We conclude that the use of JDM methods to analyze system behaviors has great potential for improving our understanding of those systems from both fundamental and applied perspectives.

**Keywords:** system models, idiographic, nomothetic, case studies

[13]

### A PURELY CONFIRMATORY REPLICATION STUDY OF STRUCTURAL BRAIN-BEHAVIOR CORRELATIONS

**WOUTER BOEKEL** *University of Amsterdam*  
**ERIC-JAN WAGENMAKERS** *University of Amsterdam*  
**LUAM BELAY** *University of Amsterdam*  
**JOSINE VERHAGEN** *University of Amsterdam*  
**SCOTT BROWN** *University of Newcastle*  
**BIRTE FORSTMANN** *University of Amsterdam*

A recent 'crisis of confidence' has emerged in the empirical sciences (Pashler & Wagenmakers, 2012). Recent studies have suggested that questionable research practices (QRPs; Simmons, Nelson, & Simonsohn, 2011) such as post-hoc storytelling and non-publication may be common in the field of psychology (John, Loewenstein, & Prelec, 2012). These QRPs can result in a high amount of false-positive findings, decreasing reliability and replicability of research findings. A potential solution is to preregister experiments prior to acquisition or analysis of data (Wagenmakers, Wetzels, Borsboom, van der Maas, & Kievit, 2012; Chambers, 2013). In this study we attempt to replicate studies that employ the commonly used approach of relating brain structure to behavior and cognition. These structural brain-behavior correlations (SBBCs) occasionally receive great scientific and media attention. Given the impact of these studies, it is important to investigate their replicability. Here, we attempt to replicate five SBBC studies comprising a total of 17 effects. To prevent QRPs such as post-hoc storytelling and non-publication, we employed a preregistered, purely confirmatory replication approach. We were unable to successfully replicate any of the 17 effects. For all of the 17 findings under scrutiny, Bayesian hypothesis tests indicated evidence in favor of the null hypothesis. The extent of this support ranged from anecdotal ( $BF < 3$ ) to strong ( $BF > 10$ ). To our knowledge, this is the first multi-study confirmatory replication of SBBCs. With this study, we hope to encourage other researchers to undertake similar replication attempts.

**Keywords:** preregistration; confirmatory; replication; brain-behavior correlations

[14]

### THE RELIABILITY AND INTERCHANGEABILITY OF RESPONSE CONTROL TASKS

**CRAIG HEDGE** *Cardiff University*  
**PETROC SUMNER** *Cardiff University*

Behavioural tasks which measure an individual's ability to inhibit unwanted response are used prominently in the literature to examine the mechanisms underlying cognitive control, decision making processes, and in the examination of individual differences and psychopathology. The interchangeable use of different tasks, and in some cases of different performance measures from the same task, entails several critical assumptions: First, that the different measures

tap in to the same underlying mechanisms; and second, that these measures reflect an underlying trait of the individual that is stable over time. Despite the importance of these assumptions, little evidence exists to support them for many widely-used tasks. Measures of response control from different tasks often show modest or inconsistent correlations, and retest reliability across multiple testing sessions is not commonly examined.

Our study seeks to address this gap in the literature, and provide guidance on the optimal use of these tasks. We examined four common measures of response control; the Stroop task, the Eriksen flanker task, the Go/No-go task, the Stop-Signal task. In addition, participants completed a self-report measure of impulsive behaviour. Forty-seven participants completed the tasks in two sessions, taking place three weeks apart. We report data on several key points. First, does each of these tasks provide an assessment of a trait that is stable over time, and second, how many trials are required for this to be achieved? Third, if the stability of these individual measures can be verified, is there then a relationship within (e.g., comparing reaction time and error rates from the same task) and between (e.g., Stroop reaction time cost and Flanker reaction time cost) the different response control measures?

Retest reliability for the various measures ranged from fair to good, and we use a method of subsampling different trial numbers to illustrate the number of trials required to achieve a given level of reliability. However, even with higher trial numbers, the correlations between the different response control measures within and between tasks ranged from low to moderate, prompting caution in assuming that such measures can be used interchangeably. We discuss the implications of these findings for the use of these tasks in both theoretical and applied contexts.

**Keywords:** Response inhibition, Cognitive control, reliability, impulsivity

[15]

### MAGNITUDE VERSUS TIME: HOW THE COMMUNICATION OF UNCERTAINTY INFLUENCES CLIMATE CHANGE RISK PERCEPTIONS

**TIMOTHY BALLARD** *University of Bristol*  
**STEPHAN LEWANDOWSKY** *University of Bristol; University of Western Australia*  
**MARK HURLSTONE** *University of Western Australia*

Uncertainty is an inherent feature of any climate change model. From a normative perspective, greater uncertainty increases the risk associated with the impact of climate change. However, uncertainty is frequently cited in political and public discourse as a reason to delay mitigating action. This failure to understand the implications of uncertainty may have considerable costs both for the climate and the global economy. It is therefore critical to communicate uncertainty in a way that better calibrates people's risk perceptions with the projected impact of climate change.

We conducted an experiment that examined whether people's concern about projected climate change outcomes was influenced by the manner in which uncertainty is expressed. Specifically, we examined whether concern

about projected outcomes depended on whether outcomes were expressed as certain in magnitude but arriving at an uncertain time or uncertain in magnitude but arriving at a certain time. We presented participants with a series of statements and graphs indicating projected increases in temperature, sea levels, ocean acidification, and a decrease in arctic sea ice. In the uncertain magnitude condition, the statements and graphs reported the upper and lower confidence bounds of the projected magnitude and the mean projected time of arrival. In the uncertain time of arrival condition, they reported the upper and lower confidence bounds of the projected time of arrival and the mean projected magnitude.

Preliminary analysis suggests that participants in the uncertain magnitude condition reported greater concern about the projected outcomes than participants in the uncertain time of arrival condition. This result suggests that the manner in which uncertainty concerning the impact of climate change is communicated affects people's interpretation of the information. This finding has important implications for effectively communicating the risks associated with climate change to policy makers and the general public.

**Keywords:** Uncertainty, Risk, Judgement, Climate Change

[16]

### HOW AN INSECT COMBINES INFORMATION FOR SPATIAL DECISION MAKING

MATTHEW COLLETT *University of Exeter*

Animals use information from multiple sources in order to navigate between goals. Ants such as *Cataglyphis fortis* use an odometer and a sun-based compass to provide input for path integration (PI). They also use configurations of visual features to learn both goal locations and habitual routes to the goals. Information is not combined into a unified representation (Collett, et al, 1998; Collett, et al, 2003; Wehner et al, 2006; Collett & Collett, 2009), but appears to be exploited by separate expert guidance systems (Collett, 2010; Cruse & Wehner, 2011). Visual and PI goal memories are acquired rapidly and provide the consistency for route memories to be formed. Do established route memories then suppress the guidance from PI? A series of manipulations putting PI and route memories into varying levels of conflict found that ants follow compromise trajectories (Collett, 2012) The guidance systems are therefore active together and share the control of behaviour. A simple model shows that observed patterns of control could arise from a superposition of the output commands from the guidance systems, potentially approximating Bayesian inference (Ma, et al, 2006). These results help show how an insect's relatively simple decision-making can produce navigation that is reliable and efficient and that also adapts to changing demands.

**Keywords:** Insect decision making, navigation, integration

[17]

### MONKEY GLOBOPHOBIA AND THE BALLOON ANALOGUE RISK TASK

DAVE E.W. MALLPRESS *University of Bristol*

ALASDAIR I. HOUSTON *University of Bristol*

BENJAMIN Y. HAYDEN *University of Rochester*

The Balloon Analogue Risk Task (BART, Lejuez et al. 2002) is a recent, but hugely influential and well-validated human behavioural task that correlates with measures of impulsivity, self-control and risk attitude. Subjects performing the task inflate a virtual balloon that increases in reward size with time, but also brings greater risk of bursting, which would result in no reward. The subject must then decide the best moment to cease inflating and accept the reward. Here we will provide a discussion of the strategies that should be employed by an optimal decision-maker as well as the conceptual issues with these personality constructs. Several novel versions of this task were developed and performed on Rhesus macaques to test its worth as an impulsivity measure across species. Our continuous-time sequential choice versions of this task share several features with the natural decision problems faced by animal foragers, particularly in trading off delay and risk with reward magnitude and our results of our experiments show that the monkeys take the balloon reward earlier than they should if behaving optimally. Finally, we will discuss why mathematically equivalent versions of the task do not necessarily elicit the same behaviour.

**Keywords:** Impulsivity, Risk, Balloon Analogue Risk Task, Optimality

[18]

### WHEN'S DADDY COMING HOME? OPTIMAL WAIT TIMES AT THE NEST FOR MONOGAMOUS BREEDING BIRD SPECIES

REBECCA PIKE *University of Bristol*

ALASDAIR HOUSTON *University of Bristol*

JOHN McNAMARA *University of Bristol*

Monogamy is the prevailing mating system in birds and much is to be understood of the functions and mechanisms behind these long term affiliations.

In species where incubation is shared by both parents, the mate must make a decision of how long to wait at the nest for their partner to return from foraging e.g. A breeding Albatross pair.

We modelled a situation with two types of individual in the population; some individuals will return to the nest whereas others will abandon it. We identified the incubating mate's optimal time to wait until abandoning the nest themselves given the initial estimate from the population of the likelihood their partner is to return,  $p$ .

The optimal wait time was found by maximising the incubating mate's payoff function. As  $p$  increases, the optimal waiting time is longer.

We then considered how gaining information about a partner affected the estimation of  $p$ . Differentiating the

optimal payoff function with respect to  $p$  allows us to quantify the value of information.

Gaining knowledge allows refinement of your estimate of  $p$  altering your waiting time, thereby increasing the expected payoff. The value of knowing your partner may explain why some species form long term relationships

**Keywords:** Pair bond, Trust, Value of Information, Statistical Decision Theory

[19]

## REASONING AND DECISION-MAKING OF ACTIVITY LEVEL PROPOSITIONS IN THE LAW USING CHAIN EVENT GRAPHS: HOW TO ADDRESS ASYMMETRIC EVIDENCE

**ANJALI MAZUMDER** *University of Warwick*

**JAMES Q. SMITH** *University of Warwick*

The case assessment and interpretation framework commonly used by forensic scientists in the UK involves a hierarchy of propositions to be addressed in casework. Forensic scientists are often focused on addressing source or sub-source level propositions, e.g. what is the source of the DNA sample or glass fragment. However, there is increased interest in addressing activity level propositions (e.g. did Mr Y break the window) to better assist the courts (without straying outside the bounds of scientific knowledge). The pairing, framing and development of at least two competing propositions must be done with careful consideration of the case circumstances and evidence.

There is a plethora of literature on the use of Bayesian Networks (BNs) for forensic science which expresses the graphical and probabilistic relationship between measured variables. BNs have been particularly useful in providing a graphical representation of the problem, calculating marginal and conditional probabilities of interest, and making inferences particularly addressing lower level propositions. To address activity level propositions, there is a need to account for different plausible explanations of suspect/perpetrator's actions and events as it relates to the evidence. In this talk, we propose the use of another class of graphical models, chain event graphs (CEGs), exploiting event tree structures to depict the unfolding of events as postulated by each side (defence and prosecution) and differing explanations/scenarios. Different scenarios can introduce different sets of relevant information affecting the dependence relationship between variables and symmetry of the structure. CEGs are a flexible class of graphical models which can model the asymmetric story structure directly in its topology. Yet because of its graph modular structure it also inherits many benefits of the BN. A BN can be represented as a symmetric CEG but the BN is not always a rich enough structure to incorporate all obtainable information.

We demonstrate how CEGs can be very useful in addressing activity level propositions and assist the courts by directly supporting the barrister's argument within the topology of a graph, particularly in complex cases. Its structure provides a graphical and probabilistic framework for decision-making regarding the validity of a set of proposed explanations in a systematic way. With the use of case examples involving transfer and persistence and

different evidence types, we further show how CEGs can assist in the careful pairing and development of propositions and analysis of the evidence by addressing the uncertainty and asymmetric unfolding of the events to better assist the courts.

**Keywords:** Bayesian networks, chain event graphs, evidence analysis, asymmetry, uncertainty

[1]

## EXPLORING DOMESTIC FOOD WASTE PRODUCTION AND REDUCTION BEHAVIOUR THROUGH JUDGMENT AND DECISION MAKING THEORY: IMPLICATIONS FOR INTERVENTION

**NICOLA BOWN** *Leeds University Business School*

Globally, we consume far less food than we produce. Estimates suggest that annually up to half of food produced is wasted, around 1.2 to 2 billion metric tons (Institute of Mechanical Engineers, 2013). In the UK, food wasted annually in the home that could have been eaten is estimated at 4.2 million tonnes, worth £12.5 billion (Waste and Resource Action Programme, WRAP, 2013). This is a global issue that is financially costly and environmentally damaging. This paper considers how our understanding of human judgment and decision making (JDM) processes can contribute to understanding and addressing this issue.

Recent UK government-backed campaigns such as Love Food Hate Waste have reported a significant 21% reduction in the amount of useable food that is wasted over the last 5 years (WRAP, 2013). What is unclear is the extent to which this is attributable to the campaign effectiveness, consumers' increasing financial constraints or other factors. Although this topic receives ongoing attention from practitioners and policy makers, relatively little is still known about the drivers of consumers' food-related perceptions or the full range of effective strategies for communicating about food waste. This is due in part to the highly complex nature of consumers' decisions around food, which include cognitive, emotional, cultural, health-related, financial, and moral dimensions (Evans, 2011; 2012). The current paper integrates evidence of food waste production and reduction interventions (WRAP, 2013; Quested et al, in press) with a review of relevant JDM literature, to provide insight into human psychological processes during the food planning, shopping, preparation, consumption and storage cycle.

I discuss JDM concepts relevant to household food planning and how these may contribute to increased waste, including focalism (usually discussed in relation to the planning fallacy, Buehler, Griffin & Peetz, 2010; Kahneman & Tversky, 1979; Lovallo & Kahneman, 2003); inter-temporal preferences (Read, Loewenstein & Kalyanaraman, 1999) and variety seeking (Read & Loewenstein, 1995; Read, Loewenstein, & Rabin, 1999; Simonson, 1990; Simonson & Winer, 1992). We also consider that those responsible for food shopping and preparation may be willing to pay for having a choice of food in order to meet individuals' preferences. In other words, some consumers will accept a level of food waste in the home.

An important belief in JDM literature is that humans demonstrate two types of thinking: System 1 and System 2 (Frederick & Kahneman, 2002; Kahneman, 2011). System 1 is fast, intuitive, heuristic, prone to biases, but often functional, and System 2 is slower, effortful, more cognitively demanding, but generally resulting in accurate judgments. Many recommendations from waste reduction campaigns (e.g., "plan what you will eat", "and buy what you need") require the consumer to rely heavily on System 2 processes. While evidence suggests these are relatively successful strategies in reducing waste generation (WRAP, 2013), we discuss as yet unexplored System 1 strategies which could contribute to

campaigns incorporating elements of both System 1 and System 2, possible to be most effective in achieving sustained behaviour change.

**Keywords:** Food waste, judgment and decision Making, System 1 and System 2, behaviour change

[2]

## INVESTIGATING THE ROLE OF DOPAMINE IN REINFORCEMENT LEARNING AND DECISION MAKING USING *DROSOPHILA* *MELANOGASTER*

**HANNAH JULIENNE** *University of Bristol*

**DAVID LESLIE** *University of Bristol*

**JAMES HODGE** *University of Bristol*

Dopamine is thought to play a crucial role in reinforcement learning, namely through signalling the difference between expected reward and actual reward (Schultz et al., 1998). There is increasing evidence that it may also play a role during decision making, but its precise role here is less clear. Parkinson's disease patients show a range of symptoms resulting from dopamine neuron degeneration, including disrupted performance in some reinforcement learning tasks. Interestingly, patients on medication to increase dopamine levels also show impaired performance in certain tasks (Frank et al., 2004), further indicating a multi-faceted role for this neurotransmitter in the learning and memory process.

The fruit fly *Drosophila melanogaster* is an ideal model organism for dissociating these different roles for dopamine. The fly dopaminergic system serves similar functions as that of mammals, including a crucial role in reinforcement learning. Reinforcement learning in flies can be assessed using an olfactory assay where flies are trained to associate an odour with a shock punishment or sugar reward. A number of fly models of Parkinson's disease have been developed that display selective dopamine neuron degeneration. Finally, a range of genetic tools have been developed that allow us to selectively increase or block neurotransmitter release from specific subsets of dopamine neurons.

We are taking both experimental and theoretical approaches to answering this question. We are testing reinforcement learning and memory in *Drosophila* models of Parkinson's disease and comparing their performance to that of wildtype controls. We are also using thermogenetic tools to facilitate or inhibit dopamine release at specific time points in the learning and decision making processes. We will fit data generated from this to mathematical models of reinforcement learning and decision making. We hope to detect differential effects of dopamine at memory acquisition, consolidation and retrieval.

**Keywords:** dopamine, reinforcement learning, *drosophila*, Parkinson's disease

[3]

### WHEN BIGGER IS NOT STRONGER: AUTO-INHIBITION OF STIMULUS-INDUCED ACTIVITY IN A TOPOGRAPHICAL NETWORK OF SPIKING NEURONS

**GEOFFREY MÉGARDON** *Cardiff University; Aix-Marseille Université*

**CHRISTOPHE TANDONNET** *Aix-Marseille Université; Université de Genève*

**PETROC SUMNER** *Cardiff University*

**ALAIN GUILLAUME** *Aix-Marseille Université; New York University*

Topographically organized structures are ubiquitous in sensory and motor regions on the brain. These neural maps have long been modelled with neural fields that incorporate lateral excitation of neighbouring units and lateral inhibition of remote units (a pattern of lateral interactions often called 'Mexican Hat'). Here our aim was to get further insight into the computational properties of this organization by testing whether and under what circumstances it could result in auto-inhibition, i.e. whether activity induced by some stimulus sizes would inhibit itself. Such an explanation has been suggested for eye movement data where increasing the size of a distractor can decrease its effect (a 'reversal pattern'; Tandonnet et al. *J Vis* 2012). We show here that a one-layer model of spiking neurons (a simplification of layers in superior colliculus, for example) containing only Mexican Hat interactions of moderate inhibition extent is sufficient to show auto-inhibition of stimuli. This auto-inhibition shapes three distinct activity patterns over time and space according to the stimulus' size. We also observe that a similar reversal pattern to that observed by Tandonnet et al. can be obtained in this kind of model, but only with an elliptic-shaped Mexican hat. Interpretations of this elliptic shape are discussed and invite to further investigation. These results suggest that auto-inhibition is a possible phenomenon at the level of a single neuronal map and that it could participate in the bottom-up processing of sensory signals and the interactions of motor plans.

[4]

### PATTERNS OF INFORMATION USE DRIVING COOPERATION AND CONFLICT IN HUMANS

**STEPHEN HEAP** *University of Jyväskylä*

**TAPIO MAPPES** *University of Jyväskylä*

**MIKAEL PUURTINEN** *University of Jyväskylä*

Human social evolution is intricately tied to cooperation and conflict at the levels of the individual and the group. Specifically, humans have a tendency to cooperate with others at personal cost when cooperation improves the performance of their own group relative to other groups. However, the relationship between an individual's coalitional psychology (which forms the proximate basis for social decisions) and the emergence of necessary conditions for cooperation to be ultimately selected for remain unclear. We aimed to experimentally investigate this relationship using a

public goods game, in which participants were repeatedly divided into random groups of four and presented with a social dilemma between costly cooperation with their group and maximising their individual payoff. Additionally, participants could vote on the manner in which their group interacted with a randomly paired opponent group. The choice of between-group interaction determined whether funds were competitively taken or cooperatively shared between groups, or whether groups operated in isolation from one another. We show that groups of randomly assembled individuals spontaneously engage in costly group competition, and that decisions promoting between-group conflict are associated with high levels of within-group cooperation. Furthermore, we found that individuals were more likely to vote for competitive interactions with other groups as the variance in cooperative behaviour within their own group decreased and variance in public good provisions between groups increased. This result conforms to multilevel selection theory, which states that between-group variation in fitness must outweigh within-group variation in order for individually costly, but group beneficial, traits to be selected for. Additionally, these results support the theoretical implication that cooperative actions within groups coincide with competitive interactions between groups. We analysed the decision-making processes of individuals using statistical decision theory to determine how participants utilised information from a dynamic social environment and their own experience in order to modify their behaviour in anticipation of future conditions. We discuss these results in terms of how the nature of information processing corresponds with the ultimate drivers of human social behaviour and organisation in an evolutionary context.

**Keywords:** cooperation, evolution, public-goods game, ecology

[5]

### PILLARS OF JUDGMENT: HOW MEMORY ABILITIES AFFECT PERFORMANCE IN RULE-BASED AND EXEMPLAR-BASED JUDGMENTS

**JANINA A. HOFFMANN** *University of Basel*

**BETTINA VON HELVERSEN** *University of Basel*

**JÖRG RIESKAMP** *University of Basel*

Making accurate judgments is an essential skill in everyday life. However, although the relation of different memory abilities to categorization and judgment processes has been hotly debated, the question is far from resolved. We contribute to the solution by investigating how individual differences in memory abilities affect judgment performance in two tasks that induce rule-based or exemplar-based judgment strategies. In a study with 279 participants, we investigated how working memory, episodic memory, and implicit memory affect judgment accuracy and strategy use. As predicted, participants switched strategies between tasks. Furthermore, structural equation modeling showed that the ability to solve rule-based tasks was predicted by working memory, whereas episodic memory predicted judgment accuracy in the exemplar-based task. Last, the probability of choosing an exemplar-based strategy was related to better episodic memory, but strategy selection was unrelated to working memory capacity. In sum, our results suggest that

different memory abilities are essential for successfully adopting different judgment strategies.

**Keywords:** Judgment; working memory; episodic memory; rule-based and exemplar-based processes

[6]

## SINGLE OR MULTIPLE STRATEGIES IN CONJUNCTION AND DISJUNCTION FALLACIES?

**JAMES TRIPP** *University of Warwick*  
**ADAM SANBORN** *University of Warwick*  
**NEIL STEWART** *University of Warwick*  
**TAKAO NOGUCHI** *University of Warwick*

People often misjudge the probability of the combination of two events, famously with the conjunction and disjunction fallacies. Classical probability theory tells us that the conjunction of two events cannot be more probable than either event on its own, and the disjunction of two events cannot be less probable than either event on its own. A long line of research shows us that people tend to violate both of these rules, and do so in a way that suggests that they are averaging the probabilities of the two events instead of combining them correctly.

Here we adopt a novel approach to explaining why people make these errors. We investigate the possibility that people consider multiple strategies when judging conjunctions and disjunctions, and use a statistically sophisticated procedure to combine these strategies: weighting each strategy according to the probability it is correct. For example, participants may consider both the normative combination rule and the mean of the two probabilities. Then the strategies are combined in a weighted average using weights that depend on the question asked. Our findings suggest that people do consider multiple strategies when judging conjunctions and disjunctions.

In Experiment 1 participants estimated how many people out of 100 had either (a) one of two attributes, (b) a conjunction of two attributes, or (c) a disjunction of two attributes. The attributes were denoted by randomly chosen letters and, following Wyer (1976), the probability of an attribute occurring was either "usually", "sometimes", or "rarely", corresponding to a high, moderate, and low probability respectively. This approach allows us to elicit the subjective probability of each adverb for a participant and calculate the corresponding predictions of each strategy for a conjunction and disjunction. These subjective probabilities of each adverb and the weights for each strategy were inferred from the data at an individual level.

Our results appear inconsistent with a single strategy account. When the adverbs are different, participant responses to conjunctions and disjunctive lie between their responses to each event occurring on its own, producing on average conjunction and disjunction fallacies and is consistent with an averaging account. However, when the adverbs are the same, responses to the conjunction and disjunction questions were more normative. Taken together, our results show that a single strategy explanation struggles to explain participant responses and question-dependent weights better explain the data. We further explored why

participants choose the weights they do using more transparent stimuli and rewards. We presented participants with events, using spinners, where the correlation between events is clear. Rewards were used to encourage participants to use the strategy they considered most likely. The results are discussed in terms of support for single or multiple strategy of information integration.

**Keywords:** Conjunction fallacy, disjunction fallacy, information integration, probability

[7]

## IN LIMBO: MARKING REGIONS IN STATISTICAL PARAMETRIC MAPS THAT ARE TOO CLOSE TO CALL

**GILLES DE HOLLANDER** *University of Amsterdam*  
**ERIC-JAN WAGENMAKERS** *University of Amsterdam*  
**LOURENS WALDORP** *University of Amsterdam*  
**BIRTE FORSTMANN** *University of Amsterdam*

Traditionally, fMRI data are analyzed using statistical parametric mapping approaches. Regardless of the precise thresholding procedure, these approaches ultimately divide the brain in regions that do or do not differ significantly across experimental conditions. This binary classification scheme fosters the so-called imager's fallacy (Henson, 2005), where researchers prematurely conclude that region A is selectively involved in a certain cognitive task because activity in that region reaches statistical significance and activity in region B does not. For such a conclusion to be statistically valid, however, a test on the differences in activation across these two regions is required (Nieuwenhuis et al, 2011). Here we propose a simple GLM-based method that defines a third "in-between" category of brain regions that are neither significantly active nor inactive, but rather "in limbo". For regions that are in limbo, the activation pattern is inconclusive: it does not differ significantly from baseline, but neither does it differ significantly from regions that do show significant changes from baseline. This pattern indicates that measurement was insufficiently precise. By directly testing differences in activation, our procedure helps reduce the impact of the imager's fallacy. The method is illustrated using concrete examples.

**Keywords:** functional magnetic resonance imaging (fMRI), statistical parametric mapping (SPM), imager's fallacy, interaction, univariate modeling, sandwich estimator

[8]

## ERRORS IN TYPING REVEAL HOW GOAL-DIRECTED AND HABITUAL SYSTEMS COORDINATE IN DECISION MAKING

**TOM STAFFORD** *University of Sheffield*  
**ÇIGIR KALFAOĞLU** *Eastern Mediterranean University*  
**COLIN BANNARD** *University of Texas at Austin*  
**LIZ MILNE** *University of Sheffield*

Errors in decision making can reveal the mechanisms by which correct performance is normally generated. We present three studies of errors in typing and discuss them in terms of how goal-directed and habitual processes coordinate in decision making.

Many studies of decision making measure responses in a series of discrete trials. Such paradigms bring an admiral clarity of what and when the decision under study is. Typing, as a decision making paradigm, sacrifices this clarity - multiple decisions are made per second, at multiple levels of description (e.g. word and letter choice) with overlapping timescales of execution. But even within discrete trial tasks, cross-over effects between notionally independent decisions are found (e.g. negative priming). The advantage of typing is that it is a real-world task for which highly skilled participants can rapidly generate a stream of data (keypress identities and timings). This data has the advantage of clearly defined errors. For our present purposes we also wish to highlight that typing is an example of a behaviour which combines both highly practised habitual or automatic elements with deliberate, controlled or goal-directed, action.

Our first study (Kalfaoglu, & Stafford, 2014) involved asking 19 touch-typists to copy-type 100 sentences without visual feedback from either their hands or the screen. Contrary to previous results using discrete trial tasks, we found no speeding or slowing of decisions prior to errors. Instead we found increased variability in inter-key intervals (IKIs) predicted error commission. This heightened variability was found for errors which were corrected (i.e. using the delete key) and errors which went uncorrected, suggesting that it was independent of goal-directed error-monitoring processes (unlike, for example, error slowing, which differentiated between corrected and uncorrected errors).

Our second study analysed EEG data recorded during the first study. Time frequency and Event Related Potential analysis were applied the components identified by an Independent Components Analysis. This analysis showed that electrophysiological signals normally associated with error detection are reliably associated with uncorrected errors. We conclude that they are insufficient to trigger error correction on their own. Synchronisation of activity in medio-frontal regions in the theta range (4-7hz) appeared necessary for errors to be corrected. This synchronisation also predicted greater error slowing, demonstrating the co-extensive nature of performance monitoring processes and decision making, in the domain of typing at least.

Finally, we conducted a computational linguistic analysis of error types produced during a second copy-typing experiment (Bannard, Kalfaoglu, Stafford, unpublished analysis). Errors were coded as simple motoric errors (e.g. pressing the key adjacent to the correct key) vs "action slips" (which involved the intrusion of some habitual element into ongoing goal-directed action; e.g. intending to type THINK but actually typing THING, derailing from the correct TH into

the high-frequency INC stem). Our analyses shows how transitional probabilities derived from a large corpus of written English show significant differences between the three decision outcome categories. We discuss this analysis in terms the potential for distinguishing the relative contributions of habitual and goal-directed processes to typing decisions between individuals and at different levels of control.

**Keywords:** Errors, Action-Slips, Habits

[9]

## DECISION MAKING UNDER COGNITIVE EXHAUSTION: USING OR OMITTING INFORMATION?

**JOHANNES RITTER** *University of Erfurt*

A basic tenet of the bounded rationality approach holds that processing of decision-relevant information is effortful and that processing resources are limited. Building on this assumption, several research strands have shown that people can selectively use pieces of information in their decision making - as opposed to exhaustive use of all information - and thereby save effort and time. In contrast, there are several experimental paradigms (e.g., hindsight bias, stroop task, etc.) suggesting that the suppression of readily available information itself is effortful. In addition, there is widespread evidence that the human cognitive system is able to integrate several pieces of information effortlessly and automatically even across different modalities. Thus, there are two opposing hypothesis concerning the consequence of cognitive exhaustion on decision making. The bounded rationality perspective predicts omission of information and employment of parsimonious decision strategies such as heuristics. The automated integration perspective predicts more exhaustive use of information available in the decision making process.

The present study was designed to test these opposing hypotheses using a cognitive exhaustion paradigm: Participants in the exhaustion condition were required to perform a color-stroop task before and in between blocks of probabilistic decision tasks. Participants in the control condition were required to perform a non-exhaustive version of the stroop task before and in between blocks of probabilistic decision tasks. Additionally to the exhaustion factor, the decision environment was manipulated with a low dispersion of cue validities and high dispersion of cue validities, resulting in a 2x2 design with the factors exhaustion (high vs. low) and decision environment (high dispersion vs. low dispersion). High dispersion environments are assumed to stipulate selective decision strategies such as non-compensatory heuristics because these environments allow omitting pieces of information without a decline in accuracy. It was therefore of special interest for the hypotheses how participants would decide in the condition with high dispersion and under high exhaustion. The decisions of the participants were analyzed using decision outcomes and confidence judgments. Participants were classified as users of either an exhaustive decision strategy (using all information) or a selective decision strategy (using only limited information). As expected, the results show that participants were more likely to use selective decision strategies in the high-dispersion environment. However, they were less likely to

do so under cognitive exhaustion. These results are more in line with the automated integration perspective and call the assumption of generally limited cognitive resources of the bounded rationality perspective into question.

**Keywords:** bounded rationality, cognitive exhaustion, probabilistic inference

[10]

## THE COLLECTIVE MARKET DYNAMICS OF POPULARITY-ORDERED SEARCH

**HRVOJE STOJIC** *Universitat Pompeu Fabra*

**PANTELOS P. ANALYTIS** *Max Planck Institute for Human Development*

**MEHDI MOUSSAI** *Max Planck Institute for Human Development*

From academic search engines like Google Scholar to news sharing tools such as Digg many of our choices take place in environments that rank the alternatives in decreasing order of popularity. The order in which the decision makers consider the alternatives has been shown to be an important element of the decision making process. What happens when individual choices alter the popularity of the alternatives, and consequently, the order in which other decision makers consider them?

In this paper, we present a simple collective behavior model in which the individual behavior is based on the principles of search theory. In our framework decision makers with correlated preferences decide sequentially which alternative to choose. The decision makers sample the alternatives in order of decreasing popularity and choose the first satisficing alternative they encounter. The popularity order is updated after each individual choice. In comparison to existing social influence accounts of multi-alternative choice our model is the first to explicitly describe the search and stopping rules followed by the decision makers.

We present analytical results for some boundary conditions that are tractable and simulation results for the rest. We simulate a market consisting of hundred alternatives and thousand agents varying the satisficing threshold employed by the agents and the degree of heterogeneity of preferences in the population. Through-out the analysis, we employ an optimal stopping model in which the decision makers sample the alternatives at random as a theoretical yardstick.

We find that in popularity search the decisions of few decision makers in the beginning of the sequence determine the search path for decision makers choosing after them and have a large impact on the overall outcome in the market. This effectively leads to rich-get-richer dynamics, which imply greater unpredictability and inequality in the market as compared to a random search environment, where the market shares of the alternatives are predictable in the long run. We show that in popularity search a lower threshold implies more inequality in the market while it has the exactly opposite effect in random search. Further we study the strength of the rich-get-richer dynamics by comparing the utility of the most popular alternatives in the market to that of the average alternative utility and the highest possible utility. Finally, in popularity search we find that some preference heterogeneity leads to higher levels of average utility as

compared to a population with perfectly homogenous preferences. We explain this result by showing how some heterogeneity implies additional search, which in turn leads to capturing some of the informational externalities.

**Keywords:** ordered-search, satisficing, collective behaviour, rich-get-richer dynamics

[11]

## STATUS AND EMPATHY: MODULATING FACTORS OF COOPERATIVE BEHAVIOURS?

**JIE-YU LV** *Queen Mary University of London*

**MICHAEL J. PROULX** *University of Bath*

**MAGDA OSMAN** *Queen Mary University of London*

“What can one do to promote cooperation?” As the most frequent prosocial behaviour, cooperation is a well-studied topic in psychology, and has led many to take up this question. Several methods have been used to promote cooperation, such as the inclusion of reinforcers (i.e. rewards and punishments) (Fehr & Gächter, 2000). Others have focused on methods designed to promote trust, identify with others, and increase people’s perspectives of long term prospects (Parks et al., 2013). Along with these factors, one which has been argued to be crucial for supporting cooperation is empathy (Batson & Ahmad, 2001; Batson et al., 1995; Batson & Moran, 1999; Beadle et al., 2013; Cohen & Insko, 2008; Rumble et al., 2010; Xu et al., 2012). In order to address how empathy and status effects cooperation, we used a classic social dilemmas paradigm Public Goods Game (PGG).

The present study included two computer-based PGG experiments. There were three fictitious confederates and one real participant involved. Experiment 1 (N= 52): We used a between-subject design with three levels of empathy (high-empathy [HE, n = 17], low-empathy [LE, n = 17] and no-empathy [NE, n=18]) and randomly allocated to each condition. The empathy manipulation involved breaking up with boyfriend, experiencing a car accident, and having a mobile phone stolen; this was similar to Batson and Moran’s (1999) empathy induction method. After this, there was a PGG, and finally a set of empathy questionnaires were presented (Interpersonal Reactivity Index (IRI) (Davis, 1983). We found that the empathy manipulations didn’t yield significant differences in the average contributions in the PGG [F (2, 49) = 0.101, p = 0.904 > 0.05]. Post hoc test showed that there were no significance differences among the three conditions; HE (M = 9.00, SD = 7.29), LE condition (M = 9.08, SD = 7.29), NE condition (M = 8.50, SD = 7.08).

In Experiment 2, just as in Experiment 1 there were three fictitious confederates and one real participant were involved (N = 49). We compared high-status [n = 26] with low-status [n = 23] conditions. There was no manipulation of empathy. Participants simply saw pictures and names of three fictitious characters. All three fictitious confederates were endowed with 20 tokens at the start of the PGG, but half of the real participants were endowed with 30 tokens (high-status) and the other half were endowed with 10 tokens (low-status). We found that consistent with our prediction, the low status group (M = 0.565, SD = 0.28) contributed more percentage than the high status group (M = 0.441, SD = 0.301) [F (1, 47), p = 0.041 < 0.05].

In summary, in Experiment 1 we found that the empathy manipulation didn't increase cooperative behaviours in the PGC. A further replication would be needed to fully establish how reliable this finding is, and whether empathy manipulations in general are an effective way of inducing higher level of cooperation. Batson and Moran (1999) and Rumble et al. (2010) have shown that empathy induces cooperation, but other work suggests that the role of empathy in cooperation is less reliable (Eisenberg & Strayer, 1990, p. 301). In Experiment 2, we found individuals in the high status position behaved less cooperatively than those in the low status position. Future work could explore potential interactions between empathy and status. It may well be that the empathy induction methodology that we used requires further finessing, and establishing a link between empathy and status as indexed by levels of cooperation would be an innovation in the domain of research on pro-social cooperative behaviours.

**Keywords:** empathy; status; cooperation; Public Goods Game

[ 12 ]

### HUMAN VISUAL CATEGORISATION IS ONLY CONSISTENT WITH BAYESIAN GENERATIVE REPRESENTATIONS

FERYAL M. P. BEHBAHANI *Imperial College London*  
 A. ALDO FAISAL *Imperial College London*

Intelligent behaviour is fundamentally tied to the ability of the brain to make decisions in uncertain and dynamic environments. In order to accomplish this task successfully, the brain needs to categorise novel stimuli in real-time. In neuroscience, the generative framework of Bayesian Decision Theory has emerged as a principled way to predict how the brain has to act in the face of uncertainty (Ernst & Banks, 2002, Kording & Wolpert, 2004, Kemp & Tenenbaum, 2008, Faisal et al., 2008). We hypothesise that the brain might also use generative Bayesian principles to implement its categorisation strategy. Previous experimental work on human categorisation shows data that is consistent with both discriminative and generative classification (Hsu & Griffiths, 2010) and did not allow confirming the implementation of one or the other. Therefore, we designed a novel experiment in which subjects are trained to distinguish two classes A and B of visual objects, while exemplars of each class are drawn from Gaussian parameter distributions, with equal variance and different means. During two different experimental paradigms, we test how the subjects representation of the categories change as a result of being exposed to outliers for only one of the categories, A, far from category B, i.e., increasing category A's variance.

Generative classifiers are by necessity sensitive to novel information becoming available during training, which updates beliefs regarding the generating distribution of each class. In contrast, discriminative classifiers are sensitive to novel information only if it affects the immediate discrimination of classes. In the first paradigm, we characterise the categorisation boundary, i.e., the point where a subject assigns with equal probability a test stimulus to either category. Subsequently, we track the shift in the boundary after the introduction of outliers. A generative

classifier will prompt to reconsider the variance of class A after being exposed to the outliers and accordingly shifts the categorisation boundary towards category B. However, the discriminative classifier will not react, as there is no new information added to the boundary itself. Our second paradigm provides an even more stringent test for generative models: again, outliers for class A are presented far away from class B. Additionally, the two classes are selected to be close enough, such that a generative classifier would assume that class A's variance has increased significantly, thereby reaching across the region occupied by B. This will result in the emergence of a second classification boundary to the distal side of class B, far away from class A. Again, the discriminative classifier would not change its behaviour.

Our results in both paradigms show that the introduction of the outliers for category A influences the subject's knowledge of the distribution associated with alternative categories. This can result in the introduction of additional boundaries only predicted by our simulations of generative classifiers. These results give clear evidence that visual categorisation is only consistent with generative and not discriminative classification mechanisms. Furthermore, our experiments provide an ideal experimental framework for neurophysiological and functional imaging investigations of the underlying neural mechanisms involved.

**Keywords:** categorisation, visual perception, generative models

[ 13 ]

### EMOTIONAL GAMES

RICHARD FAIRCHILD *University of Bath*

Behavioural economists are increasingly understanding that humans are not completely self-interested or emotionless, but often exhibit "other-regarding" behaviour. We develop a game-theoretic approach in which players gain utility from their own material payoffs, but who also develop empathetic emotions towards each other. David Sally (2001, 2002) argues that reciprocal behaviour may depend on the social context, social interaction, and psychological closeness of the players. Motivated by Sally's seminal analysis of sympathy games, we develop a two-stage, extended form, empathy model in which players simultaneously choose empathy levels in one stage, and, in a separate stage, make simultaneous strategy choices in a material game. We consider both conscious (strategic/instrumental), and unconscious (innate) empathy (you simply like someone for who they are). We demonstrate that increasing empathy can have positive or negative effects on welfare in equilibrium, and that these equilibria can be crucially affected by whether empathy is formed consciously, or unconsciously. We tentatively term our modelling approach, 'emotional game theory.'

**Keywords:** Game theory; empathy; conscious and unconscious emotions

## THE PSYCHOLOGY OF NATIONS AND STATES

**NEIL STEWART** *University of Warwick*

**TAKAO NOGUCHI** *University of Warwick*

**CHRISTOPHER Y. OLIVOLA** *Carnegie Mellon University*

**HELEN SUSANNAH MOAT** *Warwick Business School*

**TOBIAS PREIS** *Warwick Business School*

People's Internet searching reveals something of their thinking. Search engines such as Google offer access to daily aggregated frequencies for search terms at the level of nations and states. We used Google Trends to extract (a) psychological measures of people's thinking about time and (b) measures of the connotation of people's searches. These measures correlate strongly with the economic performance of nations and states, demonstrating the viability of constructing nation-level psychological measures from Google Trends data.

Time Perspective: from the frequencies of 4-digit year search terms (e.g., "2012") we constructed four psychologically inspired measures of time-perspective and examined their relationship with a widely-used measure of economic activity, per-capita gross domestic product (GDP). Past and future focus measures the extent to which people search for last year or next year compared to this year. At the level of individuals, higher future focus is associated with stronger economic performance (Zimbardo & Boyd, 1999). Past and future time horizon measures how far into the past or future people are searching. At the level of individuals, thinking further into the future is associated with being richer and more educated (for review, see Reimers, Maylor, Stewart, & Chater, 2009). We find that nations with higher per-capita GDP are more focused on the future,  $\chi^2(1) = 19.65$ ,  $p < .001$ , and less on the past,  $\chi^2(1) = 9.19$ ,  $p = .002$ , and that when these nations do focus on the past, it is more likely to be the distant past,  $\chi^2(1) = 7.77$ ,  $p = .005$  (Figure 1). Our measures of time-perspective capture 53% of the variance in per-capita GDP.

Connotations of Search Terms: the semantic differential (Osgood, 1952) is to concepts as the big five is to personality. Using ratings of how well adjectives describe concepts, Osgood recovered a three-factor solution for the measurement of meaning. We used the relative frequencies of Osgood's adjectives in search terms to construct a semantic differential for Google searches at the level of US states. The second factor (kind, nice, sharp, deep, sweet vs. crude, awful, dull, shallow, bitter) showed a strong correlation with per capita gross state product (GSP) and equality of income distribution (Gini coefficient). More kind, nice, sharp, deep, and sweet searching is associated with lower GSP ( $r = -.21$ ,  $p = .0002$ , Figure 2) and a more equal distribution of income ( $r = -.39$ ,  $p = .004$ , Figure 3).

Although the direction of causality cannot be established from our correlational analyses, the strong associations of our measures with known economic measures like GDP and inequality indicate the viability of using nation-level data, together with psychological concepts, to measure the behaviour of nations.

**Keywords:** Google trends, temporal discounting, gross domestic product

## UNIVERSITY ADMISSIONS: A GRAPH-MODULAR DECISION-SUPPORT TOOL

**ANJALI MAZUMDER** *University of Warwick*

**JULIA BRETTSCHEIDER** *University of Warwick*

The rich heterogeneity and success of academic institutions can in part be reflected by universities' admissions decision-making process. Admissions tutors are tasked with making decisions regarding placement offers to applicants with the expectation that an applicant will succeed in the programme of study. While institutions provide information about the applicant profile they seek, the actual decision-making process is often seen as less transparent. Admissions decision-making can be a complex process, often made difficult by multiple and partial sources of information which is not uniformly collected or used. A simplistic generalisation of admissions can be broadly described as (i) formulaic, relying primarily on academic results, or (ii) 'holistic', examining broader student profiles such as statements and references. Often decisions involve a more comprehensive approach in which formulaic criteria are combined with holistic judgements: The challenge then becomes assigning ratings to attributes of judgements to facilitate a structured, consistent and transparent approach.

A key challenge to making an offer is that decisions are often made on predicted grades as at the time of applying, most students will not have final grades. This introduces a new decision-making problem: Once an offer has been made, some applicants may fail to meet the threshold and subsequently a decision to rescind the offer may be made. Other supplementary results can support an offer such as STEP results (for the mathematical sciences). An added uncertainty is that there may be variability in what a grade level achieved at one school is versus another. Holistic evidence such as participation in competitions (e.g. Math Olympiads), supporting statements and references, is inherently variable. Thus, decision-making involves a mixture of measured variables and expert judgement.

While some institutions have long-established practices, entrenched in historic culture and underlying missions, admissions has become increasingly dynamic. External factors such as changing demographics, economics, and political, social and educational environment have prompted many institutions to modify their selection process. Shifts in early decisions, open days, interviews, etc. and technology e.g., online applications, have encouraged institutions to rethink how to manage applications and effect admissions process. Students' decisions to apply and accept an offer is multi-faceted: This may involve financial incentives, career prospects, location, etc. Thus, therein lies the challenge of achieving a careful balance of set targets and numbers (whether financial and/or demographic) and attracting applicants who will succeed and fit into the institutional environment.

We extract and use historically relevant data from applicants to mathematical sciences programmes at departmental, university and national level to inform a graph modular structure for an admissions decision support system. Using Bayesian Networks, we develop a decision-support system to deal with uncertain, ambiguous and incoherent evidence. The graph modular structure facilitates interrogation of sub-models, calibrating particular facets or sub-objectives. Furthermore, it aids and provides a vehicle for

communicating the decision-making process within a team, ensuring a dynamic, systematic and transparent approach is managed over a long application period.

**Keywords:** decision support systems, uncertainty, expert judgement, admissions, combining evidence

[16]

## THE ROLE OF EMOTIONAL INTELLIGENCE IN ECONOMIC DECISION MAKING

**JASMIN MAHMOODI** *University of Edinburgh*  
**ELIZABETH J. AUSTIN** *University of Nottingham*

A common decision making paradigm is delay discounting, which is a way of assessing financial decisions regarding receiving either a smaller immediate or a larger delayed monetary reward. Individuals differ in the amount of benefit required to cause them to choose the delayed options. Currently, the main candidates for accounting for these differences are mood and the personality traits impulsivity and self-control. To date, however, no study yet has examined the potential role of trait Emotional Intelligence (EI), which can be described as the ability to recognise, modify, and use own and others' emotions, in delay discounting. This study investigated the relationship of EI to individual differences in delay discounting, contrasting this with possible effects of induced mood. Contrary to previous findings, neither positive nor negative induced mood influenced discounting behaviour. Turning to the effects of EI, regression analyses showed that the "Regulation of Emotion" facet of trait EI significantly predicted lower discount rates, which is the preference for larger delayed rewards. No similar effects were found for either total EI nor other EI subscales. Results are discussed against the background of previous research and its limitations and implications for future research are addressed.

**Keywords:** delay discounting, emotional intelligence, mood, decision making, individual differences

[17]

## MEASURING RATIONAL THINKING

**ROBERT BIEGLER** *NTNU*  
**GERIT PFUHL** *NTNU*

Making good decisions depends on rationality as much as or even more than on intelligence. Stanovich (2009) pointed out that if rationality is lacking, intelligence may only provide people with more plausible reasons for believing whatever it is they want to believe. When it comes to selecting people for various positions, intelligence is often tested, but rationality at best rarely, perhaps never. One reason for that state of affairs is that intelligence tests are well developed, reliable, and familiar.

We have found only one rationality-test in the literature (Pacini and Epstein, 1999). It consists of 40 items that record attitudes towards and self-assessments of rationality. The test does not directly measure rationality itself.

Scattered across a number of papers are items that seem to measure rationality more directly.

We want to see how well these correlate with Pacini and Epstein's test, and also with measures of simple forms of metacognition: how good are people at judging how well they remember, and how good are they at estimating ambiguity?

We measured a) how well people can judge how precise their memory for locations and shapes is; b) their ability to estimate conditional probabilities given varying numbers of data points; c) how they perform on the rationality questionnaire items e.g. knights and knaves problem. We further examined how much evidence people need to draw a conclusion, and how much they dislike uncertainty in the forms of risk and of ambiguity.

Finally, we correlated the results with an IQ test. We found a significant correlation between IQ and rational engagement and Stanovich's rationality measure. Further, rational engagement was correlated with Need for Closure. There was no effect of attentional control or risk and ambiguity.

**Keywords:** metacognition, cognitive biases, uncertainty, risk, ambiguity, IQ

[18]

## INFORMATION FORAGING FOR PERCEPTUAL DECISIONS

**CASIMIR LUDWIG** *University of Bristol*  
**DAVID R EVENS** *University of Bristol*

How do we sample noisy information from multiple sources in order to make decisions? We address this general problem in the perceptual domain. Subjects make a comparative direction discrimination judgement of two random dot motion patterns that are activated only when directly fixated. The viewing time is limited and subjects have to adapt their time allocation to the quality of the two patterns. A simple foraging model assumes a pattern is sampled until the change in the precision of the direction estimate for that pattern drops below a threshold. At that point the subject switches to the other pattern and accrues information from that source. We test two predictions of the model: (i) the timing of the first switch should depend on the online accrual of information; (ii) the representation of previously sampled information decays.

We tested the first prediction by briefly "pulsing" the quality of motion information during the fixation of the first pattern. If the switch time depends on the amount of information accrued (i.e. amount of certainty about the direction of the first pattern), subjects should switch earlier when the motion coherence of the first pattern is briefly enhanced. Conversely, when the information quality is briefly degraded, subjects should delay the timing of the first switch.

We tested the second prediction by asking subjects to estimate the direction of either the first pattern or the second pattern. Subjects are directed to inspect one pattern first and when that pattern stops moving, they have to switch to the second pattern. We varied the viewing duration of both the first and the second patterns. Once both patterns have been fixated, a dial appears in the location of either the first of the second pattern. Subjects have to estimate the direction of the corresponding pattern. Decay is estimated

through the precision of the direction estimate of the first pattern, as a function of the viewing duration of the second pattern.

Experiment 1 confirms that in the comparative decision task, switch time is (at least partly) controlled by the online accrual of information. Experiment 2 demonstrates that there is surprisingly little decay in the representation of previously sampled information. However, the accrual of new information is slowed down as a result of having to maintain a representation of the previously sampled source. These findings place important constraints on our model of information foraging.

[19]

## ANXIETY DISORDER AS A PRODUCT OF AN ADAPTIVE THREAT DETECTION SYSTEM

[ANDREW HIGGINSON](#) *University of Bristol*

DANIEL NETTLE *University of Bristol*

BEN BRILOT *University of Bristol*

Anxiety disorder is usually assumed to be a product of a malfunctioning cognitive system. Adaptive reasons that a cognitive system may be disposed to anxiety disorder have rarely been considered. We modelled a Bayesian forager that learns about the presence and density of predators in its environment. The forager decides the proportion of its time to look for predators rather than food and what level of false alarms to tolerate. This optimal learning system results in a significant proportion of individuals behaving as if the environment is dangerous even when predators have been rare for a long time, which we consider to be operationally analogous to anxiety disorder in humans.

**Keywords:** Predation, Foraging, Bayesian learning, Signal detection



