Routine and Adaptive Expertise in Working Learners

Daniel L. Schwartz & Kristen P. Blair
Graduate School of Education, Stanford University

The earliest studies of workplace performance examined the speed at which people could perform repetitive tasks that ranged from teletype operations to rolling cigars (e.g., William & Harter, 1899). The research revealed the power law of learning, also known as the 80/20 rule. People can achieve about 80% of possible efficiency within the first 20% of their time on task. It takes much longer to squeeze out the remaining 20% of inefficiencies in performance. The power law of learning has informed many training efforts including the design of the earliest intelligent tutors in mathematics (e.g., Croteau, Heffernan, & Koedinger, 2004).

Early psychological research focused on the mastery of a defined body of highly cued and repeatable behaviors. The focus on mastery was well-suited to an industrial era when there were few machines that could execute what humans could. The work landscape is changing, and even if artificial intelligence does not displace droves of workers, there is an increased premium on learning new jobs, solving novel problems, handling emergent and social complexity, and keeping up with advances and changes in practice. Our goal is to provide a clarifying framework for understanding the tension between highly efficient behaviors versus abilities to adapt and innovate to new circumstances, and to indicate where we need new research.

Routine and Adaptive Expertise

The late Japanese psychologist Giyoo Hatano generated a useful distinction between *routine expertise* and *adaptive expertise* through his studies of abacus masters (Hatano & Inagaki, 1986). Abacus masters – people who win abacus competitions – have prodigious memories for numbers. Hearing a new number every 2.5 seconds, they can compute the following in their heads and in real-time *without* an actual abacus:

28,596 + 847,351,654 – 166,291 – 324,008,909 + 74,886,215 – 8,672,214 + 54,221 – 91,834 – 103,682,588 + 17,274 – 212,974,008 + 4,081,123 – 56,315,444 + 897,294 – 380,941,248

The abacus masters learned to simulate an abacus in their imaginations, which helped them keep track of the numbers (by moving the beads in their imagination). It takes about one year of deliberate practice to add each new column to a mental abacus.

Notably, the abacus masters were no better than the rest of us at memorizing lists of words or fruits. This implies they did not have a superior general memory capacity that enabled their stunning performance. Instead, they had amassed tremendous hours of focused practice. The degree of practice is a better predictor of expertise than native intelligence.

Significantly, the abacus masters did not use their skills to learn new kinds of mathematics, and they only performed in highly stable contexts of zero distraction. Hatano described the abacus masters as having a high degree of *routine expertise*. The word "routine" is not pejorative. The abacus masters could perform well-rehearsed *routines* in a stable environment to an astonishing degree of efficiency. Their performances are fitting analogs for outcomes that the earliest studies of human performance emphasized: speed, accuracy, and low variability.

Importantly, Hatano went on to propose a second kind of expertise, which he termed adaptive expertise. In addition to stable performance in routine situations, adaptive experts are willing to approach novel problems and flexibly generate new solutions. They are prepared to learn from new situations and to recognize the limits of their current knowledge. A nice example comes from Wineburg's (1998) comparison of history professors and college students. He asked historians who had expertise in a particular domain (e.g. Asian history) to solve history problems that, for them, were non-routine because they came from an unfamiliar domain. The problems involved interpreting decisions made by Abraham Lincoln. The history experts were much more likely than college students to resist making assumptions that readily came to mind based on knowledge of their current culture. The experts realized that these assumptions were indeed coming from their current cultural context rather from the context at the time of Lincoln. They therefore took the time to research these issues to learn what they needed to know to solve the problem. In contrast, the college students went merrily on their way building confidently on a set of flawed assumptions that came from their current knowledge of the world.

Many of today's jobs require continued learning, alongside flexibility, collaborative skills, and the ability to successfully deal with changing technology and organizational structures. These are hallmarks of adaptive expertise. The science of learning can already say a lot about the kinds of experiences that will support efficient routine expertise. New research is needed to better understand how to encourage and assess a trajectory towards adaptive expertise, both before starting a new job and while on the job.

Training for Efficiency

The distinctions between routine and adaptive expertise map onto two central, yet often competing, goals of education: (1) The development and use of efficient knowledge to solve recurrent classes of problems; (2) The ability to adapt to novel conditions by innovating new knowledge. These two classes of outcomes show up in many debates within education, such as the relative value of directly telling students the most efficient way to think and act versus asking students to discover new-to-them knowledge. Are these goals mutually exclusive, or is there a framework that can help learners develop knowledge that supports both efficiency and innovation?

Efficiency is characterized by rapid retrieval, accurate application, and high consistency in performance. Doctors who have frequently performed a particular type of surgery have highly efficient knowledge. They can diagnose and treat a new patient quickly and effectively. When choosing a surgeon for a procedure, it is wise to ask, "How many of these have you performed previously?" Efficient knowledge can help people turn non-routine, difficult-to-solve problems into routine problems that can be solved quickly and easily. Phrased another way, efficiency-oriented practice is often about "problem elimination" rather than about in-depth, sustained problem solving. By preparing people so that the problems they will face in life are essentially routine problems—or at worst very "near transfer" problems—people reduce the need for extensive problem solving and can perform quite effectively.

Scientists of learning have learned a lot about how to train for efficiency in knowledge and skills. Examples include well established memory strategies, such as the generation and testing effects, which show that having people practice retrieving information from memory leads to far better retention than re-reading or restudying that same information. Or that spacing learning opportunities over time, rather than massing them together, is more effective for long term retention. Research on worked examples provides insights into how to help novice learners efficiently move toward proficient performance, while studies of deliberate practice specify the kinds of learning experiences that allow experts to overcome procedural humps, such as an expert musician who reaches a plateau.

An important remaining question asks how to embed learning opportunities for becoming more efficient while on the job. One study found that 20th-year clinical psychologists were no more accurate than 1st-year psychologists in their diagnoses and selections of treatment (Sapyta, Riemer & Bickman, 2005). The reason was that the organization of the clinical practice did not include mechanisms for objective and precise feedback on successes and failures. Patients would often never return, and the psychologists could assume they were doing just fine. Subsequent efforts to change the feedback mechanisms within clinical practice improved performance. Building in performance support for continued learning should be a significant goal of designing environments for working learners.

Limitations of Efficiency

It is tempting to place efficiency and innovation at opposite poles of a continuum, given that adaptation often requires letting go of previously efficient routines. Our argument is that efficiency does not need to be the enemy of innovation. For example, it is well known that efficiency in some processes (e.g. learning to drive a car, learning to decode written words and sentences) frees attentional capacity to do other things (e.g., talking while driving, reading for meaning). Similarly, if people, confronted with a new complex problem, have solved aspects of it before, this helps make these sub-problems routine and easy to solve. This frees attentional bandwidth and enables people to concentrate on other aspects of the new situation that may require non-routine adaptation.

Yet, efficient knowledge is inadequate for variable environments and changing times, where what one knows is no longer fully applicable. In fact, highly efficient knowledge can interfere with effective performance. During a summer Olympics some years ago, female gymnasts were crashing during their vaults. As it turned out, the vaulting horse had been set 2” lower than usual, which was enough to disrupt the high efficiency of these world-class athletes.

Training that is dedicated to routine expertise makes great sense for recurrent tasks that take a similar form each time. However, by itself, this training cannot lead people to adaptive expertise, because the lessons are tuned to specific problem conditions. Another approach is to add in instruction that emphasizes general, content-free skills of critical thinking and problem solving. They, however, appear to provide a set of methods that are too inefficient for the large problem spaces found in many real-world tasks. Generic creativity will not lead to adaptive expertise, and in the end, these methods are often taught and assessed as a set of routinized scripts.

Balancing Innovation and Efficiency

While we know much about how to develop efficient knowledge, we know less about how to set up learning experiences that foster abilities to continue learning, innovate, and adapt. Research has suggested some ideas: (1) reduce the risk associated with sub-optimal performance, so people do not retreat into the safety of conventional and previously efficient routines; (2) provide situations that have sufficient variation, so that people develop general knowledge that can handle new situations; (3) help learners understand the “why” and not just the “what” or “how;” (4) create a culture that embraces understanding and experimentation, not just short-term performance (Bransford & Schwartz, 1999).

Figure 1 provides a possible framework for balancing experiences that promote efficiency and innovation. The lower-left describes the starting point of a novice, who is both low on efficient knowledge and abilities to adapt intelligently. To the lower right is routine expertise, where people have developed a set of highly efficient routines for handling familiar problems. In the upper right is adaptive expertise, where people have both highly efficient knowledge as well as foundational knowledge and dispositions that enable them to innovate solutions and adapt to new situations. Finally, there is the annoying novice. You probably know this person. They are the ones in the meeting who merrily brainstorming utterly unworkable solutions because they do not know anything of substance.

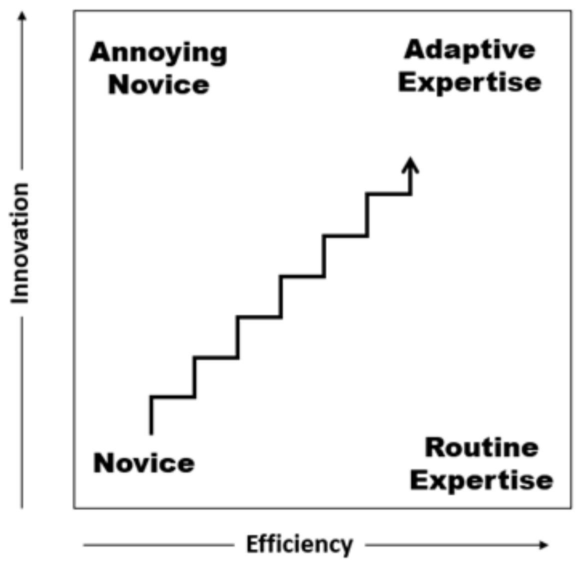

Figure 1. A framework for considering two major dimensions of learning and outcomes. (Adapted from Schwartz, Bransford & Sears, 2005)

Our reading of the learning-science literature suggests that it will not work to give learners a course of efficiency-oriented, content-filled tasks that follow the horizontal axis of Figure 1, and a separate set of strategy-training tasks that fall along the vertical axis. It is not enough to expose learners to content courses, on the one hand, and thinking courses on the other, and then help learners integrate them in a “capstone” course at the end of some educational program. This could be somewhat helpful, of course, but the conjecture is that it is far from ideal.

If learners are presented a course of study with an efficiency orientation and only late in the process are presented with opportunities to innovate, there is a good chance they will have not learned the content in a way that gives them the kinds of understandings necessary to innovate, such as a deep and flexible understanding of how features of the problem and solution space map onto each other. One of our relatives majored in astronomy, but her first 3.5 years were spent learning mathematics. She never looked through a telescope to make her own discoveries until the final semester. She had not learned the mathematics in the service of discovery, so the front-loading of mathematics courses independent of problem solving in the domain did not work very well. Not surprisingly, she did not choose a profession in astronomy.

Our belief is that instruction needs to commingle efficiency and innovation opportunities if the goal is adaptive expertise. This leads to the question of whether experiences fostering innovation versus driving towards efficiency should come first. Our claim from our own research, though the boundary conditions of this claim are unknown, is that learners should first engage in a period of innovation, and then learn the efficient solutions originally innovated by experts to address the problems. Exploring and innovating within a problem space can help people learn about the important variations and problems to be solved, as well as how to approach problem solving in the domain itself. In turn, this can set the stage for helping them understand the rationale for the efficient solution, which enables them to learn the solution more quickly and use it more appropriately. Over time, learners engage in cycles of innovation and efficiency-oriented activities that build in complexity and result in the stair-step trajectory toward adaptive expertise shown in Figure 1.

The alternative ordering, where students learn the efficient solution first, can have the inadvertent side effect of overshadowing learner abilities to recognize the need to go beyond the efficient knowledge. For instance, in several studies, learners received a worksheet that showed several variations of a common situation (Schwartz, Chase, Oppezzo & Chin, 2011). In an efficiency-first condition, learners were first told how to handle situations like these, and then applied the lesson to the worksheet. In an innovation-first condition, students were not told how to handle the situations, but instead, they had to innovate their own solutions. Afterwards, both groups received a relevant lecture on the nature of the situations and the appropriate solutions. Ultimately, both groups spent about the same amount of time overall, and they both learned the efficient solution procedures quite well. Yet, despite their equivalent trajectories towards routine expertise, there were strong differences relevant to their trajectories towards adaptive expertise. When asked to recreate the worksheet from memory, the efficiency-first students recalled the critical features of the situation to a much lesser degree than the innovation-first students. The efficiency-first students were so focused on applying the solution they had been told, they never noticed the diagnostic features of the situations. Their knowledge of the appropriate procedures overshadowed their attention to the nature of the problem. In turn, this had strong implications for future learning. When given a new class of situations that did not look like the original worksheet but were still highly relevant, the innovation-first students were more likely to use their gained knowledge to learn how to handle the new situations. In contrast, the efficiency-first students were lost, because they never saw the underlying structure that made the first and second worksheets the same problem.

Importantly, the benefits of an "innovation first" pedagogy were quite strong for lower performing students. Often when learners have difficulty, there is a tendency to retreat into efficiency-oriented instruction where we describe exactly how to do something, perhaps speaking more slowly to make sure we can get them back to basics. Yet, in the foregoing research, giving the lower-performing students the chance to innovate, before telling them the efficient answers, had a magnifying effect on their abilities to learn the content. At first, learners resisted the new way of learning, because they were afraid of failing; but by creating a classroom culture that gave them space to explore and make mistakes, they were able to deeply engage the content and ways of problem solving. This prepared them to learn the efficient answers and become more flexible in their abilities to generate new, effective ideas later on.

This research occurred in the context of an explicitly designed educational setting where it was relatively easy to create time for innovation. A major challenge for supporting working learners will be how to free up time and space for them to engage in innovative problem solving that may not yield immediate economic return in the short run, but that should be quite valuable in the long run.

Need for New Ways of Assessing for Adaptive Expertise

Many industries, and all traditional schools, apply tests of learner knowledge and abilities. Assessments can be used for many goals ranging from selecting personnel to generating feedback to instructors and learners. In all cases, most assessments of knowledge and skills emphasize the qualities of routine expertise. The tests often take a *sequestered problem-solving* format. Like a jury sequestered from contaminating influences, learners need to solve problems without resources for learning or feedback. If they have built up an efficient store of prior knowledge, they can solve relatively familiar classes of problems correctly the first time and under time pressure.

Sequestered tests that solely measure efficiency are a mismatch to the goals of adaptive expertise. An empirical example of the limitations of efficiency-driven exams comes from medical licensure. Medical exams do a poor job of predicting whether physicians have been prepared to learn on the job. A more sensitive measure, relevant to learner abilities to adapt, would use assessments that provide students with an opportunity to learn new-to-them information as part of the test (Mylopoulos, et al., 2016). We have termed these as *preparation for future learning* (PFL) assessments (Bransford & Schwartz, 1999). PFL assessments, which provide resources for learning, differ from measures of content-lite problem-solving flexibility. Tests of creative problem solving, such as "Do two people in New York City have the same number of hairs on their head," are interesting, but they do not capture people's abilities to learn new information that can help them be more effective in the future.

An example from classroom studies helps to clarify the point. We asked high school students to innovate their own ways to measure the consistency (variability) of different phenomena (Schwartz & Martin, 2004). Afterwards, students received a lecture on a standard, efficient method for computing variability. Of particular interest was whether these learning experiences would prepare students to learn to innovate solutions to highly novel problems. To find out, the high schoolers were compared to the performance of top-20 university students who had recently completed a full semester of "efficiency-oriented" statistics. On standard measures of the students' abilities to solve word problems about variability, the high school and college students looked similar. The telling finding came from an additional problem that required working with two variables, not just one. During instruction, students had only learned about working with univariate data, so determining the relation of two variables (covariance) would require innovation and learning about the structure of the situation that presented the two variables. At posttest, 34% of the high-school students invented a way to measure covariance. Though this is far below 100% (it was a difficult task), it is a high level compared to the 12% found with college students, who complained that they had not been taught how to do that yet.

It seems unlikely that the high school students had more sophisticated, content-free techniques for innovation than the college students did. It also is unlikely that they had more efficient knowledge of the statistical mathematics than the college students. Instead, the high schoolers had developed an understanding and stance towards the topic of variability that prepared them to be innovative. Without a measure of their abilities to learn new solutions to a related problem, the hidden benefit of asking the high schoolers to innovate their own solution first would have been missed.

While the example comes from a traditional school setting, it should be possible to map the spirit of PFL assessments to workforce learning. If we want to predict whether doctors are likely to learn from their practice, a useful assessment would determine whether and how they use resources for diagnosing a set of symptoms that are novel to them, such as using medical libraries, colleagues, and clinical interviews. If we want an on-going, on-the-job assessment of whether working learners are on a trajectory to adaptive expertise, we would not only evaluate how well they are doing on their routine tasks. We would also attend to how they use resources to help them handle new-to-them problems.

Going Beyond Knowledge and Skills: Learners in a Social Context

We have spent much of this brief focusing on knowledge and skills. Equally important are fostering lines of research that focus on describing the social contexts, norms, and social skills of working learners, including the material and psychological consequences for learning. Moving into a job that depends on collaborative negotiations demands a set of social skills that someone new to this kind of work may never have had a chance to learn. While there are training programs for social skills, they often have the quality of providing scripts for well-defined situations. We do not know enough about how to help people develop adaptive social expertise, so they continue to adapt and learn effective social skills and attitudes from experience.

It is also important to generate evidence on the qualities of social environments that encourage a trajectory towards adaptive expertise. Several years ago, one of us had the opportunity to consult with an Army base that was interested in improving the soldiers' abilities to innovate. Unfortunately, the base had a zero-error policy, which means they could never foster innovation among the soldiers. Adaptive expertise depends on opportunities to experiment without risk of reprimand (Hatano & Inagaki, 1986).

A growing body of research identifies how a negative or ambiguous social climate is especially damaging for learners who have been previously, or are currently, under-resourced or marginalized (e.g., Walton & Cohen, 2011, Cadaret et al., 2017). One's sense that "I belong here and can thrive" has major positive implications for learning and willingness to participate. Conversely, psychological responses, such as feeling stereotype threat, mistrust, or the sense that one does not belong, can siphon off mental resources needed for learning (and health) (Inzlicht & Schmander, 2012; Yeager et al., 2014).

An important line of research would document the particularly pernicious effects, and potential ameliorations, of a hostile social environment on a trajectory towards adaptive expertise. Our hypothesis is that negative social environments not only drain mental resources for learning routine skills, they also deeply cripple attitudes and actions needed to be on a trajectory towards adaptive expertise. Pursuing adaptive expertise depends on a willingness to engage in experimentation, take on more challenging tasks, and seek feedback. These all incur a risk of short-term failure and social approbation. The texture of one's knowledge has a major effect on one's abilities to learn and adapt, but over the long haul, the path towards adaptive expertise depends greatly on the choices and resources to do so. Being able to measure the effect of different social environments on the choices of different groups of working learners would be an important step in advancing the science and design of learning.

References

Bransford, J. D. & Schwartz, D. L. (1999). Rethinking transfer: A simple proposal with multiple implications. In A. Iran-Nejad and P. D. Pearson (Eds.), *Review of Research in Education, 24*, 61-100. Washington, D.C.: American Educational Research Association.

Cadaret, M. C., Hartung, P. J., Subich, L. M., & Weigold, I. K. (2017). Stereotype threat as a barrier to women entering engineering careers. *Journal of Vocational Behavior*, *99*, 40-51.

Croteau, E. A., Heffernan, N. T., & Koedinger, K. R. (2004). Why are algebra word problems difficult? Using tutorial log files and the power law of learning to select the best fitting cognitive model. In *International Conference on Intelligent Tutoring Systems* (pp. 240-250). Springer, Berlin, Heidelberg.

Hatano, G. & Inagaki, K. (1986). Two courses of expertise. In H. Stevenson, H. Azuma, & K. Hakuta (Eds). *Child development and education in Japan* (pp. 262-272). NY: Freeman.

Inzlicht, M., & Schmader, T. (2012). *Stereotype threat: Theory, process, and application*. Oxford University Press.

Mylopoulos, M., Brydges, R., Woods, N. N., Manzone, J., & Schwartz, D. L. (2016). Preparation for future learning: a missing competency in health professions education?. *Medical education*, *50*(1), 115-123.

Sapyta, J., Riemer, M., & Bickman, L. (2005). Feedback to clinicians: Theory, research, and practice. *Journal of clinical psychology*, *61*(2), 145-15

Schwartz, D. L., Bransford, J. D., Sears, D. L. (2005). Efficiency and innovation in transfer. In J. Mestre (Ed.), *Transfer of learning from a modern multidisciplinary perspective* (pp. 1 - 51). CT: Information Age Publishing.

Schwartz, D. L., Chase, C. C., Oppezzo, M. A., & Chin, D. B. (2011). Practicing versus inventing with contrasting cases: The effects of telling first on learning and transfer. *Journal of Educational Psychology, 103(4)*, 759-775.

Schwartz, D. L., & Martin, T. (2004). Inventing to prepare for learning: The hidden efficiency of original student production in statistics instruction. *Cognition & Instruction, 22*, 129-184.

Walton, G. M., & Cohen, G. L. (2011). A brief social-belonging intervention improves academic and health outcomes of minority students. *Science*, *331*(6023), 1447-1451.

William, L. B., & Harter, N. (1899). Studies on the telegraphic language: The acquisition of a hierarchy of habits. *Psychological review*, *6*(4), 345.

Wineburg, S. (1998). Reading Abraham Lincoln: An expert/expert study in the interpretation of historical texts. *Cognitive Science, 22*, 319-346.

Yeager, D. S., Purdie-Vaughns, V., Garcia, J., Apfel, N., Brzustoski, P., Master, A., ... & Cohen, G. L. (2014). Breaking the cycle of mistrust: Wise interventions to provide critical feedback across the racial divide. *Journal of Experimental Psychology: General*, *143*(2), 804.