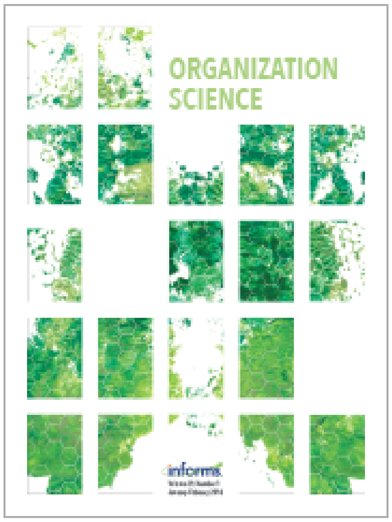

Organization Science

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

# How Do Performance Goals Influence Exploration-Exploitation Choices?

Marlo Raveendran, Kannan Srikanth, Tiberiu Ungureanu, George L. Zheng

informs
https://pubsonline.informs.org/journal/orsc

ORGANIZATION SCIENCE
*Articles in Advance*, pp. 1–23
ISSN 1047-7039 (print), ISSN 1526-5455 (online)

# How Do Performance Goals Influence Exploration-Exploitation Choices?

**Marlo Raveendran,[a,*] Kannan Srikanth,[b] Tiberiu Ungureanu,[c] George L. Zheng[d]**

[a] School of Business, Area of Management, University of California, Riverside, California 92521; [b] Department of Management and Human Resources, Fisher College of Business, The Ohio State University, Columbus, Ohio 43210; [c] Department of Management, Walker College of Business, Appalachian State University, Boone, North Carolina 28608; [d] Strategy, Innovation, and Entrepreneurship Department, College of Business, Shanghai University of Finance and Economics, Shanghai 200433, China

*Corresponding author

**Contact:** marlo.raveendran@ucr.edu, https://orcid.org/0000-0002-3117-5749 (MR); srikanth.18@osu.edu, https://orcid.org/0000-0002-5900-4503 (KS); ungureanuts@appstate.edu, https://orcid.org/0000-0001-6820-9298 (TU); hflsgeorge01@gmail.com, https://orcid.org/0000-0003-1370-2557 (GLZ)

**Received:** September 30, 2019
**Revised:** July 6, 2020; February 28, 2021; April 29, 2022; March 26, 2023; July 14, 2023
**Accepted:** September 15, 2023
**Published Online in Articles in Advance:** October 18, 2023

https://doi.org/10.1287/orsc.2019.13311



**Abstract.** Employees in organizations are frequently subject to performance goals such as sales or publication targets. However, often employees do not know what actions will allow them to meet these goals. To perform such tasks effectively, employees need to explore to quickly learn from experience which among the available alternatives offers the higher reward potential, so that they can concentrate subsequent efforts on exploiting it. Prior work models such explore-exploit problems as an adaptive learning process, where employees sequentially sample various options and learn from feedback. However, we currently do not know *how performance goals influence this adaptive learning process*. We argue that performance goals influence the adaptive learning process by modifying how feedback is perceived. Individuals subject to challenging goals are more likely to interpret feedback from poor alternatives as failures. Therefore, they quickly develop high belief strength that the inferior alternative is worse than the superior alternative, enabling them to reduce "useless exploration," but also making them slow to adapt to environmental shocks. We test our predictions in a series of laboratory experiments and find that decision makers subject to challenging goals exploit *more* (relative to those with moderate goals). We also show that such an exploitation focus, while beneficial in stable environments, is detrimental in unstable ones. Our finding that challenging performance goals improve performance in *learning* tasks stands in contrast to prior findings that such goals inhibit performance in *search* tasks, an insight that warrants further study to improve our understanding of goal setting in the knowledge economy.

**Funding:** K. Srikanth was supported by SMU Seed Funding Grant for the initial version of this paper.
**Supplemental Material:** The e-companion is available at https://doi.org/10.1287/orsc.2019.13311.

**Keywords:** goal setting • organizational learning • incentives • hot stove effect



Performance goals are ubiquitous in organizations. Employees work toward sales or revenue targets, CEOs aim to beat analyst expectations, fund managers aim to beat the benchmark index. Even in highly uncertain tasks such as R&D, engineers and scientists are often held to patenting and publishing targets. Although in all these instances employees are exhorted to perform as well as they can, falling short of these explicitly set goals usually has significant negative consequences for employees, such as foregone bonuses and promotion opportunities. Yet employees frequently do not know which actions will lead to outcomes that exceed their goals. For example, should a sales agent trying to beat their target invest effort in cultivating existing client relationships or forge new ones? To perform such tasks effectively, employees do not just need to exert more effort, they need to learn which of the available alternatives is more attractive, to decide where they should concentrate their efforts on. In such learning problems, when employees face considerable uncertainty in their decision making about which actions to choose from among various alternatives to meet their goals, they have to strike a balance between exploring new opportunities and exploiting existing knowledge (March 1991, 1996). Although performance goals encourage employees to work harder (Locke and Latham 2006), we currently do not know whether they help employees work smarter, that is, quickly learn which alternative is superior. In this paper we therefore study how performance goals influence employees' learning and therefore their explore-exploit decisions.

How to effectively make explore-exploit decisions is widely acknowledged as an important organizational problem (Denrell and March 2001, Denrell 2008, Lee and Puranam 2016). Such decisions are often modeled as employees choosing between alternatives with little prior knowledge about respective performance consequences (March 1996, Posen and Levinthal 2012). Employees have limited prior knowledge about alternatives especially in novel environments, where choices are ill-understood (cf: Song et al. 2019, Sang et al. 2020), and in changing environments, where prior experience is less useful in guiding future actions (Daw et al. 2006, Posen and Levinthal 2012).

To achieve high performance in such situations, employees engage in *adaptive learning* by repeatedly sampling the available alternatives and observing feedback (March 1996, Denrell and March 2001, Posen and Levinthal 2012). Since individuals try to reinforce successes and avoid failures, they are more likely to choose the options that previously resulted in better outcomes and avoid alternatives that previously resulted in worse outcomes. Thus, by repeated sampling, employees form beliefs about (i.e., learn) which alternative is superior. The stronger the agents' belief that one alternative is superior to another, the more likely they are to exploit (i.e., choose the alternative *they believe* to be the best) and less likely they are to explore (i.e., choose the alternative *they believe* to be inferior). For our sales agent, the higher the outcomes from discovering new customers compared with selling more to existing customers, the more pronounced the agent's preference for the former alternative compared with the latter. Yet, if the observed payoffs from these two alternatives are close, it will take more experience (sampling) before the agent can truly differentiate them, and more consistently exploit. In other words, with experience, individuals form beliefs about the relative attractiveness of the available alternatives and choose among them based on those beliefs (Sutton and Barto 1998, Daw et al. 2006, Cohen et al. 2007).[1]

How performance goals influence this adaptive learning process has not yet been investigated empirically. We argue that performance goals influence learning by changing how individuals interpret their received feedback. We theorize that performance goals divide the feedback space into successes and failures (e.g., Simon 1955). As a consequence, individuals with a challenging goal are more likely to categorize outcomes from inferior options as failures compared with those with moderate goals.[2] Since decision makers reinforce successes and avoid failures, individuals with challenging goals more quickly develop stronger beliefs that the inferior alternatives are truly inferior when compared with the superior alternatives, and therefore stop sampling them sooner (compared with individuals with moderate or do-your-best goals). While this greater belief strength is likely beneficial when the environment is stable, it can cause the individual to stick to their choices too long when the environment changes. Since individuals with challenging goals have developed stronger beliefs about the best and worst option, they will be slower in adapting to environmental shifts because they will continue to under-sample the previously worst option (which may have gotten better after the shock), and over-sample the previously best (but inferior post-shock) option.

Studying these processes in the field is exceedingly difficult (Gary et al. 2017). We therefore developed a series of laboratory studies to test our theory, in which participants chose among a limited number of alternatives with noisy payoffs. In our laboratory studies, we found that individuals with challenging goals (a) were more likely to have greater confidence in identifying the superior (and inferior) alternative; (b) which led them to choose the alternative they believed to be superior more often (i.e., exploit); and since they identified the superior alternative more confidently and exploited it sooner, (c) had higher performance, compared with those facing moderate goals. However, their performance suffered and temporarily fell below that of individuals with a moderate goal when the environment changed. We tested multiple types of goals to identify boundary conditions and found that goals improved performance only when they helped in identifying the relatively *inferior* options. Interestingly, goals designed to help reliably identify the *superior* option did not elicit similar behavior, suggesting that individuals were more vigilant about failures than successes.

Our study makes two contributions to theory. First, we contribute to the exploration-exploitation literature by showing the impact of explicit performance goals on adaptive learning, both of which are important organizational processes but surprisingly have not been jointly studied. We thus answer the call by Denrell (2008) to consider how aspiration levels influence learning rather than risk-taking. In doing so, we develop predictions on the specific mechanisms that link challenging goals to performance under stable and changing environments. Second, we add to the goal setting literature by exploring the conditions under which challenging versus moderate goals lead to higher performance in learning tasks which are particularly relevant in organizations, adding to the larger discussion on identifying boundary conditions for when goal setting improves performance. The finding that goal setting may hinder performance in search tasks but help performance in adaptive learning tasks is an important jumping-off point for future research in making goal setting more managerially relevant in the knowledge economy.[3]

## Theory

### The Impact of External Performance Goals on Task Performance Under Uncertainty

The goal-setting literature has extensively studied how different types of goals (typically, "do-your-best" or moderate goals versus challenging ones) influence task performance, mostly at the individual level. Summarizing

35 years of goal-setting research, Locke and Latham (2006) conclude that to the extent that individuals and groups are committed to goals and are able to attain them, "there is a positive linear relationship between goal difficulty and task performance" (p. 265). These scholars argue that setting challenging goals focuses attention and elicits effort and persistence, often by motivating search for high-performance strategies (Locke and Latham 2002, 2006). However, they add the important caveat that these results hold "in so far as performance is fully controllable," (Locke and Latham 2002, p. 706), that is, so long as task outcomes are solely a function of effort.

In exploring this boundary condition, scholars have considered two important contingencies (complex situations and unrealistically high goals) when effort and outcome may not be correlated, and therefore setting challenging goals may backfire. First, in complex decision-making situations, individuals do not understand how their decisions interact to produce outcomes. In these cases, setting challenging goals hurts individuals' ability to search for effective task completion strategies (Earley et al. 1989a, Kanfer and Ackerman 1989, Seijts and Latham 2005). Second, when goals are unrealistically high, goals may be more likely achievable by good luck rather than high effort. For example, studies have looked at individuals' one-time choice between a safe and a risky bet, when the goal is above the payoff for the safe bet. In such conditions, individuals revert to endogenous do-your-best goals (Locke and Latham 1990), or rely on good luck by choosing risky bets over safe ones (Heath et al. 1999, Larrick et al. 2009). Studies combining these two conditions have largely replicated these findings (Earley et al. 1989b, Gary et al. 2017).

However, empirical studies do not consider how goal setting influences learning under uncertainty when individuals have imperfect knowledge about the value of different alternatives. It is this problem we examine more closely, since it represents important decision situations in organizations. The canonical resource allocation problem where managers allocate resources among alternatives with unknown attractiveness is a good example of decision making under uncertainty (Bower 1970; March 1996, 2003). When COVID hit, vaccine researchers in pharmaceutical companies faced considerable uncertainty about which technology platform to build on—mRNA, viral vector, or weakened virus were among the available pathways. Firms invested in several options and relied on feedback to evaluate which ones to emphasize more. In other words, lack of knowledge due to novelty of the problem forced firms to rely on learning by doing; as new data about the viability of the different pathways emerged, managers had to decide between exploration and exploitation to create a viable vaccine. Would a challenging goal such as 90% efficacy for vaccines set by Johnson & Johnson[4] help vaccine developers make better decisions?

As this example illustrates, employees in firms need to make explore-exploit decisions when they do not accurately know the relative attractiveness of the options available to them—either because of novelty, or because environmental shocks made their prior knowledge a less reliable guide. Denrell and March (2001) theorize that individuals learn to prefer a safe alternative over a risky one under do-your-best goals, and that the magnitude of any performance goals may shift this preference. However, these predictions about how challenging versus moderate goals may influence choice between uncertain alternatives are not yet empirically tested. Moreover, current theory does not consider how performance goals influence the more difficult learning problem of choice between stochastic alternatives that are also changing (Daw et al. 2006), and the mechanisms underlying how goals influence adaptive learning need to be further developed.

### How Individuals Make Explore-Exploit Decisions

At the individual-level, scholarship on the exploration-exploitation tradeoff encompasses a considerable empirical literature in neuroscience and psychology that considers how individuals make decisions under uncertainty (see the reviews by Cohen et al. 2007, Toplak et al. 2010, Mehlhorn et al. 2015). This literature defines exploration and exploitation as follows: Decision makers *exploit* if they choose the option that they *believe* gives the highest payoff and *explore* otherwise, that is, if they select an option that they *believe* does not give the highest payoff (Daw et al. 2006, Laureiro-Martínez et al. 2014, Song et al. 2019).This literature conceptualizes the exploration-exploitation tradeoff in terms of two interrelated processes: (1) the *choice* process of selecting an alternative to explore to gain new knowledge, versus choosing the current best performing alternative to exploit existing knowledge; and (2) the *adaptive learning* process of translating feedback into beliefs about the relative attractiveness of the available alternatives (Sutton and Barto 1998, Daw et al. 2006, Cohen et al. 2007).

Learning from feedback is fundamental to this process. As the decision maker selects an alternative (makes a *choice*), for example, runs an experiment or engineers a protein, they receive some feedback about its efficacy (e.g., in preventing COVID). This feedback is now integrated with the feedback received from previous trials to form an overall impression about the promise of this technology platform (*adaptive learning)*. These beliefs about the relative attractiveness of the different platforms informs their decision about what experiment to run next (i.e., which platform to investigate). They may choose the platform that they believe to be the most promising thus far (*exploit*) or choose to invest in a platform that they believe they do not know enough about yet (*explore*) but may have the potential to outperform.[5] They continue this cycle of experimentation until they are confident that one of the platforms is truly superior to the others, and they can concentrate all subsequent efforts in that area.

Belief strength captures how confident the individual is that the option they identified as superior is truly superior, and is the mechanism that underlies an individual's propensity to explore versus exploit in a given problem context. Prior work suggests that the magnitude of the difference in the beliefs between two options is directly related to the probability of choosing the option with the higher belief (Posen and Levinthal 2012). The higher this difference, the greater the belief strength, and the more confident the decision maker is that the superior option is in fact superior, and therefore exploits.

Consider one decision maker, Anita, making two choices; say choice of restaurant, A and B, versus choice of coffee shop C and D. Suppose Anita scores them on a 1–10 scale of how much she likes them. Her scores for the restaurants are [A,B: 9,2], whereas for the coffee shops are [C,D: 6,5]. In this case, the choice A and C will both be recorded as exploitation—Anita chooses the options she likes best. However, her belief strength is very different. She strongly prefers restaurant A to B; the score difference of seven suggests a high belief strength that A is superior to B. In contrast, she is almost indifferent between the two coffee shops; the score difference of one suggests that she has low belief strength that coffee shop C is superior to D. Thus, for her next visit, Anita is much more likely to visit restaurant A (exploit) compared with B (explore), but likely has similar probabilities for visiting coffee shop C (exploit) versus D (explore).

How do decision makers increase their belief strength? By experiential learning. As the decision maker takes more samples from the different alternatives (for example, runs more experiments in the different vaccine technology platforms), their beliefs about their relative values stabilize. Again, consider the example from above. Suppose Anita is new to town, and initially she is indifferent between A and B. If there was a strong quality differential (perhaps the food at restaurant B is terrible), after very few visits she understands which restaurant she prefers, that is, has developed strong beliefs (in our example, in favor of restaurant A), and therefore stops visiting B quickly. In the case of coffee shops, after a few trials, she perceives them to be similar, and she has low belief strength about whether C is superior to D. It would take her much longer, that is, many more visits, to understand which one she truly prefers and she may continue to vacillate between C and D for a long time. In summary, strong beliefs that favor one alternative over the other(s) may derive from one of two pathways: (a) either the difference in beliefs about the alternatives' attractiveness is very wide; (b) or the agent's high level of experience provides greater certainty.

In learning from sequential sampling, both these pathways occur, but more importantly, they interact, that is, pathway (a) may undermine pathway (b). The wider the perceived difference between the alternatives early on, the less likely the inferior alternative is sampled in later rounds. Suppose alternatives X and Y are truly equal, but in a run of luck in the first few trials, X provided returns from the right tail of the distribution, and Y from the left tail. At this point, the decision maker perceives a wide difference between X and Y, and stops sampling Y. While the false positives from X eventually get corrected (X reverts to its mean), the false negatives from Y do not, since even at the mean level, X appears more attractive than Y, given the decision maker's (limited and skewed) experience with Y (Denrell and March 2001, Denrell 2003, Denrell and Liu 2021).

Similarly, strong beliefs may also mislead the decision maker when an environment shock improves the utility of under-sampled options. From our previous example, suppose the quality of restaurant B dramatically improved (they hire a new chef), but Anita may never enjoy this because she stopped visiting that place. In this example, what was a true negative before the shock (restaurant B was truly inferior to restaurant A), converts into a false negative after the shock (B becomes better than A, but Anita continues to believe the reverse). However, since B remains less sampled, this error is not corrected. In sum, although strong beliefs help the decision maker quickly identify and exploit the superior solution, they can also trap the decision maker into inferior solutions for longer when the environment changes.

## Performance Goals and the Adaptive Learning Process

Our novel argument is that performance goals influence this adaptive learning process by influencing individuals' belief strength about the available alternatives. We theorize that individuals subject to challenging performance goals more quickly develop strong beliefs that the inferior option performs worse than the superior option, and therefore quit sampling it sooner.

Prior empirical work on the adaptive learning process in neuroscience and psychology has used do-your-best goals, which assume that decision makers dynamically update their own endogenous goals (or aspiration level) in response to feedback received (e.g., Daw et al. 2006, Cohen et al. 2007, Mehlhorn et al. 2015). Thus, although this work has significantly improved our understanding of the belief formation process in adaptive learning, it does not consider how externally set performance goals influence belief strength and therefore decision choices.

We argue that externally imposed performance goals modify the adaptive learning process by changing how the decision maker interprets feedback. According to Simon (1955, p. 105, fig. 1), goals perform an encoding function that reduces a complex environment into a smaller number of states. He argues that this encoding function is an essential purpose of goals, partitioning the payoff space into successes and failures (see also; Hoppe 1930, Lewin et al. 1944). Applied to the choice and adaptive learning processes described above, a performance goal changes how the individual interprets the received feedback from their choice by providing them with a concrete benchmark against which to judge the feedback she receives. For example, students who score nominally close outcomes, say 49 versus 50 points, will interpret these very differently (as a failing versus passing grade), with the actions leading to the failing (passing) grade reinforced much more negatively (positively) than warranted by just the nominal difference in outcomes. Due to such asymmetric reinforcement, performance goals strongly influence the belief strength the decision maker develops about the different options. Different types of goals (moderate versus challenging) provide individuals with different benchmarks; two individuals subject to different performance goals will encode the same outcome from the same choice differently as successes versus failures. This in turn leads these individuals to making different choices next time (they sample different alternatives), their belief strength develops differently, and their subsequent performance outcomes diverge. We use a simple example to illustrate this mechanism:

Consider an employee who must choose between three investment options: A, B, or C. For simplicity, let us assume that the payoffs for each of these options are uniformly distributed, with equally wide and partially overlapping intervals, as shown in Figure 1(a).[6] When the employee chooses one of these options, they will receive a return on investment, or payoff. If the employee knew the true payoff distributions for all options, then choosing between them would be a trivial task: They would always choose option A to maximize their payoffs. In that case there is no explore-exploit tradeoff to make because there is no uncertainty about the payoffs of the investment options.

However, when employees make these choices under uncertainty, for example, choosing between novel business opportunities or technology platforms, they cannot know the options' payoff distributions and need to learn from experience. In this case, the employee may first choose option C (the objectively or *truly* worst choice), receiving a payoff of 23, and then choose option A (the objectively or *truly* best one), receiving a payoff of 18, updating their beliefs about both options in the process. If their judgment is based only on these two trials, then they would (erroneously) judge option A to be worse than option C. Yet over successive trials—assuming they choose both options repeatedly—they will develop more accurate and stronger beliefs about which is the *truly* better option. This simple example underscores the fact that learning figures prominently in a decision maker's search for better options.

As we observed earlier, such trial and error learning is subject to several pathologies (Denrell and March 2001, March 2003), and employees may not truly sample all the available options equally. In addition, such trial and error learning is costly, and the manager tries to steer the employee more quickly toward more consistently selecting a high performing option (i.e., exploiting by choosing option A in our example, unknown to both). To this end, the manager utilizes goal setting.

By setting appropriate goals, the manager manipulates the feedback the employee receives. If the feedback allows the employee to perceive one option as vastly superior (or inferior) to others with very few trials (pathway (a) as discussed earlier), then the employee has effectively short-circuited the costly trial-and-error learning process to quickly settle on the superior option. Thus, the goal setting process, if performed adequately, should align employee preferences (i.e., choose the option that meets the goal) with the manager's interests (i.e., choose the *truly* best option).

So what is the mechanism, that is, how does the given goal affect the learning and choice process? Recall that

**Figure 1.** (Color online) Feedback Encoding Based on Performance Goal

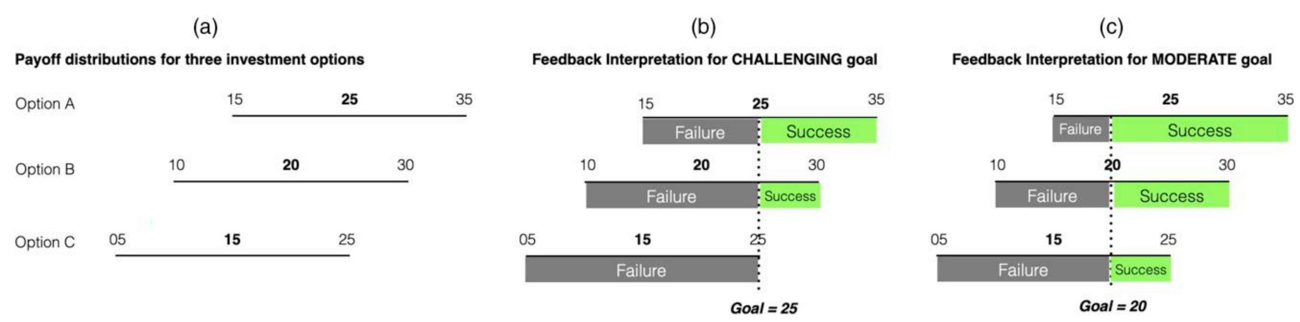

performance goals provide an encoding function that simplifies feedback received into successes and failures. The moderate or challenging goal given, presents the benchmark against which to judge the payoff received.[7] This encoding function changes the feedback perceived by the decision makers, which in turn affects their belief strength, and thus their choices and their performance consequences.

For illustration, consider two employees, Camille and Davon, who play the investment game above. We give Camille a challenging performance goal of 25 per round whereas Davon receives a moderate goal of 20 per round. With this new performance goal, Camille will now encode any payoff of 25 and above as a success, and anything below 25 as a failure (Figure 1(b)); while Davon will encode payoffs of 20 and higher as successes and anything below 20 as failures (Figure 1(c)).

This encoding has an important effect on how these employees learn from feedback. In this scenario, Camille will encode *almost all* payoffs received from option C as failures because the (objectively) worst option's highest possible payoff is 25—the same as her goal—whereas she will encode only 50% (resp. 75%) of the payoffs from option A (resp. option B) as failures. Thus, after very few samples, she expects that option C (almost always failure) is inferior to the other two investment options. In contrast, Davon, with a goal of 20, receives 25% (resp. 50%) failures from option A (resp. option B), and 75% failures for option C, making it more difficult for Davon to differentiate between the options. Consequently, Camille requires fewer feedback opportunities to identify the *truly* inferior option (option C, in this case) compared with Davon. In other words, after only a few trials, Camille's *belief strength* is higher than Davon's, and Camille stops sampling option C sooner than Davon. The corollary is that Camille will spend more of her time learning about options A and B and should therefore increase her belief strength regarding those two options more, compared with Davon.

For this mechanism to work, the performance goal should be meaningful; that is, it must allow the decision maker to distinguish between the options based on feedback. If the goal is set such that all options return mostly successes or mostly failures, the goal does not allow the decision maker to meaningfully learn about the differences between the options, and we do not expect the mechanism—that is, changes in belief strength, to operate as argued here. In sum, challenging performance goals and the resulting feedback interpretation process enables the decision maker to quickly develop higher belief strength when the decision problem consists of a limited number of options. Since higher belief strength means the decision maker is relatively more confident that the superior option is indeed superior, she will select this option, that is, she will exploit. Therefore, we hypothesize that

**Hypothesis 1.** *Decision makers with a challenging performance goal exploit, that is, select the option they* believe to be best, *more often than those with a moderate goal (so long as goals are meaningful).*

Since a challenging performance goal makes it easier for the individual to quickly identify the truly inferior options, they have more opportunity to learn about the remaining superior options and distinguish between their relative attractiveness. Thus, the decision maker with a challenging performance goal is more likely to identify the *truly* superior option from among a more limited number of choices. In addition, in a stable task environment, in which the payoff distributions do not change, identifying and exploiting the *truly* best option leads to better performance regardless of goal level. Since decision makers with a challenging goal choose the *truly* best option more frequently, they are likely to experience higher cumulative performance:

**Hypothesis 2a.** *Decision makers with a challenging performance goal select the* ***truly best*** *option* more *frequently in a stable environment than those with a moderate goal.*

**Hypothesis 2b.** *Decision makers with a challenging performance goal exhibit* ***higher*** *cumulative performance in a stable environment than those with a moderate goal.*

Recall from our discussion above that the mechanism underlying these effects is that different goals affect the development of individuals' belief strength differently. When facing a challenging performance goal, individuals encode a greater proportion of the feedback from inferior options into failures, which leads to greater confidence in their belief that the inferior option is worse than the other options, relative to an individual with a moderate goal. By eliminating that inferior option early, they have more opportunity to sample the other two options which provides more granular feedback on their relative attractiveness, hence increasing their belief strength across all options.[8] We therefore hypothesize that:

**Hypothesis 3.** *Decision makers with a challenging performance goal develop* greater *belief strength than those with a moderate one.*

## The Impact of an Environmental Shock

So far, we have only considered organizational uncertainty resulting from novel investment opportunities. However, managers also face uncertainty from environmental change that may alter the relative attractiveness of different investment options. Changing customer preferences or new technologies can profoundly alter the attractiveness of available opportunities in ways that managers may not understand or sometimes even recognize. For example, although insurance companies know that Artificial Intelligence/Machine Learning technologies can profoundly impact their business, they may not know whether the biggest impact is likely to be in customer acquisition, retention, customization, or securitization. Motorola, Nokia and later Microsoft did not recognize the profound ways the mobile phone industry changed when Apple introduced the iPhone but continued to emphasize hardware and operating systems rather than invest in platform capabilities. We next consider whether challenging performance goals make it easier or more difficult for managers to adapt to such environmental shocks.

We argue that the same mechanism that leads to faster exploitation of the truly superior option in the stable environment will likely turn into a liability when an unforeseen environmental shock makes the previously unattractive option more attractive; in other words when the environmental shock changes a true negative into a false negative. As discussed earlier, decision makers with a challenging performance goal have developed greater belief strengths, and they therefore under-sample the (previously) inferior option. After a disruptive environmental shock, these individuals continue to under-sample that option, the now false negative, and over-sample the false-positive, the still *believed to be best* (but post-shock inferior) option. This slower deviation from the previously superior choice will lead to lower overall performance:

**Hypothesis 4a.** *Decision makers with a challenging performance goal choose the newly changed best option* less *frequently* after *a disruptive environmental shock than those with a moderate one.*

**Hypothesis 4b.** *Decision makers with a challenging performance goal exhibit* worse *performance immediately* after *a disruptive environmental shock than those with a moderate one.*

However, false-positives are self-correcting with repeated sampling (as theorized by Denrell and March 2001). Over time, as this correction occurs, individuals will move away from it, and start to sample the previously neglected false-negative option. With time, they now recognize the potential of the previously inferior arm, and start sampling it again. Once again, their challenging goals lead them to form strong beliefs, but now in favor of the previously (pre-shock) inferior, but now (post-shock) superior arm. A corollary of Hypothesis 4a is that this correction will happen over a period of time, as individuals adjust their beliefs with repeated sampling. Thus, after an adjustment period post-shock, decision makers with a challenging performance goal will choose the newly changed best option *more* frequently again and exhibit better performance than those with a moderate one.

Taken together, these hypotheses also act as a mechanism test, as the same mechanism (of greater belief strength under the challenging performance goal) predicts both implications for performance and exploration behavior when decision makers face a stable versus disrupted environment. The mechanism also makes it clear that the theory developed here only applies in so far as the goal meaningfully distinguished between better and worse options. Goals that are too high or too low do not perform this function, and therefore, may not influence the adaptive learning process in the same way.

## Method

### Experimental Procedure

To test these hypotheses, we designed a behavioral laboratory experiment in which we manipulated performance goals *(challenging versus moderate)* and observed participants' choices and performance. Each participant played a single-player investment game. The laboratory experiment was set up as a between-subject design across three main studies. In each study, participants were randomly assigned to one of several goal conditions. We ensured that each participant took part in only one of the studies.

To understand how goals influence the adaptive learning processes, the experiment should include the following properties. First, it should be a task in which learning from feedback is inherent to identifying superior options. Second, the task should be able to separate out risk taking from learning. Third, to get to the heart of the mechanism, there should be a direct link between actions and feedback, without intervening complexity. Fourth, it should be easy for participants to understand their current performance relative to their goals. Finally, the choice between exploration versus exploitation should be nontrivial, but learning should be plausible in a limited amount of time. The multiarmed bandit problem satisfies these criteria and has been extensively used to study exploration-exploitation behavior (see Mehlhorn et al. 2015).[9]

The multi-armed bandit task was framed as an investment game where the decision maker chooses between multiple options. Each option has an uncertain payoff, which implies that (1) the decision maker does not know ex ante which of the different options is better; and (2) the feedback received from choosing an option is noisy. Our investment game gave participants the choice between the same three options across all studies; their payoffs were uniformly distributed and had the same variance but different means ([25, 20, 15] for options 1, 2, and 3, respectively, with a constant interval of $\pm 10$), as shown in Figure 2(a). The payoff distributions were chosen to ensure there is sufficient overlap to allow for learning to be meaningful and that there is enough unique feedback to discern payoff differences over time, while avoiding extreme payoffs and negative payoffs that may induce different psychological processes. Depending on

**Figure 2.** (Color online) Experimental Setup: Payoff Distributions and Feedback Interpretation

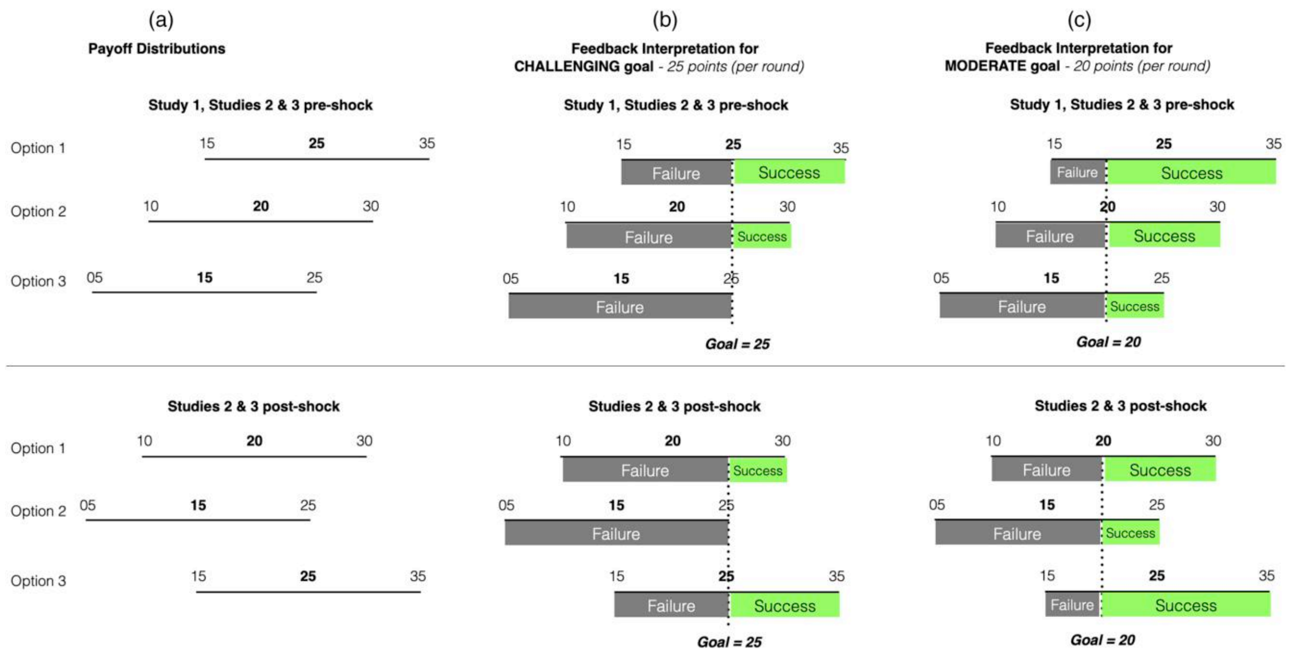


*Notes.* Panel (a) illustrates the payoff distributions of the three investment options that underlie all studies. Panels (b) and (c) overlay the "success" and "failure" interpretations made by the decision makers based on their respective goals. Panel (b) (resp. (c)) shows the interpretations for a challenging (resp. moderate) goal of 25 (resp. 20) payoff points. Studies 5a, 6a, 6c, 7a, and 7c (resp. 5b, 6b, 6d, 7b, 7d) replicate Study 1 (resp. Study 2b), using the same payoff distributions.

their goal, decision makers assigned the payoffs received into categories of "success" or "failure" with different percentages: the success percentages for the three options were 77.5%, 52.5%, and 27.5% for the moderate-goal condition; and 52.55%, 27.5%, and 2.5% for the challenging-goal condition (illustrated in Figure 2).

In order to enable learning from experience, decision makers played the investment game for 50 rounds. In each round, the decision maker chose whatever option they wanted and received a payoff. Over time, the decision maker learned which of the options yielded higher or lower payoffs. The decision maker's payoffs accumulated over the rounds of the game, and their goal was to maximize the cumulative end-of-game payoff. The task setup ensured that time pressure was not a factor: We did not limit participants' available time, and the time needed to play the game averaged just under 1.5 minutes excluding the introduction and briefing of the game (the time averaged 2 minutes for games with extended periods). The experiment was IRB-approved at the relevant institutions and did not incorporate any deception (i.e., it did not misrepresent the *purpose* or *nature* of the experiment or provide false feedback in a way that would prevent participants from giving informed consent, Cook and Yamagishi 2008). The verbatim instructions, procedures and instrument for this laboratory experiment can be found in the Online Supplement, Section 1; and our data, analyses, and preregistered analysis plans can be accessed via this link.[10]

**Task Description.** In line with our theoretical intent, we framed the task as one of investment choice under uncertainty. Each participant adopted the role of an R&D manager for a digital firm and, in each round, decided on a product platform in which to invest. All participants were informed that the payoffs are uncertain and subject to market turbulence. The task description also points out that random environmental shocks could alter the relative attractiveness of the available options. These instructions were given regardless of whether the study contained a shock.

In each period, participants selected one of the three available options for investment. Upon making an investment, a participant immediately received a payoff displayed as points earned on the given option. The feedback was clearly tied to the choice just made, and the cumulative payoffs (at the top of the screen) were updated in turn, along with the current round and progress. Participants could see at any time the payoff received on their latest investment, their cumulative payoff, the total number of trials played, and the average payoff received from each option, similar to other studies (Lee et al. 2011, Sang et al. 2020) in order to minimize distortions from different memory capacities.

**Treatment 1: Goal Setting.** Participants were randomly assigned to different performance goal conditions. In all studies, participants were informed that they should attempt to achieve their assigned performance goal. We followed best-practice in the goal setting literature to determine the goals set. Prior studies in goal setting have typically used performance achieved by the top 10% of individuals as the benchmark for setting the challenging goals, and average performance for moderate goals (please see Locke and Latham 1990 for a review). We followed a similar procedure.

We implemented the manipulation for challenging [resp. moderate] goals by instructing the participants as follows: "The previous manager had achieved total earnings of 1,250 [1,000] points over their tenure of 50 rounds, that is, 25 [20] points per round. You should aim to earn at least this amount."[11] For simplicity, we describe the specific performance goals used in each study in the results section, Table 1 provides an overview.

Before conducting our main studies, we ran pilot studies to ensure that these two goals were set such that 25 per round would be challenging but feasible while 20 per round was achieved more easily without being trivial. The statement about "the previous manager" having achieved a certain level of earnings was referencing the data from the pilot studies in a simplified way and communicating that the goal was achievable, while maintaining the clearly fictitious setting of our study.

**Treatment 2: Environmental Shock.** We implemented the environment shock treatment by reshuffling the means of the three investment options while leaving the overall payoff landscape otherwise unchanged, as shown in Figure 2.

**Boundary Conditions.** When studying the boundary conditions of our findings, we also examined the effects of other performance goals (Study 3), different environmental shock conditions (Study 4), no-goal condition (Studies 5a and 5b), and different incentive structures (Studies 6 and 7), see Online Supplement, Section 5.

**Participant Recruitment.** Because adaptive learning is a fundamental human behavior, we decided to run our studies in different countries and settings so as to increase (however slightly) generalizability of our findings. For studies 2a, 3a, and 3b, we recruited participants from a public university in the United States and the experiment was conducted in a laboratory setting (Study 2a in-person, Studies 3a & 3b online, due to the pandemic). For Study 4, students from an undergraduate course at a Singapore university participated in the experiment as part of a class exercise. For all other studies, participants were recruited through Amazon's Mechanical Turk (MTurk). Table 2 summarizes our data collection as well as participants' demographics. We used our first study to inform our power analysis and estimated that a sample size of 72 (resp. 98) would give us statistical power of 0.8 (resp. 0.9). The power analysis is reported in the Online Supplement, Section 2.

Across the different studies, MTurk participants were similar in terms of the basic demographic dimensions that we collected. On average, between 31% and 45% self-identified as female, and they reported an average age of 36 (with a wide range between 18 and 76). The majority of MTurk participants reported English as their

**Table 1.** Overview of All Study Conditions

| | Performance goals | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| Study | moderate [20] | challenging [25] | no goal | lower-edge [5] | low [15] | upper-edge [35] | Environment | Rounds | Incentives |
| Main results | | | | | | | | | |
| 1 | x | x | | | | | stable | 50 | top 5 performers |
| 2a | x | x | | | | | **shock** | 50 | top 5 performers |
| 2b | x | x | | | | | **shock** | **80** | top 5 performers |
| Boundary conditions | | | | | | | | | |
| 3 | x | x | | **x** | **x** | **x** | stable | 50 | top 5 performers |
| 4 | x | x | | | | | **Positive shock** | 50 | top 5 performers |
| 5a | x | x | **x** | | | | stable | 50 | top 5 performers |
| 5b | x | x | **x** | | | | **shock** | **80** | top 5 performers |
| 6a | x | x | | | | | stable | 50 | **fixed bonus** |
| 6b | x | x | | | | | **shock** | **80** | **fixed bonus** |
| 6c | x | x | | | | | stable | 50 | **exchange rate** |
| 6d | x | x | | | | | **shock** | **80** | **exchange rate** |
| 7a | x | x | | **x** | **x** | **x** | stable | 50 | **fixed bonus** |
| 7b | x | x | | **x** | **x** | **x** | **shock** | **80** | **fixed bonus** |
| 7c | x | x | | **x** | **x** | **x** | stable | 50 | **exchange rate** |
| 7d | x | x | | **x** | **x** | **x** | **shock** | **80** | **exchange rate** |

*Note.* The bold higlights indicate how the studies deviate from the baseline setup. *The environmental shock occurred at round 30. "top 5 performers" received a $10 gift certificate.*

**Table 2.** Sample Characteristics for All Studies

| Study | Study condition | Total number of participants | Number of participants in challenging goal condition | Location / Source | Mean age [range] | Female (%) | BART Score | English is first language (%) | Country of residence (North America %) | Education (College degree %) |
|---|---|---|---|---|---|---|---|---|---|---|
| Main results | | | | | | | | | | |
| 1 | Stable environment: no shock | 193 (220) | 98 | MTurk | 35.0 [21, 70] | 31.3 | 39.3 | 92.6 | 88.5 | 69.1 |
| 2a | Disrupted environment: replicates Study 1 *with* shock | 200 (220) | 103 | US university | 20.8 [19, 31] | 45.9 | 40 | 55 | N/A | N/A |
| 2b | Replicates Study 2 but with an additional 30 rounds post-shock | 73 (80) | 37 | MTurk | 35.4 [22, 67] | 34.6 | 40.4 | 95.1 | 92.6 | 63 |
| Boundary conditions | | | | | | | | | | |
| 3a | Replicates Study 1 but with extended goals [5, 15, 20, 25, 35] | 268 (345) | 50 | US university | 21.2 [18, 41] | 41.55 | 51.05 | 70.89 | 95.6 | 19.26 (50.68) |
| 3b | Replicates Study 2b with extended goals [5, 15, 20, 25, 35] | 247 (290) | 50 | US university | 21.4 [18, 40] | 46.2 | 49.1 | 85.81 | 91.7 | 42.21 (47.75) |
| 4 | Disrupted environment: *positive* shock: shifts average payoffs upward | 77 (90) | 39 | Singapore university | 20.7 [19, 24] | 45.5 | 40 | N/A | 66.2 (Singaporean) | 85.7 Business school (vs other department) |
| 5a | Replicates Study 1 but with a no-goal condition | 106 (120) | 35 | MTurk | 38 [18, 71] | 45.5 | 37.8 | 79.7 | 72.4 | 62.6 |
| 5b | Replicates Study 2b but with a no-goal condition | 104 (120) | 38 | MTurk | 35 [20. 69] | 44.8 | 37.8 | 81.5 | 66.1 | 75.8 |
| 6a | Replicates Study 1 but with fixed bonus incentive | 107 (128) | 42 | MTurk | 36 [20, 70] | 38 | 39.7 | 82.2 | 70.5 | 72.1 |
| 6b | Replicates Study 2b but with fixed bonus incentive | 98 (115) | 37 | MTurk | 35 [20, 67] | 31.4 | 40 | 83.1 | 69.5 | 61 |
| 6c | Replicates Study 1 but with exchange rate-based incentive | 100 (122) | 40 | MTurk | 36 [19, 69] | 26.2 | 39.8 | 83.6 | 71.3 | 66.4 |
| 6d | Replicates Study 2b but with exchange rate-based incentive | 96 (119) | 39 | MTurk | 35.5 [22, 69] | 30.7 | 38.7 | 84.9 | 71.4 | 63.9 |
| 7a | Replicates Study 6a with extended goals [5, 15, 20, 25, 35] | 264 (381) | 58 | MTurk | 35.1 [18, 75] | 35.43 | 38.24 | 88.39 | 81.79 | 62.26 |
| 7b | Replicates Study 6b with extended goals [5, 15, 20, 25, 35] | 202 (297) | 45 | MTurk | 35.76 [18, 72] | 34.71 | 43.82 | 87.41 | 76.94 | 68.81 |
| 7c | Replicates Study 6c with extended goals [5, 15, 20, 25, 35] | 147 (209) | 30 | MTurk | 36.03 [18, 76] | 34.3 | 49 | 87.5 | 84.13 | 64.9 |
| 7d | Replicates Study 6d with extended goals [5, 15, 20, 25, 35] | 149 (200) | 27 | MTurk | 34.63 [18, 72] | 36.44 | 37.84 | 77.27 | 62.62 | 77.88 |

*Notes*. The number of participants who provided useable data are provided, along with the total number of recruited participants in parentheses. The slight reduction in each study's number of subjects is due to some participants failing to complete the game or in the case of MTurk studies, accepting the study on MTurk but being screened out (on Qualtrics) as having taken part in another of the studies. The MTurk studies were not regionally restricted; the BART (balloon analogue risk task) score captures participants' risk taking propensity (see Lejuez et al. 2002).

**Table 3.** ANOVA Results for Studies 1, 2a, and 2b

| Hypothesis | Findings | Study 1 – Stable environment | | | Study 2a – Disrupted environment | | | Study 2b – Disrupted environment with extended post-shock periods | | |
|---|---|---|---|---|---|---|---|---|---|---|
| | *Compared with moderate-goal condition* | Goal | | Difference between conditions | Goal | | Difference between conditions | Goal | | Difference between conditions |
| | | 20 ($n$ = 95) | 25 ($n$ = 98) | | 20 ($n$ = 97) | 25 ($n$ = 103) | | 20 ($n$ = 36) | 25 ($n$ = 37) | |
| Across rounds | | | | | | | | | | |
| H1 | Decision makers with a challenging performance goal exploit, i.e., select the option they believe to be best, more often than those with a moderate goal. | 24.89 (5.47) | 29.43 (5.36) | $\Delta$ = 4.54 $F$ = 33.97 $p$ = 0.000 | 20.91 (7.18) | 26.30 (9.10) | $\Delta$ = 5.38 $F$ = 21.39 $p$ = 0.000 | 31.08 (5.62) | 46.83 (8.29) | $\Delta$ = 15.75 $F$ = 89.77 $p$ = 0.000 |
| Pre-shock results | | | | | | | | | | |
| H2a | Decision makers with a challenging performance goal select the truly best option more frequently in a stable environment than those with a moderate goal. | 26.35 (4.35) | 29.09 (5.36) | $\Delta$ = 2.74 $F$ = 15.18 $p$ = 0.000 | 14.22 (4.55) | 15.89 (5.13) | $\Delta$ = 1.67 $F$ = 5.94 $p$ = 0.016 | 14.22 (3) | 17.51 (4.78) | $\Delta$ = 3.29 $F$ = 12.32 $p$ = 0.001 |
| H2b | Decision makers with a challenging performance goal exhibit higher cumulative performance in a stable environment than those with a moderate goal. | 1080.57 (57.80) | 1110.96 (53.55) | $\Delta$ = 30.39 $F$ = 14.37 $p$ = 0.000 | 633.55 (44.29) | 651.82 (49.10) | $\Delta$ = 18.27 $F$ = 7.60 p = 0.006 | 640.17 (37.18) | 668.84 (43.79) | $\Delta$ = 28.67 $F$ = 9.07 $p$ = 0.004 |
| H3 | Decision makers with a challenging performance goal develop greater belief strength than those with a moderate one. | 0.447 (0.170) | 0.467 (0.149) | $\Delta$ = 0.020 $F$ = 35.89 $p$ = 0.000 | 0.414 (0.175) | 0.425 (0.168) | $\Delta$ = 0.011 $F$ = 5.69 $p$ = 0.017 | 0.394 (0.181) | 0.442 (0.163) | $\Delta$ = 0.048 $F$ = 37.78 $p$ = 0.000 |
| | … belief strength measured as the difference in success rate of the best and worst alternatives. | 0.380 (0.23) | 0.429 (0.184) | $\Delta$ = 0.049 $F$ = 126.36 $p$ = 0.000 | 0.322 (0.243) | 0.382 (0.201) | $\Delta$ = 0.059 $F$ = 96.05 $p$ = 0.000 | 0.256 (0.268) | 0.388 (0.203) | $\Delta$ = 0.132 $F$ = 150.88 $p$ = 0.000 |
| | … belief strength measured as the variance across beliefs. | 0.233 (0.087) | 0.250 (0.081) | $\Delta$ = 0.016 $F$ = 84.06 $p$ = 0.000 | 0.219 (0.092) | 0.230 (0.091) | $\Delta$ = 0.010 $F$ = 17.17 $p$ = 0.000 | 0.208 (0.095) | 0.236 (0.088) | $\Delta$ = 0.028 $F$ = 47.56 $p$ = 0.000 |
| Post-shock results (rounds 31–50) | | | | | | | | | | |
| H4a | Decision makers with a challenging performance goal choose the newly changed best option less frequently after a disruptive environmental shock than those with a moderate one. | | | | 7.45 (4.06) | 4.99 (3.70) | $\Delta$ = −2.46 $F$ = 20.08 $p$ = 0.000 | 8.02 (2.74) | 4.08 (2.95) | $\Delta$ = −3.94 $F$ = 34.93 $p$ = 0.000 |
| H4b | Decision makers with a challenging performance goal exhibit worse performance immediately after a disruptive environmental shock than those with a moderate one. | | | | 416.44 (36.21) | 406.27 (30.10) | $\Delta$ = −10.17 $F$ = 4.31 $p$ = 0.039 | 418.02 (29.96) | 399.37 (29.58) | $\Delta$ = −18.64 $F$ = 7.16 $p$ = 0.009 |
| Extended post-shock results (rounds 51–80) | | | | | | | | | | |
| H2a | Decision makers with a challenging performance goal select the truly best option more frequently in a stable environment than those with a moderate goal. | | | | | | | 15.16 (3.32) | 17.84 (5.90) | $\Delta$ = 2.67 $F$ = 5.64 $p$ = 0.020 |
| H2b | Decision makers with a challenging performance goal exhibit higher cumulative performance in a stable environment than those with a moderate goal. | | | | | | | 647.38 (42.36) | 676.78 (36.92) | $\Delta$ = 29.39 $F$ = 10.01 $p$ = 0.002 |

*Notes.* This table reports sample means with standard errors in parentheses. The "Difference between conditions" column reports the ANOVA results.

first language, and around 66% reported having at least a four-year college degree. We found a difference in country of residence for studies run before (studies 1–4) versus during (studies 5–7) the pandemic (resp. 90% versus 70% from North America). In contrast to the MTurk studies, the two university samples were younger (21 on average), had a higher percentage of female-identifying participants (46%), and included fewer students with English as their first language (in the U.S. sample; data on their first language is not available for Singapore). Across all populations, participants' risk-taking propensity, measured by the balloon analogue risk task (BART; see Lejuez et al. 2002), lay in the range of 38 to 51. Overall, this set of study populations provided us with some degree of generalizability. The fact that learning tendencies among these diverse participants were consistent, in very different exogenous conditions of uncertainty (before and during the height as well as later stages of the pandemic), appears to indicate that our studies indeed captured a fundamental human behavior.

**Incentive Structures.** Incentives matter because they make choices consequential. In studies 1 through 5, we incentivized participants to perform well by offering a $10 gift certificate to each of the top five performers. In addition, participants in studies 2a, 3a, 3b, and 4 were rewarded with course participation credits; in studies 1, 2b, 5a, and 5b, we replaced course credit with a fixed payment for participation. In setting up the initial incentive structure, we followed the goal-setting literature in using "mere goals"—goals that simply establish a reference point (Heath et al. 1999, Larrick et al. 2009)—and then observing whether such a simple manipulation results in behavioral differences.

However, research has also shown that winner-takes-all incentives (the type described above) may boost risk-taking and hence exploration (Manso 2011, Ederer and Manso 2013). If the incentive design applied to the first set of studies does indeed encourage more risk-taking and exploration, then the setup will provide a conservative test of the hypotheses that instead predict more exploitation. Yet because incentive design has a pronounced effect on learning and goals, we ran additional studies (Studies 6 and 7) to examine the effect of different incentive structures—bonus payment for meeting the goal versus piece-rate incentives—to establish boundary conditions (see Online Supplement, Section 5).

## Measures

***Exploitation.*** To test Hypothesis 1, we needed to capture exploitation choices. In line with prior work, we say that decision makers exploit if they choose the alternative *they believe* gives the highest payoff and explore otherwise (Daw et al. 2006, Song et al. 2019). Therefore, to measure *exploitation*, we needed to estimate the decision makers' beliefs about each of the alternatives.

Belief calculation involved multiple steps. First, we observed the outcome received by the participant in each round. Second, we classified this outcome as a success or failure based on whether the payoff was at least equal to (success), or lower (failure) than the goal.[12] Success (failure) was coded as 1 (0, respectively). Third, we updated the participant's beliefs using the average updating rule. With beliefs calculated in round $t$, we then coded the participant's explore-exploit choices in round $t + 1$. If in round $t + 1$ the participant chose the option that had the highest belief in round $t$, then we coded the variable exploitation as 1. Otherwise, we coded this as 0.

We checked the robustness of this measure in multiple ways. First, we also calculated beliefs using a continuous measure of outcome, as the difference between the payoff received and the decision maker's goal, again averaging the beliefs. Second, we calculated beliefs using a discounted memory updating rule following Christensen et al. (2021), such that early feedback is weighted less than recent feedback. Finally, we estimated the decision maker's beliefs about which option is the best or worst in any given round by fitting the participant's choices and payoffs into a temporal difference-learning algorithm and a softmax choice algorithm (following the extant literature, see Daw et al. 2006, Laureiro-Martínez et al. 2014; please see Online Supplement, Section 3, for details). We calculated beliefs using the learning parameter used in the best fitted model. These robustness checks are shown in Online Supplement Section 4.

***Performance* Was Measured in Two Ways.** To test Hypothesis 2a, we measured how frequently a decision maker chose the truly best option among the three choices (which were known to us but not known with certainty to the experiment's decision makers). In the stable environment (Hypothesis 2a), this option did not change over the course of the study; in the environmental shock condition (Hypothesis 4a), we measured how frequently decision makers chose the truly best *new* option after the shock in round 30. To test Hypothesis 2b, we took the cumulative points earned by a decision maker over the 50 rounds in order to compare relative performance. To test Hypothesis 4b, we measured the cumulative points in the post-shock period, rounds 31–50, and in the extended post-shock period, rounds 51–80.

***Belief strength* Captures How Confident the Decision Maker is About the Accuracy of Their Beliefs.** As discussed earlier, in prior work, the magnitude of the difference in the beliefs between two options is directly related to the probability of choosing the option with the higher belief. The higher this difference, the more 'confident' the decision maker is that the superior option is in fact superior. We use this idea to test Hypothesis 3, and measure decision maker's *belief strength* in three different ways. For our main measure, we calculate *belief strength* as the difference in the decision maker's beliefs of the best and worst options, as experienced by each decision maker (which need not be the first and third option respectively). The greater this difference, the higher the belief strength. For our robustness checks, we calculated *belief strength* as the difference in success rate of the best (option 1) and worst (option 3) alternatives. Successes (and failures) experienced by each participant were computed based on whether the payoff experienced in each round was above (or below) the assigned goal. Finally, following Posen and Levinthal (2012), we calculated *belief strength* as the variance across beliefs. Our results remain robust to these alternative specifications of *belief strength*. While the correlation between these three measures are high, the correlation between exploitation and these three measures of belief strength are low, alleviating any concerns that the measures for these two constructs may be highly correlated.[13]

## Results

In discussing the results, we will briefly summarize each of the study conditions and then document our findings. *We predicted that performance goals dynamically affect decision makers' belief strength about the relative attractiveness of different options over time, which affects their exploration-exploitation behavior and therefore their performance.* This mechanism improves performance in stable environment yet reduces adaptability and performance in the short-term when a shock alters the environment. Since we manipulate whether the goals are challenging versus moderate, we used ANOVAs to test the difference in exploration-exploitation choices and performance across the goal conditions.

### Study 1: Moderate and Challenging Performance Goals in a Stable Environment

In Study 1, participants played the three-armed bandit game in a stable environment; the payoffs were uniformly distributed, with means 25, 20, and 15 for options 1, 2, and 3, respectively, and a constant interval of $\pm 10$. The challenging (resp. moderate) goal is set at 25 (resp. 20) points per round, adding up to 1,250 (resp. 1,000) points over 50 rounds. All results for Study 1 are summarized in Table 3.

We predicted that decision makers with a challenging performance goal would *exploit* more (Hypothesis 1). We found that participants in the challenging-goal condition chose the option they believed to be best more often ($M$ = 29.4, $SD$ = 5.4) than did those with a moderate goal ($M$ = 24.9, $SD$ = 5.5; $F$-statistic = 34.0, $p$-value = 0.000)—outcomes that support Hypothesis 1. In addition, we proposed that decision makers with a challenging performance goal would be more likely than those with a moderate goal to choose the objectively best option (Hypothesis 2a). We found that participants in the challenging-goal condition chose option 1, the objectively best option more often ($M$ = 29.1, $SD$ = 5.4) than those with a moderate goal ($M$ = 26.4, $SD$ = 4.4; $F$ = 15.2, $p$ = 0.000), providing support for Hypothesis 2a. We also found that cumulative performance is significantly higher in the challenging-goal condition (mean $M$ = 1,111.0, standard deviation $SD$ = 53.6) than in the moderate-goal one ($M$ = 1,080.6, $SD$ = 57.8; $F$ = 14.4, $p$ = 0.000); these results support Hypothesis 2b. In the Online Supplement, Section 4, we show that exploitation decisions mediate the relationship between assigned goal and performance.

Driving the observed choices, we predicted that challenging goals would help decision makers form stronger beliefs about the relative attractiveness of the higher payoff options (Hypothesis 3). We found that belief strength was indeed higher for participants facing challenging performance goals ($M$ = 0.47, $SD$ = 0.15) compared with those facing moderate ones ($M$ = 0.45, $SD$ = 0.17; $F$ = 35.9, $p$ = 0.000). Table 3 reports consistent results with two alternative measures for belief strength. These results provide support for Hypothesis 3. The regression analyses shown in the Online Supplement, Section 4 are consistent with the proposed mechanism.

Taken together, the results suggest that the challenging performance goal allows decision makers to distinguish the superior options from the inferior one more easily than the moderate goal does. Yet, given adequate sampling by the decision makers, all participants should be able to distinguish between the available options. Why does setting a challenging performance goal provide an advantage? In order to explore this question we apply the idea from the tau-switch model from Lee et al. (2011). This model suggests that explore-exploit decisions in humans can be treated as a problem where participants move from an "explore" state, where they choose among the available options, to an "exploit" state, where participants stop sampling the option(s) believed to be inferior. A more detailed explanation of how we used the ideas behind the tau-switch model is presented in the Online Supplement, Section 4.2.

The results indicate that decision makers with the challenging performance goal stopped sampling the worst option sooner ($M$ = 32.7, $SD$ = 12.7) than those with a moderate one ($M$ = 40.4, $SD$ = 10.6; $F$ = 20.7, $p$ = 0.000). As a consequence, participants with a challenging goal developed greater belief strength about the relative attractiveness of their believed-to-be-best over the believed-to-be-middle option ($M$ = 0.24, $SD$ = 0.26) compared with those with a moderate goal ($M$ = 0.23, $SD$ = 0.22; $F$ = 7.08, $p$ = 0.008). This is because, participants with a challenging goal sampled the worst option fewer times (following pathway a), allowing them to sample the remaining two options more exhaustively, and distinguish between them (following pathway b).[14] Together, these results further lend support to our theorized mechanism that performance goals dynamically affect decision makers' belief formation about the relative attractiveness of different options over time, which affects their exploration-exploitation behavior and hence performance.

## Study 2: Moderate and Challenging Performance Goals in a Disrupted Environment

We next study how challenging versus moderate goals influence exploration-exploitation decisions when the environment changes, reducing the value of prior learning. In Studies 2a and 2b, participants played the same three-armed bandit game; the payoffs were uniformly distributed, with means 25, 20, and 15 for options 1, 2, and 3, respectively, and a constant interval of $\pm 10$. However, after period 30, the underlying payoffs reshuffled so that option 3 became the objectively best option ($M$ = 25) and option 2 became the worst one ($M$ = 15) (as shown in Figure 2). While Study 2a otherwise replicates Study 1 exactly, Study 2b extends the duration to 80 rounds so that we can examine decision makers' behavior and performance after an adjustment period. All results for Studies 2a and 2b are summarized in Table 3 and the verbatim instructions are available in the Online Supplement, Section 1.

We predicted that decision makers with a challenging performance goal would be at a disadvantage when faced with an environment shock compared with decision makers with a moderate goal. We found that participants in the challenging-goal condition indeed chose the newly changed best option less frequently after the disruptive environmental shock ($M$ = 4.99, $SD$ = 3.70) than did those with a moderate goal ($M$ = 7.45, $SD$ = 4.06; $F$-statistic = 20.08, $p$-value = 0.000)—outcomes that support Hypothesis 4a (replicated in Study 2b) and highlight that false negatives are not as easily corrected. We further predicted that this slower deviation from the previously superior choice will lead to lower overall performance immediately after a disruptive shock (where cumulative performance is measured between rounds 31 and 50). We found that the immediate post-shock performance for decision makers with a challenging goal were lower ($M$ = 406.27, $SD$ = 30.10) than performance for those with a moderate goal ($M$ = 416.44, $SD$ = 36.21; $F$-statistic = 4.31, $p$-value = 0.039); supporting Hypothesis 4b (replicated in Study 2b). Finally, we found that after an extensive adjustment period (rounds 51 to 80), performance and exploitation behavior of the decision makers with the challenging goal returned to pre-shock levels (as presented at the bottom of Table 3).

## Study 3: Boundary Conditions: Extended Performance Goals

The theory and results reported above rely on the performance goal to be both meaningful and achievable. In Studies 3a and 3b we examine the boundary conditions of our findings, by including additional performance goals. In Study 3a (resp. 3b) participants played the same three-armed bandit game as in Study 1 (resp. Study 2b). We again manipulated the performance goals, but we expanded these to include not just the challenging (25) and moderate goals (20) but also a symmetrical low goal (15) and two options at the extreme edges of possible payoffs (5 and 35, respectively). These additional goal manipulations and their uniformly distributed payoff ranges are shown in Figure 3. All results for Study 3a (resp. 3b) are summarized in Table 4 (resp. Table 5). All other boundary conditions are discussed in Online Supplement, Section 5.

**Feedback Interpretation with Low Goal in a Stable Environment.** The low goal condition (goal = 15) mirrors the challenging goal condition (goal = 25). In the low goal condition, the decision maker always receives a success feedback (from option 1), in contrast to the challenging goal condition where the decision maker always receives a failure feedback (from option 3 as shown in Figure 2b). Thus, we might expect that the low goal condition helps the decision maker learn faster, that is, identify the superior option (always success) sooner, similar to how the challenging goal condition

**Figure 3.** (Color online) Feedback Interpretation for the Extended Goal Conditions

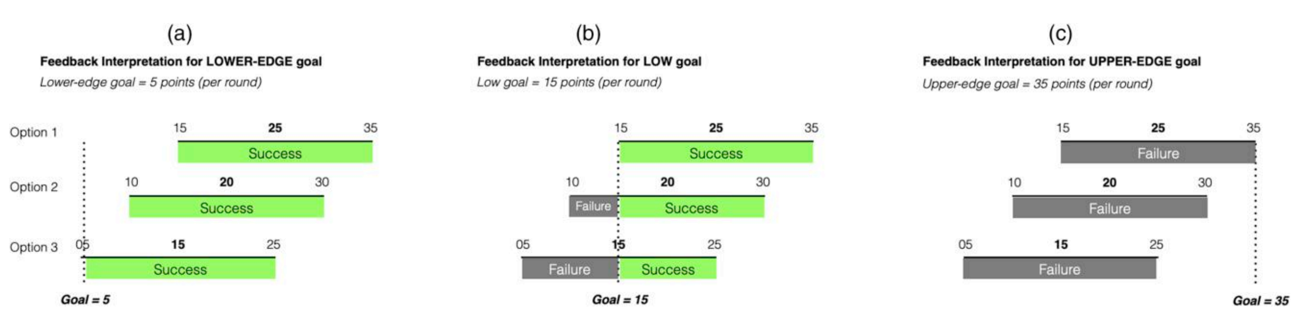

**Table 4.** ANOVA Results for Study 3a – Extended Goal Conditions in a Stable Environment

| Hypothesis | | Low goal comparison | | | | | | Lower- and upper-edge goal comparison | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | | Goal | | | Difference between conditions | | | Goal | | | Difference between conditions | | |
| *compared with moderate-goal condition* | | 15 ($n$ = 50) | 20 ($n$ = 52) | 25 ($n$ = 50) | 25–20 | 25–15 | 20–15 | 5 ($n$ = 56) | 20 ($n$ = 52) | 35 ($n$ = 60) | 20–5 | 20–35 | 35–5 |
| Across rounds | | | | | | | | | | | | | |
| H1 | Decision makers with a challenging performance goal exploit, i.e., select the option they believe to be best, more often than those with a moderate goal. | 26.14 (6.26) | 25.46 (6.53) | 30.46 (5.71) | $\Delta = 5$ $F = 16.87$ $p = 0.000$ | $\Delta = 4.32$ $F = 13.00$ $p = 0.000$ | $\Delta = -0.68$ $F = 0.29$ $p = 0.594$ | 21.55 (7.14) | 25.46 (6.53) | 24.81 (6.90) | $\Delta = 3.90$ $F = 8.76$ $p = 0.004$ | $\Delta = 0.64$ $F = 0.26$ $p = 0.614$ | $\Delta = 3.26$ $F = 6.26$ $p = 0.014$ |
| H2a | Decision makers with a challenging performance goal select the truly best option more frequently in a stable environment than those with a moderate goal. | 26.14 (6.25) | 25.75 (6.72) | 29.96 (5.76) | $\Delta = 4.21$ $F = 11.48$ $p = 0.001$ | $\Delta = 3.82$ $F = 10.08$ $p = 0.002$ | $\Delta = -0.39$ $F = 0.09$ $p = 0.763$ | 21.55 (7.14) | 25.75 (6.72) | 24.82 (6.90) | $\Delta = 4.19$ $F = 9.85$ $p = 0.002$ | $\Delta = 0.93$ $F = 0.52$ $p = 0.472$ | $\Delta = 3.26$ $F = 6.26$ $p = 0.014$ |
| H2b | Decision makers with a challenging performance goal exhibit higher cumulative performance in a stable environment than those with a moderate goal. | 1095.74 (49.81) | 1094.86 (70.08) | 1133.12 (52.03) | $\Delta = 38.25$ $F = 9.73$ $p = 0.002$ | $\Delta = 37.38$ $F = 13.46$ $p = 0.000$ | $\Delta = -0.87$ $F = 0.01$ $p = 0.942$ | 1052.36 (63.52) | 1094.86 (70.08) | 1071.73 (74.11) | $\Delta = 42.51$ $F = 10.93$ | $\Delta = 23.13$ $F = 2.85$ | $\Delta = 19.37$ $F = 2.27$ |
| | Decision makers with a challenging performance goal develop greater belief strength than those with a moderate one. | 0.419 (0.170) | 0.431 (0.169) | 0.508 (0.156) | $\Delta = 0.077$ $F = 268.47$ $p = 0.000$ | $\Delta = 0.089$ $F = 349.96$ $p = 0.000$ | $\Delta = 0.012$ $F = 6.07$ $p = 0.014$ | 0.00 | 0.431 (0.169) | 0.021 (0.041) | $\Delta = 0.431$ $F = 17154$ $p = 0.000$ | $\Delta = 0.410$ $F = 15546$ $p = 0.000$ | $\Delta = 0.021$ $F = 663.3$ $p = 0.000$ |
| H3 | … belief strength measured as the difference in success rate of the best and worst alternatives. | 0.388 (0.190) | 0.394 (0.198) | 0.486 (0.185) | $\Delta = 0.091$ $F = 275.40$ $p = 0.000$ | $\Delta = 0.098$ $F = 321.06$ $p = 0.000$ | $\Delta = 0.006$ $F = 1.26$ $p = 0.261$ | 0.00 | 0.394 (0.198) | 0.021 (0.041) | $\Delta = 0.394$ $F = 10424$ $p = 0.000$ | $\Delta = 0.373$ $F = 9529$ $p = 0.000$ | $\Delta = 0.021$ $F = 663.3$ $p = 0.000$ |
| | … belief strength measured as the variance across beliefs. | 0.221 (0.088) | 0.227 (0.089) | 0.270 (0.084) | $\Delta = 0.042$ $F = 294.74$ $p = 0.000$ | $\Delta = 0.049$ $F = 384.64$ $p = 0.000$ | $\Delta = 0.006$ $F = 6.19$ $p = 0.012$ | 0.00 | 0.227 (0.089) | 0.012 (0.024) | $\Delta = 0.227$ $F = 17186$ $p = 0.000$ | $\Delta = 0.215$ $F = 15221$ $p = 0.000$ | $\Delta = 0.012$ $F = 663.3$ $p = 0.000$ |

*Notes.* This table reports sample means with standard errors in parentheses. The “Difference between conditions” column reports the ANOVA results.

**Table 5.** ANOVA Results for Study 3b – Extended Goal Conditions in a Disrupted Environment

| Hypothesis | | Low goal comparison | | | | | | Lower- and upper-edge goal comparison | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | | Goal | | | Difference between conditions | | | Goal | | | Difference between conditions | | |
| *Compared with moderate-goal condition* | | 15 ($n$ = 44) | 20 ($n$ = 46) | 25 ($n$ = 50) | 25–20 | 25–15 | 20–15 | 5 ($n$ = 60) | 20 ($n$ = 46) | 35 ($n$ = 47) | 20–5 | 20–35 | 35–5 |
| Across rounds | | | | | | | | | | | | | |
| H1 | Decision makers with a challenging performance goal exploit, i.e., select the option they believe to be best, more often than those with a moderate goal. | 33.52 (9.01) | 33.35 (9.44) | 43.54 (7.46) | $\Delta$ = 10.19 $F$ = 34.70 $p$ = 0.000 | $\Delta$ = 10.01 $F$ = 34.71 $p$ = 0.000 | $\Delta$ = −0.17 $F$ = 0.01 $p$ = 0.929 | 31.56 (12.42) | 33.35 (9.44) | 38.59 (10.26) | $\Delta$ = 1.78 $F$ = 0.65 $p$ = 0.420 | $\Delta$ = −5.24 $F$ = 6.57 $p$ = 0.012 | $\Delta$ = 7.03 $F$ = 9.80 $p$ = 0.002 |
| Pre-shock results | | | | | | | | | | | | | |
| H2a | Decision makers with a challenging performance goal select the truly best option more frequently in a stable environment than those with a moderate goal. | 15.18 (4.12) | 14.25 (4.17) | 18.64 (3.80) | $\Delta$ = 4.29 $F$ = 27.81 $p$ = 0.000 | $\Delta$ = 3.46 $F$ = 17.86 $p$ = 0.000 | $\Delta$ = −0.83 $F$ = 0.91 $p$ = 0.343 | 14.1 (4.42) | 14.25 (4.17) | 15.28 (3.81) | $\Delta$ = 0.24 $F$ = 0.09 $p$ = 0.770 | $\Delta$ = −0.93 $F$ = 1.26 $p$ = 0.265 | $\Delta$ = 1.17 $F$ = 2.10 $p$ = 0.150 |
| H2b | Decision makers with a challenging performance goal exhibit higher cumulative performance in a stable environment than those with a moderate goal. | 644.18 (37.52) | 642.97 (47.29) | 673.86 (40.35) | $\Delta$ = 30.88 $F$ = 11.90 $p$ = 0.001 | $\Delta$ = 29.68 $F$ = 13.51 $p$ = 0.000 | $\Delta$ = −1.20 $F$ = 0.02 $p$ = 0.894 | 645.16 (48.50) | 642.97 (47.29) | 639.91 (43.92) | $\Delta$ = −2.18 $F$ = 0.05 $p$ = 0.816 | $\Delta$ = 3.06 $F$ = 0.10 $p$ = 0.746 | $\Delta$ = −5.25 $F$ = 0.34 $p$ = 0.747 |
| H3 | Decision makers with a challenging performance goal develop greater belief strength than those with a moderate one. | 0.413 (0.176) | 0.432 (0.179) | 0.497 (0.147) | $\Delta$ = 0.065 $F$ = 101.27 $p$ = 0.000 | $\Delta$ = 0.084 $F$ = 171.49 $p$ = 0.000 | $\Delta$ = 0.019 $F$ = 7.33 $p$ = 0.007 | 0 | 0.432 (0.179) | 0.021 (0.05) | $\Delta$ = 0.432 $F$ = 9476 $p$ = 0.000 | $\Delta$ = 0.412 $F$ = 6200 $p$ = 0.000 | $\Delta$ = 0.021 $F$ = 284.3 $p$ = 0.000 |
| | … belief strength measured as the difference in success rate of the best and worst alternatives. | 0.390 (0.190) | 0.311 (0.266) | 0.478 (0.168) | $\Delta$ = 0.167 $F$ = 368.51 $p$ = 0.000 | $\Delta$ = 0.088 $F$ = 151.55 $p$ = 0.000 | $\Delta$ = −0.079 $F$ = 70.53 $p$ = 0.000 | 0 | 0.311 (0.266) | 0.021 (0.050) | $\Delta$ = 0.311 $F$ = 2211.2 $p$ = 0.000 | $\Delta$ = 0.290 $F$ = 1449.9 $p$ = 0.000 | $\Delta$ = 0.021 $F$ = 284.3 $p$ = 0.000 |
| | … belief strength measured as the variance across beliefs. | 0.116 (0.078) | 0.123 (0.087) | 0.158 (0.079) | $\Delta$ = 0.035 $F$ = 115.09 $p$ = 0.000 | $\Delta$ = 0.042 $F$ = 178.74 $p$ = 0.000 | $\Delta$ = 0.006 $F$ = 4.17 $p$ = 0.041 | 0 | 0.123 (0.087) | 0.002 (0.007) | $\Delta$ = 0.123 $F$ = 3242 $p$ = 0.000 | $\Delta$ = 0.121 $F$ = 2442 $p$ = 0.000 | $\Delta$ = 0.002 $F$ = 152.98 $p$ = 0.000 |

**Table 5.** (Continued)

| Hypothesis | | Low goal comparison | | | | | | Lower- and upper-edge goal comparison | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | | Goal | | | Difference between conditions | | | Goal | | | Difference between conditions | | |
| *Compared with moderate-goal condition* | | 15 ($n$ = 44) | 20 ($n$ = 46) | 25 ($n$ = 50) | 25–20 | 25–15 | 20–15 | 5 ($n$ = 60) | 20 ($n$ = 46) | 35 ($n$ = 47) | 20–5 | 20–35 | 35–5 |
| Post-shock results (rounds 31–50) | | | | | | | | | | | | | |
| H4a | Decision makers with a challenging performance goal choose the newly changed best option less frequently after a disruptive environmental shock than those with a moderate one. | 8.09 (4.34) | 7.30 (3.85) | 5.1 (3.15) | $\Delta = -2.20$ $F = 9.46$ $p = 0.003$ | $\Delta = -2.99$ $F = 14.84$ $p = 0.000$ | $\Delta = -0.78$ $F = 0.83$ $p = 0.366$ | 7.21 (4.44) | 7.30 (3.85) | 6.42 (4.15) | $\Delta = 0.09$ $F = 0.01$ $p = 0.915$ | $\Delta = 0.88$ $F = 1.12$ $p = 0.293$ | $\Delta = -0.79$ $F = 0.88$ $p = 0.632$ |
| H4b | Decision makers with a challenging performance goal exhibit worse performance immediately after a disruptive environmental shock than those with a moderate one. | 422.06 (33.97) | 416.39 (38.45) | 406.18 (34.06) | $\Delta = -10.21$ $F = 1.90$ $p = 0.171$ | $\Delta = -15.88$ $F = 5.10$ $p = 0.026$ | $\Delta = -5.67$ $F = 0.55$ $p = 0.461$ | 410.2 (37.74) | 416.39 (38.45) | 410.21 (39.62) | $\Delta = 6.19$ $F = 0.69$ $p = 0.408$ | $\Delta = 6.17$ $F = 0.58$ $p = 0.448$ | $\Delta = 0.01$ $F = 0.00$ $p = 0.998$ |
| Extended post-shock results (rounds 51–80) | | | | | | | | | | | | | |
| H2a | Decision makers with challenging performance goals select the truly best arm more often than those with moderate ones | 15.52 (5.73) | 14.39 (5.55) | 18.34 (4.15) | $\Delta = 3.95$ $F = 15.73$ $p = 0.000$ | $\Delta = 2.82$ $F = 7.57$ $p = 0.007$ | $\Delta = -1.13$ $F = 0.90$ $p = 0.344$ | 15.47 (8.26) | 14.39 (5.55) | 15.17 (6.92) | $\Delta = -1.07$ $F = 0.58$ $p = 0.45$ | $\Delta = -0.78$ $F = 0.36$ $p = 0.55$ | $\Delta = -0.30$ $F = 0.04$ $p = 0.844$ |
| H2b | Decision makers with a challenging performance goal exhibit higher cumulative performance in a stable environment than those with a moderate goal. | 655.82 (55.47) | 648.56 (55.44) | 684.86 (52.33) | $\Delta = 36.29$ $F = 10.89$ $p = 0.001$ | $\Delta = 29.04$ $F = 6.81$ $p = 0.01$ | $\Delta = -7.25$ $F = 0.38$ $p = 0.54$ | 646.62 (61.11) | 648.56 (55.44) | 655.91 (63.85) | $\Delta = 1.95$ $F = 0.03$ $p = 0.86$ | $\Delta = -7.35$ $F = 0.35$ $p = 0.55$ | $\Delta = 9.29$ $F = 0.59$ $p = 0.45$ |

*Notes.* This table reports sample means with standard errors in parentheses. The "Difference between conditions" column reports the ANOVA results.

helps the decision maker identify the inferior option (always failure) sooner.

To test this logic we compare exploitation, performance, and belief strength between decision makers with the low (15), moderate (20) and challenging goals (25). The results, summarized in Table 4, show that decision makers in the low goal condition exhibit performance, exploitation behavior, and belief strength similar to the *moderate* goal decision makers, *not* the challenging goal condition they were set up to mirror. This is surprising when we focus solely on the way that both these performance goals allow the decision makers to discriminate between some of the options more easily. One possibility is that the goal was so low that subjects reverted to do-your-best process, where current average performance is treated as an endogenous goal.[15] Simulations suggest this process most closely resembles our data (please see Online Supplement, Section 4.3).

**The Value of "Meaningless" Performance Goals in a Stable Environment.** One of the boundary conditions in the goal setting literature is that the goals must be meaningful. Yet, when neither the manager nor her superior has a useful reference point, performance goals may be set unrealistically high (or uselessly low). Here, we examine what happens when the goals are either too high or too low relative to the feedback decision makers receive from their investment options (all coefficients and comparisons are reported in Tables 4 and 5). We find that decision makers with an upper-edge goal (of 35) exploit to a similar degree as do those with a moderate goal, while those with a lower-edge goal (of 5) exploit considerably less, compared with the moderate goal condition. We see the same pattern when we compare how often decision makers facing these goals select the truly best option. Despite that similarity, we find that decision makers with both lower-edge and upper-edge goals perform similarly and worse than those with moderate goals.[16] Finally, we find that the decision makers with a moderate goal developed higher belief strength compared with those subject to the upper-edge or lower-edge goal.[17]

Goal setting theory suggests that when subject to unrealistic goals, individuals switch back to a do-your-best mode (Locke and Latham 1990). In addition, prospect theory suggests that decision makers are less sensitive to successes compared with failures. Combining these predictions with the upper-edge goal acting as an anchor, in this context it is likely that decision makers will gravitate toward the option that reduces the magnitude of failures. These two effects together may explain why decision makers facing the upper-edge goal exploit more than the lower-edge goal, but still less than the challenging goal. It also explains why decision makers facing the upper-edge goal exploit similar to the moderate goal condition, which benchmarks performance to the average payoff across the three options.

**The Effect of "Meaningless" Performance Goals in a Disrupted Environment.** The results from Study 3b in which we replicate the disrupted environment from 2b with extended goals, show similar patterns as Study 3a in the 30 rounds prior to the shock. Exploitation by participants with upper-edge, lower-edge, and low goals is similar to those with a moderate goal (see Hypothesis 1 results in Table 5), although here (at round 30 rather than round 50) we do not find a significant difference between the lower-edge goal and the others. This pattern persists when we compare how often decision makers facing these goals select the truly best option as well as their performance. However, we do find that the decision makers with a moderate goal developed higher belief strength (already at round 30) compared with those participants with the upper-edge or lower-edge goal.

In the period immediately after the shock (rounds 31–50), we found that participants' choices in the low goal (15) condition were indistinguishable from those with a moderate goal (see Hypothesis 4a results in Table 5). Since the environmental shock is particularly detrimental for participants with a challenging goal, we examine how the edge-goal conditions compare with those: We found that participants with a low goal (of 15) switched faster and performed better than those with a challenging goal (coefficients reported in Table 5). As expected given the lower belief strength, participants with the lower-edge goal as well as those with the upper-edge goal switched to the newly best option faster than did those with a challenging goal. It is plausible that the edge goals create some anchoring, that may influence participant decisions. Further research is required to fully understand how goals that are too high or too low differ from each other and from do-your-best goals in how participants choose under uncertainty.

**Robustness Checks.** Several robustness checks are presented in the Online Supplement. In Section 4 we present additional mechanism tests and robustness to different ways of calculating beliefs. Section 5 contains a discussion of boundary conditions, including studies we ran with different incentive structures, a positive environment shock, and do-your-best/no explicit goal conditions. We find that our findings are strengthened when there is a positive environment shock, where another option performs better than the current option, which itself remains unchanged. In this case, decision makers subject to a challenging goal do not explore, since the mechanism leading to exploration in the baseline turbulence condition, where the newly false positive is corrected with further sampling does not occur in this case.[18] Furthermore, we find that our results remain the same or strengthen under an incentive scheme that makes the goal more salient, such as a bonus reward for achieving the goal, but considerably weakens when the performance goal becomes less salient such as when decision makers are subject to a piece-rate incentive scheme.[19] Section 6 shows supplementary analyses to ascertain effect sizes across our studies by employing the Single Paper Meta-Analysis method (McShane and Böckenholt 2017).

## Discussion

Performance goals play a key role in guiding the decision-making process in organizations (March and Simon 1958, Cyert and March 1963). Employees are frequently evaluated against performance goals set by their superiors, such as sales or profit targets. Failure to meet these goals often has significant negative consequences for employees. Yet employees frequently do not know which actions will lead to outcomes that exceed their goals. In other words, employees face considerable uncertainty in their decision making about which actions to choose from among various alternatives, having to strike a balance between exploring new opportunities and exploiting existing knowledge (March 1991).

Our work contributes to the exploration-exploitation literature by investigating how exogenously set performance goals influence employees' decisions to explore versus exploit, and their performance consequences. We bring a learning perspective to this problem in contrast to the risk-taking perspective that is more prevalent in the empirical work on problemistic search. We thus answer the call by Denrell (2008) for further empirical research on problemistic search using a learning lens, considering much of this work in the management literature remains theoretical. Our work builds on the empirical literature on the explore-exploit problem in neuroscience and psychology under do-you-best goals and adds exogenous performance goals to this setup.

In a series of laboratory studies, we found that decision makers subject to challenging goals exploit *more* (relative to those with moderate goals), where exploitation is conceptualized as choosing the option the decision maker currently *believes to be best*. We also showed that such an exploitation focus proves beneficial in stable environments, but detrimental in unstable ones, when a shock alters the relative attractiveness of the available options. We thus add empirical traction to the study of how performance goals influence learning in nonstationary environments, which has thus far remained understudied.

The mechanism underlying this result arises from how performance goals influence the way individuals interpret feedback (e.g., Simon 1955). Individuals subject to challenging performance goals are more likely to interpret feedback from poor alternatives as failures. Therefore, they quickly develop high belief strength that the inferior alternative is worse than the superior alternative, enabling them to reduce 'useless exploration'. This mechanism suggests a boundary condition for when goals improve performance—that is, only when they allow for such discrimination. Our empirical tests support this theory.

In addition, we show that goals enhance learning only when they allow for identifying inferior options quickly; they are less effective when they allow for identifying superior options. Similarly, when goals are set unreasonably high individuals revert to a "do your best" mode. These results suggest that there may be limited downsides to setting challenging goals from a learning perspective, although our theory does not consider other psychological influences. This suggests an interesting avenue for further deepening our understanding about how low versus high performance goals influence learning under uncertainty, although in practical settings, when mangers can perform at least some kind of benchmarking, goals are likely to be set to be more challenging rather than less.

Our work also contributes to the goal-setting literature. Although that work is predominantly focused on how concrete and measurable performance goals influence behavior (Locke and Latham 2002, 2006), how such goals influence learning is so far understudied. Scholars have shown two boundary conditions to the main proposition that challenging goals increase performance; high task complexity and situations where goal attainment depends more on luck than effort. We investigate the effect of goal setting in a learning task, where individuals need to determine quickly which alternatives result in superior payoffs. We show that *challenging goals increase belief strength about the inferiority of relatively poor outcomes, thus making learning easier, and leading to improved performance. However, this very aspect leads to inferior outcomes immediately following a disruptive shock.* This finding contributes to the larger discussion around the boundary conditions for when goal setting improves performance. The idea that challenging goals may have different consequences when tasks emphasize search versus choice (learning) behavior is an important insight that requires further study.

Our study also contributes to better understanding some of the empirical inconsistencies in the goal setting literature, as well as suggesting some directions for future research. Ordóñez et al. (2009a, b), share several anecdotes about how challenging goals led managers to make risky decisions that ultimately led to poor performance. In response, Locke and Latham (2009) cite the work by Kerr and Landauer (2004), who found that setting relatively challenging goals to middle managers improved their performance at General Electric (GE). They attributed this improved performance to greater exploration by the middle managers. Recently, however, the GE story was reinterpreted to argue that challenging goals led managers to underinvest in R&D that had low short term returns in favor of over-investing in outsourcing and in GE Capital that provided immediate high returns, but was more vulnerable to shocks (Olen 2020, The Economist 2020, Gelles 2022). In unrelated work, Noda and Bower (1996) made a similar point when they argued that under discontinuous change in the telecommunications market, middle managers with challenging performance goals were more likely to exploit, which led to long term performance declines, whereas middle managers held to comparatively moderate goals explored more, leading to improved performance longer-term.

These findings from prior work are in line with our own findings. In both cases, middle-managers faced a resource allocation problem, which is akin to a situation where managers balance the explore-exploit trade-off while learning under uncertainty about the relative attractiveness of choices. Consistent with our findings, challenging goals led these managers to quickly identify the high-performing alternative, but this learning was vulnerable to environmental turbulence. There is a need for future empirical work, perhaps qualitative in nature, to understand whether our mechanism—too-rigid mental models (too strong beliefs) arising from challenging goals—were an important explanation for this failure to adapt. The Polaroid case by Tripsas and Gavetti (2000) does suggest the importance of cognitive inertia, although whether high performance goals from their historical superior performance led to these rigid mental models is unclear. More importantly, search and choice are interrelated in managerial problems. Our findings, combined with previous studies, suggest that goal-setting may have different effects on search problems versus choice problems. Future research should further examine how goal setting influences decision making under situations where both search and choice (learning) are important. One direction to explore is whether organizations can separate search tasks from choice tasks and thus use different goal manipulations for these different populations.

Our study has interesting managerial implications. Goal setting is ubiquitous in organizations, and most employees' work is geared toward meeting their goals. Such goal setting happens regardless of the exhortation to do-your-best. Managers have sales targets, although they are incentivized to sell as much as they can. Fund managers aim to beat the index, although they also have incentives for maximizing fund returns. Academics aim to publish enough to get tenure although the more they publish the better off they are. In many instances, it may not be very difficult to identify meaningful goals, that is, goals that are achievable by identifying superior strategies. However, identifying the superior strategy is typically non-trivial, since it depends on a combination of individual skills, organizational resources, and market requirements. Frequently, environmental changes make existing best practice or strategy unattractive, and employees need to look for other alternatives that they do not know very much about, triggering the learning problem. Although in new-to-the world scenarios meaningful goals may be difficult to identify, competitive benchmarking may make it more likely that such goals can be identified even in situations that are novel to the firm.

Even when it is difficult to identify meaningful goals, it is perhaps worthwhile to set challenging goals in choice (learning) problems. For example, in the race to develop a COVID vaccine, Johnson & Johnson set itself a goal of 90% efficacy, whereas Sinopharm set a goal of 70% efficacy, although they pursued very different technologies than the one pursued by Pfizer which achieved a 95% efficacy rate.[20] In this case, although efficacy rates of 70% and 80% were identified by scientists as desirable for a vaccine candidate to prevent an epidemic versus extinguish it,[21] there was no precedent to understand whether the technology platforms pursued were capable of delivering these goals. Nevertheless, meaningful goals were set against which employees could judge the adequacy of their efforts. Our study suggests that there appears to be limited downside to setting high goals. In these cases employees' appear to revert to a do-your-best mode, compared with setting goals too low where employees appear to explore more than optimal (although our mechanism ignores other plausible psychological effects of setting unachievably high goals). However, as we suggested earlier, goal-setting needs to be performed carefully when both search and choice are involved, and in this case, as far as we know, theory currently does not offer a useful guide to practice.

### Boundary Conditions and Extensions

The mechanism and results reported in this paper are certainly sensitive to various factors. In robustness checks, we have shown that incentive designs that reinforce the performance goals strengthen the mechanism at play, while those that override or replace the goals with a strong performance focus appear to diminish their influence. They are also sensitive to whether the different goals provide meaningful ways to differentiate better from clearly inferior options, and do not appear to work symmetrically for the selection of clearly superior ones. Very high performance goals that are impossible to achieve, or very low ones that are trivial to achieve, do not provide any useful signals for the decision makers. In such cases, we expect decision makers to disregard these goals and instead to endogenously develop and update their own aspiration level. In our data, it appears that individuals with the lower-edge goal explore more than those with the higher-edge goal. This may be because the endogenous aspiration level for the higher-edge goal may be higher due to anchoring. However, the mechanism(s) underlying these effects warrant further investigation.

Finally, we have highlighted a number of open questions and avenues for future research that naturally follow from our study of how performance goals influence the adaptive learning process, explore-exploit choices, and performance. There is a need to develop theory around how organizations adapt to externally imposed aspirations, such as from investors and analysts. Our research setup implied that actors take the performance goal as a given, which holds more likely true at the individual level than at the organization level. A long tradition in behavioral research has considered how aspirations adapt to performance (March 1988), especially at the firm level (Greve 2003, Audia and Greve 2021). Managers' choices may be informed not just by feedback, but also by others' choices and preferences, such as imitative behavior or deference to the market (Brandenburger and Polak 1996, Greve 2009, Levine and Zajac 2023). For example, such forces can influence belief formation by combining experiential and social processes. Prior work has mainly considered historical and social processes as influencing aspirations or goals, rather than beliefs about the relative attractiveness of choices. As such future work can investigate these processes. Since managers often face multiple goals, considering how differing levels of multiple goals influence adaptive behavior (Gaba and Greve 2019, Audia and Greve 2021) is an important but underexplored area, although one of significant importance, since many firms use the balanced scorecard or similar approaches to goal setting. Finally, similar to prior empirical work in the explore-exploit problem, we too are unable to measure participants' beliefs directly, but instead infer them by fitting their choices and associated payoffs in a reinforcement learning model. More work is required to identify innovative techniques that may allow us to measure beliefs directly in such learning tasks.

## Conclusions

Overall, the aim of this paper was to examine whether and how performance goals matter to the adaptive learning process. We show that the performance goal affects the way individuals interpret feedback which in turn affects their belief strength about the relative superiority of the available alternatives. This mechanism helps us explain why challenging (versus moderate) performance goals improve performance in stable choice situations, but hinder adaptation and reduce performance in changing environments. By highlighting this core relationship, we hope to have opened new avenues of future enquiry into the underlying processes.

## Acknowledgments

The authors would like to thank the special issue editors and six anonymous reviewers for their many valuable suggestions. We would also like to thank Oliver Baumann, Julien Clement, Jerker Denrell, Nilanjana Dutt, Olga Hawn, Reddi Kotha, Tanya Menon, Amandine Ody-Brasier, Hart Posen, Phanish Puranam, Milind Sohoni, Elena Vidal, and Maciej Workiewicz for their valuable comments on previous versions of the paper.

## Endnotes

[1] Prior work uses terms such as adaptive learning, reinforcement learning, learning from feedback, and learning by doing to describe this iterative process; we too use these terms interchangeably.

[2] Goals that are meaningless (e.g., too high or too low relative to the payoffs received) may not help individuals differentiate between better versus worse options, since then all outcomes will appear similarly attractive or not.

[3] Note that what organization theory scholars label search tasks, goal setting scholars label as learning tasks.

[4] See https://www.jnj.com/innovation/questions-about-johnson-johnson-investigational-covid-19-vaccine, published on 01/05/2021, accessed on March 25, 2023.

[5] They could of course run multiple experiments at the same time. For explanatory convenience, we treat this as a sequential choice process.

[6] Unlike the games described in many goal-setting studies (cf: Heath et al. 1999), in this game, the three options have equal risk (i.e., their variance in payoffs is identical). This setup follows Denrell's (2008) call to treat the firms' choice between exploration-exploitation as an issue of learning rather than risk taking, the latter approach being more common in the empirical literature in the behavioral theory of the firm.

[7] This begs the interesting question of what happens when the goals are not meaningful (too high or too low). They cannot help with learning, since they do not meaningfully categorize the payoff space. We will explore this further in the boundary conditions of the study.

[8] Recall our prior discussion of the two ways belief strength increases in adaptive learning processes; pathway (a) and pathway (b). We do not repeat them here for the sake of brevity. With pathway (b), given adequate time, most individuals will achieve high belief strength that the truly inferior option is worse than the truly superior option regardless of their performance goal. However, individuals with challenging goals are likely to reach high belief strength earlier than those with moderate ones, because the challenging goal strengthens pathway (a).

[9] Bandit tasks have been used extensively to study exploration-exploitation behavior in both theoretical and empirical studies. As Simon (1947) stressed, exploration-exploitation decisions are inherently behavioral because prior probabilities are unknown and must be learned in a trial-and-error process. And because choice and learning processes are intertwined in these problems (March 1996, Sutton and Barto 1998), experimental studies can be used to make valuable inferences by precisely controlling what information is available (Sterman 1989, Edmonds 2001, Schunk 2009).

[10] See https://osf.io/atqyd/?view_only=303dbaae51da4c6bb2caa7f748fc0d17.

[11] These numbers (goals and cumulative points) were adjusted across different boundary condition studies according to the performance goal given, or the number of rounds the game was played. Please refer to Section 1 in the Online Supplement for all verbatim instructions.

[12] In the no-goal condition reported in the robustness checks we classified a pull as a success if the outcome was at least equal to the average performance until that moment, or at least equal to 20 if it was the first pull.

[13] We thank an anonymous reviewer for pointing out this potential confound. The correlation table is shown in the Online Supplement, Section 4.

[14] It is plausible that the participant with the lower goal does not choose the truly best option and has lower performance because they satisfice based on their lower goal. In that case, they should switch from the "explore" mode to "exploit" sooner, as suggested by the tau-switch model. These results directly test this idea and find that it is not supported in our data. We thank an anonymous reviewer for pointing this out.

[15] Exit surveys support this assumption; see Online Supplement, Section 1.

[16] We suspect that the non-significant difference in performance between the upper- and lower-edge goal is due to lack of power (the *p*-value is 0.13); since individuals under both conditions likely set their own endogenous goals, it is likely that any performance differences between them are smaller, and require more power to identify.

[17] The upper-edge goal also develops higher belief strength than the lower-edge goal, since some decision makers do receive some successes from option 1 (when it pays 35, which is equal to the goal). However, for the lower edge goal, the decision maker only receives successes, since the lower value they receive is equal to the goal.

[18] We thank Jerker Denrell for suggesting this test.

[19] We thank the editors for suggesting these tests.

[20] See https://www.jnj.com/innovation/questions-about-johnson-johnson-investigational-covid-19-vaccine, and https://www.businesswire.com/news/home/20201118005595/en/.

[21] See Bartsch et al. (2020) and Meng et al. (2021).

## References

Audia PG, Greve HR (2021) *Organizational Learning from Performance Feedback: A Behavioral Perspective on Multiple Goals* (Cambridge University Press, Cambridge, UK).

Bartsch SM, O'Shea KJ, Ferguson MC, Bottazzi ME, Wedlock PT, Strych U, McKinnell JA, et al. (2020) Vaccine efficacy needed for a COVID-19 coronavirus vaccine to prevent or stop an epidemic as the sole intervention. *Amer. J. Prev. Med.* 59(4):493–503.

Bower JL (1970) *Managing the Resource Allocation Process* (Harvard Business School, Division of Research, Boston).

Brandenburger A, Polak B (1996) When managers cover their posteriors: Making the decisions the market wants to see. *RAND J. Econom.* 27(3):523.

Christensen M, Sørensen SK, Knudsen T (2021) Tax bandits: Learning in an interactive environment, University of Southern Denmark Working Paper.

Cohen JD, McClure SM, Yu AJ (2007) Should I stay or should I go? How the human brain manages the trade-off between exploitation and exploration. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 362(1481):933–942.

Cook KS, Yamagishi T (2008) A defense of deception on scientific grounds. *Soc. Psych. Quart.* 71(3):215–221.

Cyert R, March JG (1963) *A Behavioral Theory of the Firm* (Prentice-Hall, Englewood Cliffs, NJ).

Daw ND, O'Doherty JP, Dayan P, Seymour B, Dolan RJ (2006) Cortical substrates for exploratory decisions in humans. *Nature* 441(7095):876–879.

Denrell J (2003) Vicarious learning, undersampling of failure, and the myths of management. *Organ. Sci.* 14(3):227–243.

Denrell J (2008) Organizational risk taking: Adaptation vs. variable risk preferences. *Indust. Corporate Change* 17(3):427–466.

Denrell J, Liu C (2021) When reinforcing processes generate an outcome-quality dip. *Organ. Sci.* 32(4):1079–1099.

Denrell J, March JG (2001) Adaptation as information restriction: The hot stove effect. *Organ. Sci.* 12(5):523–538.

Earley PC, Connolly T, Ekegren G (1989a) Goals, strategy development, and task performance: Some limits on the efficacy of goal setting. *J. Appl. Psych.* 74(1):24–33.

Earley PC, Connolly T, Lee C (1989b) Task strategy interventions in goal setting: The importance of search in strategy development. *J. Management* 15(4):589–602.

Ederer F, Manso G (2013) Is pay for performance detrimental to innovation? *Management Sci.* 59(7):1496–1513.

Edmonds B (2001) Toward a descriptive model of agent strategy search. *Comput. Econom.* 18(1):111–133.

Gaba V, Greve HR (2019) Safe or profitable? The pursuit of conflicting goals. *Organ. Sci.* 30(4):647–667.

Gary MS, Yang MM, Yetton PW, Sterman JD (2017) Stretch goals and the distribution of organizational performance. *Organ. Sci.* 28(3):395–410.

Gelles D (2022) *The Man Who Broke Capitalism* (Simon & Schuster, New York).

Greve HR (2003) A behavioral theory of R&D expenditures and innovations: Evidence from shipbuilding. *Acad. Management J.* 46(6):685–702.

Greve HR (2009) Bigger and safer: The diffusion of competitive advantage. *Strategic Management J.* 30(1):1–23.

Heath C, Larrick RP, Wu G (1999) Goals as reference points. *Cognit. Psych.* 38(1):79–109.

Hoppe F (1930) Untersuchungen zur Handlungs- und Affektpsychologie. *Psych. Forsch.* 14(1):1–62.

Kanfer R, Ackerman PL (1989) Motivation and cognitive abilities: An integrative/aptitude-treatment interaction approach to skill acquisition. *J. Appl. Psych.* 74(4):657–690.

Kerr S, Landauer S (2004) Using stretch goals to promote organizational effectiveness and personal growth: General Electric and Goldman Sachs. *Acad. Management Perspect.* 18(4):134–138.

Larrick RP, Heath C, Wu G (2009) Goal-induced risk taking in negotiation and decision making. *Soc. Cogn.* 27(3):342–364.

Laureiro-Martínez D, Canessa N, Brusoni S, Zollo M, Hare T, Alemanno F, Cappa S (2014) Frontopolar cortex and decision making efficiency: Comparing brain activity of experts with different professional background during an exploration-exploitation task. *Front. Hum. Neurosci.* 7. 10.3389/fnhum.2013.00927.

Lee E, Puranam P (2016) The implementation imperative: Why one should implement even imperfect strategies perfectly. *Strategic Management J.* 37(8):1529–1546.

Lee MD, Zhang S, Munro M, Steyvers M (2011) Psychological models of human and optimal performance in bandit problems. *Cogn. Syst. Res.* 12(2):164–174.

Lejuez CW, Read JP, Kahler CW, Richards JB, Ramsey SE, Stuart GL, Strong DR, Brown RA (2002) Evaluation of a behavioral measure of risk taking: The Balloon Analogue Risk Task (BART). *J. Exp. Psych. Appl.* 8(2):75–84.

Levine SS, Zajac EJ (2023) The other invisible hand: How markets—as institutions—propagate conformity and valuation errors. *Strategy Sci.* 8(3):323–348.

Lewin K, Dembo T, Festinger L, Sears PS (1944) Level of aspiration. Hunt JM, ed. *Personality and the Behavior Disorders* (Ronald Press, Oxford, UK), 333–378.

Locke EA, Latham GP (1990) *A Theory of Goal Setting and Task Performance* (Prentice-Hall, Inc, Upper Saddle River, NJ).

Locke EA, Latham GP (2002) Building a practically useful theory of goal setting and task motivation. *Amer. Psych.* 57(9):705–717.

Locke EA, Latham GP (2006) New directions in goal-setting theory. *Curr. Dir. Psych. Sci.* 15(5):265–268.

Locke EA, Latham GP (2009) Has goal setting gone wild, or have its attackers abandoned good scholarship? *Acad. Management Perspect.* 23(1):17–23.

Manso G (2011) Motivating innovation. *J. Finance* 66(5):1823–1860.

March JG (1988) Variable risk preferences and adaptive aspirations. *J. Econom. Behav. Organ.* 9(1):5–24.

March JG (1991) Exploration and exploitation in organizational learning. *Organ. Sci.* 2(1):71–87.

March JG (1996) Learning to be risk averse. *Psych. Rev.* 103(2):309–319.

March JG (2003) Understanding organisational adaptation. *Soc. Econom.* 25(1):1–10.

March JG, Simon HA (1958) *Organizations* (Wiley, Hoboken, NJ).

McShane BB, Böckenholt U (2017) Single-paper meta-analysis: Benefits for study summary, theory testing, and replicability. *J. Consum. Res.* 43(6):1048–1063.

Mehlhorn K, Newell BR, Todd PM, Lee MD, Morgan K, Braithwaite VA, Hausmann D, Fiedler K, Gonzalez C (2015) Unpacking the exploration–exploitation tradeoff: A synthesis of human and animal literatures. *Decision (Wash. D.C.)* 2(3):191–215.

Meng Z, Shan S, Zhang R (2021) China's COVID-19 vaccination strategy and its impact on the global pandemic. *Risk Manag. Healthc. Policy* 14:4649–4655.

Noda T, Bower JL (1996) Strategy making as iterated processes of resource allocation. *Strategic Management J.* 17(S1):159–192.

Olen H (2020) Jack Welch's toxic legacy. *Washington Post* (March 3rd).

Ordóñez LD, Schweitzer ME, Galinsky AD, Bazerman MH (2009a) Goals gone wild: The systematic side effects of overprescribing goal setting. *Acad. Management Perspect.* 23(1):6–16.

Ordóñez LD, Schweitzer ME, Galinsky AD, Bazerman MH (2009b) On good scholarship, goal setting, and scholars gone wild. *Acad. Management Perspect.* 23(3):82–87.

Posen HE, Levinthal DA (2012) Chasing a moving target: Exploitation and exploration in dynamic environments. *Management Sci.* 58(3):587–601.

Sang K, Todd PM, Goldstone RL, Hills TT (2020) Simple threshold rules solve explore/exploit trade-offs in a resource accumulation search task. *Cogn. Sci.* 44(2):1–30.

Schunk D (2009) Behavioral heterogeneity in dynamic search situations: Theory and experimental evidence. *J. Econom. Dynam. Control.* 33(9):1719–1738.

Seijts GH, Latham GP (2005) Learning vs. performance goals: When should each be used? *Acad. Management Perspect.* 19(1): 124–131.

Simon HA (1947) *Administrative Behavior: A Study of Decision-Making Processes in Administrative Organization* (Macmillan, New York).

Simon HA (1955) A behavioral model of rational choice. *Quart. J. Econom.* 69(1):99–118.

Song M, Bnaya Z, Ma WJ (2019) Sources of suboptimality in a minimalistic explore–exploit task. *Nat. Hum. Behav.* 3(4):361–368.

Sterman JD (1989) Modeling managerial behavior: Misperceptions of feedback in a dynamic decision making experiment. *Management Sci.* 35(3):321–339.

Sutton RS, Barto AG (1998) *Reinforcement Learning: An Introduction* (MIT Press, Boston).

The Economist (2020) Jack Welch transformed American capitalism as boss of GE. *The Economist* (March 5th).

Toplak ME, Sorge GB, Benoit A, West RF, Stanovich KE (2010) Decision-making and cognitive abilities: A review of associations between Iowa Gambling Task performance, executive functions, and intelligence. *Clin. Psych. Rev.* 30(5):562–581.

Tripsas M, Gavetti G (2000) Capabilities, cognition, and inertia: Evidence from digital imaging. *Strategic Management J.* 21(1011): 1147–1161.

**Marlo Raveendran** is an Associate Professor of Management at University of California, Riverside. She received her PhD from London Business School, University of London. Her research focuses on the micro-foundations of organization design and the consequences of division of labor.

**Kannan Srikanth** is an Associate Professor of Strategy in the Department of Management and Human Resources at Fisher College of Business, The Ohio State University. He received his PhD from London Business School, University of London. His research focuses on tradeoffs and complementarities in organization design choices.

**Tiberiu Ungureanu** is an Assistant Professor in the Management Department at Appalachian State University. He received his PhD from The Ohio State University. His research focuses on how organizational structure influences managerial resource allocations, especially between exploration and exploitation activities.

**George L. Zheng** is an Assistant Professor in the Department of Strategy, Innovation and Entrepreneurship at the Shanghai University of Finance and Economics. He received his PhD from Singapore Management University. His research focuses on the behavioral theory of the firm, organization learning and design, specifically using computational models.