

Small Differences in Experience Bring Large Differences in Performance

Sheen S. Levine¹, Charlotte Reypens^{1,2}

¹ University of Texas, Dallas, USA

² University of Antwerp, Belgium

Abstract. In many life situations, people choose sequentially between repeating a past action in expectation of a familiar outcome (exploitation), or choosing a novel action whose outcome is largely uncertain (exploration). For instance, in each quarter, a manager can budget advertising for an existing product, earning a predictable boost in sales. Or she can spend to develop a completely new product, whose prospects are more ambiguous. Such decisions are central to economics, psychology, business, and innovation; and they have been studied mostly by modelling in agent-based simulations or examining statistical relationships in archival or survey data. Using experiments across cultures, we add unique evidence about causality and variations. We find that exploration is boosted by three past experiences: When decision-makers fall below top performance; undergo performance stability; or suffer low overall performance. In contrast, individual-level variables, including risk and ambiguity preferences, are poor predictors of exploration. The results provide insights into how decisions are made, substantiating the micro-foundations of strategy and assisting in balancing exploration with exploitation.

Keywords: Exploration, Exploitation, Decision Making, Experiment, Protocol Analysis, Cross-culture

In many life situations – R&D investments, market entry, military campaigns, romantic choices – a decision-maker chooses an action, receives feedback, and then chooses again. The choice ranges from repeating a past action in expectation of a familiar outcome (exploitation) to a novel action whose outcome is largely uncertain (exploration). For instance, in each quarter, a manager can budget more advertising for an existing product, expecting familiar (but uncertain) sales figures; or he can spend to develop a revised version of the product, whose sales prospects are more uncertain; or he can invest in a completely new product, where prospects are even more uncertain. The optimal action is not obvious: Probabilities are unknown, feedback is ambiguous. Such decisions have been discussed across domains and species, from bees and birds foraging to organizations searching for innovations (for reviews, see [1,2]). We examined – empirically – how people decide in such situations. To do that, we created a behavioral task – an oil exploration game in which participants earn money by searching an unfamiliar landscape (cf. [3]). A participant chooses a spot for

drilling and then discovers the quantity of oil it contains. The participant can keep drilling in the same spot, earning the same quantity of oil. Or he can choose a nearby location, which likely has a similar oil quantity. Or he can jump across the landscape to a faraway location, where the oil quantity is likely very different. The participant repeats the choice a fixed number of times. When the game ends, the oil he found is converted to cash, which is paid to him. The task faithfully represents the important features of an exploration-exploitation situation: The landscape is rugged, containing “peaks” and “valleys” of oil, but there is no map that describes the terrain – a decision maker discovers it by experiencing it. Because probabilities are unknown, optimization is impossible [4,5,6,7]. And since information accumulates only with experience, initial steps are necessarily random [8]. The seeker has only limited resources, so he can sample only a fraction of the entire landscape ([9]). We studied how people decide in exploration-exploitation in four studies. First, we conducted one-on-one sessions, where we collected quantitative data as well as verbal accounts from the participants, describing their decision-making process [10]. Second, we conducted laboratory sessions in the U.S. using a web-based version of the game. Third, also using the web-based instrument, we collected data from workers in a labor market [11,12]. Fourth, to ascertain the robustness of the findings, we conducted laboratory sessions in Russia, a country whose history, culture, and institutions differ from those of the U.S. [13]. Across all studies, we find that exploration is driven more by immediate experience, less by individual characteristics. Three situations boost it: When a decision-maker falls below his or her top performance; when he or she experiences performance stability; or when he or she suffers low overall performance. Exploitation is increased by the reverse experiences: exceeding top performance, experiencing performance variance, and enjoying high overall performance. Individual traits, such as risk and ambiguity preferences, are poor predictors of exploration. These experiences have similar effects in all of the studies. Behavior is strongly influenced by experience, so two identical players that face the exact same landscape can undergo completely different experiences and end up with a wide gap in performance, all due to random differences in their early choices. In everyday life, we often attribute differences in performance to traits – “she is brilliant,” “this manager is incompetent” – and not to experience. But the results here suggest otherwise: History matters.

References

1. Hills, T.T., Todd, P.M., Lazer, D., Redish, A.D., Couzin, I.D.: Exploration versus exploitation in space, mind, and society. *Trends in Cognitive Sciences* **19**(1) (2015) 46–54
2. Mehlhorn, K., Newell, B.R., Todd, P.M., Lee, M.D., Morgan, K., Braithwaite, V.A., Hausmann, D., Fiedler, K., Gonzalez, C.: Unpacking the exploration-exploitation tradeoff: A synthesis of human and animal literatures. *Decision (Washington)* **2**(3) (2015) 191–215
3. Mason, W., Watts, D.J.: Collaborative learning in networks. *Proceedings of the National Academy of Sciences* **109**(3) (2012) 764–769

4. Alchian, A.: Uncertainty, evolution, and economic theory. *Journal of Political Economy* **58** (1950)
5. Gittins, J.C., Gittins, J.C.: Bandit processes and dynamic allocation indices. *Journal of the Royal Statistical Society, Series B* (1979) 148–177
6. Kauffman, S., Levin, S.: Towards a general theory of adaptive walks on rugged landscapes. *Journal of Theoretical Biology* **128**(1) (1987) 11 – 45
7. Nagel, R., Vriend, J.N.: An experimental study of adaptive behavior in an oligopolistic market game. *Journal of Evolutionary Economics* **9**(1) (1999) 27–65
8. Winter, S.G.: Purpose and progress in the theory of strategy: Comments on gavetti. *Organization Science* **23**(1) (2012) 288–297
9. March, J.G.: Exploration and exploitation in organizational learning. *Organization Science* **2**(1) (1991) 71–87
10. Ericsson, K.A., Simon, H.A.: Protocol analysis: Verbal reports as data (Revised ed.). Cambridge, MA: MIT Press (1993)
11. Horton, J.J., Rand, D.G., Zeckhauser, R.J.: The online laboratory: conducting experiments in a real labor market. *Experimental Economics* **14**(3) (2011) 399–425
12. Paolacci, G., Chandler, J., Ipeirotis, P.G.: Running experiments on amazon mechanical turk. *Judgment and Decision Making* (2010) 411–419
13. Henrich, J., Heine, S.J., Norenzayan, A.: Most people are not WEIRD. *Nature* **466**(7302) (July 2010) 29