A very quick primer on risky choice and prospect theory

HW3 uses large-scale choice data collected for this paper…

Corrected 11 June 2021. See full text.

COGNITIVE SCIENCE

Using large-scale experiments and machine learning to discover theories of human decision-making

Joshua C. Peterson[1]*, David D. Bourgin[1]†, Mayank Agrawal[2,3], Daniel Reichman[4], Thomas L. Griffiths[1,2]

Predicting and understanding how people make decisions has been a long-standing goal in many fields, with quantitative models of human decision-making informing research in both the social sciences and engineering. We show how progress toward this goal can be accelerated by using large datasets to power machine-learning algorithms that are constrained to produce interpretable psychological theories. Conducting the largest experiment on risky choice to date and analyzing the results using gradient-based optimization of differentiable decision theories implemented through artificial neural networks, we were able to recapitulate historical discoveries, establish that there is room to improve on existing theories, and discover a new, more accurate model of human decision-making in a form that preserves the insights from centuries of research.

Understanding how people make decisions is a central problem in psychology and economics (*1*–*3*). Having quantitative models that can predict these decisions has become increasingly important as automated systems interact more closely with people (*4*, *5*). The search for such models goes back almost 300 years (*6*) but intensified in the latter half of the 20th century (*7*, *8*) as empirical findings revealed the limitations of the idea that people make decisions by maximizing expected utility (EU) (*9*–*11*). This led to the development of new models such as prospect theory (PT) (*8*, *12*). Recently, this theory-driven enterprise has been complemented by data-driven research using machine learning to predict human decisions (*13*–*19*). Although machine learning has the potential to accel-

narios in which decision-makers face a choice between two gambles, each of which has a set of outcomes that differ in their payoffs and probabilities (Fig. 1A). Researchers studying risky choice seek a theory, which we formalize as a function that maps from a pair of gambles, A and B, to the probability $P(A)$ that a decision-maker chooses gamble A over gamble B, that is consistent with human decisions for as many choice problems as possible. Discovering the best theory is a formidable challenge for two reasons. First, the space of choice problems is large. The value and probability of each outcome for each gamble define the dimensions of this space, meaning that describing a pair of gambles could potentially require dozens of dimensions. Second, the space of possible theories is even larger, with theories

This dataset includes >30 times the number of problems in the largest previous dataset (*27*) (Fig. 1B). We then used this dataset to evaluate differentiable decision theories that exploit the flexibility of deep neural networks but use psychologically meaningful constraints to pick out a smooth, searchable landscape of candidate theories with shared assumptions. Differentiable decision theories allow the intuitions of theorists to be combined with gradient-based optimization methods from machine learning to broadly search the space of theories in a way that yields interpretable scientific explanations.

More formally, we define a hierarchy over decision theories (Fig. 1C) reflecting the addition of an increasing number of constraints on the space of functions. These constraints express psychologically meaningful theoretical commitments. For example, one class of theories contains all functions in which the value that people assign to one gamble can be influenced by the contents of the other gamble. If theories in this class are more predictive than those that belong to the simpler classes contained within it (e.g., where the value of gambles are independent), then we know that these simpler theories should be eliminated. We enforce each constraint by modifying the architecture of artificial neural networks, resulting in differentiable decision theories. This theory-driven approach to defining constraints contrasts with generic methods for constraining neural networks, such as restricting their size or the ranges of their weights (*31*). After optimizing a differentiable theory to best fit human behavior, it will ideally have picked out the optimal theory in its class

A risky choice problem

Please select option A or B.

Earning a Bonus. At the end of the experiment, one reward will be selected at random from all the rewards you earned during the experiment. A fixed proportion (10%) of this value will be paid to you as your performance bonus for the task. If the sampled reward is negative, your bonus is set to $0.00.

A	B
16 with certainty	1 with probability 0.6 44 with probability 0.1 48 with probability 0.1 50 with probability 0.2

A **B**

In this trial, you chose **B** and gained **50**
Had you chosen **A**, you would have gained **16**

Dataset of 13,005 choice problems!

Data Loading

First, let's load the dataset and do some pre-processing. Don't worry too much about what the following code is doing, but the key dataframe you need afterwards is called `df` and shown below.

In [3]:

```
c13k = pd.read_csv('https://raw.githubusercontent.com/jcpeterson/choices13k/main/c13k_se
c13k_problems = pd.read_json("https://raw.githubusercontent.com/jcpeterson/choices13k/ma
df = c13k.join(c13k_problems, how="left")
df
```

Out[3]:

	Problem	Feedback	n	Block	Ha	pHa	La	Hb	pHb	Lb	LotShapeB	LotNumB	Amb	Corr	bRate	bRate_std
0	1	True	15	2	26	0.95	-1	23	0.05	21	0	1	False	0	0.626667	0.384460
1	2	True	15	4	14	0.60	-18	8	0.25	-5	0	1	True	-1	0.493333	0.413118
2	3	True	17	4	2	0.50	0	1	1.00	1	0	1	False	0	0.611765	0.432843
3	4	True	18	3	37	0.05	8	87	0.25	-31	1	2	False	0	0.222222	0.387383
4	5	False	15	1	26	1.00	26	45	0.75	-36	2	5	False	0	0.586667	0.450185
...	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...	...
14563	13002	True	15	3	30	1.00	30	42	0.80	0	0	1	True	0	0.367619	0.302731
14564	13003	True	15	5	70	0.50	-42	18	0.80	7	0	1	False	0	0.760000	0.364104
14565	13004	True	15	5	8	0.40	-17	31	0.40	-34	1	6	False	0	0.666667	0.367747
14566	13005	True	15	2	89	0.50	-49	45	0.50	-12	0	1	False	0	0.386667	0.381476

A couple more problems from the dataset…

Problem 1, Feedback = True

n = 15, bRate = 0.6267, std: 0.3845

Gamble A

	Payout	Probability
0	26.0	0.95
1	-1.0	0.05

Gamble B

	Payout	Probability
0	21.0	0.95
1	23.0	0.05

Problem 4, Feedback = True

n = 18, bRate = 0.2222, std: 0.3874

Gamble A

	Payout	Probability
1	8.0	0.95
0	37.0	0.05

Gamble B

	Payout	Probability
0	-31.0	0.750
1	86.5	0.125
2	87.5	0.125

Let's take a computational view of the mind… How do people make these choices?

Can we explain these choices as computation? If so, what is the right algorithm?

A classic economic theory:
People choose option that maximizes expected value/utility

Expected value

x : vector of payouts

p : vector of probabilities

$$\text{EV} = \sum_{i=1}^{N} x_i p_i$$

```
Problem 4, Feedback = True
n = 18, bRate = 0.2222, std: 0.3874
```

Gamble A

	Payout	Probability
1	8.0	0.95
0	37.0	0.05

Gamble B

	Payout	Probability
0	-31.0	0.750
1	86.5	0.125
2	87.5	0.125

$$\text{EV}_A = 8 \times .95 + 37 \times .05 = 9.45$$

$$\text{EV}_B = -31 \times .75 + 86.5 \times .125 + 87.5 \times .125 = -1.5$$

Thus, people would choose Option A!

However, expected value/utility theory is a pretty poor fit to real human decisions… which brings us to prospect theory

ECONOMETRICA
VOLUME 47 MARCH, 1979 NUMBER 2

PROSPECT THEORY: AN ANALYSIS OF DECISION UNDER RISK

BY DANIEL KAHNEMAN AND AMOS TVERSKY[1]

This paper presents a critique of expected utility theory as a descriptive model of decision making under risk, and develops an alternative model, called prospect theory. Choices among risky prospects exhibit several pervasive effects that are inconsistent with the basic tenets of utility theory. In particular, people underweight outcomes that are merely probable in comparison with outcomes that are obtained with certainty. This tendency, called the certainty effect, contributes to risk aversion in choices involving sure gains and to risk seeking in choices involving sure losses. In addition, people generally discard components that are shared by all prospects under consideration. This tendency, called the isolation effect, leads to inconsistent preferences when the same choice is presented in different forms. An alternative theory of choice is developed, in which value is assigned to gains and losses rather than to final assets and in which probabilities are replaced by decision weights. The value function is normally concave for gains, commonly convex for losses, and is generally steeper for losses than for gains. Decision weights are generally lower than the corresponding probabilities, except in the range of low probabilities. Overweighting of low probabilities may contribute to the attractiveness of both insurance and gambling.

1. INTRODUCTION

EXPECTED UTILITY THEORY has dominated the analysis of decision making under risk. It has been generally accepted as a normative model of rational choice [**24**], and widely applied as a descriptive model of economic behavior, e.g. [**15**, **4**]. Thus, it is assumed that all reasonable people would wish to obey the axioms of the theory [**47**, **36**], and that most people actually do, most of the time.

The present paper describes several classes of choice problems in which preferences systematically violate the axioms of expected utility theory. In the light of these observations we argue that utility theory, as it is commonly interpreted and applied, is not an adequate descriptive model and we propose an alternative account of choice under risk.

2. CRITIQUE

Decision making under risk can be viewed as a choice between prospects or gambles. A prospect $(x_1, p_1; \ldots; x_n, p_n)$ is a contract that yields outcome x_i with probability p_i, where $p_1+p_2+\ldots+p_n=1$. To simplify notation, we omit null outcomes and use (x, p) to denote the prospect $(x, p; 0, 1-p)$ that yields x with probability p and 0 with probability $1-p$. The (riskless) prospect that yields x with certainty is denoted by (x). The present discussion is restricted to prospects with so-called objective or standard probabilities.

The application of expected utility theory to choices between prospects is based on the following three tenets.

(i) Expectation: $U(x_1, p_1; \ldots; x_n, p_n) = p_1u(x_1)+\ldots+p_nu(x_n)$.

[1] This work was supported in part by grants from the Harry F. Guggenheim Foundation and from

Prospect theory

u : utility function
x : vector of payouts
p : vector of probabilities

$$\text{Value of a gamble} = \sum_{i=1}^{N} u(x_i)p_i$$

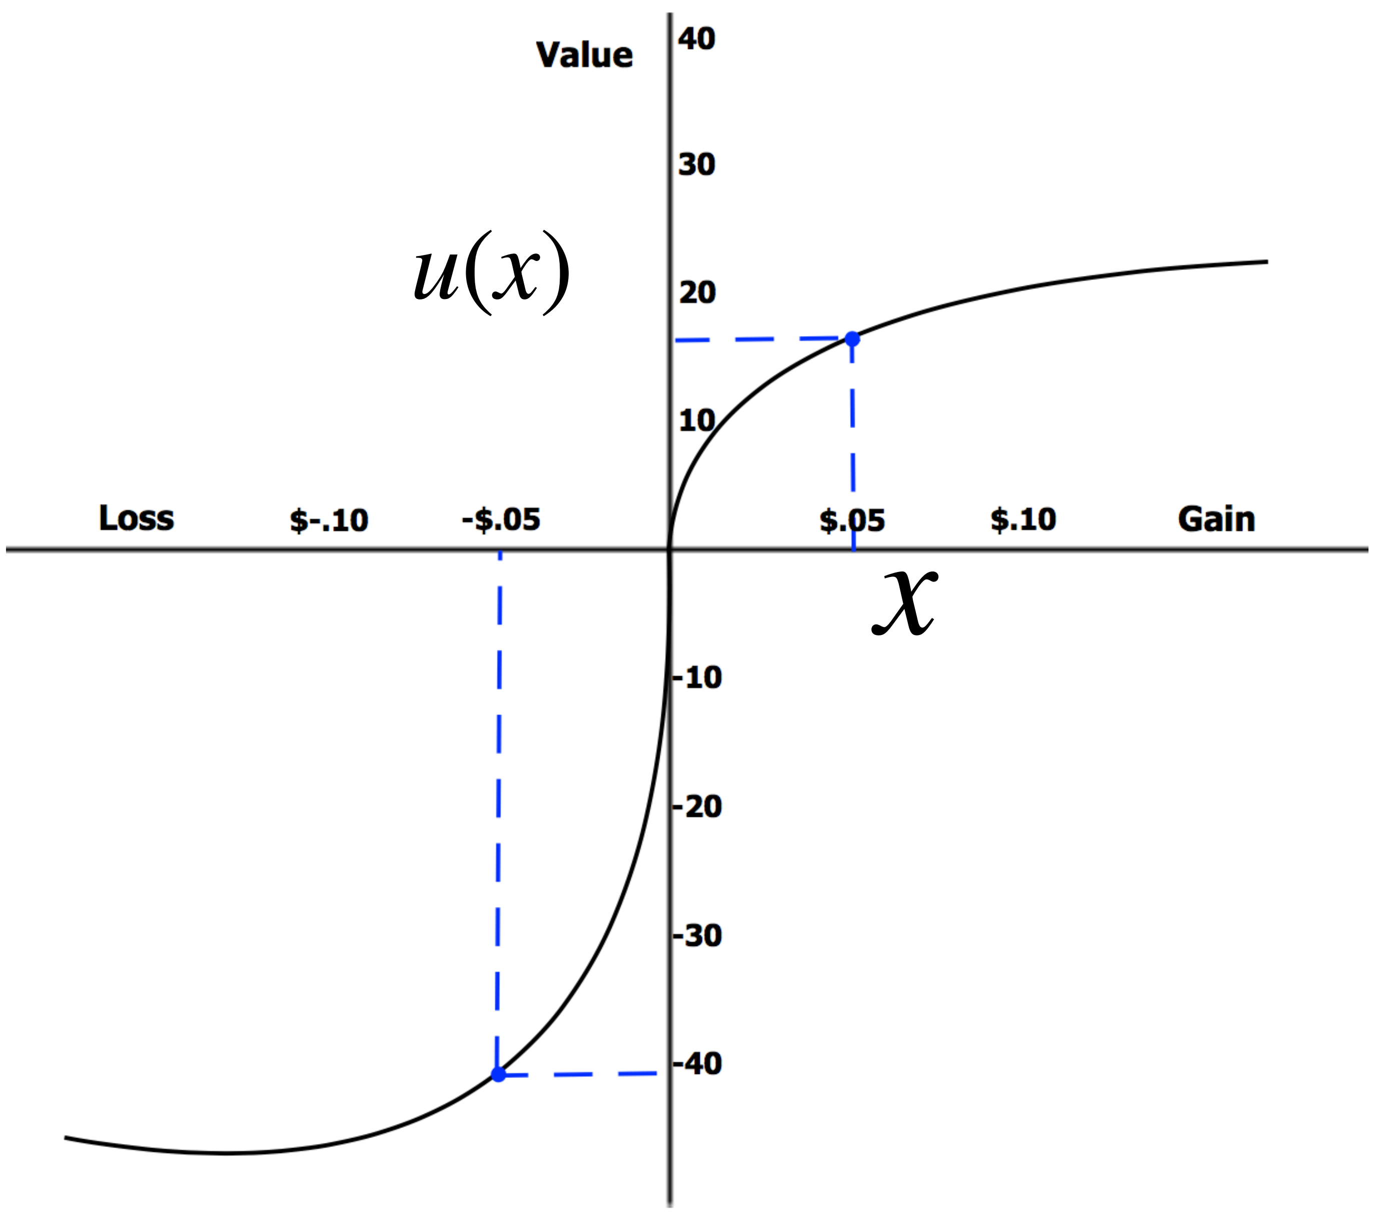

Point 1: People perceive gambles in terms of gains and losses, not their total wealth

Prospect theory: Loss aversion

u : utility function
x : vector of payouts
p : vector of probabilities

$$\text{Value} = \sum_{i=1}^{N} u(x_i)p_i$$

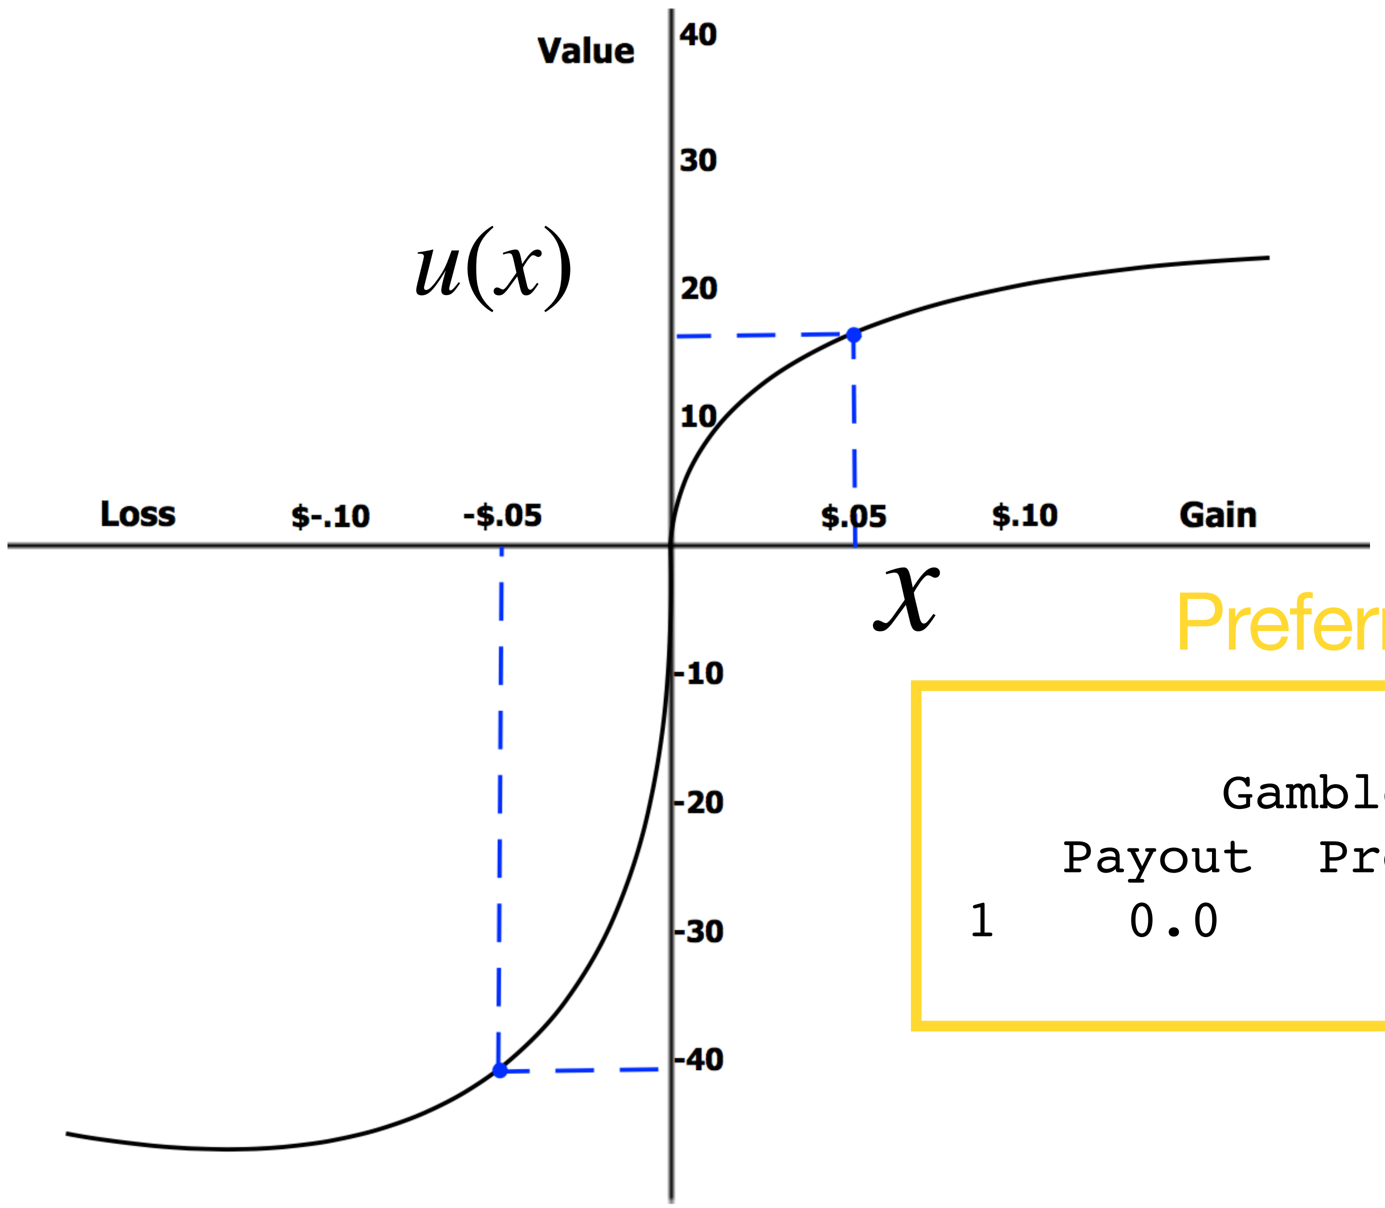

Point 2: People are **LOSS AVERSE**, preferring a sure thing to a risky bet with the same expected payoff

(Notice that the loss curve is steeper than the gain curve)

Preferred

Gamble A

	Payout	Probability
1	0.0	1.0

Gamble B

	Payout	Probability
0	-5.0	0.5
1	5.0	0.5

Prospect theory: Risk aversion for gains

u : utility function
x : vector of payouts
p : vector of probabilities

$$\text{Value} = \sum_{i=1}^{N} u(x_i)p_i$$

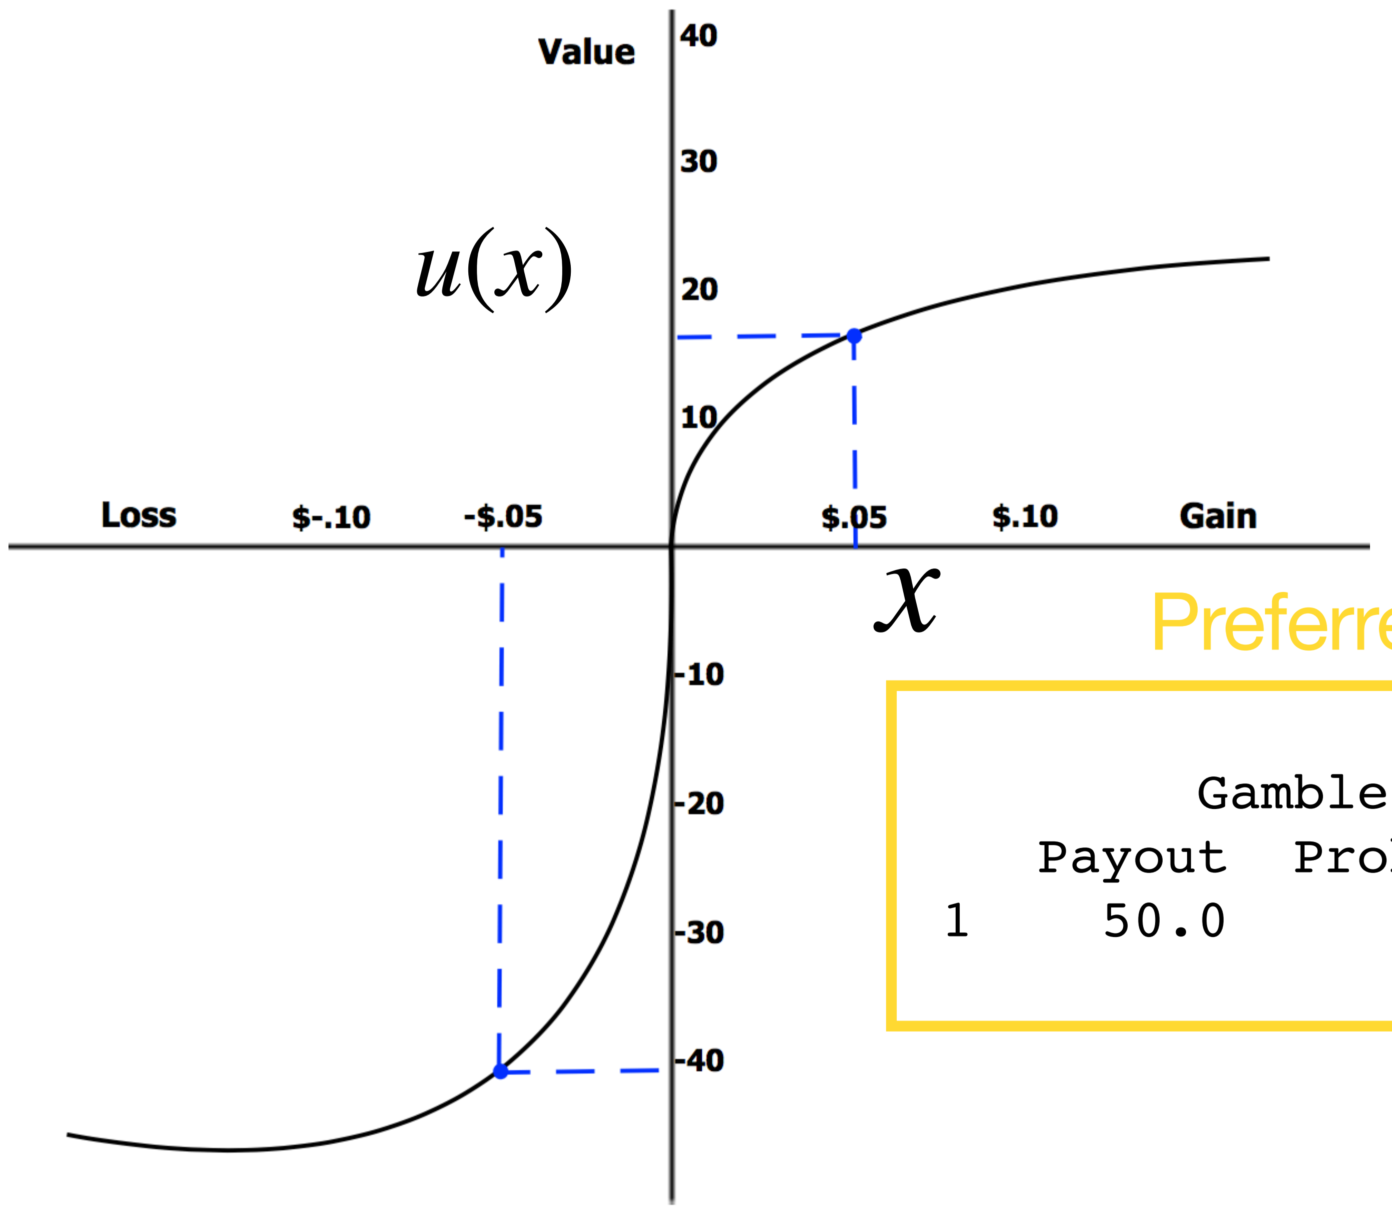

Point 3: Diminishing returns… $1B is not 1000x better than $1M

Leads to **risk aversion** for gambles with potential gains

Preferred

Gamble A

	Payout	Probability
1	50.0	1.0

Gamble B

	Payout	Probability
0	100.0	0.5
1	0.0	0.5

Prospect theory: Risk seeking for losses

u : utility function
x : vector of payouts
p : vector of probabilities

$$\text{Value} = \sum_{i=1}^{N} u(x_i)p_i$$

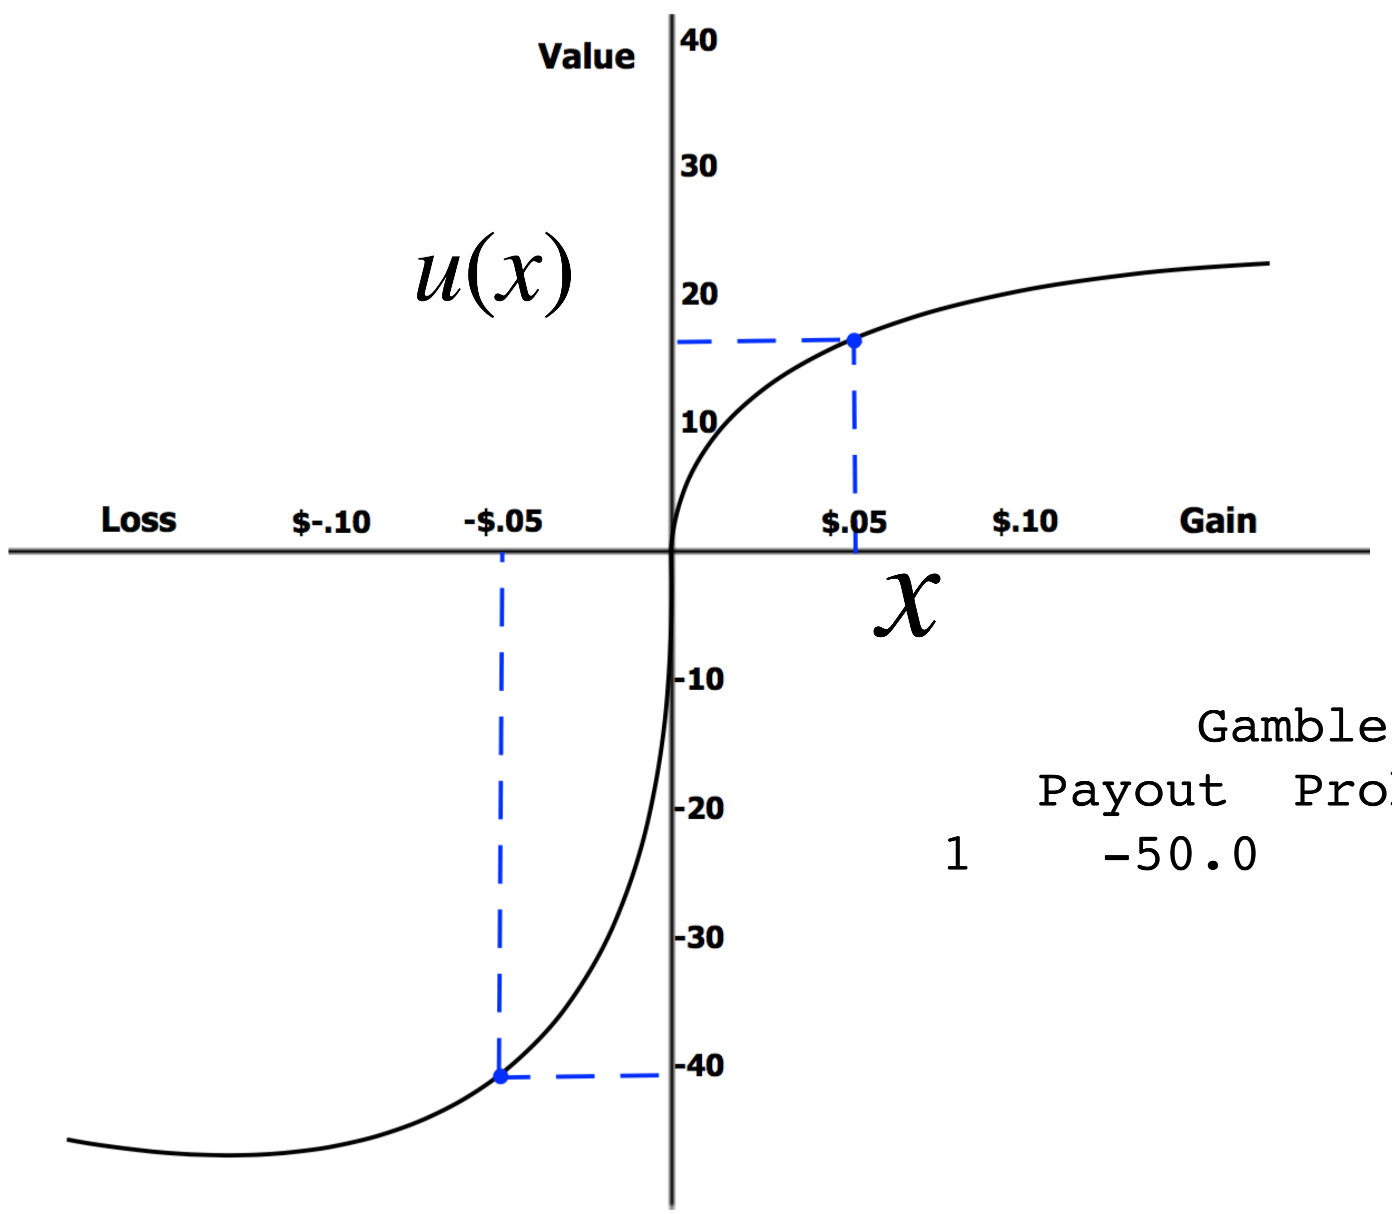

Point 4:
People are **risk seeking** for gambles with potential losses

Gamble A

	Payout	Probability
1	-50.0	1.0

Preferred

Gamble B

	Payout	Probability
0	-100.0	0.5
1	0.0	0.5