NACS 645 – Reasoning with heuristics and biases

-

Valentin Guigon

DEPARTMENT OF PSYCHOLOGY

SLD Lab

PROGRAM IN NEUROSCIENCE & COGNITIVE SCIENCE

Modern traditions of rationality

- Logical rationality (Cold War era)
- Heuristics-and-biases program (70s)
- Ecological rationality (80s-2000s)

Logical rationality

Homo economicus

- Rational agents
- Perfectly informed about options and consequences
- Maximize their subjective expected utility

Rational choice theory

- Preferences are:
  - Complete (any two options can be compared)
  - Transitive (internally consistent)
  - Stable across time and contexts
- Choices are the result of consistent utility maximization
- Rational thinking guarantees alignment between preferences and choices

Collective implications

- Aggregate behavior reflects the sum of individual rational choices
- Markets converge toward efficient allocation

In this framework, rationality is a **normative ideal** defined by formal axioms of consistency.

Mill, Pareto, von Neumann & Morgenstern, Savage, Samuelson

Logical rationality: scope, limits, and methodological generalization

- **From its origins, its normative scope was debated**:
  - Is logical rationality a universal norm, or is its validity bounded to specific environments?
  - Do deviations from axioms indicate irrationality, or limits of the theory's assumptions?
- **Early rational choice theorists explicitly recognized two fundamental obstacles to optimization**:
  - Uncertainty: future states and consequences cannot be fully specified ("unknown unknowns")
  - Intractability: optimal solutions exist in principle but cannot be computed by humans or machines (e.g., chess)
- **Key figures acknowledged these limits**:
  - Savage restricted EU to small worlds, where all states and outcomes are known
  - Allais and Ellsberg demonstrated systematic violations of Savage's choice axioms
  - Morgenstern emphasized unavoidable error in complex systems
- **Despite these, economists generalized SEU assumptions** (Friedman, 1953):
  - Psychological realism is not required
  - Agents can be modeled as if they maximize EU
  - Theories should be evaluated by whether predictions are "good enough" or better than alternatives
- **This enabled logical rationality to scale from individual choice to markets and macroeconomics** (see low D. Friedman et al., 2014& for findings on low predictive performances)

Classical rationality scaled to markets

Efficient market:

- Many agents interact within markets
- In a given market, all information is directly available (efficient market, no information asymmetry)
- Markets foster rationality -> choices reveal preferences
- Free and perfect competition -> optimal equilibrium (efficient allocation), triumph of the most efficient

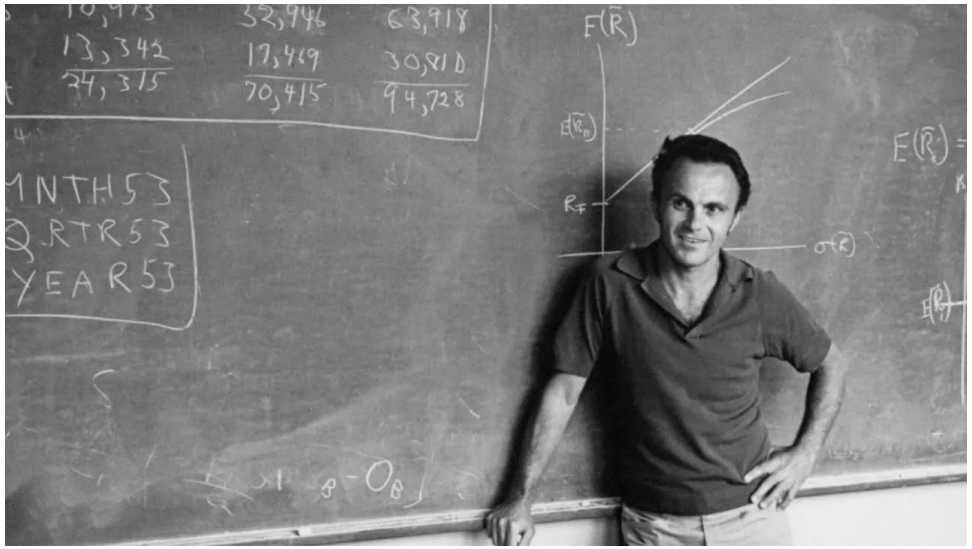

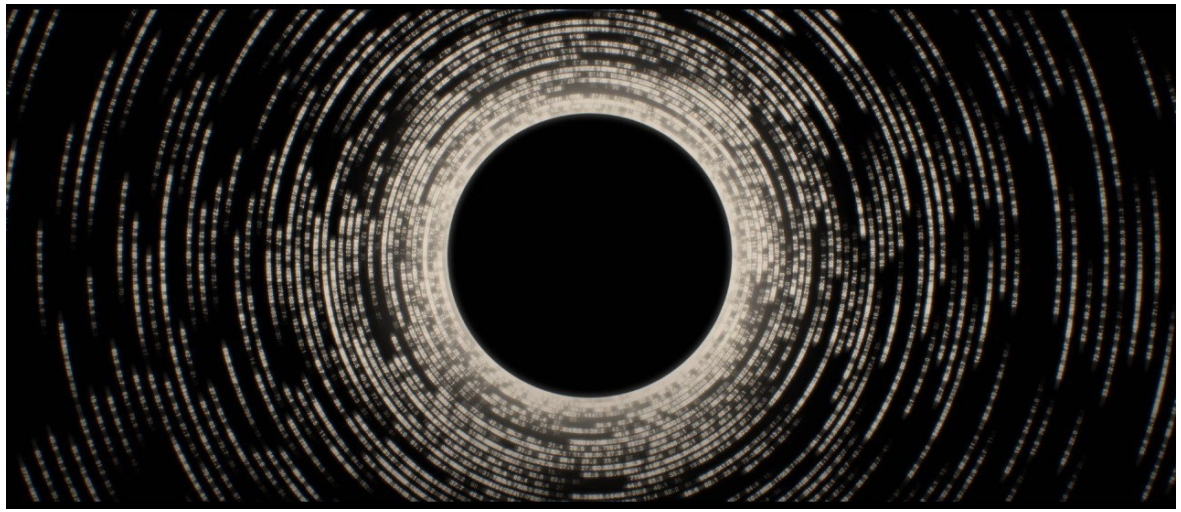

Marketplace of ideas:

Analogy of the free market applied to ideas and expression. When information is abundant and competition between ideas is free:

- Higher quality information gets the upper hand
- Truth prevails
- Quality of judgments improve

Hayek, Mill, Pareto, Fama vs Schiller

Gigerenzer, 2025. *Behavioural Public Policy*.

Modern traditions of rationality

- Logical rationality (post-WWII): **maximization of subjective expected utility, consistency axioms, Bayesian probability updating, Nash equilibrium and backward induction**.
  - Researchers sought *pure rationality, valid universally and eternally, independent of the problem at hand, and ideally to be used mechanically by a computer*.
  - Cold War strategists assumed that decision-makers would follow the axioms of logical rationality; original theorists of EU **did not intend these axioms to be universal**.

Heuristics and biases

"People do not appear to follow the calculus of chance or the statistical theory of prediction."

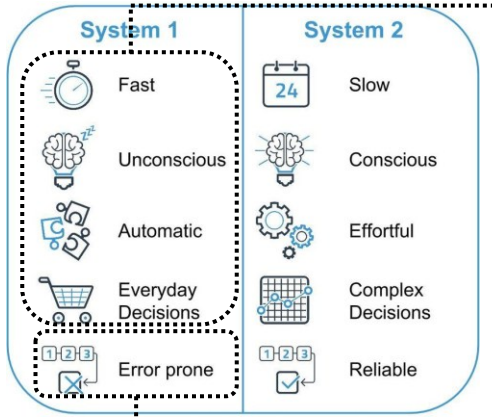

Heuristics

- **Simple, fast, frugal** cognitive strategies
- Often ignore part of the information to find a good-enough (rather than perfect) solution

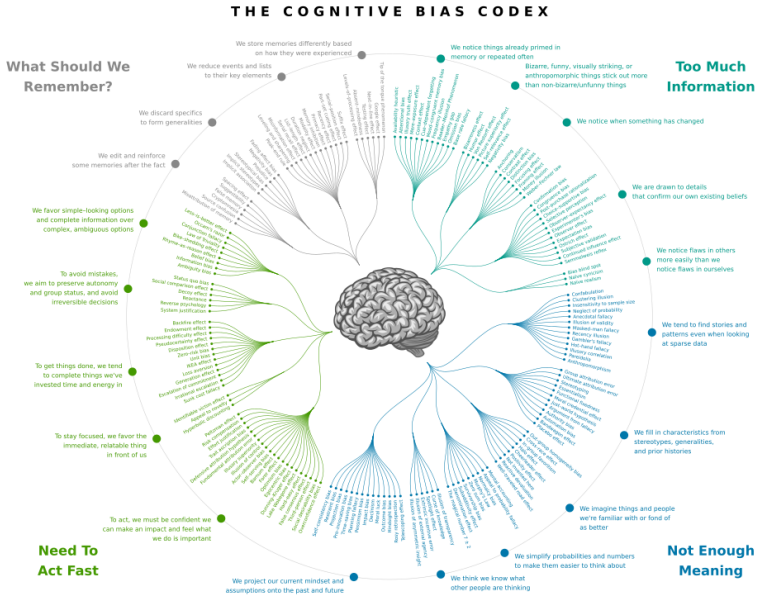

Biases

- **Error or systematic deviations** between a human judgment and a norm of rationality (e.g., laws of probabilities or logic)

Heuristics and biases: The cost-accuracy trade-off

"Heuristics and Bias": the idea of fast, frugal but low-quality mental software.

1) **Heuristics are always the 2nd best choice**
2) **We use them because of cognitive limitations**
3) **More informations, more computation and more time would always be preferable (1st choice)**

- Based on the **cost-accuracy trade-off** hypothesis: accuracy is linked to the effort expended (information, calculations, time)
- Heuristics would save resources at the cost of quality

Principle of total evidence

Cost-accuracy trade-off

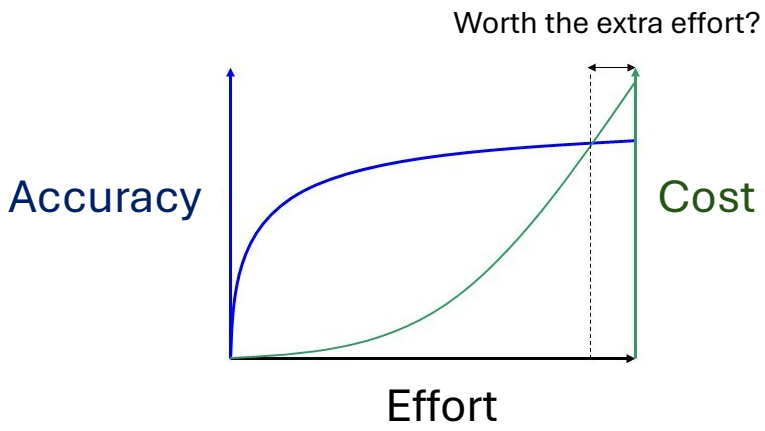

Cost-benefit ratio to be optimized

Gigerenzer & Brighton, 2009

e.g., Representativeness heuristic

Law of large numbers:
Every event is independent

$$p = \frac{1}{2}$$

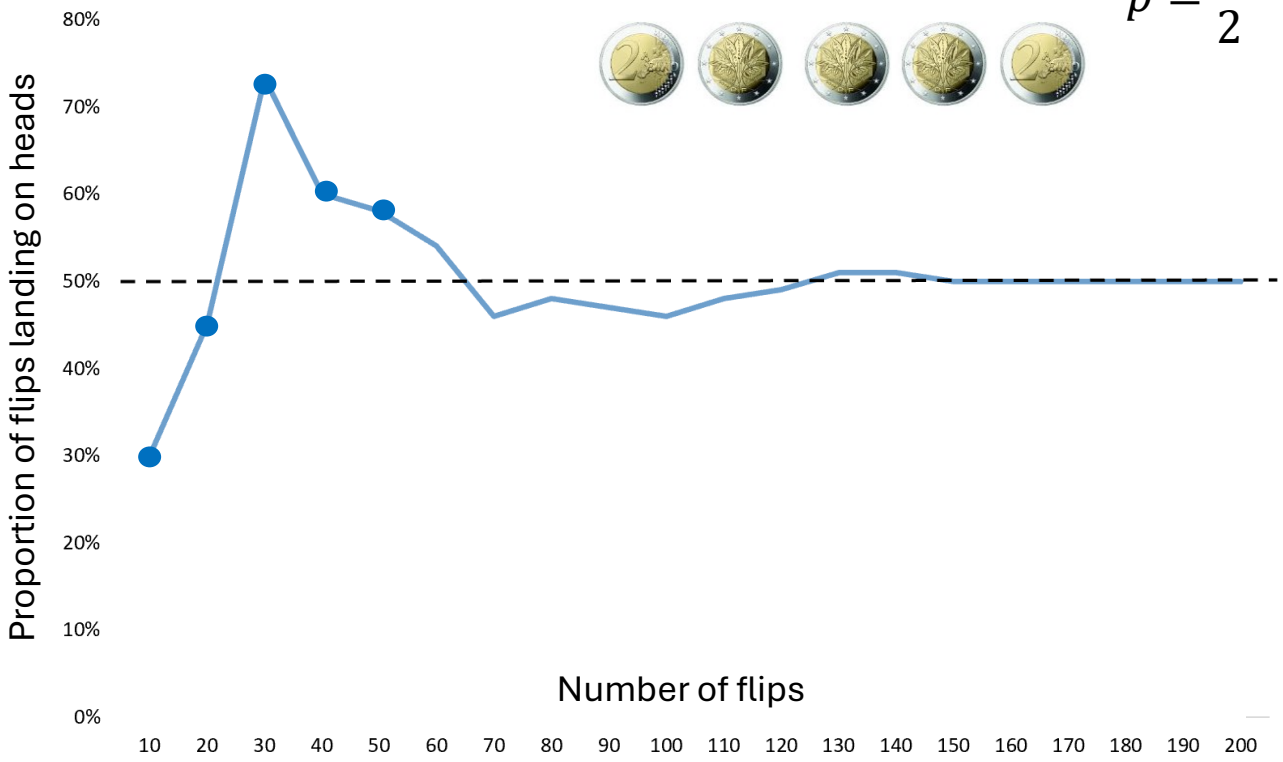

	Observed coin flips	Probability	Results
1		50/50	Heads
2		50/50	Heads
3		50/50	Heads
4		Gambler's fallacy	Predicted Tails

Representativeness heuristic: probability judgments based on how much an instance/pattern resembles the assumed generating process, rather than on base rates or sample size.

Law of small numbers: properties of short streaks are not properties of populations.
While $p(HHHH) = p(HHHT)$ for sequences of length k, **the probability of encountering at least one HHHT is higher than that of HHHH when the total number of flips *n* exceeds *k*** (i.e., $k < n$).

Modern traditions of rationality

- Logical rationality (post-WWII): **maximization of subjective expected utility, consistency axioms, Bayesian probability updating, Nash equilibrium and backward induction**.
  - Researchers sought *pure rationality, valid universally and eternally, independent of the problem at hand, and ideally to be used mechanically by a computer*.
  - Cold War strategists assumed that decision-makers would follow the axioms of logical rationality; original theorists of EU **did not intend these axioms to be universal outside small worlds**.
- Heuristics-and-biases program (70s): **deviations from logical rationality were attributed to flaws in human judgment**.
  - Humans are **systematically** and predictably **irrational**.
  - People reduce task complexity (e.g., estimating probabilities) **with heuristics that are economical and efficient** but **systematically violate logical norms**, producing stable patterns of bias.
  - Deviations were taken as evidence that human judgment is error-prone and therefore in need of external steering (e.g., nudging).
  - H-and-B authors treated **logical rationality as a universal norm** for how individuals should decide.

Are heuristics really the 2nd best choice?

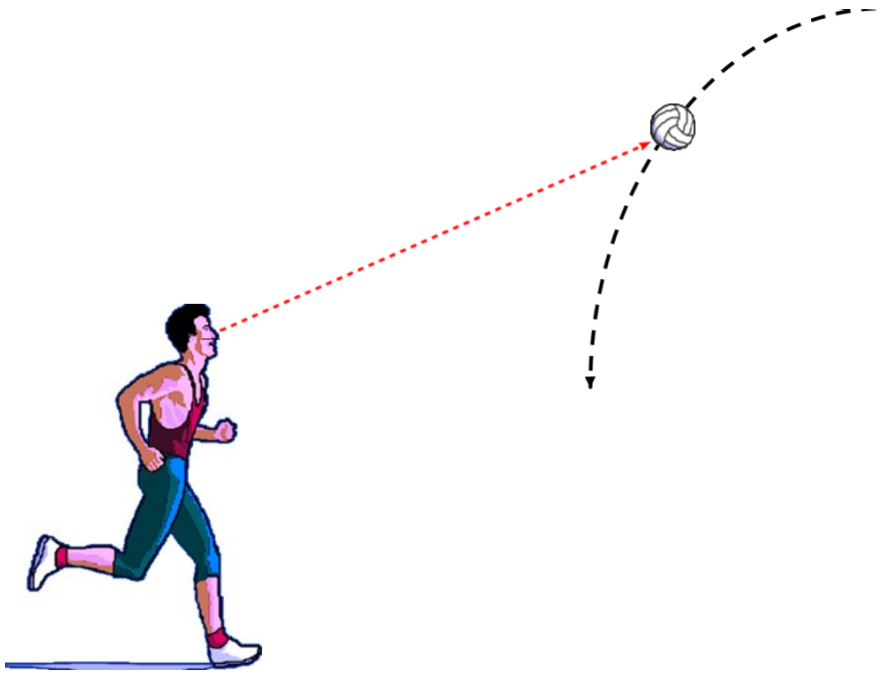

As efficient and less demanding than a calculation based on more information (e.g. differential equations)

Heuristics are functional responses to environmental uncertainty:

- Ignore information
- Are computationally efficient (neither maximization nor optimization)
- Good enough" solutions
- Are adapted (a given heuristic is optimal in certain environmental contexts, to the detriment of others)

Gigerenzer & Brighton, 2009

Less is more
- is more always preferable?

Reasons for the cognitive system's use of heuristics:

1. Cost-accuracy trade-off
   (cost savings)

2. « Less is more » (selective ignorance)

There is a point at which **more** information (indices, weights or dependencies between indices) or more calculations can become **detrimental, regardless of cost.**

More can reduce precision

Cost-accuracy trade-off

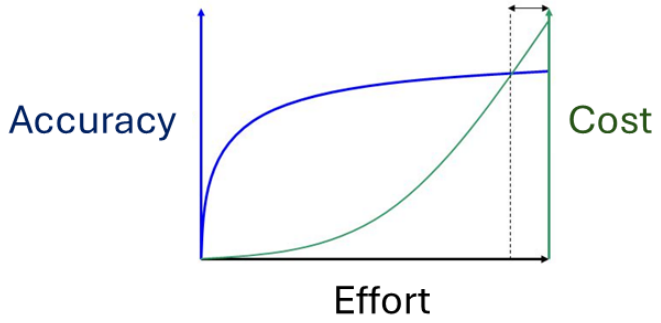

Less is More

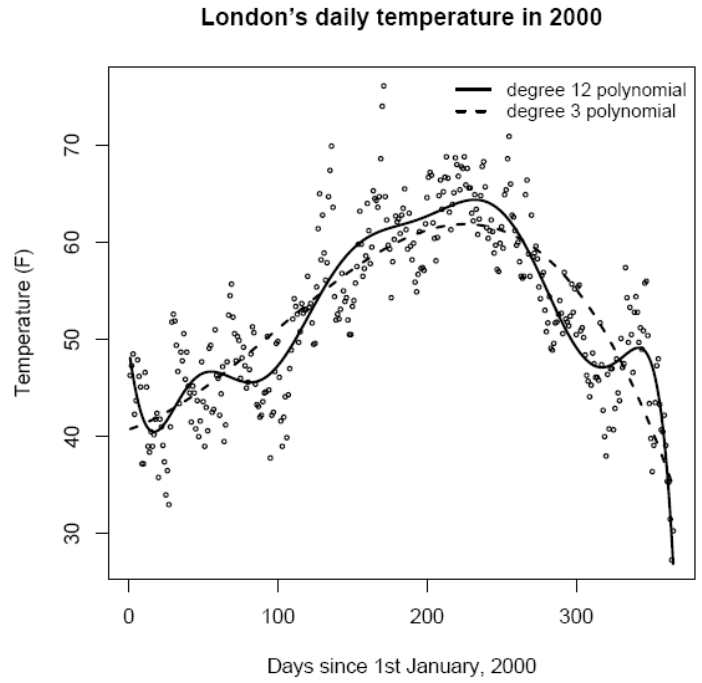

Gigerenzer & Brighton, 2009

Less is more - cognitive limitations?

A model that takes all information into account (**good *fit***) does not guarantee good performance.
The model could simply absorb **unsystematic variations**.

The ability to predict unobserved events (**good prediction**) is a better indicator.
Models are predictive **because** they primarily capture systematic regularities.

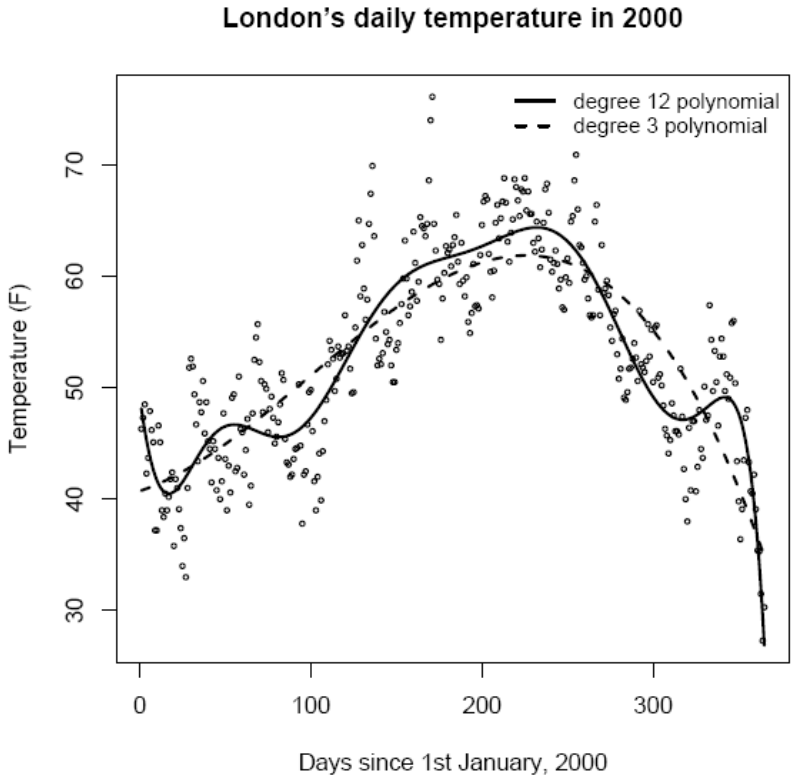

Polynomial 3 (low variance):
Low effort, medium accuracy, strong prediction

Polynomial 12 (high variance):
High effort, medium accuracy, low prediction

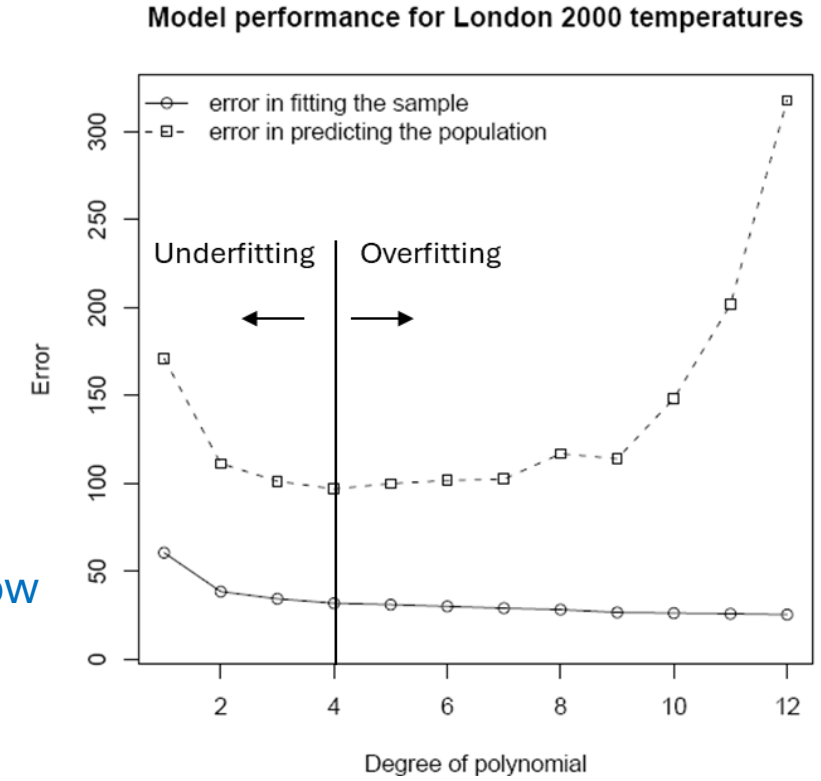

Gigerenzer & Brighton, 2009

Less is more
- example: *Take the best*

One-good-reason family heuristic:
Uses binary indices (1 vs 0) ordered by predictive validity.

Three operational rules:

1. Search rule
2. Stopping rule
3. Decision rule

City	Population	Soccer team?	State capital?	Former GDR?	Industrial belt?	License letter?	Intercity train-line?	Expo site?	National capital?	University?
Berlin	3,433,695	0	1	0	0	1	1	1	1	1
Hamburg	1,652,363	1	1	0	0	0	1	1	0	1
Munich	1,229,026	1	1	0	0	1	1	1	0	1
Cologne	953,551	1	0	0	0	1	1	1	0	1
Frankfurt	644,865	1	0	0	0	1	1	1	0	1
...	...									
Erlangen	102,440	0	0	0	0	0	1	0	0	1
Cue validities:		**0.87**	**0.77**	**0.51**	**0.56**	**0.75**	**0.78**	**0.91**	**1.00**	**0.71**

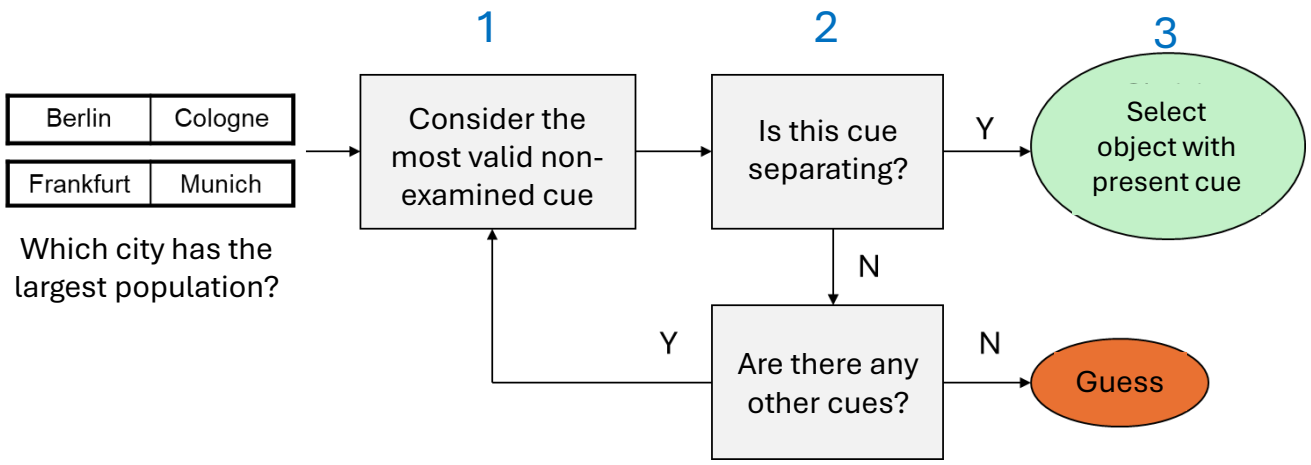

Gigerenzer & Goldstein, 1996; Gigerenzer & Brighton, 2009

Less is more
- exemple: *Take the best*

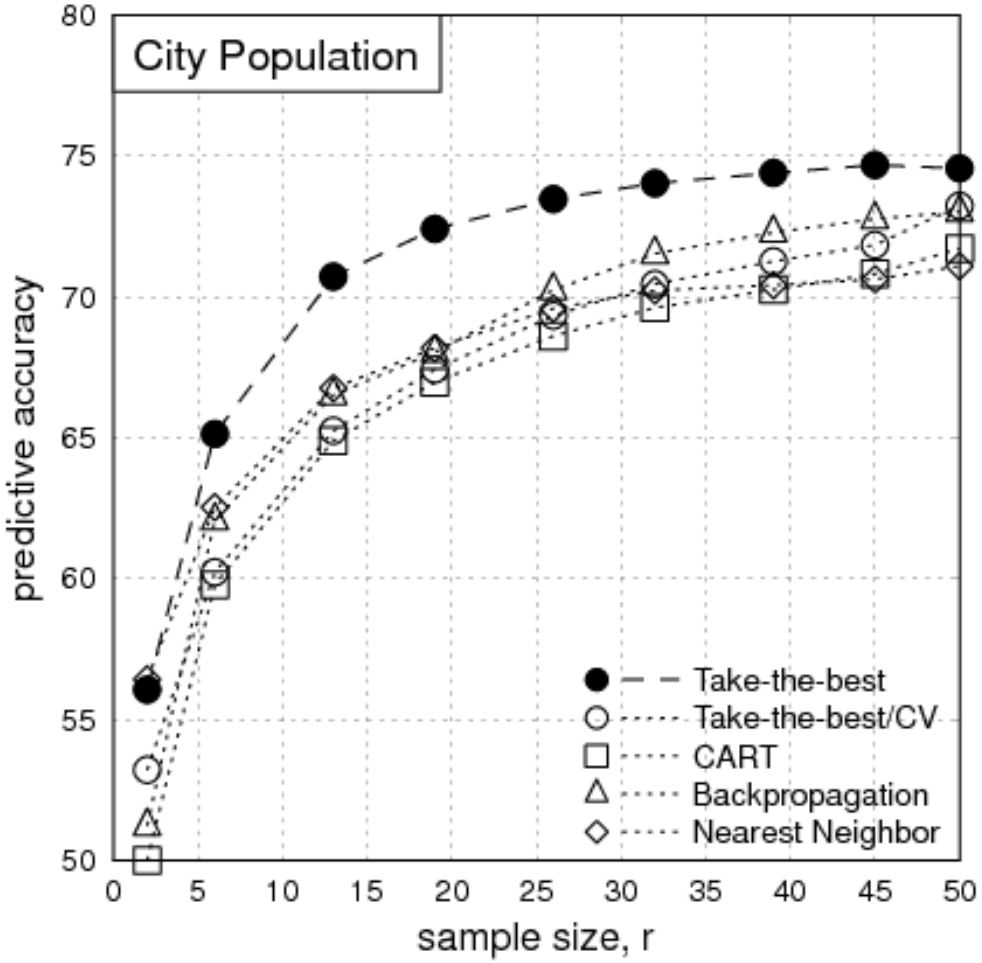

Table 5-4: Performance Across 20 Data Sets

		Accuracy (% Correct)	
Strategy	Frugality	Fitting	Generalization
Minimalist	2.2	69	65
Take The Best	2.4	75	71
Dawes's rule	7.7	73	69
Multiple regression	7.7	77	68

Performance in 20 environments

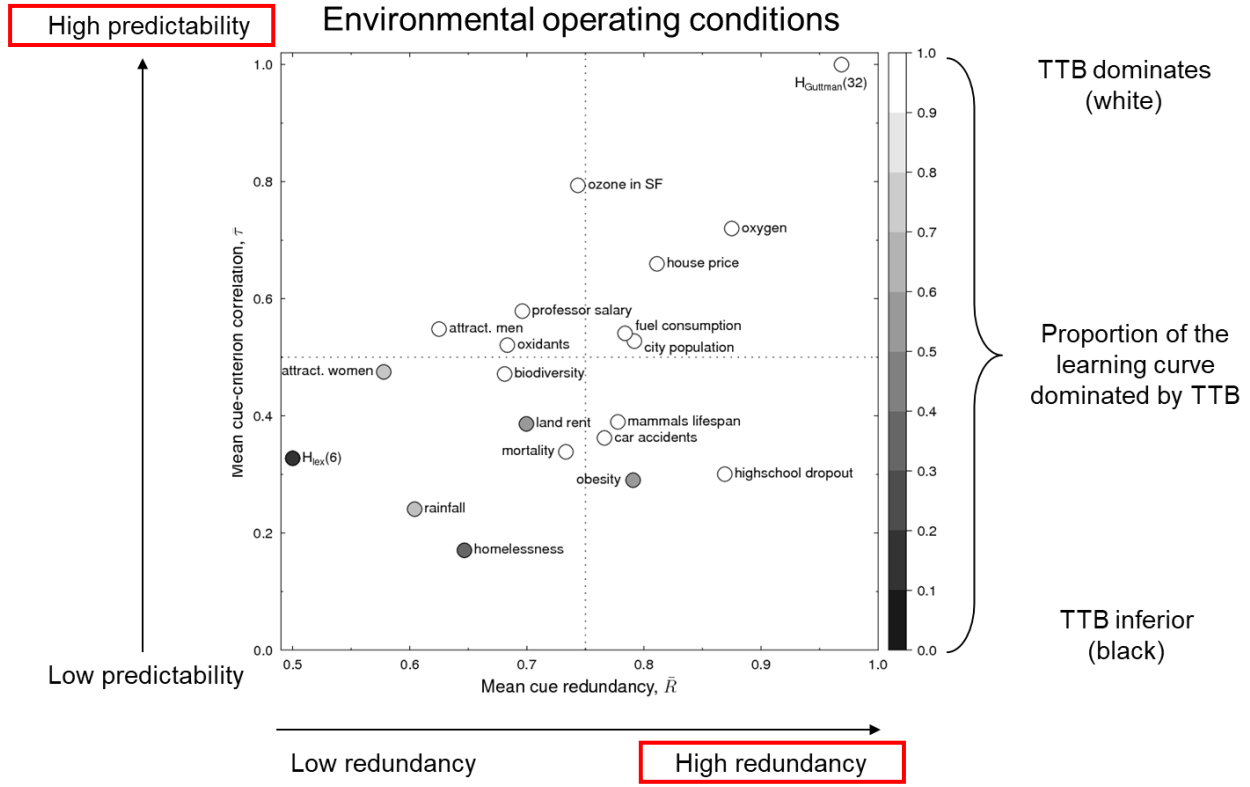

Czerlinski, Gigerenzer, & Goldstein, 1999; Gigerenzer & Brighton, 2009

Small- and large-worldedness

Gigerenzer, 2025. *Behavioural Public Policy.*

- **Small worlds:**
  - All states, outcomes, and probabilities are **known**. The problem is **well-defined**
  - **Optimization** (SEU, Bayesian updating, backward induction) **is feasible and normatively appropriate**
  - *Risk*: probabilities are known; *Ambiguity*: probabilities are unknown but states and outcomes are
- **Large worlds:**
  - States, outcomes, and probabilities are **unknown**. Problem is **ill-defined**. Optimal computations are **intractable**
  - Optimization breaks down; **simple heuristics become the rational strategy**
  - *Uncertainty:* the unknowns are unknown, probabilities can't be defined

	Small worlds		Large worlds	
Conditions	Risk	Ambiguity	Uncertainty	Intractability
Are all possible future states and consequences of all actions known?	Yes	Yes	No	Yes
Are all probabilities known?	Yes	No	No	Yes
Can optimal action be calculated?	Yes	Yes	No	No

- **Applying small-world norms to large-world decisions produces the appearance of "biases"**

When biases lead to better inferences (vs. complex models)

When decisions are made in large worlds (uncertainty and/or intractability):

- When information is scarce, degraded, noisy, or limited
- When the environment contains exploitable structure despite uncertainty

- **Predictive superiority**
- **Robustness to uncertainty**
  Ignoring information can make predictions less sensitive to noise and small samples
- **Cognitive efficiency**
  Reduces cost while maintaining sufficient performance (Martignon et al., 2008)

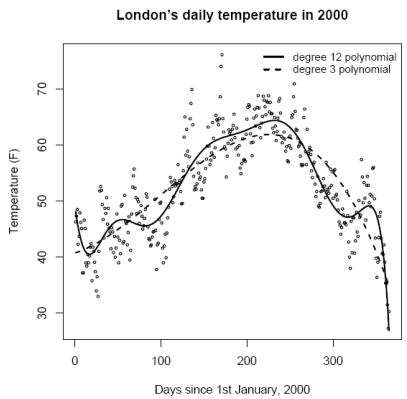

Bias-variance trade-off:

- **By simplifying**, heuristics introduce a bias:

→ This bias reduces the instability of predictions (variance)
→ Improves robustness and generalizability to similar situations, especially under uncertainty

- **By complexifying**, models reduce bias :

→ This increases the variance of predictions
→ Reduces ability to generalize to new situations

Prediction error in predictive models:

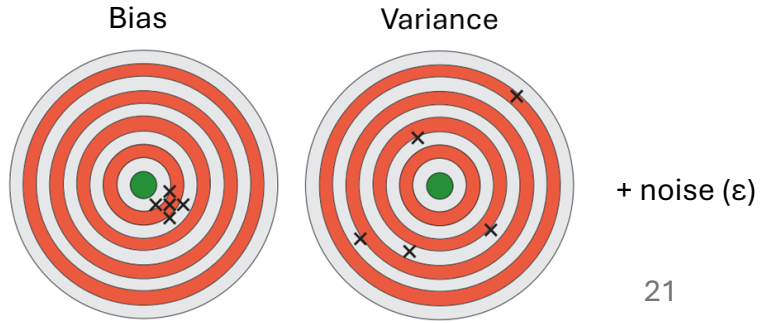

+ noise (ε)

Gigerenzer & Brighton, 2009

Modern traditions of rationality

- Logical rationality (post-WWII): **maximization of subjective expected utility, consistency axioms, Bayesian probability updating, Nash equilibrium and backward induction**.
  - Researchers sought *pure rationality, valid universally and eternally, independent of the problem at hand, and ideally to be used mechanically by a computer*.
  - Cold War strategists assumed that decision-makers would follow the axioms of logical rationality; original theorists of EU **did not intend these axioms to be universal outside small worlds**.
- Heuristics-and-biases program (70s): **deviations from logical rationality were attributed to flaws in human judgment**.
  - Humans are **systematically** and predictably **irrational**.
  - People reduce task complexity (e.g., estimating probabilities) **with heuristics that are economical and efficient** but **systematically violate logical norms**, producing stable patterns of bias.
  - Deviations were taken as evidence that human judgment is error-prone and therefore in need of external steering (e.g., nudging).
  - H-and-B authors treated **logical rationality as a universal norm** for how individuals should decide.
- Ecological rationality (90-2000): **rationality is bounded** by **biological and ecological constraints** such as uncertainty, task complexity and limited cognitive resources.
  - Heuristics function by ***purposeful*** information restriction.
  - Maximizing utility cannot handle ill-defined situations of **uncertainty** and **intractability**.
  - Biases are not cognitive failures; they appear when *purposeful information restriction* is evaluated under the **wrong *normative standard*** (i.e., assuming full information and unlimited computation).