Noisy Cognition and Economic Decision Making

Michael Woodford

Columbia University

Sloan-NOMIS Summer School
on Cognitive Foundations of Economic Behavior
June 25, 2022

The Random Element in Economic Decisions

- Standard theory implies that a given DM's choice should be a perfectly **predictable** function of the distribution of returns associated with alternative options

  — they should with certainty choose the option that implies the highest expected utility (or at any rate, the distribution of returns that is most preferred under some well-defined ordering)

The Random Element in Economic Decisions

- This postulate isn't easily testable in the case of decisions observed "in the wild"

  — hard for an observer to be sure exactly how the possible returns are understood by a given DM

- But it **can** be tested in the case of laboratory experiments, in which both possible payoffs and their probabilities are **stated** by the experimenter

  — and there choices are observed to be **random,** though with **probabilities** that vary systematically with the properties of the gambles offered

Mosteller and Nogee (1951)

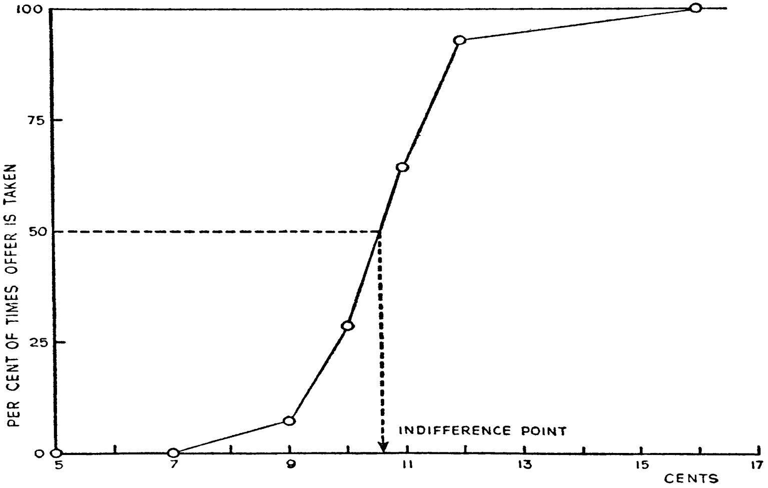

Understanding Randomness of Choice

- A common interpretation of such observations: people **have** a well-defined valuation for each possible option, which depends only on its features (and hence is invariant across contexts)

  — but instead of choosing the highest-valued option with certainty, the probability of choosing a given option depends on **how great the difference in value** relative to the alternatives

- This explains Mosteller and Nogee's method: they expect vonN-M utility to explain when two lotteries are **equally valued,** as revealed by 50-50 choice frequency

Stochastic Choice From Comparison Noise

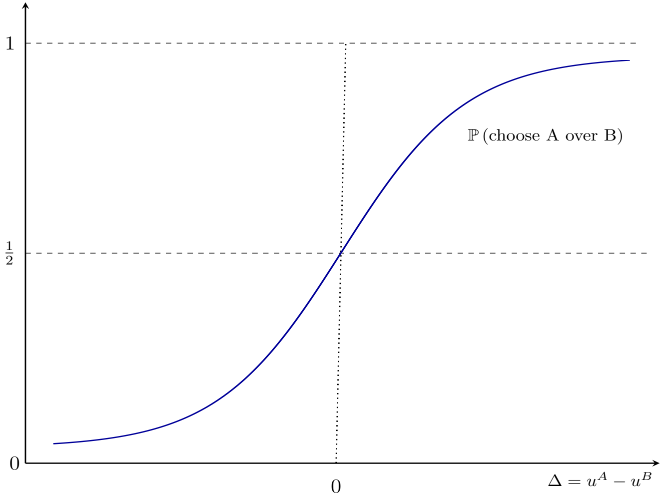

common functional form: $\Phi(\Delta) = \frac{e^{\Delta/\phi}}{e^{\Delta/\phi}+1}$

Understanding Randomness of Choice

- This interpretation consistent with many familiar models of stochastic choice:
  - Luce (1959) model
  - "softmax" choice
  - additive random utility models
  - drift-diffusion model
  - quantal-response model

"Late Noise" vs. "Early Noise"

- In such models, **noise** only enters at the **end** of the choice process, when the (accurately computed) values of the various choice options must be **compared** in order to choose between them

- But there is an alternative possible source of randomness in responses: the hypothesis that the **features** that define the available options are **corrupted by noise,** before they can be integrated to compute assessments of value

"Late Noise" vs. "Early Noise"

- This would then result in random choice [as a function of the **objective** features], even if the value estimates (and hence choices) are perfectly optimal, conditional on their being based on noisy representations

- A common interpretation of randomness of **perceptual** judgments

  - early stages of processing of many sensory features are demonstrably random [random firing of cortical neurons can be measured, and can in some cases be shown to explain randomness of judgments: e.g., Newsome *et al.*, 1989]

  - yet judgments may be modeled as **optimal,** conditional on noisy sensory data [e.g., signal detection theory, Bayesian models]

“Late Noise” vs. “Early Noise”

- But the “early noise” need not be perceptual: even in the case of data that are accurately perceived [e.g., because presented symbolically], the data may be recorded and/or retrieved with noise when they need to be **integrated** to form an overall value assessment
  - such noise in integration processes is observed even in the case of perceptual judgments [Drugowitsch *et al.*, 2016]

“Late Noise” vs. “Early Noise”

- But the “early noise” need not be perceptual: even in the case of data that are accurately perceived [e.g., because presented symbolically], the data may be recorded and/or retrieved with noise when they need to be **integrated** to form an overall value assessment

- Even when numerical data are presented symbolically, there is evidence that the brain **also** represents the **semantic** content of the number symbols in an approximate way, similar to the representations of sensory magnitudes [Dehaene, 2011]

  — can be detected from the structure of errors when responses to the numbers must be made **rapidly** [e.g., Dehaene *et al.*, 1990], or when data must be recalled after a **delay** [e.g., Dehaene and Marques, 2002]

Does the Nature of the Noise Matter?

- But is there any observationally distinguishable **difference** between models with

  - **noisy** evidence about the situation, but a **reliable** (perhaps optimal) response to the noisy data **[“early noise”]**

  VS.

  - **reliable** recognition of the situation, and computation of the values of presented choices, but a **noisy** response on basis of that info **[“late noise”]**?

Does the Nature of the Noise Matter?

Cases where the hypothesis of early noise + optimal decoding has different implications:

1. Biases in the estimation of individual **features** of a choice option, resulting from noisy encoding of the individual features, can result in estimates of its overall value that are not simply a function of the **true overall value** [i.e., the value that would be computed from the true features]

   - for example, even if decision rule is adapted to maximize DM's **average financial reward** from decisions [arguably the right objective, when all gambles are **small,** so that marginal utility of wealth should be essentially the same across outcomes], the probability of choosing one lottery over another need not be a function solely of their respective **expected values**

Illustration: Choice Between Lotteries

- Khaw *et al.* (2021): model subjects' choice between a lottery offering payoff X with probability $p < 1$, and a certain amount C
  - decision problem on a given trial defined by two (variable) amounts X and C
  - each quantity assumed to have a noisy internal representation

    $$r_Y \ \sim \ N(m(Y), \nu^2)$$

    where $Y = X$ or C
  - optimal decision rule [to max expected financial wealth], if based on noisy representations of these quantities: choose lottery iff

    $$p \cdot \mathrm{E}[X \,|\, r_x] \ > \ \mathrm{E}[C \,|\, r_c]$$

Khaw *et al.* (2021)

- Khaw *et al.* assume $m(Y) = \log Y$
  - chosen to fit observed biases in **estimation of numerosity** — assuming that similar "approximate number system" determines pattern of imprecision in both domains

- Implication: if prior (for either variable) is **also** log-normal ($\log Y \sim N(\mu, \sigma^2)$),

$$\mathrm{E}[Y \,|r_Y] \;=\; \exp(\alpha + \beta r_Y), \qquad \text{where} \quad \beta \;\equiv\; \frac{\sigma^2}{\sigma^2 + \nu^2}$$

$\Rightarrow$ optimal decision rule: accept gamble if and only if

$$\log p + \beta r_x \;>\; \beta r_c$$

Khaw *et al.* (2021)

- Predicted probability of acceptance of gamble, as function of objective data:

$$P(\text{accept}) = \Phi\left(\frac{\log(X/C) - \beta^{-1}\log p^{-1}}{\sqrt{2}\nu}\right)$$

- Continuously increasing "psychometric function," as obtained by Mosteller and Nogee (1951)
  - predicts not only **randomness** of choice, but "indifference point" above C/p
  - thus apparent **risk aversion,** even though decision actually maximizes expected value

Risk Attitude as Response to Cognitive Noise

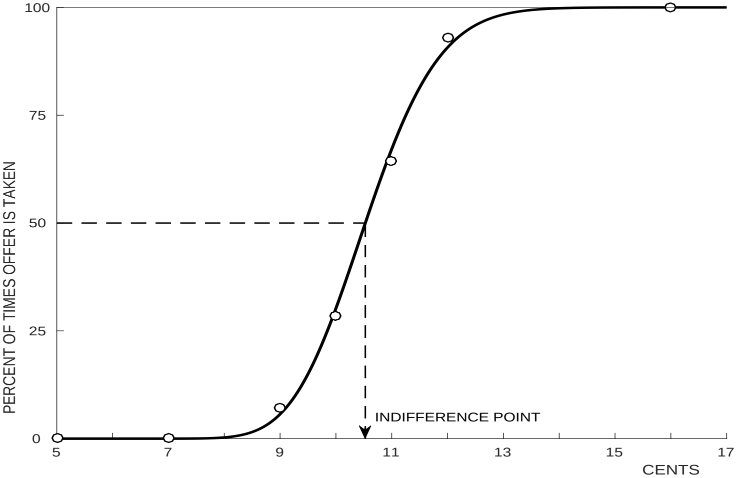

fit to Mosteller-Nogee data: if $\sigma = 0.26, \nu = 0.07$

Risk Attitude as Response to Cognitive Noise

- Are the parameters required to fit the Mosteller-Nogee figure plausible?

- Required prior uncertainty about monetary amounts: $\sigma = 0.26$ means that the difference between max and min values of $\log X$ used in the trials reported in M-N figure (1.16 log points) amounts to 4.5 standard deviations

  — a not unreasonable range if these were indeed draws from the log-normal prior

Risk Attitude as Response to Cognitive Noise

- Why not (more conventionally) explain the risk aversion as a consequence of **concave utility** for monetary payoffs, and then add **comparison noise** to explain the randomness of choice?

  1. the existence of a concave utility for the payoff from an individual small gamble only makes sense if one assumes that this source of funds isn't **integrated** with the rest of the DM's wealth

     — they must care about **this payoff** and not just their overall budget (regardless of its sources)

     - that "narrow bracketing" then requires an explanation

       — in Khaw *et al.* model, it follows naturally from **separate encoding** of the individual payoff

Risk Attitude as Response to Cognitive Noise

- Why not (more conventionally) explain the risk aversion as a consequence of **concave utility** for monetary payoffs, and then add **comparison noise** to explain the randomness of choice?

  2. Khaw *et al.* model implies **scale-invariant** choice probabilities [P(accept) depends only on X/C], for all small enough gambles

     — a comparison-noise model in which choice probability is a function of **EU difference** would not

     - experimental data in Khaw *et al.* support the prediction of scale-invariance

Testing Scale Invariance (Khaw *et al.*, 2021)

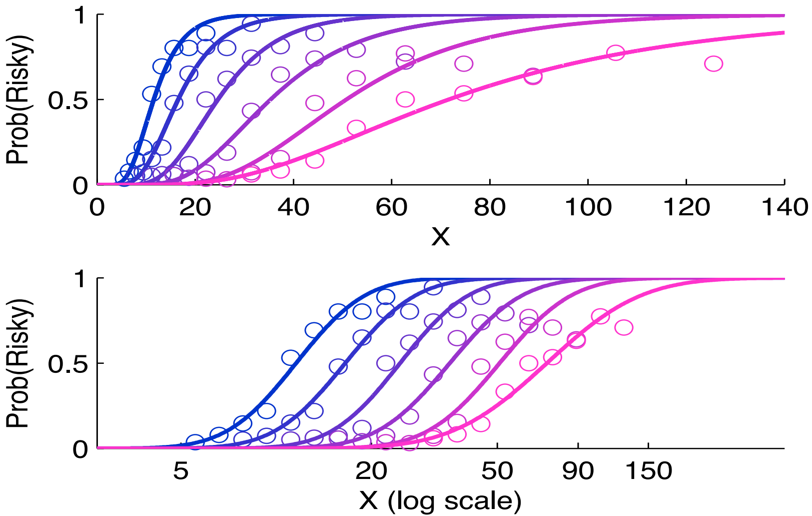

"psychometric functions" for six values of C

Testing Scale Invariance (Khaw *et al.*, 2021)

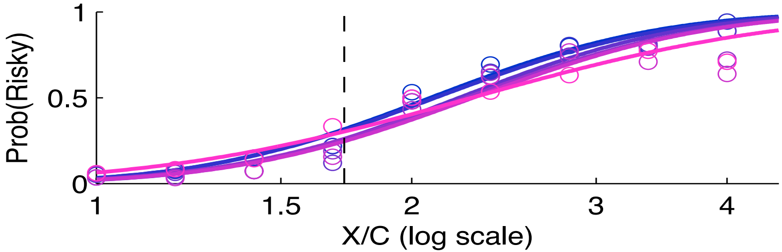

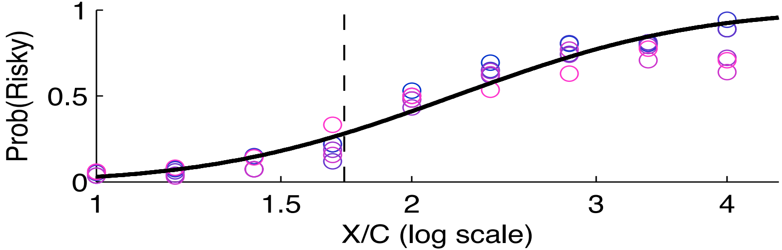

"psychometric functions" plotted against X/C

Risk Attitude as Response to Cognitive Noise

- Why not (more conventionally) explain the risk aversion as a consequence of **concave utility** for monetary payoffs, and then add **comparison noise** to explain the randomness of choice?

3. Conventional model gives no reason for degree of risk aversion and degree of stochasticity to be related to one another; Khaw *et al.* model instead ties them together

    - if parameters μ, σ of the prior are given by the range of values of $\log X$ used in experiment, then model has only a **single** free parameter ν to explain **both** degree of stochasticity (slope of psychometric function) and degree of risk aversion (location of indifference point)

    - in fact, Khaw *et al.* find that these two features of behavior are strongly **correlated** across subjects

Noise and Risk Aversion: Across Subjects

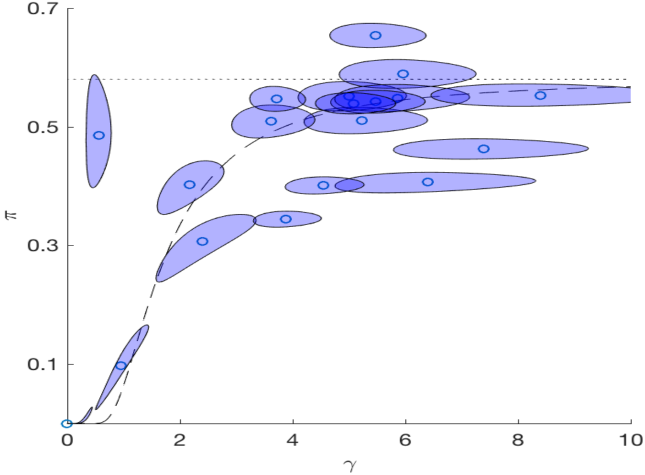

[dashed line: predicted relation if optimal decisions, common prior]

Does the Nature of the Noise Matter?

2. Noisy-coding theory implies that manipulations that change the degree of **coding precision** should change **estimation bias**

   - for example, varying **time pressure**

     — if internal evidence is a stream of noisy signals [as for example in the DDM], then less time for collecting additional signals should mean **noisier cumulative evidence**

     — Bayesian decoding of the noisier internal representation can result not just in **more variable** estimates, but in larger **average bias** in valuations

   - Polania *et al.* (2019) find that increased time pressure changes average ratings of food items

     — in a way consistent with Bayesian decoding of noisy internal representation of items' values

Does the Nature of the Noise Matter?

- Other methods of **manipulating encoding noise** can similarly influence **valuation biases**

  - Enke and Graeber (2021) elicit "certainty equivalent" values for simple lotteries [an amount x is paid with probability p; otherwise, payoff is zero]

    — look at how bias in valuation [CE/x different on average from p] varies with p

    — finding: $CE/x > p$ for small p, while $CE/x < p$ for large p, regardless of sign of x [replicating findings of Tversky and Kahneman, 1992]

Does the Nature of the Noise Matter?

- As also proposed by Khaw *et al.*, this can be interpreted as consequence of Bayesian decoding of a **noisy internal representation of p**

  — internal noise ["cognitive uncertainty"] $\Rightarrow$ estimates of p biased toward the **prior mean** [$\bar{p} = 0.5$]

- Enke and Graeber (2021) manipulate the degree of noise in the internal representation by presenting the payoff probability in a more complex form [compound lottery, rather than simply stating the implied probability p of the non-zero payoff]

  — and show that increasing noise in this way leads to **increased bias** in the elicited certainty equivalents

Enke and Graeber (2021)

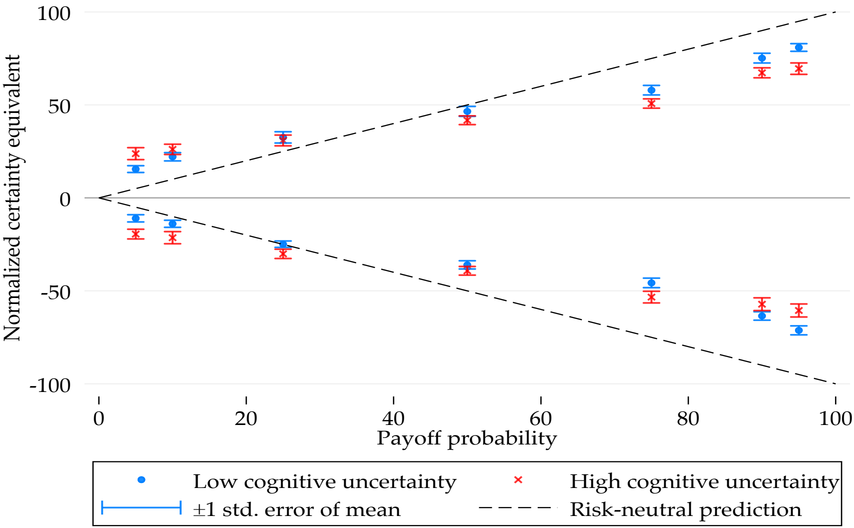

cognitive uncertainty increased by more complex presentation

Does the Nature of the Noise Matter?

3. Noisy-coding hypothesis can also explain another type of sensitivity of choice to the context in which options are encountered: **more random** choice between two given options, when they are drawn from a **wider range** of possibilities (presented on other trials)

   - this is another example of sensitivity of choice to the DM's **prior** about what the data are likely to be
   - but not because of how the prior is used in **interpreting** noisy evidence; instead, the precision of the **encoding** can vary with the prior (and hence across contexts)
   - predicted by theories of **"efficient coding"**

Efficient Coding

- Idea: the neural system used to produce internal representations of particular quantities has only a **finite capacity** to represent different amounts in sufficiently distinguishable ways

  — like the finite *capacity* of a communications channel, in Shannon's theory

Efficient Coding

- Idea: the neural system used to produce internal representations of particular quantities has only a **finite capacity** to represent different amounts in sufficiently distinguishable ways

- **Efficient coding:** the hypothesis that external stimuli are mapped into the limited variety of possible internal states in such a way as to make decisions as accurate as possible

  — worse discrimination between some states may be accepted, as the price of allowing sharper discrimination between other states, that it matters more to be able to distinguish

- This implies that the encoding scheme should depend on the **prior** (Payzan-LeNestour and Woodford, 2022)

Efficient Coding

- A particularly robust implication of efficient coding theories: **range normalization**

- Idea: the **accuracy of discrimination** between any two magnitudes will be **worse** when these two magnitudes are drawn from a prior distribution with a wider **range**

  - the larger range of objective magnitudes must be mapped into the **same** range of possible internal representations

  - hence the two magnitudes will be closer together in "psychological space" when the objective difference between them is a smaller **fraction of the overall range**

  - making the encoding noise more significant relative to the degree of difference in their internal representations

Range Normalization

- An illustration, where the internal representations can actually be observed:
  - Padoa-Schioppa (2009) measures the internal representation of the values of different choice options, by the rate of firing of certain cells in the macaque OFC ["offer cells"], when monkeys choose between offers of different quantities of two types of juice
  - the firing rate is higher when the quantity of apple juice offered is higher [this is what identifies the cells as "offer cells"]
  - but the firing rate associated with a given quantity of juice is smaller, when the range of quantities of juice that occur on different trials [in that experimental session] is greater

Padoa-Schioppa (2009)

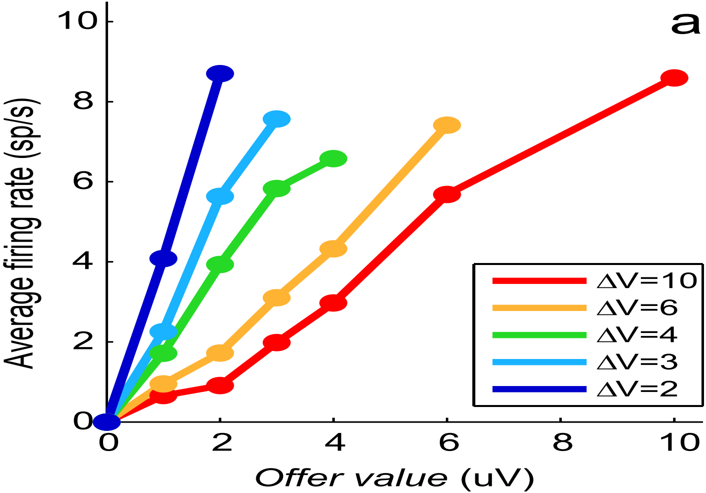

vertical axis = probability of staying

Range Normalization

- The fact that two drops of juice are differently encoded when 4 is the upper bound, than when 10 is the upper bound, doesn't mean that they are **valued more** (on average) in the former case
  - the "decoding" of the internal representation seems to adjust to the range in an efficient way as well (Rustichini *et al.*, 2017)
- But the change in encoding when the range is 10 **does** mean that two drops are not as accurately distinguished from four drops, as is the case when the range is 4
  - resulting in less predictable choices

Range Normalization

- Frydman and Jin (2022) show that the same seems to be true of the internal representation of numerical quantities in humans
  - task: a two-digit Arabic numeral is presented, and the subject must (rapidly) say whether it is greater or less than 65
  - idea: when such a judgment must be made rapidly enough, it is based on an **approximate semantic representation** of the number [triggered by recognition of the numeral], rather the kind of exact representation used in arithmetic calculations (Dehaene *et al.*, 1990)
  - consequence: more mistakes (and slower responses) when the number presented is **closer** to 65
  - but for numbers near 65, responses are **slower** (and yet more mistakes) when the numbers are drawn from a wider **range**

Number Comparison [Frydman and Jin (2022)]

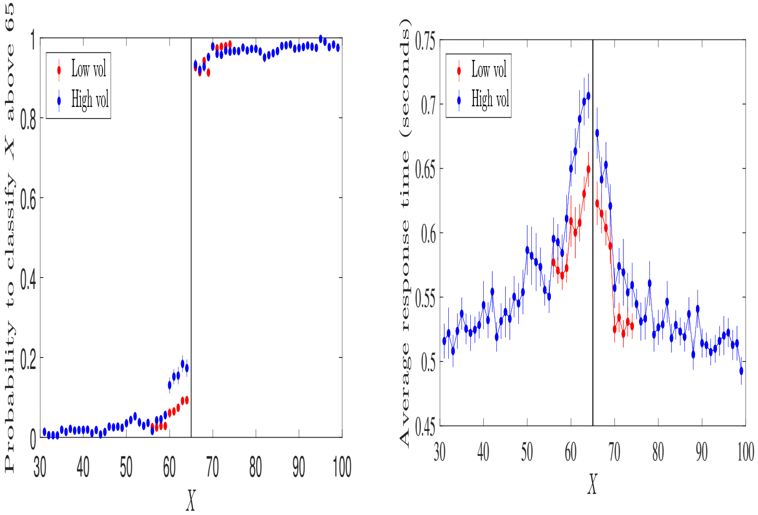

Range Normalization

- This suggests that in our model above of noisy encoding of numerical magnitudes,

$$r \sim N(\log X, \nu^2),$$

  the parameter ν should vary with the **range** of values of $\log X$

  - or alternatively: that there is a noisy representation

    $$r \sim N(m(X), \nu^2)$$

    with a fixed value for ν;

  - but the mapping $m(X)$ must adjust so that m has the **same bounded range,** regardless of the range over which $\log X$ varies

- Consequence: larger range of variation in $\log X$ should result in noisier choice between lotteries

Lottery Choice [Frydman and Jin (2022)]

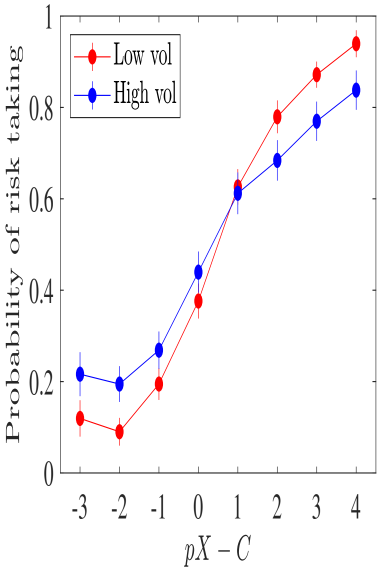

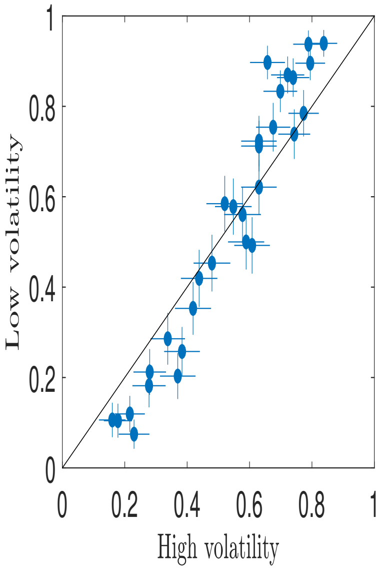

Does the Nature of the Noise Matter?

4. Noisy-encoding hypothesis also has different implications from those of comparison noise for the nature of **coordination** of different DMs' decisions in a situation of strategic interaction

A Coordination Game

- Game studied (experimentally) by Frydman and Nunnari (2021):

	leave	*stay*
leave	(θ, θ)	$(\theta, 47)$
stay	$(47, \theta)$	$(63, 63)$

- each of two players must **simultaneously** make a binary decision
- payoffs for each depend on their **joint** decision
- payoff matrix: in each cell, (a, b) means row player gets a, column player gets b
- parameter θ is different on different trials

A Coordination Game

- Game studied (experimentally) by Frydman and Nunnari (2021):

	leave	*stay*
leave	(θ, θ)	$(\theta, 47)$
stay	$(47, \theta)$	$(63, 63)$

- **Nash equilibrium:** situation in which each player's choice frequencies are **optimal** ($=$ maximize their expected payoff), given the choice frequencies of the other player
- If $\theta < 47$, only NE is for both to **stay** (with prob. 1); if $\theta > 63$, only NE is for both to **leave**

A Coordination Game

	leave	*stay*
leave	(θ, θ)	$(\theta, 47)$
stay	$(47, \theta)$	$(63, 63)$

- But if $47 \leq \theta \leq 63$, **both** of these patterns are equilibria: optimal to stay if you expect other will, optimal to leave if you expect other will

  — hence predicted outcome **not determined** by economic "fundamentals"; and (over some range) changes in θ don't reduce extent to which either outcome makes sense

A Coordination Game

- What difference does **cognitive noise** make?

- Common approach: **"quantal response equilibrium"** (McKelvey and Palfrey, 1995) generalizes NE to introduce **comparison noise**

  — each player's frequencies adapted to those of the other; but choice frequencies increasing in the **difference in expected payoff** of one's two choices

- Consequence: if only a small amount of noise, **still multiple equilibria**

  — though now the higher-expected-payoff choice isn't chosen all of the time

Quantal Response Equilibrium

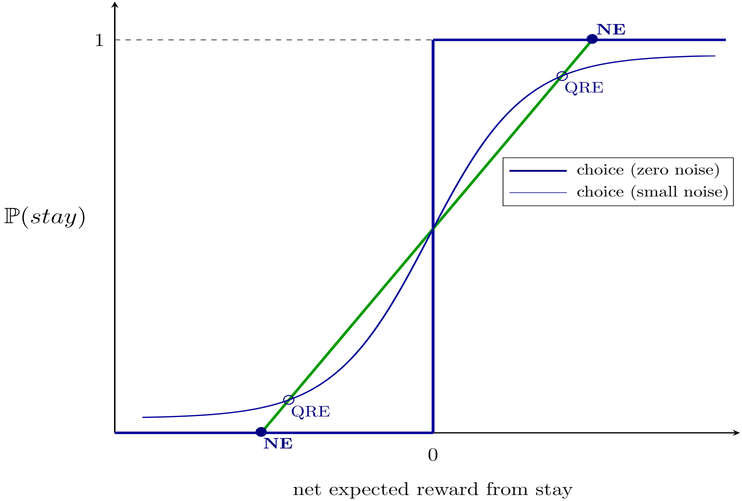

green line: net expected reward as function of other's prob of staying

Frydman and Nunnari (2021)

- Conclusions different in the case of **early noise** [the "cognitive imprecision" model of F&N]

- Suppose each player's choice must be based on noisy internal representation

$$r \sim N(m(\theta), \nu^2)$$

- Equilibrium: each player leaves if and only if $r > r^*$ (for them), where each player's threshold r^* is the point at which expected payoff from leaving is exactly equal to expected payoff from staying [given other player's decision rule, and optimal Bayesian decoding of the noisy representation r]

Frydman and Nunnari (2021)

- Result: if cognitive noise is **small enough** (though non-zero), there will be **unique** equilibrium strategies

  — unique threshold r^* such that r^* is optimal choice for a player, given that other uses threshold r^*

Determination of Equilibrium r^*

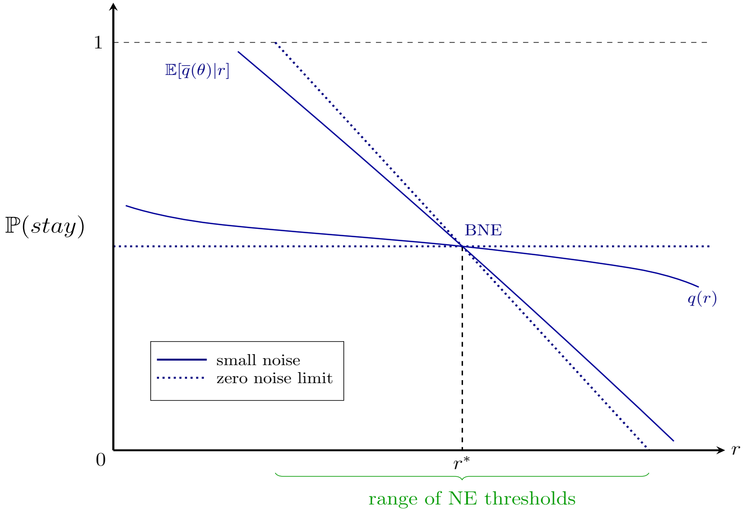

$\bar{q}(\theta) =$ required prob. other stays, to make it optimal to stay (if state θ)

$$q(r) \; = \; \mathrm{P}[r_{-i} < r \,|r_i = r]$$

Frydman and Nunnari (2021)

- Result: if cognitive noise is **small enough** (though non-zero), there will be **unique** equilibrium strategies

  — unique threshold r^* such that r^* is optimal choice for a player, given that other uses threshold r^*

- Moreover, the unique equilibrium prediction is that the probability that players leave should be an **increasing function of θ** [even for values of θ in the range where there are multiple NE]

- This prediction of **sensitivity to fundamentals** is more consistent with behavior observed in laboratory experiments (Heinemann *et al.*, 2004, 2009), and arguably with what is observed in real-world financial crises (Gorton, 1988, 2012)

Frydman and Nunnari (2021)

- A further difference: QRE implies that the possible equilibrium choice frequencies, for a given game (specified by θ), should be **independent of the prior** distribution from which the θ is drawn on different trials

- The CI model (with Bayesian decoding) instead makes the prior relevant to the equilibrium decision rules

  — and if we further assume **efficient coding,** the model predicts **greater sensitivity to θ** when the value of θ varies over a **smaller range**

- This is what Frydman and Nunnari find in their experiment

Frydman and Nunnari (2021)

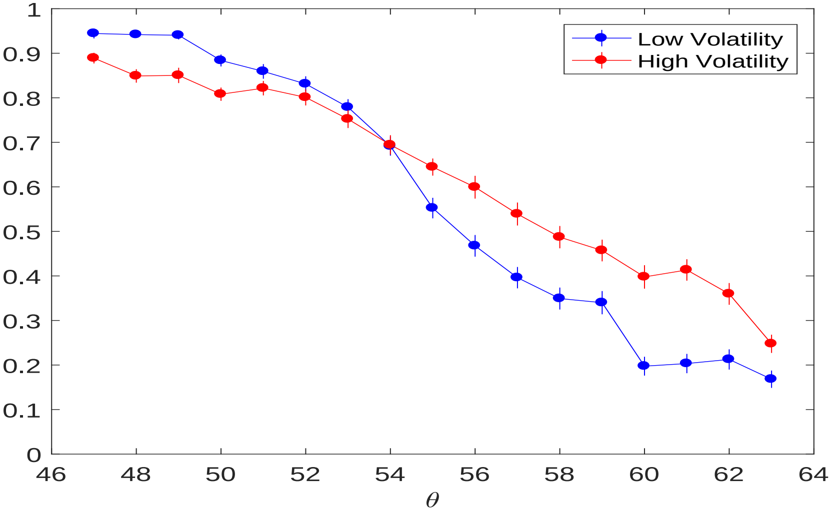

vertical axis = probability of staying

Summary

- The hypothesis that decisions are based on **noisy internal representations** of the presented data can explain phenomena that a mere assumption of **comparison noise** (or more generally, response noise) cannot
  - especially when the hypothesis of noisy coding is combined with the further assumptions of **efficient coding** and **Bayesian decoding** [often used in the literature on perceptual errors]
- This doesn't mean that there may not **also** be comparison noise

  — only that a hypothesis of comparison noise **by itself** doesn't adequately capture the role of cognitive noise in decision making
- Study of cognitive noise in other domains may help to improve economic modeling

Additional References

Dehaene, S., *The Number Sense: How the Mind Creates Mathematics,* Oxford University Press, 2d ed., 2011.

Dehaene, S., E. Dupoux, and J. Mehler, "Is Numerical Comparison Digital? Analogical and Symbolic Effects in Two-Digit Number Comparison," *J. Experimental Psychology: Human Perception and Performance* 16: 626-641 (1990).

Dehaene, S., and J.F. Marques, "Cognitive Euroscience: Scalar Variability in Price Estimation and the Cognitive Consequences of Switching to the Euro," *Quarterly J. Experimental Psychology* 55: 705-731 (2002).

Additional References (p.2)

Drugowitsch, J., V. Wyart, A.-D. Devauchelle, and E. Koechlin, "Computational Precision of Mental Inference as Critical Source of Human Choice Suboptimality," *Neuron* 92: 1398-1411 (2016).

Enke, B., and T. Graeber, "Cognitive Uncertainty," working paper, January 2021.

Frydman, C., and S. Nunnari, "Cognitive Imprecision and Strategic Behavior," working paper, October 2021.

Gorton, G.B., "Banking Panics and Business Cycles," *Oxford Economic Papers* 40: 751-81 (1988).

Additional References (p.3)

Gorton, G. B., “Reflections on the Recent Financial Crisis,” NBER Working Paper no. 18397, September 2012.

Heinemann, F., R. Nagel, and P. Ockenfels, “The Theory of Global Games on Test: Experimental Analysis of Coordination Games with Public and Private Information,” *Econometrica* 72: 1583-1599 (2004).

Heinemann, F., R. Nagel, and P. Ockenfels, “Measuring Strategic Uncertainty in Coordination Games,” *Rev. Econ. Studies* 76: 181-221 (2009).

Luce, R.D., *Individual Choice Behavior: A Theoretical Analysis,* Wiley, 1959.

Additional References (p.4)

McKelvey, R.D., and T.R. Palfrey, "Quantal Response Equilibria for Normal Form Games," *Games and Econ. Behavior* 19: 6-38 (1995).

Mosteller, F., and P. Nogee, "An Experimental Measurement of Utility," *J. Political Economy* 59: 371-404 (1951).

Newsome, W.T., K.H. Britten, and J.A. Movshon, "Neuronal Correlates of a Perceptual Decision," *Nature* 341: 52-54 (1989).

Padoa-Schioppa, C., "Range-Adapting Representation of Economic Value in the Orbitofrontal Cortex," *J. Neuroscience* 29: 14004-14014 (2009).

Additional References (p.5)

Payzan-LeNestour, E., and M. Woodford, "Outlier Blindness: A Neurobiological Foundation for Neglect of Financial Risk," *Journal of Financial Economics* 143: 1316-1343 (2022).

Polania, R., C.C. Ruff, and M. Woodford, "Efficient Coding of Subjective Value," *Nature Neuroscience* 22: 134-142 (2019).

Rustichini, A., K.E. Conen, X. Cai, and C. Padoa-Schioppa, "Optimal Coding and Neuronal Adaptation in Economic Decisions," *Nature Communications* 8: 1208 (2017).

Tversky, A., and D. Kahneman, "Advances in Prospect Theory: Cumulative Representation of Uncertainty," *J. Risk and Uncertainty* 5: 297-323 (1992).

Further Explanation of Slide 45

- If p is the probability that the other player chooses to stay, then one's own expected payoff from staying is

$$u^{stay}(p) \;=\; p \cdot 63 + (1-p) \cdot 47.$$

Instead, one's expected payoff from leaving is

$$u^{leave} \;=\; \theta,$$

regardless of the value of p. The expected payoff **differential** is therefore

$$\Delta(p) \;\equiv\; u^{stay}(p) \;-\; u^{leave} \;= (47-\theta) + 16p.$$

Further Explanation of Slide 45

- The function $\Delta(p)$ is graphed by the green line on slide 44, for some single value of θ. The vertical axis shows p and the horizontal axis the corresponding value of $\Delta(p)$. The graph of this function is an upward-sloping straight line (the location of which depends on θ).

- The graph is drawn for a case in which $47 < \theta < 63$, so that $\Delta(0) < 0$ while $\Delta(1) > 0$. (This is the case in which Nash equilibrium will be non-unique.)

- If one expects the other player to stay with probability p, **optimal behavior** requires one to stay with probability zero if $\Delta(p) < 0$, and with probability 1 if $\Delta(p) > 0$, but is consistent with any probability of staying if $\Delta(p) = 0$ exactly.

Further Explanation of Slide 45

- Thus optimal behavior requires that the value of Δ (given one's expectation about the other's choice) and one's own probability of staying must be the coordinates of a point on the **choice correspondence** $\mathcal{C}$ shown by the thick black lines on slide 44.

- A symmetric **Nash equilibrium** is a probability of staying p^* (for both players) with the property that

$$(\Delta(p^*),\ p^*)\ \in\ \mathcal{C}$$

  Thus the point (Δ, p^*) must belong to both the green line and the correspondence $\mathcal{C}$.

- When $47 < \theta < 63$, there are **multiple intersections.** The two solid dots indicate the two symmetric NE that involve *pure strategies* (deterministic behavior); these are the NE that are locally stable under plausible learning dynamics.

Further Explanation of Slide 45

- Introducing **comparison noise:** suppose that each player chooses randomly, with the probability of staying an increasing, sigmoid function $\Phi(\Delta)$

  — for example, $\Phi(\Delta) \equiv \frac{e^{\Delta/\phi}}{e^{\Delta/\phi}+1}$, where $\phi > 0$ indexes the degree of noise [low noise $\leftrightarrow$ large ϕ]

- A symmetric **quantal response equilibrium** is a probability of staying p^* (for both players) with the property that

$$p^* \;=\; \Phi(\Delta(p^*))$$

  — or alternatively, such that $\Phi^{-1}(p^*) = \Delta(p^*)$ [hence a point of intersection of the **sigmoid curve** and the **green line,** in the figure]

Further Explanation of Slide 45

- In the case of **small enough noise** [e.g., large enough ϕ in the parametric model of comparison noise proposed above], the graph of the sigmoid curve will be close to the graph, and hence will intersect the green line at points close to each of the intersections between the green line and $\mathcal{C}$

- Hence in the case of small enough noise, there will again be a **multiplicity of equilibria**

  — the points labeled "QRE" in the figure are two different quantal response equilibria, again each locally stable under learning dynamics

Further Explanation of Slide 48

- Instead introducing noise in perception of the state θ: player's decision must then be a function of internal representation r, rather than of the true value of θ [which is not accessible]

- **Optimality** of decision by player i, given the behavior of the other player $-i$: one can't define i's optimal response to representation r, taking as given how $-i$ is expected to respond to **that same** representation r

  — because when i's internal state is r, they can't be sure what the external state θ is, and hence can't be sure what $-i$'s internal state is

- Must instead consider optimal decision by i given a **rule of behavior** for player $-i$, specifying what they will they will do in the case of **any** possible internal representation r_{-i}

Further Explanation of Slide 48

- **Bayesian Nash Equilibrium:** a pair of **rules** [specifying how each player acts as a function of their internal state] with the property that each player's rule is **optimal** for them [note: no fuzziness of their choice, conditional on their internal state], given the rule of the other player, and the joint distribution of the states (θ, r_i, r_{-i})

  — optimal response for i to internal state r_i depends on the conditional distribution over actions of $-i$, given the conditional distribution $p(r_{-i} \,|\, r_i)$ and $-i$'s rule of behavior [this is where Bayesian inference comes in]

- We will further look only at equilibria where the rules of behavior are **threshold rules:** for any player i, there is a threshold $\hat{r}_i$ such that i stays if and only if $r_i < \hat{r}_i$

Further Explanation of Slide 48

- Nash equilibrium in threshold strategies, in the case of **no cognitive noise:** in this case, r reveals the value of θ with perfect precision [implied value $= \theta(r)$], and each player can be certain that the other perceives θ in the same way that they do

- It is then an NE for the two players to choose identical thresholds $(\hat{r}, \hat{r})$, where $\hat{r}$ is **any** number such that $47 \leq \theta(\hat{r}) \leq 63$

  - this is the range of possible thresholds identified as "range of NE thresholds" in the figure

  - i.e., range of values of r such that $0 < \bar{q}(\theta(r)) < 1$

- Note that this is just a translation to the discussion of equilibrium **rules of behavior** of our previous conclusions about the multiplicity of NE in the absence of cognitive noise

Further Explanation of Slide 48

- If instead a **noisy** internal representation of θ: player i must **infer** possible θ, and other's prob. of staying, from state r_i

- Given any true state θ, net reward to i from staying is

$$\Delta(p) \;=\; 16[p - \bar{q}(\theta)],$$

where we use the notation $\bar{q}(\theta) \equiv (\theta - 47)/16$ for the probability of other player's staying that is required to make it worthwhile for player i to stay

- Then given r_i, optimal for player i to stay if and only if

$$\mathrm{E}[\Delta(p)\,|r_i] \;>\; 0 \qquad \Leftrightarrow \qquad \mathrm{E}[p\,|r_i] \;>\; \mathrm{E}[\bar{q}(\theta)\,|r_i]$$

$$\Leftrightarrow \;\; \mathrm{Prob}(-i\,\text{stays}\,|r_i) \;>\; \mathrm{E}[\bar{q}(\theta)\,|r_i]$$

$$\Leftrightarrow \;\; \mathrm{Prob}(r_{-i} < \hat{r}_{-i}\,|r_i) \;>\; \mathrm{E}[\bar{q}(\theta)\,|r_i]$$

Further Explanation of Slide 48

- Hence optimal threshold $\hat{r}_i$ for player i is value such that

$$\mathrm{Prob}(r_{-i} < \hat{r}_{-i} \,|\, r_i = \hat{r}_i) \;=\; \mathrm{E}[\bar{q}(\theta) \,|\, r_i = \hat{r}_i]$$

- A symmetric **BNE** in threshold strategies is a threshold r^* such that $\hat{r}_{-i} = r^*$ makes it optimal to choose $\hat{r}_i = r^*$

— this is an r^* such that

$$\mathrm{Prob}(r_{-i} < r^* \,|\, r_i = r^*) \;=\; \mathrm{E}[\bar{q}(\theta) \,|\, r_i = r^*]$$

— i.e., an intersection of the curves

$$q(r) \;\equiv\; \mathrm{Prob}(r_{-i} < r \,|\, r_i = r) \qquad \text{and} \qquad \mathrm{E}[\bar{q}(\theta) \,|\, r_i = r]$$

shown in the figure

Further Explanation of Slide 48

- In the limit as cognitive noise of this kind becomes negligible:
  - $q(r)$ approaches constant value $1/2$ for all r [horizontal dashed line in figure: other's internal state equally likely to be above or below one's own]
  - $\mathrm{E}[\bar{q}(\theta)\,|r_i = r]$ approaches $\bar{q}(\theta(r))$ [downward sloping straight line, as shown, if $m(\theta)$ is linear]

- Hence in the case of **small enough** cognitive noise of this kind, the intersection must be **unique,** as shown
  - set of solutions for r^* as $\nu \to 0$ [single point labeled BNE] **not** the same as the set of solutions for r^* in the zero-noise model [entire interval marked by the curly bracket]

Implications of Slide 48

- In the unique equilibrium, the players' probability of staying will vary depending on θ:

$$\text{Prob}(\text{stay}\,|\theta) \;=\; \text{Prob}(r_i < r^*\,|\theta) \;=\; \Phi\left(\frac{r^* - m(\theta)}{\nu}\right),$$

  where now $\Phi(z)$ is the CDF of the standard normal distribution.

- This should be a decreasing function of θ (as seen in slide 51).

- If when the range of values of θ in the support of the prior is **wider,** the encoding function $m(\theta)$ must be **flatter** [range normalization], then the function $\text{Prob}(\text{stay}\,|\theta)$ should also be a **flatter** function of θ (as is also seen in slide 51).