Hierarchical Bayesian Models of Subtask Learning

Jeromy Anglim, Sarah K. A. Wynton¹

Citation information: Anglim, J., & Wynton, S. K. (2015). Hierarchical Bayesian Models of Subtask Learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 41, 956-974. <http://dx.doi.org/10.1037/xlm0000103>

Abstract. The current study used Bayesian hierarchical methods to challenge and extend previous work on subtask learning consistency. A general model of individual-level subtask learning was proposed focusing on power and exponential functions with constraints to test for inconsistency. To study subtask learning, we developed a novel computer-based booking task, which logged participant actions, enabling measurement of strategy use and subtask performance. Model comparison was performed using DIC, posterior predictive checks, plots of model fits, and model recovery simulations. Results showed that while learning tended to be monotonically decreasing and decelerating, and approaching an asymptote for all subtasks, there was substantial inconsistency in learning curves both at the group- and individual-levels. This inconsistency was most apparent when constraining both the rate and the ratio of learning to asymptote to be equal across subtasks, thereby giving learning curves only one parameter for scaling. The inclusion of six strategy covariates provided improved prediction of subtask performance capturing different subtask learning processes and subtask trade-offs. In addition, strategy use partially explained the inconsistency in subtask learning. Overall, the model provided a more nuanced representation of how complex tasks can be decomposed in terms of simpler learning mechanisms.

Keywords. learning, practice, skill acquisition, Bayesian hierarchical modeling, Power Law of Practice

1. Hierarchical Bayesian Models of Subtask Learning

Researchers have been modeling the effect of practice on task performance for over a century (Bryan & Harter, 1899; Heathcote, Brown, & Mewhort, 2000; Newell & Rosenbloom, 1981; Snoddy, 1926). F. J. Lee and Anderson (2001) extended such research to modeling the effect of practice on subtask performance. They examined how improvement in overall task performance could be understood in terms of improvements on subtask performance. Thus, they tested the often tacit assumption that basic research findings in cognitive psychology could be used to explain learning of complex real-world skills. In particular, they raised the question of whether the relationship between practice and subtask performance was consistent across subtasks.

¹ This article may not exactly replicate the authoritative document published in the journal. It is not the copy of record. The above doi links to the copy of record. Supplementary data and code is also available from the above doi.

Jeromy Anglim and Sarah K. A. Wynton, School of Psychology, Deakin University, Australia. We thank Simon A. Wynton for his assistance with programming the WAB Task. Correspondence concerning this article should be addressed to Jeromy Anglim, School of Psychology, Deakin University, 221 Burwood Highway, Burwood, 3125 Victoria, Australia. Email: jeromy.anglim@deakin.edu.au

Initial studies have suggested that subtask learning is often consistent (e.g., F. J. Lee & Anderson, 2001; Sohn, Douglass, Chen, & Anderson, 2005); however, we propose several improvements to modeling subtask learning that have the potential to challenge this finding. First, previous studies have analyzed learning data at the group-level, rather than the psychologically more meaningful individual-level. Second, there is a need to refine how subtask learning consistency has been parameterized. Third, studies have not formally incorporated strategy use into models of subtask learning consistency. In particular, Bayesian hierarchical modeling provides a flexible approach for estimating complete probabilistic models of individual-level subtask performance with strategy covariates. It also provides a useful model comparison framework, supporting data driven evaluation of model complexity.

To advance understanding of subtask learning consistency, we first developed a new task designed to measure strategy use and subtask performance. We then applied Bayesian hierarchical modeling to examine competing models of the effect of practice on overall and subtask performance at the individual-level. To apply such an approach we proposed several novel models of subtask learning and the effects of strategy use. Before describing the task and presenting the models, we review literature on the relationship between practice and performance first at the overall level and then at the subtask level.

1.1. Overall task learning

There are many mathematical models of the relationship between practice and task completion time (for reviews, see Heathcote et al., 2000; Lane, 1987; Newell & Rosenbloom, 1981). Newell and Rosenbloom (1981) analyzed many skill acquisition datasets finding that with practice, task completion time decreased monotonically, decelerated monotonically, and approached an asymptote. They found that learning data was well modeled by a power function, which they influentially labeled the ‘Power Law of Practice’. In its most commonly used three parameter form (hereafter referred to as the power function), the Power Law states that

$$E(y_j) = \theta^{(1)} x_j^{-\theta^{(2)}} + \theta^{(3)}$$

where expected task completion time, y_j , on trial $j = 1, \dots, J$, is a function of amount of practice, x_j ($x_1 = 1$), and parameters representing amount to be learnt $\theta^{(1)}$, rate of learning $\theta^{(2)}$, and asymptotic task completion time $\theta^{(3)}$.

Despite extensive use of the three parameter power function, researchers have continued to debate the merits of alternative models. Some alternatives include two and four parameter versions of the power function, the APEX function (Heathcote et al., 2000), and even models that imply discontinuous jumps in performance (e.g., Haider & Frensch, 2002). A particularly notable alternative is the three parameter exponential function (hereafter referred to as the exponential function)

$$E(y_j) = \theta^{(1)} \exp(-\theta^{(2)}(x_j - 1)) + \theta^{(3)}.$$

Heathcote et al. (2000) compared the power and exponential functions across a wide range of datasets. Importantly, they noted that most previous evaluations of the power function had been performed on group-level data rather than individual-level data. While group-level data is relevant to managing groups of workers, individual-level data is relevant to understanding the psychological learning process of particular individuals. As a result Heathcote et al. (2000) fit the power and exponential functions separately for each individual across a range of experimental tasks using ordinary least squares minimization. They found that while the power function

provided better fit at the group-level, the exponential function tended to fit better at the individual-level.

Despite the importance of individual-level analysis, Bayesian hierarchical modeling offers several advantages over fitting learning functions separately (Kruschke, 2010b). In particular the Bayesian approach allows data to be modeled jointly. Also, Bayesian model comparison approaches such as DIC and model recovery simulations allow for data driven penalties to model fit based on flexibility (for a textbook introduction, see Gelman et al., 2013). Averell and Heathcote (2011) used hierarchical Bayesian analysis to compare power and exponential functions in the context of the forgetting curve. They found that while the exponential model had a smaller deviance, the power model had better fit after incorporating a penalty for model complexity. Thus, as a secondary aim, this paper assesses whether a similar finding applies to Bayesian models of the relationship between practice and task completion time.

1.2. Subtask Learning

Many prominent theories of learning and cognition are concerned with relating different levels of processing (e.g., Anderson, 2002; Newell, 1994). F. J. Lee and Anderson (2001) noted the importance of understanding how the learning mechanisms studied in simple cognitive tasks relate to learning more complex tasks. In many performance domains, tasks can be decomposed into subtasks often using hierarchical approaches (e.g., Card, Moran, & Newell, 1983). F. J. Lee and Anderson (2001) proposed the reducibility hypothesis that states that learning complex tasks can be understood in terms of learning many smaller subtasks.

To study this hypothesis F. J. Lee and Anderson (2001) studied subtask learning on the Kanfer-Ackerman Air-Traffic Controller (KA-ATC) Task. The overall task was to land planes, which could be decomposed into subtasks which included moving, landing, and queuing planes. To test whether subtask learning was consistent, they fit power functions to group-level data on subtask completion time over practice. They argued that if constraining the rate parameter $\theta^{(2)}$ to be constant across subtasks did not significantly damage the fit to the data, then subtask learning was consistent across subtasks. Although there were some instances of inconsistent rates, they concluded that in general, subtask learning appeared to be consistent.

Despite being the first paper to our knowledge to model the consistency of subtask learning, there is substantial scope to build on the methods that they employed. First, their study analyzed group-level data, when individual-level data is more psychologically meaningful (for discussion, see Heathcote et al., 2000). In particular, group-level data is well known to smooth over a wide range of variation at the individual-level. Second, F. J. Lee and Anderson (2001) focused mostly on constraining the rate parameter across subtasks. Presumably this was an attempt to distinguish scale from shape, and only constrain shape. Third, the model was not a formal probabilistic model of the data generating process; instead, point estimates of parameters were estimated using least squares. Finally, the study did not explicitly model individual differences in strategy use, which are known to exist on the KA-ATC Task (John & Lallement, 1997). As will be discussed, such strategy changes may influence the consistency of subtask learning at the individual-level.

There are several reasons why subtask learning may be inconsistent, particularly at the individual-level. First, tasks differ in underlying learning mechanisms and scope for learning. For example, for some subtasks learning involves an obvious insight, whereas for others it will be the gradual accumulation of knowledge. Second, some strategy shifts alter the subtask structure of a task such that less time needs to be allocated to a particular subtask. Third, there

may be subtask trade-offs, where spending more time on one subtask reduces the time to complete a different subtask. While participants are generally asked to minimize overall task completion time, minimizing subtask completion time is not a goal in itself. Thus, part of learning a task is learning how to coordinate and allocate effort across subtasks.

Inconsistency in subtask learning may be related to strategy use. These include the coding of information into memory (Ackerman, 1992; Hertzog, Cooper, & Fisk, 1996), improvements in keyboard accuracy or mouse clicking speed (F. J. Lee & Anderson, 2001), a reduction in the time spent fixating on task irrelevant information (Haider & Frensch, 1999; Sohn et al., 2005), and the use of better, or more efficient, strategies (Yechiam, Erev, & Parush, 2004). Within a complex task it is possible that learning is impacted by all of these processes, but they contribute differently to each subtask. Processes that contribute differently to different subtasks may explain why subtask learning is sometimes inconsistent.

1.3. The Task

In order to obtain relevant data for studying the consistency of subtask learning, we created a new task, the Wynton-Anglim Booking (WAB) Task. The task represents a simplified version of a task used in actual swimming schools. The task's complexity was designed to be similar to the KA-ATC Task; i.e., more complex than a simple cognitive task, but not so complex that it would limit experimental study. The aim of the task was to book a child into a suitable swimming class. To complete the task, participants were required to (1) ask a simulated parent questions about their child, (2) filter a large timetable to show relevant times, and (3) select a time that conforms with task rules and meets the parent's needs. To correctly select a class from the timetable, participants had to follow three rules: (1) the class level must align with the child's age; (2) the class must not already be fully booked, where the number of children permitted in a class varies by level; and (3) the class must satisfy the requirements of the parent which could include day, time, or teacher limitations.

Figure 1 shows the main screen used for asking questions, accessing task rules, filtering the timetable, and triggering the display of the timetable. The timetable screen displayed available class times that also were consistent with any active filter. Classes were arranged in a grid with five columns for the days Monday to Friday and with classes sorted from earlier to later times.

SHOW TIMETABLE

NEXT

SHOW ALL CLASS LEVEL INFORMATION

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
<input type="button" value="Select All Teachers"/>	<input type="button" value="Select All Levels"/>	<input type="button" value="Select All Times"/>	<input type="button" value="Select All Days"/>	☐ Less Than 8
<input type="button" value="UNSelect All Teachers"/>	<input type="button" value="UNSelect All Levels"/>	<input type="button" value="UNSelect All Times"/>	<input type="button" value="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 viewed by participants. The speech bubbles at the bottom represent questions that can be asked. The five columns of check boxes correspond to the timetable filters. The SHOW TIMETABLE button would display the timetable consistent with the specified filters.

The task was designed to enable studying of subtask learning; thus, all mouse presses were recorded and were assigned to one of three subtasks. The first subtask was *information gathering* (I) which included asking questions and viewing the class rules. Each of these actions would bring up a modal dialog box with information. The participant would need to click to close this dialog box. The task was designed so that questions were either always relevant (age and day-time), sometimes relevant (level and teacher), or always irrelevant (name, experience, and special needs). The second subtask was *filtering* (F). There were five filters which when applied would alter the subsequent display of the timetable to only show classes that met the restrictions implied by the filters. The five filters were teacher, level, day, time, and class size. In the case of teacher, level, day, and time filters, the default status was for all options to be selected. Options could be deselected either one at a time or by pressing a button to deselect all and then selecting the one or more options of relevance. The class size filter defaulted to no restriction and participants could select a desired size restriction. The third subtask was *timetabling* (T). This involved pressing the show-timetable button and selecting a class from the timetable list. Depending on the degree of filtering applied, various degrees of visual scanning would be required. When a class was selected, it could be correct or incorrect. If it was correct, then the trial would end and feedback on trial completion time would be presented. If it was

incorrect, then one of several error messages would be displayed. If the participant overfilled a class or booked a child in the wrong level for their age, they received a message indicating the reason for the error. If the class did not meet a parent requirement (i.e., wrong time, day, or teacher), then the error message simply indicated that the class did not suit the parent.

The task could be performed in a variety of ways. We identified six strategies of particular relevance. In relation to information gathering, there was the proportion of the three possible irrelevant questions that the participant asked on a trial. By carefully following task instructions or through experience, participants were expected to learn not to ask the irrelevant questions. The second information gathering strategy was whether the rules were accessed on the trial. The rules stated how the age of the child related to the appropriate class level and stated the maximum number of children permitted in a class for the level. Thus, with sufficient exposure, participants were expected to be able to retrieve the rules from memory. This is analogous to the learning that occurs in many algorithm-retrieval studies of skill acquisition (e.g., Delaney, Reder, Staszewski, & Ritter, 1998). We also examined four filtering strategies: use of the level, time, and day filters, and whether, when selecting the level filter, the deselect all option was used. With regards to the filters, the level filter was assumed to be the most useful for simplifying timetable search. In contrast, day and time filters were considered to be less relevant, as the timetable was already arranged into days as columns and sequentially by time. In general, each of these three filters was expected to increase the time spent on filtering and reduce time spent on timetabling. In contrast, using the deselect-all level-filter button was expected to reduce the time required to apply the level filter. Instead of having to deselect the five levels that did not apply, using deselect-all required only two mouse clicks, one on deselect all, and a second to select the appropriate level.

Overall, subtasks differed in terms of how they were learnt and whether trade-offs were present. Information gathering required many more gradual learning processes to understand task rules and identify irrelevant information. Filtering involved some aspects of efficiency, but also required an appreciation of the optimal amount of filtering in order to facilitate subsequent timetabling. Finally timetabling was a relatively simple subtask, if information gathering and filtering had been performed effectively. It only required basic skills in scanning the timetable that would presumably be acquired fairly rapidly. Thus, the task enabled the measurement of overall and subtask performance along with strategy measurement to support the modeling of individual differences in learning.

1.4. Models

We now present our proposed models of overall and subtask performance. In simple terms, the models state that expected performance is a function of practice. The functions have parameters that describe important characteristics of the learning curve such as final performance, amount of learning, and rate of learning. Observed performance deviates from expected performance, and these deviations are also incorporated into the model and quantified. The models are hierarchical in that individuals have their own parameters (i.e., individuals vary in amount learnt, rate of learning, final performance, and variability around expected performance) and these parameters (the lower level) are assumed to be drawn from a distribution with hyperparameters (the upper level). Means and standard deviations of these distributions can then characterize the way that individuals vary in their learning curves (e.g., typical rate of learning, variation in rate of learning, etc.). For a useful introduction to Bayesian hierarchical model notation, see Kruschke (2010a).

Overall task performance. Overall task completion time y_{ij} for participant i ($i = 1, \dots, n$), at time point j ($j = 1, \dots, J_i$), where $J_i = 15$ for all individuals in the present study was modeled as follows:

$$\begin{aligned}
 y_{ij} &\sim \text{Gamma}(\alpha_{ij}, \beta_{ij}) \\
 \alpha_{ij} &= \frac{\mu_{ij}^2}{\sigma_i^2} \\
 \beta_{ij} &= \frac{\mu_{ij}}{\sigma_i^2} \\
 \mu_{ij} &= f(\theta_i, x_{ij}) \\
 \theta_i^{(q)} &\sim \text{Gamma}(\alpha_{\theta^{(q)}}, \beta_{\theta^{(q)}}) \\
 \alpha_{\theta^{(q)}} &= \frac{\mu_{\theta^{(q)}}^2}{\sigma_{\theta^{(q)}}^2} \\
 \beta_{\theta^{(q)}} &= \frac{\mu_{\theta^{(q)}}}{\sigma_{\theta^{(q)}}^2} \\
 \sigma_i &\sim \text{Gamma}(\alpha_\sigma, \beta_\sigma) \\
 \alpha_\sigma &= \frac{\mu_\sigma^2}{\sigma_\sigma^2} \\
 \beta_\sigma &= \frac{\mu_\sigma}{\sigma_\sigma^2}
 \end{aligned}$$

where x_{ij} is the practice block number indexed $1, 2, \dots$, and θ_i is a vector of parameters for each individual i . The formula $y_{ij} \sim \text{Gamma}(\alpha_{ij}, \beta_{ij})$ states that y_{ij} is distributed ($\sim$) according to a Gamma distribution with parameters α_{ij} and β_{ij} . The Gamma distribution was used because it is strictly positive and is well suited to modeling task completion times. A shift parameter on the Gamma distribution was not necessary given that minimum task completion times were expected to be well above zero. A common parameterization of the Gamma distribution involves α (shape) and β (rate) parameters. We reparameterize the Gamma distribution (i.e., $\alpha = \mu^2 / \sigma^2$ and $\beta = \mu / \sigma^2$) in terms of means μ and standard deviations σ , which better aligns with theoretical properties of interest (i.e., expected performance and variability around expected performance).

We focus on two variants of this model based on the choice of learning function $f(\theta_i, x_{ij})$. The power function is

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

and the exponential function is

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

The learning curve parameters correspond to the amount of learning $\theta_i^{(1)}$, the rate of learning $\theta_i^{(2)}$, and asymptotic performance $\theta_i^{(3)}$ for participant i . Each element of θ , indexed $q = 1, \dots, Q$, varies over individuals and is assumed to be drawn from a Gamma distribution. In addition, the standard deviation of scores around expected scores σ_i is assumed to be Gamma distributed across individuals.

Subtask performance. We propose a similar model for the effect of practice on subtask performance time:

$$\begin{aligned} y_{ijk} &\sim \text{Gamma}(\alpha_{ijk}, \beta_{ijk}) \\ \alpha_{ijk} &= \frac{\mu_{ijk}^2}{\sigma_i^2} \\ \beta_{ijk} &= \frac{\mu_{ijk}}{\sigma_i^2} \\ \mu_{ijk} &= f(\theta_{ik}, x_{ijk}) \times \exp\left(\sum_{p=1}^P \lambda_{pk} z_{ijp}\right) \\ \theta_{ik}^{(q)} &\sim \text{Gamma}(\alpha_{\theta_k^{(q)}}, \beta_{\theta_k^{(q)}}) \\ \alpha_{\theta_k^{(q)}} &= \frac{\mu_{\theta_k^{(q)}}^2}{\sigma_{\theta_k^{(q)}}^2} \\ \beta_{\theta_k^{(q)}} &= \frac{\mu_{\theta_k^{(q)}}}{\sigma_{\theta_k^{(q)}}^2} \\ \sigma_i &\sim \text{Gamma}(\alpha_\sigma, \beta_\sigma) \\ \alpha_\sigma &= \frac{\mu_\sigma^2}{\sigma_\sigma^2} \\ \beta_\sigma &= \frac{\mu_\sigma}{\sigma_\sigma^2} \end{aligned}$$

where y_{ijk} is subtask completion time for participant i on block j for subtask k ($k = 1, \dots, K$), where $K = 3$ in this study ($k = 1$ for information gathering, $k = 2$ for filtering, and $k = 3$ for timetabling, also denoted, I , F , and T , respectively), and z_{ijp} denotes scores on P strategy time-varying covariates ($p = 1, \dots, P$) for individual i in block j .

Regression coefficients λ_{pk} allow the effect of z_{ijp} to vary over covariates and subtasks. This parameterization of strategy covariates ensures that predicted subtask completion time does not go below zero. It also expresses our expectation that strategy effects will be greater when expected subtask completion time is longer. While it would be possible to estimate separate

strategy covariate parameters for amount learnt and asymptote parameters, preliminary analyses suggested that this led to identification issues and reduced interpretability.

The parameter σ_i was allowed to vary over individuals, but was held constant over subtasks within individuals. Variation over individuals was allowed in order to capture the effect of differential numbers of trials per block and the large overall amount of individual differences. Variation was held constant over subtasks because each subtask was on a similar timescale, and preliminary investigations indicated that having a separate σ_{ik} for each subtask resulted in estimation issues. Note that σ_i in the subtask model is distinct from the corresponding parameter in the overall task model.

We also examine $f(\theta_{ik}, x_{ijk})$ in terms of a power function

$$\theta_{ik}^{(1)} x_{ijk}^{-\theta_{ik}^{(2)}} + \theta_{ik}^{(3)},$$

and an exponential function

$$\theta_{ik}^{(1)} \exp(-\theta_{ik}^{(2)}(x_{ijk} - 1)) + \theta_{ik}^{(3)}.$$

We note that participants are asked to optimize overall performance and not subtask performance. Also, some strategy shifts can lead to increased time spent on one subtask in order to improve overall performance. This could occasionally result in some non-standard relationships between practice and subtask completion time, such as abrupt or gradual increases. Thus, an aim of this paper is to highlight the limitations of traditional functions in modeling subtasks, and test whether the inclusion of strategy covariates can capture these non-monotonically improving subtask learning curves.

To examine subtask learning consistency, we examine the effect of two constraints applied to both the power and exponential functions. First, we define a rate constraint which constrains $\theta_{ik}^{(2)}$ to be equal across subtasks (i.e., $\theta_{ik}^{(2)} = \theta_i^{(2)}$ for all k). Second, we define a gamma constraint whereby $\theta_{ik}^{(3)}$ is a constant proportion of $\theta_{ik}^{(1)}$ (i.e., $\theta_{ik}^{(3)} = \gamma_i \theta_{ik}^{(1)}$ for all k). One or both of these constraints can be applied to both power and exponential functions. For example, both constraints applied to the power function is denoted

$$\theta_{ik}^{(1)} x_{ijk}^{-\theta_i^{(2)}} + \gamma_i \theta_{ik}^{(1)}.$$

The constraints in the above model differ somewhat from those applied in F. J. Lee and Anderson (2001). Whereas they constrained rate $\theta^{(2)}$ and asymptote $\theta^{(3)}$ parameters to be equal across subtasks, we think that constraining the asymptote is inappropriate when subtasks are on different time scales. Instead we prefer to constrain the ratio of learning to asymptote (i.e., the gamma constraint). Combining the rate and gamma constraints allows the shape of subtask learning curves to be consistent while allowing the scale to vary over subtasks. Thus, if learning is consistent, then a subtask that initially takes twice as long as another subtask, should still take twice as long after practice, even if task completion time is dramatically reduced. While it is interesting to study the effect of each constraint separately, showing that applying a single constraint does not substantially damage fit, may not be a particularly strong claim given the degree of flexibility that remains for two parameters to adapt to any monotonic decreasing, decelerating data. Thus, applying both constraints provides a stronger test of subtask learning consistency.

1.5. Bayesian Modeling

The above models go by various names including hierarchical, mixed, and random-effects models. They provide a complete specification of the likelihood. Thus, they are superior to procedures that only model the mean function using optimization criteria such as least squares. They also explicitly model the variation in parameters across individuals, making them superior to approaches that model each individual separately.

Bayesian methods provide an effective framework for estimating parameters for such hierarchical models, and have been increasingly adopted in psychology (Averell & Heathcote, 2011; M. D. Lee, 2011; Scheibehenne, Rieskamp, & Wagenmakers, 2013; Scheibehenne & Studer, 2014; Shiffrin, Lee, Kim, & Wagenmakers, 2008). This reflects a transition in many fields from reliance on frequentist methods to an appreciation of the benefits of Bayesian inference. The transition in psychology has been aided by several factors. Introductory texts (Gelman & Hill, 2007; Kruschke, 2010a; M. D. Lee & Wagenmakers, 2012; Shiffrin et al., 2008) have been written; computers have increased in power; and algorithms required for estimating Bayesian models have become more accessible using various high level modeling languages such as WinBUGS, OpenBUGS, JAGS, and Stan. Bayesian hierarchical modeling is particularly suited to understanding repeated measures psychological data as demonstrated in recent studies on the forgetting curve (Averell & Heathcote, 2011) and reasoning (M. D. Lee, 2008).

The Bayesian approach provides several advantages for the present study. First, Bayesian modeling permits clear and direct inferences about parameter uncertainty. Second, parameter estimates incorporate the uncertainty in all other parameters. Third, Bayesian modeling is a very flexible framework. Thus, modifications can easily be incorporated such as allowing error variance to vary across individuals or modeling data or individual parameters using non-normal distributions. Finally, Bayesian modeling provides tools for model comparison that incorporate model flexibility and complexity.

1.6. The Current Study

The current study examined the effect of practice on overall and subtask performance on the WAB Task. First, we tested the relative support for power and exponential functions to model learning on overall and subtask performance. Second, we estimated the effect of strategy use on subtask performance to evaluate the presence of subtask trade-offs and specificity. Third, we assessed the degree of subtask learning consistency at the individual-level. Finally, we assessed the degree to which strategy use could explain the inconsistency in subtask learning.

2. Method

2.1. Participants

Twenty-five adult participants (17 female (68%), 8 male (32%)), aged 18-51 ($M=24.2$, $SD=7.8$) completed the study. Participants were recruited through an Australian university and from workers at a local IT firm. As an incentive all individuals who participated were placed in a draw to win one of ten AUD\$100 gift cards. All participants reported spending at least 5 hours per week using a computer.

2.2. Wynton-Anglim Booking (WAB) Task

The WAB Task was developed using Microsoft Excel and VBA. The task was run on desktop computers with the Windows 7 operating system and LCD displays (at least 1680 x 1050 pixels). The task can be obtained from Figshare, "Wynton-Anglim Booking Task (The

WAB Task)", <http://dx.doi.org/10.6084/m9.figshare.938242> for use and modification by other researchers.

Each trial involved a different random configuration of the timetable, question responses, and correct answers. The timetable always included 120 classes with 24 classes per day and with three classes scheduled every 30 minutes between 3pm and 7pm. The three classes at a given time were each taught by one of the three teachers. Class level was randomly drawn from the set of possible levels. The number of existing bookings was drawn randomly from one to the maximum possible bookings which meant that most classes were not fully booked.

Answers to parent questions were randomly drawn from a weighted probability distribution. If there were not at least five correct classes in the timetable, the answers to parent questions were regenerated. For teachers, 70 percent of the time any teacher was sufficient, 15 percent of the time a specific teacher was required, and 15 percent of the time the teacher could be one of two possible teachers. For child age, the responses were evenly distributed across the age requirements for the classes. For time of day, 19 percent could be any time, 8 percent was one specific time, with the remainder reflecting various constrained time ranges. For day of the week (Monday to Friday), 20 percent was any day, 20 percent was a specific day, and the remainder was a subset of days. For the level question, 80 percent of the time the parent indicated that they did not know, for the remainder a level response was provided consistent with the child's age. The program logged every user action, the time since the previous action, and characteristics of the trial.

Trials commenced by pressing the next trial button and ended when the participant selected an appropriate class from the timetable. The time between actions was assigned to the action following it. Trial level subtask measurement was the sum of time spent on that subtask. Trial level strategy use was coded zero if the strategy was not used, and one if it was used, with the exception of the irrelevant questions strategy which was coded as the proportion of the three irrelevant questions asked. To equate time across participants, trials for each participant were grouped into one of fifteen blocks. This blocking strategy emphasizes time as the unit of practice, and is analogous to the method used in F. J. Lee and Anderson (2001). Block completion time and strategy use were the arithmetic mean of the corresponding included trial measures.

2.3. Procedure

Participants completed the study in small groups and were seated so as to prevent interaction with other participants. After completing a demographics questionnaire, participants read through general task instructions. They then completed as many trials of the task as they could within 50 minutes. The first trial of the task was considered a practice trial and the data from this trial was not used in analysis. Based on density plots, trials longer than 100 seconds were treated as outliers and removed (14 trials; 0.5% of all trials).

2.4. Data Analysis

All raw data and data analysis code will be made available on Figshare, "Analysis for Hierarchical Bayesian Models of Subtask Learning", <http://dx.doi.org/10.6084/m9.figshare.938243>. All models were estimated using Bayesian methods. Diffuse uniform priors with a lower bound of zero were used to model the μ and σ hyperparameters of each Gamma distribution. Specifically upper bounds on the uniform priors for the overall task and for the subtask were $\mu_{\theta(1)} = 50$, $\sigma_{\theta(1)} = 50$, $\mu_{\theta(2)} = 2$, $\sigma_{\theta(2)} = 2$, $\mu_{\theta(3)} =$

30, $\sigma_{\theta^{(3)}} = 30$, $\mu_{\sigma} = 20$, $\sigma_{\sigma} = 10$, $\mu_{\gamma} = 5$, and $\sigma_{\gamma} = 5$. Upper bounds were chosen that were well above the range of plausible values so that they did not substantively truncate the likelihood. The prior on strategy covariates was a normal distribution with a mean of zero and standard deviation of one. A normal distribution is relatively uninformative for the strategy covariates given that the effect of strategy covariates is being exponentiated.

Model comparison was performed using DIC, posterior predictive checks, and plots of model fits against the data. Posterior parameter estimates and DIC (Spiegelhalter, Best, Carlin, & Van Der Linde, 2002) was obtained using MCMC methods in JAGS. Four chains were run. Trace-plots of MCMC sampling were examined to identify any model estimation issues and determine an appropriate burn-in period. Four chains of 25,000 iterations each were used to sample from the posterior, which ensured that results did not change substantively from one simulation to the next. Model recovery simulations were run to assess the degree to which models could reliably recover specified parameters, and to assess the power of the DIC to discriminate between competing models. For plotting purposes, model fits were estimated as the expected value integrating over the posterior density of all parameters.

3. Results and Discussion

3.1. Overview of Results

First, we present plots of the raw data for overall performance, subtask performance and strategy use at the group- and individual-levels. The results are then structured around our four research questions: (1) Does a power or exponential function provide a better fit to the effect of practice on overall and subtask performance? (2) What are the effects of strategy use on subtask performance? (3) What effect do subtask parameter constraints have on model fits? (4) How does the inclusion of strategy covariates influence the effect of constraints on model fits?

Overall task performance. The relationship between practice and overall task completion time is shown at the group-level in Figure 2 and at the individual-level in Figure 3. The group-level plot shows the usual monotonically decreasing and decelerating completion times. The group on average took a little less than half the time to complete trials in the final block relative to the first block. The individual-level plots show that almost all participants show meaningful learning. There are substantial individual differences in average performance. In some cases, much of this learning occurred between the first and second blocks. In other cases, the learning was more gradual. Longer times are associated with greater trial-to-trial variance, but this is largely attributable to the blocking structure whereby time of the experiment was fixed and thus, participants who completed fewer trials, have blocks that aggregate over fewer trials.

Figure 2. Overall task completion time by block at the group-level with power and exponential model fits. Error bars show the standard error of the mean.

Figure 3. Overall task completion time by block at the individual-level with power and exponential model fits.

Subtask performance. Subtask completion time by block for each subtask is shown in each cell of Figure 4 for the group-level. Despite a little noise, the group-level plot shows how

task completion time for each subtask decreased monotonically. The group-level plot also shows how subtask completion time for filtering declined less with practice than for the other subtasks. Information gathering and timetabling show a similar ratio of time spent from initial block to final block, although speed-up on timetabling appears to be a little more rapid.

Figure 5 shows four illustrative cases at the individual-level along with model fits to be discussed later. At the individual-level, there is some diversity in the relative shape of subtask performance data. A feature of the individual data is that performance over time was mostly monotonically decreasing. However, filtering was fairly flat in several cases and two participants showed systematic increases in filtering time with practice. Clearly, traditional monotonically decreasing exponential and power functions will provide a poor fit for these two cases. However, these two participants also increased their use of filtering strategies suggesting that a model that included strategies should capture this increase in subtask completion time. Taken together the data hints at the expected trade-off whereby time spent filtering reduced time spent timetabling.

Figure 4. Subtask completion time by block at the group-level for models using the Power function with and without strategies and with and without constraints.

Figure 5. Subtask completion time by block at the individual-level for four illustrative participants (i.e., the four rows) with model fits overlaid for four models (i.e., the four columns) from models with and without strategies and with and without constraints.

Strategy use. The relationship between practice and strategy use is presented in Figure 6. At the group-level, all strategies changed with practice and appeared to decelerate and approach an asymptote. In particular, asking irrelevant questions and accessing the rules decreased substantially with practice while use of the deselect-all button on the level filter increased substantially with practice. Meanwhile, the other filtering strategies either showed a small increase or remained fairly stable. Consistent with the expected benefits of the level filter, it was used more than the other filters. Theoretically, using the deselect-all button made applying a level filter more efficient, but this was not necessarily obvious.

While the group-level relationships between practice and strategy use were smooth and monotonic, the relationships at the individual-level were highly varied. Some participants never used a strategy, others used it from the first trial. Some individuals changed gradually on a strategy, others changed abruptly. Thus, at this low level of methods of task execution, many diverse learning processes were taking place. As will be shown, the model suggested that many of these changes translated into small changes in task performance. It also highlights how optimizing time per strategy as indexed by strategy use is one element of learning.

Figure 6. Strategy use by block at the group-level.

3.2. Power versus Exponential Models

Overall task performance. At the overall task level, power and exponential models had similar fits with a slightly better penalized fit for the power. Table 1 reports Bayesian posterior means and 95% credible intervals for model parameters along with DIC for models of overall task performance, while Figure 2 and Figure 3 present model fits for group- and individual-level data respectively. Plots of model fits indicated a similar convergence with the data for both functions. To check that a shift parameter on the distribution of overall task completion times was not required we estimated the probability implied by the model for obtaining a task completion time less than one second. Over the $25 \text{ cases} \times 15 \text{ blocks}$ (375) the maximum probability was .00010 suggesting that a shift parameter was not required. While the individual-level plots show some diversity in the learning curves, power and exponential fits were both able to model these diverse learning curves. The power function provided a slightly better DIC than the exponential function. This larger DIC was due to the larger penalty for model complexity and in spite of the mean deviance for the exponential being lower than for the power. This larger penalty for the exponential is interesting given that both the power and exponential ostensibly have the same number of random coefficients. However, as Pitt and Myung (2002) observe, even when two models have the same number of parameters, models can differ in functional form complexity. Furthermore, the hierarchical nature of the models also influences the functional form complexity. Thus, the greater penalty for the exponential model may indicate that it is more flexible in fitting learning curve data than the power model.

Given that calculation of penalty and DIC are based partly on the mean deviance, we also examined the distribution of deviance values for non-normality or outliers. Skewness (power=0.25; exp=0.27) and kurtosis (power=0.12; exp=0.09) were close to zero and similar for both power and exponential models. This suggested that it was not the distributional form of the deviances that was causing the larger DIC for the exponential model.

Table 1

Parameter Estimates and Model Fit for Overall Task Completion Time by Block at the Individual-Level

Parameter	Power M (95% CI)	Exponential M (95% CI)
$\mu_{\theta(1)}$	33.05 (26.45, 41.46)	27.92 (22.60, 34.73)
$\mu_{\theta(2)}$	1.07 (0.76, 1.52)	1.03 (0.62, 1.63)
$\mu_{\theta(3)}$	13.32 (11.14, 16.10)	17.08 (14.78, 20.48)
μ_{σ}	4.01 (3.07, 5.31)	3.96 (3.02, 5.24)
$\sigma_{\theta(1)}$	17.59 (11.96, 26.27)	13.83 (9.14, 21.08)
$\sigma_{\theta(2)}$	0.71 (0.41, 1.23)	1.04 (0.52, 1.84)
$\sigma_{\theta(3)}$	3.97 (2.06, 7.49)	5.03 (2.59, 9.12)
σ_{σ}	2.52 (1.65, 3.95)	2.50 (1.63, 3.93)
Mean Deviance	1919	1911
Penalty	101	116
DIC	2020	2027

While providing similar model fits, there were some notable differences in parameter estimates. The power function had much lower mean asymptotic performance $\mu_{\theta(3)}$ and a much larger mean amount of learning parameter $\mu_{\theta(1)}$ than the exponential function. Also, the standard deviation of individual differences in asymptote $\sigma_{\theta(1)}$ and the amount of learning $\sigma_{\theta(3)}$ showed similar differences between power and exponential functions. The asymptote for the exponential model was only slightly less than mean performance on the final block, whereas the asymptote for the power model predicted that substantial improvement remained to be attained after the final block. A good model of practice should accurately predict performance improvements beyond the range of practice explicitly studied. While improvements in observed performance in the final blocks were very small relative to improvements in earlier blocks (see Figure 2), the asymptotic prediction relates to expected performance after infinite practice. The power model implies that greater gains would be attained after infinite practice. Given that performance after infinite practice can not readily be predicted, differences in the power and exponential asymptotes can not easily be used as evidence for the superiority of the power or exponential model.

To examine whether the models were over-parameterized (see Scheibehenne & Pachur, 2014), we examined the correlations of the posterior density estimates of the hyperparameters shown in Table 1 for both power and exponential functions. Each correlation between μ and σ for a given learning curve parameter [i.e., θ_1 , θ_2 , θ_3 , σ ; e.g., $\text{cor}(\mu_{\theta_1}, \sigma_{\theta_1})$, $\text{cor}(\mu_{\theta_2}, \sigma_{\theta_2})$] was large (mean r for exponential was .72 and for power was .74). Correlations across learning curve parameters were generally either close to zero or fairly small: $\text{cor}(\mu_{\theta_1}, \mu_{\theta_2})$ -.09 (power) and .01 (exponential), $\text{cor}(\mu_{\theta_1}, \mu_{\theta_3})$ -.12 (power) and -.25 (exponential), $\text{cor}(\mu_{\theta_2}, \mu_{\theta_3})$.28 (power) and .12 (exponential). This suggests that the models are not over-parameterized.

Subtask performance. The relative support for power and exponential models at the subtask level was similar to that at the overall task level. Fits were similar for power and exponential models with a slightly better penalized fit for the power function. Table 2 reports fit statistics and Table 3 reports parameter estimates for subtask models but for now we focus on comparing power and exponential functions for unconstrained models without strategies. DIC for the power function (DIC=3862) was lower than for the exponential function (DIC=3885). This difference is somewhat larger than what was seen at the overall task level. However, as with models of overall task performance, the exponential function had lower mean deviance but greater model complexity than the power function. This is very similar to the findings that Averell and Heathcote (2011) found when studying forgetting curves (i.e., a monotonically decreasing function), where deviance was less for the three parameter exponential, but after penalizing for complexity, the power had the better fit. Also, similar to the overall task, asymptotes were larger and amount of learning was smaller for exponential than for power at the subtask level.

Table 2

Model Fit Statistics for Models of Subtask Learning with and without Strategy Predictors		
Model	Without Strategies DIC (Penalty; Deviance)	With Strategies DIC (Penalty; Deviance)
Exponential		
No Constraints	3885 (226, 3659)	3506 (239, 3267)
Rate constraint	3937 (207, 3730)	3538 (218, 3320)
Gamma constraint	3963 (211, 3752)	3602 (223, 3379)
Both constraints	4497 (177, 4320)	3794 (183, 3611)
Power		
No Constraints	3862 (196, 3666)	3507 (214, 3293)
Rate constraint	3915 (185, 3730)	3527 (198, 3330)
Gamma constraint	3907 (188, 3719)	3562 (206, 3356)
Both constraints	4486 (156, 4329)	3807 (164, 3642)

To assess the power of the DIC comparison procedure to distinguish between power and exponential models in the present sample, a model recovery simulation was run. First, we simulated 20 datasets each from unconstrained power and unconstrained exponential models using obtained parameter estimates. Second, we fit both unconstrained power and exponential models to each of the 40 datasets. Finally, we obtained a DIC difference where for the power generated data we subtracted DIC power from DIC exponential, and for the exponential generated data we subtracted DIC exponential from DIC power.

The simulation showed that DIC reliably identified whether the data was generated from a power or an exponential model. DIC power was less than DIC exponential in 100% of simulations when data was generated from the power model ($DIC_{exp} - DIC_{power}$: $M = 181.3$, $SD=91.8$). Likewise, DIC exponential was less than DIC power in 100% of simulations when data was generated from the exponential model ($DIC_{power} - DIC_{exp}$: $M = 271.2$, $SD=134.0$). Furthermore, mean ($DIC_{power} - DIC_{exp}$) was substantially larger than mean ($DIC_{exp} - DIC_{power}$). This indicates that the exponential model is better at fitting power data than the power model is at fitting exponential data. This further supports the validity of the increased complexity penalty that DIC applies to the exponential model.

Finally, to ensure that our results were robust to the choice of prior, we conducted a prior sensitivity analysis. Specifically, we examined the effect of multiplying the upper bound of the uniform distributions on the hyperparameters by factors of 2 and 10 for both exponential and power models. Altering the prior did not substantively change parameter estimates or differences in DIC between power and exponential models.

Thus, in summary, power and exponential models provide a reasonable basis for examining subtask learning. Exponential and power functions share the core features of being monotonically decreasing, monotonically decelerating, and approaching an asymptote. Both power and exponential models provided flexibility to fit different rates, initial levels, and asymptotes. More fundamentally, there was minimal visual evidence for discontinuities in learning, and both the power and exponential functions provided a reasonable basis for assessing subtask consistency. Given the slightly better fits for the power function, subsequent analysis of the effect of strategy use and constraints on model fits mostly focus on the power function.

Table 3

Parameter Estimates and Model Fit for Subtask Completion Time by Block at the Individual-Level

Parameter	Power Unconstrained M (95% CI)	Exponential Unconstrained M (95% CI)	Power Constrained M (95% CI)	Exponential Constrained M (95% CI)
$\mu_{\theta_I^{(1)}}$	20.67 (15.99, 26.87)	17.32 (13.53, 22.33)	16.61 (13.07, 21.22)	13.05 (10.28, 16.46)
$\mu_{\theta_F^{(1)}}$	5.42 (2.72, 10.75)	3.74 (1.74, 8.38)	8.53 (6.48, 11.32)	6.58 (5.12, 8.58)
$\mu_{\theta_T^{(1)}}$	10.13 (7.78, 13.25)	9.07 (7.07, 11.68)	7.45 (5.75, 9.73)	5.90 (4.57, 7.83)
$\mu_{\theta_I^{(2)}}$	1.00 (0.76, 1.33)	0.85 (0.55, 1.35)		
$\mu_{\theta_F^{(2)}}$	0.32 (0.18, 0.53)	0.21 (0.07, 0.68)		
$\mu_{\theta_T^{(2)}}$	1.64 (1.18, 1.98)	1.34 (0.86, 1.79)		
$\mu_{\theta_I^{(3)}}$	6.01 (4.34, 8.89)	8.25 (6.83, 10.25)		
$\mu_{\theta_F^{(3)}}$	4.02 (2.07, 5.76)	5.19 (3.68, 6.81)		
$\mu_{\theta_T^{(3)}}$	2.42 (2.13, 2.80)	2.88 (2.50, 3.45)		
μ_{γ}			0.43 (0.28, 0.67)	0.77 (0.62, 0.99)
$\mu_{\theta^{(2)}}$			0.70 (0.48, 1.03)	0.76 (0.41, 1.32)
μ_{σ}	2.14 (1.57, 2.95)	2.13 (1.56, 2.95)	2.57 (2.00, 3.32)	2.51 (1.98, 3.22)
$\sigma_{\theta_I^{(1)}}$	12.58 (8.39, 19.40)	10.36 (6.92, 15.92)	9.09 (6.09, 13.86)	7.27 (4.91, 11.05)
$\sigma_{\theta_F^{(1)}}$	7.08 (3.13, 18.34)	5.54 (2.08, 16.40)	5.49 (3.66, 8.57)	4.16 (2.83, 6.41)
$\sigma_{\theta_T^{(1)}}$	6.24 (4.11, 9.73)	5.23 (3.48, 8.05)	4.43 (2.90, 6.91)	3.61 (2.33, 5.85)
$\sigma_{\theta_I^{(2)}}$	0.58 (0.35, 0.96)	0.75 (0.41, 1.41)		
$\sigma_{\theta_F^{(2)}}$	0.08 (0.00, 0.37)	0.18 (0.00, 1.02)		
$\sigma_{\theta_T^{(2)}}$	1.54 (0.99, 1.97)	1.61 (1.00, 1.99)		
$\sigma_{\theta_I^{(3)}}$	3.63 (1.74, 7.94)	3.45 (2.20, 5.75)		
$\sigma_{\theta_F^{(3)}}$	2.69 (1.48, 4.71)	2.86 (1.84, 4.46)		
$\sigma_{\theta_T^{(3)}}$	0.41 (0.11, 0.85)	0.75 (0.40, 1.36)		
σ_{γ}			0.28 (0.09, 0.64)	0.32 (0.17, 0.53)
$\sigma_{\theta^{(2)}}$			0.51 (0.29, 0.91)	0.92 (0.42, 1.82)
σ_{σ}	1.58 (1.03, 2.55)	1.58 (1.02, 2.55)	1.53 (1.03, 2.36)	1.47 (0.99, 2.21)

3.3. Effect of Strategy Use on Subtask Performance

Table 4 reports the parameter estimates for the effect of the six strategy covariates on performance for each subtask when estimated as part of the unconstrained power model. The $\exp(\lambda)$ parameters indicate the amount to which subtask completion time in seconds is expected to be multiplied when the strategy is used. Thus, a value less than one indicates faster and a value greater than one indicates slower completion times when using the strategy. So for example, the level filter increased filtering time by a factor of 3.09, but reduced timetabling by a factor of 0.59.

Overall, parameters for the strategy covariates showed a meaningful pattern of differential effects across subtasks. The effects included a mix of subtask specificity and subtask trade-offs. For instance, the magnitude of the effect of accessing rules and irrelevant questions was much greater for information gathering than for other subtasks. Also, use of the deselect all level filter specifically benefitted filtering. Interestingly, use of the three filters (level, day, and time) led to an increase in filtering time that was greater than the decrease in timetabling time. While use of these filters appears to be sub-optimal, there may have been cognitive benefits that justified their use. In particular using filters reduced the need for participants to store parent requirements in working memory. Filters also reduced the need to visually scan the timetable and may have reduced errors in timetable selection.

Table 4

Parameter Estimates ($\exp(\lambda)$) for Strategy Covariates by Subtask for Unconstrained Power Model of Subtask Performance

Parameter	Information Gathering	Filtering	Timetabling
	M (95% CI)	M (95% CI)	M (95% CI)
1. Filter level	0.97 (0.78, 1.19)	3.09* (2.45, 3.99)	0.59* (0.44, 0.80)
2. Irrelevant questions	1.45* (1.32, 1.60)	1.20* (1.04, 1.36)	1.31* (1.02, 1.69)
3. Access rules	1.51* (1.40, 1.62)	1.10 (1.00, 1.20)	1.24* (1.06, 1.46)
4. Filter day	1.18* (1.10, 1.27)	1.54* (1.39, 1.70)	0.97 (0.81, 1.18)
5. Deselect all level filter	1.05 (0.97, 1.13)	0.82* (0.76, 0.89)	0.96 (0.83, 1.11)
6. Filter time	1.01 (0.94, 1.08)	1.25* (1.13, 1.38)	0.79* (0.67, 0.95)

Note. $\exp(\lambda)$ values represent the predicted multiple of subtask completion time that results from using the strategy. Strategy use increases subtask completion time when $\exp(\lambda) > 1$, decreases when $\exp(\lambda) < 1$, and has no effect when $\exp(\lambda) = 1$.

* 95% credible interval does not include one.

3.4. Effect of Constraints on Model Fit without Strategies

Overall, for models without strategies applying either constraint on its own damaged fit somewhat and applying both constraints led to a major damage to fit. The application of two constraints provided a more challenging, and arguably more valid, test of subtask learning consistency. When both constraints were applied, only one parameter remained for scaling with

shape parameters constrained to be equal across subtasks within individuals. By this reasoning, subtask learning was inconsistent.

Table 2 presents DIC values for all subtask models and for now we focus on the power models without strategies. Constraining $\theta_k^{(2)}$ to be constant across subtasks (i.e., the rate constraint) noticeably decreased fit as did constraining $\theta_k^{(3)}$ to be an individually varying multiple of $\theta_k^{(1)}$ (i.e., the gamma constraint). The gamma constraint resulted in a greater damage to fit than did the rate constraint. However, applying both constraints increased DIC more than the sum of the increase in DIC that occurred for each of the single constraints. For example, for the power model without strategies, a model with both constraints increased DIC by 624 (i.e., 4486-3862) whereas the models with a single constraint only increased DIC by 53 for the rate constraint (3915-3862) and 45 for the gamma constraint (3907-3862). As an aside, the same pattern of results emerged for the exponential model without strategies.

We now compare plots of fits, parameters estimates and posterior predictive checks to understand the features of the data not captured by the constrained model (i.e., the model with both constraints). The discrepancy between model fits and data in Figure 4 and Figure 5 show how the constrained model without strategies fails to adequately predict performance on the initial and final blocks. Specifically, task completion time for the constrained model on the initial block is under predicted for information gathering and timetabling, and over predicted for filtering, and the asymptote is particularly over predicted for information gathering. In contrast, the unconstrained model does not have these problems.

The parameters for the constrained and unconstrained power models (see Table 3) reveal several reasons for the inconsistency in subtask learning. For the power model, the three subtasks have quite different mean learning rate parameters of 1.00, 0.32, and 1.64 yet when constrained it was 0.70 for all subtasks. Likewise, the mean ratio of amount of learning to asymptote also varied substantially across subtasks (6.01/20.67=.29; 4.02/5.42=.74; 2.42/10.13=.24) yet when constrained it was 0.43 for all subtasks. This was reflected in the data (see Figure 4 and Figure 5) where speed up in task completion is gradual and extensive for information gathering, gradual and minimal for filtering, and rapid and extensive for timetabling.

In order to investigate the capacity of subtask models to capture important features in the data, we performed posterior predictive checks. We simulated 10,000 datasets for each model where parameter values were sampled from the posterior distribution. For the real data and each simulated dataset, we calculated 16 statistics. These statistics along with the results of the posterior predictive checks are shown in Table 5. We then calculated the mean and 0.025 and 0.975 quantiles for each statistic in the simulated datasets. To index the convergence of a given statistic to the sample data, we calculated a posterior predictive p-value (Gelman et al., 2013). We first calculated the proportion of simulated statistics that were (a) less than sample statistic, and (b) greater than the sample statistic. The posterior predictive p-value was defined as two times the smaller of these two proportions. Thus, small p-values (e.g., $p < .01$) indicate that the model is failing to capture the feature of the data captured by the statistic.

We present Figure 7 to graphically illustrate the meaning of posterior predictive checks using the statistic of block 15 minus block 1 filtering task completion time. A density is shown for each of the four models and a vertical line indicates the sample data. If the distribution indicates that the sample statistic is unlikely, this suggests a limitation with the model. So for example, when data is generated from the unconstrained model without strategies, the proportion of mean filtering improvements less than the statistic for the actual data is .295. Therefore the

posterior predictive p-value is $.295 \times 2 = .59$. This indicates that this model captures the sample statistic reasonably well (i.e., posterior predictive $p > .05$).

Figure 7. Density plots of mean filtering improvement based on simulated data from the posterior distribution of the four models (with and without strategies; constrained and unconstrained). The dotted line indicates the value of the sample statistic for the data.

Table 5 presents the equations for the 16 posterior predictive check statistics. Statistics 1 to 3 capture the mean amount of learning (block 1 to block 15) on each subtask respectively. A good model should be able to capture the diversity of amount of learning across the three subtasks. Statistics 4 to 6 capture the standard deviation of amount of learning (block 1 to 15) on each subtask respectively. A good model should also be able to capture the degree to which individuals vary in how much they learn across each subtask. Statistics 7 and 8 start by

calculating the ratio of final completion time to initial completion time for each individual for each subtask (e.g., Filtering time block 15 = 5 seconds; Filtering time block 1 = 10 seconds; the ratio is $5/10=0.5$). We then calculate the standard deviation of this ratio for each individual over subtasks (e.g., information gathering: $15/20 = .75$; filtering: $5/10 = .50$; timetabling = $2/10 = .20$; $SD(.75, .50, .20) = .28$). A larger standard deviation of ratios indicates greater inconsistency between subtasks. Thus, Statistics 7 and 8 capture the mean and standard deviation over individuals of these standard deviations of ratios. Statistics 9 and 10 start by calculating the correlation of task completion time over blocks for each individual between each pair of subtasks (i.e., correlations between IG and F, IG and T, and F and T). We then obtain the mean of these three correlations for each individual. This statistic provides another index of the consistency of subtask learning curves, but also captures transient fluctuations in performance that affect all subtasks. Thus, Statistics 9 and 10 capture the mean and standard deviation of these mean correlations over individuals. Statistics 11 to 16 are included to assess the ability of the models to capture the variability in the data. Statistics 11 to 13 start by calculating the standard deviation of task completion time at each block for each subtask. We then take the mean of these 15 standard deviations to get a measure of average task completion time variability for each subtask. Statistics 14 to 16 operate similarly except that instead of calculating the standard deviation of each block, they calculate the 10th percentile of each block. Overall, Statistics 1 to 10 index subtask learning consistency and Statistics 11 to 16 index the distribution of performance.

Focusing initially on the models without strategies, the posterior predictive checks show that the constrained model fails to capture several important features of the data. Notably, the inclusion of constraints led to a failure to capture differences in how the amount of learning varies across subtasks. For instance, Statistics 1 to 3, which capture the mean amount learnt on each subtask respectively, are reproduced by the unconstrained model, but not the constrained model. Likewise the constrained model does not capture the mean standard deviation of the ratio of block 15 performance to block 1 performance over subtasks (see statistic 7) yet the unconstrained model does. This is likely due to the constrained model failing to capture the variability in amount learnt relative to the scale of subtask completion time. Statistics 4 to 6 relating to the standard deviation of amount learnt were captured by both constrained and unconstrained models. This may be due to variation in standard deviations relative to the mean being fairly similar on this particular task. Interestingly, Statistic 9 indicated that the correlation between subtask performance over time was slightly greater in the data than was implied by either the constrained or unconstrained model. This may be due to a range of factors including transient effects that affect multiple subtasks (e.g., changes in effort, distraction, etc.) and systematic learning not captured by the power function. Overall, it seems that the constrained model fails to capture several statistics relating to subtask learning consistency.

Statistics 11 to 16 all focus on the variability in the data and are relevant to appraising the distributional assumptions of the models. Overall, both constrained and unconstrained models did reasonably well in capturing the distribution of task completion times. The main failure to reproduce the data occurred on the filtering subtask. The constrained model overestimated the standard deviation of filtering, and the unconstrained model overestimated the 10th percentile. These two statistics are connected because for a given mean, larger standard deviations will predict lower 10th percentiles. These statistics highlight the challenge of modeling the diversity of learning curves that can occur at the subtask level. It is likely that a few outlier participants who infrequently used the filtering strategy are explaining these 10th percentiles, and that a parametric power model is only an approximation for these cases.

Table 5

Posterior Predictive Checks for Power Models of Subtask Learning with and without Strategy Predictors and with and without Constraints

Statistic	Equation	Actual Data	Without Strategies		With Strategies	
			Unconstrained (UWOS) ^b	Constrained (CWOS) ^b	Unconstrained (UWS) ^b	Constrained (CWS) ^b
1. Mean RT Improvement (Information Gathering)	$\frac{1}{N} \sum_{i=1}^N y_{i,1,I} - y_{i,15,I}$	15.93	15.96 (14.27 - 17.65) p=0.97	11.76 (9.86 - 13.61) p=0.00	15.84 (14.27 - 17.39) p=0.92	14.28 (12.56 - 15.95) p=0.05
2. Mean RT Improvement (Filtering)	$\frac{1}{N} \sum_{i=1}^N y_{i,1,F} - y_{i,15,F}$	1.99	2.35 (0.81 - 3.90) p=0.59	6.10 (4.37 - 7.75) p=0.00	1.99 (0.57 - 3.39) p=0.99	2.60 (1.11 - 4.10) p=0.37
3. Mean RT Improvement (Timetabling)	$\frac{1}{N} \sum_{i=1}^N y_{i,1,T} - y_{i,15,T}$	8.05	7.49 (5.79 - 9.16) p=0.53	5.11 (3.34 - 6.84) p=0.00	8.04 (6.47 - 9.60) p=0.99	6.69 (5.05 - 8.27) p=0.09
4. SD RT Improvement (Information Gathering)	$sd_i(y_{i,1,I} - y_{i,15,I})$	9.21	9.81 (7.82 - 12.08) p=0.59	8.95 (6.69 - 11.38) p=0.80	9.68 (7.79 - 11.84) p=0.65	8.60 (6.65 - 11.02) p=0.55
5. SD RT Improvement (Filtering)	$sd_i(y_{i,1,F} - y_{i,15,F})$	4.38	4.01 (2.33 - 8.61) p=0.49	5.33 (3.56 - 9.68) p=0.48	4.20 (2.66 - 8.35) p=0.56	4.45 (2.89 - 8.52) p=0.74
6. SD RT Improvement (Timetabling)	$sd_i(y_{i,1,T} - y_{i,15,T})$	6.40	5.52 (3.93 - 8.15) p=0.35	5.31 (3.49 - 8.03) p=0.31	5.39 (3.92 - 7.81) p=0.28	5.25 (3.65 - 7.65) p=0.24
7. Mean of the SD of the subtask final/initial ratios	$\frac{1}{N} \sum_{i=1}^N sd_k(\frac{y_{i,15,k}}{y_{i,1,k}})$	0.27	0.27 (0.23 - 0.31) p=0.84	0.19 (0.15 - 0.23) p=0.00	0.27 (0.24 - 0.31) p=0.68	0.24 (0.20 - 0.27) p=0.13
8. SD of the SD of the subtask final/initial ratios	$sd_i \left[sd_k \left(\frac{y_{i,15,k}}{y_{i,1,k}} \right) \right]$	0.13	0.12 (0.09 - 0.14) p=0.55	0.12 (0.09 - 0.15) p=0.78	0.12 (0.10 - 0.15) p=0.85	0.13 (0.10 - 0.15) p=0.93

Table 5 continued...

Statistic	Equation	Actual Data ^a	Without Strategies		With Strategies	
			Unconstrained	Constrained	Unconstrained	Constrained
			d (UWOS) ^b	(CWOS) ^b	d (UWS) ^b	(CWS) ^b
9. Mean of mean intercorrelations between subtasks over time	$\frac{1}{N} \sum_{i=1}^N \left[\frac{1}{3} \sum_{\substack{k_1, k_2 \in \{I, F, T\}, \\ k_1 \neq k_2}} \text{cor}(y_{i, k_1}, y_{i, k_2}) \right]$	0.54	0.40 (0.34 - 0.47) p=0.00	0.42 (0.36 - 0.48) p=0.00	0.42 (0.36 - 0.48) p=0.00	0.41 (0.35 - 0.47) p=0.00
10. SD of mean intercorrelations between subtasks over time	$\text{sd}_i \left[\frac{1}{3} \sum_{\substack{k_1, k_2 \in \{I, F, T\}, \\ k_1 \neq k_2}} \text{cor}(y_{i, k_1}, y_{i, k_2}) \right]$	0.29	0.27 (0.22 - 0.32) p=0.35	0.28 (0.22 - 0.33) p=0.86	0.30 (0.25 - 0.34) p=0.06	0.30 (0.25 - 0.35) p=0.24
11. Mean of block RT SDs (Information Gathering)	$\frac{1}{15} \sum_{j=1}^{15} \text{sd}_i(y_{ij, I})$	9.30	9.22 (8.34 - 10.17) p= 0.85	9.28 (8.41 - 10.24) p= 0.95	9.23 (8.39 - 10.14) p= 0.85	9.25 (8.37 - 10.19) p= 0.89
12. Mean of block RT SDs (Filtering)	$\frac{1}{15} \sum_{j=1}^{15} \text{sd}_i(y_{ij, F})$	3.37	3.72 (3.30 - 4.53) p=0.13	3.96 (3.50 - 4.76) p=0.01	3.61 (3.24 - 4.37) p=0.30	3.80 (3.36 - 4.60) p=0.06
13. Mean of block RT SDs (Timetabling)	$\frac{1}{15} \sum_{j=1}^{15} \text{sd}_i(y_{ij, T})$	6.42	5.71 (4.79 - 6.73) p=0.16	5.81 (4.90 - 6.83) p=0.23	5.58 (4.70 - 6.55) p=0.08	5.73 (4.83 - 6.72) p=0.17
14. Mean of block RT 10th percentiles (Information Gathering)	$\frac{1}{15} \sum_{j=1}^{15} 10\text{th percentile}(y_{ij, I})$	6.07	6.02 (5.65 - 6.38) p=0.81	5.93 (5.52 - 6.35) p=0.52	6.07 (5.73 - 6.38) p=0.98	6.00 (5.65 - 6.35) p=0.73
15. Mean of block RT 10th percentiles (Filtering)	$\frac{1}{15} \sum_{j=1}^{15} 10\text{th percentile}(y_{ij, F})$	0.84	1.44 (0.92 - 2.01) p=0.02	1.16 (0.70 - 1.67) p=0.20	1.62 (1.11 - 2.15) p=0.00	1.49 (1.02 - 2.01) p=0.00
16. Mean of block RT 10th percentiles (Timetabling)	$\frac{1}{15} \sum_{j=1}^{15} 10\text{th percentile}(y_{ij, T})$	2.09	1.90 (1.69 - 2.10) p=0.06	1.75 (1.50 - 2.01) p=0.01	1.94 (1.76 - 2.11) p=0.10	1.85 (1.66 - 2.05) p=0.02

Note. $\text{sd}_i(\cdot)$ indicates the calculation of the sample standard deviation over subscript i . Posterior predictive p-values less than .05 indicate failure of the model to recover the sample statistic and are bolded. RT=Task completion time; I=information gathering; F=Filtering; T=Timetabling; y_{ij} =block task completion time for individual i in block j .

^aStatistic calculated on the actual data.

^bMean, 0.25 quantile, 0.975 quantile, and posterior predictive p-value of the statistic when applied to data simulated from the posterior from the given model.

3.5. Effect of Constraints on Model Fit with Strategies

Inclusion of strategies in subtask models improved model fit and reduced the damage to fits caused by constraints. To examine the influence of strategies on the effect of constraints on model fits we compared four power models: unconstrained without strategies (UWOS),

constrained without strategies (CWOS), unconstrained with strategies (UWS), and constrained with strategies (CWS). As seen in Table 2 including strategy covariates in the unconstrained model substantially improved model fits even after incorporating the penalty for additional model parameters (i.e., $DIC_{UWOS} = 3862$; $DIC_{UWS} = 3507$). Also, the increase in DIC due to applying constraints was less when strategies were included ($DIC_{CWS} - DIC_{UWS} = 3807 - 3507 = 300$) than when they were not included ($DIC_{CWOS} - DIC_{UWOS} = 4486 - 3862 = 624$). This suggests that the inconsistency of subtask learning is partially explained by changes in strategy use, and their differential effects across subtasks.

To assess the statistical power of the differential reduction in DIC, we performed another model recovery simulation. First, we simulated 20 datasets each from the UWS and CWS models. Then, we fit all four models (UWOS, CWOS, UWS, and CWS) to each of the 40 datasets.

The simulation showed that DIC reliably identified whether the data was generated from a constrained or unconstrained model. When data was generated from the unconstrained model, DIC was less for the unconstrained than for the constrained model both when the estimator was with strategies ($DIC_{CWS} > DIC_{UWS}$: 100%; $DIC_{CWS} - DIC_{UWS}$: $M = 579.7$, $SD = 194.9$) and without strategies ($DIC_{CWOS} > DIC_{UWOS}$: 100%; $DIC_{CWOS} - DIC_{UWOS}$: $M = 672.9$, $SD = 225.2$). Similarly when data was generated from the constrained model, DIC was less for the constrained model than for the unconstrained model both when the estimator was with strategies ($DIC_{UWS} > DIC_{CWS}$: 90%; $DIC_{UWS} - DIC_{CWS}$: $M = 34.8$, $SD = 23.2$) and without strategies ($DIC_{CWOS} > DIC_{UWOS}$: 100%; $DIC_{CWOS} - DIC_{UWOS}$: $M = 39.0$, $SD = 15.8$).

Importantly, the model recovery simulation showed that the increase in DIC resulting from the application of constraints was reliably reduced when strategies were included in the model. For both estimators we calculated the increase in DIC caused by constraints and the effect of strategies on reducing this damage,

$$DIC_{DIFF} = (DIC_{CWOS} - DIC_{UWOS}) - (DIC_{CWS} - DIC_{UWS})$$

Results indicated that DIC_{DIFF} was systematically positive when the generator was unconstrained ($DIC_{DIFF} > 0$: 100%; DIC_{DIFF} : $M = 93.2$, $SD = 69.2$). In contrast, DIC_{DIFF} was not systematically different from zero when the generator was constrained ($DIC_{DIFF} > 0$: 42.1%; DIC_{DIFF} : $M = 4.2$, $SD = 21.5$). Thus, if data was generated from a constrained model, then inclusion of strategy variables would not affect differences in fit between unconstrained and constrained models.

The fits superimposed on group- and individual-level data in Figure 4 and Figure 5, reinforce the findings above. At the group-level inclusion of strategies in the constrained model somewhat reduced problems with capturing the asymptote and initial block performance. At the individual-level, while the UWOS model was able to model the data well, the error around the model would normally be seen as random variation. However, the UWS model was able to capture many of these non-monotonic features in the data and showed that some of the seemingly random variation was actually due to changes in strategy use. For example, changes in strategy use captured filtering times for participant 9 and 10. Also, participant 10, had a very abrupt shift in timetabling completion time that was not captured in the constrained model.

The posterior predictive checks (see Table 5) show how by including strategies, the CWS model removed several problems present in CWOS model in capturing features of the data. Specifically, the CWOS model had seven statistics with a p-value less than .05 compared to the CWS model which only had three. In particular, the CWS model was able to recover Statistics 1,

2, 3 and 7 which were particularly concerned with subtask learning consistency. Taken together, these results show that strategy use partially explained the inconsistency in subtask learning.

4. General Discussion

The main aim of this study was to examine the consistency of subtask learning at the individual-level. To address this question we proposed and tested several Bayesian hierarchical models of overall and subtask performance on a novel computer-based task. Overall, power and exponential models both provided a good fit for the data; this was true both at the overall and the subtask level. However, the power function provided a slightly more parsimonious model using DIC at overall and subtask levels. Both at the individual- and group-level, subtask learning curves showed a degree of inconsistency. Even after adjusting for general scaling, the relative amount of speed up and the rate of that speed up with practice varied across subtasks. Incorporating strategies in the model partially explained the inconsistency in subtask learning. In general these results suggest a refinement to the Lee and Anderson's (2001) reducibility hypothesis. In particular, our results suggest that subtask learning consistency is unlikely to hold at the individual-level, and does not hold on the present task.

4.1. Power versus Exponential Models

The present study makes several contributions to the debate about the relative merits of power and exponential functions in modeling learning data. First, we evaluated the two models using a hierarchical Bayesian approach which overcomes issues both with the mostly group-level analysis of Newell and Rosenbloom (1981), and the fitting of models separately by individuals as was done in Heathcote et al. (2000). Furthermore, the use of DIC and model recovery simulations allowed for more nuanced penalties for model flexibility. Second, our findings suggest that the superior unadjusted fit of the exponential as compared to the power model to individual-level data found by Heathcote et al. (2000) may be due to the exponential model's greater flexibility. This finding that the exponential model has lower mean deviance but slightly higher DIC was also found by Averell and Heathcote (2011) when studying the forgetting curve. Thus, this may be a more fundamental feature of the exponential model. To investigate the generality of this finding it would be interesting for future research to re-examine existing learning datasets using Bayesian hierarchical modeling to see whether the exponential model would still have the best fit after a penalty for model flexibility was incorporated. It may also be useful to study longer periods of practice in order to more substantially sample asymptotic performance.

4.2. Inconsistency in Subtask Learning Curves

Building on the work of F. J. Lee and Anderson (2001), our study makes several contributions to understanding subtask learning consistency. Our findings present a multi-layered perspective on consistency. At a basic level, subtask performance showed consistency at the broadest level whereby it was generally monotonically decreasing and decelerating, and approaching an asymptote. However, constraining parameters across subtasks presented a more nuanced perspective, highlighting substantial inconsistencies. In particular, by constraining two parameters to be equal across subtasks and only allowing one parameter to capture variation in scale, the substantial inconsistency could be observed.

Our study suggests several reasons for subtask learning inconsistency. First, subtasks can involve different learning processes leading to variation in the rate of learning and the amount to be learnt. Second, overall task learning requires learning subtasks and learning how to effectively combine those subtasks. For instance, individuals aim to optimize performance on the overall

task, not the subtask; as such, subtask trade-offs mean that greater performance on one subtask can be achieved by poor performance on another subtask. Third, it is likely that similar mechanisms which cause between-person variability in learning parameters may also cause within-person variability in subtask learning parameters.

Our findings contrast with Lee and Anderson's (2001) findings that subtask learning was consistent. Our study and their study differed in several ways. First, they modeled data at the group-level, we modeled data at the individual-level. Second, they used least squares estimation with chi-square model comparison, whereas we used Bayesian methods with DIC model comparison. Third, they applied rate $\theta^{(2)}$ and asymptote $\theta^{(3)}$ constraints whereas we applied rate $\theta^{(2)}$ and gamma (i.e., ratio of $\theta^{(3)}$ to $\theta^{(1)}$) constraints. Finally, they studied subtask learning consistency using an air traffic control task whereas we used a booking task.

Our task had quite prominent subtask trade-offs, whereby failing to spend sufficient time gathering information and filtering could substantially increase time spent on the timetabling subtask. It is unclear whether the KA-ATC Task involved subtask trade-offs, although trade-offs were not obviously present in group-level learning curves. Nonetheless, it seems that many tasks have some degree of trade-off between subtasks. In many tasks, strategy shifts can alter the subtask structure or importance of different subtasks. Likewise, often activities such as planning, gathering information, and checking can make subsequent tasks more accurate and efficient. Therefore, task structure will influence the degree of subtask trade-off, and these trade-offs are typically amplified at the individual-level.

4.3. Strategies: Explaining Subtask Inconsistency

Results suggest that the six strategies are influenced by several different learning processes. The reduction in accessing rules occurs as a result of memorizing the task rules; learning which questions and filters are most useful occurs by reasoning about the task or through acquiring experience of their effectiveness; in contrast the deselect all level filter requires a degree of insight to discover. Execution of these strategies are also driven by broader meta-strategies including (1) filtering versus scanning, (2) certainty versus trial and error, and (3) actively trying to memorize class rules versus acquiring memory gradually over time.

A final aim of this paper was to examine whether strategies could explain inconsistency in subtask learning. Strategy coefficients showed meaningful differential prediction of subtask performance consistent with the interpretation that strategies were measuring different learning processes. Some strategies captured learning on a particular subtask, whereas others captured trade-offs, including trade-offs between subtask performance, and trade-offs regarding the use of cognitive resources. Thus, incorporating strategy use improved prediction of subtask performance. More importantly, this improvement in prediction was greater for the constrained model. This supports the claim that strategy use partially explains inconsistency in subtask learning.

4.4. Conclusion

In summary, the present study makes several core contributions to the literature. First, we developed and applied Bayesian hierarchical modeling to individual differences in the effect of practice and strategy on subtask performance. The present study joins several others (e.g., Averell & Heathcote, 2011; M. D. Lee, 2008) in highlighting how longstanding questions in cognitive psychology can be approached in novel ways using Bayesian methods. Second, a novel task was developed for studying subtask performance, of which the task, the data, and the analysis code will be made available for other researchers to use. Third, the study refined

conceptions of and methods for evaluating subtask learning consistency. The results challenge and extend previous findings by showing that subtask learning can be inconsistent and that this inconsistency can be partially explained by strategy use.

5. References

- Ackerman, P. L. (1992). Predicting individual differences in complex skill acquisition: dynamics of ability determinants. *Journal of Applied Psychology*, 77(5), 598-614.
- Anderson, J. R. (2002). Spanning seven orders of magnitude: A challenge for cognitive modeling. *Cognitive Science*, 26(1), 85-112.
- Averell, L., & Heathcote, A. (2011). The form of the forgetting curve and the fate of memories. *Journal of Mathematical Psychology*, 55(1), 25-35.
- Bryan, W. L., & Harter, N. (1899). Studies on the telegraphic language: The acquisition of a hierarchy of habits. *Psychological Review*, 6(4), 345-375.
- Card, S. K., Moran, T. P., & Newell, A. (1983). *The psychology of human computer interaction*. Hillsdale, NJ: Erlbaum.
- 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.
- Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., & Rubin, D. B. (2013). *Bayesian data analysis*: CRC press.
- Gelman, A., & Hill, J. (2007). *Data analysis using regression and multilevel/hierarchical models*. New York: NY: Cambridge University Press.
- Haider, H., & Frensch, P. A. (1999). Eye movement during skill acquisition: More evidence for the information-reduction hypothesis. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 25(1), 172.
- 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). *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 28(2), 392-406.
- Heathcote, A., Brown, S., & Mewhort, D. (2000). The power law repealed: The case for an exponential law of practice. *Psychonomic Bulletin and 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), 497-520.
- 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.
- Kruschke, J. K. (2010a). *Doing Bayesian Data Analysis: A Tutorial Introduction with R*. Burlington, MA: Academic Press.
- Kruschke, J. K. (2010b). What to believe: Bayesian methods for data analysis. *Trends in Cognitive Sciences*, 14(7), 293-300.
- Lane, N. E. (1987). *Skill acquisition rates and patterns: Issues and training implications*: Springer-Verlag Publishing.
- 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.
- Lee, M. D. (2008). Three case studies in the Bayesian analysis of cognitive models. *Psychonomic Bulletin & Review*, 15(1), 1-15.
- Lee, M. D. (2011). How cognitive modeling can benefit from hierarchical Bayesian models. *Journal of Mathematical Psychology*, 55(1), 1-7.

- Lee, M. D., & Wagenmakers, E. (2012). *Bayesian cognitive modeling: A practical course*: Cambridge: Cambridge University Press.
- Newell, A. (1994). *Unified theories of cognition*. Cambridge, MA: Harvard University Press.
- Newell, A., & Rosenbloom, P. S. (1981). Mechanisms of skill acquisition and the law of practice. *Cognitive skills and their acquisition*, 1-55.
- Pitt, M. A., & Myung, I. J. (2002). When a good fit can be bad. *Trends in Cognitive Sciences*, 6(10), 421-425.
- Scheibehenne, B., & Pachur, T. (2014). Using Bayesian hierarchical parameter estimation to assess the generalizability of cognitive models of choice. *Psychonomic bulletin & review*, 1-17.
- Scheibehenne, B., Rieskamp, J., & Wagenmakers, E.-J. (2013). Testing adaptive toolbox models: A Bayesian hierarchical approach. *Psychological review*, 120(1), 39.
- Scheibehenne, B., & Studer, B. (2014). A hierarchical Bayesian model of the influence of run length on sequential predictions. *Psychonomic Bulletin & Review*, 21(1), 211-217.
- Shiffrin, R. M., Lee, M. D., Kim, W., & Wagenmakers, E. J. (2008). A survey of model evaluation approaches with a tutorial on hierarchical Bayesian methods. *Cognitive Science*, 32(8), 1248-1284.
- Snoddy, G. S. (1926). Learning and stability: a psychophysiological analysis of a case of motor learning with clinical applications. *Journal of Applied Psychology*, 10(1), 1-36.
- Sohn, M.-H., Douglass, S. A., Chen, M.-C., & Anderson, J. R. (2005). Characteristics of fluent skills in a complex, dynamic problem-solving task. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 47(4), 742-752.
- Spiegelhalter, D. J., Best, N. G., Carlin, B. P., & Van Der Linde, A. (2002). Bayesian measures of model complexity and fit. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 64(4), 583-639.
- 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), 89-96.