

When are Real-Time Feedback Loops Most Valuable? New Insights from Bandit Simulations of Decision Making in Turbulent Environments

Joseph W. Clark
Arizona State University
joeclark.phd@gmail.com

Abstract

Businesses are increasingly encouraged to invest in technologies to support real-time feedback loops between decision making and analysis, but research has little examined when real-time feedback is most and least valuable. In the present study we extend the canonical model of decision making under uncertainty to examine how the value of real-time feedback varies with the turbulence of the decision environment. Additionally, we explored the complementarity of real-time feedback with a knowledge renewal capability that serves to unlearn out-of-date knowledge. Simulation results reveal an unexpected nonlinearity: a real-time feedback loop has a markedly greater positive impact on performance at moderate levels of turbulence than at high or low levels of turbulence. The impacts of knowledge renewal capacity, as we simulated it, are less clear-cut, and the two capabilities were not found to interact. Implications for research and practice, and recommendations for further study are discussed in the paper's final section.

1. Introduction

In fast-paced environments, businesses are urged to speed up the feedback loops between action and knowledge by using, for example, real-time business intelligence and it is often assumed simply that “faster BI is better BI” [1]. Yet, real-time information systems are difficult to implement, both technologically and organizationally [2]. Legacy systems and batch processes must be upgraded or replaced. Managers and employees need to learn how to deal with data that arrives more rapidly and at a different grain, and may suffer growing pains [3]. Moreover, managerial attention is finite, and there must be a limit to how many real time data streams any individual can monitor and effectively process at once. An important question for research is: when (i.e., for what types of problems and decisions) are real-time feedback loops enabled by information

systems such as BI, data streams, and analytics most valuable? And conversely: when can they be omitted? The answers to these questions would be valuable in helping practitioners focus their investments; in this we hope to answer the call to relate the affordances of BI technology to the types of decision environments in which they are most effectively used [4].

A related capability that has received less attention is that of “organizational forgetting”, that is, the managed unlearning of outdated knowledge in order to make way for the new [5]. Real-time business intelligence can be classified as part of a knowledge-based dynamic capability [6]. To be a dynamic capability is to be part of “the firm’s ability to integrate, build, and reconfigure internal and external competencies to address rapidly changing environments” [7,8,9], but new knowledge cannot be turned into competitive advantage unless an organization is willing to let go of its investments in old knowledge [10]. Abandoning hard-won knowledge is, of course, not to be taken lightly, and there is a need for new research to confirm how this knowledge renewal process works. Our second question, then, is how and when does this knowledge renewal capability complement the value of real-time feedback?

Simulation research in organization theory has an excellent tradition of allowing us to model different aspects of the business environment that may distinguish different types of problems, aspects such as turbulence, munificence [11], velocity, ambiguity, complexity, and unpredictability [12]. These environmental dynamics are difficult to measure and control in the outside world, but can be controlled precisely in the virtual world of a stochastic processes or agent-based simulation [13, 14].

In this study, we extended a canonical simulation model—the multi-armed bandit—to examine how the speed (or slowness) of a feedback loop affects performance, and furthermore to show how that performance effect differs as environmental turbulence changes. We also added to the simulation a model of knowledge renewal capability, that is, the

dynamic capability of discarding outdated knowledge about the environment, to examine this capability's effect on performance and its interaction with the rapidity of the feedback loop.

Our theory-derived hypotheses, simulation model, findings, and implications for theory, research, and practice are presented below.

2. Theory Development

Many researchers have argued that real-time business intelligence and analytics are (part of) an important dynamic capability allowing firms to adapt to rapidly-changing business environments [15]. Brown & Eisenhardt [16] found that businesses on “the edge of chaos” adapted by making frequent, small probes of the future, necessitating the ability to quickly sense and respond [17]. From a pragmatic standpoint it is often argued that speed of the feedback or OODA loop is essential to getting ahead of the competition [18,19]. We infer that the faster the organization's or decision maker's feedback loop, the better able it or he is to achieve successful decision outcomes; and moreover, that this effect should only increase as environmental turbulence increases.

Hypothesis 1: As latency in a decision maker's feedback loop decreases (i.e. as it approaches real-time feedback), decision making performance in turbulent environments increases.

Hypothesis 2: As turbulence increases, the positive impact of real-time feedback on decision making performance will increase.

Additionally, it is necessary that decision makers un-learn knowledge that has become outdated due to environmental change [5]. Research shows that there are a number of pathologies stemming from the path-dependent nature of learning, such as ‘success traps’, ‘failure traps’ [20], and learned risk-aversion [21, 22]. This is a necessary but rarely explicit part of the concept of a dynamic capability—discarding old knowledge complements generation of new knowledge [10]. This process has not been clearly identified in typologies of knowledge-based dynamic capabilities [6] so we refer to it as *knowledge renewal capability*, for want of a better term. Of course, knowledge must not be thrown out too soon. The amount of knowledge pruning that needs to occur is likely dependent on the rate of change.

Hypothesis 3: If a decision maker discards old knowledge, decision making performance in turbulent environments increases. A nonlinear relationship is hypothesized: some optimal rate of “knowledge renewal” will have a greater positive impact on performance than a too-high or too-low rate.

We expect there will be an interaction between real-time feedback and active knowledge renewal, particularly at very high levels of turbulence. We see in practice that many businesses are leveraging data streams instead of data warehouses in fast-paced environments: analyzing data from a recent time window instead of historical databases, and doing so rapidly and continually [23]. We expect that the decision maker in a very turbulent environment who processes data quickly and is not biased by outdated knowledge will perform as well as possible on “the edge of chaos”.

Hypothesis 4: The positive impacts of low-latency feedback loops and active knowledge renewal will interact, producing a performance impact greater than the sum of the two effects.

3. Simulating the Multi-Armed Bandit

The “multi-armed bandit” is a well-established model for the study of sequential decision making behavior under uncertainty [11,21]. In the model, a decision maker (or organization) is presented with a number of options to choose from (we may visualize a slot machine with many arms), each with a different, unknown, probability distribution over rewards. By picking options one at a time, the organization accumulates rewards or losses while forming beliefs from experience about the expected values of each option.

In the bandit model, decision making cannot be separated from sampling. Each choice leads to profit or loss but also helps the decision maker learn only about the option that was chosen. The recurring decision therefore is characterized by a meta-decision which we may call the strategic aspect of the model: should the decision-maker take the option currently believed to have the highest expected payoff, or should he choose a different option in order to gain more information about its (potentially greater, but yet unknown) rewards?

Some of the key findings are that even in a stable environment, the organization achieves the best performance with a strategy that mixes regular *exploitation* of the best-known option with some

amount of *exploration* of other options [24]; that in a changing environment, which can be represented by intermittently altering the probability distributions over rewards, the optimal balance of the two activities (i.e., the optimal strategy) changes [11]; and that the path-dependent nature of learning in this model can generate risk-aversion endogenously [21,22].

We selected the multi-armed bandit model for this research because its long tradition in peer-reviewed research establishes its legitimacy, and allows us to compare our findings meaningfully with prior research. Bandit simulations can reproduce some of the same findings as evolutionary models of competitive selection and adaptation, another popular form [22]. Bandit models also have some real-world relevance—they are being used in software such as Google Analytics for A/B testing and other online experiments [25, 26].

In order to further ground our research in validated literature, we began by replicating a recent example of the multi-armed bandit by Posen and Levinthal [11], a particularly well-explained implementation which includes a model of environmental turbulence, before extending it for our purposes. Replicating the results of their main experiments (see Figures 1 and 2, below) validated that our code was faithful to theirs and free of software bugs. In addition, it helps assure that our findings resulted from intentional changes we made to the model’s assumptions, rather than quirks of our implementation. Finally, we believe that replicating other researchers’ results is a small but meaningful contribution to science.

Our program was coded in Python and all simulations were run in Python 3.4.2. It is designed to be replicated and extended by other researchers, and released as an open-source project. The code is available at <https://github.com/joeclark-phd/bandito> and is open to forks or new contributions through the GitHub platform. All results reported below were generated with version 2.0 of the “bandito” project (commit #22aad83858).

3.1. Program specifics

At the heart of the model are a set of N alternatives from which an organization or decision maker must choose, each with its own probability distribution over rewards. In each time-step of the process the organization makes a choice, receives a reward drawn from the probability distribution of the chosen option, and updates its beliefs about the relative or absolute favorability of the alternatives. Replicating the Posen and Levinthal implementation

as exactly as possible, we used the following parameters:

- $N = 10$
- Probability distributions over rewards are Bernoulli. Rewards are +1 or -1, with the probability of a favorable reward for each option $P = [p_1, \dots, p_N]$ determined at the initialization of the program by independent draws from a Beta distribution with $\alpha=2$ and $\beta=2$. This is a symmetrical Normal-like distribution of values with mean 0.5 and standard deviation 0.22, but is strictly bounded between 0 and 1, making it a good choice for generating random probabilities.
- Rewards are cumulative. In each turn the organization earns a reward of +1 or -1. The organization begins with a “stock” of zero and the simulation runs for 500 turns.
- At each experimental condition we replicate the simulation 1000 times, computing the mean and variance of the stock in turn 500 as the main outcome variable (“performance”).¹
- The decision maker’s beliefs $Q = [q_1, \dots, q_N]$ about the N options are estimates of p_1, \dots, p_N derived from the feedback the organization has received on its choices. Initial beliefs are set to 0.5 across the board and once the decision maker has some experience with an option, its belief is an average of a 1 for each positive reward and a 0 for each negative reward experienced with that option.²
- Turbulence is simulated as probabilistic shocks to the payoffs probabilities P , with the probability of a shock in any given turn determined by the parameter η . In the event of a shock, some of the payoff probabilities p_1, \dots, p_N are re-drawn from the initial beta

¹ [11] used 25,000 replications.

² Two additional trials with one positive outcome are averaged in with this history. This provides the initial belief and prevents the belief from jumping all the way to 0.0 or 1.0 after the first experience. Instead the first real feedback is “averaged in” with the prior belief.

distribution, with an independent probability of 0.5 for each to be re-drawn.

The decision maker’s sequential choices among the N alternatives are determined by its strategy. A number of strategies can be, and have been, articulated for bandit models. A myopic or “greedy” strategy chooses, in every time period, the option believed to have the greatest probability of reward. A more common model from the exploitation-exploration literature is an “ ϵ -greedy” rule in which the believed-best option is chosen most of the time, but with probability ϵ the organization chooses one of the other options at random.

Posen and Levinthal [11] adopt a more sophisticated implementation of strategy characterized by the “softmax” algorithm in which the probability of choosing an option i is m_i , a function of its relative rank among the organization’s beliefs $Q = [q_i, \dots, q_N]$. Specifically, m_i is given by:

$$m_i = \frac{e^{(q_i/(\tau/10))}}{\sum_{i=1}^N e^{(q_i/(\tau/10))}}$$

In this model, the parameter τ characterizes the organization’s strategy. At very small levels of τ , differences in beliefs are weighted heavily and the organization will almost always choose the believed-best option. Even when “exploring” it will likely choose the believed-second-best option and is very unlikely to visit the believed-worst option. At higher levels of τ , the organization is increasingly indifferent to the differences in beliefs and more likely to try the alternatives that it does not have positive experiences with. The benefit of the softmax algorithm is its sensitivity to the strength of the organization’s opinions about the differences between options.

3.2. Measurement of Attainable Performance

Research using bandit simulations has shown that the optimal strategy (ϵ or τ) differs in different environmental conditions. In order to find the best-attainable performance for a given experimental condition, we followed [11]’s procedure: first we simulated the model at five levels of τ : 0³, 0.25, 0.5, 0.75, and 1.0, and from the results we determined the mean performance at each level of τ .

³ Instead of a true zero we used 0.02 to avoid a division-by-zero problem. This is true to the replication of [11].

Figure 1 replicates their first experiment in a stable environment, revealing the nonlinear relationship of strategy to performance.

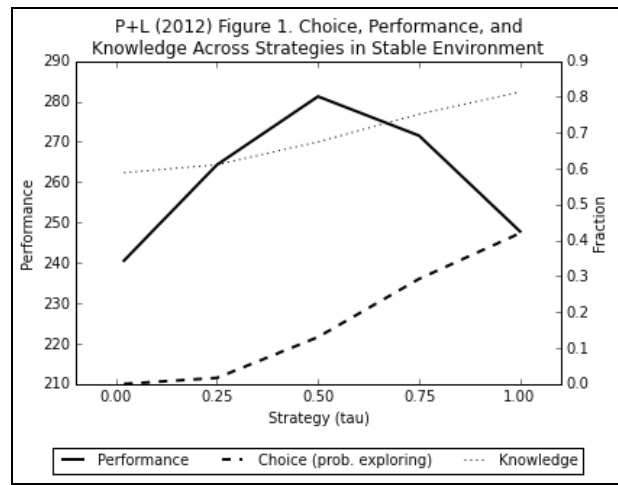


Figure 1. [11], Fig 1, replicated

With this type of simulation output for each experimental condition, they fit a third-order polynomial to the data and calculated the optimal strategy by calculus. This method saved them the necessity of simulating dozens or hundreds of levels of τ , which would have been computationally prohibitive. This method was used to generate their second main experiment’s findings, which we replicated as Figure 2.

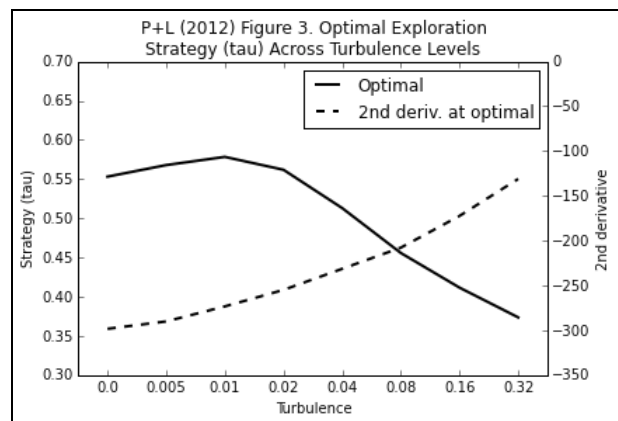


Figure 2. [11], Fig 3, replicated

Our experiments used the same method to calculate the optimal strategy, and the expected level of performance attainable with that strategy, from the least squares fit polynomial function. This measure is called “attainable performance” in our figures below.

3.2. Validation

We argue that our modeling approach is validated in three ways: first, the bandit model is considered a canonical model of decision-making and organizational learning under uncertainty, and used in a number of fields; second, bandit models have of late attained real-world relevance in applications such as A/B testing; and third, faithful replication of other researchers' results validates the quality of our code.

4. Extending the Model

One feature of most bandit simulations that stands out, from the perspective of an information systems scholar, is that the decision-maker is assumed to receive feedback and learn from the outcomes of his decisions instantaneously and without error. To relax this assumption, we introduced a latency parameter, L . The decision maker receives and learns from a decision outcome L turns later (after having made L choices in the meantime). We experimented with values of L from 0 (i.e., real-time) to 16 turns.

Second, we modeled knowledge renewal as a simple time window of memory: the decision maker renews his beliefs about the payoffs of his options each turn based on data received in the past M turns. If a particular arm i had not been chosen during that time window, the belief q_i would revert to 0.5. We experimented with values of M from 20 to 320, as well as the control condition $M=500$ in which no data was ever forgotten by the decision maker.

5. Experiments with Latency

Hypothesis 1 states that in turbulent environments, attainable performance decreases as latency in the feedback loop (L) increases. We experimented with turbulence (η) of 0, 0.005, 0.01, 0.02, 0.04, 0.08, 0.16, and 0.32. Figure 3 illustrates the observed effect of latency on performance at a moderate level of turbulence ($\eta=0.04$).

At least at this level of turbulence, the performance cost of latency is dramatic and unambiguous. With a real-time feedback loop ($L=0$), the decision maker achieved markedly greater performance than with a delayed feedback loop ($L=16$), and the function is monotonic. Every decrease in latency corresponded to an increase in attainable performance.

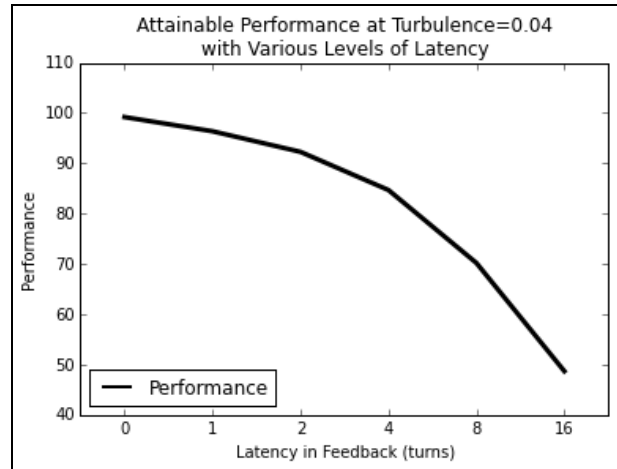


Figure 3. Attainable performance decreases as latency in feedback loop increases

Confirming Hypothesis 1, this pattern is seen quite clearly at every level of turbulence simulated, except the stable case ($\eta=0$). The impacts of latency on performance can be seen in Figure 4, with the results at each level of turbulence made relative by subtracting them from their performance at $L=0$.




Figure 4. Impacts of latency on attainable performance at various levels of turbulence

Interestingly, the difference is greatest at moderate levels of turbulence: 0.04 and 0.08. This contradicts what we expected, as in Hypothesis 2 we predicted that the value of real-time feedback would scale monotonically with turbulence. To view this another way, we subtract attainable performance at $L=16$ from attainable performance at $L=0$ to create a measure we may call “cost of latency” or “value of real-time”. As Figure 5 illustrates, it is plainly nonlinear.




Figure 5. Nonlinear effect of turbulence on the value of real-time feedback

To try to explain the nonlinearity, we observe the absolute performance at each level of turbulence—i.e., the raw data that went into Figure 4. What we see in Figure 6 is that, at the higher levels of turbulence ($\eta=0.16$, $\eta=0.32$), performance tends to decline to an average of zero, the theoretical long-run average of completely random guesses.

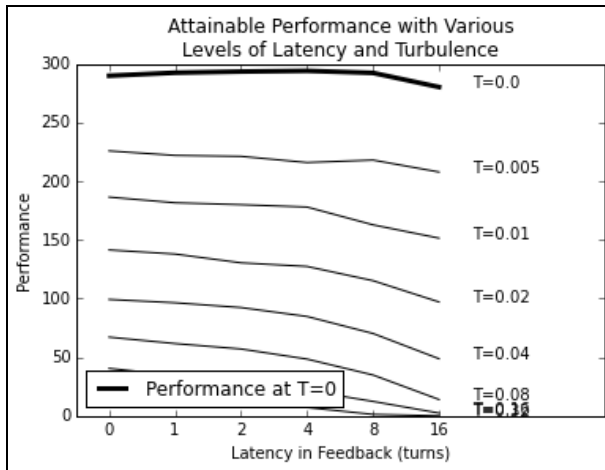


Figure 6. Attainable performance at various levels of latency and turbulence

We may conjecture that in extremely turbulent environments, knowledge becomes outdated so quickly that the decision maker simply cannot benefit from his learning no matter how quickly he processes the data. Therefore, we find that real-time feedback loops are most impactful at a moderate level of turbulence—where the environment is stable enough that knowledge can be exploited for profit, but changes quickly enough that there is a real advantage to rapid learning from feedback.

5. Experiments with Knowledge Renewal

Hypothesis 3 supposes that at some optimal level of knowledge renewal (that is, the discarding of outdated knowledge), there will be a positive impact on attainable performance. To seek this pattern, we experimented with several levels of the memory parameter M . Recall that M determines the time window of memory, with lower values meaning that old knowledge is discarded sooner. $M=500$ is the control case in which no knowledge is discarded.

In Figure 7, we see that the inverted-U shaped pattern we expected is not clearly seen. The effect of M on performance is ambiguous, if it exists at all, and seems to be different at different levels of turbulence. Note that in no case does a change in M affect performance by more than 8 points; compare that to Figure 4 in which we saw that latency moved performance by up to 50 points in some cases.

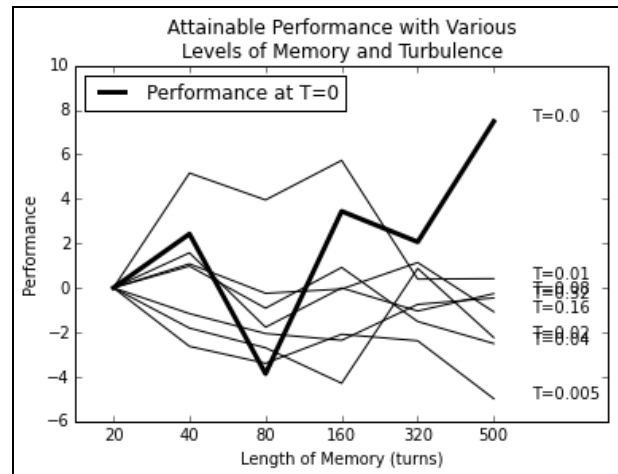


Figure 7. Effects, maybe, of knowledge renewal on attainable performance

Viewed another way, a graph of the best-performing M for each level of turbulence tells us little except that we do not yet understand the effect of knowledge renewal, if it exists (Figure 8).

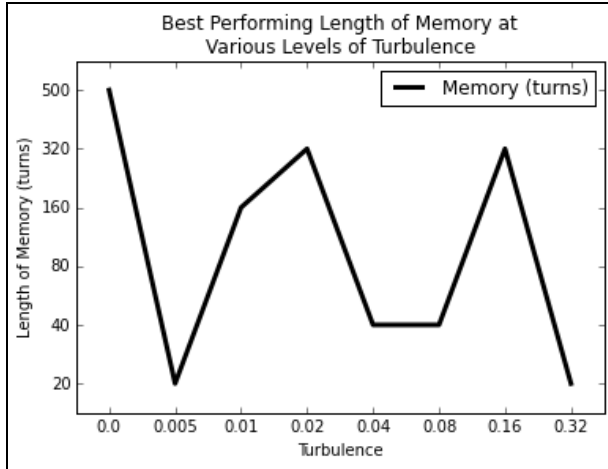


Figure 8. Optimal M for each turbulence level

In Hypothesis 4, we predicted a complementarity between a low-latency feedback loop and an optimal level of active knowledge renewal. Given the ambiguity of the previous experiment, we now simply wish to know if there is any interaction effect at all between the two variables. To generate Figure 9, we replicated Figure 5 at each level of M .

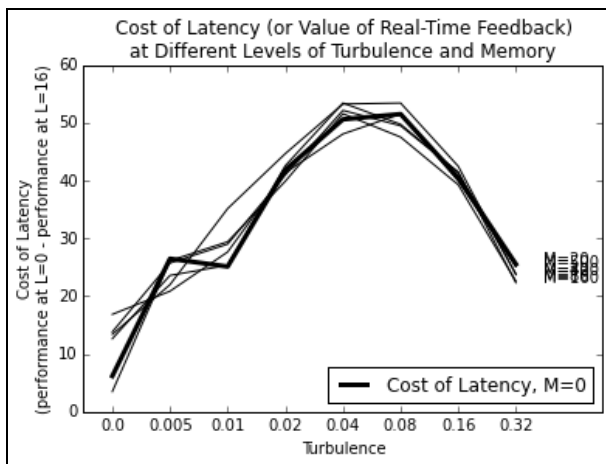


Figure 9. Knowledge renewal has little effect on the cost of latency (value of real-time)

Clearly no meaningful effect of knowledge renewal on the cost of latency (value of real-time) is seen. Hypotheses 3 and 4, therefore, are not supported by these findings.

5.1. Limitations

That we observed no clear effect of knowledge renewal (via the discarding of old knowledge) is difficult to explain, and may be due to limitations of our simulation model. For example, we may need to sample different levels of M and perhaps run the

simulation for a longer number of turns in order to see long-run effects of this process.

Alternatively, it may be the case that knowledge renewal in practice cannot be achieved by simply putting an expiration date on every chunk of experience. We may need to further develop a model of how old knowledge can be managed and renewed more intelligently.

In future research, our simulation may be extended by developing diagnostics to quantify the staleness of the decision maker's knowledge, and the harm it causes him, and compare these to the harmful side effects of any knowledge renewal approach. This would help us to test better models and to identify better levels of parameters to experiment with.

6. Discussion

Motivated by the intention to better understand the value of real-time feedback loops, and how it may differ in different decision making environments, we extended the canonical multi-armed bandit model and carried out a series of simulation experiments.

We confirmed that decision making performance in turbulent environments increases as latency in the feedback approaches zero, but our most important new finding is that the value of real-time feedback (which may also be called the cost of latency in feedback) varies *nonlinearly* with turbulence. At moderate levels of turbulence, real-time feedback confers a markedly greater advantage to a decision maker than in mostly-stable environments or highly turbulent environments. We conjecture that in stable environments, there is little penalty to using "old" knowledge, and at the other extreme, change is so rapid that knowledge becomes stale before it can be exploited [11,27].

The implications for practice are that real-time feedback loops, supported by business intelligence, data streams, and other analytics technologies, are not equally valuable in all decision environments [4]. Because these technologies are not implemented without significant challenges and costs in money, time, risk, and attention, businesses may use this research as a guide to determine which types of problems and solutions are the best candidates for such investment. Our findings indicate that it is not as simple as saying that the more turbulent the problem space, the more value one will get from real-time data. When operating in extreme turbulence, "far from equilibrium", some other approach may need to be taken [28, 29]. Organizations facing these types of problems may need to invest in

improvisational capabilities [30, 31] rather than dynamic capabilities.

Further research may build on our findings by considering more types of environmental dynamism. We replicated the model of turbulence as random shocks of varying frequency from [11] but did not vary other environmental dynamics such as complexity, ambiguity, or (un)predictability which has been simulated by [12]. Extensions to our model to examine these dimensions may allow us to better understand the value of real-time feedback.

Regarding knowledge renewal capability, our findings did not confirm our expectations, and we are left with a mystery for further research. The questions we have regarding our simulation are questions that businesses themselves need answers to: how simple or complicated does knowledge renewal need to be? How can they measure the cost of retaining stale data? And how can this be compared to the risk of throwing away still-relevant data? Further modeling and experimentation of this capability will be a meaningful avenue for future research.

10. References

- [1] D. Stodder, "Chasing the Real-Time BI Dream", TDWI, Oct. 21, 2014, <http://www4.tdwi.org/TDWI/TDWI-Staging/Articles/2014/10/21/Chasing-Real-Time-BI-Dream.aspx>.
- [2] Harvard Business Review Analytic Services, "Making the Real-Time Enterprise a Reality" [TDWI white paper], 2015, <http://tdwi.org/whitepapers/2015/03/making-the-real-time-enterprise-a-reality>.
- [3] J.W. Clark, O.A. El Sawy, "Business Intelligence at Guthy-Renker: The Promise and Challenges of Sensing the Pulse", SIGDSS/TUN BI Congress II, St. Louis, 2010.
- [4] O. Isik, M.C. Jones, and A. Sidorova, "Business intelligence success: The roles of BI capabilities and decision environments", *Information & Management* 50, 2013, pp. 13-23.
- [5] P.M. de Holan, N. Phillips, and T.B. Lawrence, "Managing Organizational Forgetting", *MIT Sloan Management Review* 45(2), 2004.
- [6] J.S. Denford, "Building knowledge: developing a knowledge-based dynamic capabilities typology", *Journal of Knowledge Management* 17(2), 2013, pp. 175-194.
- [7] D.J. Teece, G. Pisano, and A. Shuen, "Dynamic capabilities and strategic management", *Strategic Management Journal* 18(7), 1997, pp. 509-533.
- [8] K.M. Eisenhardt and J.A. Martin, "Dynamic capabilities: what are they?", *Strategic Management Journal* 21(10/11), 2000, pp. 1105-1121.
- [9] M. Zollo and S.G. Winter, "Deliberate learning and the evolution of dynamic capabilities", *Organization Science* 13(3), 2002, pp. 339-351.
- [10] A. Hermann, O. Gassmann, and U. Eisert, "An empirical study of the antecedents for radical product innovation and capabilities for transformation", *Journal of Engineering and Technology Management* 24, 2007, pp. 92-120.
- [11] H.E. Posen and D.A. Levinthal, "Chasing a Moving Target: Exploitation and Exploration in Dynamic Environments", *Management Science* 58(3), 2012, pp. 587-601.
- [12] J.P. Davis, K.M. Eisenhardt, and C.B. Bingham, "Optimal Structure, Market Dynamism, and the Strategy of Simple Rules", *Administrative Science Quarterly* 54(3), 2009, pp. 413-452.
- [13] J.P. Davis, K.M. Eisenhardt, and C.B. Bingham, "Developing Theory Through Simulation Methods", *The Academy of Management Review* 32(2), 2007, pp. 480-499.
- [14] K.M. Carley, "Computational Approaches to Sociological Theorizing", in J.H. Turner (ed.), *Handbook of Sociological Theory* (pp. 69-83), Kluwer Academic/Plenum Publishers, New York, 2002.
- [15] V. Sambamurthy, A.S. Bharadwaj, and V. Grover, "Shaping agility through digital options: Reconceptualizing the role of IT in contemporary organizations", *MIS Quarterly* 27(2), pp. 237-263.
- [16] S.L. Brown and K.M. Eisenhardt, "The art of continuous change: Linking complexity theory and time-paced evolution in relentlessly shifting organizations", *Administrative Science Quarterly* 42, 1997, pp. 276-304.
- [17] Y. Park and O.A. El Sawy, "Discovering the Multifaceted Roles of Information Technology with a Holistic Configurational Theory Approach", *Proceedings of HICSS*, 2012.
- [18] M. Rivera, "Real-Time Business Intelligence in the Real World", TDWI, February 3, 2015, <http://tdwi.org/articles/2015/02/03/real-time-bi-real-world.aspx>.
- [19] R. Houghton, O.A. El Sawy, P. Gray, C. Donegan, and A. Joshi, "Vigilant Information Systems for Managing Enterprises in Dynamic Supply Chains: Real-Time Dashboards at Western Digital", *MIS Quarterly Executive* 3(1), 2004, pp. 19-35.

- [20] D.A. Levinthal and J.G. March, "The Myopia of Learning", *Strategic Management Journal* 14, 1993, pp. 95-112.
- [21] J.G. March, "Learning to Be Risk Averse", *Psychological Review* 103(2), 1996, pp. 309-319.
- [22] J. Denrell and J. G. March, "Adaptation as Information Restriction: The Hot Stove Effect", *Organization Science* 12(5), 2001, pp. 523-538.
- [23] L. Zliobaite, A. Bifet, M. Gaber, B. Gabrys, J. Gama, L. Minku and K. Musial, "Next challenges for adaptive learning systems", *ACCM SIGKDD Explorations Newsletter* 14(1), 2012, pp. 48-55.
- [24] J.G. March, "Exploration and exploitation in organizational learning", *Organization Science* 2(1), 1991, pp. 71-87.
- [25] J.M. Whyte, *Bandit Algorithms for Website Optimization*, O'Reilly Media, Sebastopol, CA, 2013.
- [26] S. Hanov, "20 lines of code that will beat A/B testing every time", *Steve Hanov's Blog*, 2012, <http://stevehanov.ca/blog/index.php?id=132>.
- [27] J.W. Selsky, J. Goes, and O.N. Baburoglu, "Contrasting Perspectives of Strategy Making: Applications in 'Hyper' Environments", *Organization Studies* 28(1), 2007, pp. 71-94.
- [28] F.E. Emery and E.L. Trist, "The Causal Texture of Organizational Environments", *Human Relations* 18(1), 1965, pp. 21-32.
- [29] A.D. Meyer, V. Gaba, and K.A. Colwell, "Organizing Far from Equilibrium: Nonlinear Change in Organizational Fields", *Organization Science* 16(5), pp. 456-473.
- [30] P.A. Pavlou, O.A. El Sawy, "The 'Third Hand': IT-Enabled Competitive Advantage in Turbulence Through Improvisational Capabilities", *Information Systems Research* 21(3), 2010, pp. 443-471.
- [31] S.A. Carlsson, O.A. El Sawy, "Managing the five tensions of IT-enabled decision support in turbulent and high-velocity environments", *Information Systems and eBusiness Management* 6.3, 2008, pp. 225-237.