Abrupt Strategy Change Underlies Gradual Performance Change: Bayesian Hierarchical Models of Component and Aggregate Strategy Use

Sarah K. A. Wynton, Jeromy Anglim¹

Abstract. While researchers have often sought to understand the learning curve in terms of multiple component processes, few studies have measured and mathematically modeled these processes on a complex task. In particular, there remains a need to reconcile how abrupt changes in strategy use can co-occur with gradual changes in task completion time. Thus, the current study aimed to assess the degree to which strategy change was abrupt or gradual, and whether strategy aggregation could partially explain gradual performance change. It also aimed to show how Bayesian methods could be used to model the effect of practice on strategy use. To achieve these aims, 162 participants completed 15 blocks of practice on a complex computer-based task—the Wynton Anglim Booking (WAB) Task. The task allowed for multiple component strategies (i.e., memory retrieval, information reduction, and insight) that could also be aggregated to a global measure of strategy use. Bayesian hierarchical models were used to compare abrupt and gradual functions of component and aggregate strategy use. Task completion time was well-modeled by a power function, and global strategy use explained substantial variance in performance. Change in component strategy use tended to be abrupt, whereas change in global strategy use was gradual and well-modeled by a power function. Thus, differential timing of component strategy shifts leads to gradual changes in overall strategy efficiency, and this provides one reason for why smooth learning curves can co-occur with abrupt changes in strategy use.

Keywords: strategy use, skill acquisition, Bayesian hierarchical modeling, learning decomposition, memory retrieval

---

¹ **Citation Information:** Wynton, S. K. A., & Anglim, J. (2017). Abrupt Strategy Shifts Underlie Gradual Performance Changes: Hierarchical Bayesian Models of Decomposed Strategy Use. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 43, 1630-1642.. <http://dx.doi.org/10.1037/xlm0000404>

© 2017, American Psychological Association. This paper is not the copy of record and may not exactly replicate the final, authoritative version of the article. If you cite the paper, please follow the doi to check for updated issue and page information. The final article will be available, upon publication, via its doi: <http://dx.doi.org/10.1037/xlm0000404>

Author Note: Sarah K. A. Wynton and Jeromy Anglim, School of Psychology, Deakin University, Australia. We thank Simon A. Wynton for his assistance with programming the WAB Task. Data, scripts, and the task are available at <https://osf.io/5krfq>. Correspondence concerning this article should be addressed to Sarah K. A. Wynton, School of Psychology, Deakin University, Locked Bag 20000, Geelong, 3220, Victoria, Australia. Email: skaw@deakin.edu.au

The time it takes to complete a cognitive or psychomotor task generally decreases gradually with practice (e.g., Newell & Rosenbloom, 1981; Thorndike, 1908). One way that individuals improve with practice is by adopting more efficient strategies (e.g., Crossman, 1959; F. J. Lee & Anderson, 2001; Logan, 1988; Tenison & Anderson, 2016). Some research suggests that gradual changes in performance are an artifact of group-level analysis and that individual-level learning curves may show discontinuities when abrupt changes in strategy use occur (Gaschler, Marewski, & Frensch, 2015; Haider & Frensch, 2002; Rickard, 1997, 2004). However, even at the individual-level, learning curves rarely display abrupt changes (Anglim & Wynton, 2015; Heathcote, Brown, & Mewhort, 2000). Thus, there remains a need to reconcile how abrupt changes in strategy use can co-occur with gradual changes in task completion time. One explanation is that tasks involve multiple strategy components. Some component strategies may change abruptly, while others change gradually; the timing of abrupt strategy changes may also vary between components. Thus, the effect of practice on an overall strategy efficiency may be more gradual and begin to mirror the functional forms typically seen for task completion time. The present study investigates this theory using novel methods for comparing abrupt and gradual strategy change on a complex task.

Practice, Strategy Use, and Performance

Practice leads to improvement in strategy use and performance. A general finding in skill acquisition research is that task completion time decreases with practice and that the rate of change decelerates monotonically towards an asymptote (Heathcote et al., 2000; Newell & Rosenbloom, 1981). The three-parameter power function provides a good fit to data and has been widely adopted (Ghisletta, Kennedy, Rodrigue, Lindenberger, & Raz, 2010; F. J. Lee & Anderson, 2001; Logan, 1988; Newell & Rosenbloom, 1981). This Power Law of Practice (Newell & Rosenbloom, 1981) states that

$$E(y) = \theta^{(1)} x^{-\theta^{(2)}} + \theta^{(3)},$$

where expected task completion time, y , is a function of practice, $x = 1, 2, \dots$, and where the three parameters represent amount of learning $\theta^{(1)}$, rate of learning $\theta^{(2)}$, and asymptotic performance $\theta^{(3)}$. While researchers have debated whether the learning curve is better fit by an exponential or a power function (Anglim & Wynton, 2015; Heathcote et al., 2000; Myung, Kim, & Pitt, 2000), both functions imply that change in task completion time with practice involves monotonic reduction, monotonic deceleration, and an approach to an asymptote.

Strategy use influences performance and strategy change partially explains the effect of practice on performance (e.g., Anglim & Wynton, 2015; F. J. Lee & Anderson, 2001; Logan, 1988). While, strategy use can be defined as an optional method of task completion, for the present study we focus on three types of strategies linked to faster task completion times: memory retrieval, information reduction, and insight. Memory retrieval involves retrieving a solution from memory instead of applying a multi-step algorithm (e.g., Bajic & Rickard, 2009, 2011; Compton & Logan, 1991; Crowley, Shrager, & Siegler, 1997; Logan, 1988; Logan & Etherton, 1994; Rickard, 1997; Siegler, 1988b; Tenison & Anderson, 2016). Information reduction involves ignoring task-irrelevant information (see, for example Haider & Frensch, 1996, 1999; F. J. Lee & Anderson, 2001). Finally, insight strategies involve discovery where the key to using the strategy is knowing that it exists, and then choosing to use it (e.g., Anglim & Wynton, 2015; Blessing & Anderson, 1996; Yechiam, Erev, & Parush, 2004).

Research suggests that practice can lead to both abrupt and gradual changes in strategy use at the individual-level. Abrupt changes seem to be more prevalent when strategy change

involves an insight or simple discrete change. For example, the transition to memory retrieval may occur abruptly on a single item but not on set of items (Rickard, 2004), and learning to ignore a single irrelevant piece of information (Haider & Frensch, 2002) can occur more abruptly than a gradual refinement of how visual attention is allocated on a complex computer-based task (F. J. Lee & Anderson, 2001). Formal instruction may also increase the prevalence of abrupt strategy shifts (Alibali, 1999).

Some researchers have explicitly sought to reconcile abrupt strategy shifts with models of the learning curve. Haider and Frensch (2002) noted that if the timing of strategy shifts differ across individuals, group-level learning curves will appear gradual even when individual-level curves show discontinuities. While some research has obtained abrupt changes in performance when strategy shifts occur (Gaschler et al., 2015; Haider, Frensch, & Joram, 2005; Rickard, 2004), most studies show that even for tasks that involve abrupt changes in strategy use, individual-level learning curves remain gradual (e.g., Anglim & Wynton, 2015; Heathcote et al., 2000). Second, several researchers have suggested that learning curves are strategy-specific (Delaney, Reder, Staszewski, & Ritter, 1998; Palmeri, 1999; Siegler, 1987), meaning that individuals get faster at implementing a particular strategy with practice. This suggests that the performance benefits of a strategy change may take time to accrue. Third, for some tasks, strategy change is an item-specific process (e.g., memory retrieval on multiplication problems). Thus, for tasks that include multiple items, performance may improve gradually due to different item-specific timings for strategy change (e.g., Palmeri, 1997; Rickard, 2004).

A Theory of Strategy Aggregation and Change

We propose that on complex tasks changes in task completion time will be influenced by multiple component strategies. These component strategies will often change abruptly with practice; that is, individuals will abruptly shift to retrieving one item from memory, ignoring a single piece of irrelevant information, or using a particular short-cut that makes task completion more efficient. These component level strategies can then be aggregated to domains and to an overall level of strategy efficiency. In line with research that shows that aggregated curves smooth over discontinuities in underlying processes (Estes, 1956; Haider & Frensch, 2002; Heathcote et al., 2000), we propose that when examined at the global level, strategy change will be gradual at the individual-level. Furthermore, because strategy change is a driving factor of performance change, we propose that when there are a reasonable number of component strategies, changes in global strategy efficiency should broadly mirror changes in performance; i.e., they should follow a learning curve similar to a power function.

Bayesian Hierarchical Models of Strategy Change

While several researchers have fit mathematical models of the effect of practice on strategy use (e.g., Gaschler et al., 2015; F. J. Lee & Anderson, 2001), this literature generally shares limitations with research examining practice and performance. First, many studies have analyzed strategy change at the group-level (e.g., Touron & Hertzog, 2009; Touron, Hertzog, & Frank, 2011). Group-level curves can, and often do, have distinct functional forms to individual-level learning curves (Estes, 1956; Haider & Frensch, 2002; Heathcote et al., 2000), and individual-level curves are of most relevance when developing a psychological theory of skill acquisition. Issues of aggregation are greater for strategy curves because individual differences in functional form are greater than they are for performance; and when strategy use shifts abruptly, the timing of the shift varies across individuals, which can yield smooth power-like functions at the group-level (Haider & Frensch, 2002). Second, researchers have rarely mathematically modeled both gradual and abrupt functions to assess which fits best to the data,

although see Gaschler et al. (2015) for an exception. Third, researchers have rarely used formal statistical models of individual differences in strategy change. Researchers have often relied on descriptive statistics (John & Lallement, 1997; Siegler, 1988a, 1988b, 1991) and piece-wise analyses where models are fit separately for each individual (e.g., Gaschler et al., 2015). This often overfits the data, fails to pool parameter estimates over individuals, and does not model individual differences in strategy change. Finally, studies have rarely used methods that enable an appropriate penalization for sources of model complexity (see Pitt & Myung, 2002 for discussion of sources of model complexity).

Bayesian hierarchical methods are particularly well-suited to overcoming the above limitations related to modeling the effects of practice on strategy use and performance (for further discussion and recent applications, see, Anglim, Weinberg, & Cummins, 2015; Anglim & Wynton, 2015; Averell & Heathcote, 2011; Dennis, Lee, & Kinnell, 2008; Farrell & Ludwig, 2008; Jaeger, Moulding, Anglim, Aardema, & Nedeljkovic, 2015; M. D. Lee, 2008; M. D. Lee, 2011; Wasserman, 2000). Bayesian methods are flexible. It is easy to incorporate non-linear functions such as the power function and non-normal distributions suitable for skewed response time distributions and binary strategy use data. Bayesian methods also offer several approaches to model comparison that penalize different forms of model complexity that are well-suited to comparing abrupt and gradual models of strategy change.

The Current Study

In summary, little research has formally modeled both component and overall strategy change on a complex task at the individual-level; little research has formally tested competing models of gradual and abrupt strategy change; and finally, little research has sought to reconcile abrupt changes in strategy use with gradual changes in performance. Thus, the current study aimed to assess the degree to which strategy change was abrupt or gradual, and assess the degree to which strategy aggregation could explain power function learning. It also aimed to show how Bayesian methods could be used to model the effect of practice on component and aggregate strategy use on a complex task. To achieve these aims, we had participants practice a computer-based timetable booking task. The task permitted objective strategy measurement on each trial. Task analysis identified a range of component-level strategies which could be aggregated up to domains and global levels. Bayesian hierarchical modeling was used to model abrupt and gradual functions of the effect of practice on strategy use at the individual-level. In particular, the models tested our theory of strategy aggregation whereby strategy efficiency may be abrupt at the component level, but approximate a power-function at the global level.

Method

Participants

A sample of 162 participants (64% female) aged 18-51 years ($M = 25.3$, $SD = 8.2$) completed the study. Participants were recruited through an Australian university and received a AU\$25 gift card for their participation. From an original sample of 163 participants, data for one participant was removed due to not completing all practice blocks. Ethics approval for this study was granted by Deakin University Faculty Human Ethics Advisory Group.

The Task

The study used the Wynton-Anglim Booking (WAB) Task (Anglim & Wynton, 2015). The version used in the present study is available at <<https://osf.io/5krfq>>. This computer-based task is designed to be more complex than simple cognitive tasks but simple enough for meaningful learning to take place in a lab setting. The task is a simulated version of that

performed in swimming schools. The participant plays the role of a receptionist whose job it is to book children into swimming lessons that meet the requirements of the parent and child.

Each trial requires the participant to identify an appropriate class for a given child. To achieve this goal, participants will typically need to ask the simulated parent questions about their child, apply filters to reduce the number of classes displayed on the timetable, and select an appropriate class from the timetable. Appropriate classes need to meet three criteria: (a) each level can only contain a certain number of children, so classes cannot be overbooked, (b) the class must adhere to the day, time, and teacher requirements of the parent, and (c) the child must be booked into the right level for their age. There are six levels, Beginner (ages 0-3), Tadpole (ages 4-6), Seals (ages 7-9), Dolphins (ages 10-12), Advanced (ages 13-16) and Squad (17-21). If participants ask the parent what level the child is, 80% of the time the answer is “I don’t know”, so participants typically identify the level by asking for the child’s age and looking up the appropriate level in the class rules. There are 22 different answers to “How old is your child?” that range from 6 months to 21 years.

Figure 1 displays the main screen presented to participants. This screen allows participants to ask questions of the parent, view the class rules (information regarding number of bookings allowed per level and the age requirements for each level), apply filters and view the timetable. When asking questions of the parent or accessing the class rules, a popup message box appears that displays the parents response or class rules. The participant must click “ok” to close this dialog box, before they can perform another action. The timetable consisted of 120 classes (unfiltered) arranged in five columns by day, Monday to Friday, and sorted in rows by time. This also appears as a pop-up window that must be closed (or the correct class selected) before another action can be performed. If filters are applied, the timetable will only display classes consistent with the specified filters.

SHOW TIMETABLE

SHOW ALL CLASS LEVEL INFORMATION

NEXT

Teachers	Levels	Times	Days	Class Size
☒ Racheal	☐ Beginners	☒ 3 pm	☒ Monday	☒ No Restriction
☒ Paul	☐ Tadpoles	☒ 3:30 pm	☒ Tuesday	☐ Less Than 15
☒ Sonia	☒ Seals	☒ 4 pm	☒ Wednesday	☐ Less Than 14
	☐ Dolphins	☒ 4:30 pm	☒ Thursday	☐ Less Than 13
	☐ Advanced	☒ 5 pm	☒ Friday	☐ Less Than 12
	☐ Squad	☒ 5:30 pm		☐ Less Than 11
		☒ 6 pm		☐ Less Than 10
		☒ 6:30 pm		☐ Less Than 9
Select All Teachers	Select All Levels	Select All Times	Select All Days	☐ Less Than 8
UNSelect All Teachers	UNSelect All Levels	UNSelect All Times	UNSelect All Days	☐ Less Than 7
				☐ Less Than 6
				☐ Less Than 5
				☐ Less Than 4

What is Your Child's Name?

Do they have any Previous Swimming Experience?

How Old is Your Child?

What Days or Times suit

What level is your child?

Do You Have a Teacher in Mind?

Do they have any Special Needs?

Figure 1. Task screen as seen by participants. The blue speech bubbles are the question participants can ask of the parent. There are 5 filters which reduce the number of classes displayed on the timetable. The default is to have all options selected (as shown for all but “Levels”) and “No Restriction” on “Class Size”. The orange class level button displays the age requirements and maximum size of each level and the red star “Show Timetable” button displays the timetable in accordance with the filters applied.

Because the WAB Task records all participant actions, strategy use can be measured objectively. Strategy use was represented in terms of global, domain, and component levels. Component level strategy use was calculated as the proportion of trials within each block that a component strategy was used. Domain and global strategy use was calculated as the mean of included component strategies. Task performance was measured as the average time taken to complete a trial within a block, so shorter times represent better performance. To align performance and strategy metrics so that lower strategy values indicated faster performance and more efficient strategy use, all strategy variables represented the proportion of trials in the block that an inefficient strategy was used. If the action was an efficient strategy, then inefficient strategy use was the proportion of trials where the strategy was not used. Thus, inefficient strategy use ranged from 0 (efficient) to 1 (inefficient). Figure 2 presents a hierarchical aggregation of task strategies. For the current study, we aggregated component strategies to domain levels based on strategy type. Inefficient information gathering reflects the amount of irrelevant information being processed, inefficient filtering reflects a lack of insight into the task, and accessing rules reflects the use of a look-up rather than a memory-retrieval strategy. Inefficient information gathering and inefficient filtering each consisted of several component-level strategies.

Figure 2. Hierarchical aggregation of strategy use. All strategies are defined as inefficient strategies.

For inefficient information gathering there were three questions that individuals could ask parents for which the information provided was of no relevance to the task. Thus, we expected that over practice participants would stop asking these questions. Inefficient filtering was broken down into two separate measures of insight. The filters reduce the number of classes displayed on the timetable and therefore the amount of time participants spend scanning the timetable for the correct class. While all the filters reduced the number of classes displayed, the day/time filters had less impact on scanning time, because the timetable is already organized by day and time. Therefore, the first component measure of inefficient filtering was not using the level filter, as the level filter has the largest impact on the number of classes displayed on the timetable. Because the default for the level filter is to have all levels selected, filtering for a specific level requires five mouse clicks to deselect all but one level. However, a participant can reduce this to two clicks by pressing “deselect all” and then the desired level. Therefore, our second measure of inefficient filtering was not using the deselect all button.

Accessing rules was also defined as a domain-level strategy as the age of the student, and thus the item-level memory retrieval required, differs from trial-to-trial. These item-level shifts reflect component-level shifts. However, as the age of the simulated child was random and the number of trials completed in a block differed between individuals, we did not examine these shifts within the models. Instead, these item-level memory retrieval shifts were examined separately (see online supplement).

The WAB Task was created using Visual Basic programming with Microsoft Excel. The task was delivered on desktop computers running Windows 7. The computers had LCD screens with a resolution of at least 1,680 x 1,050 pixels.

Procedure

Participants completed the study in a computer laboratory in small groups and were seated apart to prevent any interactions. Before completing the task, participants completed measures of individual differences that are not the focus of the present study. The skill acquisition task first involved reading instructions about the task. Participants were instructed to

identify a class for each child in the shortest amount of time. Participants then completed 15 3-minute blocks of trials. In each block participants completed as many trials as possible. Blocks ended once trial completion coincided with the passage of at least three minutes since the beginning of the block. At the end of each block, participants were given feedback on their average trial completion time for that block. Trial characteristics such as the age of the child, teacher preferences, and time preferences were randomized according to a schedule (for more information on the task see Anglim & Wynton, 2015).

Bayesian Models

All models were estimated using Bayesian methods. The recent increase in Bayesian methods within psychology has been facilitated by increases in computer processing capabilities, software that simplifies Bayesian analysis, and the publication of introductory textbooks and tutorial papers (e.g., Gelman, Carlin, Stern, & Rubin, 2014; Gelman & Hill, 2007; Kruschke, 2010). In this section we first set out the model of task completion time, followed by the model of strategy use. In broad terms, task completion time was modeled using a three-parameter power function with parameters that varied over individuals. Strategy use was modeled using both an abrupt and a power function. The abrupt function involved a single step (e.g., from always using a strategy to not using a strategy), and the power function was broadly equivalent to the model of task completion time. Various technical details were required to capture the positively skewed nature of task completion time and the use of proportions for strategy use. However, the important point is that the abrupt and power functions provide an empirical test of whether strategy change was abrupt or gradual.

The models state that an individual's expected task completion time or expected strategy use on a particular block comes from a distribution. The mean of the distribution is then defined using a function that describes the effect of practice. Each function includes parameters that reflect specific features of the change (i.e., in the power function $\theta^{(1)}$ reflects the amount of change). All models are hierarchical in that each individual has his or her own function parameters and these parameters are assumed to be drawn from a population distribution. This population distribution can then be described using means and standard deviations (called hyperparameters) to characterize how the practice–performance (or practice–strategy) relationship varies over individuals. In Bayesian analysis, the probability of every potential value for each hyperparameter is calculated, providing a full posterior distribution.

Practice and performance. In order to check that the relationship between practice and performance was similar to that of previous research with this and other tasks, the fit of a power function was examined. In the model below, expected task completion time, y_{ij} , for participant i ($i = 1, \dots, n$) on block j ($j = 1, \dots, 15$) is distributed according to a Gamma distribution with parameters α_{ij} and β_{ij} . The Gamma distribution has characteristics suitable for modeling task completion time, in that it is strictly positive and can readily incorporate a positive skew. For ease of interpretation these parameters were reparameterized into mean, μ_{ij} , and standard deviation, σ_i . The mean is modeled as a power function where x_{ij} represents block number. Each power function parameter ($\theta^{(q)}$, where $q \in (1, 2, 3)$ and indexes the parameter number) has a Gamma distribution that is also reparameterized to be described by mean and standard deviation hyperparameters that describe the distribution across individuals.

$$y_{ij} \sim \text{Gamma}(\alpha_{ij}, \beta_{ij})$$

$$\alpha_{ij} = \frac{\mu_{ij}^2}{\sigma_i^2}$$

$$\beta_{ij} = \frac{\mu_{ij}}{\sigma_i^2}$$

$$\mu_{ij} = \theta_i^{(1)} x_{ij}^{-\theta_i^{(2)}} + \theta_i^{(3)}$$

$$\theta_i^{(q)} \sim \text{Gamma}(\alpha_{\theta^{(q)}}, \beta_{\theta^{(q)}})$$

$$\alpha_{\theta^{(q)}} = \frac{\mu_{\theta^{(q)}}^2}{\sigma_{\theta^{(q)}}^2}$$

$$\beta_{\theta^{(q)}} = \frac{\mu_{\theta^{(q)}}}{\sigma_{\theta^{(q)}}^2}$$

$$\sigma_i \sim \text{Gamma}(\alpha_\sigma, \beta_\sigma)$$

$$\alpha_\sigma = \frac{\mu_\sigma^2}{\sigma_\sigma^2}$$

$$\beta_\sigma = \frac{\mu_\sigma}{\sigma_\sigma^2}$$

The effect of inefficient strategy use on performance was then examined by comparing the fit of the above model to a model where

$$\mu_{ij} = \left(\theta_i^{(1)} x_{ij}^{-\theta_i^{(2)}} + \theta_i^{(3)} \right) + \gamma z_{ij}$$

where z_{ij} represents global inefficient strategy use for individual i on block j and γ is a parameter representing the influence of inefficient strategy use on performance.

Practice and strategy use. To provide a formal test of the degree to which strategy change was abrupt or gradual, abrupt and power models were fit for each strategy. For descriptive purposes, we focus on the observed proportion of trials that a participant, i , used a given strategy, z_{ij} , in a particular block, j . However, because trial-level strategy use was a binary variable, we modeled the number of times a strategy was used z'_{ij} using a binomial distribution. In this way we are modeling the latent probability of strategy use, μ_{ij} , rather than the observed proportion, which includes error. In the model below, w_{ij} represents the number of trials in the block, and k represents the number of strategy components (for example, for inefficient filtering, $k = 2$). Strategy use was modeled as

$$z'_{ij} \sim \text{Binomial}(\mu_{ij}, w_{ij}k)$$

$$z'_{ij} = z_{ij} w_{ij} k$$

$$\mu_{ij} = f(\theta_i, x_{ij})$$

where the abrupt function was

$$\mu_{ij} = \begin{cases} \theta_i^{(a3)} + \theta_i^{(a1)}, & x_{ij} \leq \theta_i^{(a2)} \\ \theta_i^{(a3)} & , x_{ij} > \theta_i^{(a2)} \end{cases}$$

and the power function was

$$\mu_{ij} = \theta_i^{(p1)} x_{ij}^{-\theta_i^{(p2)}} + \theta_i^{(p3)}$$

and where the a and p in the parameter index indicates abrupt and power parameters respectively. To ensure that μ_{ij} remained within the bounds of a proportion (i.e., between 0 and 1), constraints were placed, for both power and abrupt functions, on the population distributions for $\theta_i^{(1)}$ and $\theta_i^{(3)}$ to ensure that $0 < \theta_i^{(1)} < 1$, $0 < \theta_i^{(3)} < 1$, and $0 < (\theta_i^{(1)} + \theta_i^{(3)}) < 1$. The asymptote parameters for each function were drawn from a beta distribution (which only includes values between 0 and 1), i.e., $\theta_i^{(3)} \sim \text{Beta}(\alpha_{\theta^{(3)}}, \beta_{\theta^{(3)}})$ where $\alpha_{\theta^{(3)}}$ and $\beta_{\theta^{(3)}}$ are hyperparameters describing the distribution over individuals. All beta distributions were truncated 0.001 and 0.999 to prevent numerical estimation issues. Parameters representing amount of change for both functions, $\theta_i^{(1)}$, were equal to $\lambda_i(1 - \theta_i^{(3)})$, where $\lambda_i \sim \text{Beta}(\alpha_\lambda, \beta_\lambda)$. Because individuals could not shift strategy before the first block or after the last block, the distribution of the shift point for the abrupt function was distributed on a truncated normal constrained to be between 1.5 and 14.5, i.e., $\theta_i^{(a2)} \sim N_c(\mu_{\theta^{(a2)}}, \sigma_{\theta^{(a2)}}) \in (1.5, 14.5)$. The rate parameter for the power function, $\theta_i^{(p2)}$, was $1.5 \times \bar{\theta}_i^{(p2)}$, where $\bar{\theta}_i^{(p2)}$ was drawn from a beta distribution, i.e., $\bar{\theta}_i^{(p2)} \sim \text{Beta}(\alpha_{\bar{\theta}^{(p2)}}, \beta_{\bar{\theta}^{(p2)}})$, and thus the values of $\theta_i^{(p2)}$ ranged from 0 to 1.5. Constraining the upper bound of $\theta_i^{(p2)}$ limited the ability of the power function to mimic abrupt changes in the first few blocks of practice, thus ensuring it mirrored the learning curve function rather than an abrupt, or rapid, shift.

Estimating Bayesian models requires the specification of priors on the hyperparameters. Relatively uninformative priors were used that respected natural constraints of the problem domain. Hyperparameters for amount of change and asymptotic performance in both abrupt and power functions had gamma priors: $\alpha_\lambda \sim \text{Gamma}(4, 1)$, $\beta_\lambda \sim \text{Gamma}(4, 4)$, $\alpha_{\theta^{(3)}} \sim \text{Gamma}(4, 4)$, and $\beta_{\theta^{(3)}} \sim \text{Gamma}(4, 1)$. These priors are relatively uninformative but also suggest that strategy use in the first block is closer to one, and asymptotic strategy use is close to zero. The priors on hyperparameters for the change point in the abrupt function were uniform: $\mu_{\theta^{(a2)}} \sim \text{Unif}(1, 15)$, $\sigma_{\theta^{(a2)}} \sim \text{Unif}(0, 7)$. Priors for the hyperparameters for the rate parameter in the power function were gamma. $\alpha_{\bar{\theta}^{(p2)}} \sim \text{Gamma}(3, 1)$, $\beta_{\bar{\theta}^{(p2)}} \sim \text{Gamma}(3, 3)$. Additional analyses where priors were varied suggested that results were broadly robust to the exact choice of prior. Model fitting was

conducted using MCMC methods in JAGs with 4 chains of 20,000 iterations, each thinned by 10 with a burn-in of 8000. Model evaluation and comparison was done by inspecting the DIC of each model, examining the plots of model fits and running posterior predictive checks. DIC is a measure of model fit based on how much the model deviates from the data after applying a penalty for model complexity (Spiegelhalter, Best, Carlin, & Linde, 2014). Thus, lower DIC values represent better fit. We also performed model recovery simulation where we generated simulated strategy data from either a power or abrupt function. Results indicated that when data was generated using a power function, DIC was lower for the power model than the abrupt model, and similarly the abrupt model had a lower DIC for data generated from an abrupt function (see online supplement for details).

General Data Analytic Approach

The results section is structured as follows. First, to ensure that performance on the WAB Task was similar to most skill acquisition tasks, the relationship between practice and task completion time was examined graphically and by analyzing the fit of a power function. Second, we examined the relationship between global inefficient strategy use and performance using a second power function with an extra parameter representing inefficient strategy use. Third, we examined plots of the relationship between practice and strategy use at each level of strategy decomposition at the group- and individual-level and examined the fit of the Bayesian models. Finally, we conducted posterior predictive checks to further assess the abruptness of strategy change and model fit.

Results

Practice and Performance

Average task completion time by block is shown at the group level in Figure 3 and at the individual level for a random sample of nine participants in Figure 4 (see online supplement for individual-level plots for additional participants). Both figures overlay power function model fits. As expected the figures show that task completion time decreased substantially in the first few blocks and that the rate of change tended to decrease with practice approaching an asymptote. The individual-level graphs show that the general pattern of monotonically decreasing and decelerating task completion time applied across individuals. However, there was variation between individuals in initial performance, rate of speed-up and block-to-block variation in performance. Table 1 presents parameter estimates and model fits for the power function models of task performance both with and without global strategy use as a covariate. The model fits superimposed on the data show that the power function provides reasonable fit for this data. Parameter estimates indicate that the expected change in task completion time from the first block to asymptotic levels was 39.6 seconds and average asymptotic task completion time was 11.5 seconds. They also indicate that there was more variability between individuals in amount of learning than asymptotic performance.

Figure 3. Group-level task completion time by block with power function model fit.

Figure 4. Individual-level task completion time by block with power function model fits for nine example individuals. Some individuals took longer than 60 seconds on block one, however, for ease of viewing the y-axis was truncated.

Table 1

Parameter Estimates and Model Fit Statistics for Models of Task Completion Time

	Power Function Without Strategy M [95% CI]	Power Function With Strategy M [95% CI]
Parameters		
Mean amount of learning $\mu_{\theta^{(1)}}$	39.61 [36.50, 42.99]	30.08 [27.08, 33.38]
SD amount of learning $\sigma_{\theta^{(1)}}$	20.41 [17.83, 23.47]	19.18 [16.54, 22.36]
Mean rate of learning $\mu_{\theta^{(2)}}$	1.06 [0.98, 1.14]	1.35 [1.22, 1.50]
SD rate of learning $\sigma_{\theta^{(2)}}$	0.46 [0.40, 0.54]	0.66 [0.55, 0.80]
Mean asymptotic performance $\mu_{\theta^{(3)}}$	11.50 [10.67, 12.37]	10.49 [9.83, 11.19]
SD asymptotic performance $\sigma_{\theta^{(3)}}$	4.22 [3.52, 5.06]	3.42 [2.89, 4.03]
Mean standard deviation μ_{σ}	3.38 [3.05, 3.75]	3.22 [2.90, 3.58]
SD standard deviation σ_{σ}	2.10 [1.80, 2.47]	2.05 [1.74, 2.42]
Effect of inefficient strategy use γ		17.68 [16.08, 19.22]
Fit statistics		
Deviance	11735	11439
Penalty	658	647
DIC	12393	12087

Note: For each parameter we present mean and 95% credible intervals for their respective posterior distributions.

Several analyses were conducted to examine the relationship between strategy use and performance. First, adding global strategy inefficiency as a predictor improved the power function model of task completion time (DIC went from 12393 without strategy to 12087 with strategy; see Table 1). The strategy parameter, γ , indicates that task completion time was predicted to be 17.7 seconds faster when strategy use was the most efficient (i.e., global strategy inefficiency = 0) compared to when it was least efficient (i.e., global strategy inefficiency = 1). Second, when strategy inefficiency and task completion time were averaged over the 15 blocks for each participant, a strong correlation was obtained between these two variables $r = .70$, $p < 0.001$. Thus, participants who used more efficient strategies tended to complete the task more quickly.

Practice and Strategy Use

The effect of practice on strategy use for each strategy is shown at the group level in Figure 5, and at the individual level in Figure 6 for six random participants (see online supplement for individual-level plots for additional participants). Both figures have model estimates from abrupt and power functions overlaid. As expected, group-level strategy change

was smooth. While the rate and amount of change varied between strategies at the group-level, all strategies showed monotonic reduction and monotonic deceleration in inefficiency. The individual-level data, however, displayed substantial variation between individuals and between strategies, with many instances of abrupt shifts. In particular component-level strategy use typically changed abruptly at a single point in practice while global measures tended to change more gradually.

Figure 5. Group-level strategy inefficiency by block for the nine strategies with abrupt and power function model fits

Figure 6. Individual-level strategy inefficiency by practice block for the nine strategies with abrupt and power function model fits for six sample participants. Global = inefficient strategy use, Info = inefficient information gathering, Filt = inefficient filtering, Rules = accessing rules, Name = asking for name, Exp = asking about experience, Spec = asking about special needs, Level = not using the level filter, Desel = not using the deselect all button.

Table 2 presents model fit statistics for the abrupt and power function for each measure of strategy use. The power function provided a better fit for global strategy use (DIC difference of -392) while the abrupt function fit better at the component and domain levels. Among the component strategies, the largest difference in DIC was seen for the deselect all strategy and the smallest difference in DIC observed for the level filter strategy. The similar fit of power and

abrupt functions for the level filter strategy may be due to the large proportion of participants who either always or never used the level filter (76%). At the domain level, the largest DIC difference was observed for inefficient filtering. Due to the consistency of not using the level filter, changes in this domain measure often reflected changes in not using deselect all. The smallest DIC difference at the domain level was observed for accessing rules. The graphs for this strategy indicate that although changes were generally gradual, they do not follow a power function.

Table 2

Model Fit Statistics for Abrupt and Power Functions of the Practice-Strategy Use Relationship

Strategy	Power			Abrupt			Difference
	Deviance	Penalty	DIC	Deviance	Penalty	DIC	DIC
Global							
Inefficient strategy use	12479	309	12788	12314	866	13180	-392
Domain							
Inefficient information gathering	7159	263	7422	5701	604	6305	1117
Inefficient filtering	6297	215	6511	4080	305	4385	2126
Accessing rules	8652	201	8853	7525	641	8165	688
Component							
Asking for name	2665	207	2872	1295	352	1647	1225
Asking about experience	3063	217	3279	1955	414	2369	910
Asking about special needs	4525	230	4755	2911	539	3450	1305
Not using the level filter	1040	128	1168	781	150	931	237
Not using deselect all	4880	174	5054	1531	283	1814	3240

Note: Lower DIC values indicate better model fit. DIC difference was calculated by subtracting abrupt DIC from power DIC. Thus, a positive DIC difference indicates that the abrupt model fit better than the power function and a negative DIC difference indicates that the power function fit better than the abrupt function.

Posterior Predictive Checks

To evaluate the degree to which the competing models captured features of strategy change in the data, posterior predictive checks were performed (for general discussion, see Anglim & Wynton, 2015; Gelman et al., 2014; Lynch & Western, 2004). Posterior predictive checks involve defining sample statistics of interest and then assessing how well a model is able to generate simulated data that captures these statistics. We assessed two sample statistics: (a) mean shift point and (b) proportion abrupt. To obtain the first statistic we calculated the point of greatest change between two blocks. Shift point was defined as the block before this point. We then calculated the mean shift point across individuals. Proportion abrupt provided a sample estimate of the proportion of individuals engaging in an abrupt strategy shift. To obtain this statistic, we first identified the block after which the greatest change in strategy use occurred. Here, we define the shift point as half way between this block and the next. We then obtained the variance in strategy use before and after the shift point, weighted by the number of blocks in each half (within variance), as well as the total variance. If the ratio of within variance to total

variance was less than 0.15 then that participant was considered to have shifted strategies abruptly. Excluding participants who did not change strategies at all during practice, we then calculated the proportion classified as abrupt shifters.

To calculate posterior predictive checks, 1000 simulated datasets were generated for both abrupt and power models and for each strategy. These data sets were generated by fitting each model with 50,000 iterations of 1 chain. Every 50th estimate was saved, thus generating 1000 datasets. Then, the sample statistics for the actual dataset and the mean, .025 quantile, and .975 quantile for the simulated statistics were calculated.

Table 3 presents the results of the posterior predictive checks for both statistics. The data statistics show that the shift point varied substantially between component strategies. This explains how aggregating over these components leads to more gradual changes. Table 3 also shows that both the abrupt and power functions are able to recover the mean shift point with some accuracy. However, data produced by the abrupt function tends to shift later in practice than what is seen in the actual data. This later shift could be the result of the truncation applied to the distribution of the shift parameter for the abrupt function.

Table 3

Posterior Predictive Checks – Proportion Abrupt Average Block of Greatest Strategy Change Across Individuals for Collected Data (Dataset) and Data Simulated From Each of the Three Models

Strategy	Proportion Abrupt			Mean Shift Point		
	Dataset	Power	Abrupt	Dataset	Power	Abrupt
Inefficient strategy use	.30	.11	.43*	2.73	2.56*	4.39
Domain						
Inefficient information gathering	.61	.30	.69*	2.12	2.16*	3.06
Inefficient filtering	.68	.15	.34*	4.07	3.88*	4.62
Accessing rules	.06	.03*	.16	6.94	5.65	6.73*
Component						
Asking for name	.89	.25	.66*	2.10	3.07*	3.16
Asking about experience	.70	.27	.65*	2.36	2.62*	2.81
Asking about special needs	.64	.25	.62*	2.86	2.74*	3.34
Not using the level filter	.38	.14	.24*	3.69	5.30	4.98*
Not using deselect all	.77	.16	.58*	4.20	4.07*	5.13

Note: Values presented here for the simulated data are the mean across the 1000 simulated datasets. The asterisk (*) highlights the most successful model at reproducing the sample statistic.

The proportion of individuals classified as abrupt was lowest for the global measure of inefficient strategy use as well as for the domain measure of accessing rules. This low proportion for accessing rules supports the assertion that while changes in this measure are not better modeled by a power function, they still occur gradually rather than abruptly. The highest values for proportion abrupt were observed for three of the component measures, supporting the assertion that at least some component strategy changes are abrupt. Generally, the power

function underestimates the proportion abrupt and is thus only superior to the abrupt function at capturing this statistic for accessing rules. When the proportion of abrupt shifters in the actual data was low, the proportion was overestimated by the abrupt function. However, when the proportion of abrupt shifters in the actual data was high, the proportion was underestimated by the abrupt function.

Discussion

The present study examined the abruptness of strategy change at component and aggregate levels. It sought to reconcile how abrupt changes in strategy use can coexist with gradual changes in task completion time. Consistent with past research (Anglim & Wynton, 2015; Newell & Rosenbloom, 1981), performance improvement on the WAB Task was gradual and well modeled by a power function, and strategy use was strongly correlated with performance. In particular, including global strategy use as a predictor in a model of the effect of practice on performance resulted in improved model fit. Importantly, the relationship between practice and strategy use varied greatly between individuals and between strategies. However, Bayesian hierarchical modeling was able to capture whether changes were generally abrupt or gradual while incorporating individual differences. Supporting expectations, component strategies often displayed abrupt shifts. Because these shifts typically varied in timing within individuals, changes in global strategy efficiency were relatively gradual and were often well approximated by a power function. Thus, the differential timing of component strategy shifts at the individual-level provides one explanation for why abrupt strategy shift can co-occur with gradual power function learning.

Practice, Strategy and Performance

Overall, the results are consistent with a multi-faceted explanation for how abrupt strategy shifts can co-occur with gradual power function learning. First, for a given individual the timing of abrupt changes in component strategies varies across components, which when aggregated to a global-level can lead to gradual strategy change. Second, the timing of strategy shifts was broadly distributed like a power-function where the probability of a strategy change decreased with practice. When strategy changes occur at the beginning of practice they coincide with a wide range of other learning processes, and variance in task completion time is also greater at the start of practice. Thus, the effect of a single strategy shift is harder to distinguish from other sources of variation early in practice. Third, some component strategy changes are less abrupt, which amplifies the gradual nature of global strategy change. Fourth, although not the focus on the present study, individuals can improve in how they execute a strategy (e.g., Delaney et al., 1998; Siegler, 1987); thus, the performance benefits of a new strategy may take time to accrue. Finally, performance metrics like task completion time are not direct reflections of latent skill. Variations in task features, chance, and the noise that accompanies executing cognitive, perceptual and motor processes leads to variations in task performance. In contrast, many of the strategies examined in this study are relatively discrete and deterministic choices. Thus, even if the strategy shift leads to an abrupt change in latent skill, the noise in performance execution will often obscure the latent change.

Our results support theories that suggest that the overall learning curve is composed of multiple component learning processes that have different functional forms to the overall learning curve. Research has shown that task performance can be decomposed into subtasks and strategies, and that the functional form of changes on subtasks and strategies does not necessarily reflect the trial-to-trial changes in overall task performance (e.g., Anglim & Wynton, 2015; Delaney et al., 1998; Kirsner & Spelman, 1996; F. J. Lee & Anderson, 2001; Rickard, 1997).

Often these studies have shown how the power function can occur as a result of summing components that change exponentially (e.g., Heathcote et al., 2000; Neves & Anderson, 1981; Rickard, 1997). More generally, several seminal papers have outlined the mathematics of how changes in component processes can yield a power function (Crossman, 1959; Haider & Frensch, 2002; Newell & Rosenbloom, 1981). The present study extends this past theorizing by measuring and formally modeling the changes in component strategies that underlie performance change. While mathematical modeling has shown what is possible, the present study provides, for a given task, an empirical estimate of the prevalence of the functional forms of several component learning processes.

The present study also contributes to an understanding of the factors that moderate the effect of practice on strategy use. Results suggest that abrupt strategy shifts are more likely when the strategy involves a single clear insight that is consistently useful across trials of a task. For example, for the current task, recognizing that certain questions never provide useful information could be derived from experience or an understanding of the task. Thus, the shift to using these strategies tended to happen early and abruptly. In contrast, use of a memory retrieval strategy, implied by not accessing the class rules, required knowledge that varied across trials (i.e., the age of the child determined the level). While supplementary analyses suggested that the shift for a given age tended to be abrupt, the large number of possible ages, meant that the overall shift to memory retrieval was less abrupt. This is similar to other studies examining item-level shift to memory retrieval (e.g., Palmeri, 1997, 1999; Rickard, 2004; Touron, 2006).

More generally, it is interesting to consider the broader context and the relevance of factors such as practice structure, task complexity, and strategy features on the generalization of the present results. First, the present study involved strategies that were clearly useful and easy to apply; in some cases new strategies can be difficult to apply (e.g., using a script to automate a task) and can lead to decreases in performance as the individual learns how to appropriately apply a new strategy (Yechiam et al., 2004). Second, simple tasks commonly studied in cognitive psychology, such as the noun-pair lookup task (e.g., Hertzog, Cooper, & Fisk, 1996) and the alphabet verification task (e.g., Haider & Frensch, 1996; Haider & Frensch, 1999) may not have the same multi-component structure as the current task. Skill acquisition on these simple tasks is defined more by the large number of stimulus-response pairs that need to be memorized in order to transition to a more efficient retrieval strategy. Similarly, very complex tasks such as chess, computer programming, or using sophisticated computer software may involve hundreds or thousands of component strategies. Finally, the structure of practice should influence the abruptness of strategy shift. The present study used massed practice on a highly repetitive task. When practice involves greater spacing, processes relating to forgetting may yield more gradual strategy acquisition and reversions to simpler or less memory-demanding strategies.

Modeling Strategy Change

The present study sought to show the benefits of a Bayesian approach to modeling skill acquisition, in general, and strategy use, in particular. Previous modeling approaches have had various limitations (Crossman, 1959; Gaschler et al., 2015; F. J. Lee & Anderson, 2001). In addition to the benefits to statistical inference (see Wagenmakers, 2007; Wetzels et al., 2011; Zyphur & Oswald, 2013), the present study illustrates the flexibility of the Bayesian approach, particularly in performing multilevel modeling of non-linear functions with non-normal data. The Bayesian approach also provided a more objective metric for judging the degree to which strategy changes were abrupt or gradual, and provided a means for integrating strategy use into models of the learning curve.

The formality of the modeling approach also helped to identify challenges in modeling strategy change at the individual-level. First, component level strategy use is more discrete than task completion time. At a trial level an individual either does or does not apply a strategy. This requires models that incorporate the natural constraints of probabilities. In earlier modeling efforts, we considered employing a logistic transformation, but this led to issues such as the power function being able to mirror the abrupt function. In addition, while the binomial distribution provided a formal model of counts of block-level strategy use, data often implied probabilities of strategy use close to zero or one. Second, the substantial variability in individual-level strategy change makes summarizing individual-level patterns more challenging. In particular, the aim of a task is to optimize performance (e.g., completion time); strategy use is typically subservient to this goal, which can lead to multiple functional forms of strategy change including constant strategy use (which we saw in this study could be approximated by the power and abrupt functions) and back and forth shifting, particularly for strategies with a weaker relationship with performance. While we focused on abrupt and power functions, and these worked quite well given the nature of the strategies being studied and the focus on differentiating between rapid/abrupt changes and learning curve changes, future research could consider a broader range of strategy change functions. Future research could also attempt to identify subgroups of individuals or component strategies that change in a similar pattern across practice and fit a range of different models to each subgroup, then identify the individual-level factors that could explain the different groupings.

Conclusion

In summary, the present research represents the first instance of the use of hierarchical models to examine strategy change at different levels of aggregation and contributes to the study of strategy change and how it relates to task completion time. The models allowed for a formal test of abruptness in individual-level strategy change. The task used mirrored many real world tasks where strategy use could be examined at component, domain, and global levels. Results showed that the Power Law of Practice holds at the group and individual level even when abrupt shifts occur in strategy use. This is partially because the changes in global strategy efficiency that underlie speed-up in task completion time occur gradually even when component strategy changes are abrupt.

References

- Alibali, M. W. (1999). How children change their minds: Strategy change can be gradual or abrupt. *Developmental Psychology*, 35(1), 127. doi:10.1037/0012-1649.35.1.127
- Anglim, J., Weinberg, M. K., & Cummins, R. A. (2015). Bayesian hierarchical modeling of the temporal dynamics of subjective well-being: A 10 year longitudinal analysis. *Journal of Research in Personality*, 59, 1-14. doi:10.1016/j.jrp.2015.08.003
- Anglim, J., & Wynton, S. K. A. (2015). Hierarchical Bayesian Models of Subtask Learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 41(4), 957-974. doi:10.1037/xlm0000103
- Averell, L., & Heathcote, A. (2011). The form of the forgetting curve and the fate of memories. *Journal of Mathematical Psychology*, 55(1), 25-35. doi:10.1016/j.jmp.2010.08.009
- Bajic, D., & Rickard, T. C. (2009). The temporal dynamics of strategy execution in cognitive skill learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 35(1), 113. doi:10.1037/a0013647
- Bajic, D., & Rickard, T. C. (2011). Toward a generalized theory of the shift to retrieval in cognitive skill learning. *Memory & cognition*, 39(7), 1147-1161. doi:10.3758/s13421-011-0114-z
- Blessing, S. B., & Anderson, J. R. (1996). How people learn to skip steps. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 22(3), 576.
- Compton, B. J., & Logan, G. D. (1991). The transition from algorithm to retrieval in memory-based theories of automaticity. *Memory & cognition*, 19(2), 151-158. doi:10.3758/BF03197111
- Crossman, E. (1959). A Theory of the Acquisition of Speed-Skill. *Ergonomics*, 2(2), 153-166. doi:10.1080/00140135908930419
- Crowley, K., Shrager, J., & Siegler, R. S. (1997). Strategy discovery as a competitive negotiation between metacognitive and associative mechanisms. *Developmental Review*, 17(4), 462-489. doi:10.1006/drev.1997.0442

- Delaney, P. F., Reder, L. M., Staszewski, J. J., & Ritter, F. E. (1998). The strategy-specific nature of improvement: The power law applies by strategy within task. *Psychological science*, 9(1), 1-7. doi:10.1111/1467-9280.00001
- Dennis, S., Lee, M. D., & Kinnell, A. (2008). Bayesian analysis of recognition memory: The case of the list-length effect. *Journal of Memory and Language*, 59(3), 361-376. doi:v
- Estes, W. K. (1956). The problem of inference from curves based on group data. *Psychological bulletin*, 53(2), 134. doi:10.1037/h0045156
- Farrell, S., & Ludwig, C. J. (2008). Bayesian and maximum likelihood estimation of hierarchical response time models. *Psychonomic Bulletin & Review*, 15(6), 1209-1217. doi:10.3758/PBR.15.6.1209
- Gaschler, R., Marewski, J. N., & Frensch, P. A. (2015). Once and for all—How people change strategy to ignore irrelevant information in visual tasks. *The Quarterly Journal of Experimental Psychology*, 68(3), 543-567.
- Gelman, A., Carlin, J. B., Stern, H. S., & Rubin, D. B. (2014). *Bayesian data analysis* (Vol. 2): Taylor & Francis.
- Gelman, A., & Hill, J. (2007). *Data analysis using regression and multi-level/hierarchical models*. Hillsdale, NJ: Erlbaum.
- Ghisletta, P., Kennedy, K. M., Rodrigue, K. M., Lindenberger, U., & Raz, N. (2010). Adult age differences and the role of cognitive resources in perceptual-motor skill acquisition: Application of a multilevel negative exponential model. *The Journals of Gerontology Series B: Psychological Sciences and Social Sciences*, gbp126. doi:10.1093/geronb/gbp126
- Haider, H., & Frensch, P. A. (1996). The role of information reduction in skill acquisition. *Cognitive psychology*, 30(3), 304-337. doi:10.1006/cogp.1996.0009
- Haider, H., & Frensch, P. A. (1999). Information reduction during skill acquisition: The influence of task instruction. *Journal of Experimental Psychology: Applied*, 5(2), 129. doi:10.1037/1076-898X.5.2.129
- Haider, H., & Frensch, P. A. (2002). Why aggregated learning follows the power law of practice when individual learning does not: Comment on Rickard (1997, 1999), Delaney et al. (1998), and Palmeri (1999). doi:10.1037/0278-7393.28.2.392
- Haider, H., Frensch, P. A., & Joram, D. (2005). Are strategy shifts caused by data-driven processes or by voluntary processes? *Consciousness and Cognition*, 14(3), 495-519. doi:10.1016/j.concog.2004.12.002
- Heathcote, A., Brown, S., & Mewhort, D. (2000). The power law repealed: The case for an exponential law of practice. *Psychonomic Bulletin & Review*, 7(2), 185-207.
- Hertzog, C., Cooper, B. P., & Fisk, A. D. (1996). Aging and individual differences in the development of skilled memory search performance. *Psychology and aging*, 11(3), 497. doi:10.1037/0882-7974.11.3.497
- Jaeger, T., Moulding, R., Anglim, J., Aardema, F., & Nedeljkovic, M. (2015). The Role of Fear of Self and Responsibility in Obsessional Doubt Processes: A Bayesian Hierarchical Model. *Journal of Social and Clinical Psychology*, 34(10), 839. doi:10.1521/jscp.2015.34.10.839
- John, B. E., & Lallement, Y. (1997). *Strategy use while learning to perform the Kanfer-Ackerman air traffic controller task*. Paper presented at the Proceedings of the Nineteenth Annual Conference of the Cognitive Science Society.
- Kirsner, K., & Speelman, C. (1996). Skill acquisition and repetition priming: One principle, many processes? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 22(3), 563. doi:10.1037/0278-7393.22.3.563
- Kruschke, J. K. (2010). *Doing Bayesian data analysis: A tutorial introduction with R*. Burlington, MA: Academic Press.
- Lee, F. J., & Anderson, J. R. (2001). Does learning a complex task have to be complex?: A study in learning decomposition. *Cognitive psychology*, 42(3), 267-316. doi:10.1006/cogp.2000.0747
- Lee, M. D. (2008). Three case studies in the Bayesian analysis of cognitive models. *Psychonomic Bulletin & Review*, 15(1), 1-15. doi:10.3758/PBR.15.1.1
- Lee, M. D. (2011). How cognitive modeling can benefit from hierarchical Bayesian models. *Journal of Mathematical Psychology*, 55(1), 1-7.
- Logan, G. D. (1988). Toward an instance theory of automatization. *Psychological review*, 95(4), 492. doi:10.1037/0033-295X.95.4.492
- Logan, G. D., & Etherton, J. L. (1994). What is learned during automatization? The role of attention in constructing an instance. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 20(5), 1022. doi:10.1037/0278-7393.20.5.1022
- Lynch, S. M., & Western, B. (2004). Bayesian posterior predictive checks for complex models. *Sociological methods & research*, 32(3), 301-335. doi:10.1177/0049124103257303
- Myung, I. J., Kim, C., & Pitt, M. A. (2000). Toward an explanation of the power law artifact: Insights from response surface analysis. *Memory & cognition*, 28(5), 832-840. doi:10.3758/BF03198418
- Neves, D. M., & Anderson, J. R. (1981). Knowledge compilation: Mechanisms for the automatization of cognitive skills. In J. R. Anderson (Ed.), *Cognitive skills and their acquisition* (pp. 57-84). Hillsdale, NJ: Erlbaum.
- Newell, A., & Rosenbloom, P. S. (1981). Mechanisms of skill acquisition and the law of practice. In J. R. Anderson (Ed.), *Cognitive skills and their acquisition* (Vol. 1, pp. 1-56). Hillsdale, NJ: Erlbaum.
- Palmeri, T. J. (1997). Exemplar similarity and the development of automaticity. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 23(2), 324. doi:10.1037/0278-7393.23.2.324
- Palmeri, T. J. (1999). Theories of automaticity and the power law of practice. doi:10.1037/0278-7393.25.2.543
- Pitt, M. A., & Myung, I. J. (2002). When a good fit can be bad. *Trends in cognitive sciences*, 6(10), 421-425. doi:10.1016/S1364-6613(02)01964-2
- Rickard, T. C. (1997). Bending the power law: A CMPL theory of strategy shifts and the automatization of cognitive skills. *Journal of Experimental Psychology: General*, 126(3), 288. doi:10.1037/0096-3445.126.3.288

- Rickard, T. C. (2004). Strategy execution in cognitive skill learning: an item-level test of candidate models. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 30(1), 65. doi:10.1037/0278-7393.30.1.65
- Siegler, R. S. (1987). The perils of averaging data over strategies: An example from children's addition. *Journal of Experimental Psychology: General*, 116(3), 250. doi:10.1037/0096-3445.116.3.250
- Siegler, R. S. (1988a). Individual differences in strategy choices: Good students, not-so-good students, and perfectionists. *Child Development*, 833-851.
- Siegler, R. S. (1988b). Strategy choice procedures and the development of multiplication skill. *Journal of Experimental Psychology: General*, 117(3), 258. doi:10.1037/0096-3445.117.3.258
- Siegler, R. S. (1991). Strategy choice and strategy discovery. *Learning and Instruction*, 1(1), 89-102. doi:10.1016/0959-4752(91)90020-9
- Spiegelhalter, D. J., Best, N. G., Carlin, B. P., & Linde, A. (2014). The deviance information criterion: 12 years on. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 76(3), 485-493.
- Tenison, C., & Anderson, J. R. (2016). Modeling the distinct phases of skill acquisition. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 42(5), 749-767.
- Thorndike, E. L. (1908). The effect of practice in the case of a purely intellectual function. *The American Journal of Psychology*, 374-384.
- Touron, D. R. (2006). Are item-level strategy shifts abrupt and collective? Age differences in cognitive skill acquisition. *Psychonomic Bulletin & Review*, 13(5), 781-786. doi:10.3758/BF03193997
- Touron, D. R., & Hertzog, C. (2009). Age differences in strategic behavior during a computation-based skill acquisition task. *Psychology and aging*, 24(3), 574.
- Touron, D. R., Hertzog, C., & Frank, D. (2011). Eye movements and strategy shift in skill acquisition: Adult age differences. *The Journals of Gerontology Series B: Psychological Sciences and Social Sciences*, 66(2), 151-159. doi:10.1093/geronb/gbq076
- Wagenmakers, E.-J. (2007). A practical solution to the pervasive problems of p values. *Psychonomic Bulletin & Review*, 14(5), 779-804. doi:10.3758/BF03194105
- Wasserman, L. (2000). Bayesian model selection and model averaging. *Journal of Mathematical Psychology*, 44(1), 92-107. doi:10.1006/jmps.1999.1278
- Wetzels, R., Matzke, D., Lee, M. D., Rouder, J. N., Iverson, G. J., & Wagenmakers, E.-J. (2011). Statistical evidence in experimental psychology: an empirical comparison using 855 t tests. *Perspectives on Psychological Science*, 6(3), 291-298. doi: 10.1177/1745691611406923
- Wilkins, N. J., & Rawson, K. A. (2011). Controlling retrieval during practice: Implications for memory-based theories of automaticity. *Journal of Memory and Language*, 65(2), 208-221. doi: 10.1016/j.jml.2011.03.006
- Yechiam, E., Erev, I., & Parush, A. (2004). Easy first steps and their implication to the use of a mouse-based and a script-based strategy. *Journal of Experimental Psychology: Applied*, 10(2), 89. doi:10.1037/1076-898X.10.2.89
- Zyphur, M. J., & Oswald, F. L. (2013). Bayesian probability and statistics in management research: A new horizon. *Journal of Management*, 39(1), 5-13.

Online Supplement

Model Recovery Simulation

To ensure our models were able to distinguish between abrupt and learning curve data, we conducted a model recovery simulation. This involves generating data from the two functions and ensuring that the model fitting and comparison process is able to identify the generating function (Gelman et al., 2014). To do this we fit each model using 10,000 iterations of one MCMC chain to inefficient strategy use data. Using the thinning function we saved every 1000th sample. This produced 10 data sets generated by each model for a total of 20 simulated datasets. Both abrupt and power functions were then fit to each simulated dataset, and the DIC values calculated. Then we calculated a DIC difference. For data generated from the power function we subtracted DIC power from DIC abrupt and for data generated from the abrupt function we subtracted DIC abrupt from DIC power. For power generated data the DIC for the power function was lower than for the abrupt for all 10 datasets ($DIC_{abrupt} - DIC_{power}$: $M = 2094.6$, $SD = 123.1$) and for the abrupt generated data the DIC for the abrupt function was lower than for the power function for all 10 datasets ($DIC_{power} - DIC_{abrupt}$: $M = 3994.4$, $SD = 133.7$). This suggests that the DIC is able to identify the data-generating model. The mean DIC difference for the power generated data was lower than the mean difference for abrupt generated data. This suggested that the abrupt function is better at capturing power data than the power is at fitting abrupt data.

Decomposing Accessing Rules

Within our task analysis we defined accessing the class rules as a domain level strategy. Although accessing the rules requires a single action on each trial (clicking on the “Show all class level information” button and reading the rules), the rule participants must use differs on each trial depending on the age of the simulated child for that trial. Thus accessing rules could not be easily decomposed into component strategies measured on each trial. However, as noted by Rickard (2004) a true understanding of memory retrieval can best be achieved by examining this strategy at the item-level. Therefore we examined how the use of memory retrieval changed over practice as a function of experience with each specific stimulus response pairing, the correct level for each age. Furthermore, we examined how memory retrieval changed over practice as a function of experience with each level (each level was associated with at least three ages). This enabled us to determine if participants were trying to remember a specific rule, or were just automatically remembering the stimulus-response pairing.

As the number of trials an individual completes in a block depends on how quickly they are completing the trials, and the age of the child on a particular trial was random, all individuals had a different level of experience with each rule and stimulus. Thus using hierarchical modeling to test for abrupt and gradual changes in component learning would have resulted in challenges in estimation. Instead, the mathematical function used in the posterior predictive checks to calculate the proportion of individuals classified as abrupt shifters, was used to assess the degree of abruptness of shift in memory retrieval. Only 6% of individuals were classified as abrupt shifters using this function when accessing rules was examined as an aggregated measure (see Table 3). For item-level shifts there were 22 different answers to the age question (i.e., 22 different stimulus-response pairings). We calculated the proportion of participants classified as abrupt for each question, which ranged from 0.38 to 0.87 ($M = 0.61$, $SD = 0.14$). Furthermore, when we examined how memory retrieval changed as a function of experience with each specific level we found a much lower proportion of abruptness ranging from 0.07 to 0.20 ($M = 0.13$, SD

= 0.05). Thus, participants often shifted abruptly from algorithm to memory retrieval based on their experience with a certain stimulus but when stimuli were grouped together into specific levels or all the rules together gradual change was observed.

While some research has suggested that volition is involved in using a previously created memory trace (Wilkins & Rawson, 2011), minimal research has examined whether volition is involved in creating the memory. In the current task, individuals could remember which class a child should be in based on their age either by (a) reading the rules, noticing the pattern and committing the entire set of rules to memory, (b) committing each rule to memory, or (c) encoding each age-level pairing to memory. As abrupt shifts were mainly seen at the component level, our results suggest that it is likely that participants were not choosing to remember the rules but were instead automatically remembering the correct response for each stimulus (age).

Detailed Graphs of Individual-Level Performance and Strategy Use

The following figures provide additional examples of individual performance and strategy curves.

Figure 7. Individual-level task completion time by block for 55 participants. Some individuals took longer than 70 seconds on block one, however, for ease of viewing the y-axis was truncated.

Figure 8. Individual-level global inefficient strategy use by block for 55 participants.

Figure 9. Individual-level inefficient information gathering by block for 55 participants.

Figure 10. Individual-level inefficient filtering by block for 55 participants.

Figure 11. Individual-level accessing rules by block for 55 participants.

Figure 12. Individual-level proportion of asking for name by block for 55 participants.

Figure 13. Individual-level proportion of asking about special needs by block for 55 participants.

Figure 14. Individual-level proportion of asking about experience by block for 55 participants.

Figure 15. Individual-level proportion of not using the level filter by block for 55 participants.

Figure 16. Individual-level proportion of not using the deselect all option by block for 55 participants.