Axiomatic Models in Economics and Neuroscience

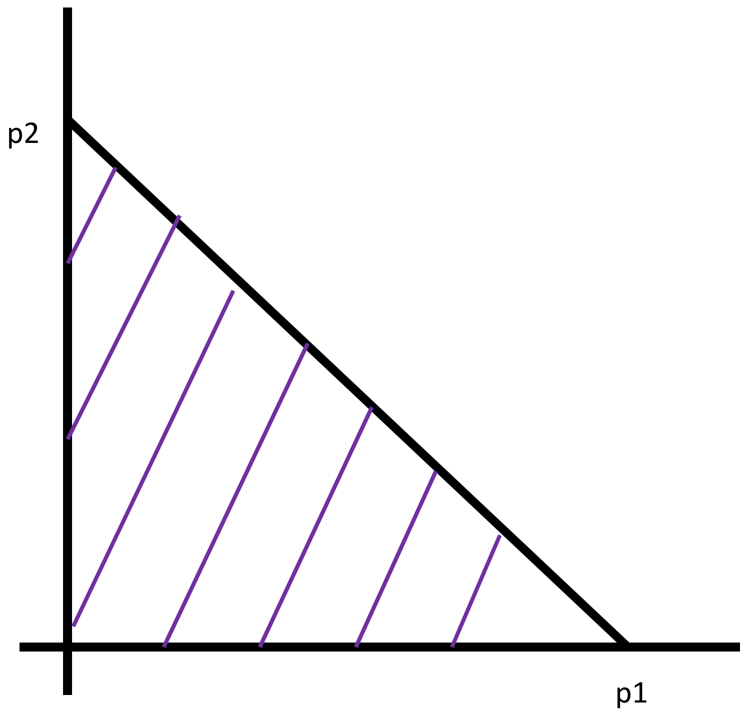

Society for Neuroeconomics
October 2010

Mark Dean
Department of Economics

Brown University

Introduction

- Economists have an antisocial habit
- Sometimes, we describe our models in terms of axioms
  - A set of logical propositions that are equivalent to that model
- This has been considered
  - Dogmatic/anti scientific
  - Overly complicated
  - Overly simplistic
  - The pointless ramblings of an mathematically-obsessed discipline that is completely divorced from reality

My Claim

- These criticisms are not *always* justified
- Axioms provide a handy approach for addressing problems within economics
  - How to test our models
- This approach has aided model development
- Neuroscience now facing some of the same problems

Outline

- Example 1: Utility Maximization
- Discussion:
  - What just happened?
  - Why was it a good idea?
- Example 2: Reward Prediction Error

Example 1: Utility Maximization

How do we tell if people are maximizing utility?

Stage 1: The Data

- We observe:
  - The choices someone makes
  - What they were choosing from
- Example: choices from different sets of snack foods
  - Cf. Rangel Labs 2009, 2010, 2010, 2011, etc, etc ad nauseum

Available Snacks	Chosen Snack
Jaffa Cakes, Kit Kat	Jaffa Cakes
Kit Kat, Lays	Kit Kat
Lays, Jaffa Cakes	Jaffa Cakes
Kit Kat, Jaffa Cakes, Lays	Jaffa Cakes

Stage 2: The Model

- We want to test the model of utility maximization
- Every object has a fixed utility value attached to it
- For example:
  - U(jaffa cakes)=10
  - U(kit kat) =5
  - U(lays)=2
- In any choice set, choose the object with highest utility

The Question

- Is our data set consistent with the model of utility maximization?
- Problem: Our model contains ‘unobservables’
  - We do not observe utilities
  - Kit Kats do not come with utility numbers stamped on them
  - Model says that people maximize utility, but I experimenter does not observe utility
- How can we proceed?

Two Approaches

- The 'Standard Approach':
  - Pick a particular utility function
  - Test whether this utility function can explain the data
  - Failure rules out that utility function
- The 'Axiomatic Approach':
  - Ask whether there exists any utility function that can explain the data?
  - Identify patterns of data that cannot be explained by any utility function
  - Failure rules out utility maximization

Algenon's Choices

Choice	Available Snacks	Chosen Snack
1	Jaffa Cakes, Kit Kat	Jaffa Cakes
2	Kit Kat, Lays	Kit Kat
3	Lays, Jaffa Cakes	Lays
4	Kit Kat, Jaffa Cakes, Lays	Jaffa Cakes

- Is there any utility function that can explain Algenon's choices
- No!
  - Choice 1 implies u(jaffa cake)>u(kit kat)
  - Choice 2 implies u(kit kat)>u(lays)
  - Choice 3 implies u(lays)>u(jaffa cakes)
- Implies u(jaffa cake)>u(jaffa cake): Contradiction

Brittney's Choices

Choice	Available Snacks	Chosen Snack
1	Jaffa Cakes, Kit Kat	Jaffa Cakes
2	Kit Kat, Lays	Kit Kat
3	Lays, Jaffa Cakes	Jaffa Cakes
4	Kit Kat, Jaffa Cakes, Lays	Kit Kat

- What about Brittney's Choices?
- No!
  - Choice 1 implies u(jaffa cake)>u(kit kat)
  - Choice 4 implies u(kit kat)>u(jaffa cakes)
- Contradiction

Colvin's Choices

Choice	Available Snacks	Chosen Snack
1	Jaffa Cakes, Kit Kat	Jaffa Cakes
2	Kit Kat, Lays	Kit Kat
3	Lays, Jaffa Cakes	Jaffa Cakes
4	Kit Kat, Jaffa Cakes, Lays	Jaffa Cakes

- How about Colvin's Choices?
- Yes!
  - u(jaffa cakes)>u(kit kat)>u(lays)
- Eg
  - u(jaffa cakes)=3
  - u(kit kat)=2
  - U(lays)=1

A General Rule

- Question: Is there a general rule that differentiates data sets that can be explained by some utility function from those that can't?

A General Rule

- Question: Is there a general rule that differentiates data sets that can be explained by some utility function from those that can't?

The Independence of Irrelevant Alternatives

Say x is chosen from a set of alternatives A

B is a subset of A that contains x

Then x must be chosen from B

A General Rule

The Independence of Irrelevant Alternatives

Say x is chosen from a set of alternatives A

B is a subset of A that contains x

Then x must be chosen from B

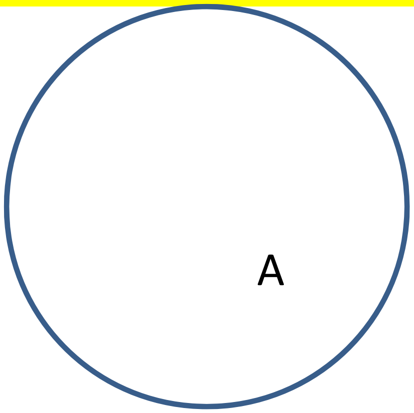

A General Rule

The Independence of Irrelevant Alternatives

Say x is chosen from a set of alternatives A

B is a subset of A that contains x

Then x must be chosen from B

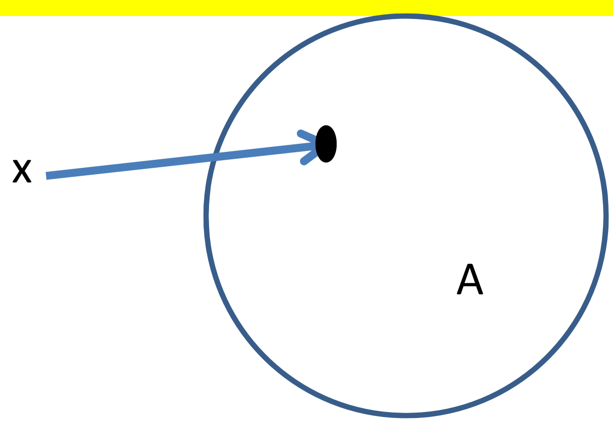

A General Rule

The Independence of Irrelevant Alternatives

Say x is chosen from a set of alternatives A

B is a subset of A that contains x

Then x must be chosen from B

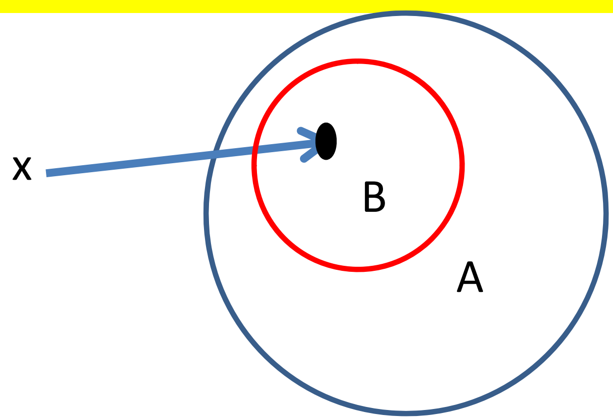

Independence of Irrelevant Alternatives

Choice	Available Snacks	Chosen Snack
1	Jaffa Cakes, Kit Kat	
2	Kit Kat, Lays	
3	Lays, Jaffa Cakes	
4	Kit Kat, Jaffa Cakes, Lays	

- In our example, whatever is chosen in set 4 must always be chosen when it is available

Algenon's Choices

Choice	Available Snacks	Chosen Snack
1	Jaffa Cakes, Kit Kat	Jaffa Cakes
2	Kit Kat, Lays	Kit Kat
3	Lays, Jaffa Cakes	Lays
4	Kit Kat, Jaffa Cakes, Lays	Jaffa Cakes

- Algenon's choices violate these condition
  - – Jaffa cakes chosen in set 4
  - – Lays chosen in set 3

Brittney's Choices

Choice	Available Snacks	Chosen Snack
1	Jaffa Cakes, Kit Kat	Jaffa Cakes
2	Kit Kat, Lays	Kit Kat
3	Lays, Jaffa Cakes	Jaffa Cakes
4	Kit Kat, Jaffa Cakes, Lays	Kit Kat

- Also violated by Brittney's choices
  - Kit Kat chosen in set 4
  - Jaffa cakes chosen in set 1

Colvin’s Choices

Choice	Available Snacks	Chosen Snack
1	Jaffa Cakes, Kit Kat	Jaffa Cakes
2	Kit Kat, Lays	Kit Kat
3	Lays, Jaffa Cakes	Jaffa Cakes
4	Kit Kat, Jaffa Cakes, Lays	Jaffa Cakes

- Colvin’s choices satisfy IIA
  - Jaffa cakes chosen in 4
  - Also chosen in 3 and 1

A Necessary Condition

The Independence of Irrelevant Alternatives

Say x is chosen from a set of alternatives A

B is a subset of A that contains x

Then x must be chosen from B

- If we observe a utility maximizer, then they must satisfy IIA
  - If x is chosen from A, must have a higher utility than anything in A
  - B is a subset of A
  - X must have higher utility than anything in B
  - Should be chosen from B

A Sufficient Condition

The Independence of Irrelevant Alternatives

Say x is chosen from a set of alternatives A

B is a subset of A that contains x

Then x must be chosen from B

- If IIA holds then subject is a utility maximizer*
  - There exists some utility function such that choices maximize utility according to that utility function
  - Not at all obvious

* Assuming that the set of all objects is finite, and we see choices from every subset

A (very quick) Outline of the Proof

- Look at binary choices i.e. between two objects x and y
- Define a binary preference relation P as xPy if is chosen when offered a choice from x and y
- Independence of Irrelevant Alternatives ensures that
  - P is transitive (xPy and yPz implies xPz)
  - P represents choices
    - Take any set of alternatives A={x,y,z,w..}
    - If x is chosen from A then xPy, xPz, xPw…
- Any complete, transitive preference relation (on a finite set) can be represented by a utility function
  - u(x)>=u(y) if and only if xPy

Our First Representation Theorem

For any set of choices*, there is a utility function such that the chosen object is the available option with the highest utility

if and only if

those choices satisfy the independence of irrelevant alternatives

* Assuming that the set of all objects is finite, and we see choices from every subset

Our First Representation Theorem

For any set of choices*, there is a utility function such that the chosen object is the available option with the highest utility

if and only if

those choices satisfy the independence of irrelevant alternatives

- Testing the model of utility maximization is the same thing as testing IIA

* Assuming that the set of all objects is finite, and we see choices from every subset

A Note on Uniqueness

- If choices satisfy the independence of irrelevant alternatives, is the utility function that explains those choices unique?

Choice	Available Snacks	Chosen Snack
1	Jaffa Cakes, Kit Kat	Jaffa Cakes
2	Kit Kat, Lays	Kit Kat
3	Lays, Jaffa Cakes	Jaffa Cakes
4	Kit Kat, Jaffa Cakes, Lays	Jaffa Cakes

- No! Take Colvin’s Choices
  - u(jaffa cakes)=3, u(kit kat)=2, u(lays)=1
  - u(jaffa cakes)=354, u(kit kat)=0, u(lays)=-110020

A Note on Uniqueness

- Utilities identified up to a strictly positive transformation
  - i.e. if u represents choices, and v is another utility function such that $v(x)=T(u(x))$ for some strictly increasing function, then v will also represent choices
- Any utility function that preserves ordering will do the trick
- No point arguing about whether
  - Utility of x is twice as big as the utility of y
  - Utility of x minus utility of y is 7
  - Utility of x is negative

Discussion

What did we just do, and did it make sense

What Did We Just Do?

1. We defined a class of models
   - Utility maximization is a model class
2. We identified a data set on which to test this class
   - Choices from different set of objects
3. We identified a set of necessary and sufficient conditions on the data set for a model in the class to explain data
   - If and only if the independence of irrelevant alternatives holds there is a utility functions that can explain the data
4. We described the size of the subset of the class of models that can explain the data
   - Utilities unique up to a strictly positive transformation

Is This a Useful Strategy in General?

- We can test the whole class of models
- Provides exact implications of the model
- Tells us how ‘seriously’ to take unobservables

Is This a Useful Strategy in General?

- We can test the whole class of models
- Provides exact implications of the model
- Tells us how ‘seriously’ to take unobservables

Testing the Class of Models

- Standard approach: pick one element in the class and test that
  - i.e. define utility in some particular way, and test that
- Or some subset of class of models and test them
  - e.g. Assume utility is a function of some properties of the object
  - Estimate parameters
  - Take that as your candidate utility function

The ‘Standard’ Approach

- Sometimes this is easy to do
  - E.g. choosing over amounts of money
- But how do we define utility over chocolate bars?
  - Function of size, calories, cocoa content, nougat-iness?
  - Involves making arbitrary assumptions
  - What if utility function needed to explain choices between jaffa cakes and a spaniel?

U(


)>U(

)

The ‘Standard’ Approach

- Resulting test is now of two assumptions
  - Person maximizes utility
  - Utility is as specified by researcher
- If test fails, cannot reject class of utility maximizing models, only this specific example
- Axioms mean that we don’t need to make additional assumption
  - Axioms are a substitute for assumptions about unobservable elements in our model

Why Was This A Good Idea?

- We can test the whole class of models
- Provides exact implications of the model
- Tells us how ‘seriously’ to take unobservables

Provides Exact Empirical Implications

- Axiomatic approach provides necessary and sufficient conditions for model to be true
  - If conditions hold, model is explains data
  - If they don't, it doesn't
- There are at least three reasons why this might be useful
  - Tells you if your model has any predictions
  - Tells you if the predictions are different from other models
  - Can tell you which bits of your model are driving predictions

An Example: The Satisficing Model

- Consider the alternative model of choice [Simon, 1955]
  - Search through alternatives one by one
  - Stop searching when reach an alternative that is ‘good enough’
  - Choose that alternative
- What are the predictions of this model?

Implications of Satisficing Model

- Depends on the assumptions made about the search process
- Search order can change between choices
  - Model has no predictions
  - Any set of choices can be explained this way
  - Assume that all objects are 'good enough', and that the chosen object is the first thing searched
- If search order is fixed between choices
  - Predictions are exactly the same as utility maximization
  - People act in line with the satisficing model if and only if they satisfy the independence of irrelevant alternatives

An Example: The Satisficing Model

- If choices are our only data no point arguing whether subject is satisficing or utility maximizing
  - Either assume search order is fixed, in which case predictions are the same…
  - ….or assume it can vary, in which case we can explain any choice
- (Herbert Simon knew this)

What About Regressions?

- Significant regression coefficients neither necessary nor sufficient for utility maximization
  - If coefficients are not significant, could have wrong definition of utility
  - If they are significant still could be lots of violations of utility maximization
    - Say that jaffa cakes have more calories than kit kats which have more calories that lays
    - On weekdays I am healthy, and so prefer fewer calories
    - On weekends I am unhealthy, and prefer more calories
    - Regression of choice on calories could give significant negative coefficient

Why Was This A Good Idea?

- We can test the whole class of models
- Provides exact implications of the model
- Tells us how ‘seriously’ to take unobservables

How seriously do we take unobservables?

- Axiomatic approach treats unobservables as output of rather than input to the modeling process
  - Utilities are derived from choices, rather than assumed
- This approach tells us how ‘seriously’ to take these numbers
  - For utilities – only ordering matters, actual numbers do not
  - No more information can be extracted from choices

When Axioms Go Wrong

- A common criticism of axioms: they are sensitive souls
- One 'bad' observation is enough to declare model a failure
- For example
  - we observe someone choosing jaffa cakes over kit kats every day for a year
  - On one day, they choose a kit kat over a jaffa cake
  - Violate axioms – model of utility maximization rejected

Are Axioms 'Nearly' Satisfied?

- One response
  - This is correct – utility maximizing model is an incorrect description of these choices
  - Not a very useful response - all our models will be wrong sometimes
- Can we have some measure of whether a data set is close to satisfying axioms?
  - Active area of research
  - Some measures have been developed

Example: The Houtman Maks (HM) Index

- Houtman-Maks [1985] Index
  - Largest subset of data that satisfies a set of axioms
  - e.g., remove choice observations until the remaining set of data is consistent with utility maximization
- Benchmarking [Bronars 1987]
  - Need some measure of power – was it possible for axioms to be violated in a data set?
  - Use random choice – compare HM index to the distribution generated by a population generated by random choosers

Do People Utility Maximize?

- Experimental Data: Choi, Fisman, Gale and Kariv [2007]
- Subjects made 50 choices from different budget sets
- Find largest subset of their choices that satisfy rationality

Do People Utility Maximize?

- 22 of 93 subjects rational
- 91 out of 93 subjects above the 95th percentile of random choices

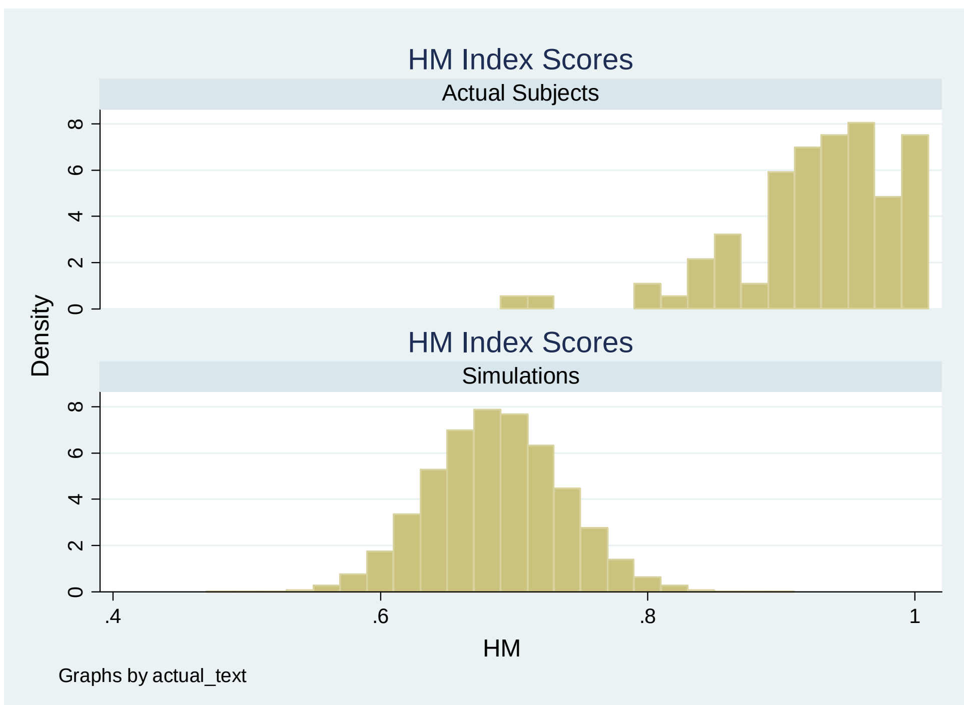

Example 2: Reward Prediction Error

Does dopamine encode a reward prediction error

Claim

- Neuroscience starting to use models with variables that are not directly observable:
  - Rewards
  - Beliefs
  - Salience
  - Incentive Salience
- Approach that has proved useful in economics may also prove useful in neuroscience
- Example: Reward Prediction Error

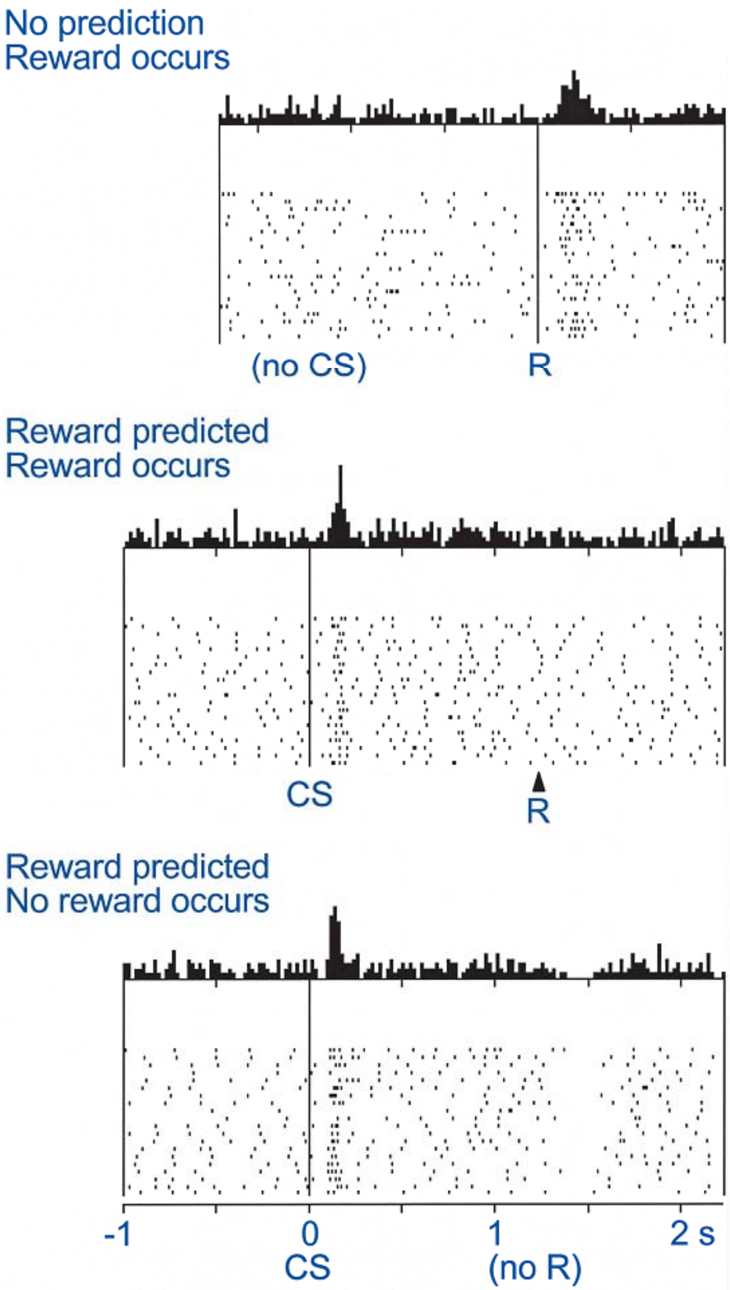

Schultz et al., Science, 1997

- Midbrain dopamine neurons are thought to encode a reward prediction error (**RPE**) used for learning

RPE = experienced reward – predicted reward

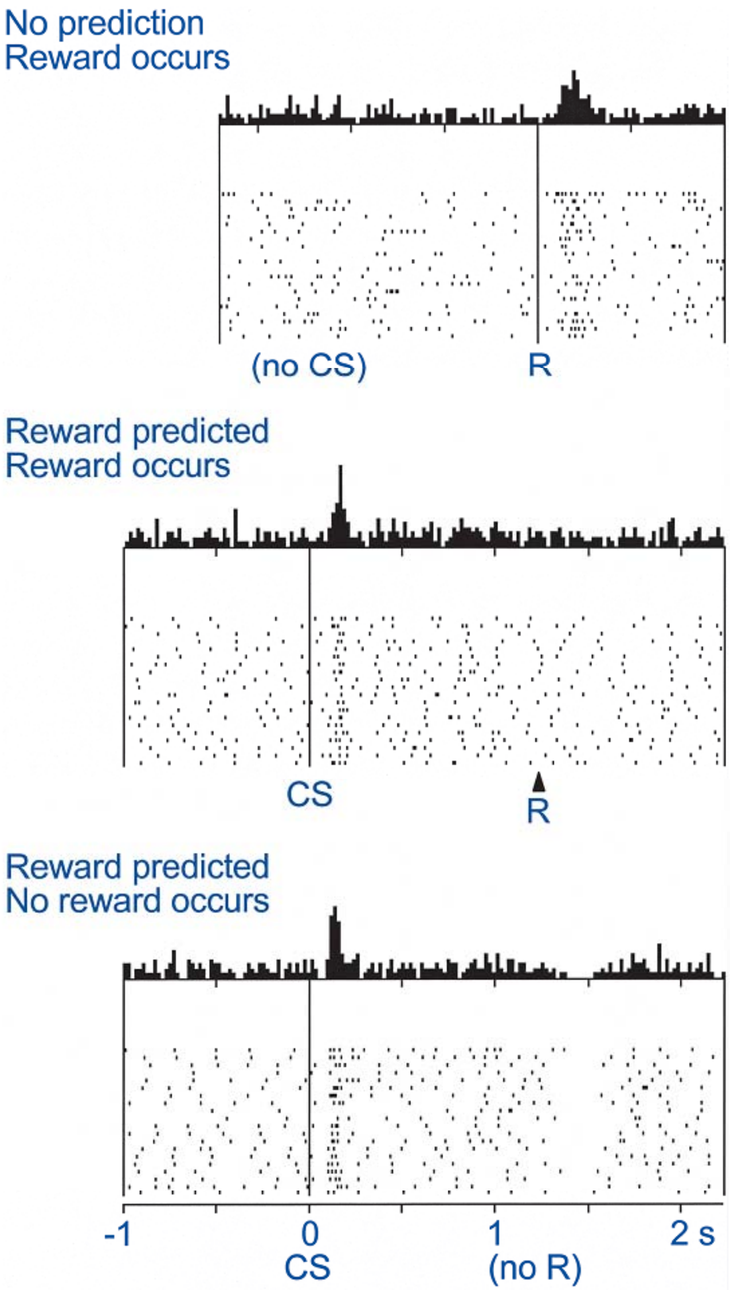

Schultz et al., Science, 1997

- Midbrain dopamine neurons are thought to encode a reward prediction error (**RPE**) used for learning

RPE = experienced reward – predicted reward

- **But:** How is ‘experienced reward’ and ‘predicted reward’ defined?

RPE = experienced reward – predicted reward

- **Experienced Reward**
  - Money?
  - Fruit juice?
  - Monotonic function?
  - Linear function?

RPE = experienced reward – predicted reward

- **Predicted Reward**
  - Bayesian Updating?
  - Reinforcement Learning?
  - Priors?
  - Parameter Values?

RPE = experienced reward – predicted reward

- **Difference**
  - Experienced Reward – Predicted Reward?
  - Log difference?

A Familiar Problem?

- RPE model contains unobservables
  - Rewards and beliefs equivalent to utility in our model of choice
- Really a class of models
  - Elements in class defined by assumptions about beliefs and rewards
- Standard approach: Pick one element in the class and test that

Example: O’Doherty et al. [2003]

- Experienced Reward:
  - Equal to subjective reported liking ratings of fruit juice

Example: O'Doherty et al. [2003]

- Experienced Reward:
  - Equal to subjective reported liking ratings of fruit juice
- Predicted Reward:
  - Temporal Difference Learning

$$V(t) = \sum_{t} w(t)x(t)$$

Example: O'Doherty et al. [2003]

- Experienced Reward:
  - Equal to subjective reported liking ratings of fruit juice
- Predicted Reward:
  - Temporal Difference Learning

$$V(t) = \sum_t w(t)x(t)$$

- Reward Prediction Error

$$\delta(t) = r(t) + \gamma V(t+1) - V(t)$$

The problem

RPE
model class

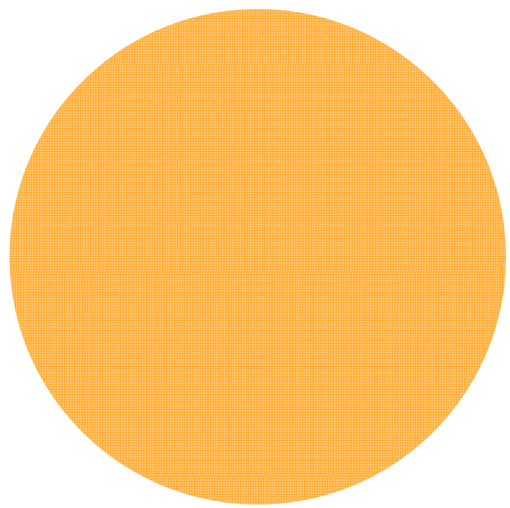

The problem

RPE
model class

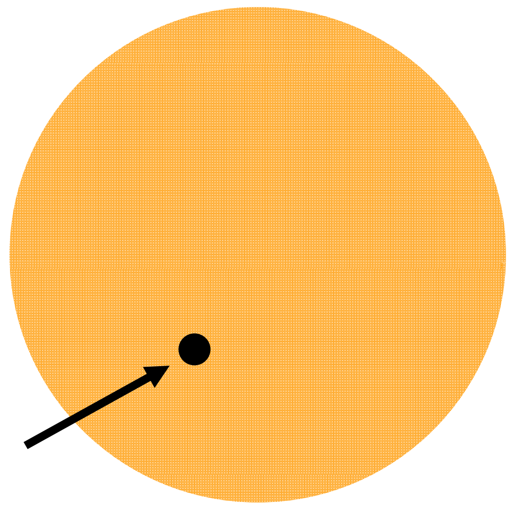

O'Doherty et al. (2003)

The problem

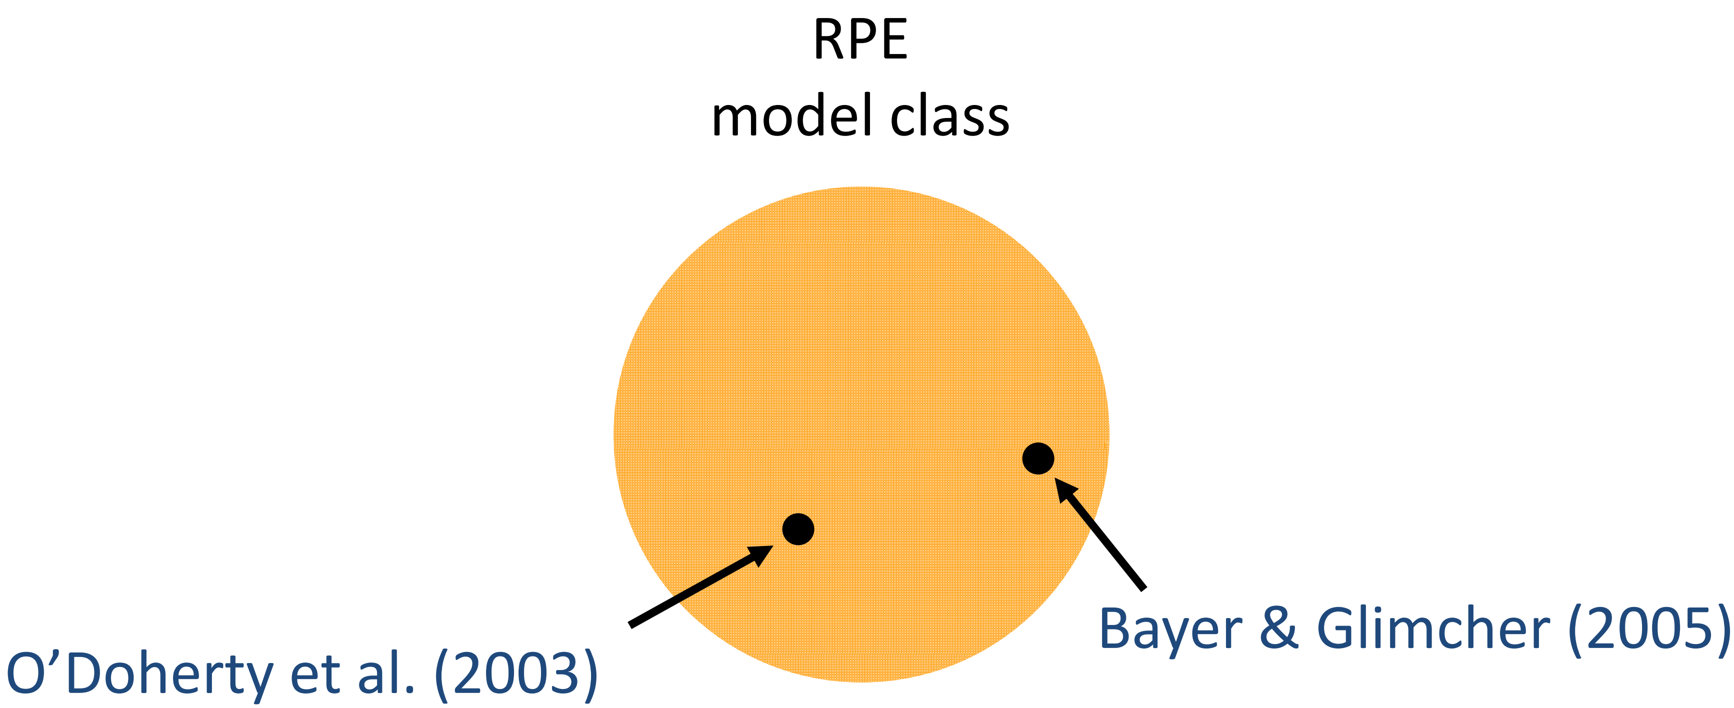

Axiomatic Approach

- The ‘Axiomatic Approach’:
  - Ask whether there exists any reward and belief functions that can explain the data?
  - Identify patterns of data that cannot be explained by any reward and belief functions
  - Failure rules out RPE

The Axiomatic Approach

- **Consistent Experienced Reward:**

$$Given\ (z, p), (z', p), (z, p'), (z', p') \in A,$$
$$\delta(z, p) > \delta(z', p) \Rightarrow \delta(z, p') > \delta(z', p')$$

- **Consistent Predicted Reward:**

$$Given\ (z, p), (z', p), (z, p'), (z', p') \in A,$$
$$\delta(z, p) > \delta(z, p') \Rightarrow \delta(z', p) > \delta(z', p')$$

- **No Surprise Equivalence:**

$$Given\ z, z' \in Z,$$
$$\delta(z, z) = \delta(z', z')$$

The Axiomatic Approach

- **Consistent Experienced Reward:**

$$Given\ (z,p),(z',p),(z,p'),(z',p') \in A,$$
$$\delta(z,p) > \delta(z',p) \Rightarrow \delta(z,p') > \delta(z',p')$$

- **Consistent Predicted Reward:**

$$Given\ (z,p),(z',p),(z,p'),(z',p') \in A,$$
$$\delta(z,p) > \delta(z,p') \Rightarrow \delta(z',p) > \delta(z',p')$$

- **No Surprise Equivalence:**

$$Given\ z,z' \in Z,$$
$$\delta(z,z) = \delta(z',z')$$

Are exactly equivalent to the RPE hypothesis

The Axiomatic Approach

- **Consistent Experienced Reward:**

$$Given\ (z,p),(z',p),(z,p'),(z',p') \in A,$$
$$\delta(z,p) > \delta(z',p) \Rightarrow \delta(z,p') > \delta(z',p')$$

- **Consistent Predicted Reward:**

$$Given\ (z,p),(z',p),(z,p'),(z',p') \in A,$$
$$\delta(z,p) > \delta(z,p') \Rightarrow \delta(z',p) > \delta(z',p')$$

- **No Surprise Equivalence:**

$$Given\ z,z' \in Z,$$
$$\delta(z,z) = \delta(z',z')$$

Are exactly equivalent to the RPE hypothesis

Testing class of RPE models the same as testing axioms

Environment

- Observe dopamine activity when prizes are obtained from lotteries

Environment

- Observe dopamine activity when prizes are obtained from lotteries
  - Prize: Win $5, Lose $5

Environment

- Observe dopamine activity when prizes are obtained from lotteries
    - Prize: Win $5, Lose $5
    - Lottery : e.g. 25% chance of winning $5, 75% chance of losing $5

Environment

- Observe dopamine activity when prizes are obtained from lotteries
    - Prize: Win \$5, Lose \$5
    - Lottery : e.g. 25% chance of winning \$5, 75% chance of losing \$5
- Allows us to vary 'experienced reward' and 'predicted reward' independently

Environment

- Observe dopamine activity when prizes are obtained from lotteries
    - Prize: Win \$5, Lose \$5
    - Lottery : e.g. 25% chance of winning \$5, 75% chance of losing \$5
- Allows us to vary ‘experienced reward’ and ‘predicted reward’ independently
    - Varying prizes: varying experience reward

Environment

- Observe dopamine activity when prizes are obtained from lotteries
    - Prize: Win $5, Lose $5
    - Lottery : e.g. 25% chance of winning $5, 75% chance of losing $5
- Allows us to vary 'experienced reward' and 'predicted reward' independently
    - Varying prizes: varying experience reward
    - Varying lotteries: varying predicted reward

The axiomatic model

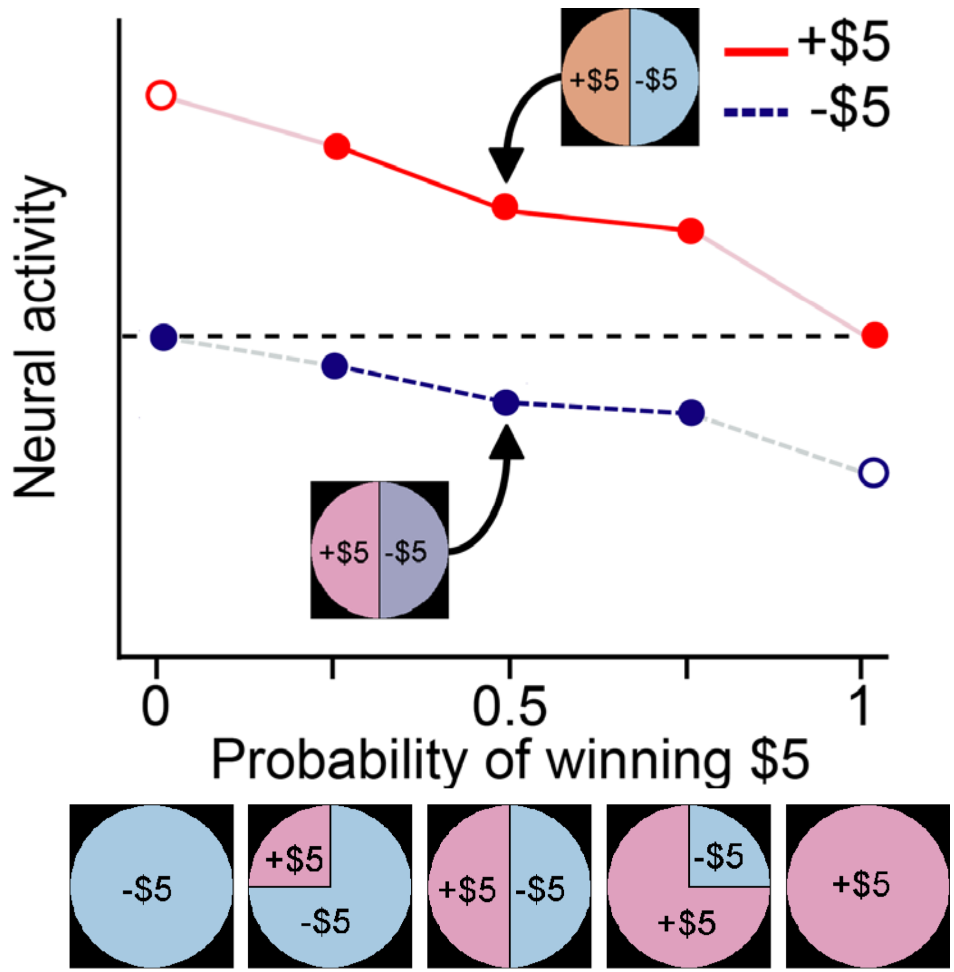

- Neural activity is plotted against probability of winning $5 from five different lotteries

Axiom 1: Consistent experienced reward

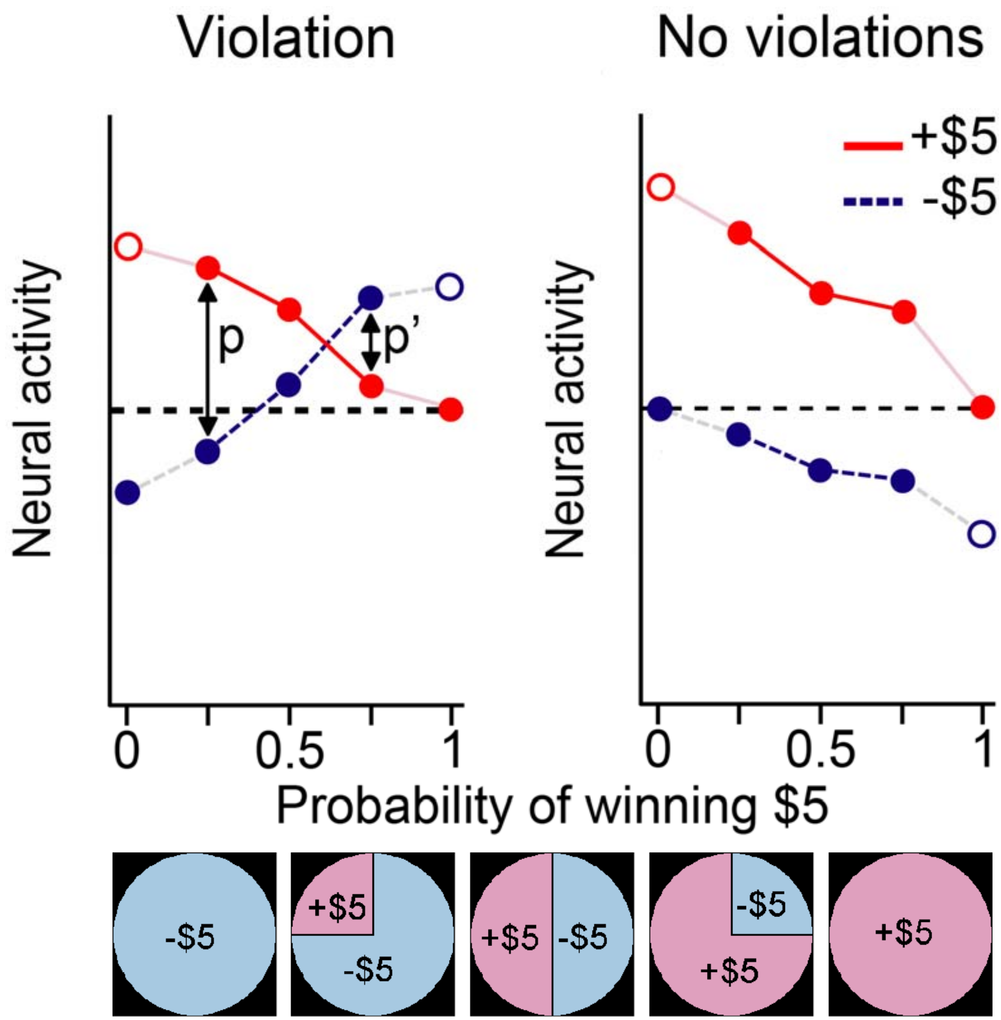

- $Given\ (z,p),(z',p),(z,p'),(z',p') \in A,$
  $\delta(z,p) > \delta(z',p) \Rightarrow \delta(z,p') > \delta(z',p')$
- For a fixed lottery, a better prize leads to more neural activity
- Lines do not cross

Axiom 2: Consistent predicted reward

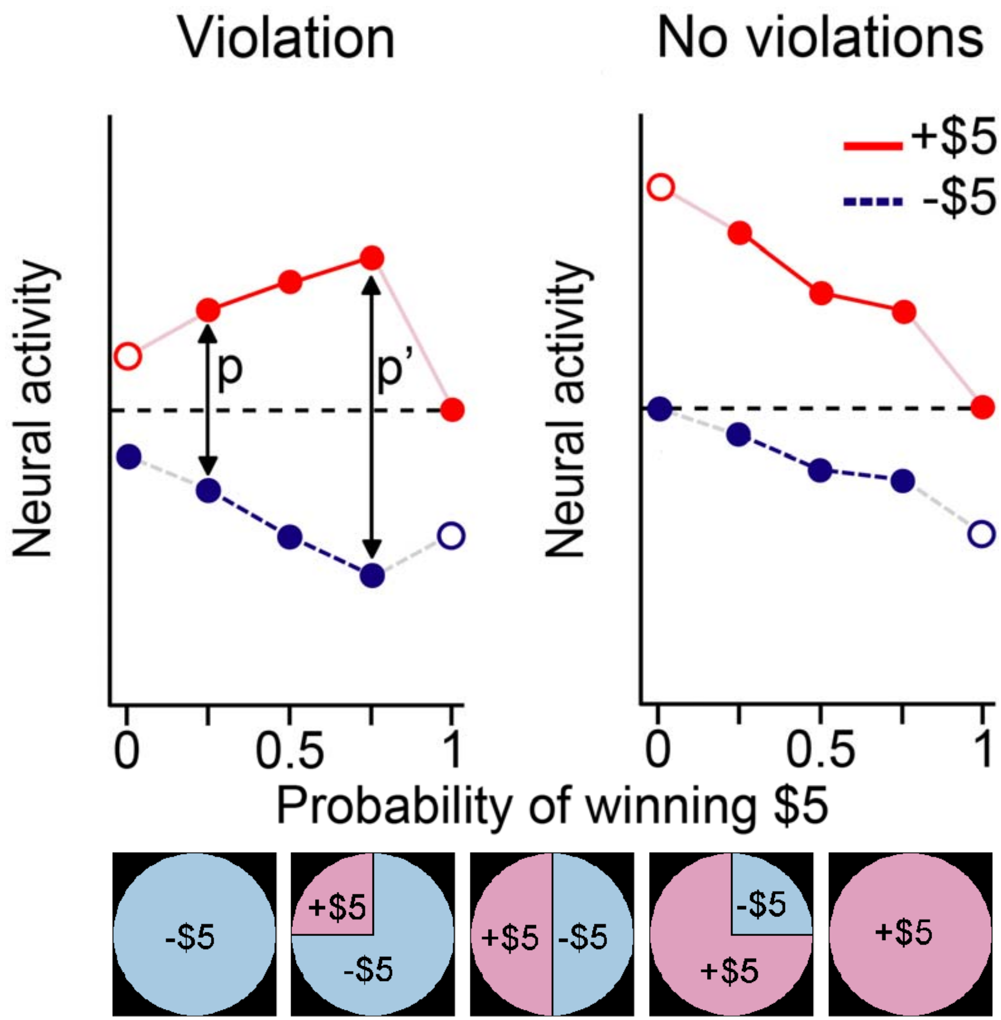

- $Given\ (z,p),(z',p),(z,p'),(z',p') \in A,$
  $\delta(z,p) > \delta(z,p') \Rightarrow \delta(z',p) > \delta(z',p')$
- For a fixed prize, a better lottery leads to less neural activity
- Lines are downward sloping (co-monotonic, more specifically)

Axiom 3: No surprise equivalence

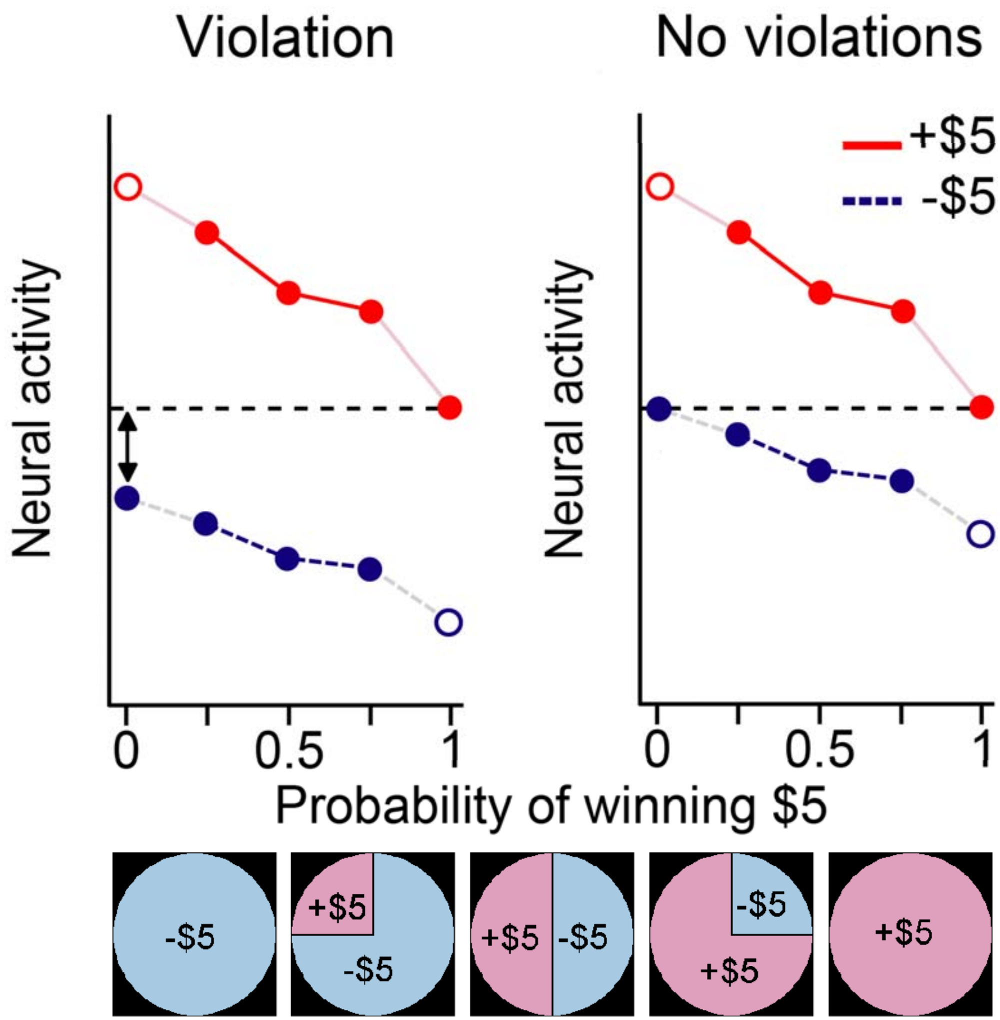

- *Given* $z, z' \in Z$, $\delta(z,z) = \delta(z',z')$
- Any fully anticipated prize should lead to the same neural activity
- Endpoints match

Summary

- If and only if
  - Lines do not cross
  - Lines have same direction of slope
  - Endpoints line up
- Then
  - Data is consistent with RPE model

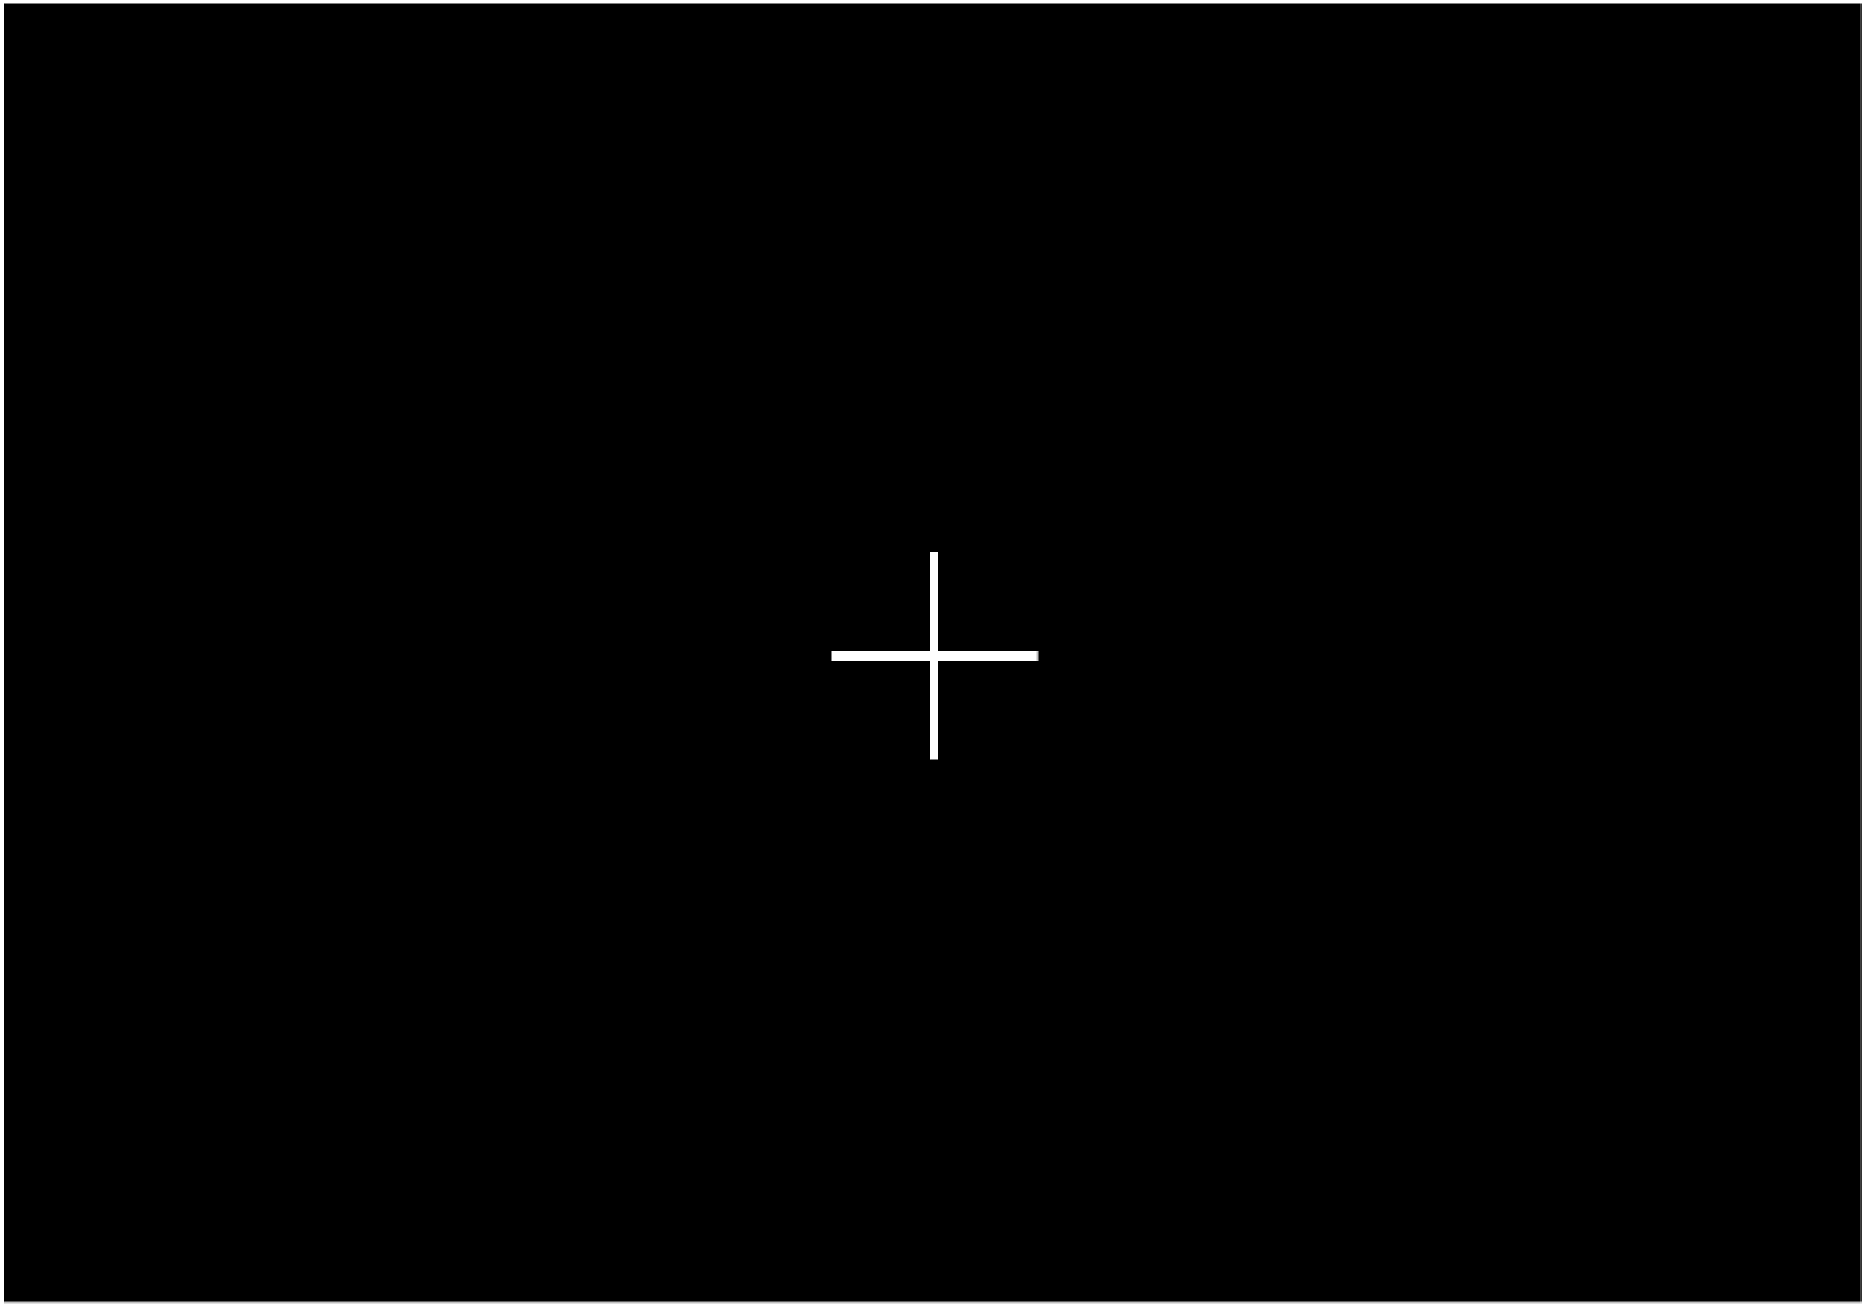

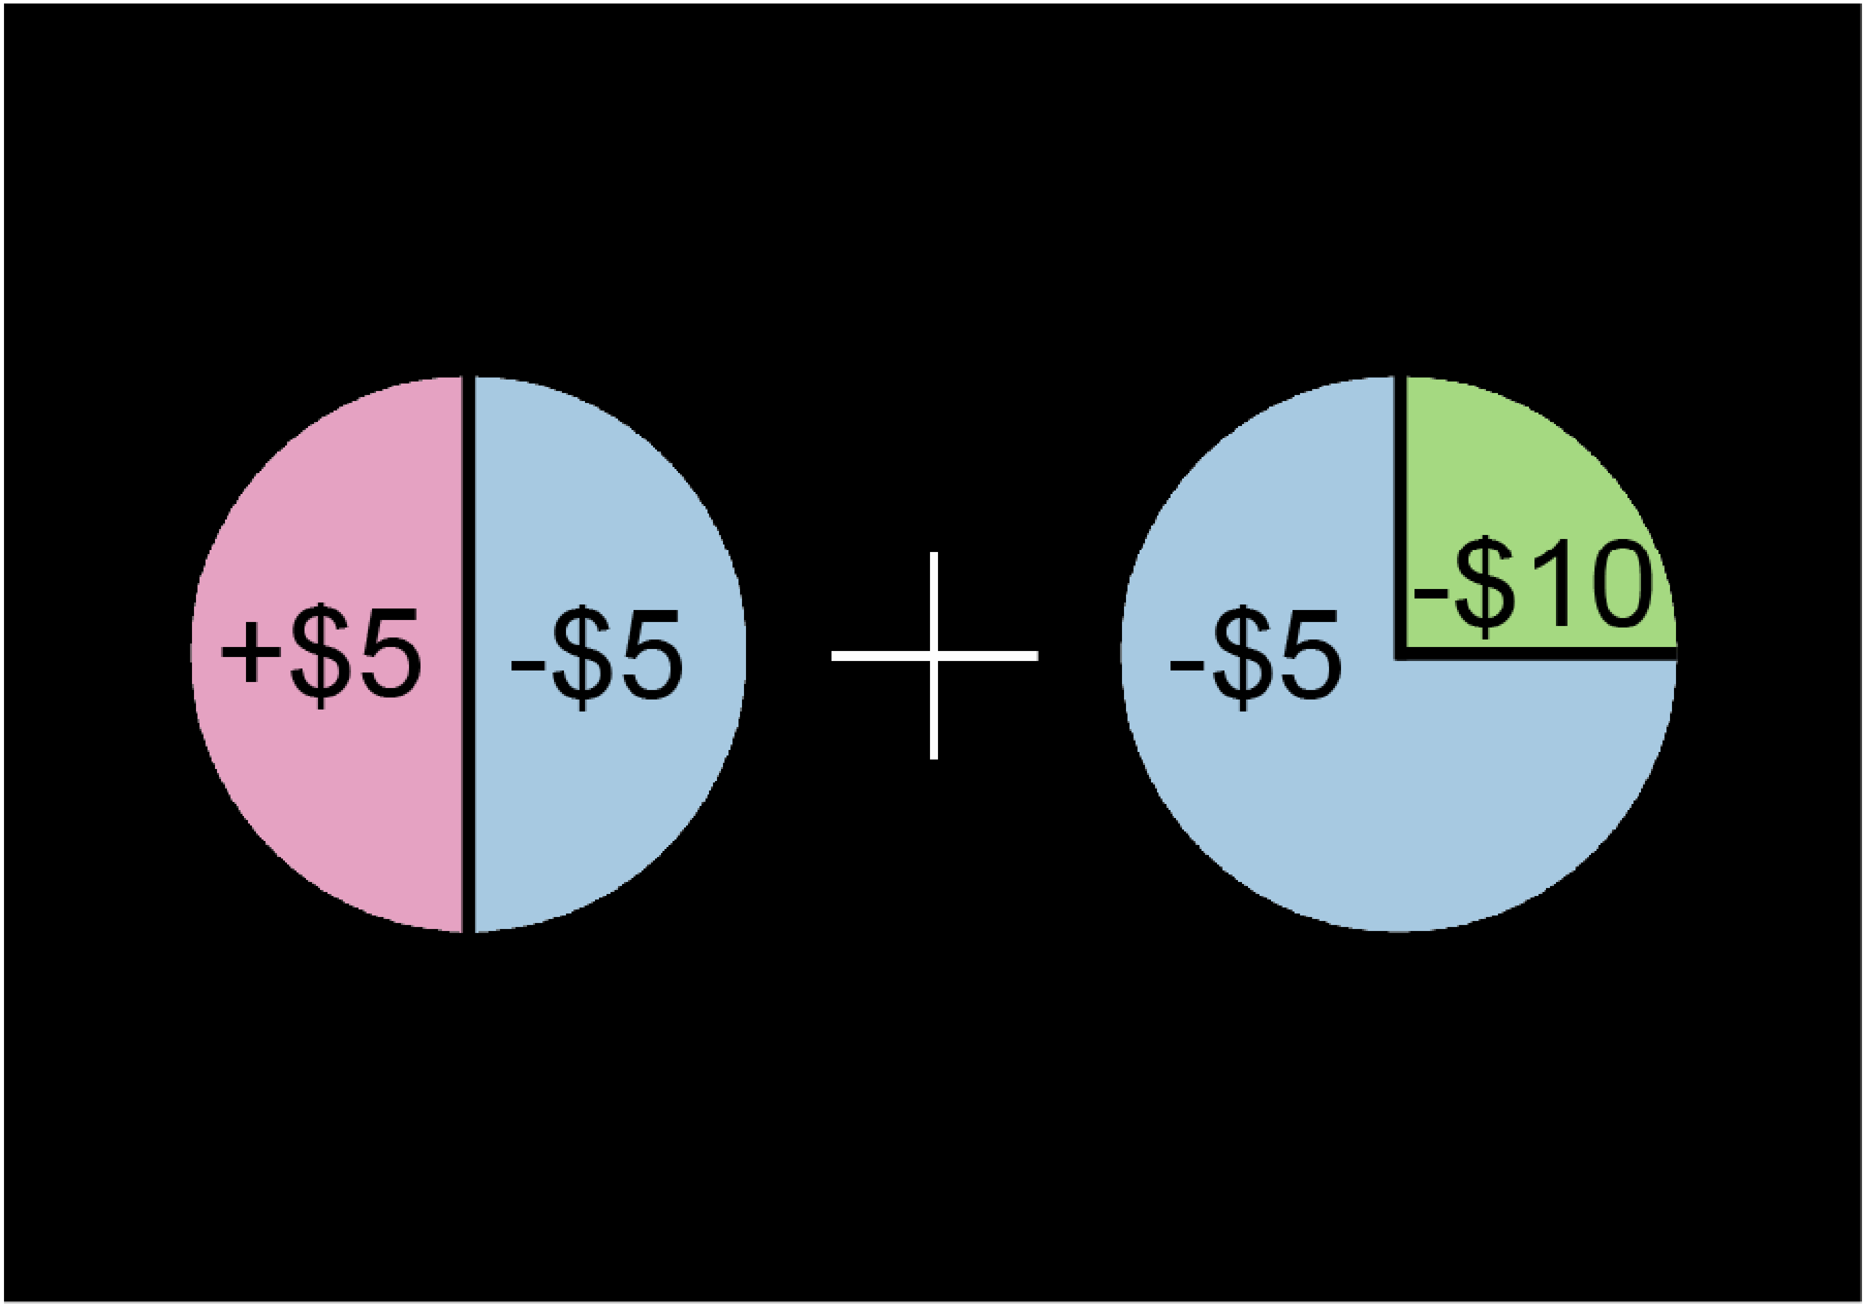
+$5
-$5
-$5
-$10

+$5
-$5
-$10
-$5

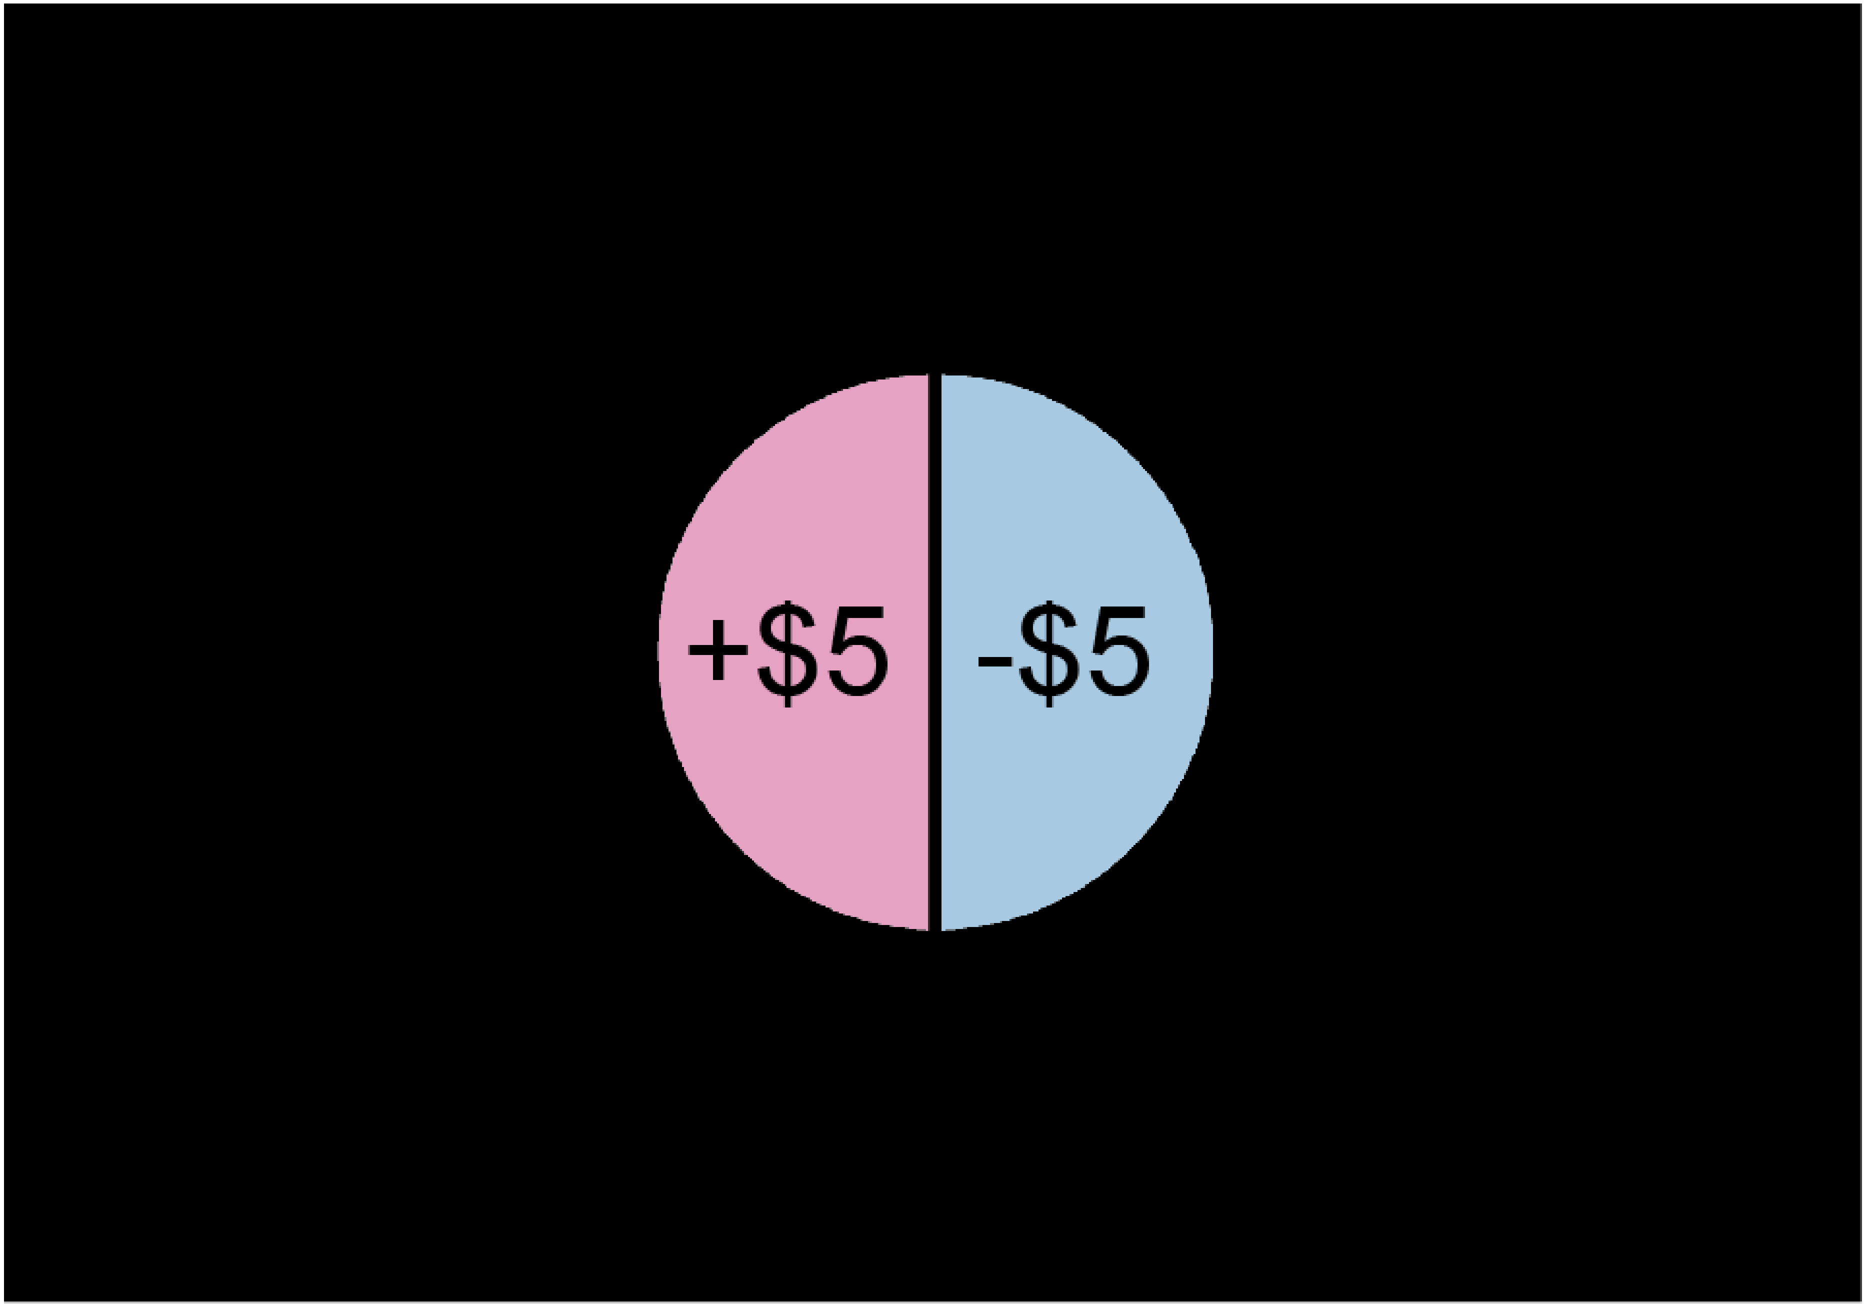
+$5
-$5

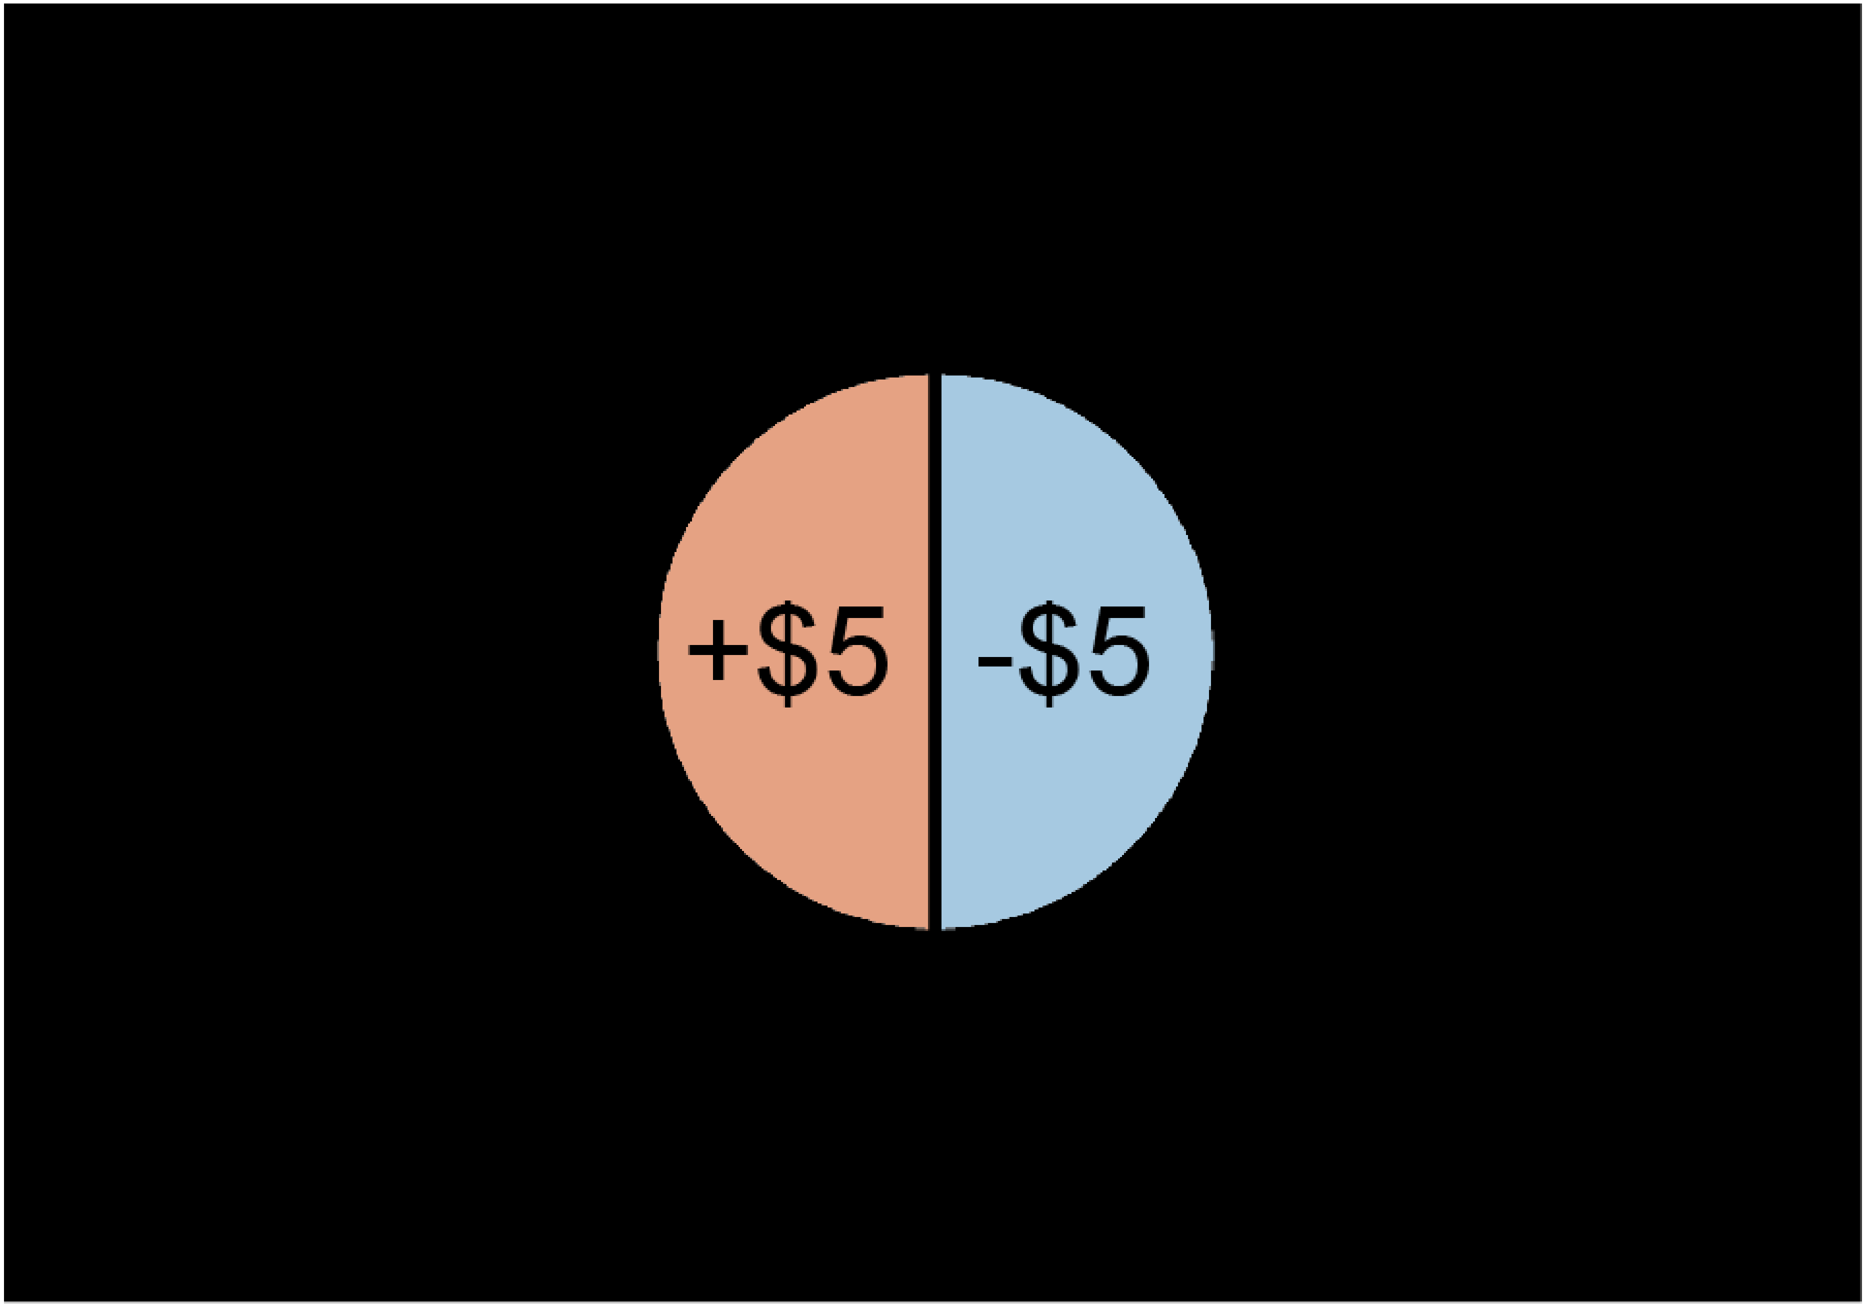
+$5
-$5

Gambling task

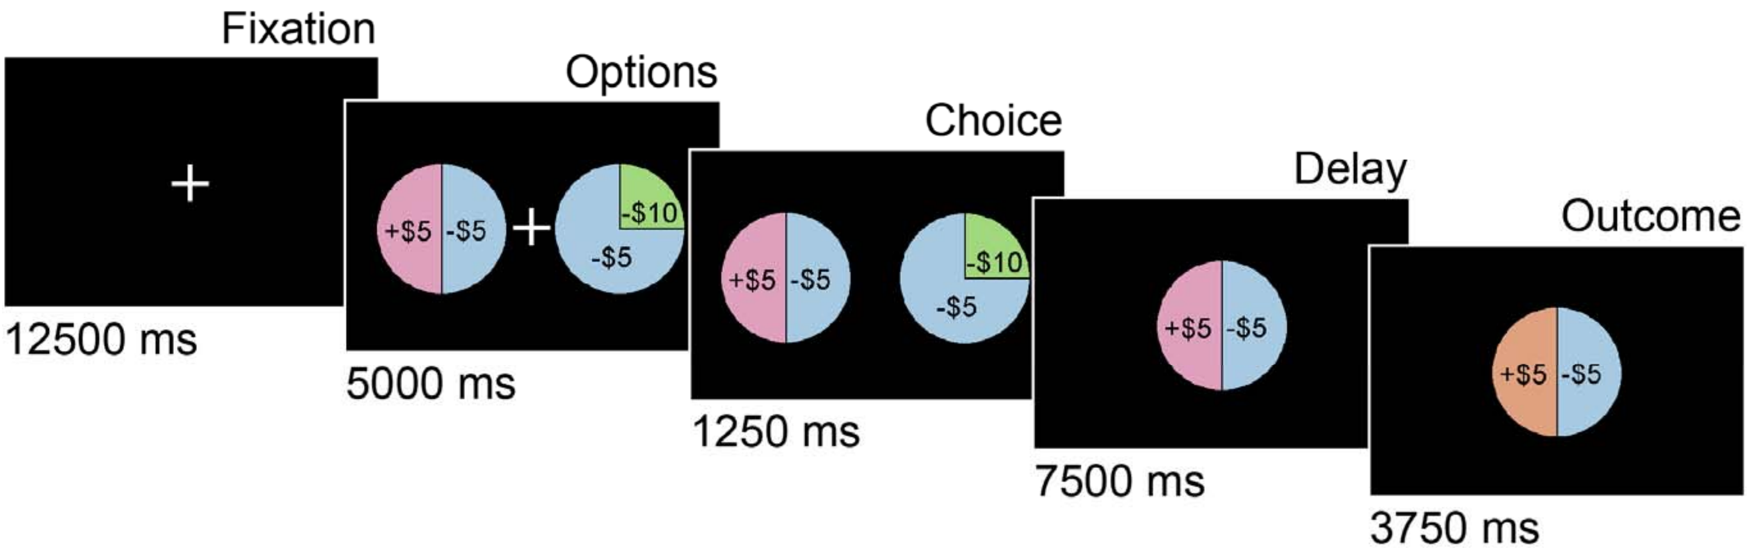

12 subjects, scanned 3 hours each, earned $100-$230.

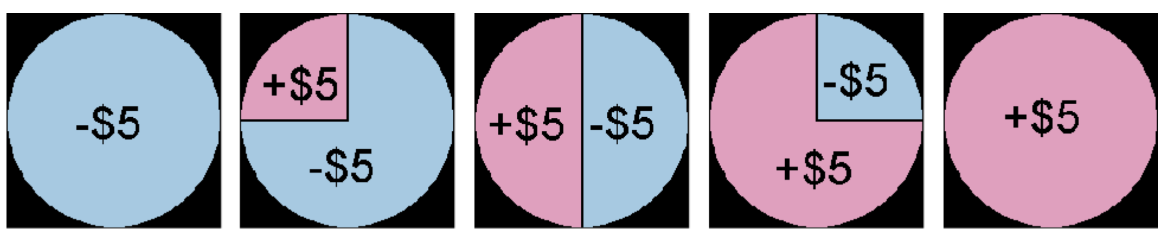

Anatomically defined region of interest

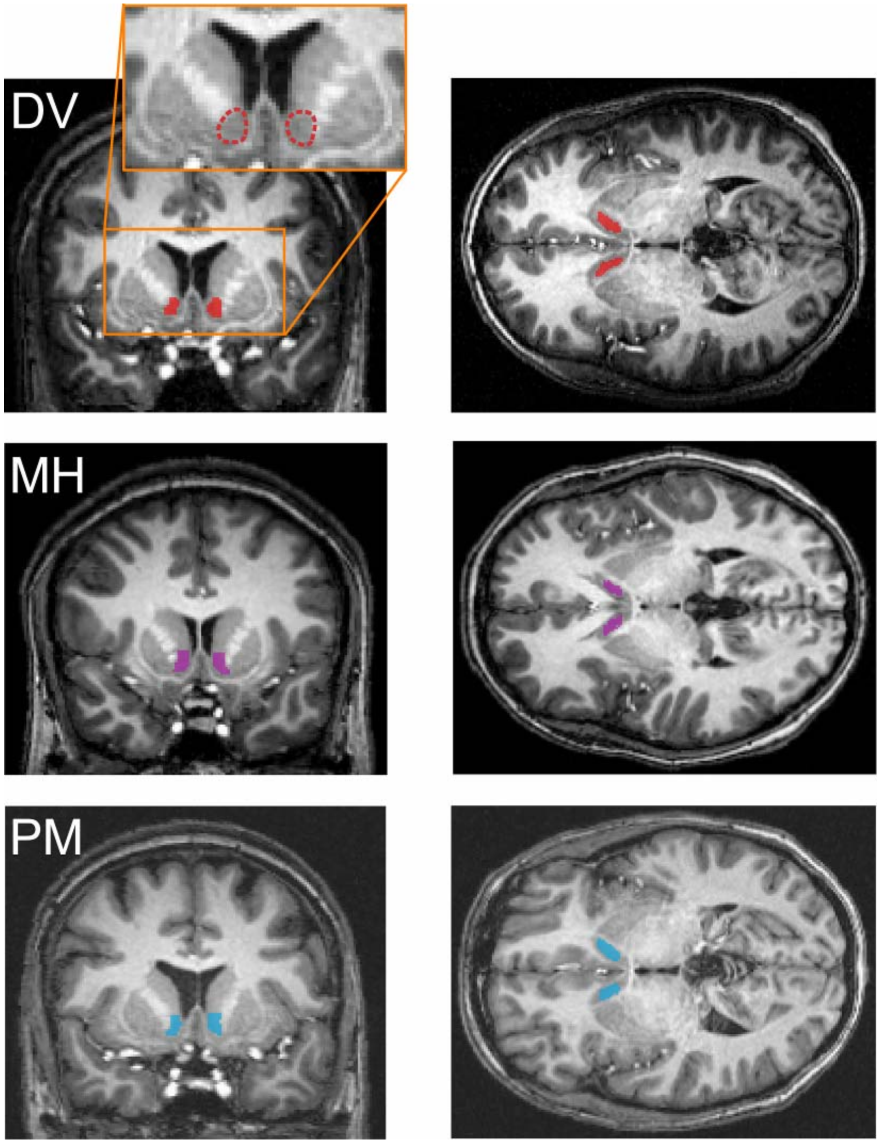

We anatomically defined the nucleus accumbens in subjects.

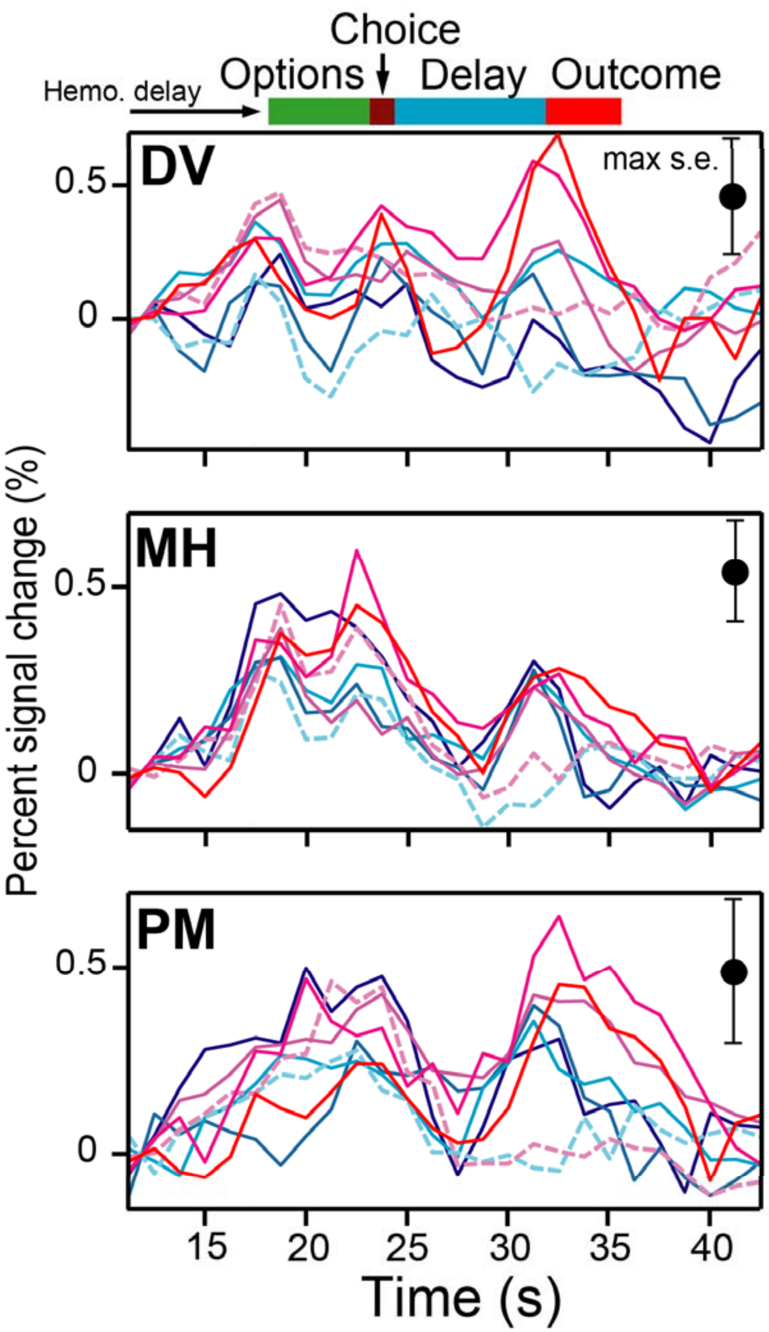

BOLD time series were extracted from the nucleus accumbens of each subject.

BOLD activity in nucleus accumbens

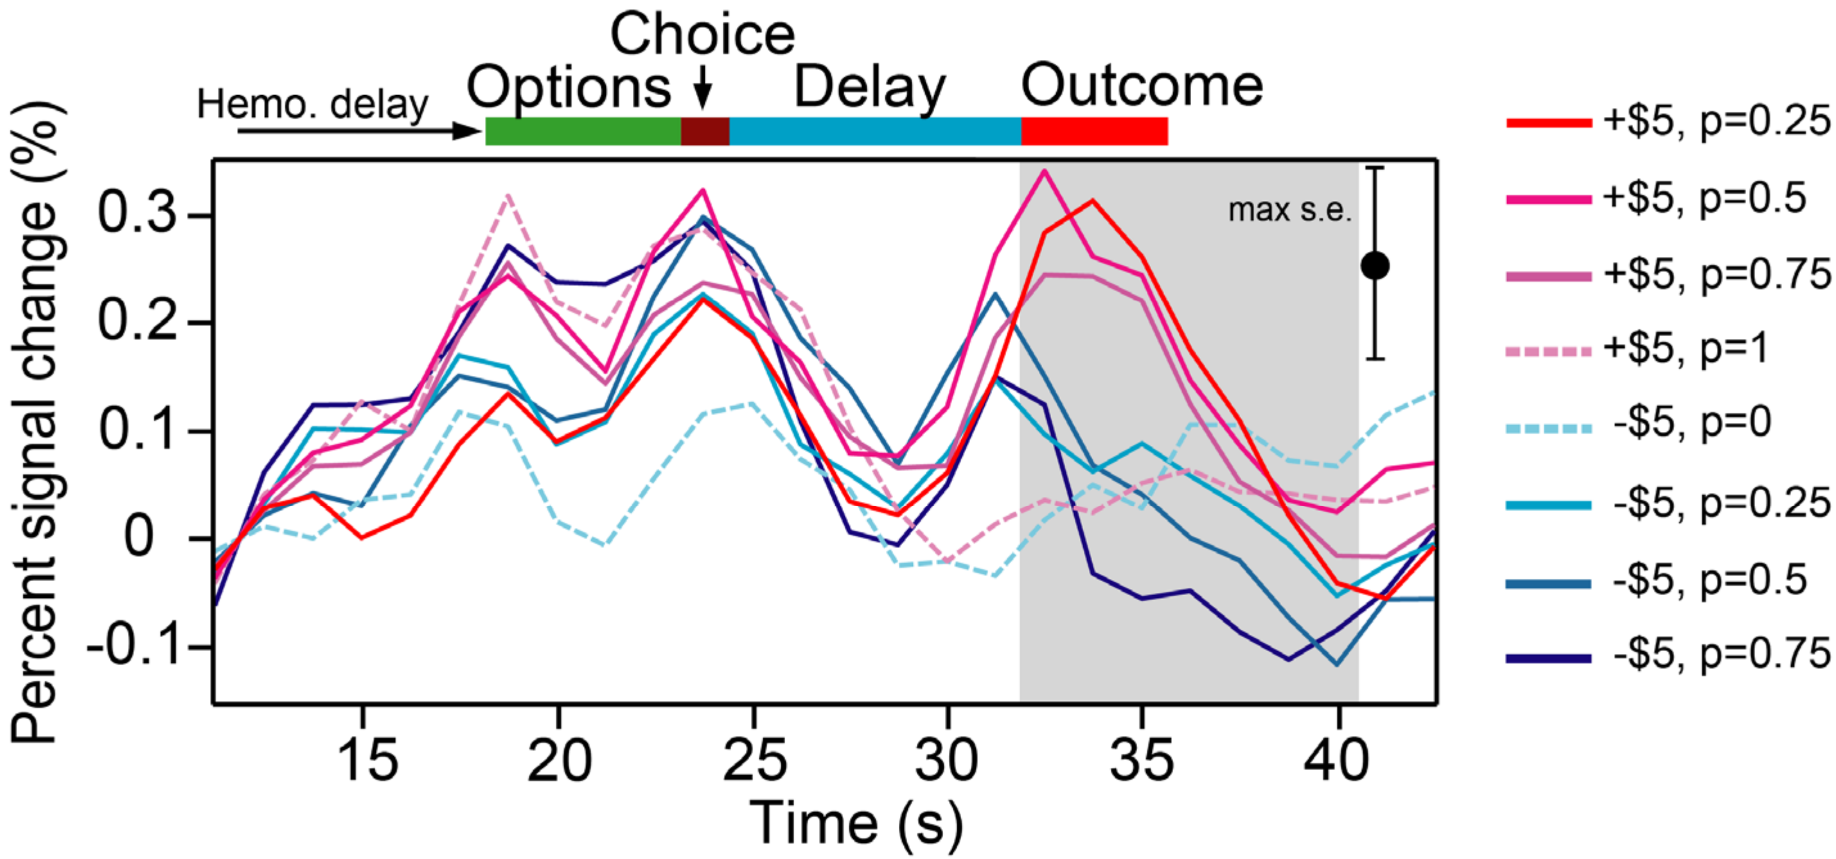

n = 12 subjects, 2975 trials

Testing the axioms

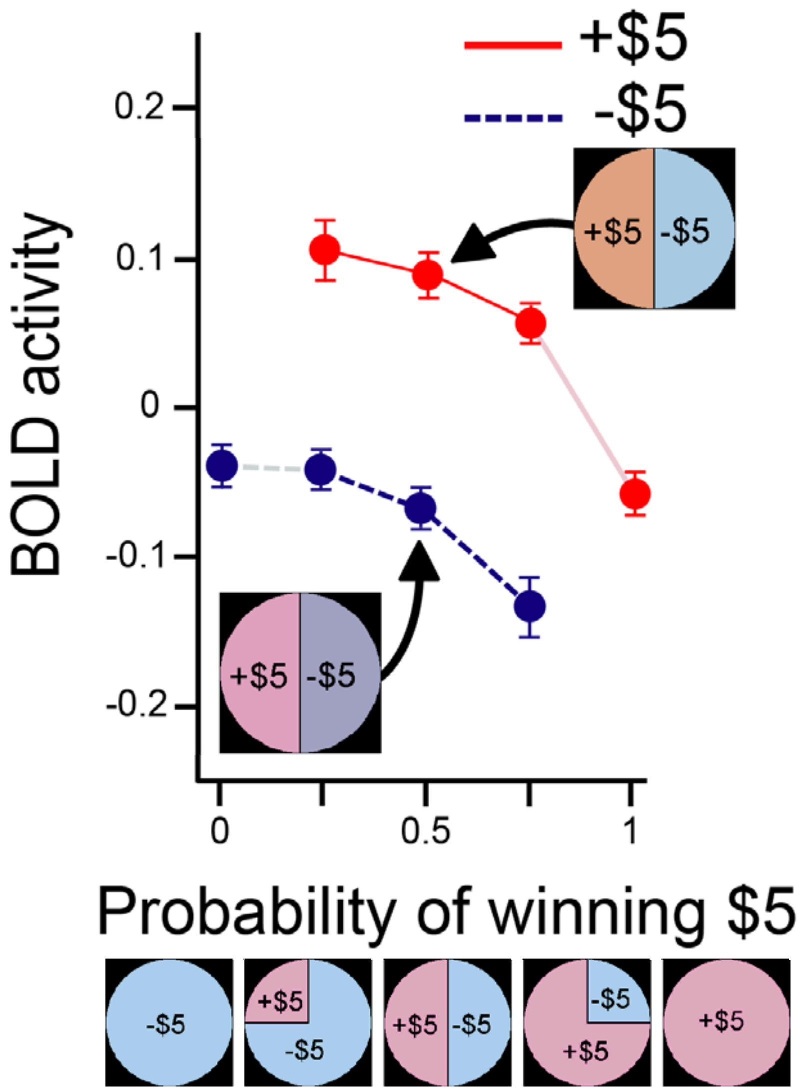

Testing the axioms

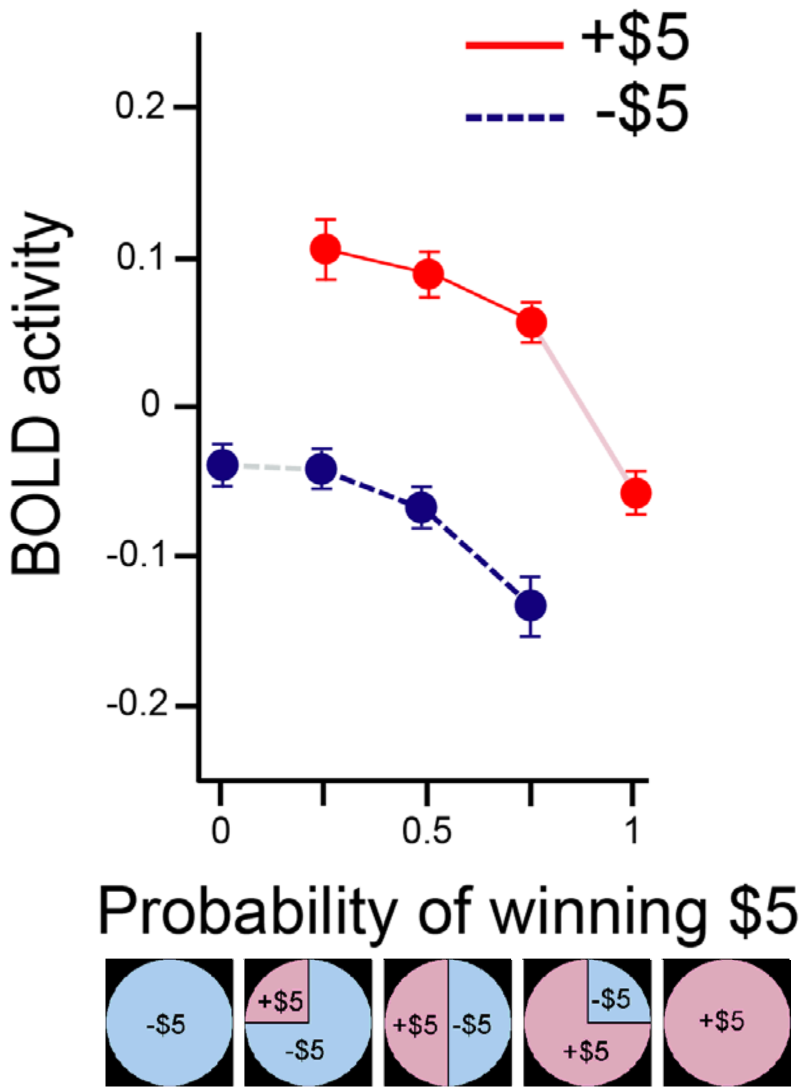

A1: Lines do not cross

A2: Lines have same direction of slope

A3: Endpoints match

Testing the axioms

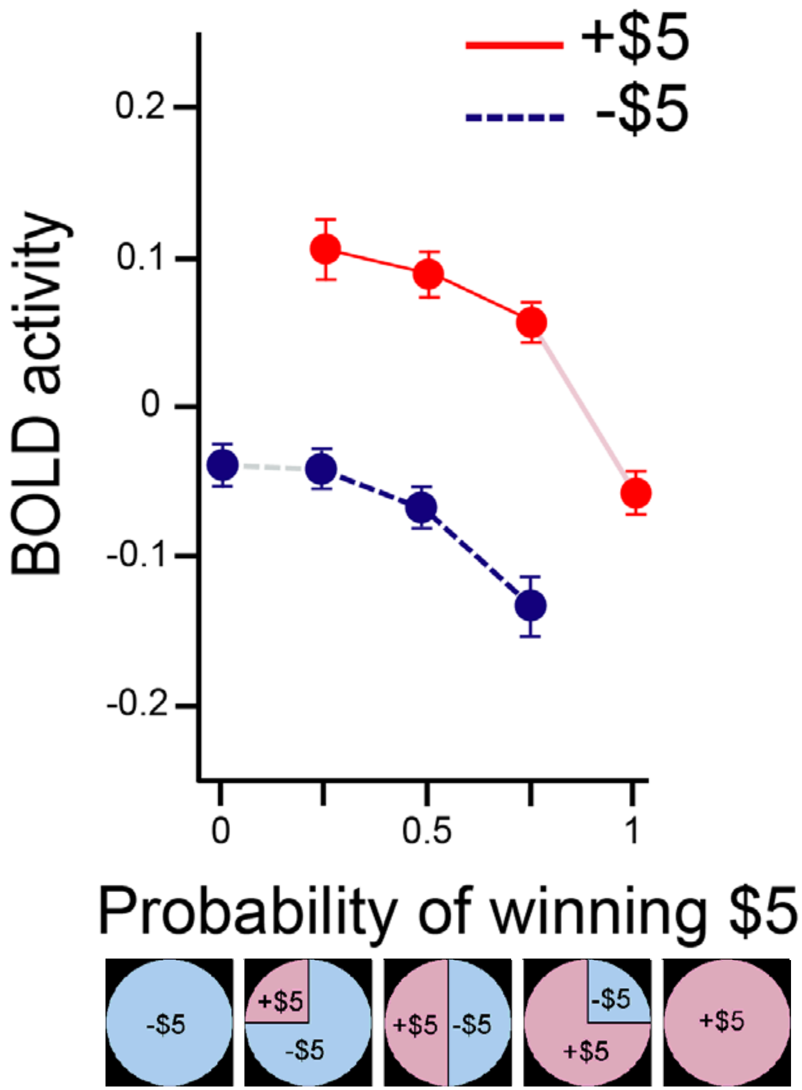

A1: Lines do not cross

A2: Lines have same direction of slope

A3: Endpoints match

Testing the axioms

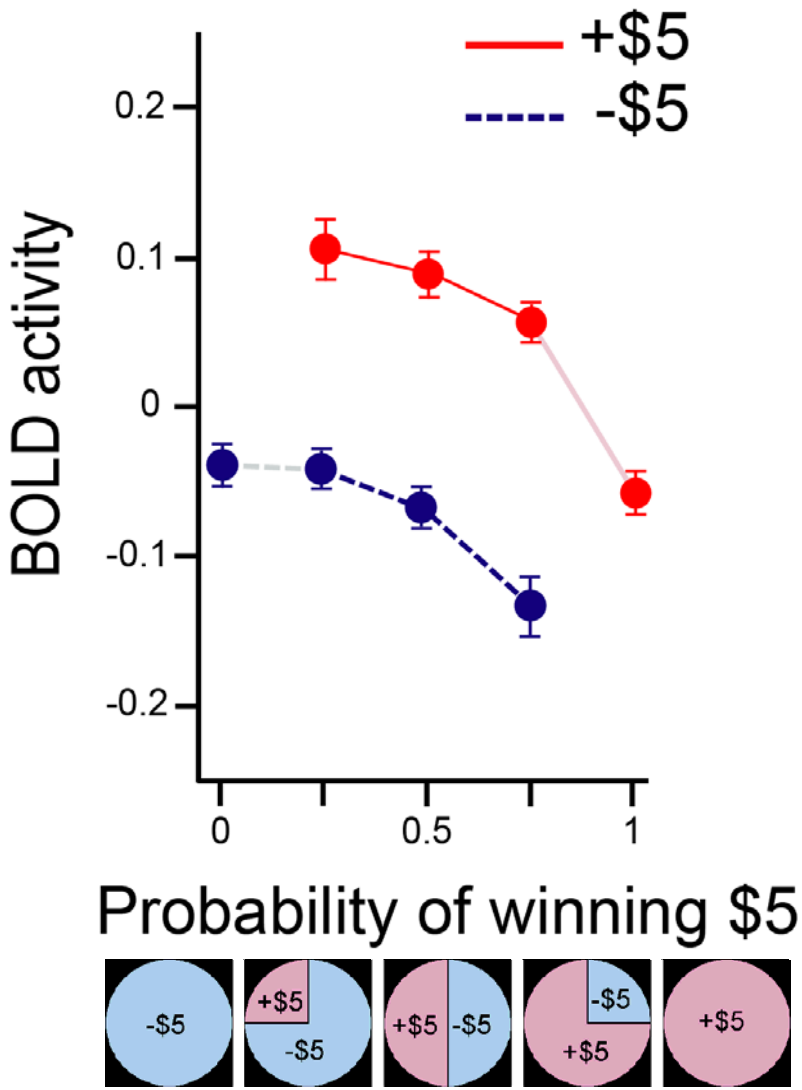

A1: Lines do not cross

A2: Lines have same direction of slope

A3: Endpoints match

Testing the axioms

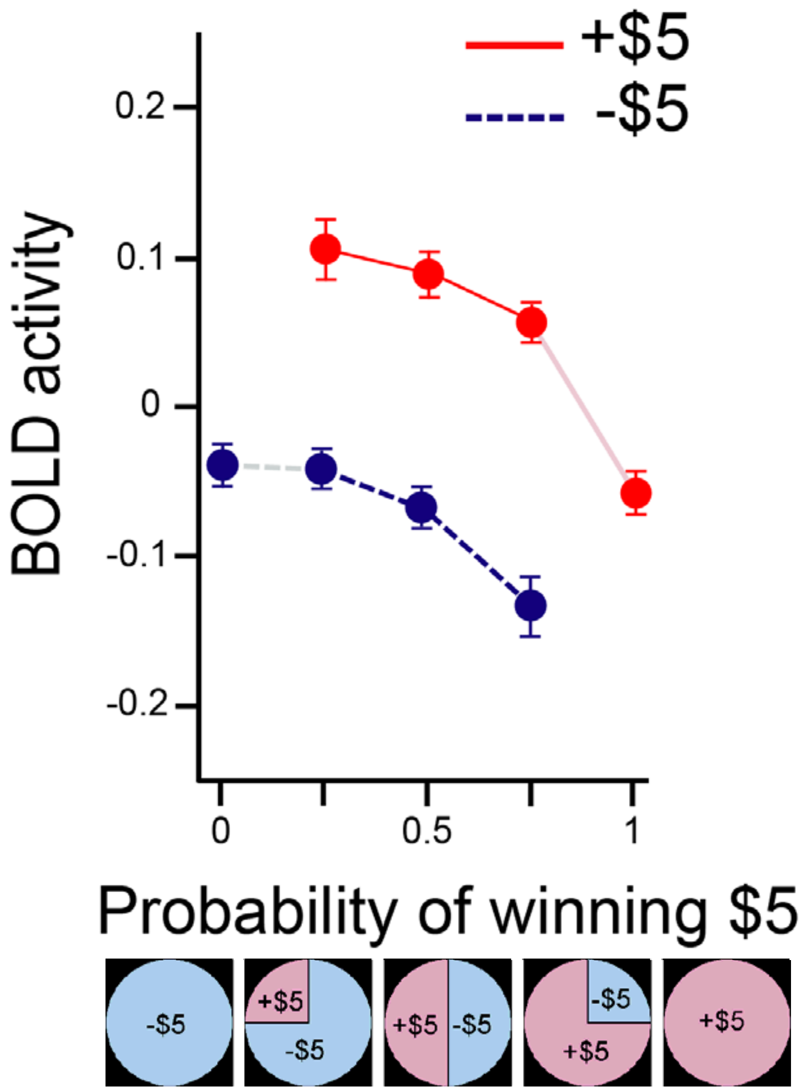

A1: Lines do not cross

A2: Lines have same direction of slope

A3: Endpoints match

Testing the axioms

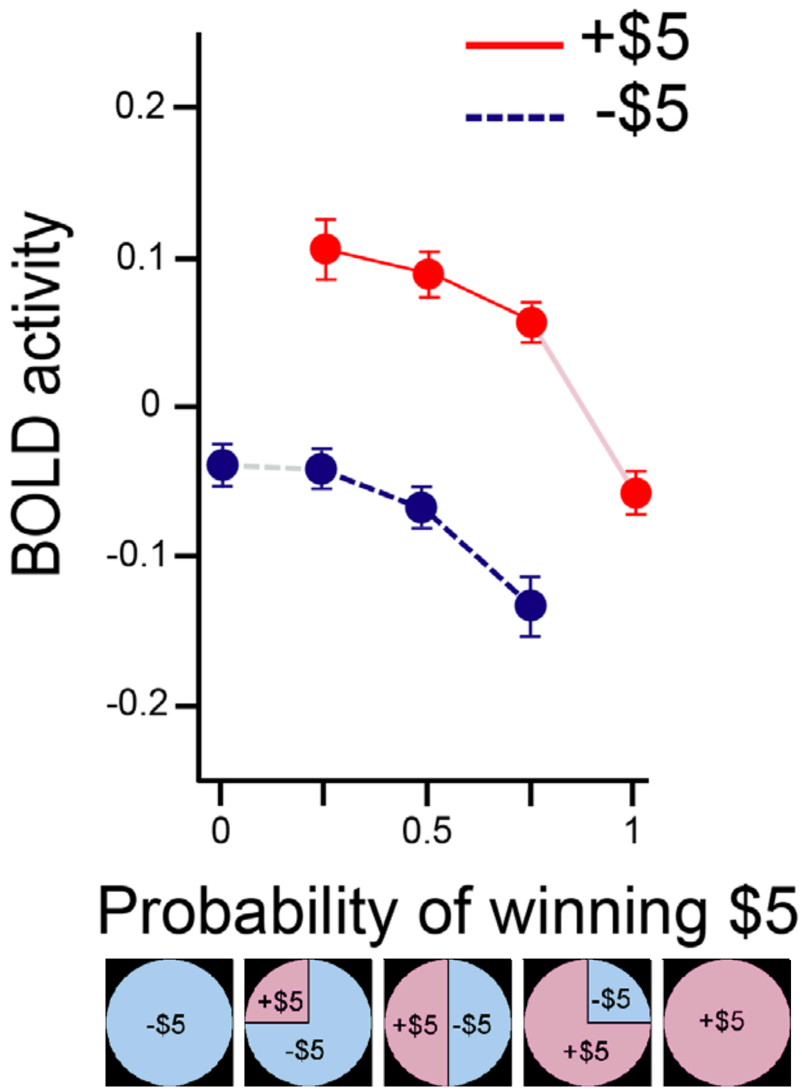

We cannot reject the class of **RPE** models.

BOLD activity in nucleus accumbens *can* serve as a **RPE** signal.

What About Other Brain Areas?

insula

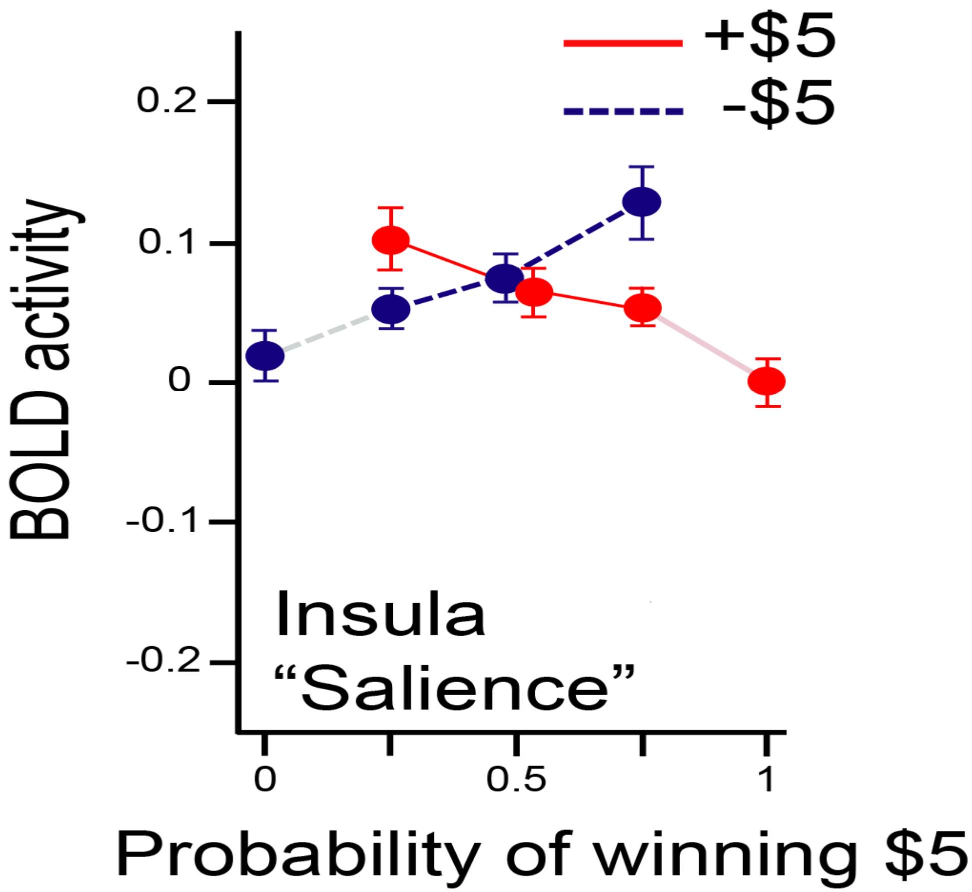

Summary

- Axioms allow us to identify necessary and sufficent conditions to test class of RPE model
- Test these axioms in nucleus accumbens using fMRI
- Axioms hold
- There is some definition of experienced and predicted reward such that measured activity encodes an RPE

Comparison to other models

- Redish [2004] uses a different approach to think abount dopamine activity (described in Sutton and Barton [1998])
- Adapted to this setting:

$$\partial(z, p) = R(z) - V(p)$$

- Where
  - States=lotteries
  - V(p) is the value of state p
  - R(z) is the reward of receiving prize z
- Clearly, up to this point, same as the class of models we have defined
  - Without making assumptions on R and V, need our approach to test

Redish [2004]

$$V_t(p) = \alpha\delta_t(z, p) + V_{t-1}(p)$$

- However, this model contains a rule for determining V(p) at time t
- What are the predictions of this model?
- Depends on the data set

Redish [2004]

$$V_t(p) = \alpha\delta_t(z, p) + V_{t-1}(p)$$

- In data set used so far, then this model can explain any observed dopamine activity
  - We assume data is static: do not observe history of rewards
  - Need to assume experienced and predicted rewards fixed
  - Redish model allows predicted reward over time
- In order to test model need to enrich the data set

Redish [2004]

- Need to assume we observe histories of reward/lottery pairs
- Data is now

$$\delta_t(z_t, p_t, h(p_t))$$

- Z is prize received
- P is lottery
- h(p) is history of previous rewards from lottery p

Redish [2004]

- For this data, model predicts

$$\delta_t(z_t, p_t, h(p_t)) = R(z_t) - \sum_{j \in h(p_t)} \alpha^j R(z_j)$$
$$= R(z_t) - V(h(p_t))$$

- Is this model 'axiom free'?
- By making specific assumption, has removed one unobservable
  - 'Beliefs' now observable
  - Moved towards testing specific model in RPE class
  - 'Assumptions substitute for Axioms'

Redish [2004]

- For this data, model predicts

$$\delta_t(z_t, p_t, h(p_t)) = R(z_t) - \sum_{j \in h(p_t)} \alpha^j R(z_j)$$
$$= R(z_t) - V(h(p_t))$$

- Model still contains unobservables
- Rewards are unobservable
- Can either become completely specific (and make assumptions about R(.))
- Or develop new axioms

Pseudo Axioms

- **Consistent Prize Ordering:** If, after some history, prize z leads to higher dopamine response than x, this must be true for all histories
- **Sequence Consistency:** If z is a better prize that x (in the sense above), then substituting z for x in a history should lower dopamine activity
- **Forgetting:** For any prizes z and x, switching z for x in a history should have a smaller effect the further back in time the switch occurs

Summary

Summary

- Economists often have models which contain unobservable features
- Axioms provide a way of testing these models, without having to make additional assumption
- Neuroscience is now using models with unobservables
- Axiomatic tools may be helpful in testing these models

Thank You

Thanks to Andrew Caplin, Robb Rutledge, Paul Glimcher and the Glimcher Lab

Problem #1: Model predictions were correlated