Part-Time Bayesians:
Behavioral Heterogeneity in Financial Decisions

Michele Garagnani*[1]

[1]University of Zurich

If a stock went up by 60% last year, is it a good idea to invest on it now? If following a fund manager's advice worked well last time, is it a good idea to trust her next time? If the competition's profits are higher than those of your firm, should you fire the CEO? All these situations, and many others, share two elements. First, the outcome of a past decision is described in terms of wins or losses. Second, a new decision has to be made, where the past behavior could be repeated. In such cases, humans (and other animals) are hardwired to blindly repeat choices if they led to success in the past and change them otherwise. This tendency goes by the name of reinforcement behavior and has been shown to be one of the most fundamental processes underlying learning and decision making in the brain (e.g., Schultz, Dayan, and Montague, 1997; Holroyd and Coles, 2002), to the extend that reinforcement-based models are the dominant paradigm in neuroscience (e.g. Daw, Niv, and Dayan, 2005). Decisions in management and finance might be often affected by this tendency, because they typically rely on information that includes explicit win-lose feedback: gains and losses, success and failure, beating the competition or not, etc. Indeed, both individual investors and mutual fund holdings are more likely to repurchase stocks that were previously sold for a gain rather than for a loss (Strahilevitz, Odean, and Barber, 2011; Du, Niessen-Ruenzi, and Odean, 2018), and new corporate acquisitions are strongly influenced by whether recent, previous acquisitions were successful or not (Bharath, Cho, and Choi, 2019).

The problem with reinforcement behavior is that it is a simple rule of thumb which can and often does contradict optimal, rational prescriptions. When confronted with new information on an uncertain event (whether a stock will go up or down, whether a CEO is the best choice or not), there is one and only one way to update previous beliefs: the mathematical result known as Bayes' rule. This result tells us how we should we adjust our previous beliefs in the face of new evidence based, by precisely weighing that evidence with the one contained in the previous beliefs. The importance of understanding and properly applying Bayes' rule cannot be overstated, as it is the normative principle capturing how optimal decisions should be made in the financial

*Zurich Center for Neuroeconomics (ZNE), Department of Economics, University of Zurich. Blümlisalpstrasse 10, CH-8006 Zurich, Switzerland.

domain, but also in everyday's life. Yet, there is overwhelming evidence showing that human decision makers have a very limited grasp of probabilities and, especially, of how beliefs should be updated in the face of new information (e.g., Kahneman and Tversky, 1972; Grether, 1980; Camerer, 1987). Thus, it is not surprising that humans rely on simple rules of thumb as "win-stay, lose-shift" reinforcement behavior whenever the shoe appears to fit.

This is not to say that reinforcement is always wrong. In particular, in financial markets, there might be a number of confounds partially explaining why repeating previously-successful behavior might be a good idea. But reinforcement can go wrong in spectacular ways. In economic and financial contexts, in particular, reinforcement might lead to "irrational" behavior such as an excessive focus on past performance (*outcome bias*)(Baron and Hershey, 1988; Dillon and Tinsley, 2008; Choi et al., 2009; Ater and Landsman, 2013) where, for instance, the performance of managers who just "got lucky" is in practice misattributed and information on market conditions, fundamentals, etc, is neglected in favor of the pure win-loss outcome.

To isolate the problem, in my research I rely on laboratory experiments where all the relevant features of financial and managerial decisions can be reproduced, but all undesired confounding factors can be eliminated. Specifically, in my work presented in Alós-Ferrer and Garagnani (2021), participants made decisions where previous beliefs on an uncertain event could be updated on the basis of new information, but that information was whether a previous choice had resulted in a monetary payoff or not. Hence, as in many financial decisions, the information derived from previous choices included win-loss feedback. The task was specifically designed to ensure that rational decisions as derived from Bayes' rule sometimes agreed with the prescriptions of a win-stay, lose-shift rule, but sometimes clashed with those. Therefore I can disentangle these two behavioral tendencies (and other economically-relevant ones, as decision inertia).

Using statistical classification techniques related to machine learning technique, I am able to determine which rules every decision maker relied the most on, and whether they switched between different ways of thinking. The results are striking. Less than 50% of participants can be classified as mostly following Bayes' rule, around 30% are essentially reinforcers, blindly repeating the same choices if they led to success in the past and changing them otherwise, and the rest follow other, non-rational rules. Even more strikingly, this heterogeneity between individuals is coupled with a large heterogeneity within individuals, meaning that most of the participants switched across different approaches to solve the same problem during the course of the experiment. That is, even subjects who were mostly adopting the correct way of thinking were actually only "part-time Bayesians." Therefore, when it comes to belief updating, and despite the fact that people are often treated as homogeneous in standard economic models, we are led to the conclusion that one size does not fit all for describing actual behavior, but, even worse, one size might not even fit one, because people rely on multiple decision rules for the same problem.

By varying incentives across individuals, the experiment also demonstrated a worrying "reinforcement paradox." People who behaved mostly according to reinforcement performed worse when paid more. This is because the win-lose cue that triggers reinforcement behavior was even more relevant to them (as it involved more money) when incentives were higher, and hence swayed them away even further. This paradox, which reflects observations from neuroscience (Achtziger et al., 2015) might be very relevant for financial and managerial contexts, where information is often monetary and communicated in terms of win-lose feedback. Counter-intuitively, paying more to elicit higher performance might actually backfire by eliciting higher levels of reinforcement behavior. Needless to say, this is in stark contrast with many employee compensation plans (e.g., repeated performance boni) and with the standard economic view that higher monetary incentives induce higher effort and improve performance. As it is so often the case, reality is less straightforward, because people are highly heterogeneous in how they update their beliefs. Therefore, in financial contexts where beliefs need to be frequently updated, ranging from corporate decisions to personal finance, some people might perform better when paid more, but others might feel they are putting more effort just to obtain worse results. Thus, these results speak in favor of re-thinking the coupling between incentives and performance in financial and managerial decisions to better understand their relation and how to promote the latter when the reactions to the former are highly heterogeneous.

References

Achtziger, Anja, Carlos Alós-Ferrer, Sabine Hügelschäfer, and Marco Steinhauser. 2015. "Higher Incentives Can Impair Performance: Neural Evidence on Reinforcement and Rationality." *Social Cognitive and Affective Neuroscience* 10 (11):1477–1483.

Alós-Ferrer, Carlos and Michele Garagnani. 2021. "Part-Time Bayesians: Incentives and Behavioral Heterogeneity in Belief Updating." Working Paper, University of Zurich.

Ater, Itai and Vardit Landsman. 2013. "Do Customers Learn From Experience? Evidence from Retail Banking." *Management Science* 59 (9):2019–2035.

Baron, Jonathan and John C. Hershey. 1988. "Outcome Bias in Decision Evaluation." *Journal of Personality and Social Psychology* 54 (4):569–579.

Bharath, Sreedhar T., DuckKi Cho, and Lyungmae Choi. 2019. "Win-Stay, Lose-Shift: A Strategy of Serial Acquirers." In *European Financial Management Association 2019 Annual Meetings*.

Camerer, Colin F. 1987. "Do Biases in Probability Judgment Matter in Markets? Experimental Evidence." *American Economic Review* 77 (5):981–997.

Choi, James J., David Laibson, Brigitte C. Madrian, and Andrew Metrick. 2009. "Reinforcement Learning and Savings Behavior." *Journal of Finance* 64 (6):2515–2534.

Daw, Nathaniel D., Yael Niv, and Peter Dayan. 2005. "Uncertainty-Based Competition Between Prefrontal and Dorsolateral Striatal Systems for Behavioral Control." *Nature Neuroscience* 8:1704–1711.

Dillon, Robin L. and Catherine H. Tinsley. 2008. "How Near-Misses Influence Decision Making Under Risk: A Missed Opportunity for Learning." *Management Science* 54 (8):1425–1440.

Du, Mengqiao, Alexandra Niessen-Ruenzi, and Terrance Odean. 2018. "Stock Repurchasing Bias of Mutual Funds." SSRN 3247066.

Grether, David M. 1980. "Bayes Rule as a Descriptive Model: The Representativeness Heuristic." *Quarterly Journal of Economics* 95:537–557.

Holroyd, Clay B. and Michael G. Coles. 2002. "The Neural Basis of Human Error Processing: Reinforcement Learning, Dopamine, and the Error-Related Negativity." *Psychological Review* 109:679–709.

Kahneman, Daniel and Amos Tversky. 1972. "Subjective Probability: A Judgment of Representativeness." *Cognitive Psychology* 3:430–454.

Schultz, Wolfram, Peter Dayan, and P. Read Montague. 1997. "A Neural Substrate of Prediction and Reward." *Science* 275 (5306):1593–1599.

Strahilevitz, Michal Ann, Terrance Odean, and Brad M. Barber. 2011. "Once Burned, Twice Shy: How Naive Learning, Counterfactuals, and Regret Affect the Repurchase of Stocks Previously Sold." *Journal of Marketing Research* 48:S102–S120.