UNIVERSITY OF CAMBRIDGE

Investments and Asset Pricing in a World of Satisficing Agents

Tony Berrada* Peter Bossaerts[o] Giuseppe Ugazio*

*University of Geneva [o]University of Cambridge

August 2024

The Contribution

- In 1955, Herbert Simon proposed that economic agents do not optimize, but instead satisfice [Simon, 1955].
- We provide a formal theory built on Reference-Model Based Adaptive Control (MRAC) in robust control engineering.
- The MRAC agent aims to interact with markets in order to produce return distributions that minimize surprise with respect to a desired (target) reference distribution.
- The satisficing agent mostly acts "as if" optimizing, but we discover important – and realistic – deviations; asset pricing predictions change accordingly.

Introductory Remarks

1.
There is little work on formally modeling satisficing; Caplin et al. [2011] interpret satisficing as sequential search towards a level of reservation utility; Murawski and Bossaerts [2016] shows that humans continue searching way beyond the point where value increases, which casts doubt on the search-towards-reservation-utility hypothesis.

2.
There are various ways to motivate MRAC. Best is to motivate MRAC as in engineering: robustness for a nonstationary world.

MRAC: A Standard Textbook

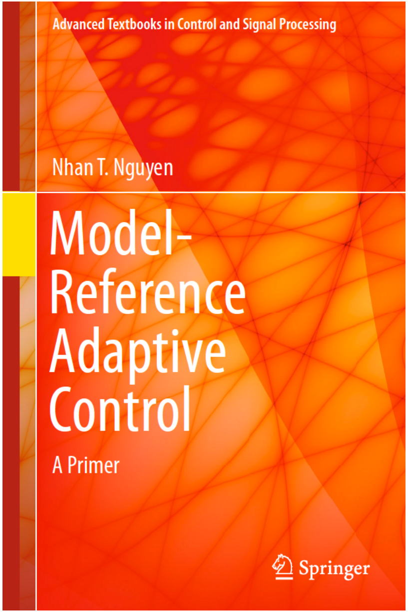

1. The Theory: Setup

Traditional Control Theory (As in Reinforcement Learning/Computational Neuroscience; and Core of AI)

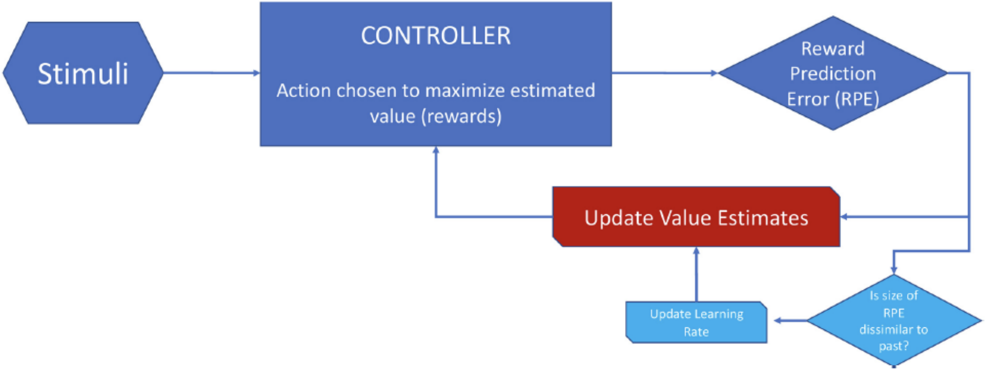

Value = "Q Value" and updating is based on "Temporal Difference" (TD Learning)

MRAC (Model-Reference (Based) Adaptive Control)

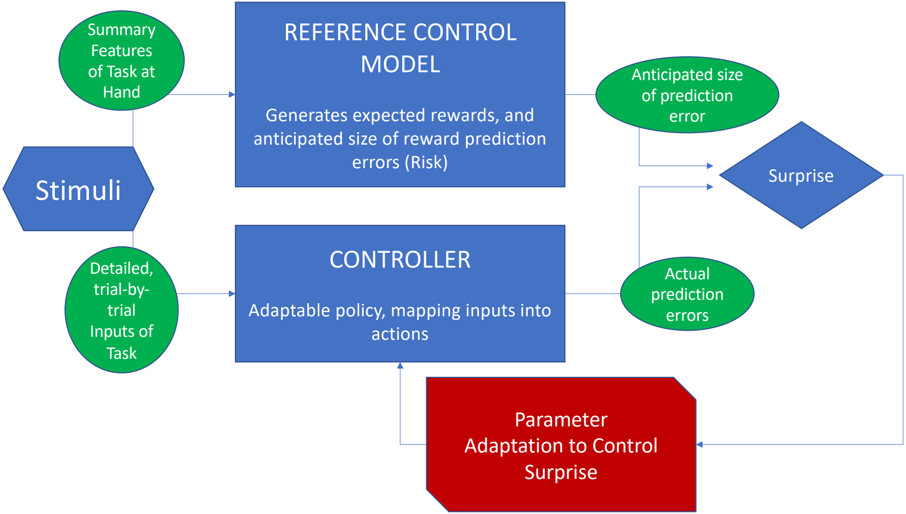

Situating MRAC in Behavioral Economics & Neuroeconomics

Behavioral economists and neuroeconomists have recently embraced the use of reference models, but for a different reason:

- Decoding information under neural constraints: Yamada et al. [2018], Azeredo da Silveira and Woodford [2019], Polanía et al. [2019], Vieider [2024].
- Value encoding under cognitive constraints: [Louie et al., 2013, Glimcher and Tymula, 2023, Payzan-LeNestour and Woodford, 2022, Frydman and Jin, 2022, Glimcher, 2022].
- Extensions of variational Bayes analysis (reference model = "recognition model"): Samuelson and Steiner [2024], Bossaerts and Rayo [2023].

Roughly speaking, these approaches do not concern robustness, but information loss.

MRAC and Investments

- We will not be interested in adaptation dynamics *per se*, but in the nature of the control, and hence, the choices, that obtain *at a particular point in time, given beliefs at that time*. Why we can do this? Most of our results hold for large classes of beliefs. For testing the theory, it suffices to identify which class the agent's beliefs belong to when she is making choices.
- There is a second dimension in which we simplify analysis: the principal has full control over the investment; the principal directly controls the optimal investment policy given the desire to minimize expected surprise

MRAC and Investments

Define Expected Surprise S as traditional in engineering:

$$S = (E(\tilde{R} - \mu)^2 - \sigma^2)^2.$$

$\tilde{R}$ is the return on the chosen portfolio, E is taken with respect to the perceived ("true") distribution, and μ and σ are the desired return mean and volatility, respectively. The agent chooses an investment strategy with (true) volatility x and expected return y.

$$\begin{aligned} S =& (E(\tilde{R} - y + y - \mu)^2 - x^2 + (x^2 - \sigma^2))^2 \\ =& ((y - \mu)^2 + (x - 0)^2 - \sigma^2))^2. \end{aligned} \tag{1}$$

Graphical Representation

Lemma

The set of portfolios which generate surprise $S = 0$ is a circle with center $(0, \mu)$ and radius σ. We label it the **0-surprise circle***.*

The 0-Surprise Circle

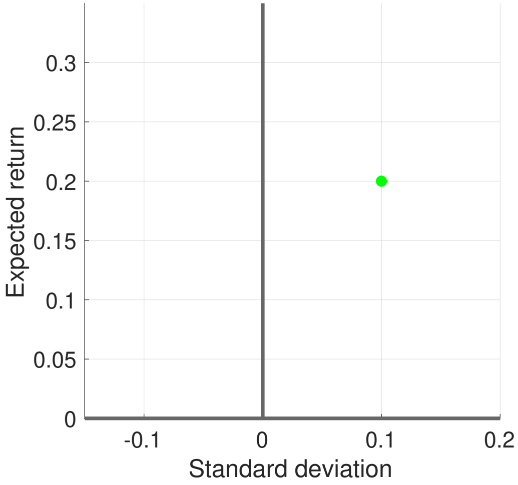

- Agent's desired (μ, σ)

The 0-Surprise Circle

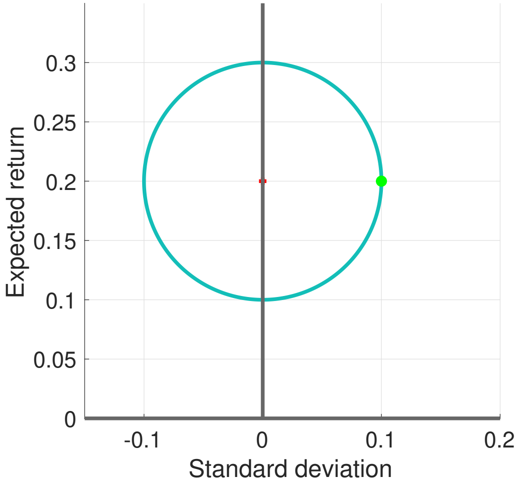

- Agent's desired (μ, σ)
- $0-$surprise circle

2.
Theoretical Predictions

Case 1: Unique Solution

Define θ = Market Sharpe Ratio (maximal mean excess return per unit of volatility)

Proposition

When the efficient frontier is such that

$$0 < \theta \leq \frac{y_{tan} - R_F}{x_{tan}}$$

where

$$y_{tan} = \frac{\mu^2 - \mu R_F - \sigma^2}{\mu - R_F}$$
$$x_{tan} = \sqrt{(y_{tan} - \mu)(y_{tan} - R_F)}$$

the surprise minimizing portfolio is at the tangency point of an ε-**surprise circle** *and EF, and is unique.*

Case 1: Unique Solution

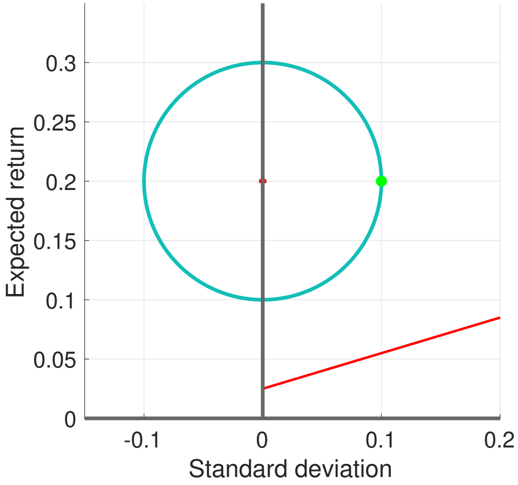

- Agent's desired (μ, σ)
- 0−surprise circle
- Efficient frontier *EF* (Capital Market Line; slope $= \theta$; intercept $=$ riskfree rate R_F)

Case 1: Unique Solution

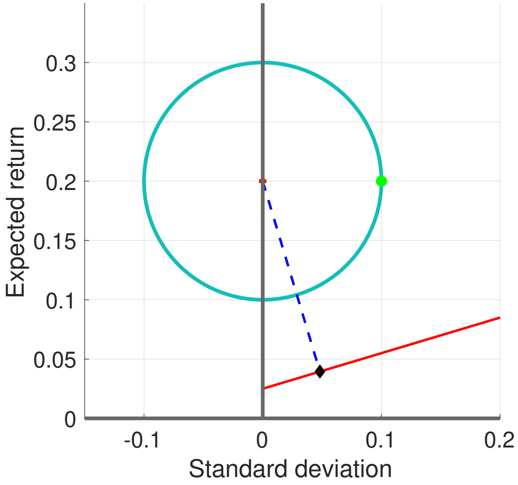

- Surprise-minimizing choice: orthogonal projection on *EF*

Case 1: Unique Solution

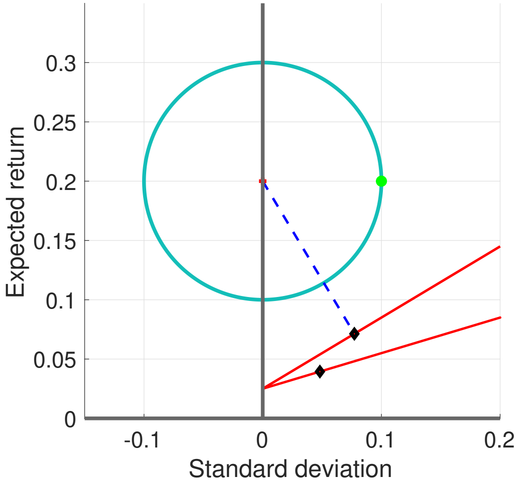

- Orthogonal projection on *EF*

Case 1: Unique Solution

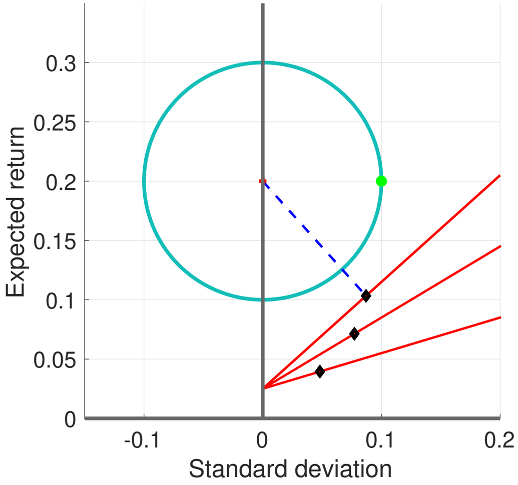

- Orthogonal projection on *EF*

Case 1: Unique Solution

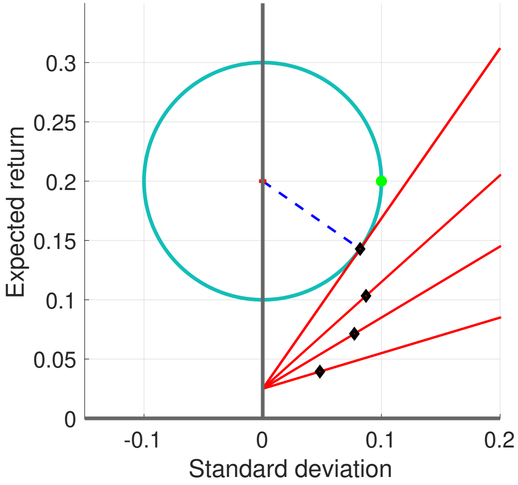

- Orthogonal projection on *EF*
- Tangency

Case 1: Unique Solution – *as if* MV-optimizers...

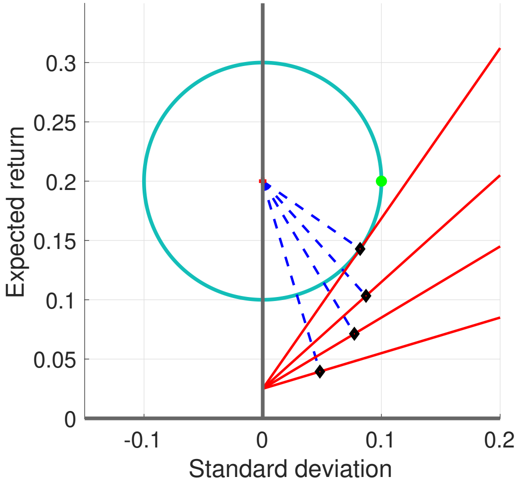

- Efficient allocations
- Market portfolio is efficient
- CAPM must hold

This happens when there is low tolerance for volatility

Case 2: Multiple Solutions

Proposition

When the efficient frontier is such that

$$\theta > \frac{y_{tan} - R_F}{x_{tan}}$$

where

$$y_{tan} = \frac{\mu^2 - \mu R_F - \sigma^2}{\mu - R_F}$$
$$x_{tan} = \sqrt{(y_{tan} - \mu)(y_{tan} - R_F)}$$

there exist multiple surprise minimizing portfolios. They are on the arc of the **0-surprise circle** *inferior to EF and superior to nEF.*

Case 2: Multiple Solutions

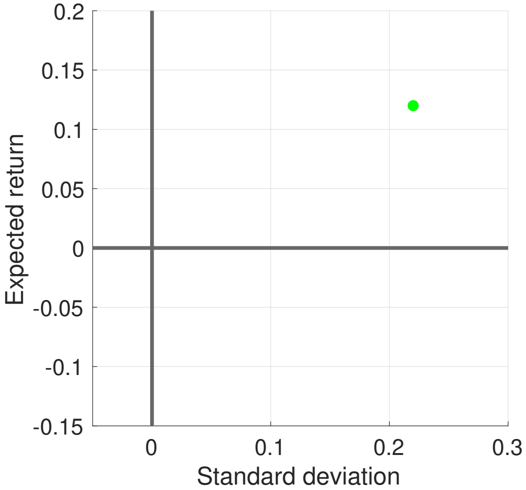

- Agent's desired (μ, σ)

Case 2: Multiple Solutions

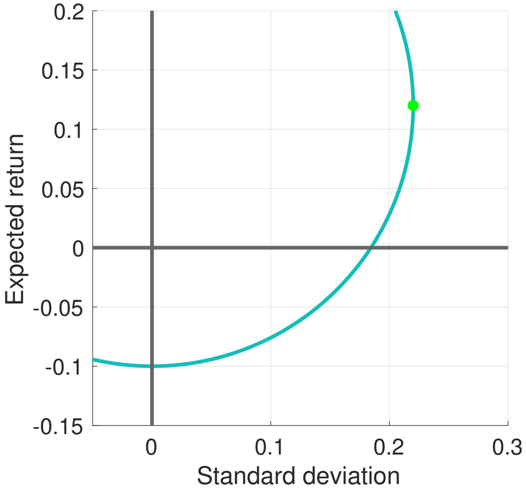

- Agent's desired (μ, σ)
- $0-$surprise circle

Case 2: Multiple Solutions

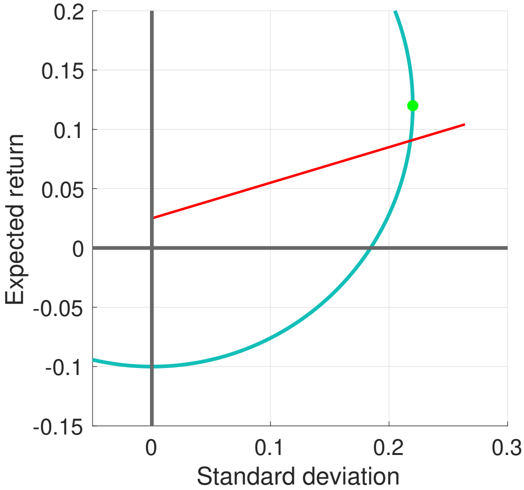

- Agent's desired (μ, σ)
- $0-$surprise circle
- Efficient frontier *EF*: Intersects with 0-surprise circle

Case 2: Multiple Solutions and non-MV behavior

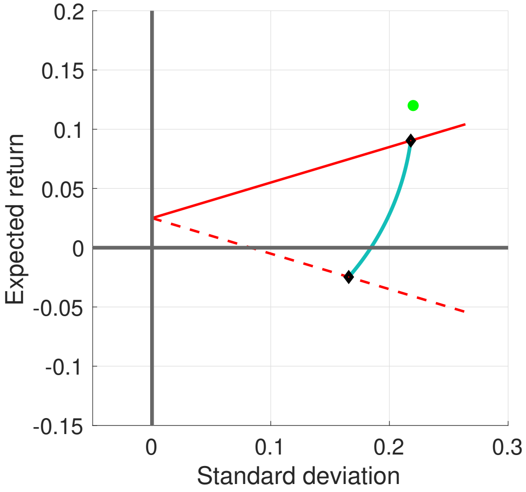

- Efficient frontier *EF*
- Inefficient frontier *nEF*
- Agent chooses any portfolio on arc between *EF* and *nEF*
- CAPM fails generically

This happens when there is high tolerance for volatility; "irrational exuberance" [Shiller, 2015]?

Why would anyone insist on buying an inefficient portfolio if we can do better?

ROBUSTNESS!

Excessive specialization for one environment may prove disastrous when the environment changes (too slow to adapt). MRAC modeling is to ensure one's reference model affords converging adaptation in a wide variety of contexts. One notable failure: X-15 crash, 1967.

Case 2: Limits to Risk Taking

Corollary

When the desired pair (μ, σ) satisfies the conditions of proposition 3, the surprise minimizing portfolios have risk inferior to the desired level of risk σ.

As a result, there may be equilibrium existence problems, as in credit rationing [Stiglitz and Weiss, 1981].

Risk Taking Despite Zero Risk Premium

Corollary

When the desired Sharpe ratio is strictly inferior to one, i.e. $\frac{\mu - R_F}{\sigma} < 1$*, the demand for risky assets is generically positive when the risk premium is 0.*

The mean-variance trade-off is only in the mind of the (financial) economist; the MRAC agent is not averse to risk, on the contrary; she is averse to surprise, i.e., to risk and return that is "unusual," and absence of risk is surprising.

Demand Analysis

N (number of risky assets) $= 1$.
(The agent trivially holds a mean-variance optimal portfolio)

Demand as a Function of Price

Expected vol s and return m change with price; Sharpe ratio of risky security (location: plus sign) decreases as price increases, as below

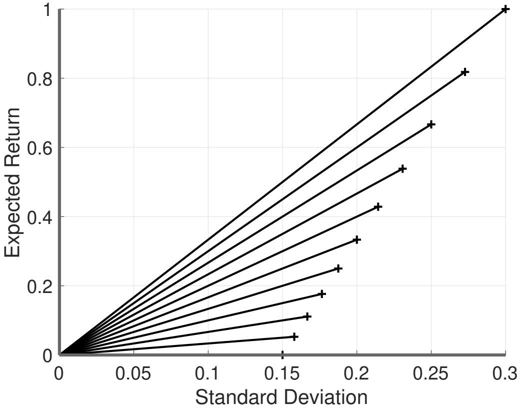

Demand as a Function of Price

Proposition

Let P_b denote the price boundary level defined as $P_b = (M - \Sigma \frac{y_{tan} - R_F}{x_{tan}})(1 + R_F)$. The demand function of the risky asset for the surprise minimizing agent with desired (μ, σ) is given by

$$\overline{d}(P) = \frac{(\mu - R_F)\theta(P)}{1 + \theta(P)^2} \frac{P}{\Sigma}$$

when $P \geq P_b$, and, when $P < P_b$, by

$$\underline{d}(P) = \frac{(\mu - R_F)\theta(P) \pm \sqrt{\Delta(P)}}{1 + \theta(P)^2} \frac{P}{\Sigma},$$

where $\theta(P) = \frac{\frac{M}{P} - 1 - R_F}{\frac{\Sigma}{P}}$ and $\Delta(P) = \theta(P)^2(\mu - R_F)^2 - \left((\mu - R_F)^2 - \sigma^2\right)(1 + \theta(P)^2)$.

Demand as a Function of Price

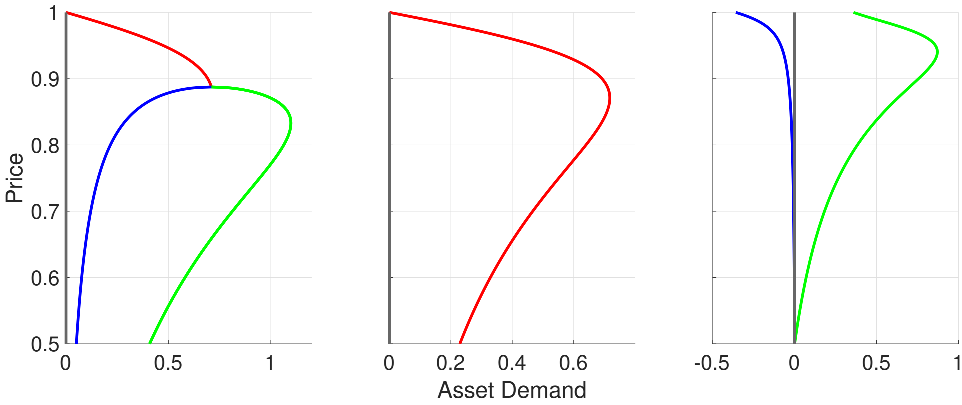

Left: desired mean return (μ) equals 0.25, desired volatility (σ) equals 0.20; middle: $\mu = 0.25$, $\sigma = 0.05$; right: $\mu = 0.14$, $\sigma = 0.15$ (notice demand for risk at price $= 1$)

3. Conclusion

Conclusion

We have presented a mathematical model of a satisficing agent based on a core concept in robust control theory, namely, Model Reference Based Adaptive Control (MRAC). The agent aims at minimizing expected surprise when targeting a particular expected return and return volatility. The agent often ends up choosing *as if* mean-variance optimizing, but the situations in which she does not provide a unique opportunity to re-visit the many empirical anomalies recorded since the emergence of the first formal asset pricing model, the CAPM.

Future Work

1. Bring in adaptation (will require extending traditional RMAC mathematics!); a simple version is in Bossaerts [2018].
2. Allow the "controller" to be at arms length from investor; the controller can be a traditional reward optimizer, like a mean-variance optimizer; a simple version is in Bossaerts [2018].
3. The framework can also be used to represent delegated portfolio management, including delegation to an algorithmic (robotic) trader.
4. Experimental tests.
5. Neurobiological foundations: re-interpret *anterior insula* neural activation in response to surprise using sEEG; change behavior through *anterior insula* stimulation; based on application of MRAC (adaptive version) for a simple stochastic target-centering game [Bossaerts, 2018] (with Fabienne Picard [HUG] and Nina Sooter [UNIGE]).

References

R. Azeredo da Silveira and M. Woodford. Noisy memory and over-reaction to news. In *AEA Papers and Proceedings*, volume 109, pages 557–561. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203, 2019.

P. Bossaerts. Formalizing the function of anterior insula in rapid adaptation. *Frontiers in integrative neuroscience*, 12:61, 2018.

P. Bossaerts and L. Rayo. Modeling behavior when inference is complex. *Cambridge Working Paper*, 2023.

A. Caplin, M. Dean, and D. Martin. Search and satisficing. *American Economic Review*, 101(7):2899–2922, 2011.

C. Frydman and L. J. Jin. Efficient coding and risky choice. *The Quarterly Journal of Economics*, 137(1):161–213, 2022.

P. W. Glimcher. Efficiently irrational: deciphering the riddle of human choice. *Trends in cognitive sciences*, 26(8):669–687, 2022.

P. W. Glimcher and A. A. Tymula. Expected subjective value theory (esvt): a representation of decision under risk and certainty. *Journal of Economic Behavior & Organization*, 207:110–128, 2023.

K. Louie, M. W. Khaw, and P. W. Glimcher. Normalization is a general neural mechanism for context-dependent decision making. *Proceedings of the National Academy of Sciences*, 110(15): 6139–6144, 2013.

C. Murawski and P. Bossaerts. How humans solve complex problems: The case of the knapsack problem. *Scientific reports*, 6:34851, 2016.

E. Payzan-LeNestour and M. Woodford. Outlier blindness: A neurobiological foundation for neglect of financial risk. *Journal of Financial Economics*, 143(3):1316–1343, 2022.

R. Polanía, M. Woodford, and C. C. Ruff. Efficient coding of subjective value. *Nature neuroscience*, 22(1):134–142, 2019.

L. Samuelson and J. Steiner. Constrained data-fitters. *Yale university working paper*, 2024.

R. J. Shiller. Irrational exuberance. In *Irrational exuberance*. Princeton university press, 2015.

H. A. Simon. A behavioral model of rational choice. *The quarterly journal of economics*, pages 99–118, 1955.

J. E. Stiglitz and A. Weiss. Credit rationing in markets with imperfect information. *The American economic review*, 71(3):393–410, 1981.

F. M. Vieider. Decisions under uncertainty as bayesian inference on choice options. *Management Science*, 2024.

H. Yamada, K. Louie, A. Tymula, and P. W. Glimcher. Free choice shapes normalized value signals in medial orbitofrontal cortex. *Nature communications*, 9(1):162, 2018.