Delay discounting as a tool for computational psychiatry

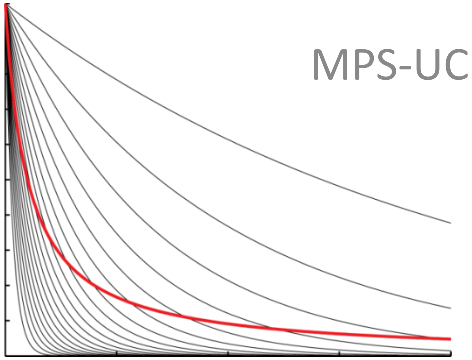

MPS-UCL Symposium and Advanced Course on Computational Psychiatry and Ageing Research

2012-Sep-19

Zeb Kurth-Nelson
Wellcome Trust Centre for Neuroimaging


BOOK SHOP
I'VE JUST GOT TO BUY THIS!
SELF HELP
HOW TO STOP IMPULSE BUYING
CHRIS MADDEN.
COPYRIGHT ©
IF THIS MESSAGE IS PRESENT THIS IMAGE IS BEING USED WITHOUT PERMISSION

Outline

1) Delay discounting is important

2) Designing a task & analyzing the data

3) What delay discounting measures

4) Modelling discounting

Outline

1) Delay discounting is important

2) Designing a task & analyzing the data

3) What delay discounting measures

4) Modelling discounting

Computational psychiatry: A basis for psychiatric disorders that reflects the underlying structure of the problems

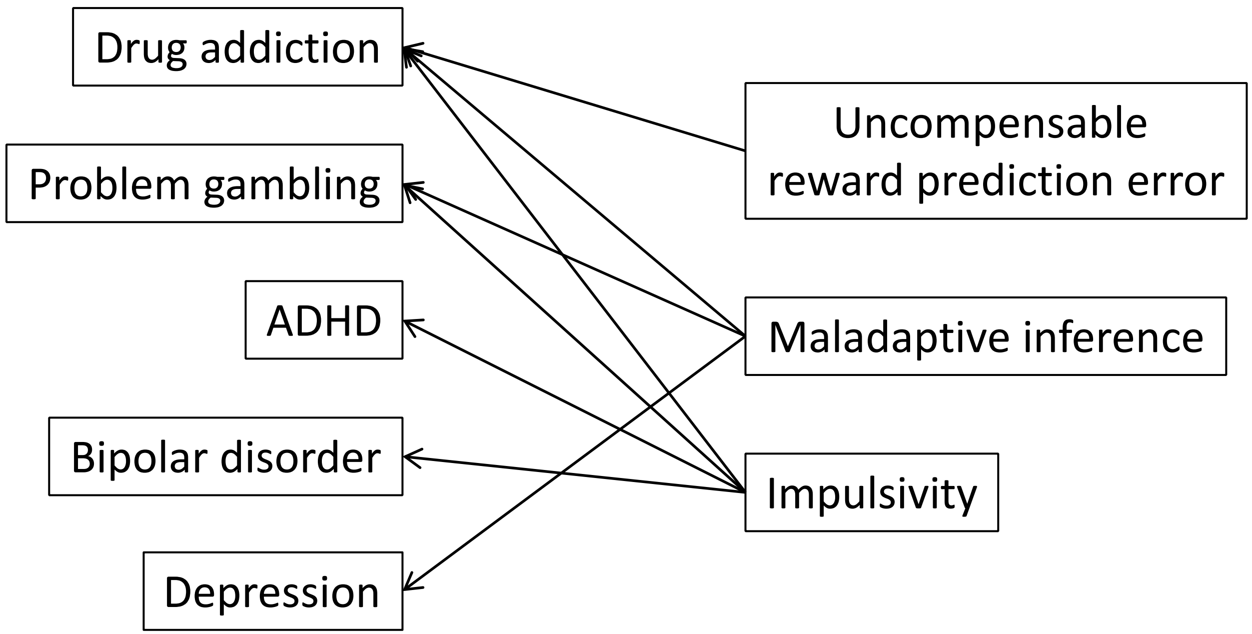

How much is $1000 worth if you have to wait for it?

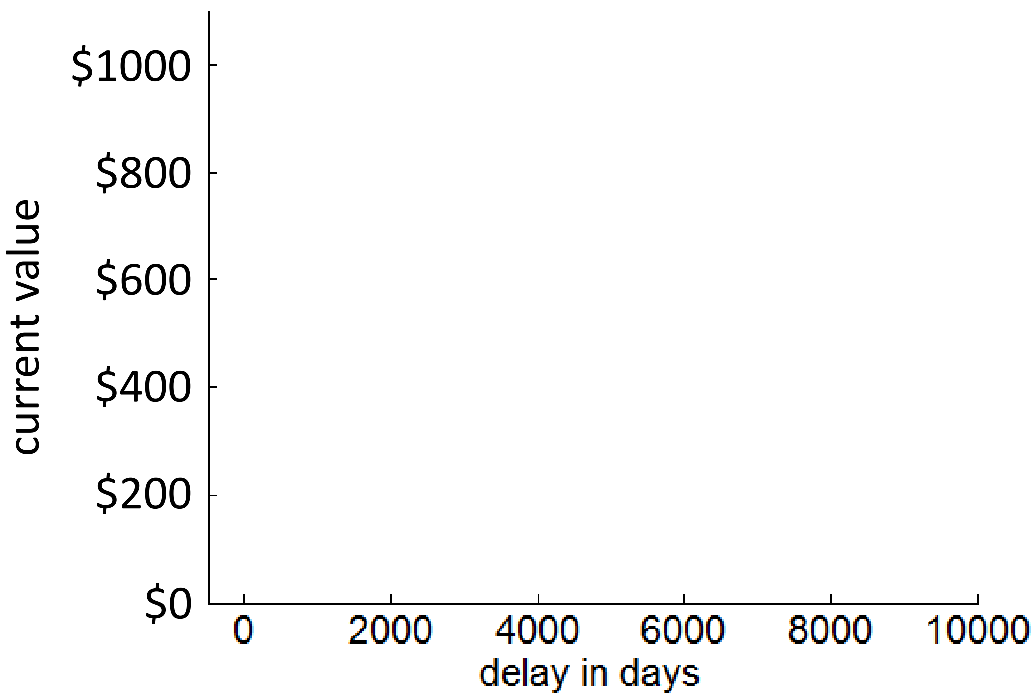

Measure[a]	Edu	Income	BIS NON	BIS MTR	BIS COG	IQ	DD
Age	−.02	.18□	−.15□	− .12	−.09	.02	−.04
Education		.24□	−.25□	−.10	−.26□	.51□	−.27□
Income			−.38□	−.02	−.12	.25□	−.27□
BIS NON				.32□	.44□	−.15□	.26□
BIS MTR					.56□	−.06	.05
BIS COG						.26□	−.16□
IQ							−.37□

□ $p < .001$.

de Wit et al (2007) Person Indiv Diff 42:111

Drug addicts discount more steeply than healthy controls

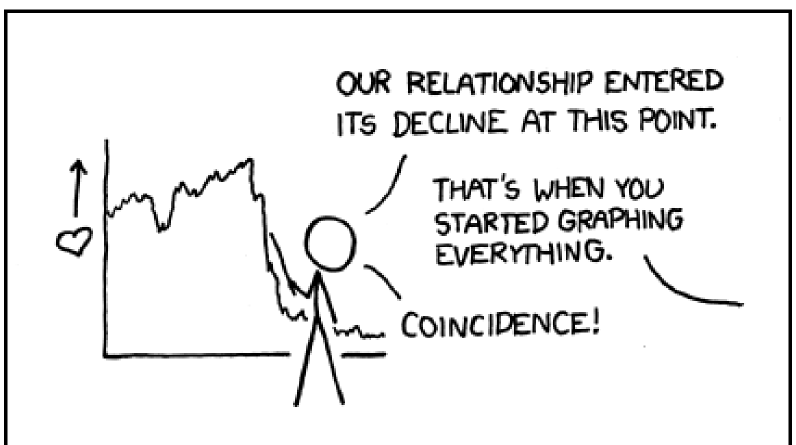

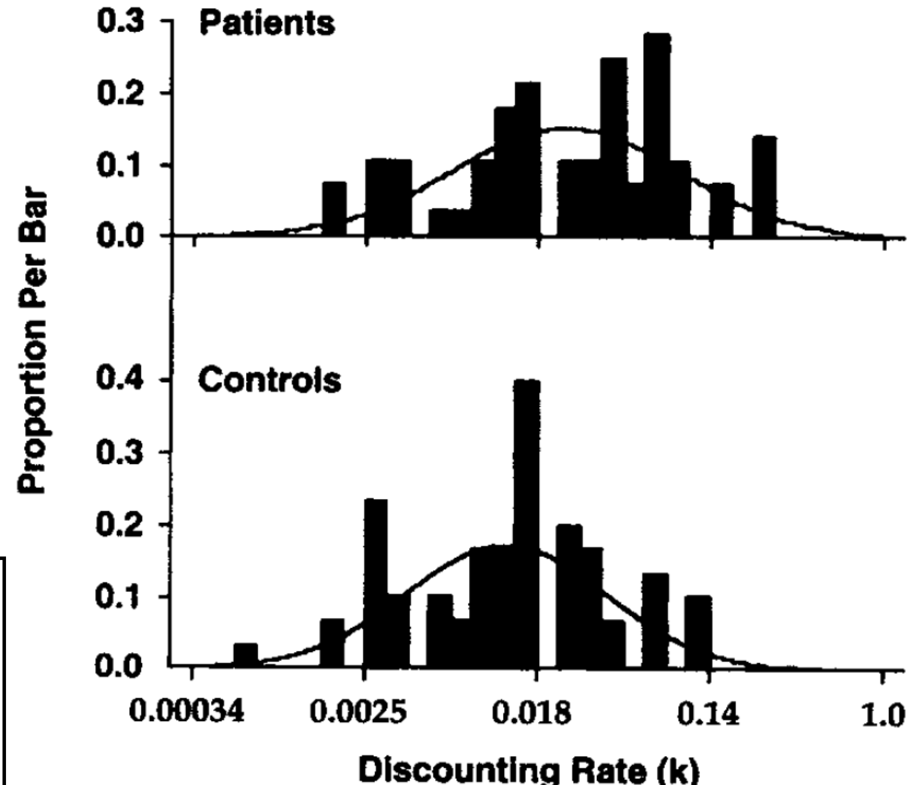

Kirby KN, Petry NM, Bickel WK (1999) JEP:G 128:78

Steeper delay discounting in…

- Opiate addicts Madden et al (1997) Exp Clin Psychopharm 5:256
- Cocaine addicts Coffey et al (2003) Exp Clin Psychopharm 11:18
- Methamphetamine addicts Hoffman et al (2006) Psychopharm 188:162
- Alcoholics Dom et al (2006) Addiction 101:50–59
- Smokers Bickel et al (1999) Psychopharm 146:447
- Obese Weller et al (2008) Appetite 51:563–569
- Gamblers Petry (2001) Abnorm Psych 110:482
- ADHD Wilson et al (2011) J Child Psych&Psych 52:256
- Boderline personality disorder Coffey et al (2011) Person Disord 2:128
- People with low credit scores Meier and Sprenger (2012) Psych Sci 23:56

People who discount steeply at the beginning of treatment are less likely to see a benefit of treatment

	Number of negative urine drug screens			Continuous abstinence			4 Weeks abstinence		8 Weeks abstinence	
Predictor	*B*	*SE*	β	*B*	*SE*	β	OR	95% CI	OR	95% CI
Model 1										
$100 money	−0.69	0.35	−0.15*	−0.26	0.16	−0.12	0.90	[0.79, 1.02]	0.88	[0.77, 1.01]
$1,000 money	−0.95	0.36	−0.20*	−0.43	0.17	−0.20*	0.87	[0.75, 0.99]*	0.82	[0.71, 0.95]*
$100 marijuana	−0.08	0.25	−0.03	−0.00	0.12	−0.00	0.97	[0.89, 1.06]	0.96	[0.87, 1.06]
$1,000 marijuana	−0.35	0.23	−0.12	−0.13	0.11	−0.09	0.91	[0.84, 0.99]*	0.93	[0.85, 1.02]
Model 2										
$100 money	−0.39	0.31	−0.09							
$1,000 money	−0.24	0.34	−0.05	−0.13	0.16	−0.06	0.97	[0.83, 1.14]	0.88	[0.74, 1.04]
$1,000 marijuana							0.93	[0.85, 1.03]		

Stanger C et al (2012) Exp Clin Psychopharm 20:205

Adolescents who discount steeply are more likely to take up smoking

	Level				Trend			
	β	SE	z	p-value	β	SE	z	p-value
Regular smoking								
Delay discounting level	–	–	–	–	.08	.04	2.16	.03
Delay discounting trend	–	–	–	–	−.24	.38	−.64	.53

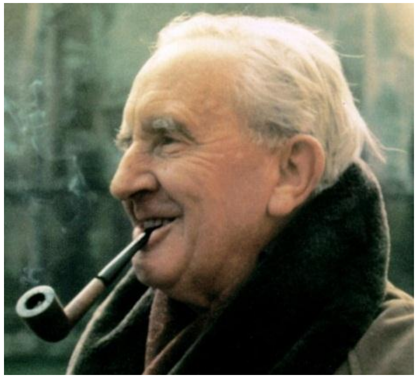

Audrain-McGovern et al (2009) Drug Alc Depend 103:99

Outline

1) Delay discounting is important

2) Designing a task & analyzing the data

3) What delay discounting measures

4) Modelling discounting

exponential

k^d

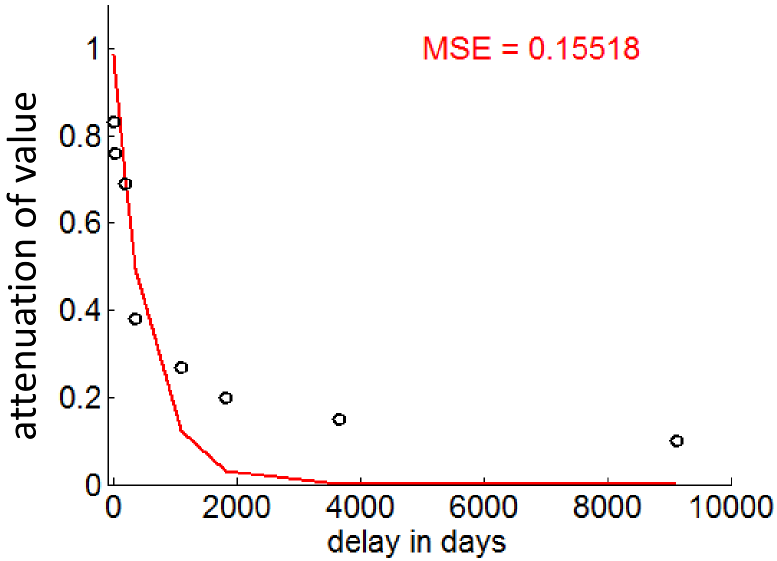

Data from Vuchinich and Simpson (1998) Exp Clin Psychopharm 6:292

exponential

$$k^d$$

hyperbolic

$$\frac{1}{1+kd}$$

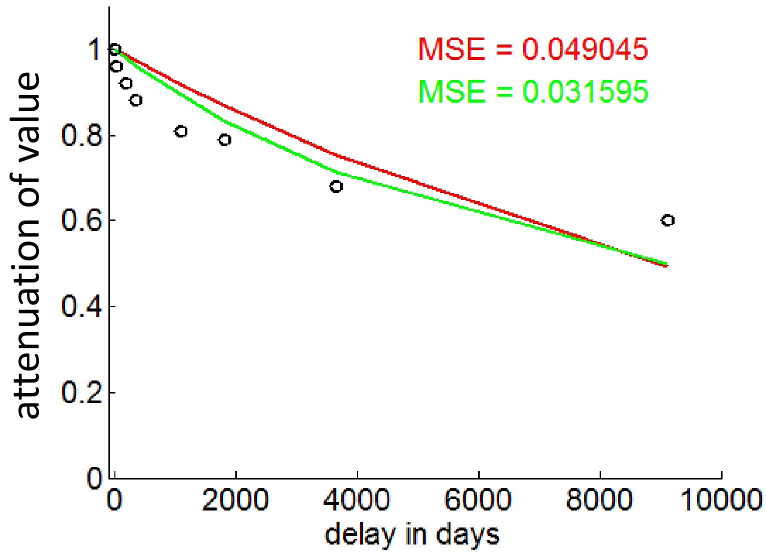

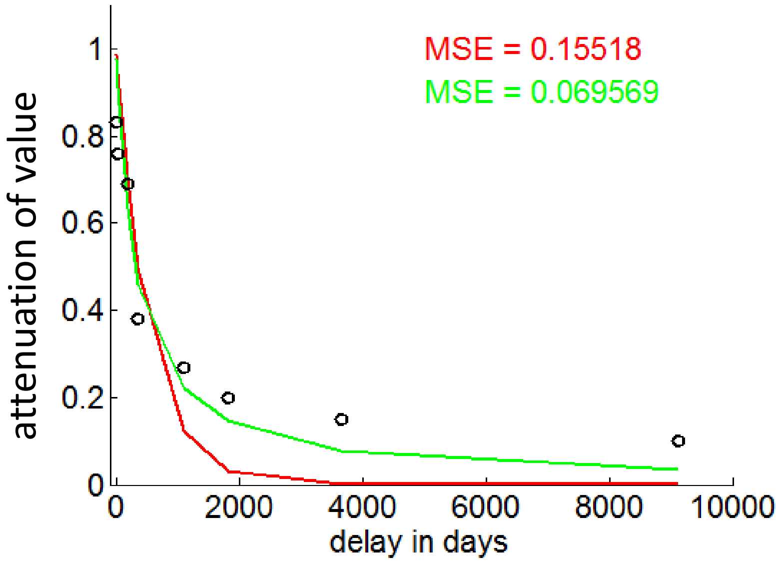

Data from Vuchinich and Simpson (1998) Exp Clin Psychopharm 6:292

exponential

$$k^d$$

hyperbolic

$$\frac{1}{1+kd}$$

power law

$$d^{-k}$$

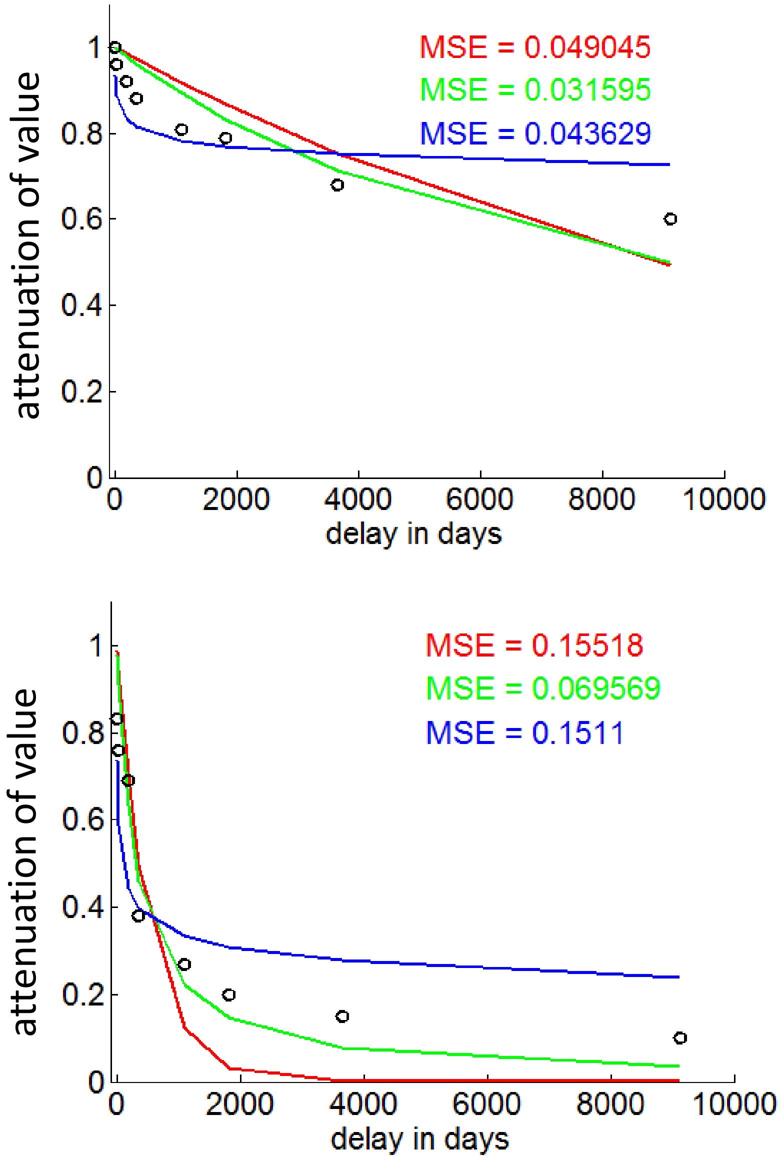

Many studies now show the superiority of hyperbolic fits for human and animal discounting data

Is hyperbolic significantly better?

-Bayesian model comparison

- Rank-sum test on MSEs across subjects

Data from Vuchinich and Simpson (1998) Exp Clin Psychopharm 6:292

exponential

$$k^d$$

hyperbolic

$$\frac{1}{1+kd}$$

power law

$$d^{-k}$$

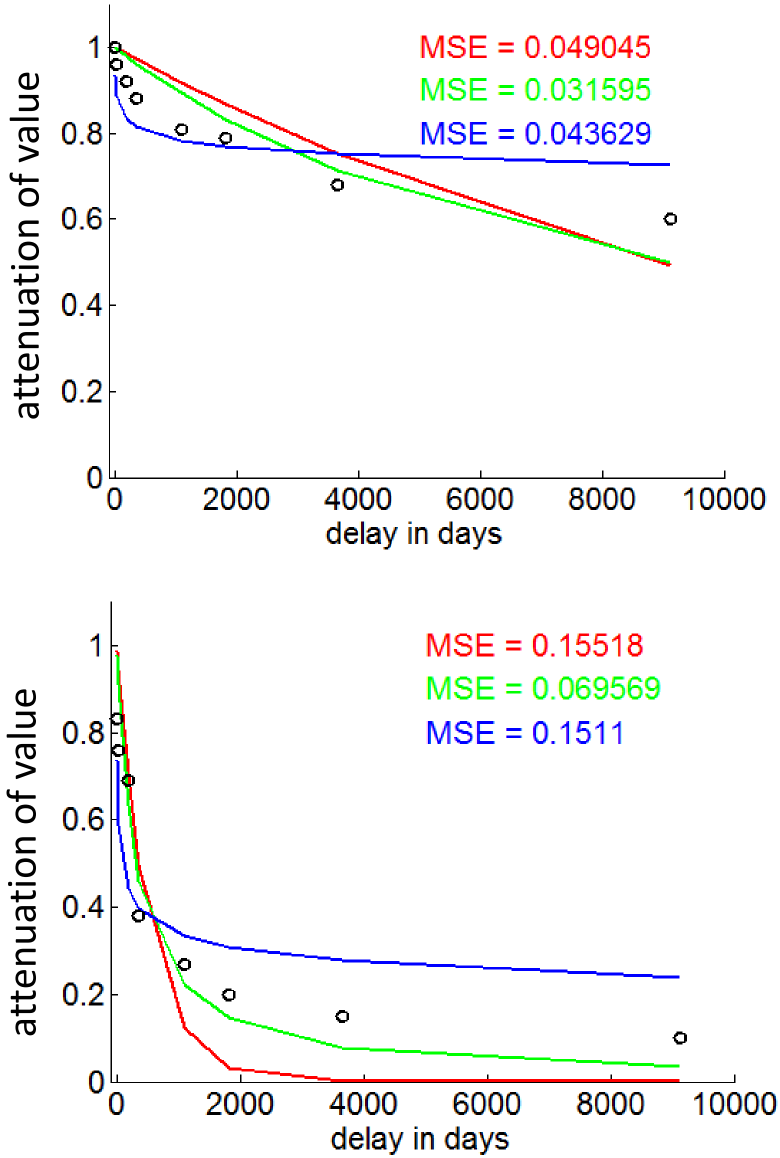

Why hyperbolic?

- Uncertain hazard rates (Sozou)
- Two or more processes with different time scales (Laibson, Kurth-Nelson&Redish)
- Non-linear time estimation (Bossaerts)

Data from Vuchinich and Simpson (1998) Exp Clin Psychopharm 6:292

Important:

You should fit subjects individually, rather than fitting averaged data.

If the individual data are exponential, the averaged data *will* be hyperbolic!

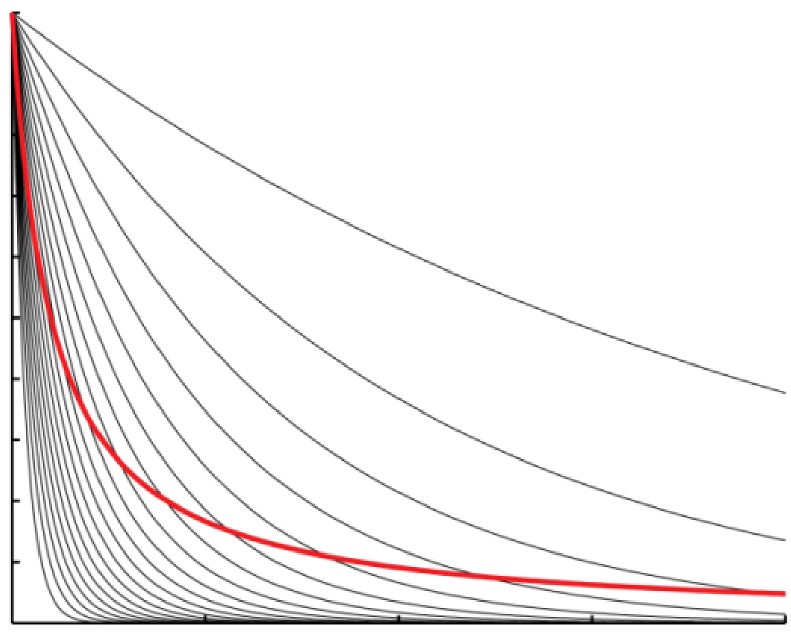

Non-exponential discounting

How to measure discounting?

What would you prefer?

$500 right now	OR	$1000 in a week
$750 right now	OR	$1000 in a week
$875 right now	OR	$1000 in a week
$937 right now	OR	$1000 in a week
$969 right now	OR	$1000 in a week
$984 right now	OR	$1000 in a week

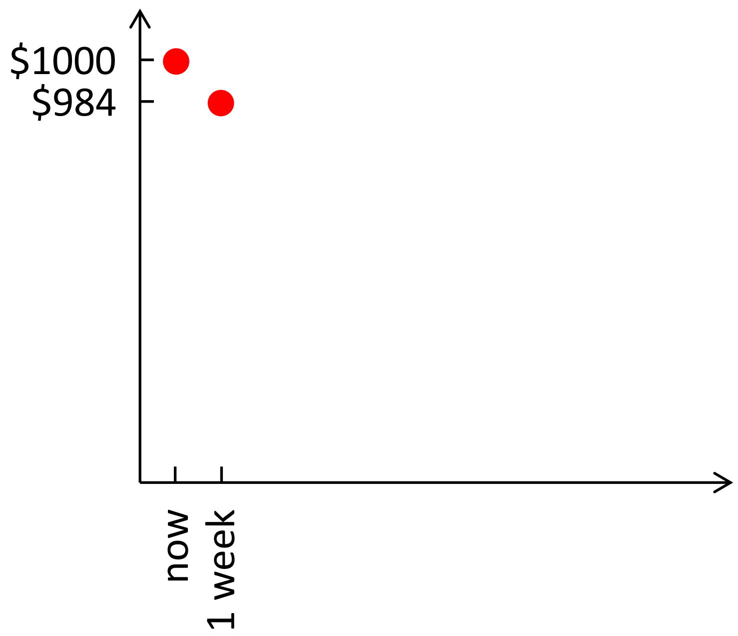

How to measure discounting?

What would you prefer?

$500 right now	OR	$1000 in 5 years
$250 right now	OR	$1000 in 5 years
$375 right now	OR	$1000 in 5 years
$437 right now	OR	$1000 in 5 years
$406 right now	OR	$1000 in 5 years
$391 right now	OR	$1000 in 5 years

$$\frac{1}{1+kd}$$

best fit k = 0.0009

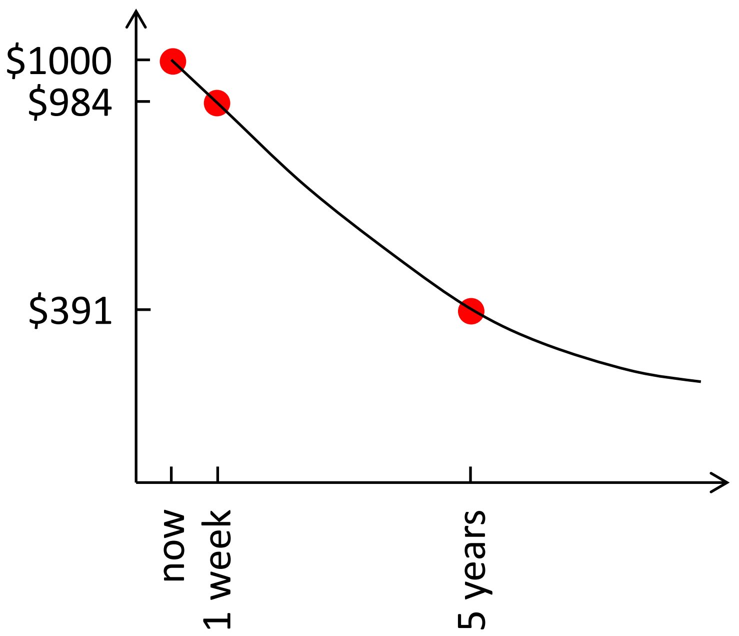

Area under the curve (AUC)

A non-parametric alternative to function fitting

AUC = $(7\ days - 0\ days) \cdot \frac{\$1000+\$984}{2} + \cdots$

Useful if an experimental manipulation could make discounting more or less hyperbolic!

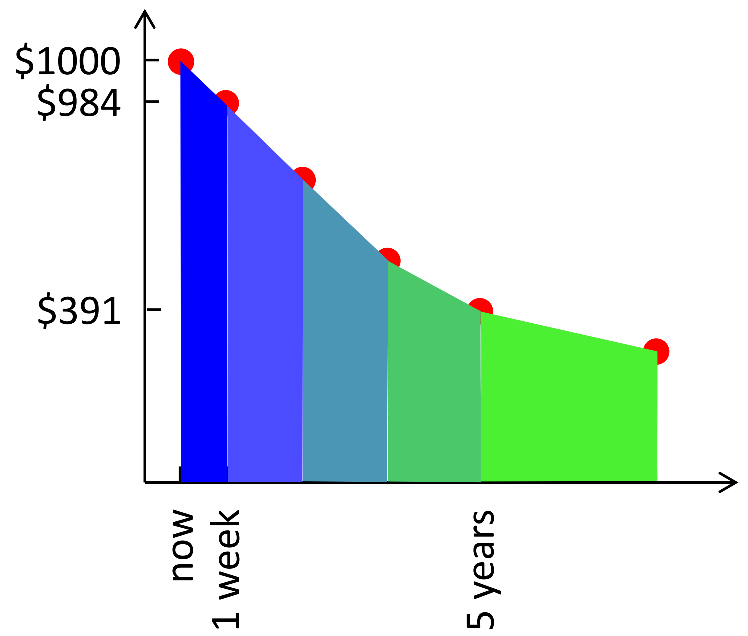

Subject makes a sequence of choices, *D*

We assume they're using hyperbolic discounting with rate *k*

What is the value of *k* that maximizes $P(D|k)$?

Maximum likelihood

The *subjective* value, *V*, of a reward is the magnitude, *R*, discounted by the delay, *d*

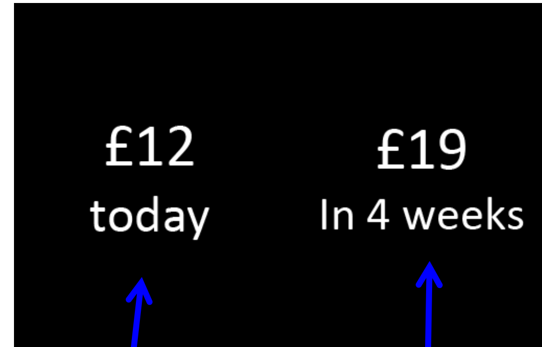

$V_1 = 12$ $\quad$ $V_2 = 14.8$

$$V = R \cdot \frac{1}{1 + k \cdot d}$$

$$V_1 = 12 \cdot \frac{1}{1 + 0.01 \cdot 0} = 12$$

$$V_2 = 19 \cdot \frac{1}{1 + 0.01 \cdot 28} = 14.8$$

$$V_2 = 19 \cdot 0.78 = 14.8$$

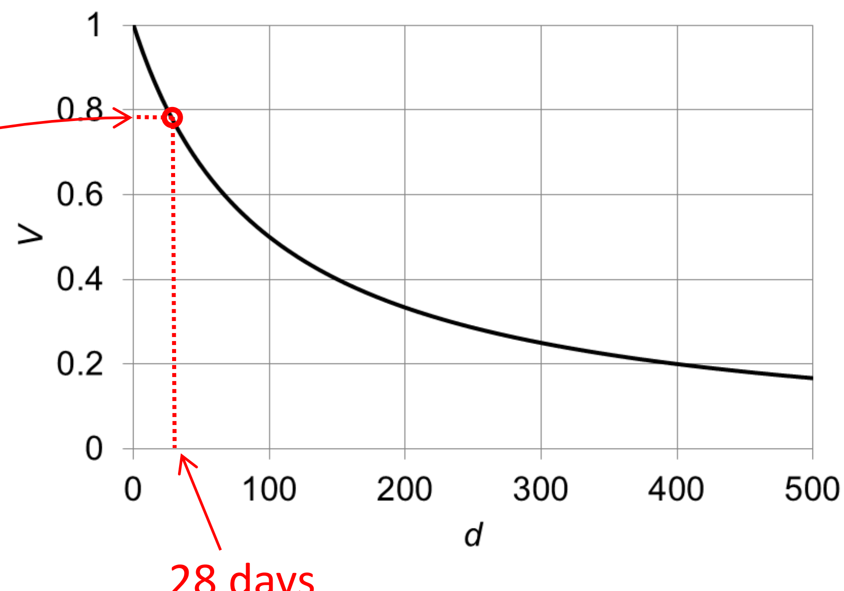

So how likely is each choice?

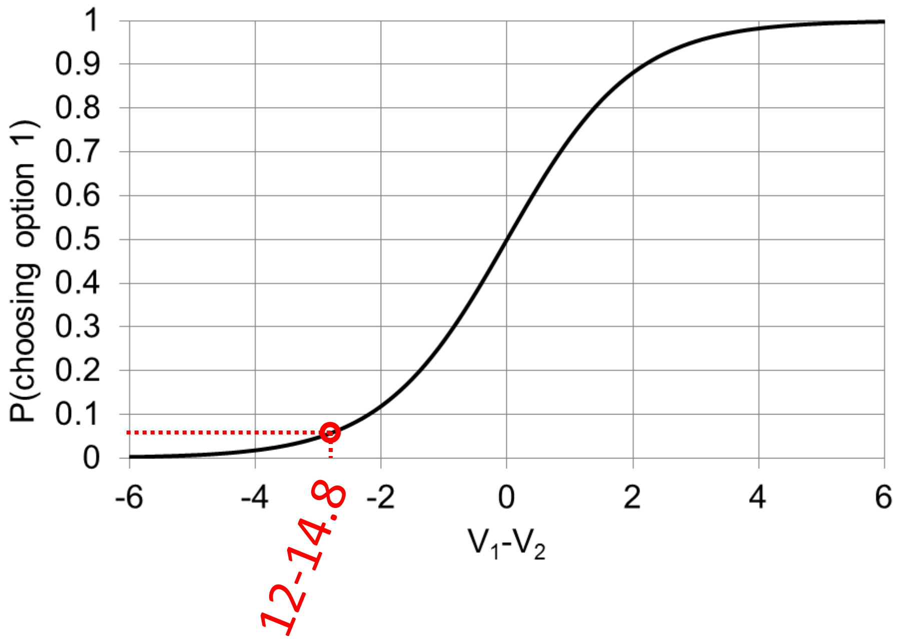

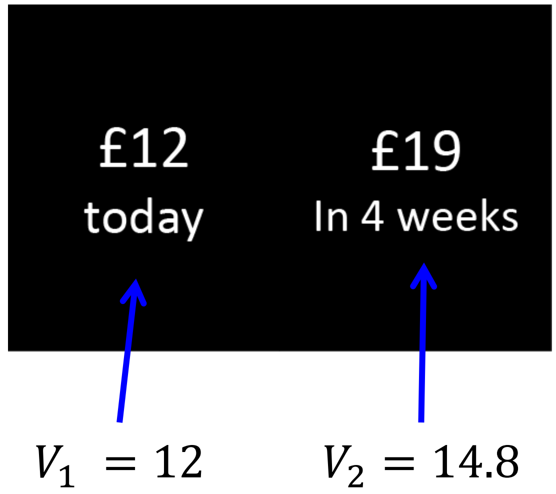

$$P(\text{choosing option 1} \mid k{=}0.01) = \frac{1}{1+e^{-\beta\cdot(V_1-V_2)}} = 0.06$$

$$P(\text{choosing option 2} \mid k{=}0.01) = \frac{1}{1+e^{-\beta\cdot(V_2-V_1)}} = 0.94$$

Maximum likelihood

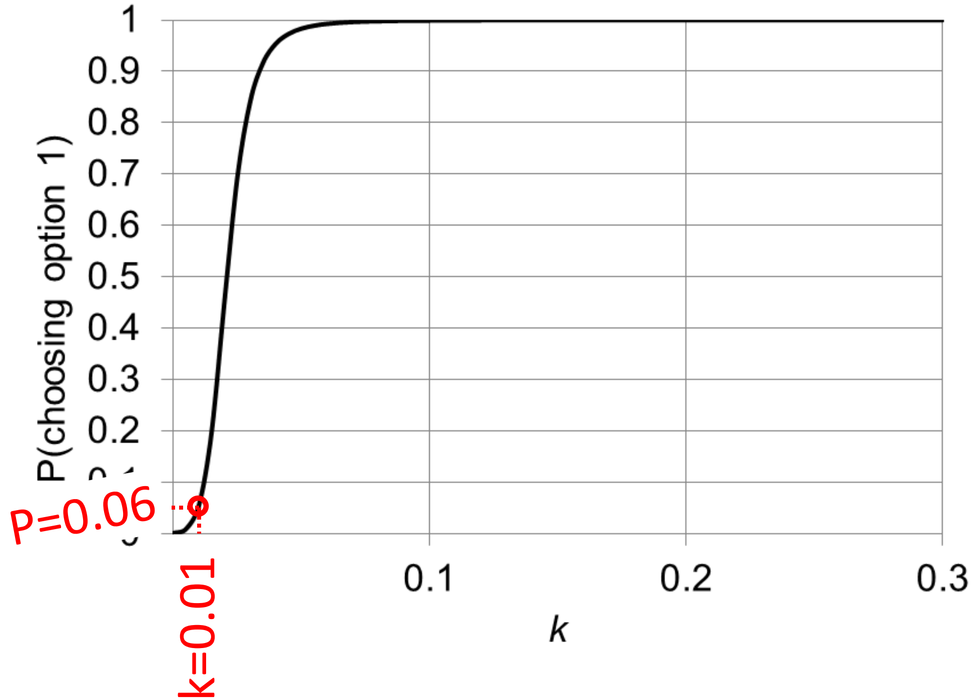

P(choosing option 1 | k=0.01) = $\frac{1}{1+e^{-\beta\cdot(V_1-V_2)}}$ = 0.06

Let's suppose the subject did choose option 1. What *k* did they probably have?

Maximum likelihood

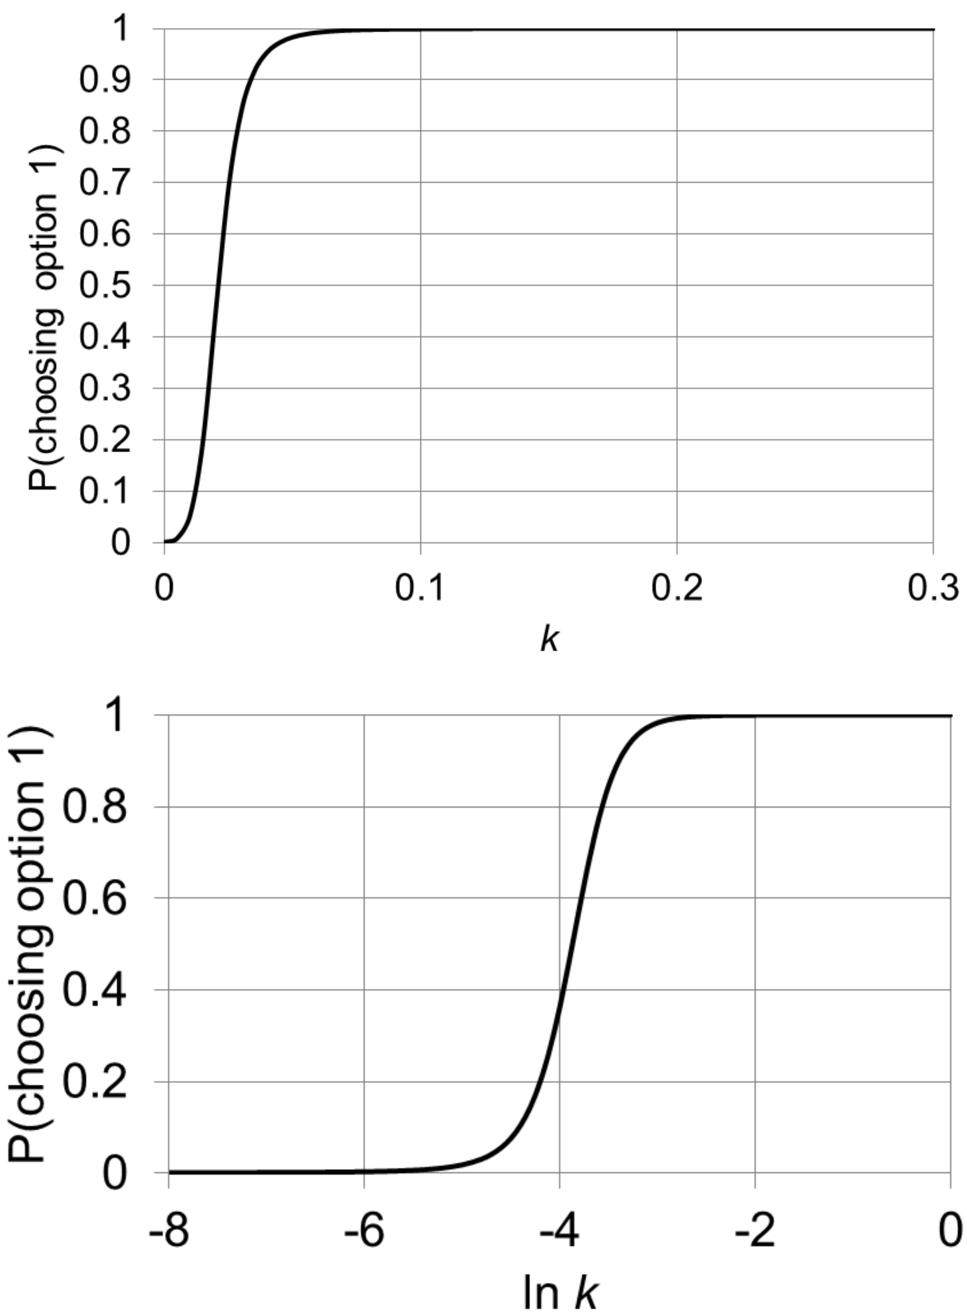

Maximum likelihood

The *most likely* ln k is +∞

So we need to observe multiple choices to make a good guess about the subject's real discount rate

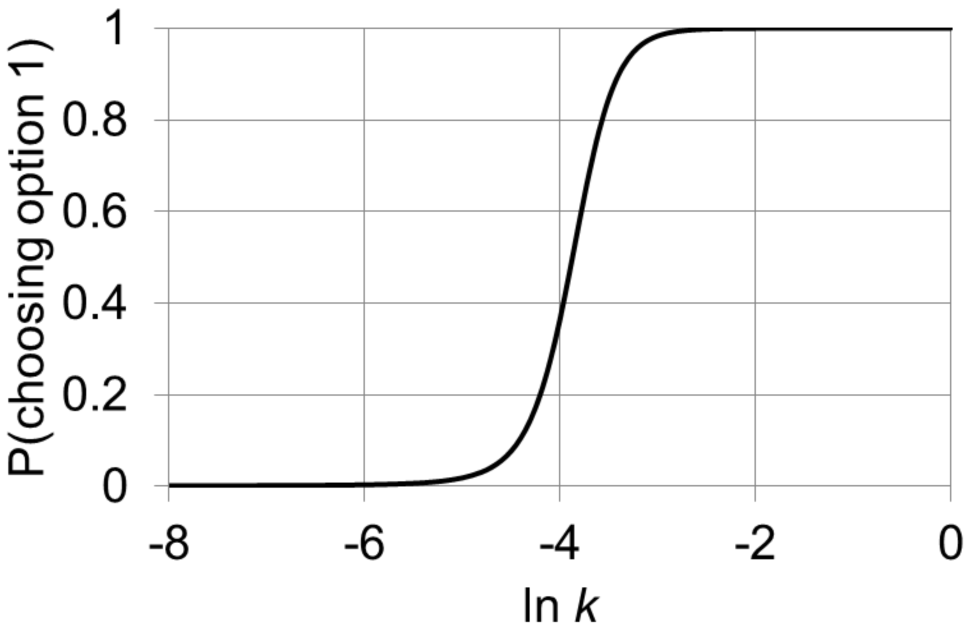

Maximum likelihood

Subject makes a sequence of choices, *D*

We assume they're using hyperbolic discounting with rate *k*

What is the value of *k* that maximizes $P(D|k)$?

Maximum likelihood

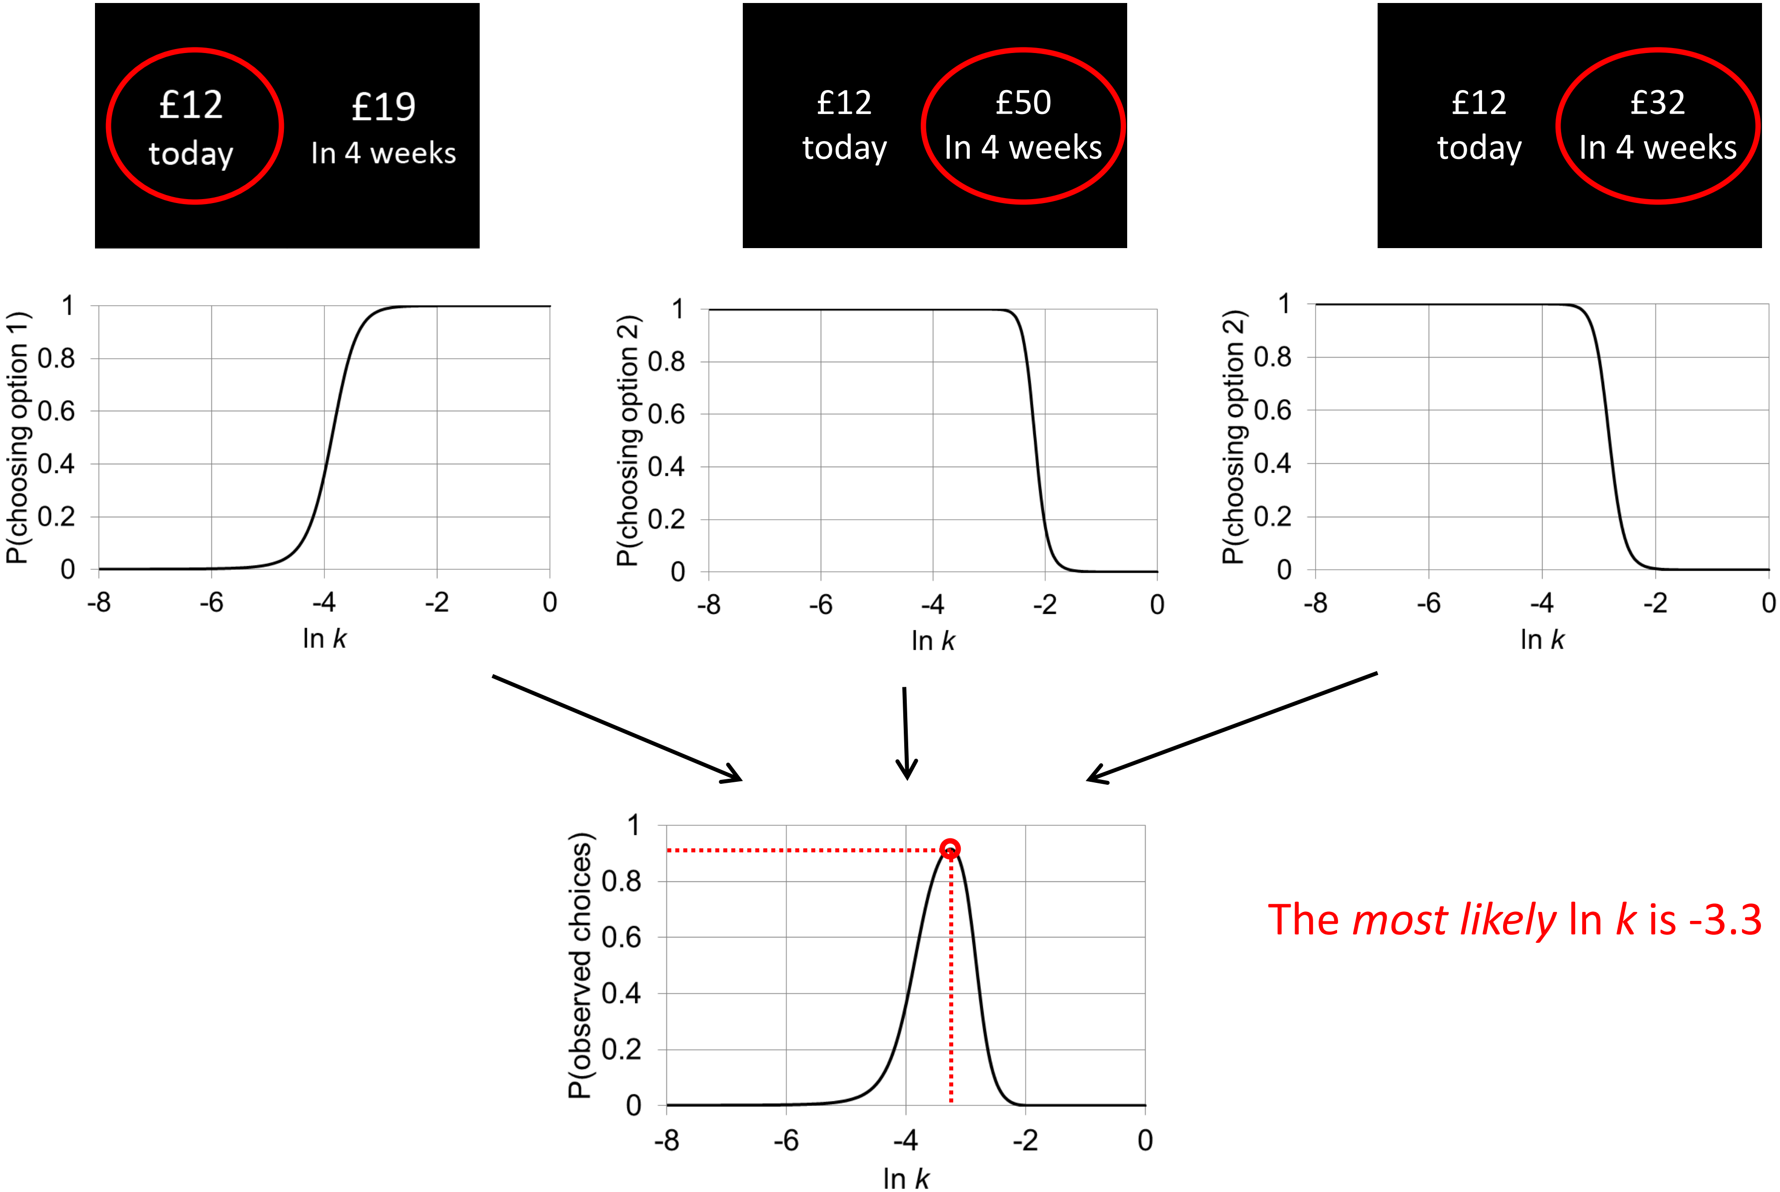

The *most likely* ln k is -3.3

Maximum likelihood

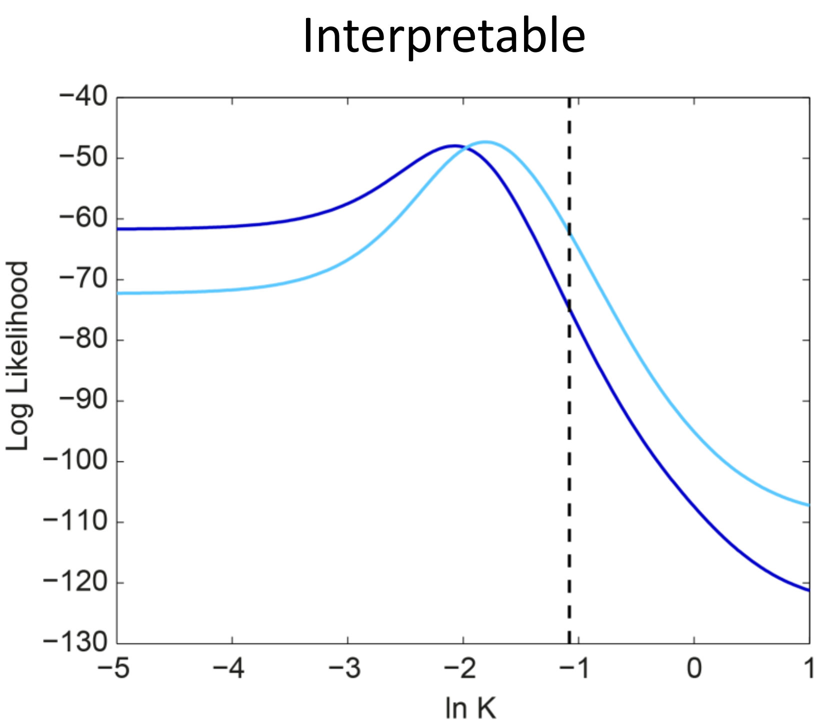

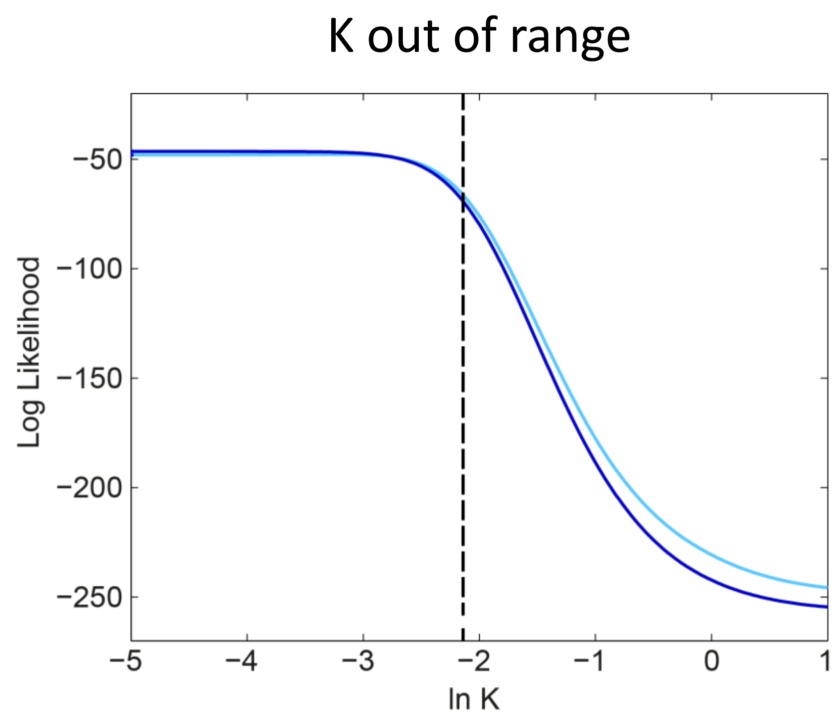

How can we design questions to get the most information out of the fewest questions?

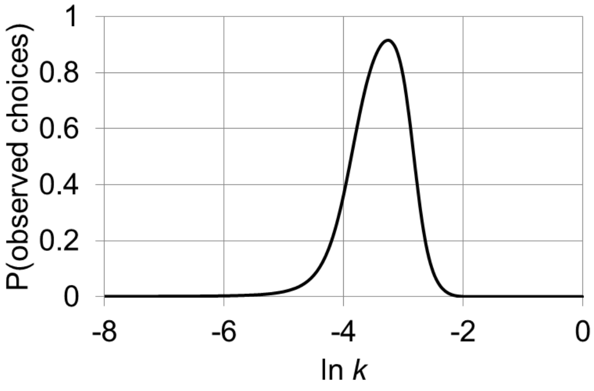

The expected value of ln *k* is -3.6

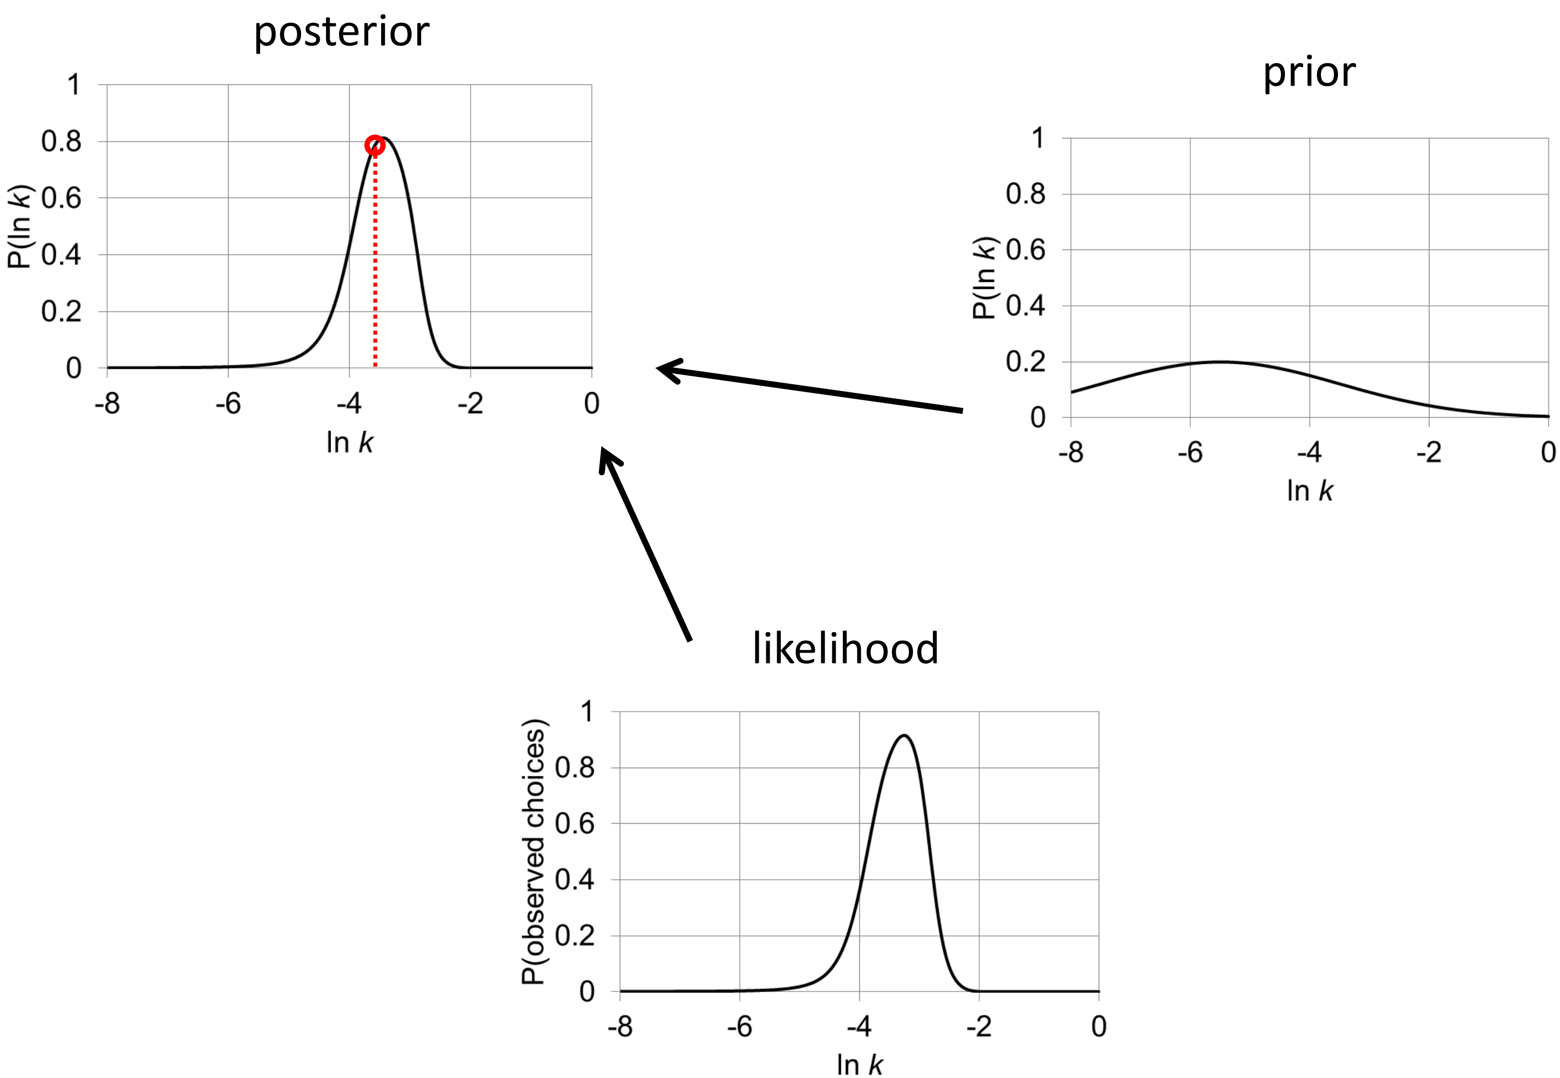

The expected value of ln *k* is -3.6

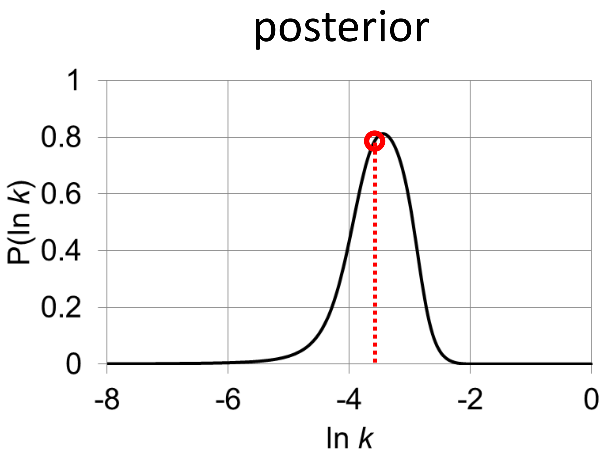

£? today | £? in ?

choose a random delay and delayed amount

£? today | £21 in 2 weeks

£15 today | £21 in 2 weeks

so if the subject's ln *k* is really -3.6 (i.e., our current best estimate), then this should be the hardest question to answer

$$V_2 = \frac{21}{1 + e^{-3.6} \cdot 14} = 15$$

$$V_1 = V_2 = 15$$

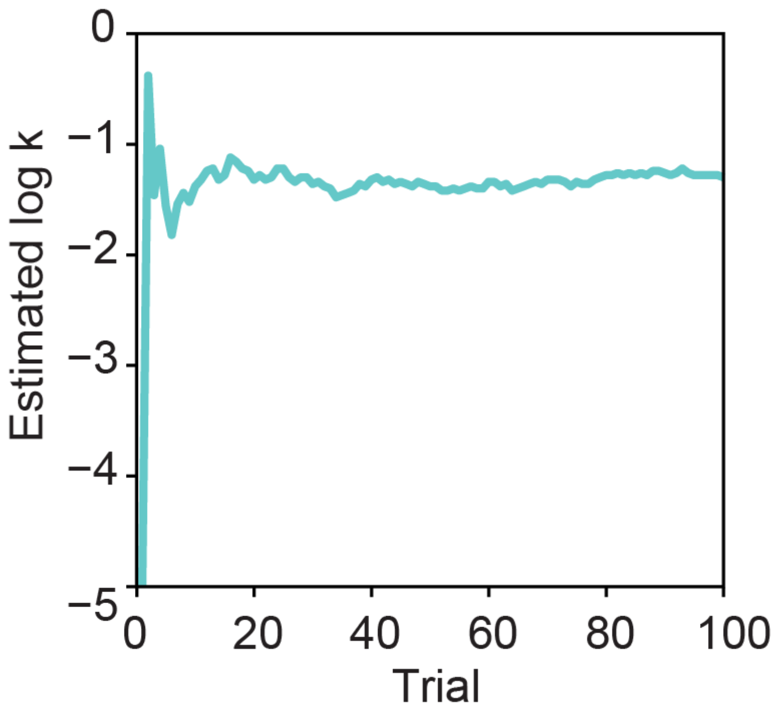

Not incentive compatible

Can instead use random questions or optimized random questions

Garvert MM, Moutoussis M, Kurth-Nelson ZL, Behrens TE, Dolan RJ (in preparation)

Fitting beta

$$P(\text{choosing option 1} \mid \beta) = \frac{1}{1+e^{-\beta\cdot(V_1-V_2)}}$$

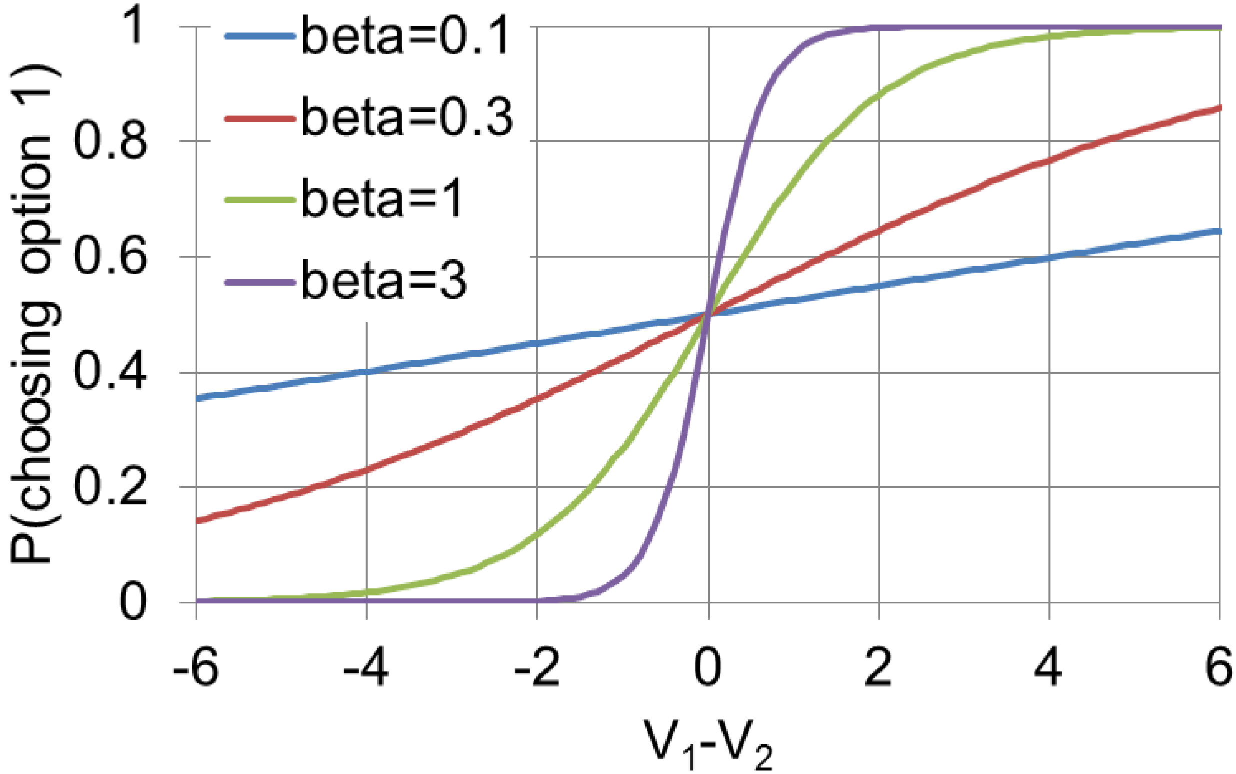

Fitting beta

- When beta is allowed to be small, *k* can be contaminated

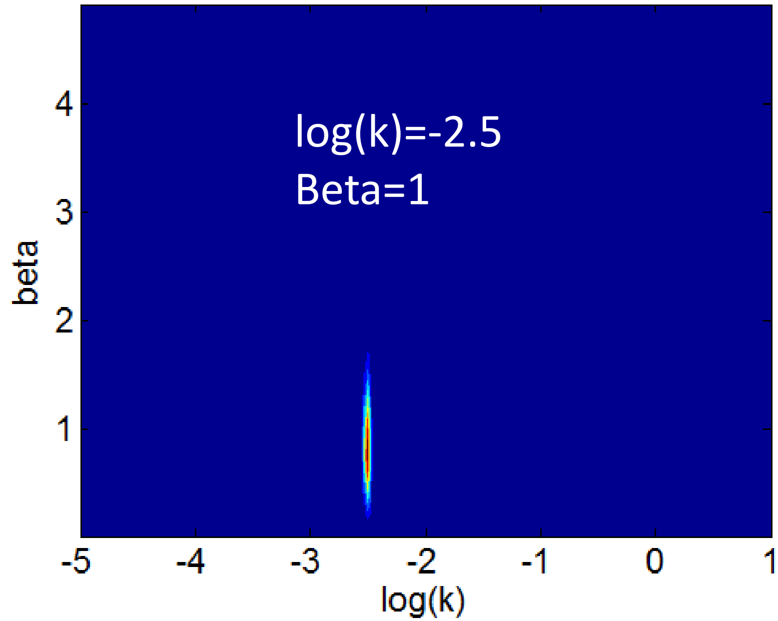

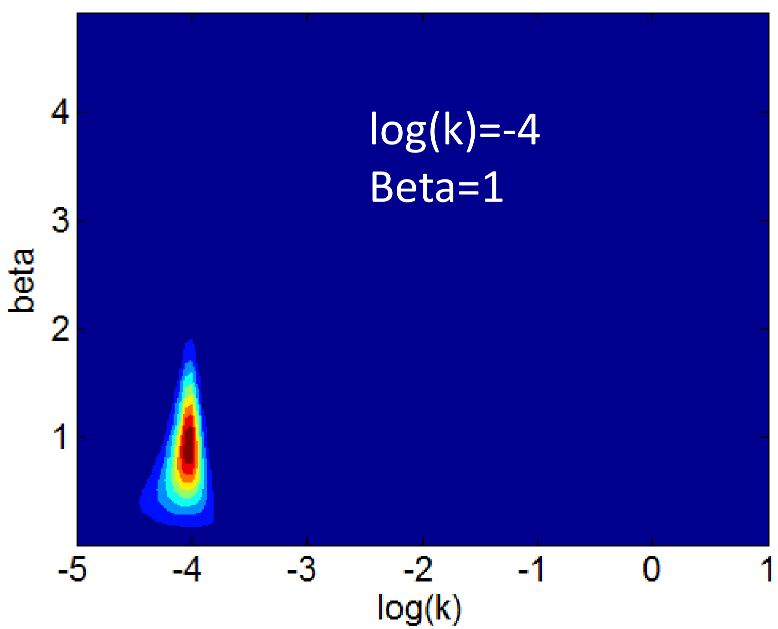

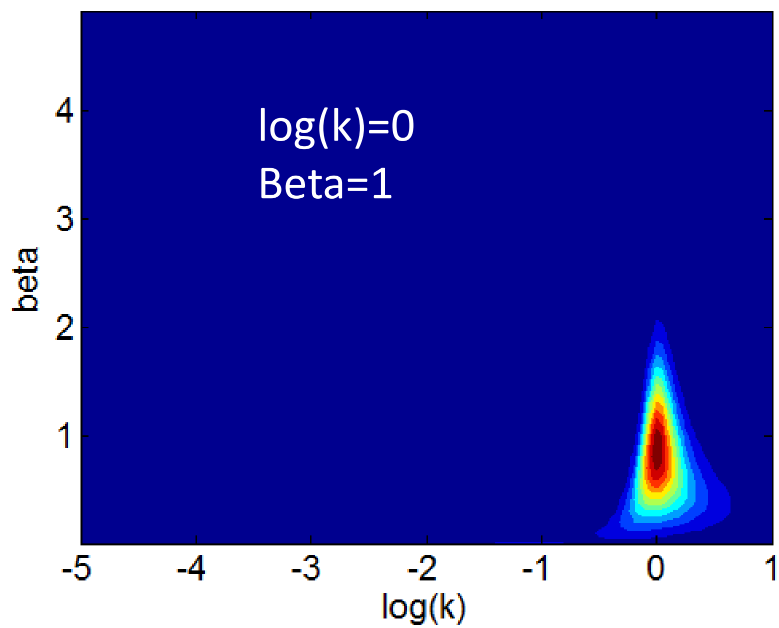

Fitting beta

- When beta is allowed to be small, *k* can be contaminated
- beta can take lots of trials to converge

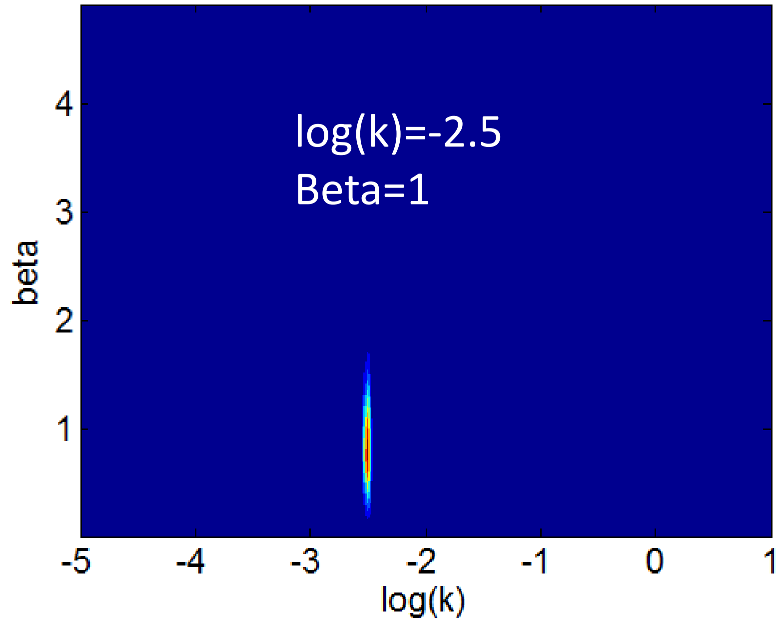

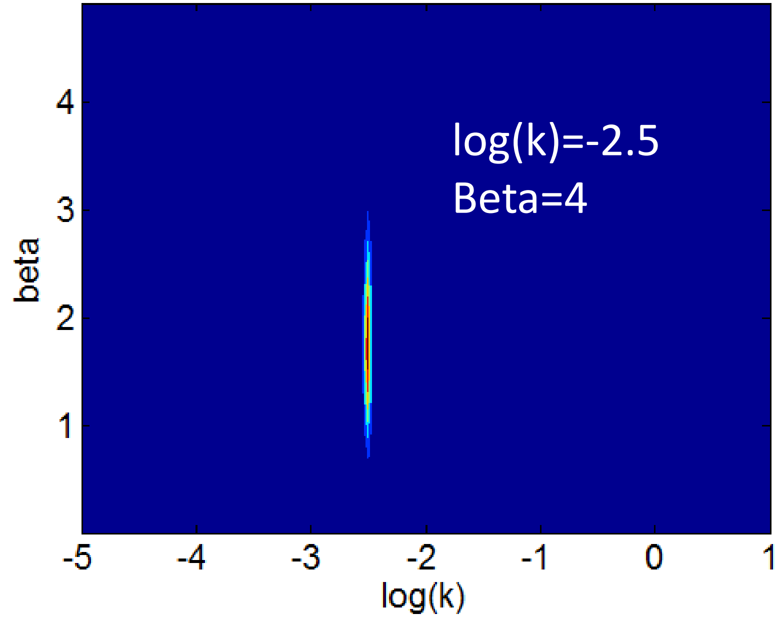

Utility curvature

$$V = R \cdot \frac{1}{1 + k \cdot d}$$

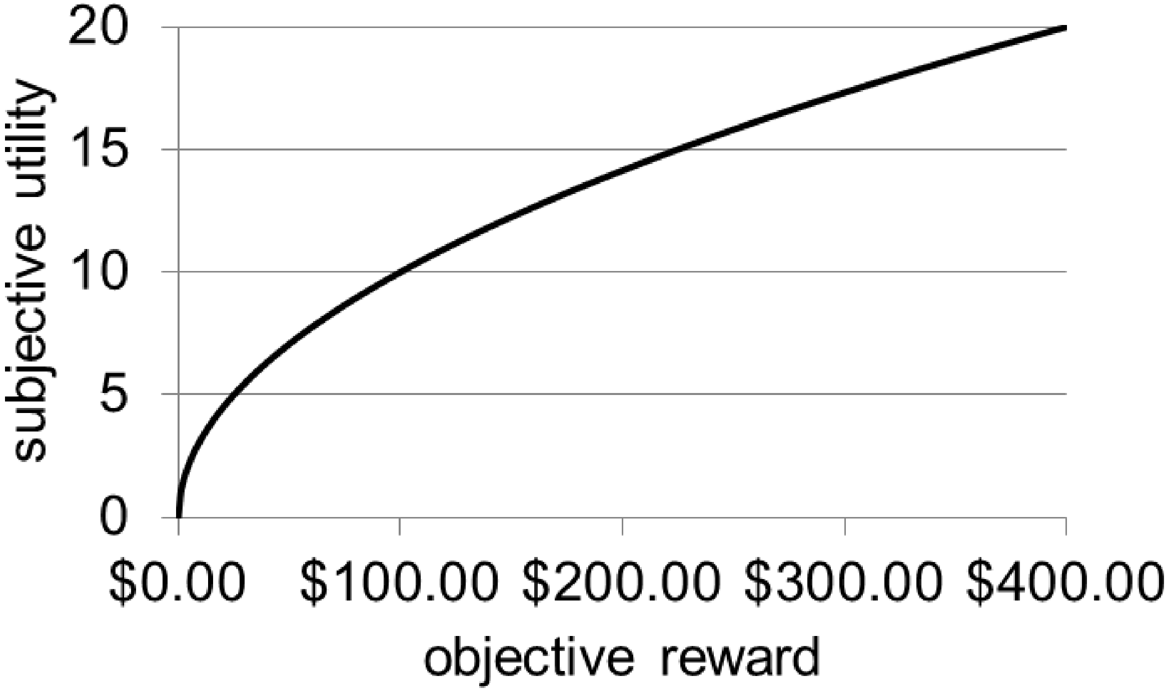

$100 now OR $200 in a year

Utility curvature

Model number (Eq.)	Sum AIC	Delta AIC	Akaike weight
2, (4)—Hyperbolic discounting of utility	3595	0	1
1, (2)—Hyperbolic discounting of magnitude	3630	35	2.51E-08

A change in utility curvature can look like a change in discount rates!

Pine A (2010) J Neurosci 29:9575

Other task design issues

- Primary vs. secondary rewards
- Real vs. hypothetical rewards
- Large vs. small rewards

Primary vs. secondary rewards

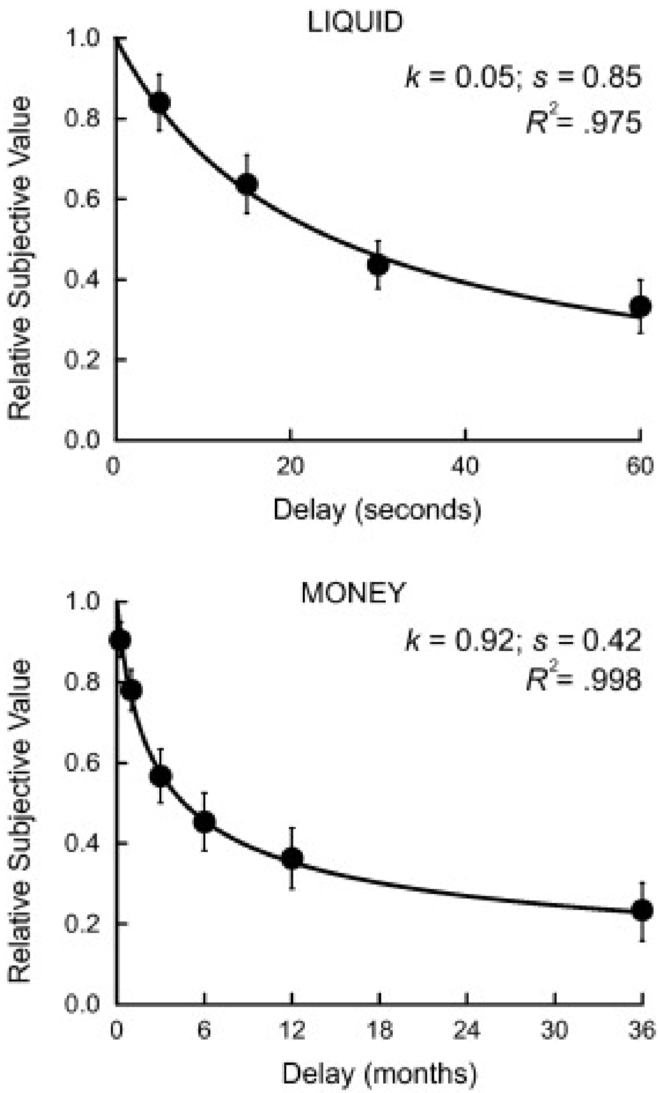

Little-to-no correlation between discounting for juice and money

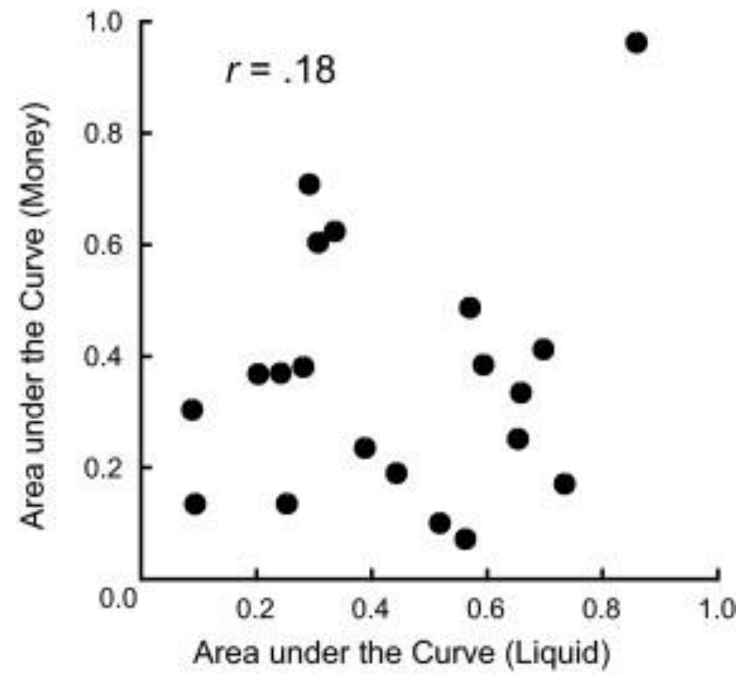

Jimura K et al (2011) Behav Process 87:253

Real and hypothetical rewards discount the same

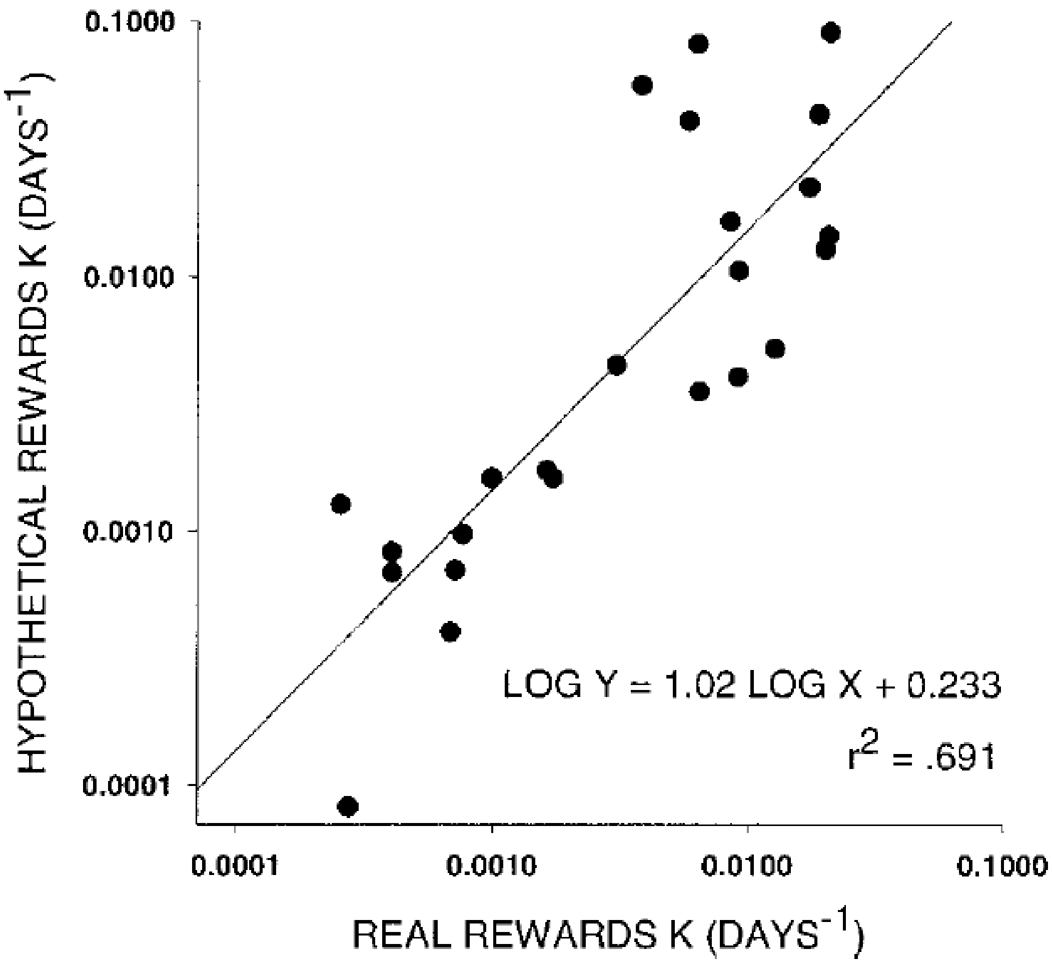

Johnson and Bickel (2002) JEAB 77:129

Larger rewards discount less steeply

$10 in the future

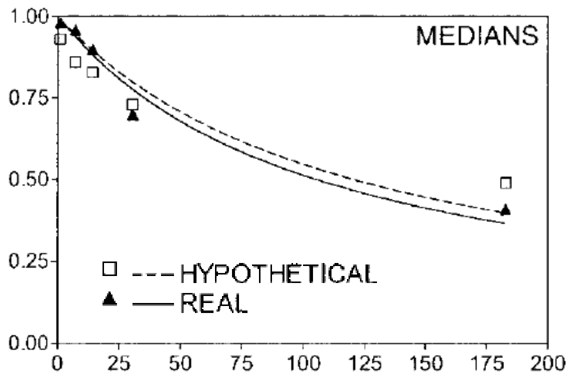

delay in days

$25 in the future

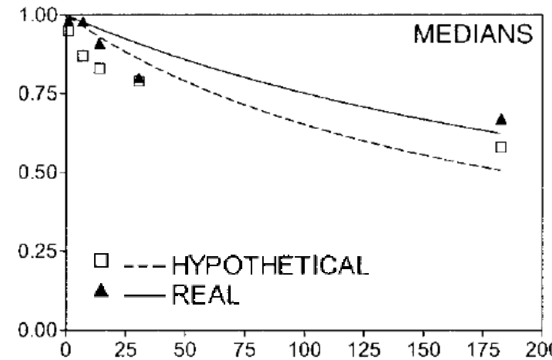

delay in days

$100 in the future

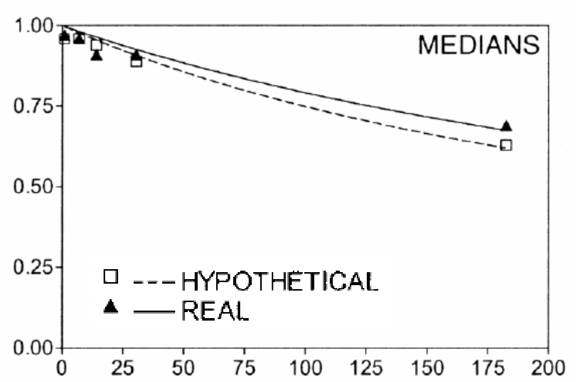

delay in days

$250 in the future

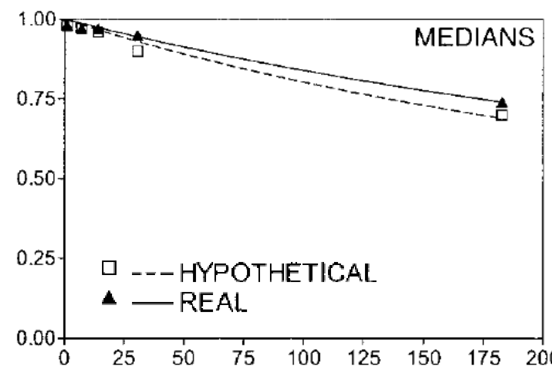

delay in days

Johnson and Bickel (2002) JEAB 77:129

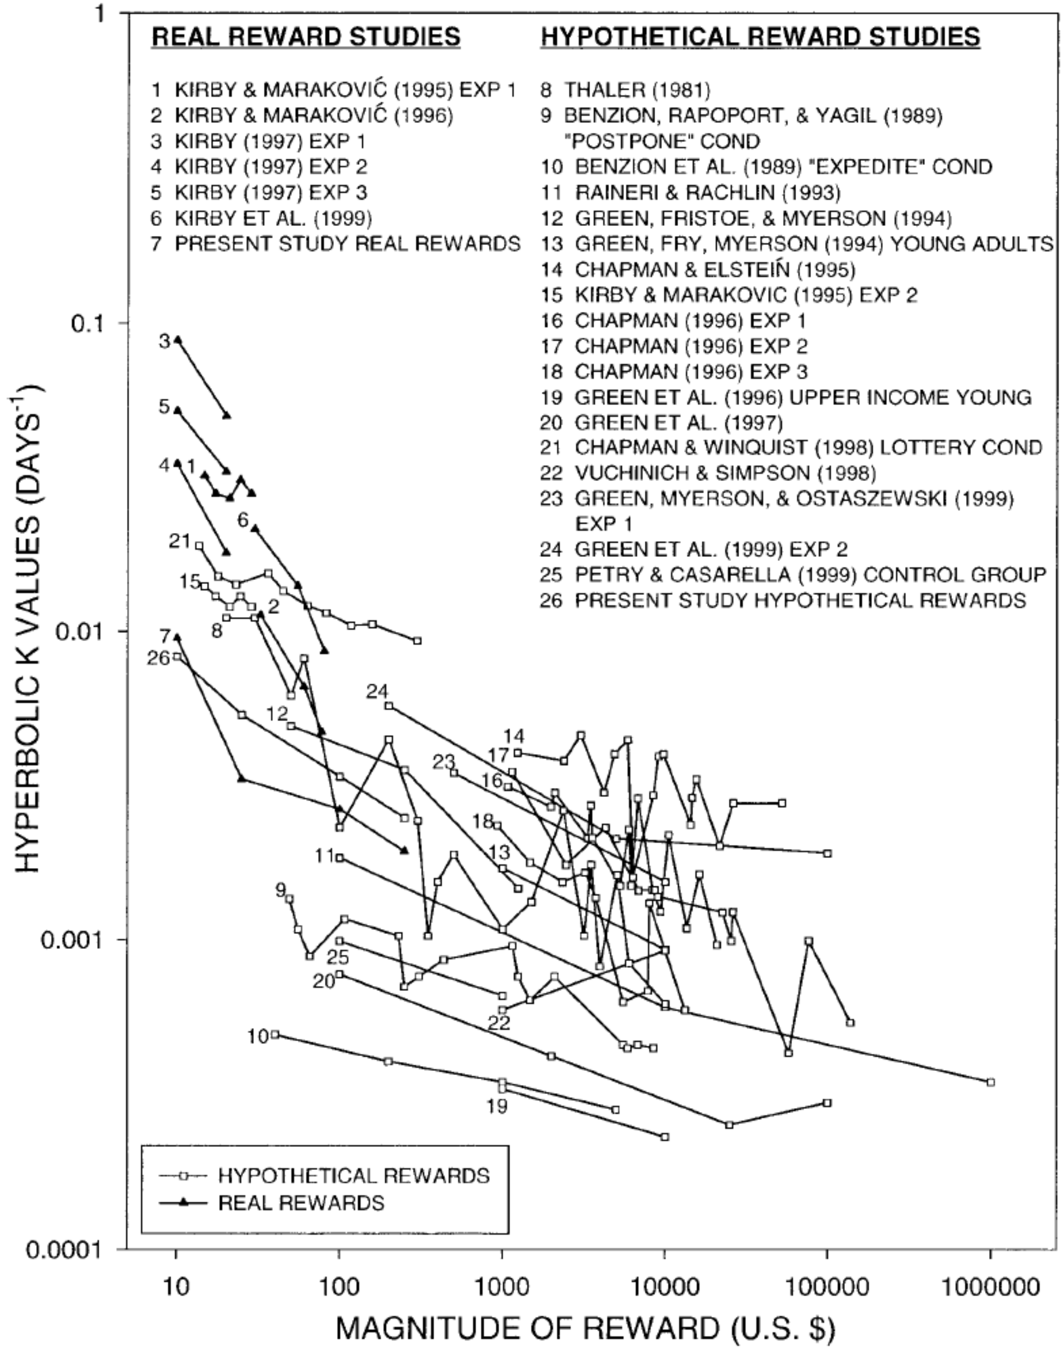

Larger rewards discount less steeply

Johnson and Bickel (2002) JEAB 77:129

Outline

1) Delay discounting is important

2) Designing a task & analyzing the data

3) What delay discounting measures

4) Modelling discounting

What are we measuring?

- Discounting is normally stable, but also surprisingly labile
- Paradoxes of discounting
  - Violation of valuation model
  - Reverse discounting

Stability of discounting

Stability over two weeks

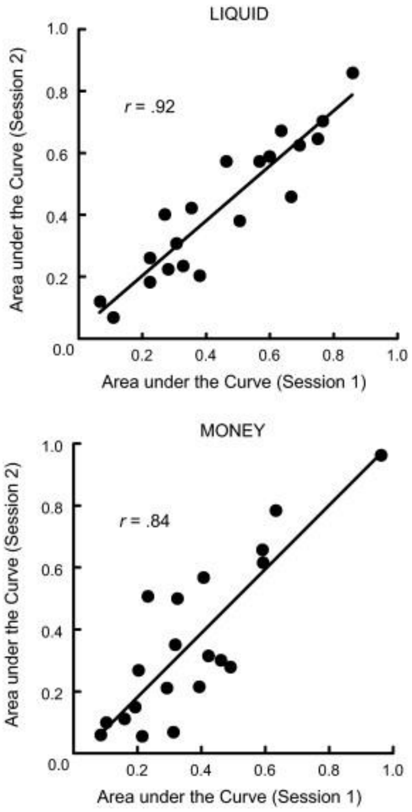

Jimura K et al (2011) Behav Process 87:253

Stability over three months

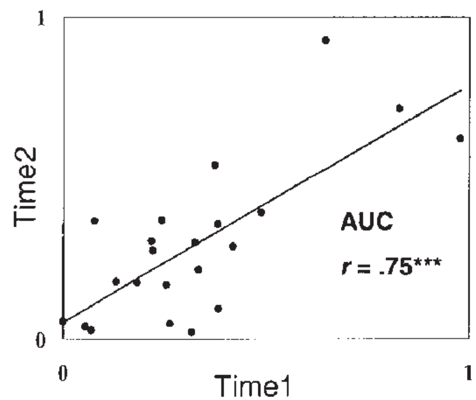

Ohmura Y at al (2006) Exp Clin Psychopharm 14:318

Stability over one year

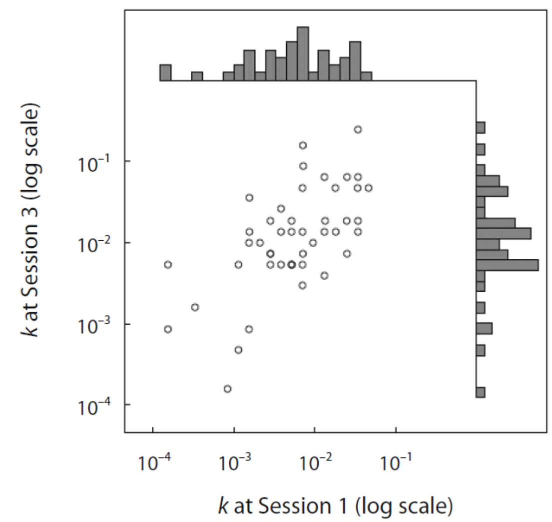

Kirby KN (2009) Psych Bull 16:457

Discounting is modulated by social conformity

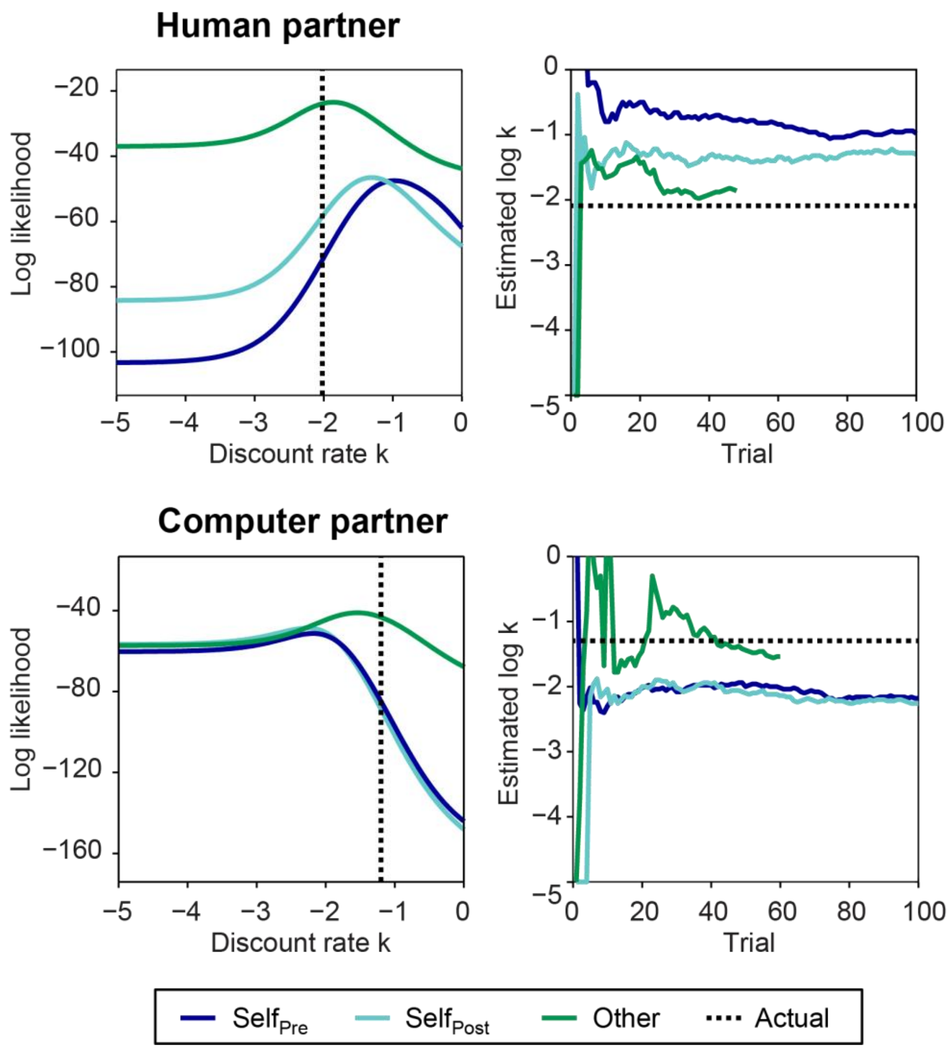

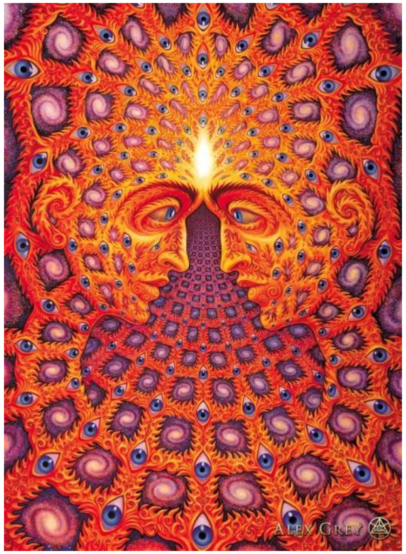

Garvert MM, Moutoussis M, Kurth-Nelson ZL, Behrens TE, Dolan RJ (in preparation)

Discounting is modulated by social conformity

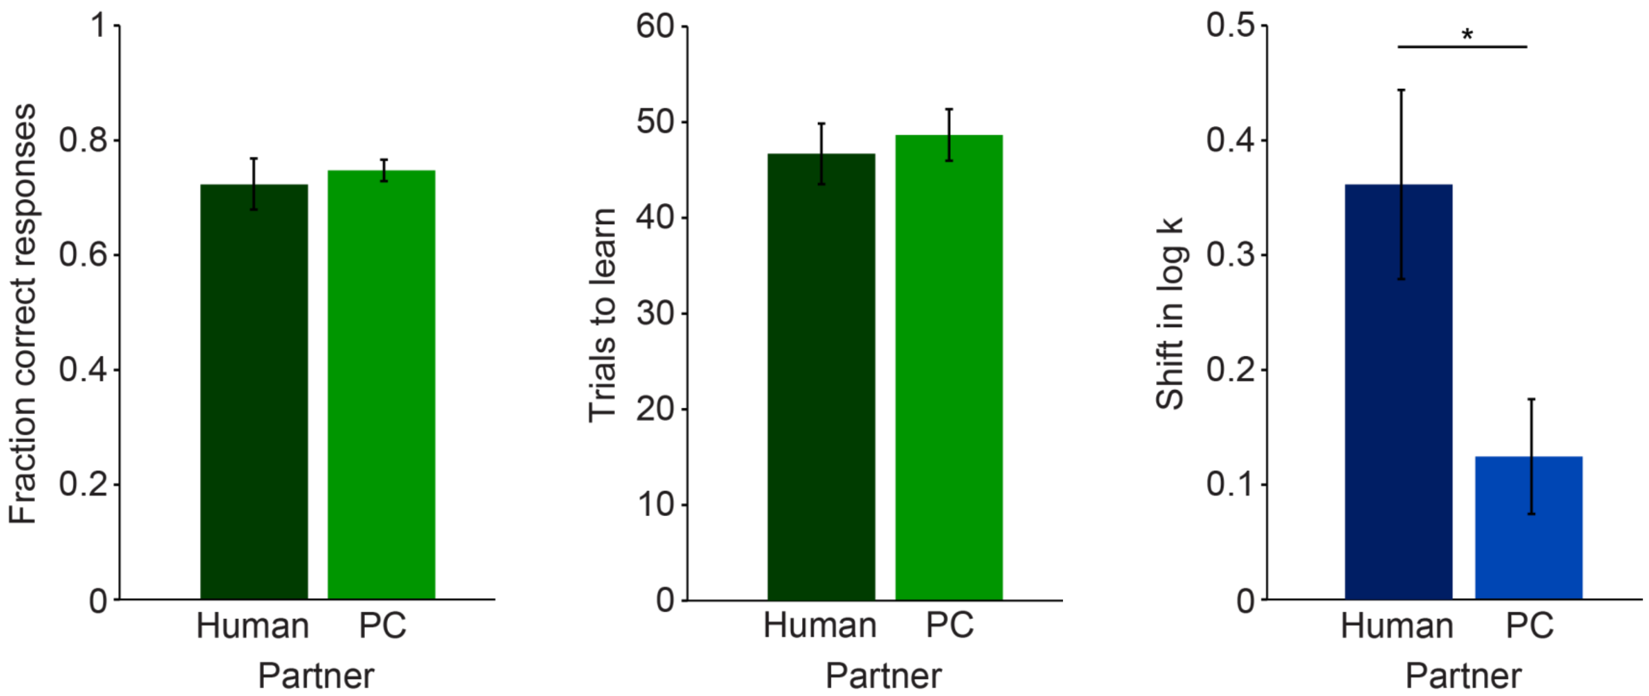

Garvert MM, Moutoussis M, Kurth-Nelson ZL, Behrens TE, Dolan RJ (in preparation)

Vivid imagination slows discounting

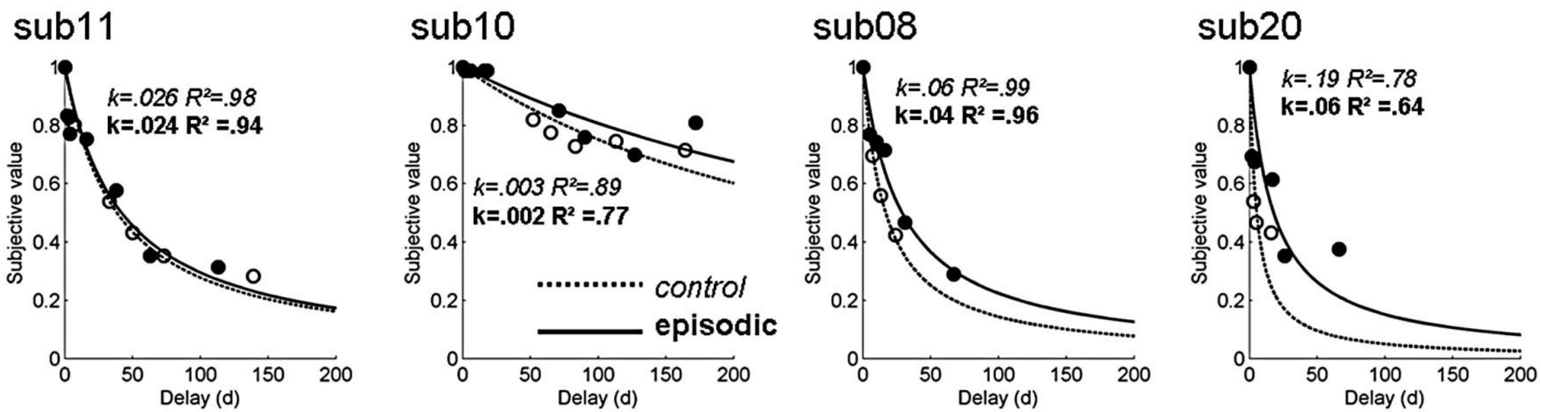

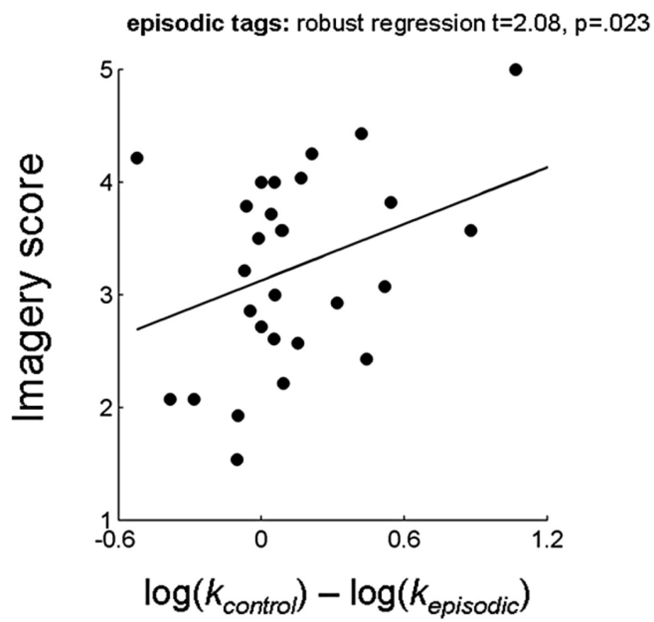

Peters J, Buchel C (2010) Neuron 66:138

Serotonin depletion makes discounting steeper

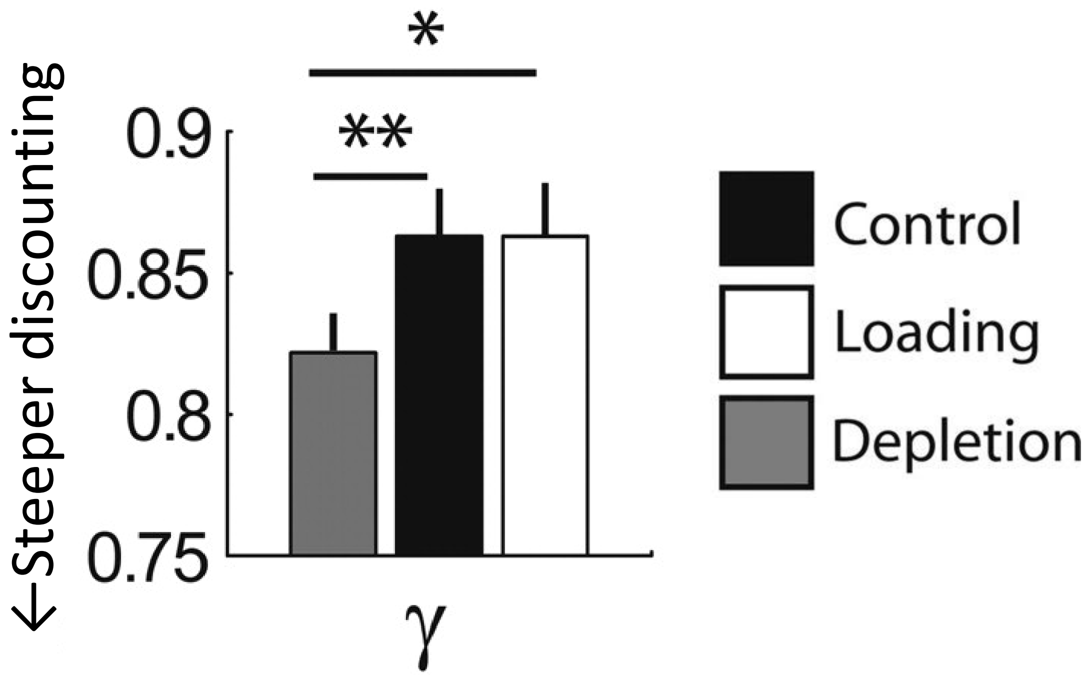

Schweighofer N et al (2008) J Neurosci 28:4528

Working memory training slows discounting

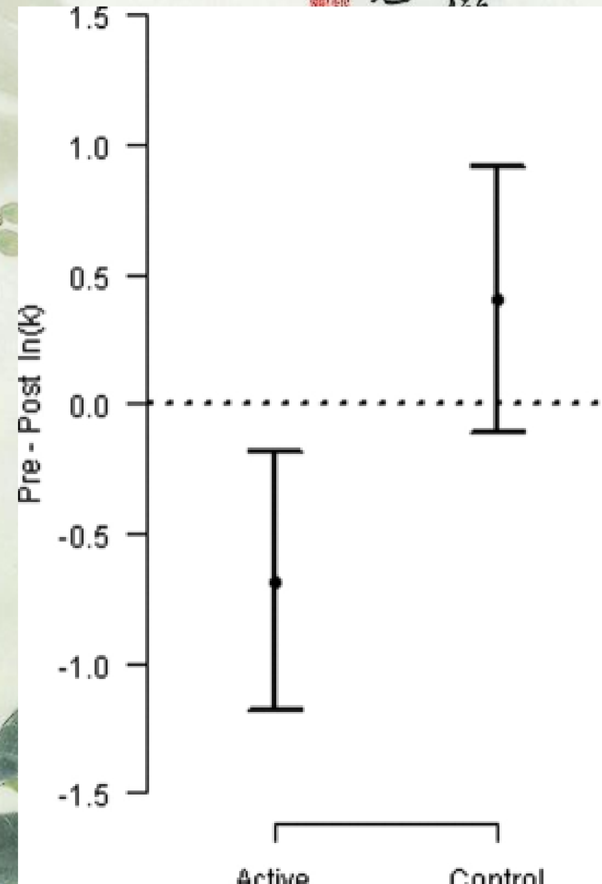

Bickel WK et al (2011) Biol Psych 69:260

Are choices evaluated independently?

$$V_1 = R_1 \cdot \frac{1}{1+k \cdot d_1} \quad \text{and} \quad V_2 = R_2 \cdot \frac{1}{1+k \cdot d_2}$$

Cross-commodity discounting

immediate-by-delayed
interaction, P=0.039

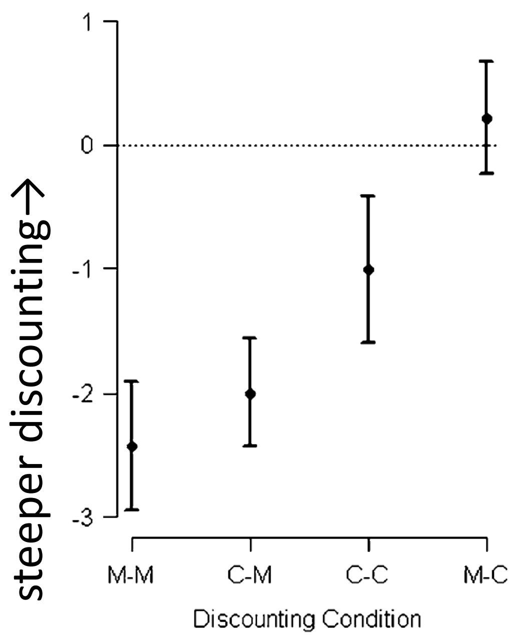

M-M: money now vs. money later
C-M: cocaine now vs. money later
C-C: cocaine now vs. cocaine later
M-C: money now vs. cocaine later

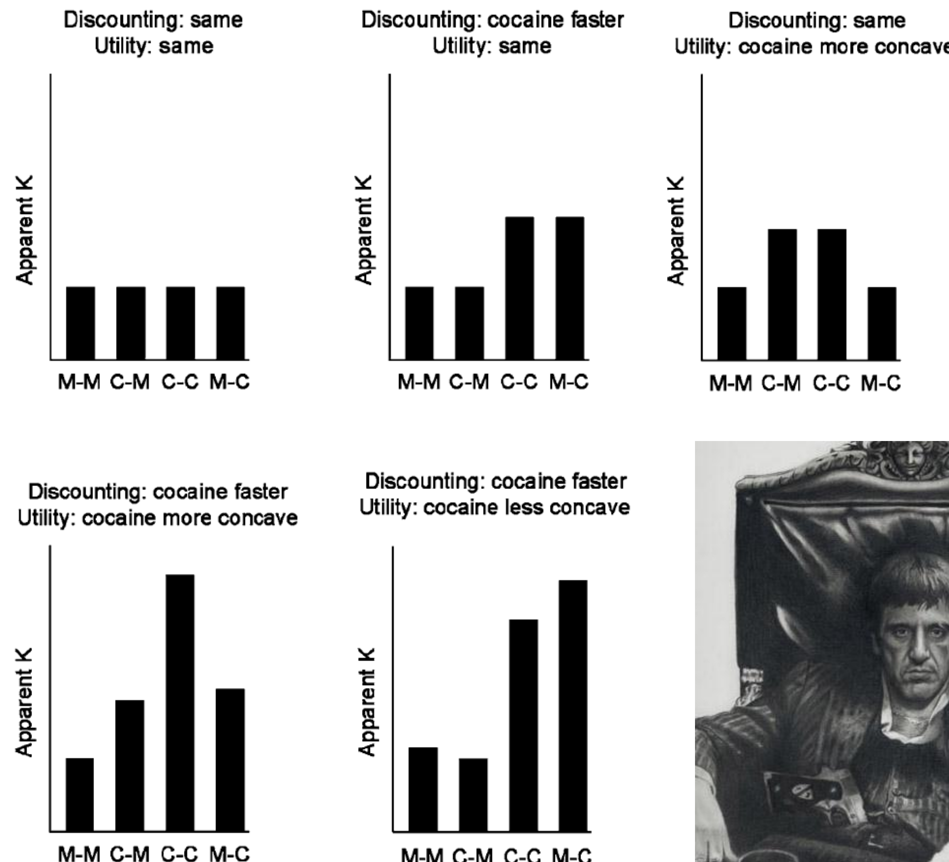

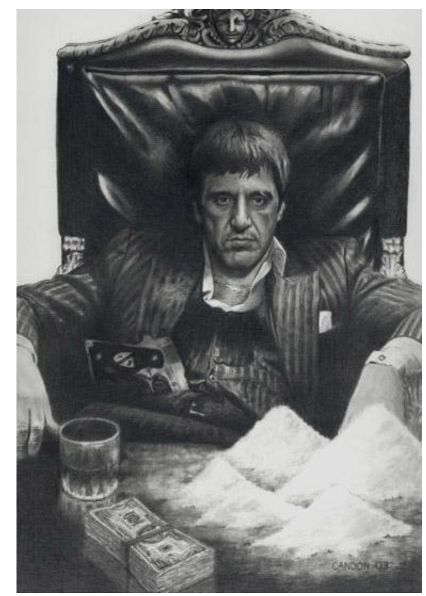

Bickel WK, Landes RD, Christensen DR, Jackson L, Jones BA, Kurth-Nelson ZL, Redish AD (2011) Psychopharmacology 217:177

Is the earliest outcome treated as "now"?

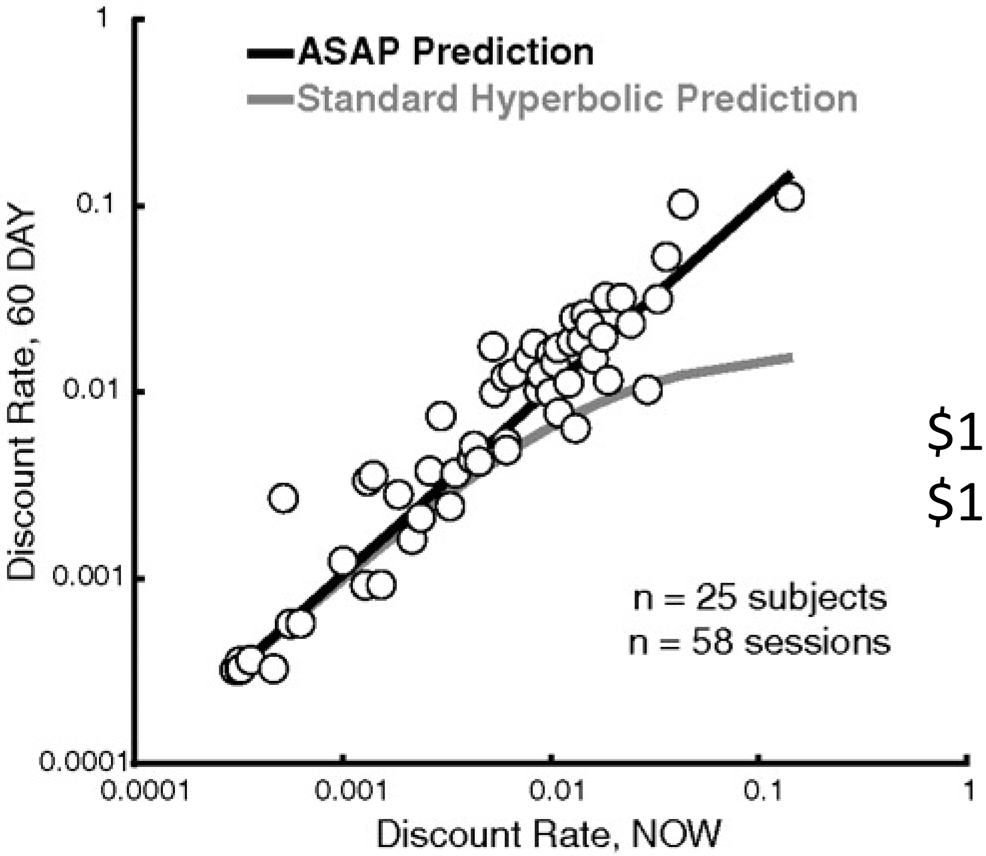

$10 NOW vs $15 in 20 days
$10 in 60 days vs $15 in 80 days

Kable JW and Glimcher PW (2010) J Neurophys 103:2513

Is the earliest outcome treated as "now"?

But,

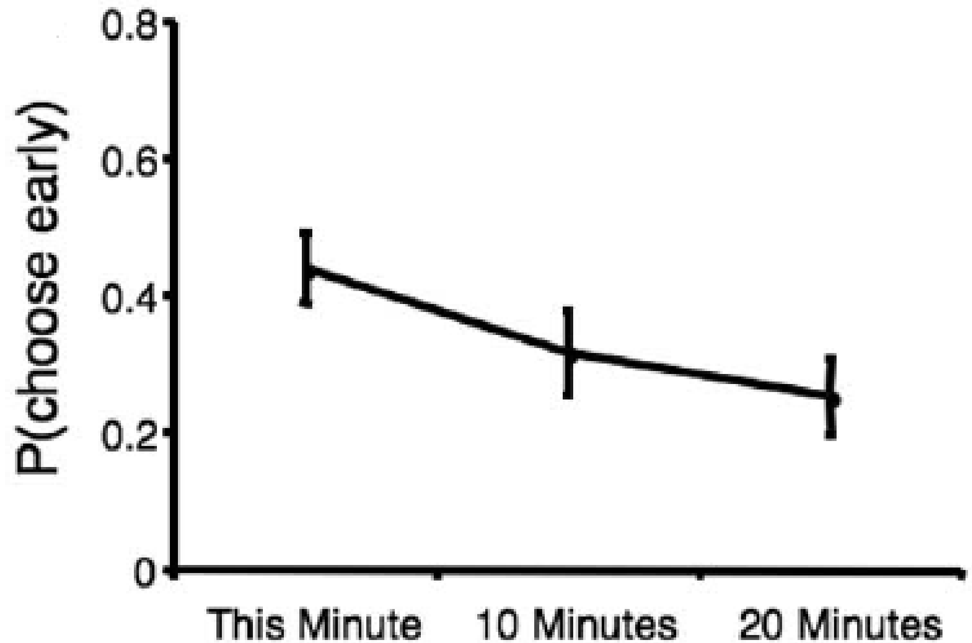

delay to earlier outcome

McClure SM et al (2007) J Neurosci 27:5796

Savoring and dread

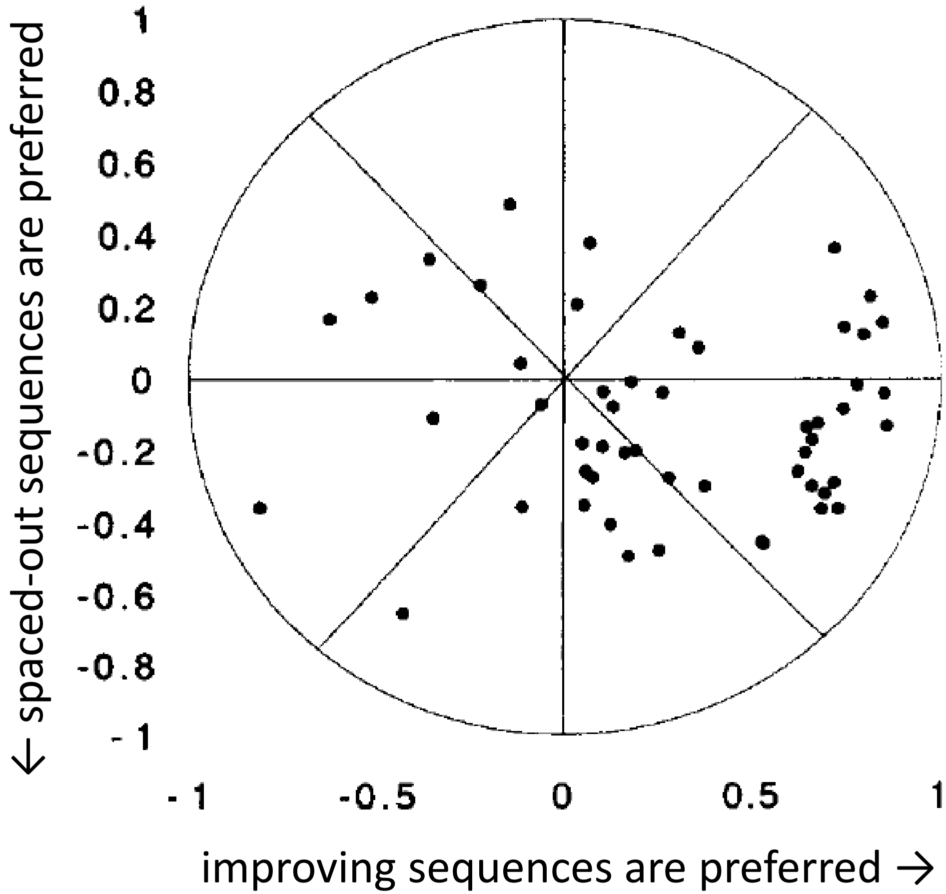

Loewenstein GF, Prelec D (1993) Psych Review 100:91

Savoring and dread

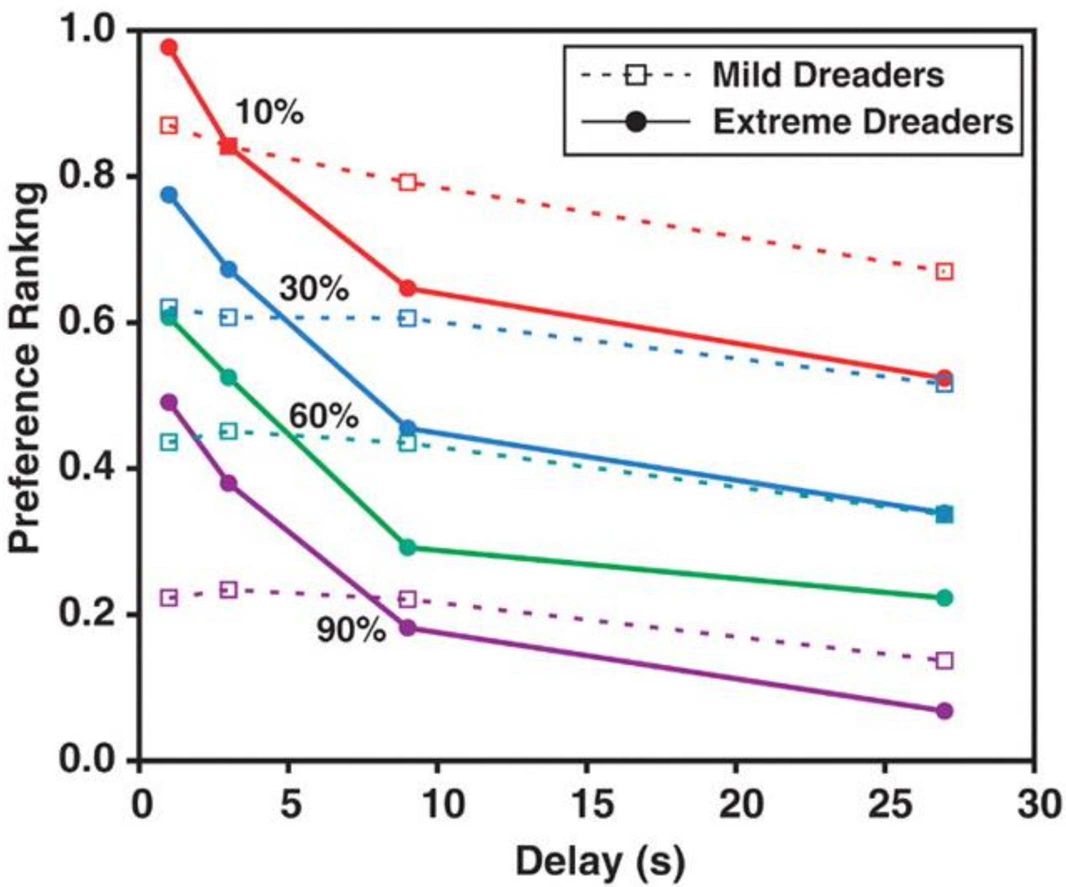

Berns GS, Chappelow J, Cekic M, Zink CF, Pagnoni G, Martin-Skurski ME (2006) Science 312:754

Discounting the past

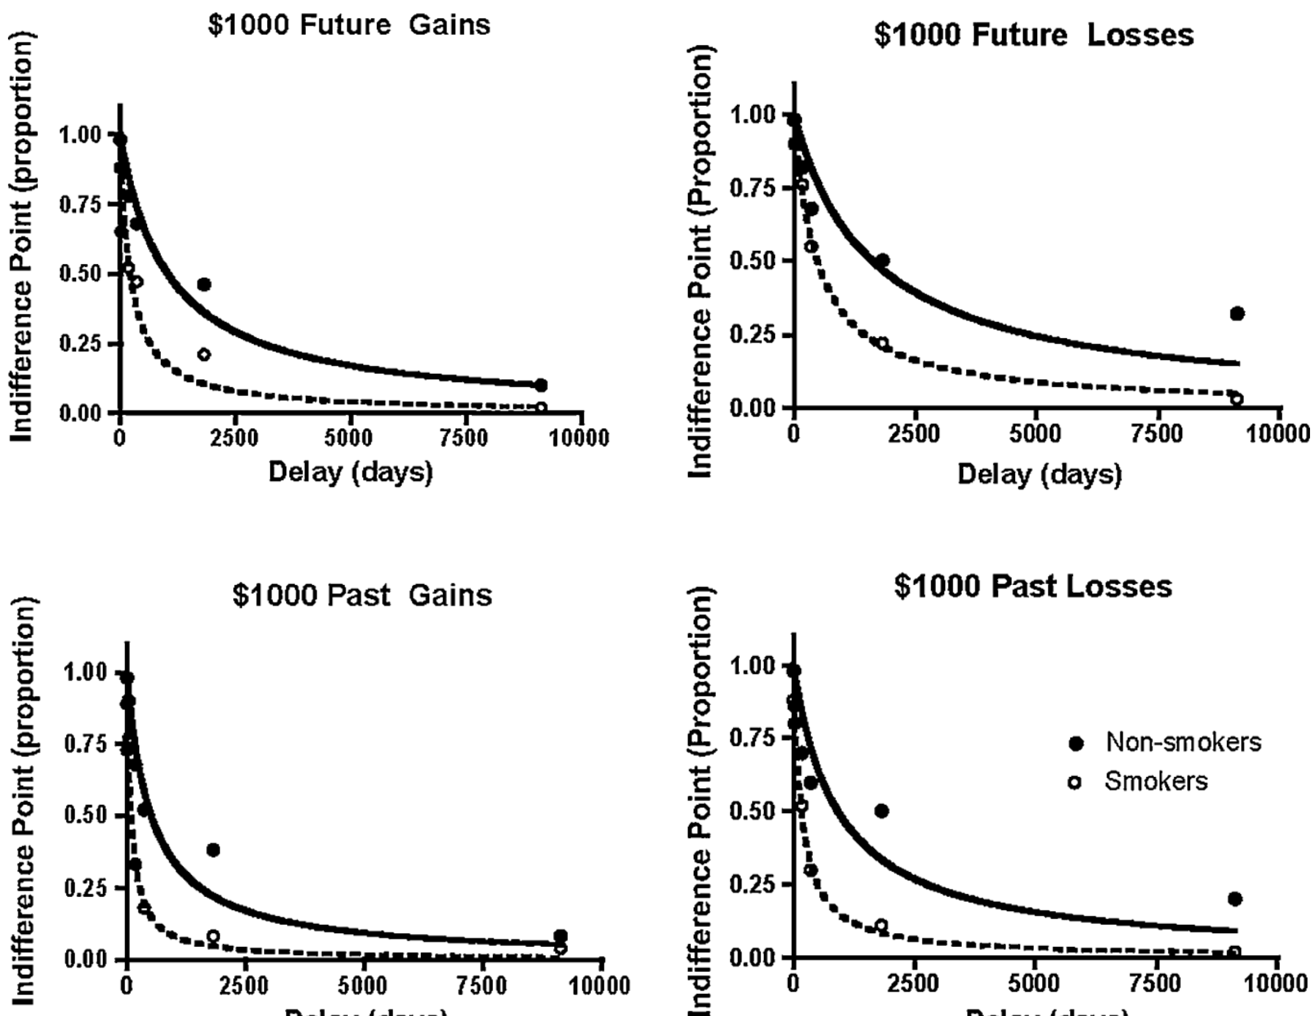

Bickel WK, Yi R, Kowal BP, Gatchalian KM (2008) Drug Alcohol Depend 96:256

Discounting the past

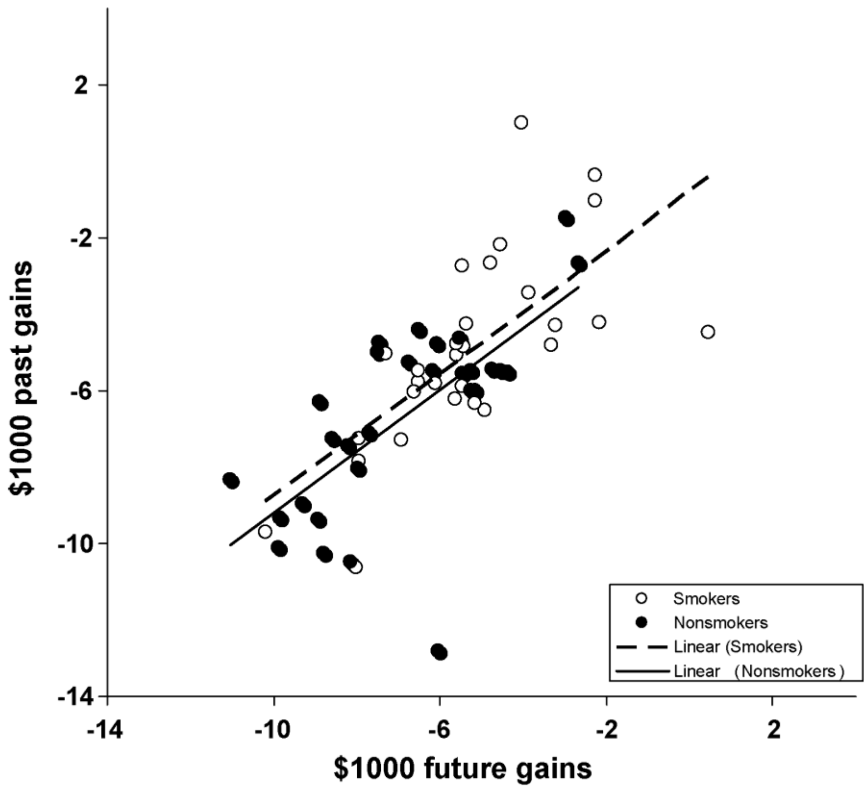

WHILE IT'S TECHNICALLY TRUE, I WISH SHE'D STOP PREFACING EVERY SENTENCE WITH THAT.

Bickel WK, Yi R, Kowal BP, Gatchalian KM (2008) Drug Alcohol Depend 96:256

Outline

1) Delay discounting is important

2) Designing a task & analyzing the data

3) What delay discounting measures

4) Modelling discounting

Hyperbolic discounting in temporal difference learning

TD models can predict behavioral and neural data.

But standard TD models can only accommodate exponential discounting.

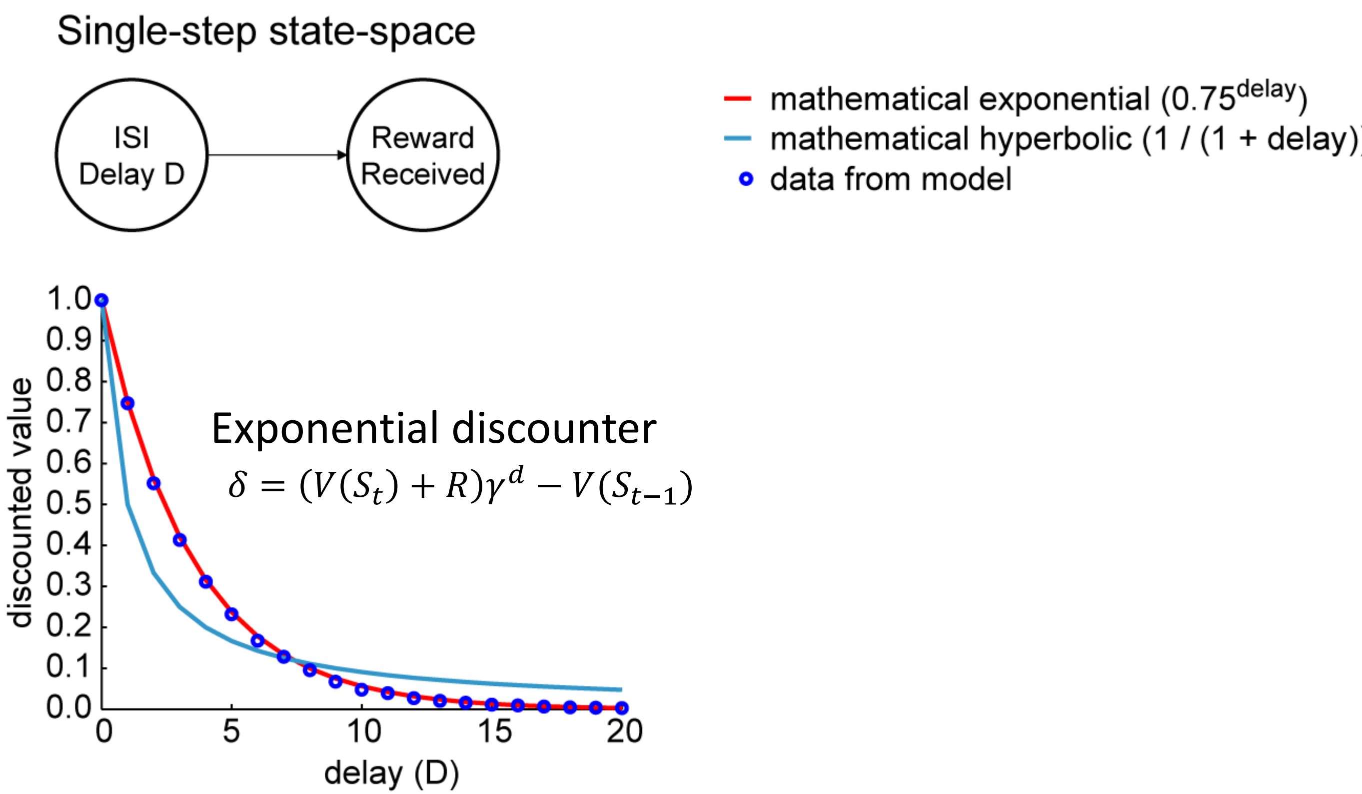

Kurth-Nelson and Redish (2009) PLoS ONE 4:e7362

Hyperbolic discounting in temporal difference learning

Across a multi-step state-space, standard TD cannot produce hyperbolic discounting.

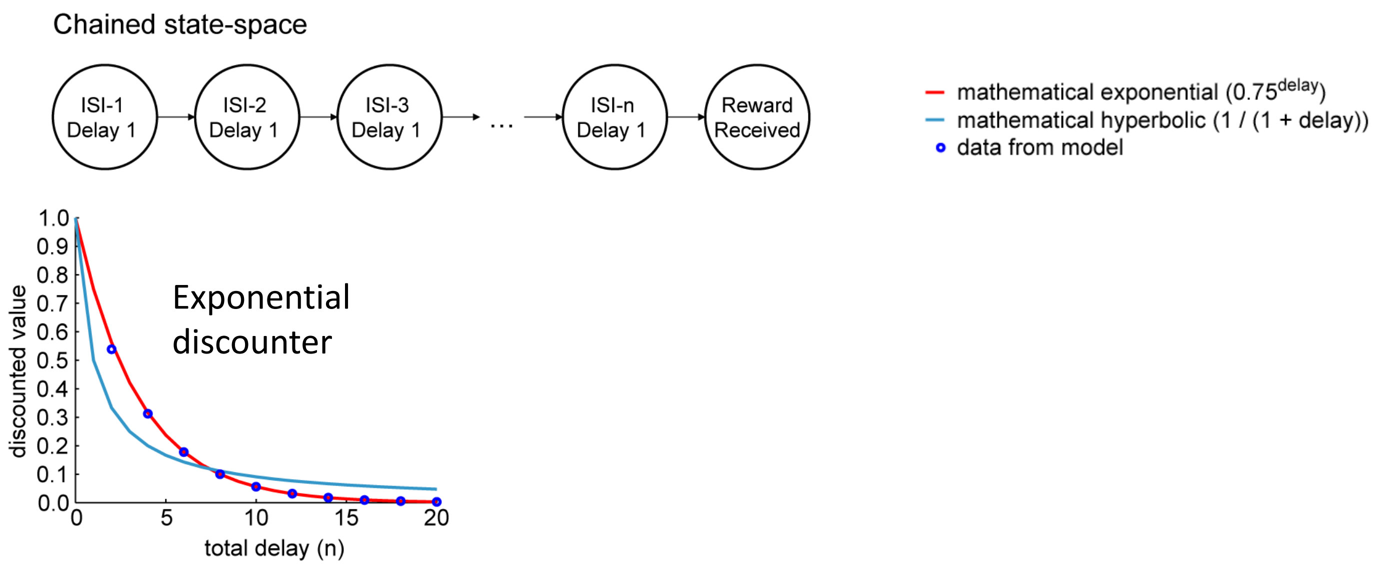

$$\gamma^{d_1} \cdot \gamma^{d_2} = \gamma^{d_1 + d_2}$$

$$\frac{1}{1+d_1} \cdot \frac{1}{1+d_2} \neq \frac{1}{1+(d_1+d_2)}$$

Kurth-Nelson and Redish (2009) PLoS ONE 4:e7362

μAgents model

Each μAgent learns its own estimate of the value function.

For action selection, value estimates are averaged across μAgents.

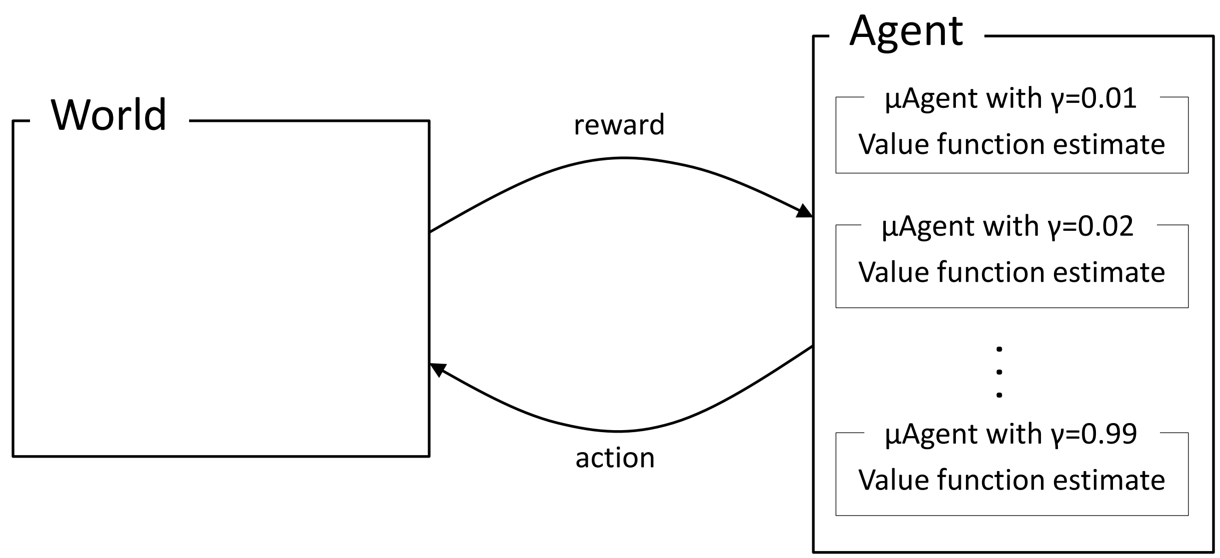

$$\delta_i = (R(S_t) + V_i(S_t))\gamma_i - V_i(S_{t-1})$$

$$V_i(S_{t-1}) \leftarrow V_i(S_{t-1}) + \delta_i \alpha$$

Update rules for μAgent *i*

Hyperbolic is the average of exponentials

μAgents have exponential discount rates (γ) uniformly spread from 0 to 1.

Average across μAgents approximates hyperbolic discounting.

$$\int_0^1 \gamma^d d\gamma = \frac{1}{1+d}$$

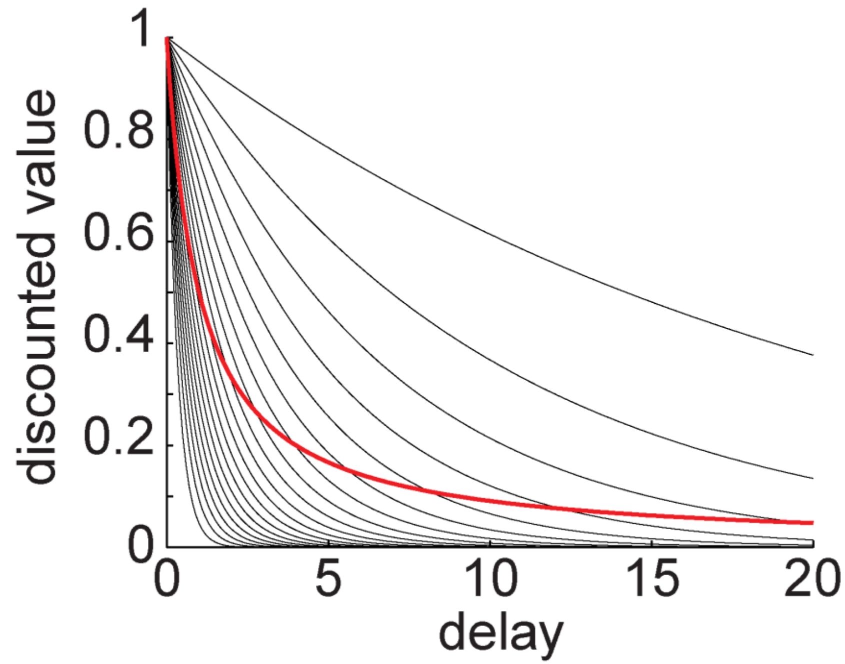

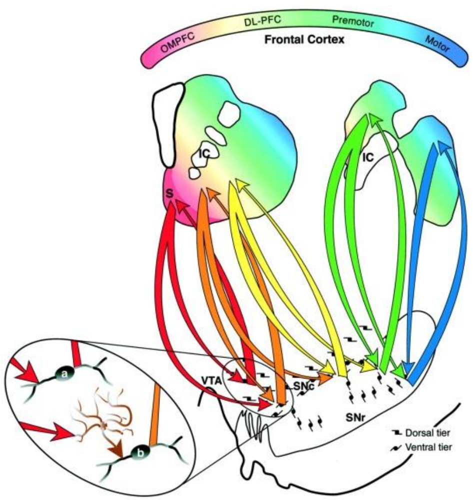

Haber et al (2000) J Neurosci 20:2369

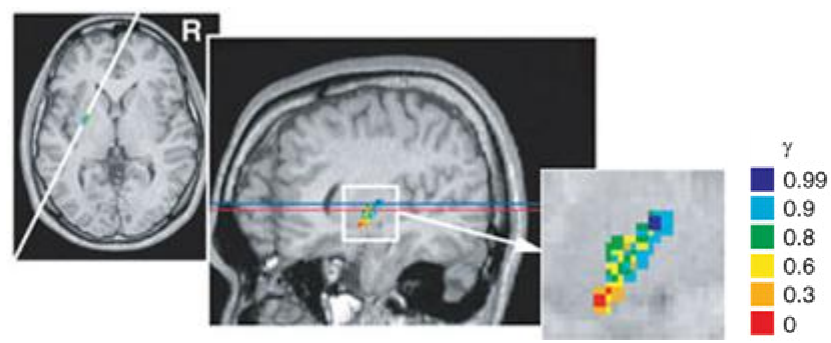

Tanaka et al (2004) Nat Neurosci 7:887

µAgents allows hyperbolic discounting across multiple transitions

Across a multi-step state-space, standard TD cannot produce hyperbolic discounting.

The µAgents model does produce hyperbolic discounting in this state-space.

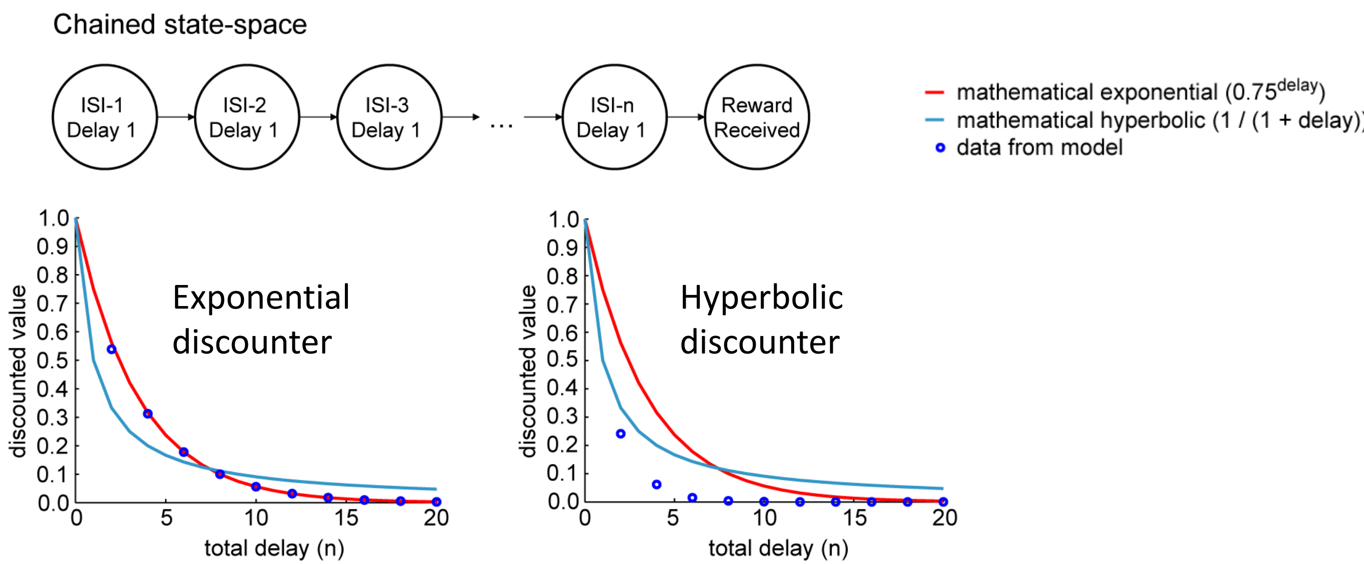

Kurth-Nelson and Redish (2009) PLoS ONE 4:e7362

Precommitment

In exponential discounting, adding the same delay to both outcomes doesn't change their relative values.

In hyperbolic discounting, preferences can reverse as you view the choice from a distance.

exponential

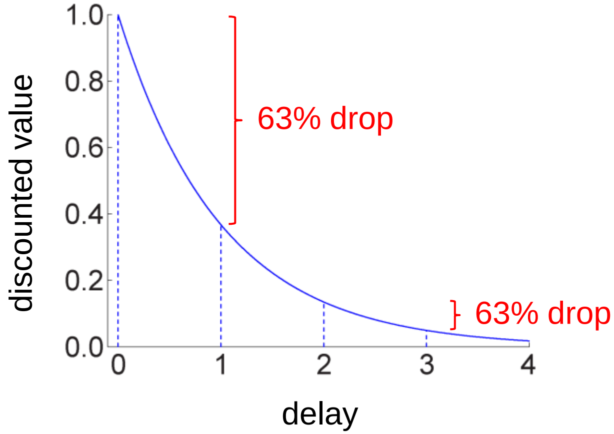

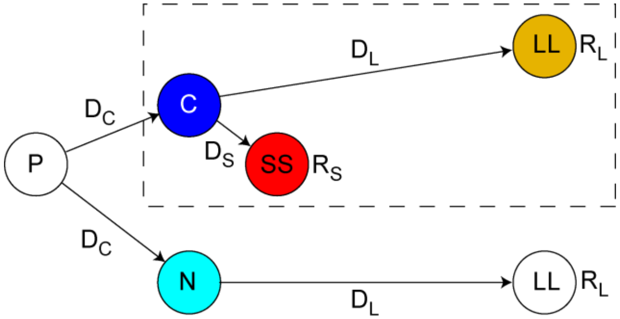

hyperbolic

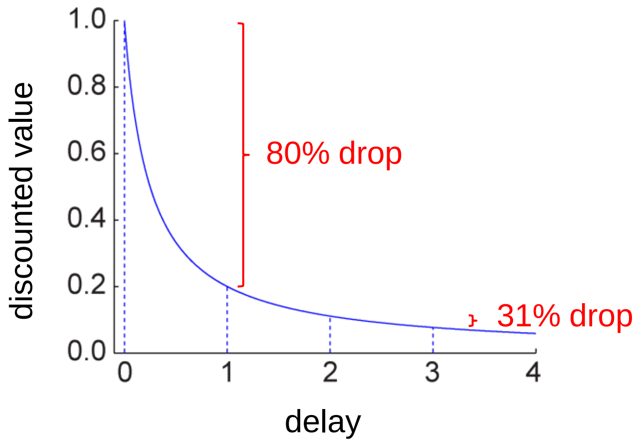

Kurth-Nelson and Redish (2010) Front Behav Neurosci 4:184

μAgents model precommits

At C, SS is preferred.
But at P, N is preferred.

The same average value can be encoded by different distributions. Distributions with more value carried by the more impulsive μAgents will discount faster.

Thus, average values can cross as discounting progresses.

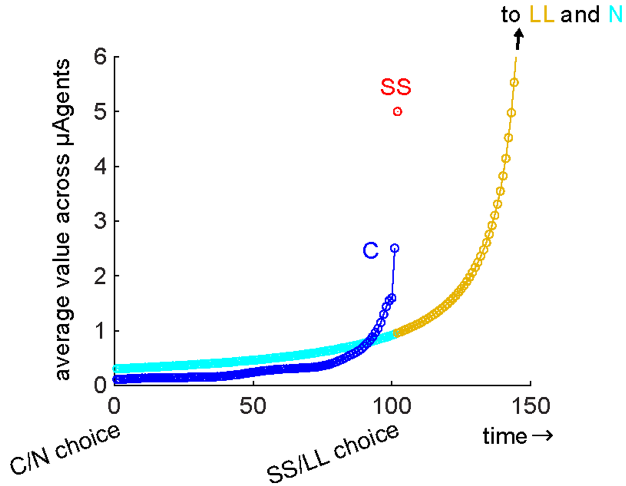

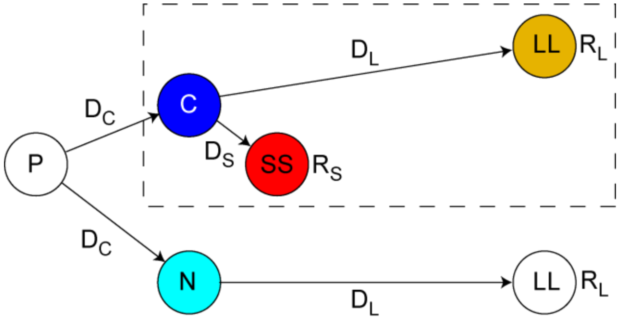

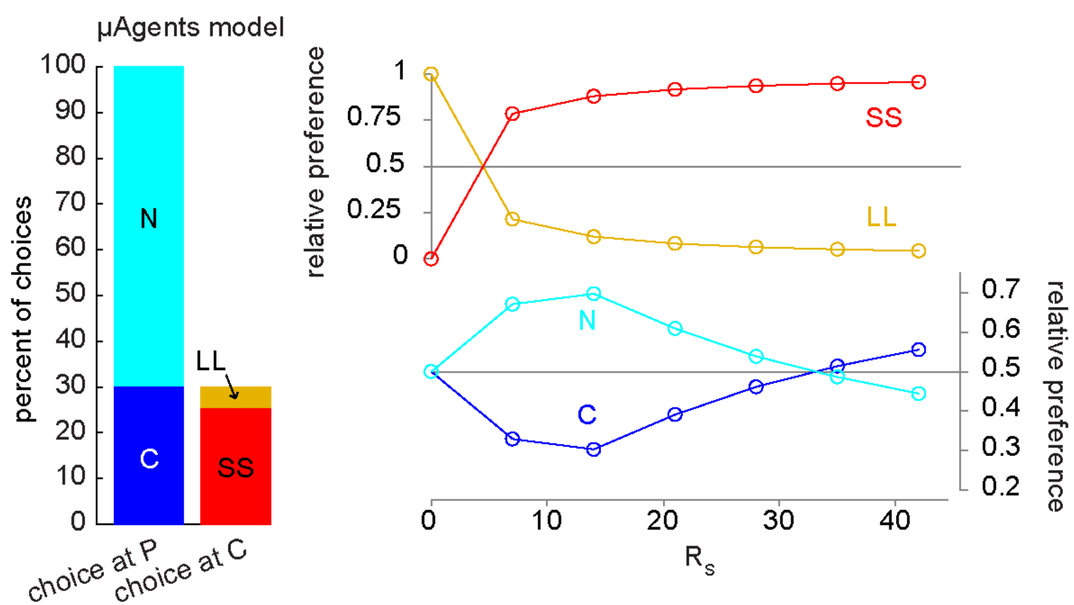

Kurth-Nelson and Redish (2010) Front Behav Neurosci 4:184

Cognitive search

At the choice point, rats project their hippocampal place representation ahead toward the feeders, suggesting a search process.

Ventral striatum also fires during this deliberation.

Johnson and Redish (2007) J Neurosci

Discounting arises from a search process

Three assumptions:

1. A reward that is easy to find is attributed more value
2. A reward that is closer in search space is easier to find
3. A reward that is closer in time is also closer in search space

Random diffusion from the origin.

The delay to an outcome is defined as its distance from the origin.

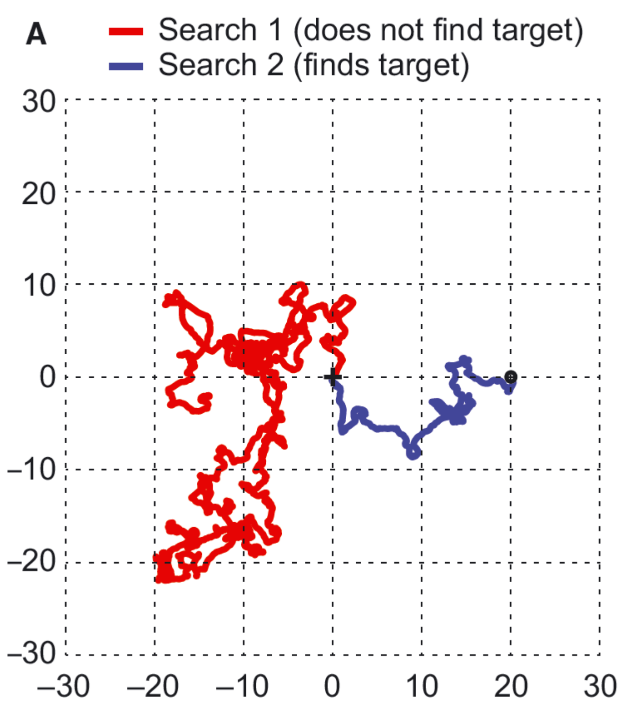

Kurth-Nelson Z, Bickel W, Redish AD (2012) EJN

Longer search time produces slower discounting

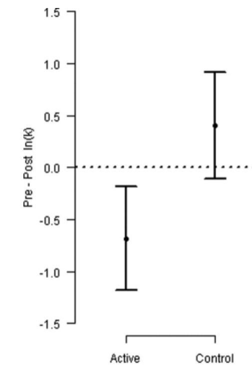

With more search time, it is more likely that the reward will be found, even if it is further away.

Search time is a stand-in for overall search resources:

- Working memory
- Cognitive load
- IQ

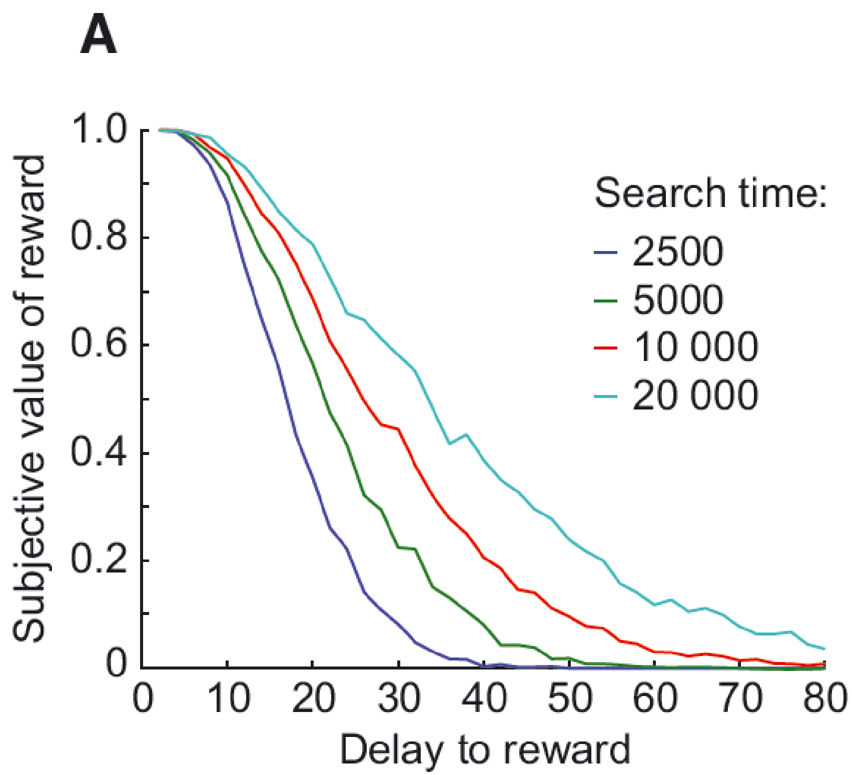

Kurth-Nelson Z, Bickel W, Redish AD (2012) EJN

Deeper basins produce slower discounting

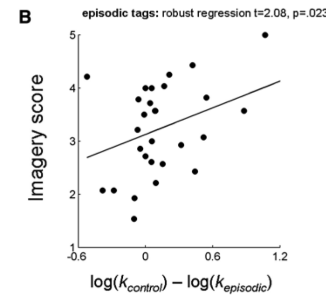

Deeper basins attract searches, making them more likely to find the outcome.

Deeper basins are hypothesized encode more episodic representations.

The form in which a state is represented is important to how decisions about it are made.

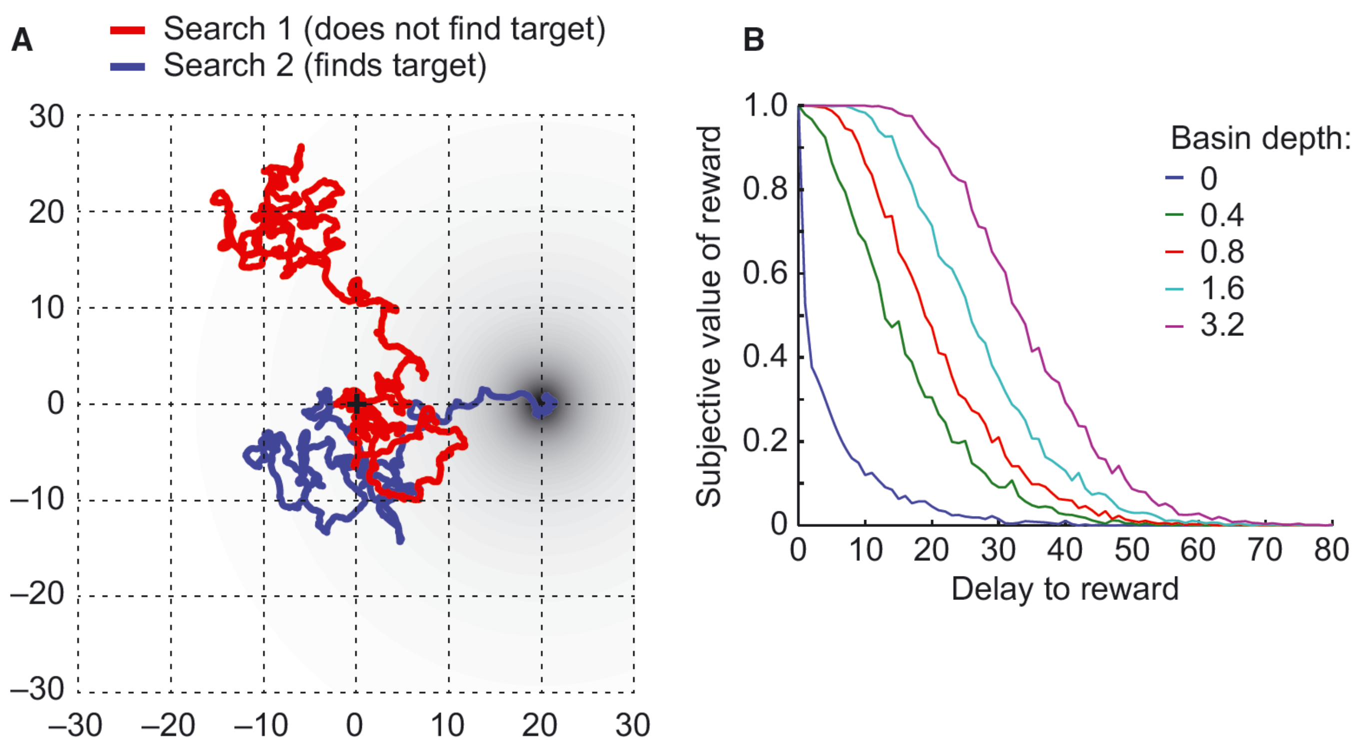

Kurth-Nelson Z, Bickel W, Redish AD (2012) EJN

More basins cause more impulsivity

If the representational space is dense with distractors, then it becomes harder to search through extra distance.

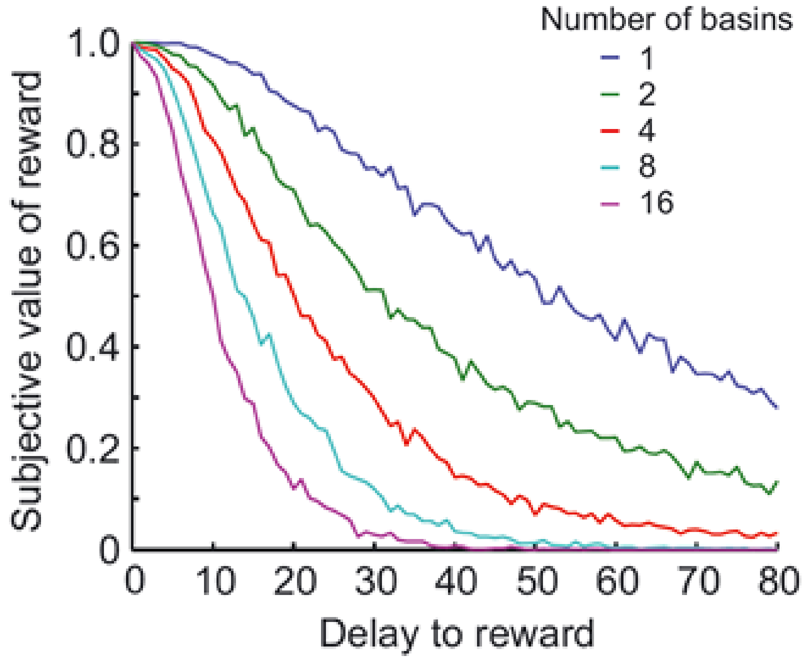

Kurth-Nelson Z, Bickel W, Redish AD (2012) EJN

Thanks!

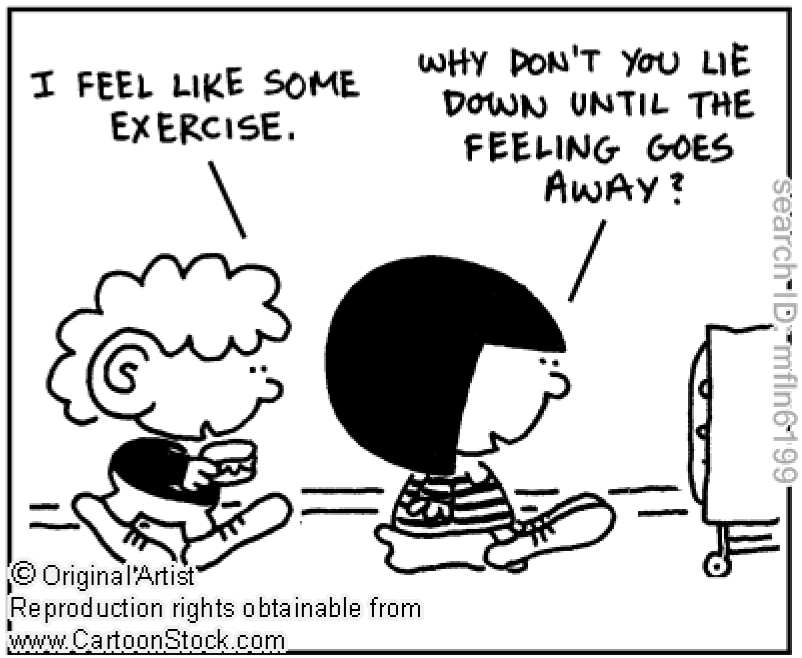
I FEEL LIKE SOME EXERCISE.
WHY DON'T YOU LIE DOWN UNTIL THE FEELING GOES AWAY?
search ID: mfln6199
© Original Artist
Reproduction rights obtainable from
www.CartoonStock.com