

Decision-making under uncertainty: An Introduction to Neuroeconomics and Neurofinance

Lecturers:

Prof. Kerstin Preuschoff

Dr. Marco Lehmann

TA:

Jean-Paul Nohra

15.09.2026




What is a decision — and why is it hard to study?

Decisions science understands best: visual decisions (what do you see?), whether to reach left or right, which of two chocolate bars to buy.

Have been studied intensely for decades; small enough to dissect — two or three options, a handful of properties, and usually plenty of time to decide.

Real decisions rarely look like that. They come with more options than we can list, more factors than we can weigh, and often not nearly enough time or information to sort through it all.

That is what this course is about.



Three more decisions

💰 **The investment** → a friend calls: a new company, not public yet, no track record, but “this could be huge.” They want you to put in real money. Do you say yes?

🌱 **The sustainable choice** → someone consistently chooses sustainable options often paying extra — even when no one's watching. What's going on there?

🩺 **The health screening** → you're offered a test for a disease you currently can't do anything about. Do you take it?

Three decisions, three domains — two of these we're coming straight back to today. But there's a fourth, much closer to home, that some of you are making right now.






A decision you're making right now

Most of you haven't committed to this course yet — you have a few more weeks weeks to decide. What are you actually weighing?

Workload. What you'll walk away knowing. Whether it matters five years from now. Whether you'll enjoy showing up. The grade you're likely to get. The number of homework assignments. The stress of a final exam in January. Who else is taking the course.


 **Moodle — two questions:** → a number from from 0–10 for how likely you are to take this course, and the main factors driving your decision. decision.

On a scale of 0 to 100%, how likely are you to enrol in this course?



A decision you're making right now

Which factors do you take into account when choosing a course? Select the 3 that apply most



How every week is structured

Section	What it contains
Description + Learning Objectives	A short abstract of the week's topic, and what you should be able to identify, name, and critically argue by the end of it.
In-Class Activity (90 minutes)	What happens in the room. This lecture is Week 1's in-class activity.
Self-Study (every week)	Videos, articles, and webpages. Every block ends with a few questions you should be able to answer if you did it properly. Expect roughly 2–3 hours.
Homework Assignments	The four pass/fail assignments (details on the next slide) — not every week, only Weeks 3, 5, 7, and 9.
Optional: Feed Your Curiosity	Extra material, not examined, not provided every week — for anyone who wants more.

Assessment at a glance

Element	Detail
Assignments	Four pass/fail homework assignments, handed out in Weeks 3, 5, 7, and 9.
Deadlines	Full pass: submitted by the following Monday at noon. Half pass: submitted between Monday noon and Tuesday's class at 10:00. No credit at all after Tuesday 10:00.
Grading	Passing 3 of the 4 assignments = 30% of the grade. Attendance = 10% (11 classes needed for full credit). Final exam = 60%.
Role of this course	Foundation for the spring follow-up course, going deeper into recent research and applied case studies in neuroeconomics and neurofinance.

General remarks

💰 What this course is not

→ a standard course covering only well established facts, models and methods whose contents has not really changed in 20 years.

🌱 What this course is

→ A recently restructured course, constantly evolving, covering topics we sometimes don't know much about.

→ An interdisciplinary exploration, an attempt to bring together insights from different fields that have studied similar questions independently of one another.

→ An exploration of current questions and emerging topics.



Questions about the course?

Disciplinary perspectives

Economics

Neuroscience

Computational
Neuroscience

Neurofinance
Neuroeconomics

Psychology

Machine
learning

Clinical
neuroscience

Psychiatry

Economics

Neuroscience

Computational
Neuroscience

Neurofinance
Neuroeconomics

Psychology

Machine
learning

Psychologist:

“How is the investor’s decision affected by his emotions?”



Neuroscientist:

“What are the underlying neural circuits and computational mechanisms?”



Biologist:

“How does this contribute to the survival of the organism and the species?”



Economist:

“What is the simplest theory that can explain the choices I observe?”



Financial economist:

“How do individual investments lead to the formation of market prices?”

Psychologist:

“How is the investor’s decision affected by his emotions?”



Neuroscientist:

“What are the underlying neural circuits and computational mechanisms?”



Biologist:

“How does this contribute to the survival of the organism and the species?”



Economist:

“What is the simplest theory that can explain the choices I observe?”



Financial economist:

“How do individual investments lead to the formation of market prices?”


The decisions each field has classically studied

Field	Classic domain of study
Economics	Choices between goods and money — easily observable; clean, quantifiable trade-offs.
Psychology	Stimulus-response behavior, reactions to stimuli — controllable, repeatable, easy to manipulate in the lab (the signal-detection/psychophysics tradition).
Neuroscience	Visual and motor decisions — the underlying circuitry is well-mapped and easiest to isolate (e.g. random-dot motion discrimination tasks; Gold & Shadlen, 2007).

What we're about to do — walk a real investment decision and a real sustainability decision through all three lenses — is precisely the harder case: not a lab-reduced choice, but the messy, real kind from the preamble.


Psychology

- sees the individual as a sentient being guided primarily by feelings, emotions, and physiological states.
- our choices seek to maximise our happiness, or fulfilment
- provides an understanding of how we respond to our external world and how we learn about it



Pavlov, 1901

Decision making in psychology: classical and instrumental conditioning



Observing and inferring from choice



>



<



<



$$a \succ b, b \succ c \Rightarrow a \succ c$$

Decision making in economics and finance

- Is about maximising wealth (economics) and managing risk/uncertainty (finance)
- Decision making: levels of analysis
 1. Normative
What constitutes “rational behavior”? i.e., What does it mean to maximize self-interest?
 2. Descriptive
What kinds of judgments and decisions do people actually make in practice?
 3. Prescriptive
How can people go about making higher quality decisions?

Decision making in biology & neuroscience:

Two levels of explanation

Proximal

Neural circuits and computational algorithm,
i.e., What neural circuits are active during a particular behaviour? How do neural circuits communicate? How do they learn?


Ultimate

Maximize fitness,
i.e., maximize the number of genes that occur in future generations



Modern Neurobiology

- provides a wealth of constraints and restrictions that shape neuroeconomic theory
 - the basics of information coding in the nervous system
 - constraints on the cost of this encoding
 - fairly complete and (surprisingly sparse) map of what connects to what in the human brain
 - explain biophysical/molecular constraints on the form and costs of learning



What does each field study, in general?

Economics

Psychology

Neuroscience

What they study	How scarce resources get allocated — production, exchange, distribution of goods, services, and wealth across individuals, firms, and societies.	Mind and behavior broadly — cognition, emotion, perception, development, personality, social behavior.	The nervous system — from molecules and cells through circuits to whole-brain systems, and how these produce behavior.
Core premise	Agents have preferences and face constraints; resources are scarce relative to wants; a rational, optimizing agent is the usual starting baseline.	Behavior and mental life follow lawful regularities discoverable empirically, even though much of what drives behavior isn't consciously accessible.	Behavior and mental processes are implemented by physical processes in the brain; the biological substrate offers a mechanistic, causal layer of explanation.
Goal	Explain and predict choices, how resources are allocated and prices form; often prescriptive (policy, welfare).	Describe, explain, and predict thought, emotion, and behavior; in applied psychology, also to change it.	Map mental functions onto neural structures and computations; identify causal mechanisms, often with clinical application in mind.
Approach	Formal mathematical/axiomatic models of choice; predictions tested against market data, surveys, experiments; emphasis on equilibrium, incentives, trade-offs.	Controlled lab and field experiments, observational studies, psychometrics, case studies; theories validated against behavioral data.	Direct measurement (EEG, fMRI, single-unit recording) and causal manipulation (lesions, stimulation, pharmacology) of neural activity, increasingly tied to computational models.

Same three lenses, aimed at a decision

	Economics / Finance	Psychology	Neuroscience
Core question	Given constraints and information, what choice maximizes expected value/utility?	What cognitive and emotional processes produce this choice, and where does it depart from rational calculation?	What neural computations and circuits implement the choice?
What counts as “value”	Money, wealth, expected returns — one comparable currency.	Subjective well-being, satisfaction, felt importance — often domain-specific.	Reward and prediction-error signals — a candidate common neural currency.
Typical methods	Formal modeling, revealed preference from market/choice data, econometrics, natural experiments.	Controlled experiments, self-report/psychometric scales, behavioral observation.	fMRI, EEG, single-unit recording, lesion studies, pharmacological manipulation.
What counts as an explanation	The choice fits — or measurably deviates from — an optimizing model.	A named heuristic or bias accounts for the pattern (loss aversion, framing, anchoring...).	A specific circuit's activity correlates with, or causally drives, the choice.
Where evidence comes from	Markets, natural experiments, incentivized lab choices.	Lab experiments, surveys, observed behavior.	Brain imaging/recording during decisions, animal models.

Two decisions, three lenses each

Lens	Investment		Sustainable choice
			
Economic	Real expected future value (would this actually pay off) vs. a purely psychological return.		Real expected future value (durability, reputation, resale) vs. a purely psychological return.
Psychological	Loss aversion and framing; the friend's recommendation as a social cue that shifts the decision independent of its informational content.		Values/identity-driven vs. status/signaling.
Neuroscience	Reward and valuation circuitry, measurable neural signature, named later in Week 4.		Does “ feeling good” about a sustainable choice show up anywhere measurable?

A unifying frame: Marr's levels of analysis (over-simplified, use with caution)

Marr's level	Question it asks	Fits most naturally	For our investment decision
Computational	What problem is being solved, and why this solution?	Economics	Given expected returns and risks, what allocation of savings maximizes expected utility?
Algorithmic	What representations and steps actually produce the choice?	Psychology	Loss aversion, framing, heuristic shortcuts — the actual procedure people use.
Implementation	What physical hardware carries out the algorithm?	Neuroscience	Dopaminergic reward-prediction-error signals, vmPFC valuation circuitry, etc.

Marr, D. (1982). Vision. — originally proposed for a single system (vision) computing one function, borrowed here as a metaphor for granularity across three research communities.


What is uncertainty?

Naming the fog: what is uncertainty?

	What you know	Investment example	Sustainability example
Risk	The possible outcomes and their odds — real historical data.	Comparable past ventures give rough base rates.	Well-studied, measurable environmental costs of an established product.
Ambiguity	The outcomes, but not reliable odds — no track record for this case (Elsberg, 1961).	A brand-new company with no history of its own.	A genuinely novel material or technology with no track record.
Deep / Knightian uncertainty	You can't even fully specify the possible outcomes (Knight, 1921).	A macro shock nobody saw coming.	Climate tipping points — long-run consequences that resist even probabilistic modeling.

Intuitive vocabulary only — the formal treatment is Weeks 7–8.

There are levels



1. You know exactly what will happen.
2. You don't know the outcome, but you know the odds cold.
3. You don't know the odds, but you could find out with enough work or data.
4. You don't know the odds, and only some of that gap is closeable.
5. You don't know the odds, and — as far as anyone can tell — you never fully will.

Maps onto Lo & Mueller's five-level uncertainty taxonomy (2010).

Types of uncertainty (while deciding for or against this course)

Risk (known probabilities):

- "Based on past grades, there's probably a 15% chance I'll get below 5"

Ambiguity (unknown probabilities):

- "How strict will the grading be? Could be lenient, could be harsh"
- "Will the online content be engaging or boring? No way to know"

Model Uncertainty (don't know the structure):

- "What if interdisciplinary approaches don't actually help my career?"
- "Maybe the whole field of neuroeconomics will become irrelevant"

Strategic Uncertainty (depends on others):

- "If all the smart students drop out, discussions will be worse"
- "If everyone else works harder than me, I'll look bad in participation grades"

Temporal/Dynamic Uncertainty:

- "My interests might change as I learn more"
- "The job market could shift dramatically by graduation"

Personal/Preference Uncertainty:

- "I don't know if I actually enjoy neuroscience or just think I should"
- "Will I regret not taking the 'safer' traditional course?"

One more layer

Something can sit at 'low' uncertainty and still catch you off guard — or sit at a 'high' uncertainty and feel completely mundane. How surprising something feels is its own separate layer on top of how uncertain it technically is.

Many of the classifications assume you at least know what could happen. What about the case where you don't even know that?




Populations vs. individuals

Two people can be standing in the exact same spot on that uncertainty map — same information, information, same odds, same fog — and still make completely different calls. Why?

Average adult height in a population is a solid, well-established number. But if you asked me to guess your neighbor's height from that average alone, I'd probably be wrong.

A population-level average can be extremely accurate and still tell you almost nothing about any one individual. Standard decision-making models often have the same property.

A population distribution and a single individual



Population level vs. individual level

	Population level	Individual level
What it tells you	A well-estimated average and spread across many people.	How this particular person is likely to behave.
Height example	Average adult height in a population, precisely known.	Your neighbor's exact height — unknown from the average alone.
Risk-taking example	Average risk-aversion parameter across investors, useful for pricing markets in aggregate.	Whether you personally will take a given bet.
Course-choice example	The overall spread of factors driving the room's decisions.	Which single factor is decisive for you.
Most useful for	Policy, market-level pricing, designing for a “ typical” user.	Personalized advice, understanding one person's choice.

Let's test this on ourselves

Look at the results from our survey earlier — you didn't converge on the same answer.

Same story with money. Offer two people an identical bet — a coin flip for a meaningful sum — and one says yes without blinking, the other says absolutely not. Same odds, same payout, same information.

And one of the things that differ is a personal parameter economists call risk aversion.

Why we start with value, not uncertainty

You might expect a course called 'Decision-Making Under Uncertainty' to dive straight into uncertainty. We're not going to — not yet.

Before we can usefully ask 'how uncertain is this,' we need to answer a more basic question first: uncertain about what?

What's actually at stake — what are we trying to protect, grow, or gain? That's value, and it's Weeks 2 through 4.

The same models and methods we build for value get reused, and later extended, once uncertainty gets added back in.



>



The semester roadmap

Weeks	Block	Main move
2–4	Value	What “this matters to me” formally means, before any uncertainty is added — econ, psych, and neuro vocabulary, then methods, then the same question with brain data. (Week 2 leans heavily into classical economic theory — Bernoulli to von Neumann–Morgenstern)
5–6	Learning	How you adjust your model of the world when new information arrives — both what you believe, and how you expect the world itself to behave.
7–8	Uncertainty	Only now, with value and learning in hand, do we get to uncertainty itself — more parameters, more factors, more ways for a model to be wrong. Today's fuzzy risk/ambiguity/uncertainty vocabulary becomes precise and testable.
9–10	Modulators	Hold value and uncertainty completely constant, and people still decide differently: mood, the pressure of being watched, deciding fast vs. carefully, deciding for someone else.
11–13	Applications	The payoff: real analysis on real scenarios, investment behavior and trading, sustainable decision-making, and strategic foresight. Same domain as today's cases, a sharper version — not a promise to resolve today's exact examples.

Back to where we started


The friend's tip — real money, no track record. We now have a word for that: ambiguity, not just risk. Still not telling you whether to say yes.

The person who pays more to choose sustainably — three disciplinary lenses, three partial stories, no way yet to say which one's right.

The course choice — some of you are still deciding. Your own numbers just showed this room doesn't agree on what matters most.

And the visual images, the reach, the chocolate bars, from the very start — today's four decisions were deliberately not that simple. That's the whole point of the next twelve weeks.


 **Moodle — same question, same scale:** → post your 0–10 likelihood number again.



Back to where we started

Beginning of session

On a scale of 0 to 100%, how likely are you to enrol in this course?



End of session

On a scale of 0 to 100%, how likely are you to enrol in this course?

