LIBR
Laureate Institute for Brain Research

Decision-Making and Computational Psychiatry: An Explanatory and Pragmatic Perspective

Martin P. Paulus M.D.
Scientific Director and President
Laureate Institute for Brain Research
6655 S Yale Ave
Tulsa, OK 74136-3326
email: mpaulus@laureateinstitute.org
web: http://www.laureateinstitute.org

Disclosures

- Royalties from UpToDate
- Funding support:
  - 1R01MH101453
  - U01 DA041089
  - P20 GM121312

Outline

- An Explanatory Or Pragmatic Framework
- Computational Failure Models
- **Four Examples Of Computational Failures Modes In Anxiety**:
  - Differentiating Signal From Noise During Change Point Detection
  - Adjusting To Current And Future Error In A Movement Task
  - Perception Versus Decision-making In An Uncertain Environment.
  - Differentiating Decision Uncertainty From Emotional Conflict In An Approach/Avoidance Situation.

What are the decision-making dysfunctions in individuals with psychiatric disorders?

How does the brain contribute to these dysfunctions?

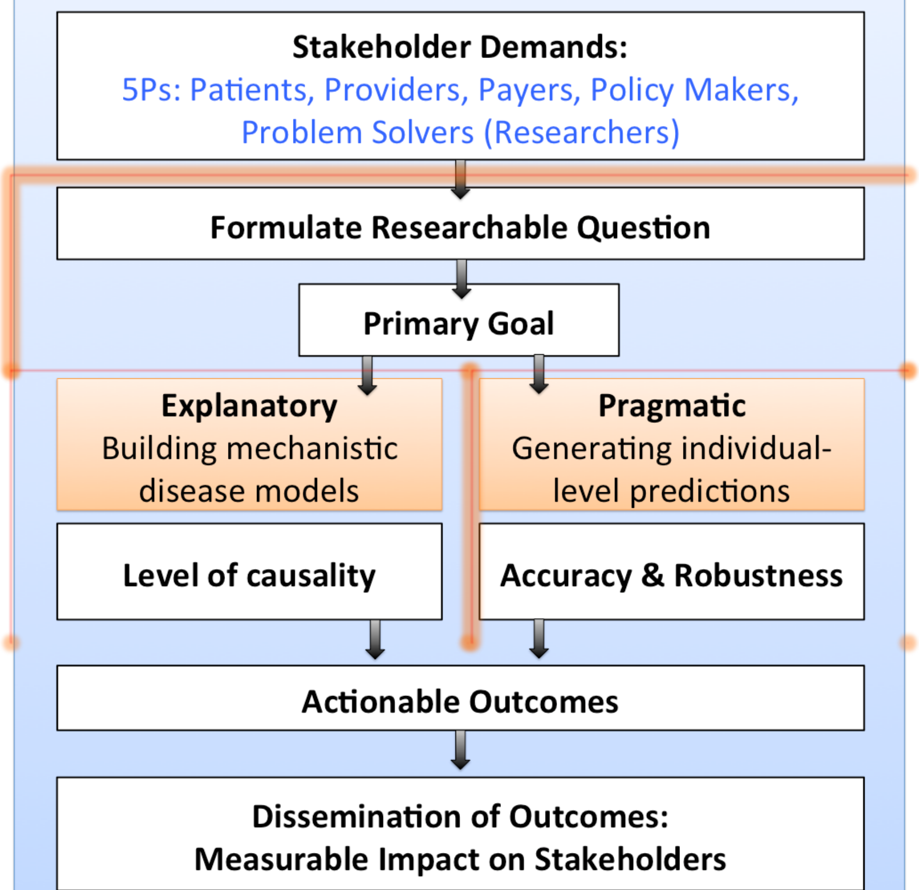

Criteria for Evidenced-Based Pragmatic Psychiatry:

- Problem (what do stakeholder want?)
- Experimental Design (intervention vs. observation)
- Levels of Analysis (is it explanatory or predictive?)
- Statistical Framework (how reproducible?)
- Outcome Variables (how valid and actionable?)

Can we use biological or other variables to predict clinically meaningful outcomes?

Evidenced-Based Pragmatic Psychiatry—A Call to Action

Martin P. Paulus, MD

Computational Psychiatry

Old Approach:

Behavior

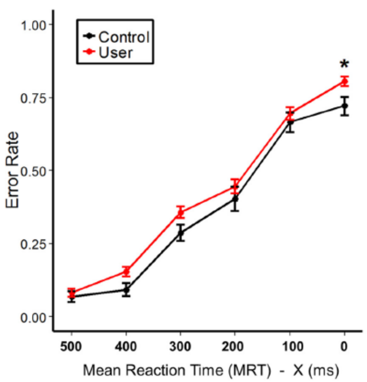

Brain Processing

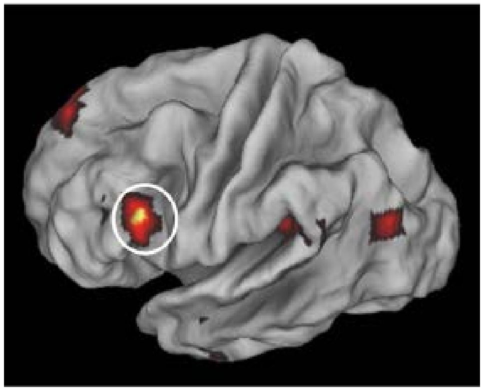

New Approach:

Behavior

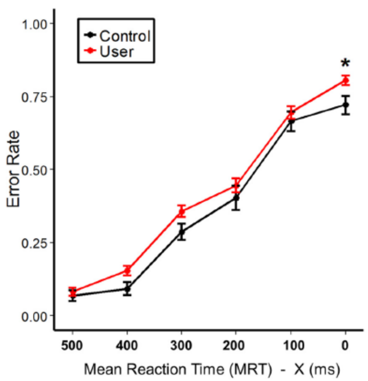

Processing Model:

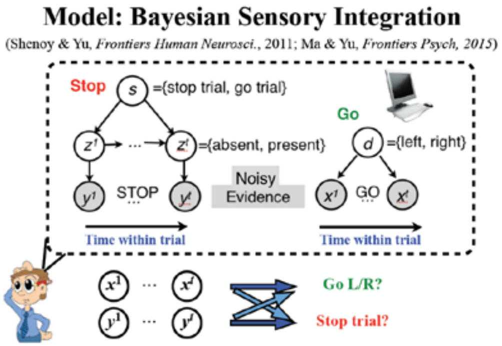

Model-derived Brain Processing

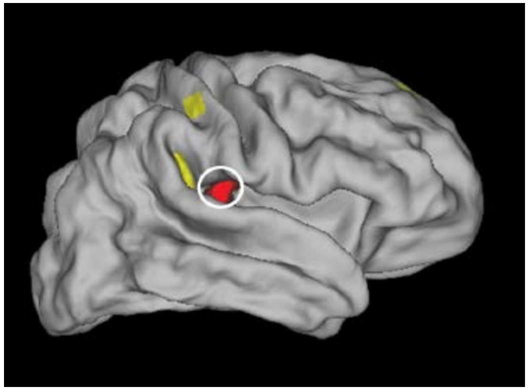

Computational Approaches to Aversion-Related Decision-Making In Psychiatry (ARDM)

Driven by Pain, Not Gain: Computational Approaches to Aversion-Related Decision Making in Psychiatry

Martin P. Paulus*
Laureate Institute for Brain Research, Tulsa, Oklahoma
Department of Psychiatry, University of California, San Diego, La Jolla, California

Anxiety: Computational Failure Modes and ARDM

- Reinforcement Learning framework:
  - Sensitivity to reward or punishment
  - Slower updates to aversive prediction errors
  - Overwhelming Pavlovian biases
  - Altered reference points- framing, counterfactuals
- Active Inference framework:
  - Habitual predictions that are computationally less effortful
  - Excessive response cost
  - Altered prior beliefs about state-observation relationship

Value-based Decision-Making in Addiction and Anxiety

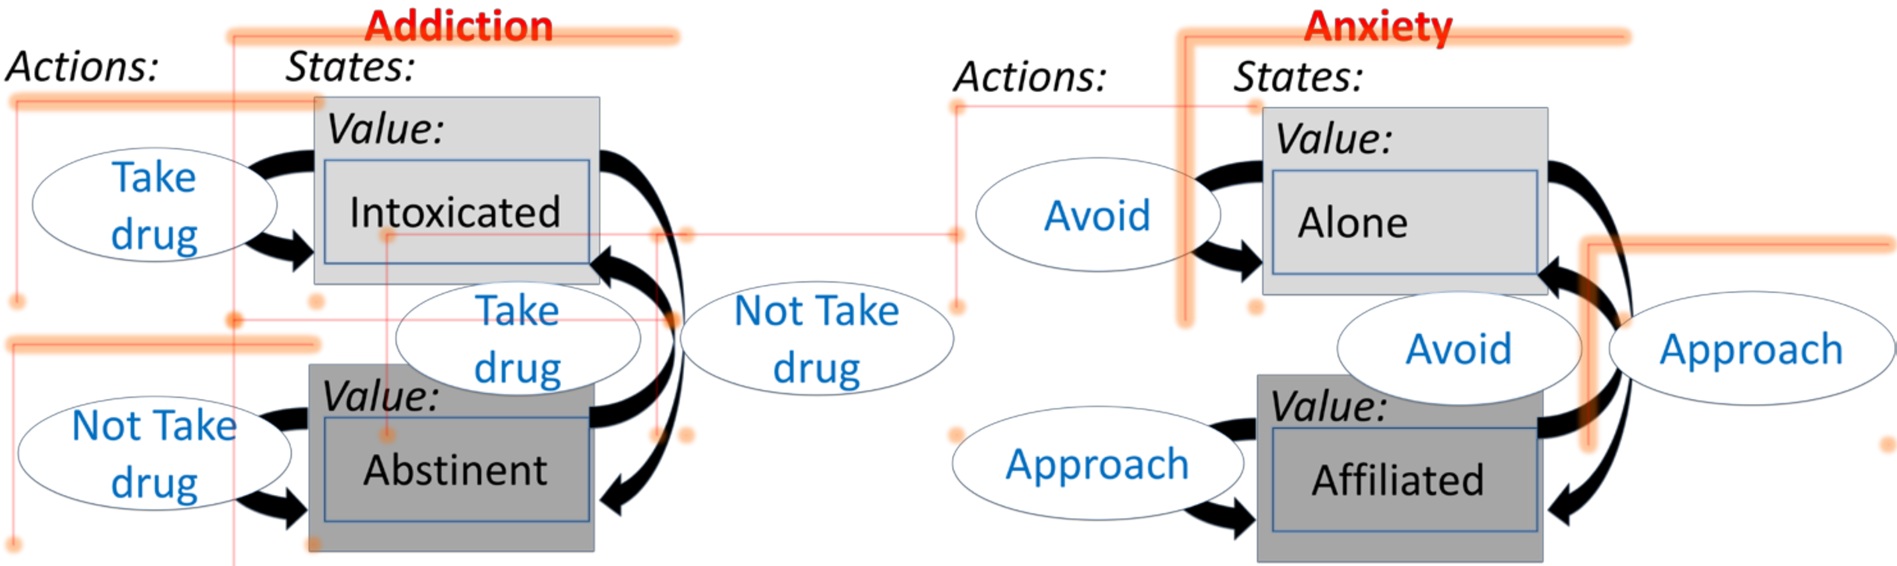

Use Value-Based Action Selection: select actions that increase the future value of the states (i.e. intoxicated/alone or abstinent/affiliated)	
Positive Reinforcement	*Negative Reinforcement*
Select an action that make positively valued states **most** likely.	Select an action that makes negatively valued states **least** likely.

Active Inference in Addiction and Anxiety

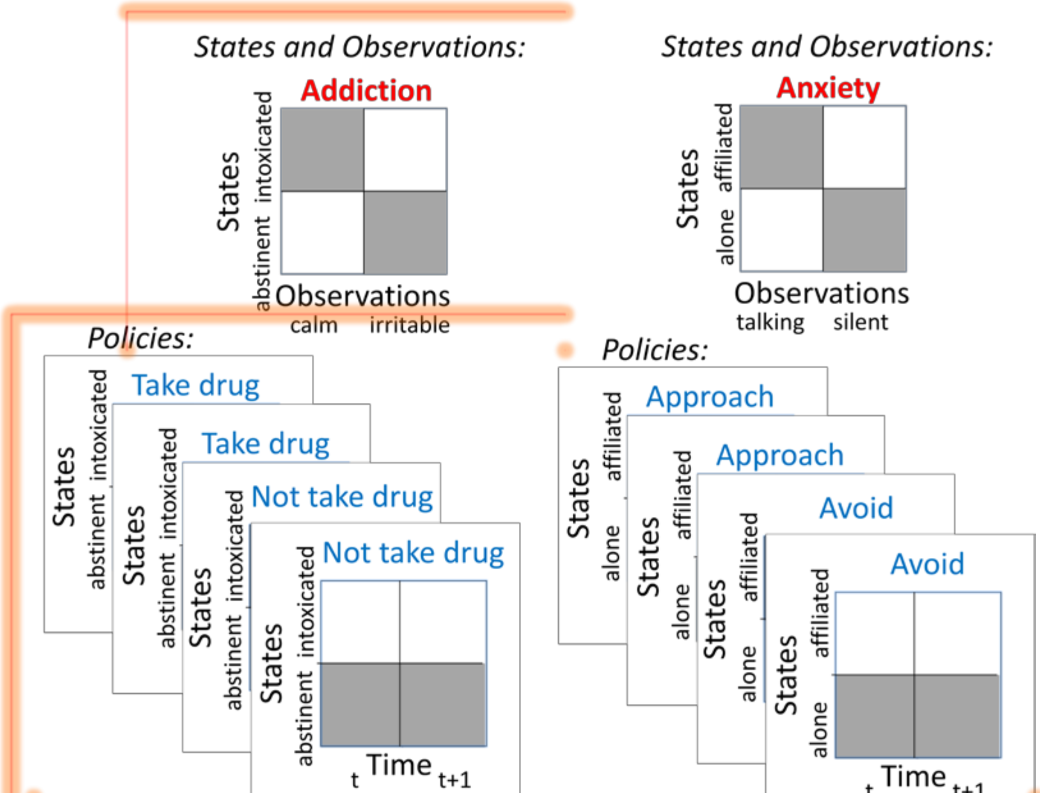

Use Free Energy Principle: minimizing the difference between the states that are expected and observed and the amount of surprise of making an observation (i.e. calm/irritable or talking/silent)

Positive Reinforcement	*Negative Reinforcement*
Select an action from the policies that makes the observation (calm, talking) most likely.	Select an action from the policies that makes the observation (irritable, silent) least likely

Ccomputational Failure Modes:

- Hyper-precise Priors (prior probability)
  - The expectation of afferent information is so precise that incoming evidence does not significantly alter the expectation.
- Context Rigidity
  - The individual is unable to adjust the prior expectation of information to a different context.

Annual Review of Clinical Psychology

An Active Inference Approach to Interoceptive Psychopathology

Martin P. Paulus,[1] Justin S. Feinstein,[1,2] and Sahib S. Khalsa[1,2]

Computational Failure Modes and Interoception:

- Hyper-precise Priors (prior probability)
  - The expectation of afferent information is so precise that incoming evidence does not significantly alter the expectation.
- Context Rigidity
  - The individual is unable to adjust the prior expectation of information to a different context.

Annual Review of Clinical Psychology

An Active Inference Approach to Interoceptive Psychopathology

Martin P. Paulus,[1] Justin S. Feinstein,[1,2] and Sahib S. Khalsa[1,2]

Hyper-precise Priors

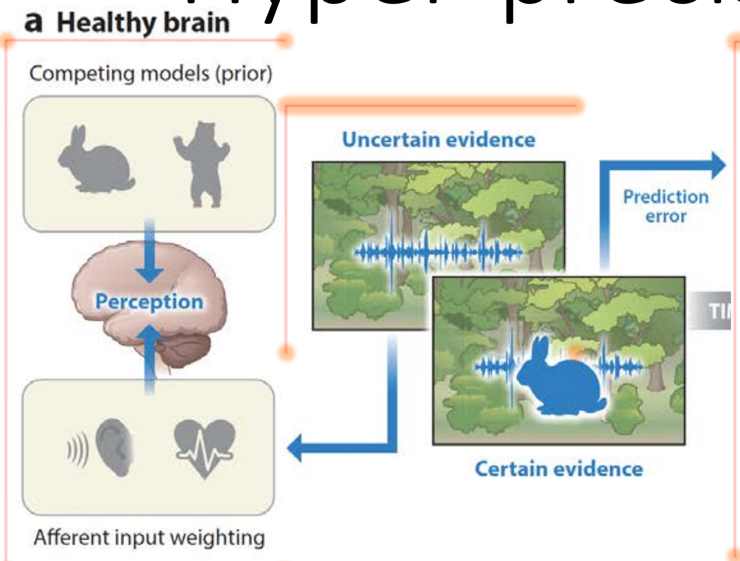

Context Rigidity

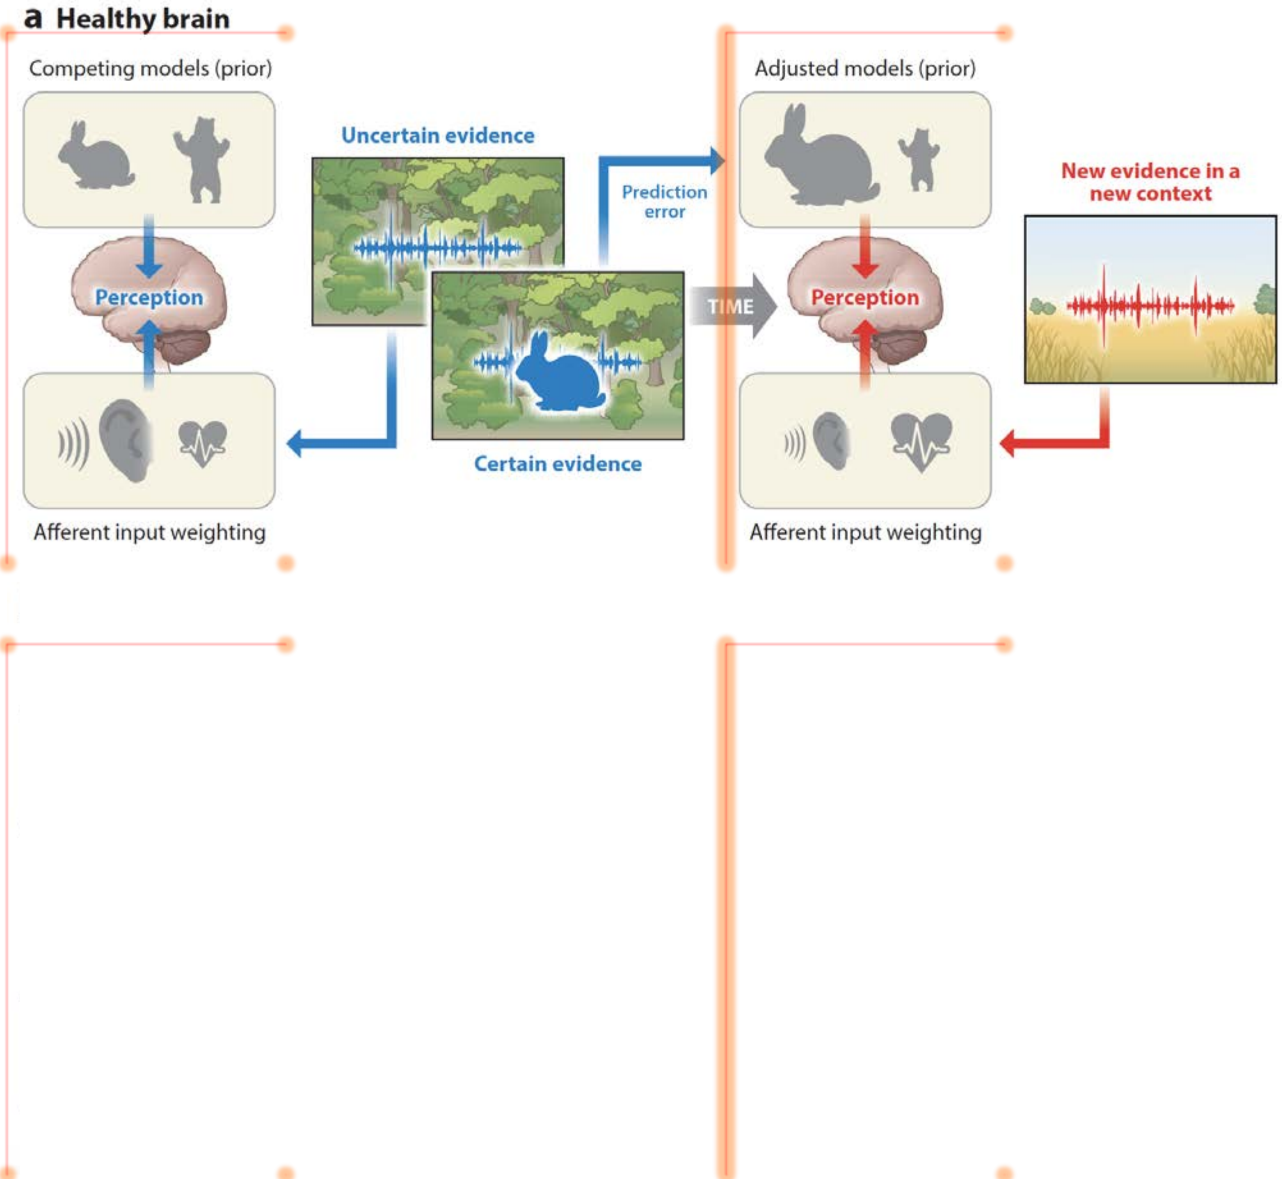

Computational Approaches to Identify Processing Dysfunctions in Anxious Individuals

Computational Dysfunctions in Anxiety: Failure to Differentiate Signal From Noise

He Huang, Wesley Thompson, and Martin P. Paulus

Change Point Detection

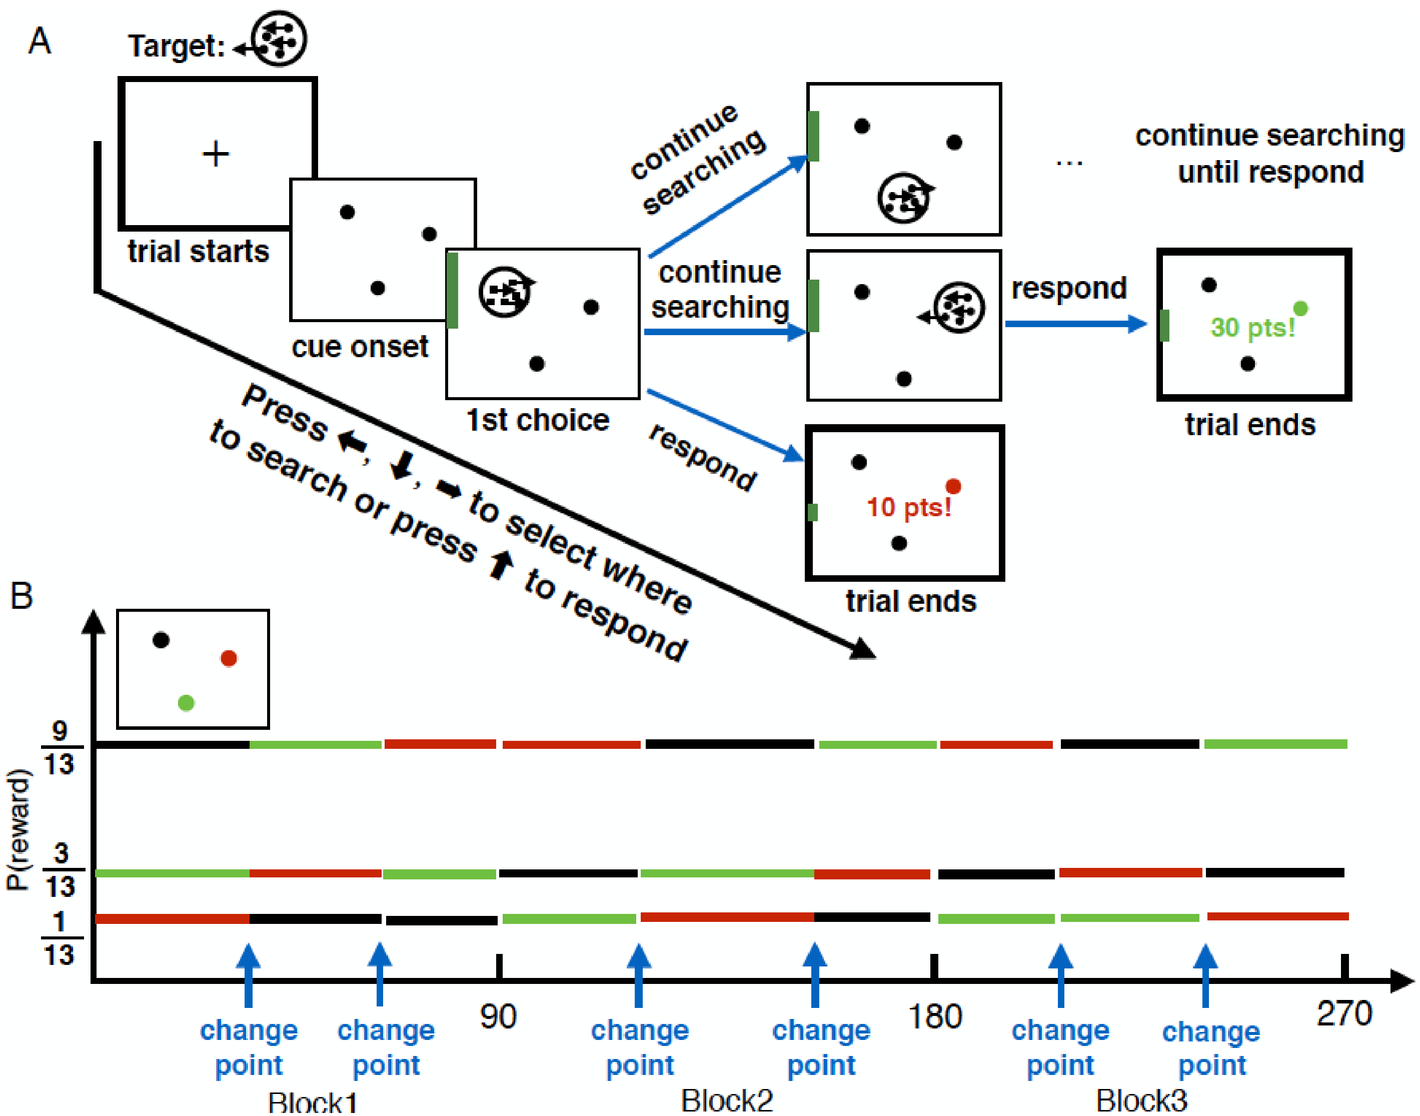

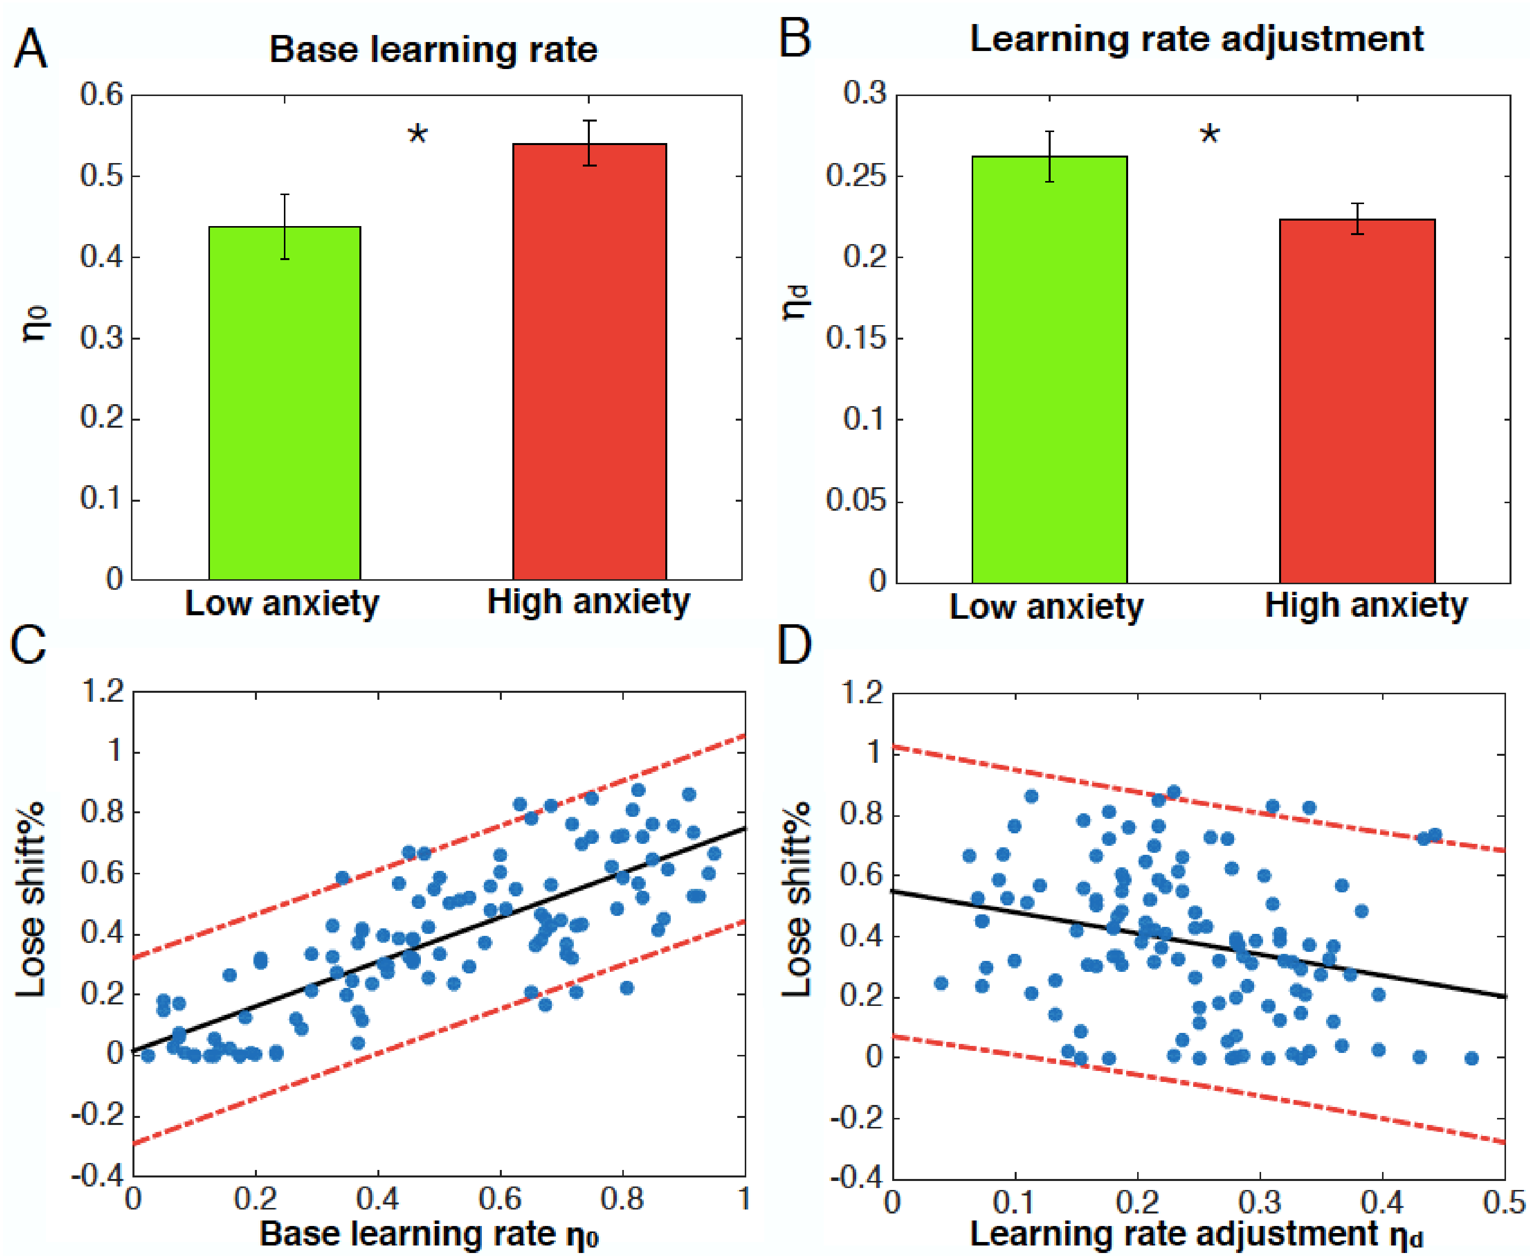
A
Base learning rate
0.6
0.5
0.4
0.3
0.2
0.1
0
η0
*
Low anxiety
High anxiety
B
Learning rate adjustment
0.3
0.25
0.2
0.15
0.1
0.05
0
ηd
*
Low anxiety
High anxiety
C
1.2
1
0.8
0.6
0.4
0.2
0
-0.2
-0.4
Lose shift%
0
0.2
0.4
0.6
0.8
1
Base learning rate η0
D
1.2
1
0.8
0.6
0.4
0.2
0
-0.2
-0.4
Lose shift%
0
0.1
0.2
0.3
0.4
0.5
Learning rate adjustment ηd

Conclusions

Why is this important?

- Both anxious and non-anxious in duals learn about the underlying statistics of the c e point detection task.
- **Anxious individuals tical fluctuations ge.**

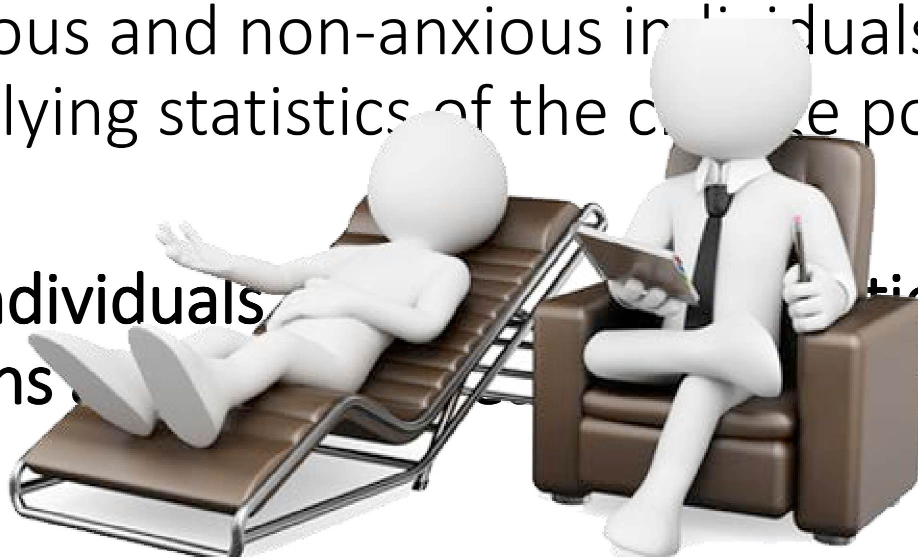

- Modern treatments of anxiety are based on re-learning fear-related content.
- New behavioral or pharmacological strategies need to be developed if anxious individuals do not learn appropriately.
- The computational approach allows us to precisely quantify the degree of learning dysfunction and to determine how much intervention correct it.

Evidence for Underweighting of Current Error and Overestimation of Future Error in Anxious Individuals

Jonathan Howlett

Start – Stop Task

- The position of a virtual car was controlled using a gaming joystick.
- Each subject completed 30 trials. In each trial, subjects were instructed to drive the car as quickly as possible and stop as close as possible to a stop sign without crossing the stop-line.
- At each time point within a trial, acceleration was modeled as a linear combination of current error (goal position minus current care position) and derivative of the error, with coefficients K_p and K_d, respectively.

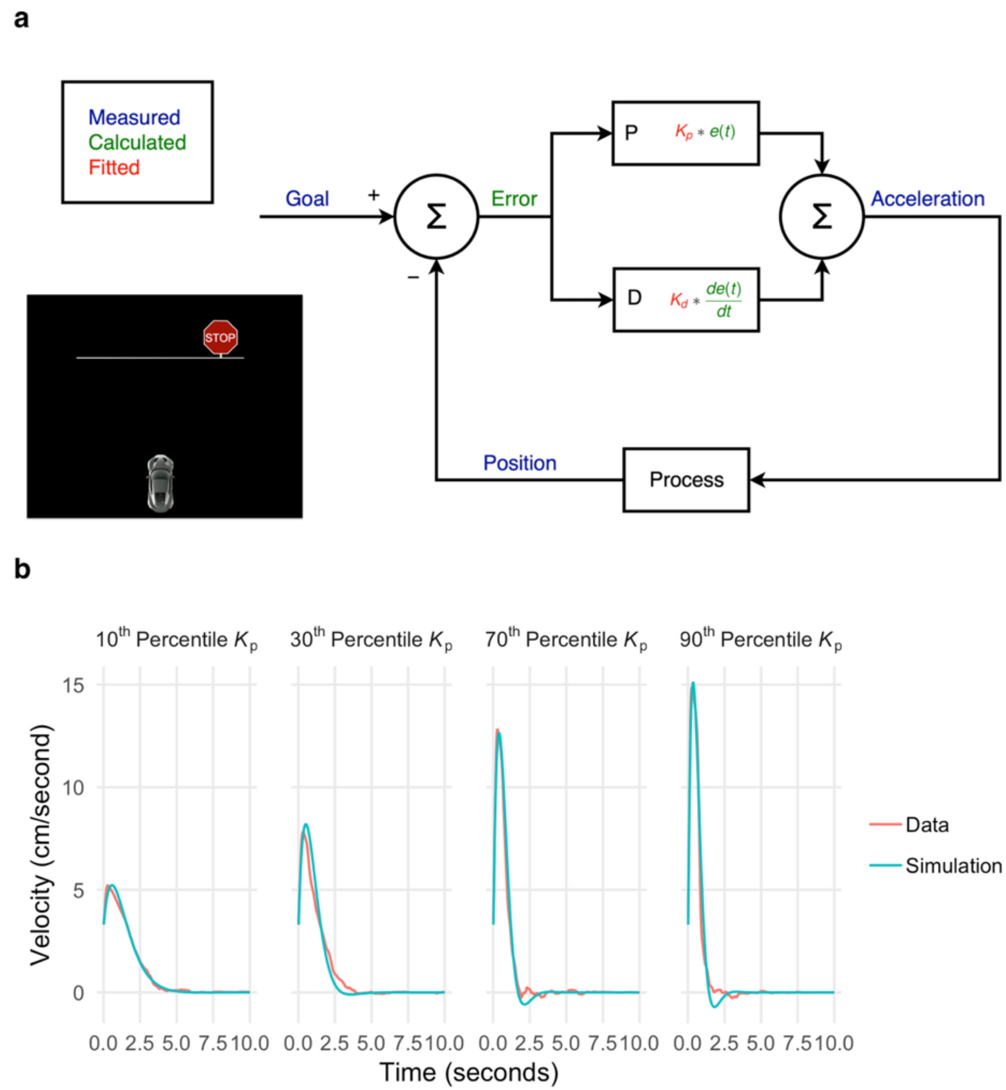

Anxiety- Related Findings

- Subjects reporting high levels of fear displayed:
  - decreased weighting on current error (consistent with inhibited goal approach)
  - and also decreased weighting on the rate of change of error (leading to overcorrecting oscillations around the goal).
- These findings were specific to fear after controlling for general negative affect.
- The experimental approach is easy, robust, and yields reliable motor trajectories and can be conducted on mobile platforms.

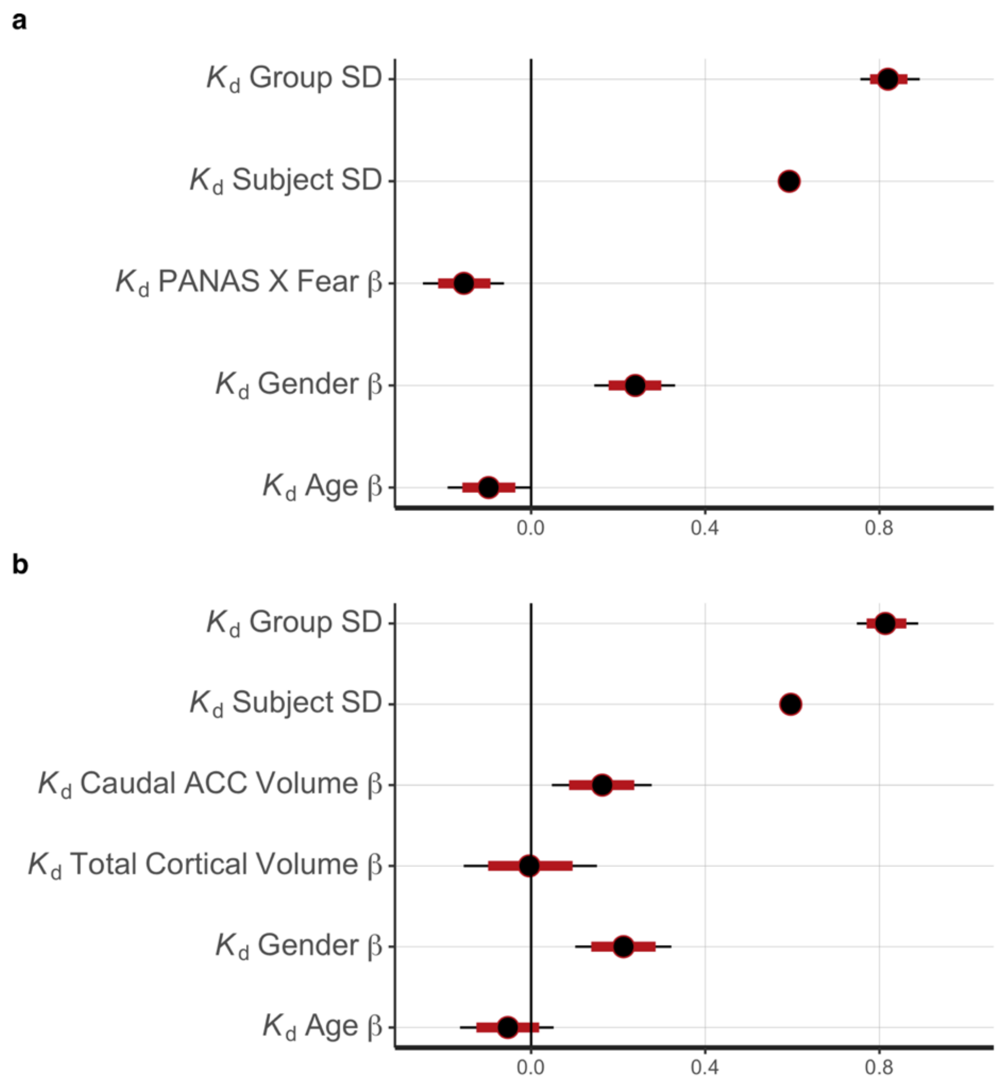

Conclusions

- Using a proportion-integral-derivative control framework we can parse **altered error control in individuals with anxiety-related problems**:
  - Anxious individual underestimate the error of current motor actions consistent with increased inhibition
  - Anxious individuals underestimate the rate of change of the error which results in oscillatory behavior ("should I stay or should I go now")

  - These parameters have direct relevance to treatment targets in behavioral interventions.
    - Direct brain modulation to increase error sensitivity.

Evidence for Slower Updating of Visual Expectations in Anxious Individuals

Jonathan Howlett

Background

- **Surprising events** are important **sources of internal model updating** which adjusts expectations for how we perceive available options and select among them.
- Based on previous work, we hypothesized that anxious individuals experienced **exaggerated surprise to predictable events, which imbues them with undue salience**.
- We applied a hybrid Rescorla Wagner (RW)/Drift Diffusion Model (DDM) to a change point detection task in a transdiagnostic group of individuals with mood and anxiety disorders.

Model Approach

- The model assumes that expectations regarding target location influences both:
  - the initial location choice on a trial and
  - the response and reaction time to the random-dot stimulus
- The updating of location expectations based on the true target location on each trial was modeled using an RW model.
- RW expectations influenced either the DDM bias parameter, DDM drift rate parameter, or both.

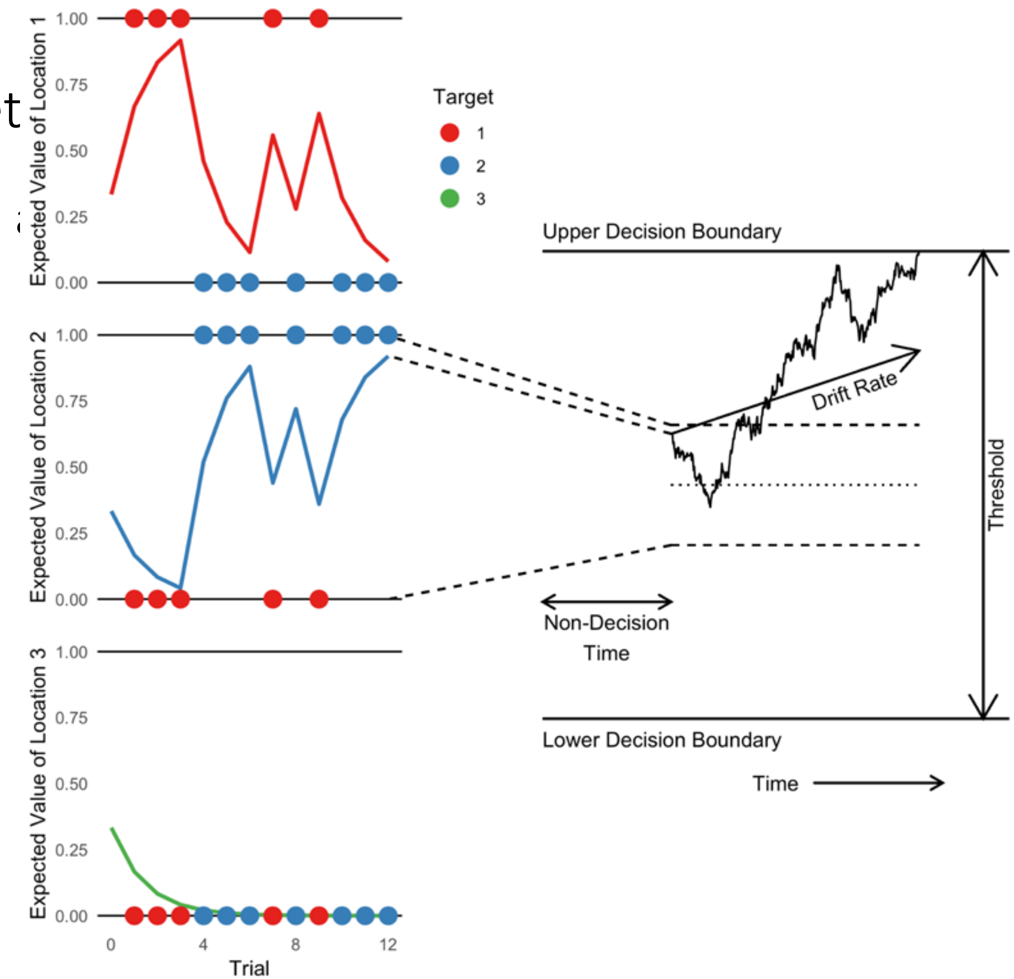

Model Comparisons

- We performed a model comparison of six models:
  - bias-only single α model,
  - bias-only dual α model,
  - drift-only single α model,
  - drift-only dual α model,
  - bias and drift single α model,
  - and bias and drift dual α model.
- All models predicted both categorical location choices and random-dot reaction times

Hierarchical Statistical Model

- To determine the relationship between fear and model parameters, we constructed a hierarchical model in which both subject-level learning rates depended on scaled age, gender, and PANAS X Fear.

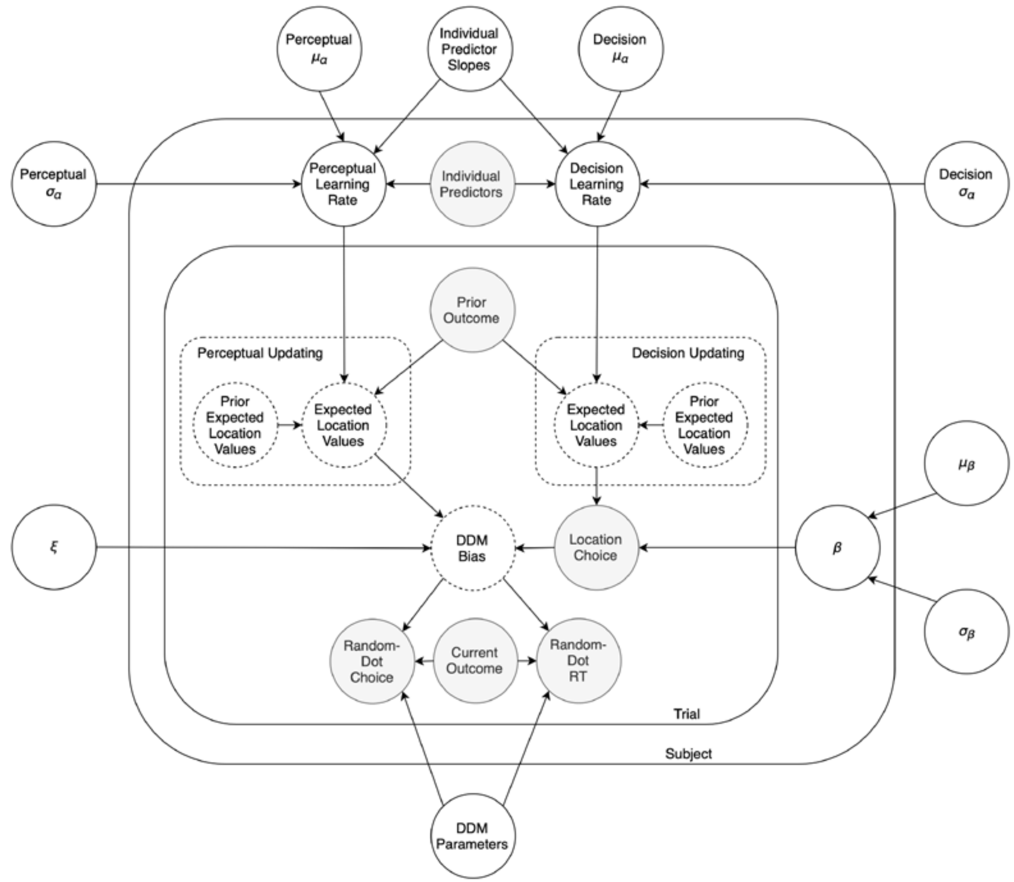

Results

- Model comparison using WAIC indicated that the bias and drift dual α model provided the best fit for the observed data.
- Individuals who reported the highest fear scores showed the lowest rate of perceptual updating
- Older individuals showed slower perceptual but not decisional updating.
- For the decision learning rate, median ICC was .62
- For the perceptual learning rate, median ICC was .80

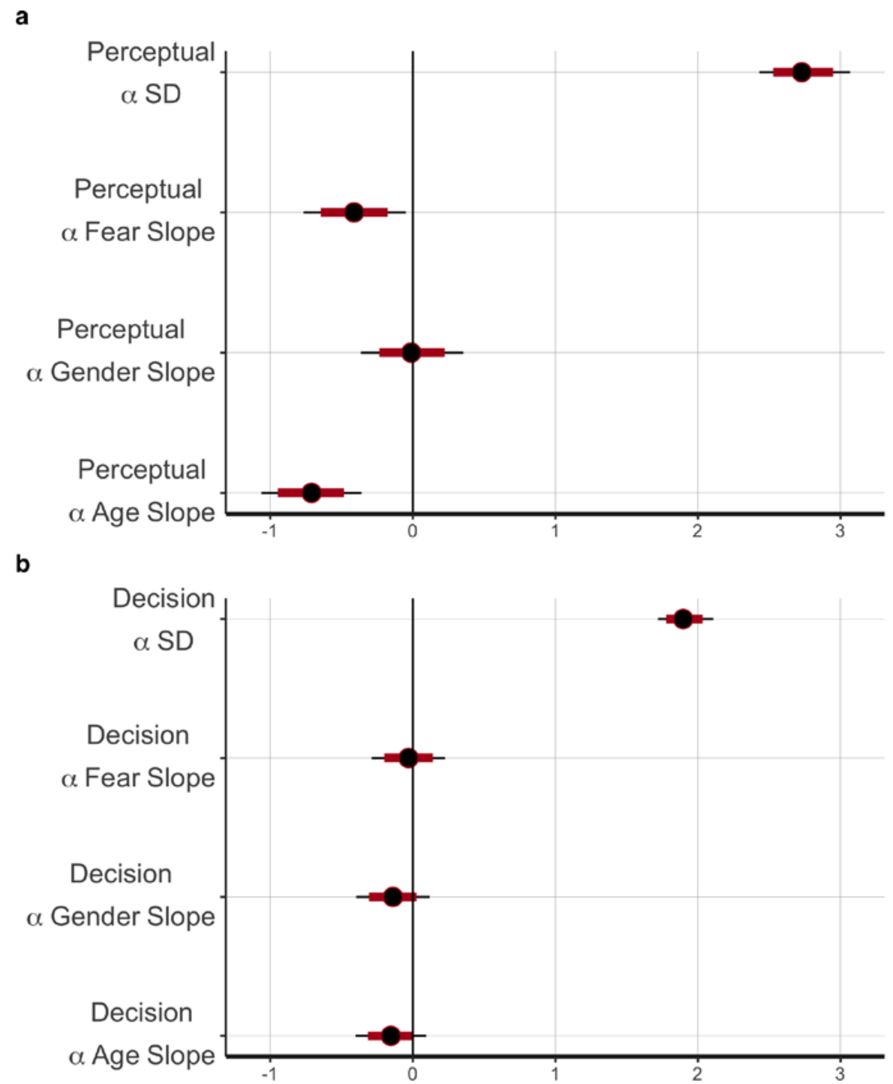

Conclusions

- Anxious (and older) individuals exhibit **slower updating of the internal model** that influences perceptual processing, but not the model that influences decision-making.
- The two models employ separate updating processes with separate learning rates (a decision learning rate and a perceptual learning rate), which are only weakly correlated.
- Taken together, **anxious individuals have difficulty updating their expectations relayed to perceptual circuits**, rather than those relayed to decision-making circuits.

Greater decision uncertainty but not emotional conflict during approach-avoidance conflict

Ryan Smith

Background

- Imbalances in the **decision to approach or avoid** when both positive and negative consequences are expected (i.e., approach-avoidance conflict; AAC) is often problematic in mental health conditions.
- AAC paradigms create conflict between the receipt of monetary rewards and either monetary punishments, pain or aversive affective stimuli.
- This study aimed to examine the difference between **decision uncertai**nty and the **emotional conflict** arising from an individual's relative sensitivity to negative affective stimuli vs. reward.

Approach Avoidance Conflict Task

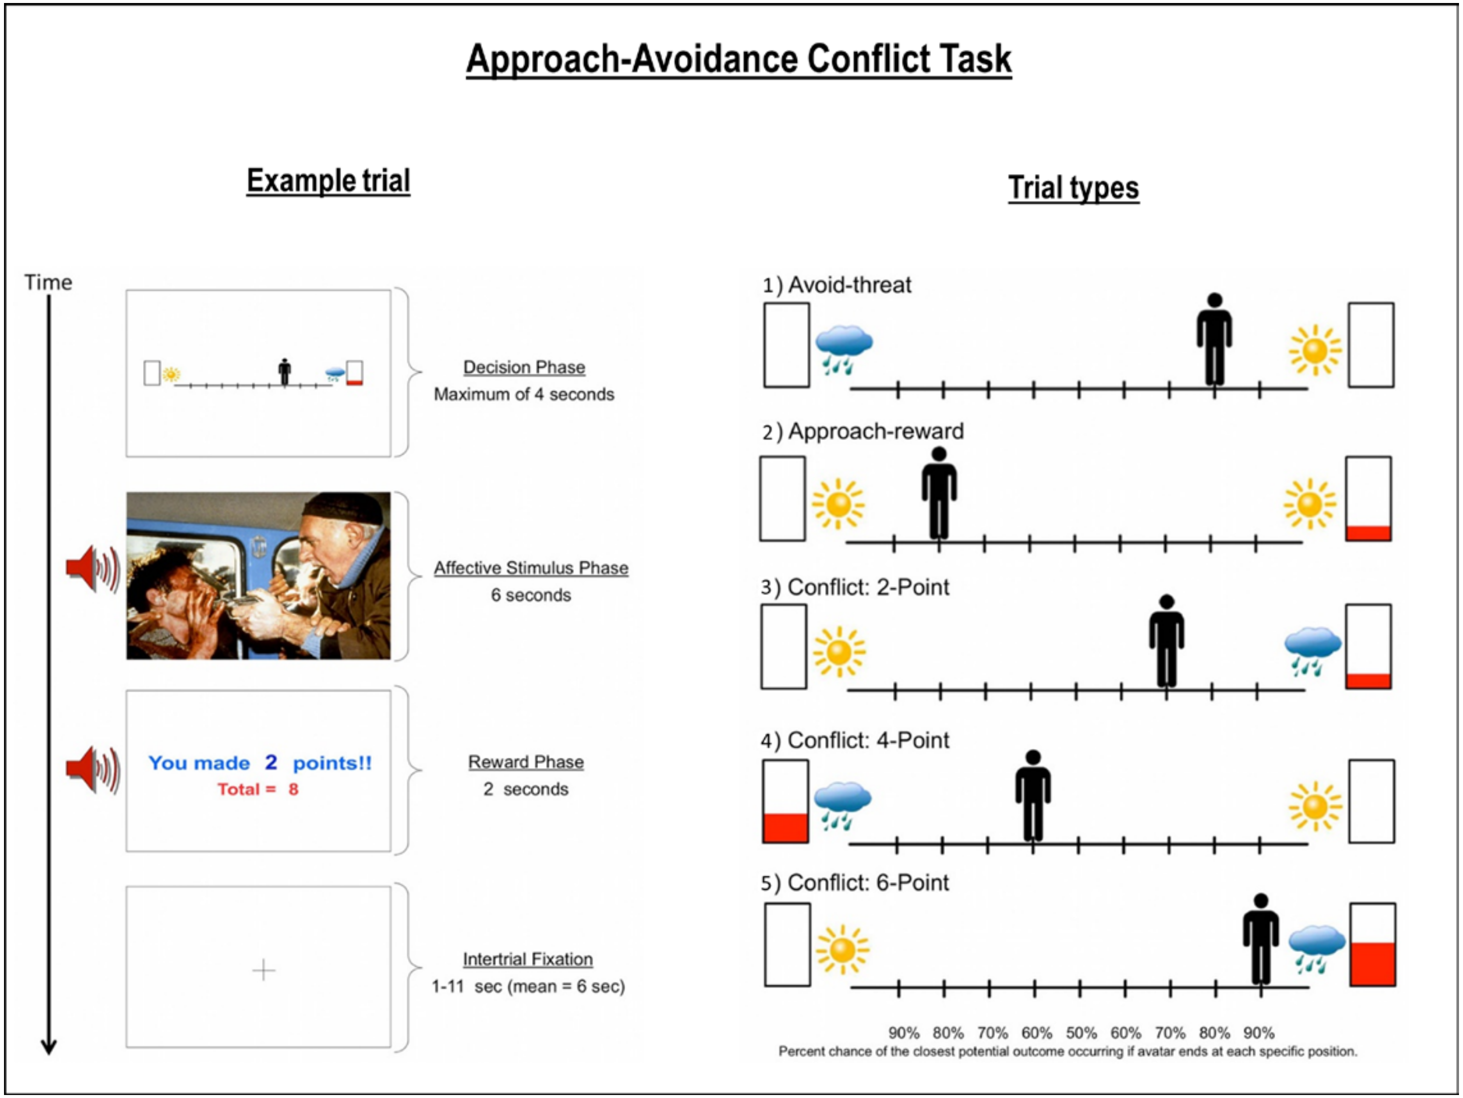

Model Approach

- Active Inference Model:
  - A: Relationship between Observations (o) and hidden states (s)
  - B: Relationship between current and previous states.
  - C: The prior preferences of the agent.
- Free parameters:
  - β decision uncertainty
  - EC emotional conflict

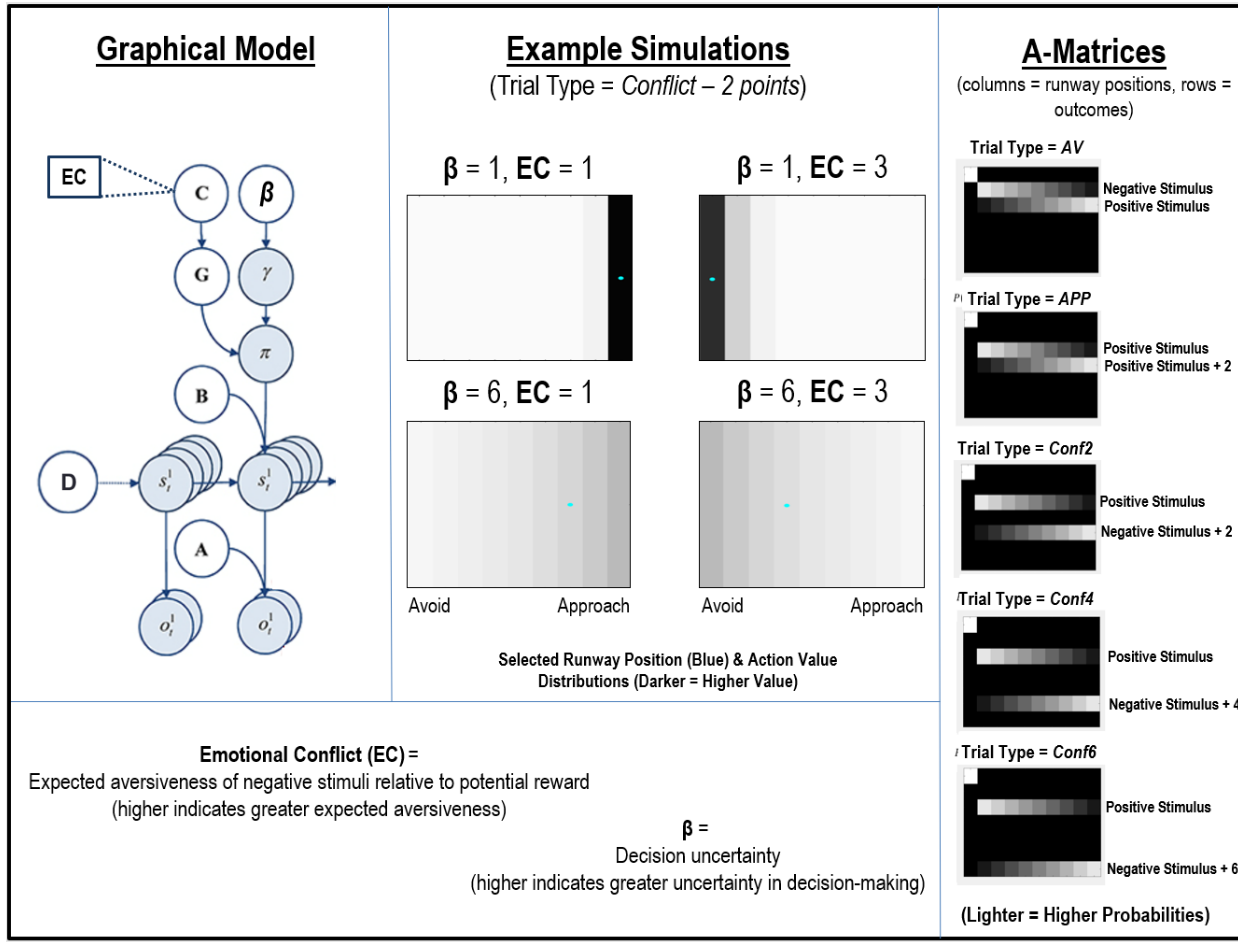

Population

Full Sample	HCs (N = 59)	DEP/ANX (N = 260)	SUDs (N = 159)	p value
Age	32.14 (11.13)	35.89 (11.30)	33.93 (9.09)	0.024
Sex (male)	28 (48%)	70 (27%)	74 (47%)	<0.001
PHQ	0.90 (1.36)	12.63 (5.14)	6.50 (5.66)	<0.001
OASIS	1.27 (1.88)	9.80 (3.42)	5.78 (4.66)	<0.001
DAST-10	0.12 (0.38)	0.67 (1.41)	7.48 (2.20)	<0.001
WRAT	62.37 (5.06)	63.53 (4.76)	58.49 (5.65)	<0.001
Propensity Matched	HCs (N = 59)	DEP/ANX (N = 161)	SUDs (N = 56)	P value
Age	32.14 (11.13)	35.11 (10.84)	32.67 (10.26)	0.119
Sex (male)	0.47 (0.50)	0.25 (0.44)	0.62 (0.49)	<0.001
PHQ	0.90 (1.36)	12.64 (5.38)	7.95 (6.50)	<0.001
OASIS	1.27 (1.88)	9.78 (3.42)	6.80 (5.15)	<0.001
DAST-10	0.12 (0.38)	0.62 (1.26)	7.45 (2.65)	<0.001
WRAT	63.53 (4.76)	62.58 (4.53)	61.89 (4.43)	0.15

Results

- Individuals with depression and anxiety related problems show greater uncertainty in decision-making relative to comparison subjects.

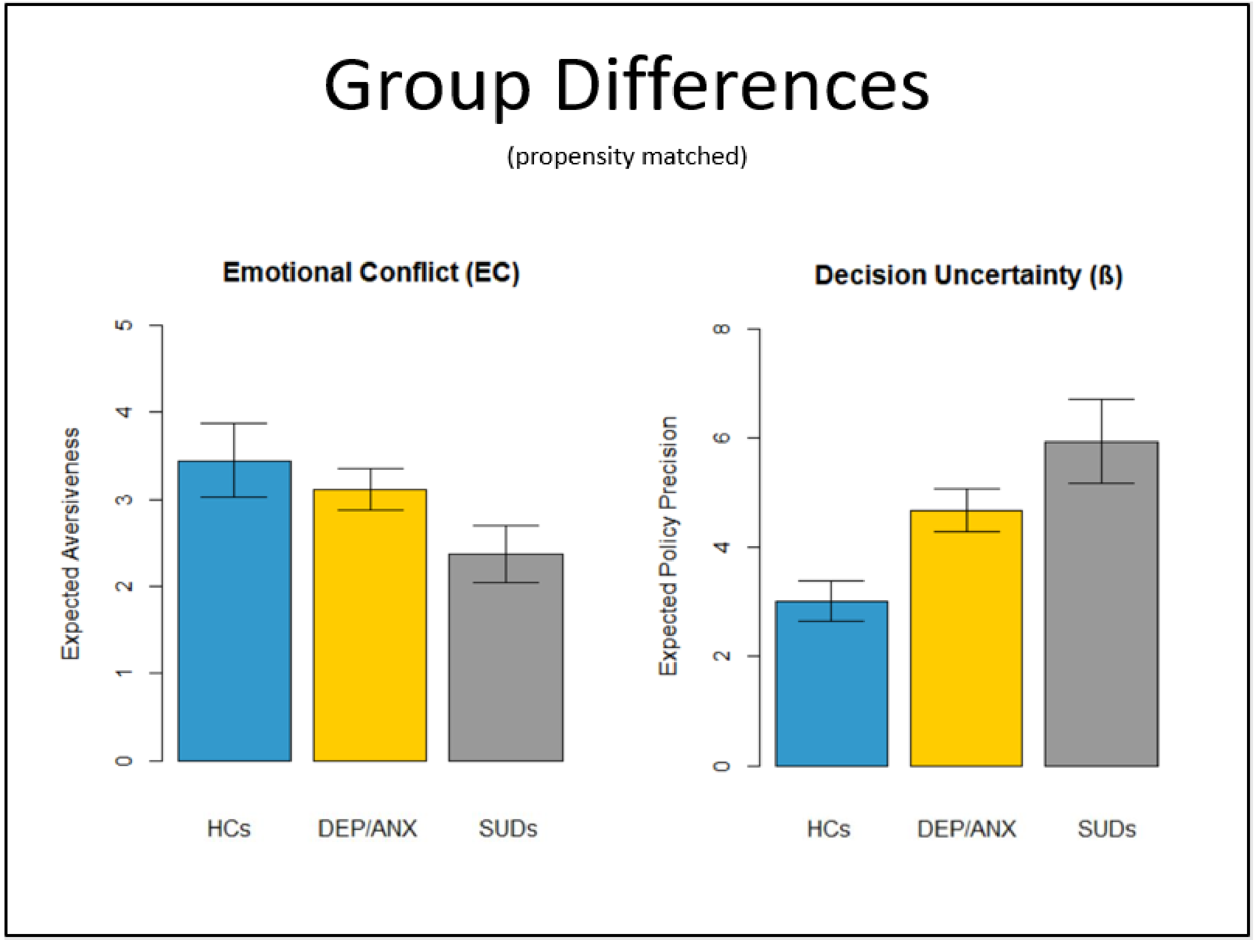

Relationship to subjective report

Post-Task Self-Report Questions (Likert Scale: 1 = not at all; 7 = very much)	EC	β
1. I found the POSITIVE pictures enjoyable:	.07	.02
2. The NEGATIVE pictures made me feel anxious Or uncomfortable:	.32**	.06
3. I often found it difficult to decide which outcome I wanted:	.10*	.45**
4. I always tried to move ALL THE WAY TOWARDS the outcome with the LARGEST REWARD POINTS:	-.74**	-.48**
5. I always tried to move ALL THE WAY AWAY FROM the outcome with the NEGATIVE PICTURE/SOUNDS:	.67**	.37**
6. When a NEGATIVE picture and sound were displayed, I kept my eyes open and looked at the picture:	-.37**	-.17**
7. When a NEGATIVE picture and sound were displayed, I tried to think about something unrelated to the picture to distract myself:	.29**	.11*
8. When a NEGATIVE picture and sound were displayed, I tried other strategies to manage emotions triggered by the pictures	.32**	.05

Conclusions

- The model showed high accuracy in predicting behavior.
- Parameter estimates showed strong relationships with RTs and participants’ self-reported feelings/motivations during the task.
- EC was uniquely associated with self-reported anxiety on the task.
- β was uniquely associated self-reported difficulty making decisions on the task.
- EC and β were not highly correlated and showed distinct relationships with psychopathology.

General Conclusions

- Computational Failure Modes in Anxiety:
  - Attenuated Error Control
  - Attenuated Updating of incoming sensory information
  - Exaggerated processing of uncertainty

- These failure modes can be:
  - Readily assessed with behavioral paradigms
  - Associated with distinct neural circuits
  - Used to develop specific failure mode interventions

Applied Computational Psychiatry:

A Roadmap for the Development of Applied Computational Psychiatry

Martin P. Paulus, Quentin J.M. Huys, Tiago V. Maia

DOI: http://dx.doi.org/10.1016/j.bpsc.2016.05.001

Michael Browning

Goals

- Identify mechanistically interpretable parameters.
  - “how does the system work?”
- Integrate measurements across units of analysis.
  - “relate behavior to circuits”
- Classify individuals into different classes.
  - “separate health from pathology”
- Predict class membership (current and future)
  - “identify individuals at risk for bad outcomes”

Future Directions

	Preclinical	Phase I(a)	PhaseI(b)	Phase II	Phase III	Phase IV
Drug Development Analog	Target (a) identification, (b) optimization	Safety / Tolerability	experimental medicine / target engagment	Small Scale Efficacy	Large Scale Efficacy	Post-marketing
ORBIT(61)		Define	Refine	Proof of Concept / Pilots	Efficacy Trial	Effectiveness
Time Line	Discovery (1-6 years)		Development (6 - 12 years)			
Goals	**"to identify probe(s) / measure(s) / model(s) / intervention(s)"**	**"to establish a reliable / robust probe(s) / measure(s) / model(s) / intervention(s)"**	**"to establish target process and engagement / model application / intervention engagement"**	**"to establish clinical efficacy and validity"**	**"to confirm clinical validity and demonstrate outcome improvement"**	**"new applications"**
Stages	Identification		Validation		Launch Readiness / Release	
Population	Healthy Volunteers (HV)	HV	HV, Target Population(s) TP	TP	TP	new TP
Study Type	cross-sectional (cs)	cs, longitudinal (l)	cs, l, experimental design(s)	Randomized Controlled Trial (RCT)	RCT	cs, l, RCT
Sites	single / few sites	single - multi-site	single / few sites	single - multi-site	multi-site	single / few sites
Study Size	small n	small to large n	medium n	large n	large n	small n

A Roadmap for the Development of Applied Computational Psychiatry

Martin P. Paulus, Quentin J.M. Huys, Tiago V. Maia

DOI: http://dx.doi.org/10.1016/j.bpsc.2016.05.001

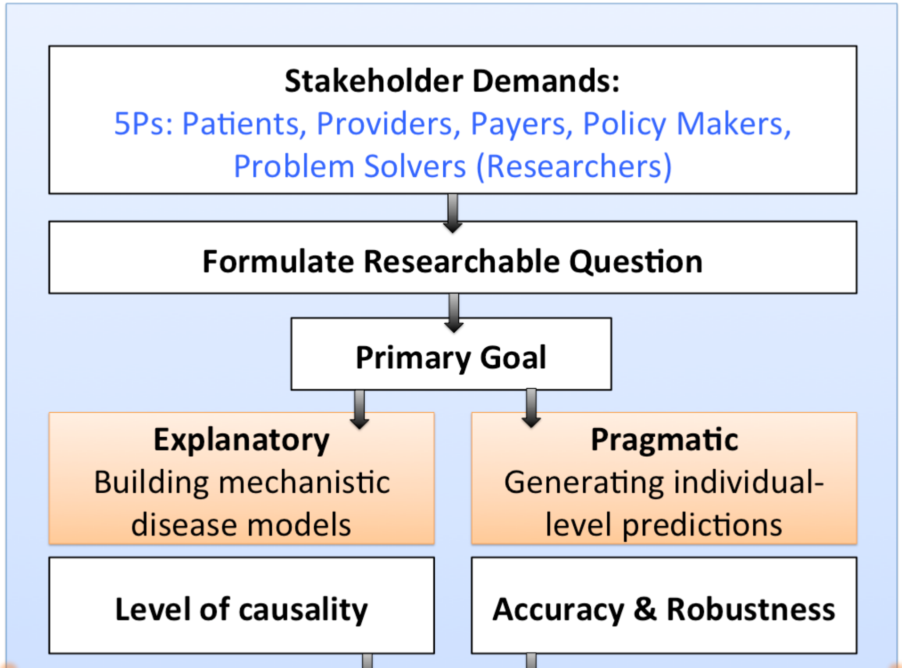

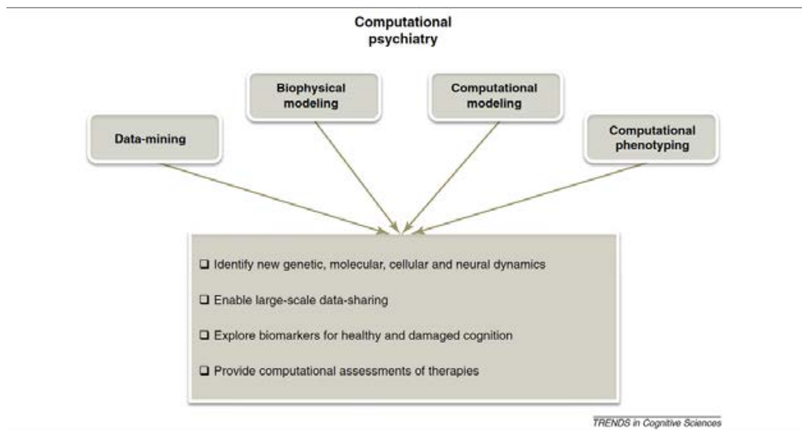

Computational psychiatry

P. Read Montague[1,2], Raymond J. Dolan[2], Karl J. Friston[2] and Peter Dayan[3]

[1] Virginia Tech Carilion Research Institute and Department of Physics, Virginia Tech, 2 Riverside Circle, Roanoke, VA 24016, USA
[2] Wellcome Trust Centre for Neuroimaging, University College London, 12 Queen Square, London, WC1N 3BG, UK
[3] Gatsby Computational Neuroscience Unit, Alexandra House, 17 Queen Square, London, WC1N 3AR, UK