Measuring Implicit and Explicit Sequence Learning

Inauguraldissertation

zur

Erlangung des Doktorgrades

der Humanwissenschaftlichen Fakultät

der Universität zu Köln

nach der Prüfungsordnung vom 10. Mai 2010

vorgelegt von

Marius Barth

aus

Düsseldorf

Tag der Abgabe: 03. Dezember 2018

1. Berichterstatter: Prof. Dr. Christoph Stahl (Köln)
2. Berichterstatter: Prof. Dr. Hilde Haider (Köln)
Diese Dissertation wurde von der Humanwissenschaftlichen Fakultät der Universität zu Köln im Juni 2019 angenommen.
Tag der mündlichen Prüfung: 12. Juni 2019

Mein besonderer Dank gilt meinen beiden Betreuern Christoph Stahl und Hilde Haider, von denen ich in den vergangenen Jahren unschätzbar viel lernen durfte. Ich danke meinen Kolleginnen und Kollegen Roscoe Araujo, Frederik Aust, Karoline Bading, Julia Haaf, Tobias Heycke, Anita Jain, Joachim Radt und Philine Thomasius, die mit vielen hilfreichen Gedanken und wichtiger Unterstützung zum Gelingen dieser Arbeit beigetragen haben.

Für die Datenerhebung sowie Hilfe bei vielen kleineren und größeren Aufgaben danke ich unseren Hilfskräften Pauline Anton, Nora Becker, Hannuh Frings, Julia Haaf, Laura Henze, Fabia Högden, Melina Köchling, Julia Krasko, Katharina Mattonet, Lucy Moerschbacher, Philipp Musfeld, Marius Schmitz, Lisa Spitzer, Margret Vössing und Leif Wolters. Für die Datenerhebungen in Düsseldorf danke ich Conni Lebbing, Imge Ürer, Jan Czarnomski und Friederike Neugebauer, für die Bereitstellung der dortigen Untersuchungsräume danke ich Ute J. Bayen.

Mein besonderer Dank gilt zudem Benjamin Haarhaus für viele hilfreiche Kommentare und Anregungen zum inhaltlichen Aufbau dieser Arbeit. Zuletzt danke ich meiner Familie und meinen Freunden, die mich in den vergangenen Jahren immer unterstützt haben.

Chapter II basiert auf dem Manuskript:

Stahl, C., Barth, M., & Haider, H. (2015). Distorted estimates of implicit and explicit learning in applications of the process-dissociation procedure to the SRT task. *Consciousness and Cognition*, *37*, 27–43. doi: 10.1016/j.concog.2015.08.003

CS und HH entwickelten die Fragestellung; MB, CS und HH planten die Experimente; MB setzte die Experimente um; MB und CS analysierten die Daten; CS, MB und HH verfassten das Manuskript.

Chapter III basiert auf dem Manuskript:

Barth, M., Stahl, C., & Haider, H. (2018). Assumptions of the process-dissociation procedure are violated in implicit sequence learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition.* doi: 10.1037/xlm0000614

CS und HH entwickelten die Fragestellung; MB, CS und HH planten die Experimente; MB setzte die Experimente um; MB und CS analysierten die Daten; MB, CS und HH verfassten das Manuskript.

Contents

Abstract

Sequence learning in the serial reaction time task (SRTT) is considered a key demonstration of learning that may proceed both implicitly or explicitly. However, claims about the implicit or explicit nature of the acquired knowledge were frequently based on identifying tasks with one of these two types of learning. The present studies assume that probably no task is a pure measure of implicit or explicit learning; instead, within each task, it is necessary to disentangle the contributions of both types of learning by means of measurement models of task performance.

In the first two studies, we scrutinized a measurement model that has already been used to dissociate implicit and explicit sequence knowledge: The process-dissociation (PD) approach as applied to the generation task yields separate estimates of implicit and explicit sequence knowledge that are derived from two variants of the same task. It therefore avoids many problems of previous measurement approaches, and studies utilizing this approach have yet provided the most convincing evidence in favor of dissociable types of sequence knowledge. However, the PD approach comes with its own set of critical assumptions. In two studies, we investigated the assumptions underlying the PD approach and found that they are violated in applications to sequence learning. Taking these limitations into account, studies that utilized the PD approach do not provide firm evidence for two dissociable types of learning.

In a third study, we investigated the processes that are involved in the expression of implicit and explicit sequence knowledge in the SRTT. We found that stimulus encoding, response execution, and response selection mediated the expression of implicit sequence knowledge. An involvement of response selection indicates that implicit sequence learning is mediated by representations containing both stimulus and response features. In contrast, the acquisition of explicit sequence knowledge resulted in a decision bias towards regular responses, and other effects of sequence learning disappeared in the course of training, indicating that participants switched from stimulus-based to plan-based action control.

Zusammenfassung

Sequenzlernen in der seriellen Wahlreaktionsaufgabe (SRTT) gilt als entscheidende Demonstration von Lernen, das sowohl implizit als auch explizit ablaufen kann. Die Annahme, dass Sequenzwissen entweder implizit oder explizit vorliegt, beruht jedoch häufig auf der Gleichsetzung einer experimentellen Aufgabe mit einer dieser beiden Arten von Lernen. Die vorliegenden Studien gehen davon aus, dass wahrscheinlich keine Aufgabe ein reines Maß für implizites oder explizites Lernen ist; stattdessen ist es innerhalb jeder Aufgabe erforderlich, die Beiträge beider Arten von Lernen mit Hilfe von Messmodellen der Aufgabenbearbeitung zu trennen.

In den ersten beiden Studien untersuchten wir ein Messmodell, wie es bereits genutzt wurde, um implizites und explizites Wissen zu dissoziieren: Die Prozessdissoziationsprozedur, wie sie auf die Generierungsaufgabe angewendet wird, liefert getrennte Schätzer für implizites und explizites Sequenzwissen, die auf der Bearbeitung von zwei Varianten ein und derselben Aufgabe beruhen. Sie vermeidet daher viele Probleme früherer Messansätze. Studien mit diesem Ansatz haben bisher die überzeugendsten Befunde für dissoziierbare Arten von Sequenzwissen geliefert. Der PD-Ansatz beruht jedoch auf eigenen kritischen Annahmen. In zwei Studien untersuchten wir die Annahmen, die dem PD-Ansatz zugrunde liegen, und fanden Verletzungen dieser Annahmen. Unter Berücksichtigung dieser Einschränkungen liefern Studien, die den PD-Ansatz verwenden, keinen eindeutigen Nachweis für zwei dissoziierbare Formen von Lernen.

In einer dritten Studie untersuchten wir die Prozesse, die in der SRTT mit dem Erwerb impliziten und expliziten Sequenzwissens einhergehen. Implizites Sequenzwissen ging mit Veränderungen auf den Ebenen der Stimulusenkodierung, der Reaktionsausführung, und der Reaktionsauswahl einher. Veränderungen auf der Ebene der Reaktionswauswahl deuten darauf hin, dass implizitem Sequenzlernen Repräsentationen zugrunde liegen, die sowohl Merkmale der Stimuli als auch der Reaktionen beinhalten. Im Gegensatz dazu ging der Erwerb expliziten Sequenzwissens mit einer Antworttendenz in Richtung der regulären Reaktion einher; andere Auswirkungen von Sequenzlernen verschwanden im Verlauf der Aufgabe. Dieser Befund deutet darauf hin, dass der Erwerb expliziten Sequenzwissens mit einem Wechsel von stimulusbasierter zu planbasierter Handlungssteuerung einhergeht.

Chapter I

General Introduction

> *The lack of definitional and conceptual clarity in the study of the unconscious stems from the implicit or explicit association of certain tasks with characteristics of observers or rememberers such as intentionality or phenomenal awareness.* (Reingold & Merikle, 1990, p. 20)

Consider a jazz trumpet player practicing a new tune. While he may acquire a rich body of knowledge about the tune that he can readily verbalize (e.g., the sequence of harmonic changes), he will frequently be unable to verbally express knowledge about the coordinated movements of lips, hand, and tongue that are required for generating a specific pitch, timbre, or a melody. Rehearsing the tune with a band, he will adapt to many other sequential regularities of his band mates' playing: For instance, he may adapt to weaknesses in their playing, such as his drummer playing too laid back, or his bass guitar player always forgetting to play the last fermata. In many of these instances, he will likely remain unaware of his behavioral changes; moreover, he will sometimes not even recognize that he adapted to their playing. Capturing the intuition that learning may occur with our without awareness, theories of human learning distinguish between two types of learning: Explicit learning that is accompanied by awareness of its contents, and implicit learning that proceeds independently of awareness (Abrahamse et al., 2010; Shanks & St. John, 1994).

Implicit learning has been demonstrated using the serial reaction time task (SRTT): In its basic appearance, participants respond to stimuli presented at horizontally aligned locations on a computer screen by pressing keyboard keys assigned to each location. Unbeknownst to participants, locations of stimuli follow a regular sequence. With practice, participants tend to respond faster to regular compared to nonregular stimuli. Critically, participants are frequently not able to express any explicit knowledge of the sequence, which has been interpreted as evidence that learning proceeded in the absence of awareness (Cohen, Ivry, & Keele, 1990; Nissen & Bullemer, 1987).

Theories of Implicit and Explicit Sequence Learning

Based on the observation that sequence learning may proceed in the absence of awareness, multiple theories of implicit and explicit sequence learning have been proposed. These can be subdivided by their assumptions about the number of systems that are involved in sequence learning: while multiple-systems accounts assume that independent knowledge bases subserve implicit vs. explicit learning, the single-system account assumes that a unitary knowledge base subserves both implicit and explicit learning.

The multiple-systems view

Probably the most elaborate multiple-system account is the *dual-system* model by Keele, Ivry, Mayr, Hazeltine, and Heuer (2003), which assumes that there are two independent mechanisms involved in learning. A *unidimensional* system consists of a set of multiple *encapsulated* modules, where each is restricted to a single *dimension* of stimulus or response features. The authors already acknowledged that the term dimension is not well-defined by the theory, noting that the term dimension may or may not be used interchangeably with modality, even though modalities may be subdivided (e.g., visual stimuli into shape, location or color). Learning in this system is considered to have three more properties: (1) it occurs automatically and without attention, (2) because of the encapsulated nature of the acquired representations, it is restricted to implicit knowledge, and (3) the input to this system is uninterpreted information, and representations therefore contain only uninterpreted information. A second, *multidimensional* system forms representations of events from multiple dimensions; it therefore provides the ability to form much more complex representations. Only information that is attended (i.e., that is part of the relevant *task set*) and provides predictive value will be represented in this system – still, if these two criteria are met, learning occurs automatically. Representations in this system substantially differ from those of the unidimensional system: in order to allocate attention to an event (e.g., a stimulus location), it has to be categorized – the input to this system is therefore categorized information, and representations are also formed on this categorized input. Knowledge in this system may be implicit or explicit, and both unidimensional or multidimensional information may be learned.

The single-system view

A contrasting, single-system view was provided by Cleeremans and Jiménez (2002). This account is rooted in connectionist models of cognition (e.g., Cleeremans & McClelland, 1991; Thomas & McClelland, 2008) and applies their basic principles to sequence learning: According to this view, the cognitive system may be represented as a large set of processing modules, each consisting of many simple processing units; processing is considered to be highly parallel and distributed. Long-term knowledge is acquired by changes of the connectivity between and within processing modules – current results of processing are captured by *transient patterns of activation* within each module. Processing of each module is constantly and continuously influenced by processing in other modules, depending on the activity of the other modules and connectivity between modules – as a consequence, processing is graded and continuous. Information is exclusively represented by transient patterns of activation; these representations vary by quality of representation (including factors such as strength, stability in time, and distinctiveness) in a graded fashion. Learning occurs as a mandatory effect of information processing that involves high-quality representations. Awareness is correlated with quality of representation in an inverse U-shaped fashion: with increasing representational quality, representations are more likely to become available to awareness – with even higher representational quality, representations may again be less accessible to awareness – representations may eventually have become automatic. The function of awareness is to provide flexible control over behavior.

How explicit knowledge emerges

In recent years, Cleeremans and colleagues (Cleeremans, 2011; Cleeremans, Timmermans, & Pasquali, 2007; Pasquali, Timmermans, & Cleeremans, 2010; Timmermans, Schilbach, Pasquali, & Cleeremans, 2012) considerably extended this model by proposing a possible mechanism for the emergence of explicit knowledge. Building on the distinction between *access consciousness* and *phenomenal consciousness* (Block, 1995, 2007), Cleeremans and colleagues noted that representations as conceptualized by Cleeremans and Jiménez (2002) are only knowledge *in the system*, not *for the system* – it is therefore knowledge that may enable access consciousness, but not the phenomenal experience of knowing that one possesses knowledge of a specific content (i.e., phenomenal consciousness). Therefore, it is assumed that, in addition to a *first-order* network such as the learning mechanism proposed by Cleeremans and Jiménez (2002), a *second-order* network constantly observes the states of the first-order network, and learns about the states of the first-order network. By doing so, it develops *metarepresentations*, redescriptions of the first-order representations. These metarepresentations inform an agent about its inner states or reproduce the first-order network's output. While this mechanism is compatible with higher-order theories of consciousness (Dienes & Perner, 1999; Rosenthal, 1990), metarepresentations are still conceptualized in the same way as first-order representations, that is, transient patterns of activation – their very nature, therefore, is gradual.

Conversely, Frensch et al. (2003) proposed a multiple-systems account of the emergence of explicit sequence knowledge (see also, Esser & Haider, 2017; Haider & Frensch, 2005, 2009; Rünger & Frensch, 2008) that conceptualizes awareness as an all-or-none phenomenon: The *Unexpected-Event Hypothesis* posits that, in an incidental learning situation, implicit learning precedes explicit learning, and is mediated by strengthening of associations in encapsulated, highly specialized modules. Such implicit learning results in behavioral changes, such as faster responding or feelings of fluency. It is assumed that individuals are able to detect such changes; if such a change is not expected, an additional explicit learning mechanism (i.e., hypothesis testing to ascribe a cause to the *unexpected event*) is triggered. Such hypothesis testing might or might not result in the detection of the regularity; if detected, the individual has acquired explicit knowledge for the regularity in an all-or-none fashion.

The question whether explicit sequence knowledge emerges gradually or in an all-or-none fashion directly points to the question whether awareness should be conceptualized as a gradual or an all-or-none phenomenon. In recent years, it has been argued that the all-or-none conceptualization of awareness, combined with empirical dissociative methods (e.g., comparing a group of participants that could verbalize the sequence with a group of participants who could not) and dichotomous thresholds of consciousness, might have artifactually generated the distinction between implicit and explicit knowledge (Augusto, 2016, 2018). Consciousness should instead be conceptualized as a more complex phenomenon, such as being multidimensional (Bayne, Hohwy, & Owen, 2016), and graded (Fazekas & Overgaard, 2016). A third option has been proposed by Cleeremans and colleagues (Anzulewicz et al., 2015; Windey & Cleeremans, 2015), who propose that consciousness may be both graded and dichotomous.

The above-described theories reflect many of the debates that have emerged in the study of implicit and explicit sequence learning: the role of attention (e.g., Shanks et al., 2005; Rowland & Shanks, 2006), awareness, and representational code (for a review, see Abrahamse et al., 2010). However, both the single-system and the multiple-systems accounts assume that sequence learning may proceed both implicitly or explicitly. While explicit learning is unanimously a real phenomenon, some authors argue that strong empirical support for the very existence of implicit learning—learning that occurs both incidentally and without awareness—has yet to be shown (e.g., Shanks, 2005) Their critique is largely founded on methodological considerations that notoriously plague the study of implicit learning and, more generally, the study of unconscious cognition – precisely, that a dissociation of some performance measure such as the SRTT and some measure of awareness cannot unambiguously interpreted as evidence for learning without awareness.

The Logic of Dissociation and the Need for Models of Task Performance

Studies aimed to demonstrate learning without awareness frequently adopted a simple *logic of dissociation* (Augusto, 2018; Reingold & Merikle, 1990; Shanks & St. John, 1994): Performance benefits in the SRTT were compared to a subsequently administered measure of awareness. Learning without awareness is demonstrated if performance in the SRTT shows sensitivity to the sequence, while there is no sensitivity to the sequence in the measure of awareness. In implicit sequence learning, multiple measures of awareness have been proposed, including verbal reports (i.e., recall of the sequence), recognition, prediction, and generation tests, and dissociations with RT advantages in the SRTT have been reported.

However, the logic of dissociation has been criticized for several interrelated reasons (Reingold & Merikle, 1990; Shanks, 2005; Shanks & St. John, 1994): First, the measure of awareness typically has to be administered after performing the SRTT; that is, SRTT and measure of awareness differ with regard to immediacy (Newell & Shanks, 2014). Therefore, a dissociation between both measures might indicate—instead of some sequence knowledge being processed implicitly in the SRTT—that sequence knowledge was acquired explicitly, but simply forgotten until the measure of awareness was administered.

Second, it has to be assumed that the information that mediates performance gains in the SRTT is the same information that also drives performance in the measure of awareness (the *information criterion*). If, instead, some related information facilitates performance on the SRTT, but not the measure of awareness, this might artifactually indicate learning without awareness. For instance, consider that a participant performing an SRTT explicitly learns that there are no reversals (i.e., 2–3–2) or no trivially-obvious runs of locations (i.e., 1–2–3) in the sequence. This knowledge might enable the participant to perform better on the SRTT. Still, when asking the participant on a subsequent awareness test what transitions he recognized, he might still not be able to produce an answer that is correct by the criteria of the experimenter, and be classified as being unaware of the sequence. In other words, the participant explicitly learns some information I^* (the absence of some specific patterns in the sequence) that is correlated to information I (specific transitions that are present in the sequence); however, the measure of awareness is not constructed for capturing I^*.

Third, it has to be assumed that both measures are equally sensitive to the relevant information (the *sensitivity criterion*); it amounts to the sometimes estranging assertion that two different measures have exactly the same reliability. Fourth, it has to be assumed that the measure of awareness exclusively measures only explicit, but no implicit sequence knowledge (the *exclusivity criterion*). If, instead, implicit knowledge influenced the measure of awareness, this resulted in biased estimates of explicit knowledge. Fifth, it has to be assumed that the measure of awareness exhaustively measures all aspects of explicit knowledge that might influence SRTT performance (the *exhaustiveness criterion*). If the assumption is violated, explicit knowledge is underestimated.

Sixth, asserting that participants did not acquire explicit sequence knowledge is typically based on comparing a participant's or a group of participants' performance on the measure of awareness with some theoretically or empirically derived chance level by means of null-hypothesis significance testing. The conclusion relies on accepting the null hypothesis that there is no deviation from chance performance. At-chance performance might therefore be frequently explained by a lack of statistical power (power analyses for this type of test are rarely reported in the sequence learning literature), or inadequate assumptions about chance levels (see Chapter II). Bayesian statistics may provide a solution to this *null-sensitivity problem* of awareness measures (Dienes, 2015); for instance, Rouder and colleagues (Morey, Rouder, & Speckman, 2008; Rouder, Morey, Speckman, & Pratte, 2007) provided an ingenious solution in subliminal priming.

Still, given the multitude of assumptions, it might be more fruitful to accept the possibility that there is probably no measure that complies with all of them, and that all measures may be sensitive to both implicit and explicit processes (Timmermans & Cleeremans, 2015). Therefore, Reingold and Merikle (1988, see also, 1990) proposed an alternative to the logic of dissociation that makes less assumptions: Performance on an *indirect* measure, where instructions make no reference to the regularity, is compared with performance on a *direct* measure, where instructions explicitly refer to the regularity. By asserting that the direct measure is at least as sensitive to explicit knowledge as is the indirect measure, any knowledge that remains undetected by the direct test, but is detected by the indirect measure, can be assumed to be implicit. However, the authors already cautioned that both direct and indirect measure must be matched on any other aspect than instructions to ensure that this conclusion is warranted – otherwise, it would still be possible that both measures differ in difficulty, demands, or contexts (Destrebecqz, Franco, Bertels, & Gaillard, 2015). Attempts to use this approach in implicit sequence learning have been reported (e.g., Jiménez, Méndez, & Cleeremans, 1996; Reed & Johnson, 1994; Stadler, 1989), but these few studies have been criticized for methodological (Shanks & St. John, 1994) or statistical reasons (Shanks & Perruchet, 2002). As noted by Reingold (2004) and Destrebecqz et al. (2015), the utility of this approach is severely limited because direct and indirect measures have to be *fully comparable*, a condition that is difficult to satisfy.

A notable exception to comparing performance between tasks is to to analyze performance on one task with measurement models that allow for disentangling multiple processes that are involved in task performance. For instance, Buchner et al. (1997b) developed a multinomial processing tree (MPT) model of performance in the recognition task that was administered after having participants perform the SRTT. The model assumes that recognition judgments are driven by three underlying cognitive processes: recollection, familiarity, and guessing. Using this model, the authors found that performance in the recognition task is driven by all of these processes. Still, this finding did not provide unequivocal evidence for *implicit* learning: as Shanks and Johnstone (1999) convincingly argued, fluency-based judgment cannot be equated with an implicit learning mechanism at work, because fluency and its detection can be easily assumed to be an expression of explicit sequence knowledge.

Destrebecqz and Cleeremans (2001) applied the process-dissociation (PD) procedure (Jacoby, 1991) to the generation task. After training participants on an SRTT, participants were instructed that stimuli had followed a regular sequence. They were then asked, under inclusion instructions, to generate a sequence as similar as possible to the trained regularity; subsequently, under exclusion instructions, participants were asked to generate a sequence that is as *dissimilar* as possible. Assuming that only explicit knowledge is under a participant's control, a comparison of inclusion and exclusion performance then yields estimates of implicit and explicit sequence knowledge as expressed in this task: Differences between both conditions can only be attributed to knowledge that is under participants control, and hence explicit; if exclusion performance is above baseline, this is assumed to reflect knowledge that is not under participants' control, and hence implicit. Destrebecqz and Cleeremans (2001) found substantial implicit sequence learning using this approach. Importantly, and in contrast to the PD model by Buchner et al. (1997b), the model directly targets the implicit/explicit distinction and provides separate estimates for both processes. Using this PD approach, one does not rely on comparing performance in the SRTT with performance on a measure of awareness; instead, performance is compared only within the PD generation task. For these reasons, the PD approach has yet been the most evidential result in favor of implicit sequence learning in the SRTT. However, the PD comes with its own set of critical assumptions that might be violated; this issue will be extensively elaborated in Chapters II and III.

The process-purity problem

More generally, the problem that dissociations between *tasks* cannot be interpreted as dissociations between *processes* or *systems* is a direct consequence of the problem that performance on a task can typically not be equated with a specific process. Instead, on a given task, almost always multiple processes and cognitive operations are involved – in other words, tasks are not *process pure* (e.g., Jacoby, 1991; Yonelinas & Jacoby, 2012).

To address this issue, in many areas of (cognitive) psychology, it has been proposed to disentangle the contributions of these multiple processes to task performance by means of (cognitive) measurement models. Such models have been successfully applied to a wide range of tasks, including recognition memory, reasoning, perception, and attitude measurement (Erdfelder et al., 2009). For instance, multinomial processing tree models (Batchelder & Riefer, 1999; Erdfelder et al., 2009) are a class of measurement models for categorical data that are specifically tailored to a given task. For each process that is considered to be involved in task performance (e.g., in a recognition test: remembering a word, familiarity-based "old" responses, and guessing "old"), the probability of its occurrence is estimated. The process-dissociation model of Destrebecqz and Cleeremans (2001) and the model by Buchner et al. (1997b) can both be subsumed by this class of models. Other highly successful models are signal-detection models (e.g., Wixted & Mickes, 2010), and models of speeded choice such as the diffusion model (Ratcliff, 1978; Ratcliff, Smith, Brown, & McKoon, 2016; Wagenmakers, 2009) or the linear ballistic accumulator (e.g., Donkin, Averell, Brown, & Heathcote, 2009). Another interesting new approach are MPT models where not only the probability of each process, but also the duration of each process is estimated (Klauer & Kellen, 2018; see Heck & Erdfelder, 2016 for a similar approach).

An important property of such models is that they provide a means to link theory to data (Rouder, Morey, & Wagenmakers, 2016); different theoretical viewpoints can be instantiated as models, and models can be compared empirically. Moreover, assumptions about the data-generating process are explicitly made, and model adequacy can be evaluated.

The need for hierarchical models

When analyzing experimental data, it is common practice to aggregate over participants, items, or both, ignoring possible variability between participants and items. The unaccounted-for variability is known to inflate Type I error in ANOVAs (Clark, 1973), but Wickens and Keppel (1983) showed this is not a big concern in well-balanced designs using linear models. Linear models are appropriate for a variety of domains, but are typically inadequate for cognitive processes or tasks; instead, cognitive measurement models are typically nonlinear. Combining a nonlinear model with substantial variability between participants or items may lead to biased estimates and, hence, false conclusions regarding the studied processes. For instance, nonlinear measurement models have been applied to the analysis of response times that necessitate aggregation over participants, items, or both. From this line of research stems the seminal *power law of practice*, which states that response times may be described by a power function of practice trials. However, this is only true for response times aggregated over participants; individual learning curves are better approximated by an exponential function (Heathcote, Brown, & Mewhort, 2000). While, at first glance, it may seem overly meticulous to draw on the distinction whether performance follows an exponential or power law, such a distinction might have implications for distinguishing between different theoretical viewpoints (see, e.g. Haider & Frensch, 2002; Logan, 1988, 1990, 1992). A possible solution might be to estimate individual functions, but a large number of observations is necessary to test distributional properties (Rouder et al., 2005).

In many contexts, aggregation seems unavoidable: Consider, for instance, a recognition test that follows an SRTT to assess explicit sequence knowledge. The quantity of interest for such a task is the proportion of recognized items (i.e., transitions from the training sequence). For each participant-item combination, such a measure will produce only one independent replicate; measuring the same participant-item pair multiple times would only generate dependent replicates (which might cause more problems than it solves).

An increasingly popular alternative to aggregation is the use of hierarchical models. These models account for participant and item variability by estimating separate parameters for each participant-item combination. Assumptions about the distributions of these parameters provide additional constraint on individual-level parameters, resulting in higher accuracy for individual-level parameters (Efron & Morris, 1977; Katahira, 2016), and providing population-level parameters that are the natural target of inference.

The problem that aggregation across participants, items, or both might lead to biased parameter estimates (and inflated or deflated Type I error rates for tests of model fit) has been highlighted in MPT modeling. As a remedy, hierarchical extensions of the models have been proposed (Klauer, 2010; Rouder, Lu, Morey, Sun, & Speckman, 2008; Smith & Batchelder, 2010).

In the following studies, we adopt a Bayesian instead of a frequentist framework for analyzing hierarchical models. One reason is a pragmatic one – the hierarchical models used in the following chapters have thousands of parameters, and estimation within a frequentist framework might be intractable. Instead, within a Bayesian framework, model specification and parameter estimation are straightforward, and general-purpose software solutions exist for specifying and estimating these models (e.g. Stan, Carpenter et al., 2016; JAGS, Plummer, 2015). Moreover, inference on parameters is straightforward (e.g., Wagenmakers, Lodewyckx, Kuriyal, & Grasman, 2010). A second reason is a philosophical one – Bayesian probability provides a unified, logically consistent framework for defining, estimating, and comparing models, and to provide evidence for or against theoretical viewpoints (cf., Rouder, Morey, & Pratte, 2013). In the following section, we will briefly introduce frequentist and Bayesian conceptualizations of probability, and show how beliefs are updated within a Bayesian framework. Subsequently, we illustrate the utility of a hierarchical model. The presentation closely follows Rouder and Lu (2005) and Rouder et al. (2013).

Bayesian basics

Within a frequentist framework, probability is conceptualized by the frequency of an event that may happen or not. Consider a coin flip, where Y denotes the frequency of *heads* and N denotes the number of times the coin is flipped. The probability of heads is then defined by the proportion of times the coin showed heads when it was tossed arbitrarily often, i.e.,

$$p = \lim_{N \to \infty} \frac{Y}{N}$$

With this conceptualization, the coin has a true probability of showing heads when it is tossed; given a finite sample of coin tosses, it can be measured to an arbitrary precision by increasing N.

A model of such an experiment contains an unknown parameter, the probability p, that is estimated from data of N coin flips. The coin flip may be modeled as a binomial random variable

$$Y \sim \text{Binomial}(p, N)$$

Frequentists treat parameters as unknown fixed values that are estimated from data – however, because a frequentist probability is only defined in the infinite-sample case, it does not provide a strict rationale for estimation in the finite-sample case. Most frequently, least-squares (LS) or maximum-likelihood (ML) methods are used to estimate parameters, but these methods are not equivalent and may sometimes lead to different estimates; for instance, the maximum-likelihood estimator of the variance is defined by $\hat{\sigma}^2 = \sum(y_i - \bar{y})/N$, while the least-squares estimator is defined by $\hat{\sigma}^2 = \sum(y_i - \bar{y})/(N-1)$.

In contrast to the frequentist conceptualization of probability, the Bayesian perspective treats probabilities as statements of an observer's subjective belief in the occurrence of an event. Such a belief may be expressed in terms of a probability distribution, and rationally updated in the light of data. The rule used for updating beliefs is *Bayes' rule.*

Bayes' Rule. Bayes' rule may be written as

$$\pi(p|Y) = \frac{Pr(Y|p)}{Pr(Y)}\pi(p)$$

The quantity on the left-hand side $\pi(p|Y)$ is the *posterior distribution*, and represents the probability distribution of p after observing data Y. The mean of this distribution, the *posterior mean*, is typically a reasonable point estimator for p. The term $\pi(p)$ is the *prior distribution*, and can be viewed as the observer's a-priori beliefs about the true value of p.

The term $Pr(Y|p)$ serves two purposes in statistics: For any Y given a fixed parameter value p, it is the probability mass (or probability density) function; for any p given a fixed value of Y, it is the *likelihood function.* In this case, it is the likelihood function, and describes the likelihood of any parameter p given a fixed value of Y. For the binomial model, the likelihood function is given by

$$Pr(Y|p) = \binom{N}{Y} p^Y (1-p)^{N-Y}$$

The term $Pr(Y)$ is the probability distribution of the data conditional on the model. It is given by

$$Pr(Y) = \int_0^1 Pr(Y|p)\pi(p)dp$$

Fortunately, its evaluation is frequently not necessary; its purpose is to normalize $\pi(p|Y)$ so that $\int_0^1 \pi(p|Y)dp = 1$. Because the value of $Pr(Y)$ does not depend on p, it is convenient to write Bayes' theorem only by those terms that depend on the parameter of interest and to combine all other terms as a proportionality constant. It then takes the form

$$\pi(p|Y) \propto Pr(Y|p)\pi(p)$$

Rationally updating beliefs. To illustrate how beliefs can be updated using Bayes' rule, consider that an observer holds the a-priori belief that the true probability of the coin showing heads may be anything in the interval [0, 1] with equal probability. Such a belief may be expressed by specifying the prior distribution as a beta distribution

$$\pi(p) \sim \text{Beta}(a, b)$$

with $a = b = 1$.

Such a belief may now be rationally updated after having observed some coin flips. The left panel of Figure 1 shows the prior and the posterior distribution for having observed $Y = 7$ heads out of $N = 10$ coin tosses; the right panel shows the same prior, but the posterior for observing observing $Y = 21$ heads out of $N = 30$ coin tosses.

In both cases, the density of the prior distribution is flat in the interval [0, 1]; the posterior distribution is more narrow than the prior distribution, indicating that the observer has considerably learned from observing the data. It can also be seen that the posterior on the right panel is even more narrow than on the left panel, reflecting the fact that the rational observer learned more from observing 30 coin tosses than from observing 10 coin tosses.

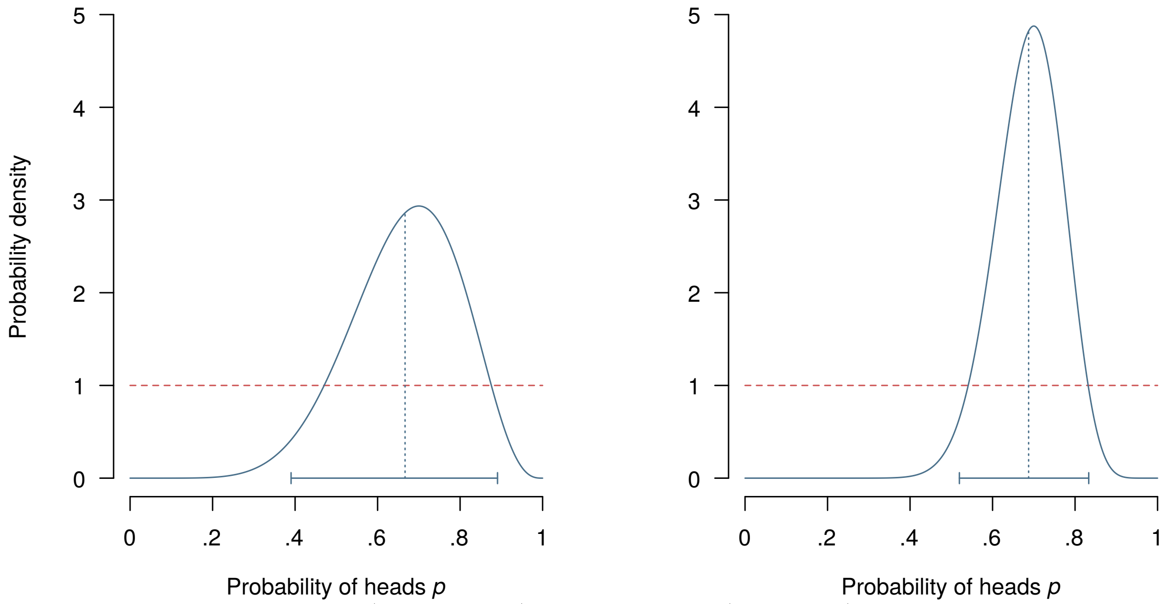

Figure 1. Left panel: Prior (dashed line) and posterior (solid line) density after observing $Y = 7$ for $N = 10$. Right panel: Prior and posterior density after observing $Y = 21$ for $N = 30$. Dotted vertical lines represent posterior means, intervals represent 95% credible intervals.

A hierarchical extension. To develop a hierarchical model, we now turn to a more psychologically interesting example. Consider a sample of participants performing a free-recall task after having learned a list of words. Each participant i has a true probability of recollecting a word p_i, which is typically estimated by calculating $\hat{p}_i = y_i/N_i$, where y_i denotes the number of recollected words and N_i denotes the number of words on the list. The following section briefly describes how a minimal hierarchical model for this task can be constructed, and illustrates how estimates from the hierarchical model are more accurate than individually derived estimates.

It is not convenient to use the beta distribution for developing a hierarchical model. Therefore, Rouder and Lu (2005), among many others, proposed to use a probit link together with a normally distributed prior density on transformed parameters z. The left panel of Figure 2 shows the probit transform. The probit Φ^{-1} is the inverse cumulative distribution function of the standard normal distribution

$$p = \Phi(z)$$
$$\Phi^{-1}(p) = z.$$

Some of the charms of this link function are obvious from the right panel of Figure 2: Using

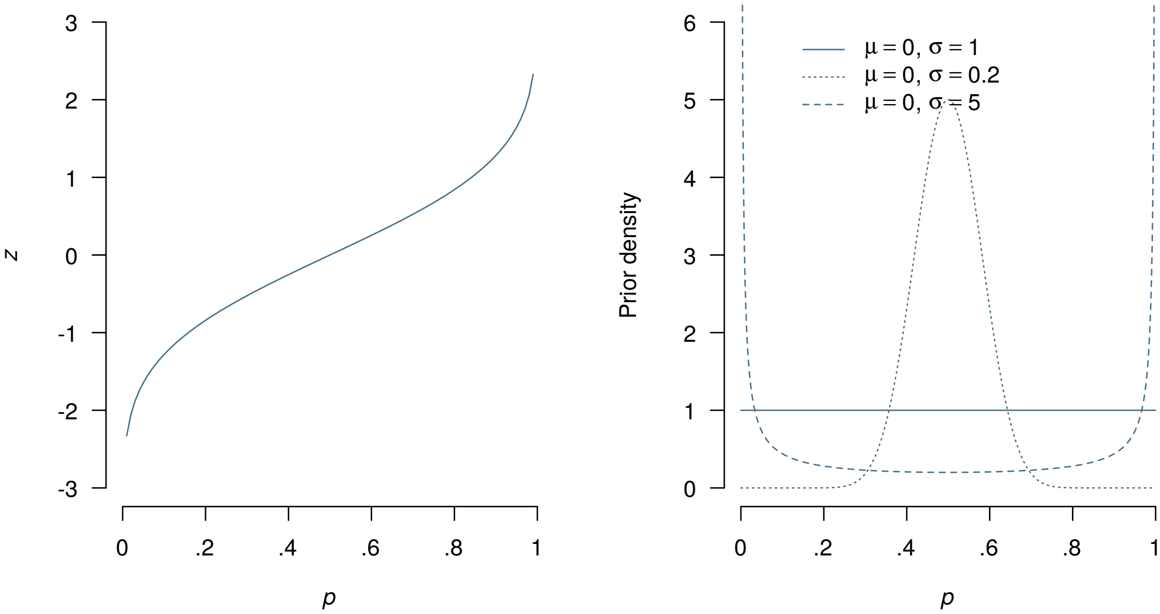

Figure 2. Left panel: Probit transform of probabilities p into z scores. Right panel: Resulting prior densities for three different distributions of z.

a normally distributed prior density results in different shapes of prior densities for the transformed values (i.e., probabilities p_i); a standard-normal distribution of z results in a flat prior for p. If the prior on z has $\sigma > 1$, the prior density is bimodal, if $\sigma < 1$, the prior density is unimodal.

The number of words participant i recollected y_i may then be modeled as binomial draws

$$y_i|z_i \overset{iid}{\sim} \text{Binomial}\left(N_i, \Phi(z_i)\right)$$

where N_i is the number of trials participant i responded to and z_i is a normally distributed random variable drawn from a *parent distribution.* The parent distribution may be modeled as

$$z_i|\mu, \sigma^2 \overset{iid}{\sim} \text{Normal}(\mu, \sigma^2)$$

Priors are needed for the parent distribution. A good choice for the prior on μ may be the normal distribution

$$\mu \sim \text{Normal}(\mu_0, \sigma_0^2)$$

A suitable choice for the variance is an inverse-gamma distribution

Table 1
True probabilities, simulated data, and parameter estimates for a probability of success

Participant i	p_i	y_i	$\hat{p}_{\text{individual}}$	$\hat{p}_{\text{hierarchical}}$
1	.766	27	.670	.669
2	.775	34	.840	.817
3	.514	13	.330	.369
4	.732	25	.621	.627
5	.657	30	.743	.733
6	.608	24	.598	.605
7	.695	29	.719	.711
8	.454	22	.549	.562
9	.663	27	.670	.669
10	.682	24	.598	.604
11	.583	21	.524	.541
12	.688	25	.622	.627
13	.774	32	.791	.775
14	.502	19	.475	.498
15	.585	31	.767	.753
16	.776	29	.718	.711
17	.791	37	.913	.880
18	.447	15	.379	.412
19	.590	19	.476	.498
20	.624	24	.597	.604
RMSE			.090	.079
MAD			.075	.063

Note. *RMSE* denotes the root mean-squared error, *MAD* denotes the mean average distance.

$$\sigma^2 \sim \text{Inverse Gamma}(a, b)$$

Values for μ_0, σ_0^2, a, and b are chosen before the analysis, and reflect the observer's a-priori beliefs about the plausible distribution of parameters.

To illustrate the usefulness of hierarchical estimation, consider the following example (adapted from Rouder & Lu, 2005). Table 1 shows data from 20 participants performing 40 trials. Each participant had a true probability of recollecting a word p_i (drawn from a uniform distribution with range [.4, .8]), and the number of recollected words y_i was randomly chosen from a binomial distribution. These data were analyzed both individually and with the hierarchical model just specified.

As can be seen from Table 1, both the mean average distance *MAD* and the root mean-squared error *RMSE* are smaller for the hierarchical-model estimates, $\hat{p}_{\text{hierarchical}}$, than for the individually derived estimates, $\hat{p}_{\text{individual}}$. The estimator $\hat{p}_{\text{hierarchical}}$ is superior because it shrinks extreme estimates toward the mean $\bar{p}$; on average, these estimates are closer to the true probabilities p_i (Efron & Morris, 1977; Rouder & Lu, 2005).

Overview of the Present Studies

The present studies embrace the notion that probably no task is process pure, and dissociations between tasks do not provide evidence for dissociable processes or systems. To study implicit and explicit sequence learning, it is therefore necessary to adopt a strategy that uses (hierarchical) measurement models of task performance to disentangle the contributions of implicit and explicit knowledge.

The PD approach as applied to sequence learning represents a simple measurement model of performance in the generation task, and studies utilizing this PD approach have yet been the most evidential result in favor of implicit sequence learning. However, the PD approach comes with its own set of critical assumptions that may be violated; these assumptions remained untested in the domain of (implicit) sequence learning. Therefore, in Chapters II and III, we tested the assumptions underlying the PD approach as applied to sequence learning. To foreshadow, we found substantial violations of the underlying assumptions, questioning the validity of the conclusion that sequence learning may proceed in the absence of awareness.

In Chapter IV, we propose a measurement model of SRTT performance that allows to disentangle the processes that are involved in the expression of sequence knowledge in the SRTT. While processes are rarely emphasized in the sequence learning literature (Schwarb & Schumacher, 2012), we believe that investigating the processes involved in the expression of learning in the SRTT may further our understanding of sequence learning in general. To foreshadow, our diffusion-model analyses are indicative of a common representational basis of stimulus and response features in implicit learning. Moreover, the acquisition of explicit sequence knowledge was accompanied by a qualitative shift in task processing.

Chapter II

Distorted Estimates of Implicit and Explicit Learning in Applications of the Process-Dissociation Procedure to the SRT Task

We investigated potential biases affecting the validity of the process-dissociation (PD) procedure when applied to sequence learning. Participants were or were not exposed to a serial reaction time task (SRTT) with two types of pseudorandom materials. Afterwards, participants worked on a free or cued generation task under inclusion and exclusion instructions. Results showed that preexperimental response tendencies, nonassociative learning of location frequencies, and the usage of cue locations introduced bias to PD estimates. These biases may lead to erroneous conclusions regarding the presence of implicit and explicit knowledge. Potential remedies for these problems are discussed.

Implicit learning refers to the ability to adapt to regularities inherent in the environment in the absence of conscious awareness about the ongoing learning process itself or about the outcome of what is learned. This ability is fundamental for human beings as it allows us to act optimally in stable environments with relatively little effort.

One of the most frequently utilized paradigms in the field of implicit learning is the serial reaction time task (SRTT) originating from Nissen and Bullemer (1987). In this standard SRTT, participants respond to locations on the screen which are mapped to spatially corresponding keys. Participants are instructed to press the appropriate response key whenever an asterisk occurs at a certain screen location. Unbeknownst to the participants, the locations of the asterisk follow a regular sequence. After several blocks of practice, the sequence is replaced by either a new but also regular sequence, or by a random sequence. In this transfer block, performance shows a decrement that disappears almost immediately when the original regularity is reintroduced, reflecting learning of the regularity. Importantly, participants are not able to explicate their acquired knowledge when asked to do so. Even with more sensitive tests including the recently introduced wagering task (Dienes & Seth, 2010; Haider, Eichler, & Lange, 2011; Persaud, McLeod, & Cowey, 2007) or the process-dissociation procedure (Destrebecqz & Cleeremans, 2001; Haider et al., 2011; Jacoby, 1991), explicit knowledge of the sequence is rare. This dissociation between performance and expressible knowledge is generally assumed to indicate implicit learning.

A central defining feature of explicit knowledge is that it is controllable, whereas implicit knowledge is thought not to be under conscious control. Destrebecqz and Cleeremans (2001) utilized this distinction and applied the process-dissociation (PD) procedure to measure sequence learning. In process dissociation, performance in two conditions of the same task is contrasted: An inclusion condition, in which explicit and implicit knowledge both produce the same response, and an exclusion condition in which explicit and implicit knowledge produce opposing responses (for applications of the PD to sequence learning with the recognition task see Buchner et al., 1997b; Buchner, Steffens, & Rothkegel, 1998).

In their application to sequence learning, Destrebecqz and Cleeremans (2001) used a generation task: After SRTT training, participants are asked to generate a sequence of responses that is either as similar as possible to the learned sequence (in the inclusion condition), or a sequence as dissimilar as possible (exclusion condition). To the degree that explicit knowledge is available, the proportion of generated responses that match the learned sequence should differ between inclusion and exclusion. To the degree that implicit knowledge is available, the proportion of matching responses in the exclusion condition should be greater than a chance baseline or control condition. In one group (i.e., RSI = 0 ms), participants were better than chance in their ability to reproduce the regularity in the sequence, even under exclusion instructions (i.e., performance under exclusion condition, E, was above a chance baseline B, $E > B$), a finding that was interpreted as reflecting sequence knowledge. However, performance under inclusion (I) and exclusion instructions was identical (i.e., $I = E$), a finding that is interpreted as indicating the absence of explicit knowledge and instead suggests that the sequence knowledge was fully implicit.[1]

The PD procedure is a simple and elegant way to disentangle controllable and uncontrollable processes, which has been widely used across a wide range of research questions (Yonelinas & Jacoby, 2012) and has the potential to address many open questions in the domain of implicit learning and memory. However, some authors have raised concerns suggesting that the assumptions underlying the PD may sometimes turn out to be overly simplified, which, in turn, would threaten the validity of PD results (e.g., Buchner, Erdfelder, & Vaterrodt-Plünnecke, 1995; Curran & Hintzman, 1995; Hirshman, 1998; Klauer, Dittrich, Scholtes, & Voss, 2015; Rouder et al., 2008).

For instance, it has been argued that the automatic process may be confounded with extra-experimental influences (Buchner et al., 1995; Rouder et al., 2008) or response tendencies (Stahl & Degner, 2007). In sequence learning, for example, participants may bring pre-experimental knowledge to the lab that interacts with task properties, or they may attempt to

[1] Note that the $E > B$ pattern could not be replicated in other studies (Norman, Price, & Duff, 2006; Wilkinson & Shanks, 2004) that, as a baseline, compared generation performance with that for a *control sequence* instead of an a-priori fixed value.

strategically generate non-regular or random sequences especially under exclusion conditions (e.g., Boyer, Destrebecqz, & Cleeremans, 2005). These participants may be influenced by their subjective theories about randomness, which may thereby affect generation performance, and perhaps distort PD estimates of implicit and explicit knowledge.

The type of task may furthermore affect the validity of PD estimates. In applications to the SRTT, discrepancies have been observed between free and cued versions of the generation task (Destrebecqz & Cleeremans, 2001; Wilkinson & Shanks, 2004): Whereas Destrebecqz and Cleeremans (2001), using free generation, have obtained evidence for implicit knowledge (i.e., $E > B$), Wilkinson and Shanks (2004), using a cued generation task, report the absence of implicit knowledge, with exclusion performance not distinguishable from baseline. However, as argued by Fu, Dienes, and Fu (2010), the failure to replicate the $E > B$ pattern could be due to a lower sensitivity of the cued generation task.

Furthermore, different types of random control conditions have been used that differ with regard to simple frequency information or other sequence-unrelated properties (Reed & Johnson, 1994; Stadler, 1992). For instance, participants who are trained on randomly selected permutations of a fixed-length sequence might learn that the entire set of response positions is used up before any position is repeated (negative recency), whereas participants confronted with fully random material during learning might learn that response positions are independent (Boyer et al., 2005). Although this type of knowledge is unspecific and sequence-unrelated, it may nevertheless affect performance in the generation task and distort conclusion about sequence knowledge if not properly controlled for.

Goal of the present study

To gauge these threats to the validity of the PD in sequence learning, and to examine the conditions under which PD can help to understand the processes underlying sequence learning, we investigated the possiblity of differential effects of material (permuted vs. random), task format (free vs. cued generation), and response tendencies on generation performance under inclusion and exclusion conditions.

Simple frequency information and sequence-unrelated properties. Generation performance may be contaminated by knowledge unrelated to the sequence such as the frequency of repetitions and reversals (Reed & Johnson, 1994). Repetitions and reversals are often not included into the regular sequence because they are considered especially salient and prone to lead to explicit knowledge (Stadler, 1992). If participants (explicitly or implicitly) pick up this absence of reversals in the learning phase, they may use this knowledge in the generation task, for instance, to generate fewer reversals across both conditions, or even to generate fewer reversals in the inclusion than the exclusion task. In other words, if the sequence does not contain reversals, then a generated sequence that reflects this property of few reversals is scored as above-chance performance. The finding that generation performance is above baseline will suggest the presence of implicit sequence knowledge. If the strength of this effect furthermore differs between the inclusion and exclusion instructions, it may artificially produce an $I > E$ finding and lead to erroneous conclusions suggesting the presence of explicit sequence knowledge.

Another type of sequence-unrelated information that may affect generation performance is the frequency of response locations: When some responses are more frequent than others in the learning phase – for instance, in mixed first-/second-order sequences – participants may pick up this information (explicitly or implicitly) and use it in the generation task (Reed & Johnson, 1994). More importantly, if the PD instruction – inclusion or exclusion – can affect the expression of this knowledge, generation performance may be differentially affected, results may be distorted or artifactual, and substantive conclusions might be erroneous.

The present study investigates the degree to which learning of such sequence-unspecific properties (reversals, zero-order frequencies) may distort estimates of explicit and implicit second-order sequence knowledge.

Free and cued generation tasks. In applications of the PD to sequence learning, two variants of the generation task have been used: free generation (Destrebecqz & Cleeremans, 2001) and cued generation (Wilkinson & Shanks, 2004). In the free generation task, participants are asked to generate a longer stretch of responses without interruption (e.g., 96 trials in Destrebecqz & Cleeremans, 2001). In cued generation, on the other hand, on each trial, a small sequence of stimuli and responses is given as a cue by the experimenter, after which the participant is asked to generate the response that would occur next in the sequence. The results obtained with both tasks have been found to diverge: Using free generation task, Destrebecqz and Cleeremans (2001) reported evidence for implicit knowledge (i.e., exclusion performance was above baseline). In contrast, Wilkinson and Shanks (2004) could not replicate this $E > B$ finding using a cued generation task. The type of generation task may explain this discrepancy if cued generation artificially lowers (or free generation artificially boosts) exclusion performance (but see Fu, Fu, & Dienes, 2008, for a reward-based explanation). Alternatively, the failure to find E>B in cued generation may be taken as evidence for the suboptimal sensitivity of the cued generation task. The present study compares free and cued generation tasks and their ability to detect learning of sequence-unspecific properties.

Response tendencies and subjective intuitions about randomness. In a control condition of an SRTT study, sequence information is typically not available (e.g., Haider et al., 2011). Participants are nevertheless asked to generate transitions reflecting some 'regular' sequence under inclusion conditions, and to avoid generating such a 'regular' sequence under exclusion conditions. In this situation, subjective notions of regularity and randomness are likely used to generate what participants regard as more regular sequences under inclusion conditions and less regular (or more random) sequences under exclusion conditions. If these subjective notions deviate systematically from the researcher's notion of randomness that is used to determine the chance baseline, they may bias the pattern of results and distort substantive conclusions. The present research addresses the question whether response tendencies may bias generation performance, whether response tendencies are acquired during the learning phase or reflect pre-experimental biases, and whether such a bias differentially affects inclusion versus exclusion conditions.

Interactions between these factors. The factors discussed above may occur in combination, with the effect of creating potentially more complex distorting influences on PD estimates. For instance, on top of subjective notions of randomness, sequence-unspecific properties of the training materials may be used to inform participants' response tendencies and perhaps lead to systematic differences between inclusion and exclusion. In addition, response tendencies informed by frequency information may interact with variants of the generation task. For instance, if participants have a tendency to avoid generating recent response locations (Boyer et al., 2005), the simple frequency information about reversals or high- versus low-frequency response location may differentially affect performance under different generation task variants (free or cued). The present study explores such potential interactive effects of simple frequency information, task format, and reponse tendencies.

The current study

In the present study, which is part of a project aiming at evaluating the validity of the PD procedure in sequence learning, we were interested in the ability of the PD procedure to signal the absence of (both implicit and explicit) sequence knowledge when such knowledge is in fact absent. In addition, we wanted to identify appropriate control conditions to serve as a baseline with which to compare experimental conditions.

We aimed at exploring effects of response tendencies, simple frequency information and unspecific properties of the material, and task properties as well as their interactions. Focus of the present study is their potential of distorting PD estimates of implicit and explicit learning. The present study realized three different 'control' conditions without any sequence information:

- a training phase with randomly drawn permutations of a second-order 8-item sequence,

- a training phase with randomly drawn response locations (from a uniform distribution),
- a no-learning condition in which participants merely familiarized themselves with the task.

Orthogonally, we implemented the two different versions of the generation task (free vs. cued).

We investigated whether inclusion and exclusion performance differed under the three control conditions and in the free versus cued task variants. If the PD model yields valid measures of implicit and explicit sequence knowledge, generation performance should be at chance level in all conditions. If, however, simple frequency properties of learning materials also affects generation performance, we would expect this to be reflected in the permuted condition when compared to the random condition. To the degree that response tendencies lead to deviations from chance level, this should be evident from the no-learning condition but reflected in all three conditions.

Method

Design

The study realized a 3 (*material*: permuted vs. random vs. no-learning) $\times$ 2 (*generation task*: free vs. cued generation) $\times$ 2 (*PD instruction*: inclusion vs. exclusion) $\times$ 2 (*block order*: inclusion first vs. exclusion first) design with repeated measures on the instruction factor.

Participants

190 participants (143 women, with a mean age of 25 years, range 18-59 years) completed the study (data from 8 participants could not be used due to a programming error, one participant failed to follow instructions during SRTT). Most were undergraduates from University of Cologne. Participants received either course credit or 7 Euro for their participation and were randomly assigned to experimental conditions.

Materials

We used two different types of pseudo-random material:

- A *permuted* sequence was randomly generated for each participant anew by drawing with replacement from the set of all possible permutations of a second-order 8-item sequence. For a given participant, each sequence was construed by randomly selecting two of the six locations to occur twice in the sequence. Per block, 18 permutations were drawn from the set.

- A *random* sequence was randomly generated for each participant anew by drawing with replacement from a uniform distribution of six response locations.

In a third *no-learning* condition, participants performed 20 responses drawn randomly in order to familiarize themselves with the task. They were instructed, prior to the generation task, to imagine they had just worked on a learning phase and to generate the sequence they may have encountered there.

In all conditions, the sequence adhered to the following (additional) restrictions: (1) there were no direct repetitions of response locations, and (2) there were no response location reversals (i.e., A-B-A). As a consequence of the random generation process, frequencies of response locations, first-order transitions, and second-order transitions varied across participants. To determine correct responses in the generation task, we computed an individual criterion for each participant based on their individual transition frequencies.

Procedure

The experiment consisted of three consecutive parts: First, participants worked on a SRTT (the *training phase*), followed by a *generation task* and a postexperimental interview. In the learning phase, participants in the permuted and random conditions performed a SRTT consisting of 6 blocks with 144 trials each (total of 864 responses). Participants in the no-learning condition performed only 20 random trials to familiarize themselves with the SRTT. SRTT and generation task were run on 17" CRT monitors with a screen resolution of 1024px × 768px. The viewing distance was approximately 60cm. A horizontal sequence of six white squares (56px × 56px) was presented on a grey screen. The distance between squares was 112px. Each screen location corresponded to a key on a QWERTZ keyboard (from left to right Y, X, C, B, N, M). Participants had to respond whenever a square's color changed from white to red by pressing the corresponding key. They were instructed to place the left ring-, middle- and index fingers on the keys Y, X and C. The right index-, middle- and ring fingers were to be placed on keys B, N and M. There was no time limit for responses in the learning phase (nor in the generation phase). A warning beep indicated an incorrect response. The response-stimulus interval (RSI) was 250 ms.

Following the SRTT phase, participants were told that stimulus locations during the SRTT followed some underlying sequential structure (participants who were not exposed to the SRTT phase were asked to imagine that they had experienced an SRTT in which locations followed some underlying sequential structure). The generation instructions were presented next, with order of inclusion vs. exclusion task counterbalanced across participants. Under inclusion (exclusion) instructions, participants were told to generate a sequence that is as similar (dissimilar) as possible to the sequence from the learning phase. For both instructions, participants were instructed to follow their intuition if they had no explicit knowledge about the underlying sequence. Direct repetitions were explicitly discouraged and were followed by a warning beep.

In the *free* generation task, after an initial sequence of three cue locations, participants freely generated 120 consecutive response locations. Participant were instructed that three squares would appear to which they had to respond; subsequently, question marks appeared at all locations and participants' key presses were reflected by the corresponding square's color changing to black. In the *cued* generation task, in each of 120 trials, 3 to 5 stimulus locations (taken from learning materials) were presented as cues, and participants had to respond with the corresponding key, in order to activate any sequence knowledge, after which the next response location had to be generated by the participant. Participants were instructed that a few squares would first appear to which they had to respond; subsequently, the question marks appeared and participants were asked to freely choose the trial's final response location.

Upon completing the computerized task, participants were asked to complete a debriefing questionnaire containing the following items: "Did you notice anything special during the task? Please note everything that comes to mind.", "One of the tasks mentioned a sequence in which the squares lit up during the first part of the study. In one of the experimental conditions, the squares did indeed follow a specific sequence. Do you think you were in this condition or not?", "How confident are you (in %)?", "Can you describe the sequence in detail?". Subsequently, participants were asked to indicate, for each of the six response keys, the next key in the sequence on a printed keyboard layout. Finally, participants were thanked and debriefed.

Results

For all analyses, a significance criterion of $\alpha = .05$ was used. If sphericity was violated in repeated-measures ANOVAs, Greenhouse-Geiser-corrected degrees of freedom and p values were used. The raw data (as well as additional supplemental materials) are available from https://github.com/methexp/pdl1.[2]

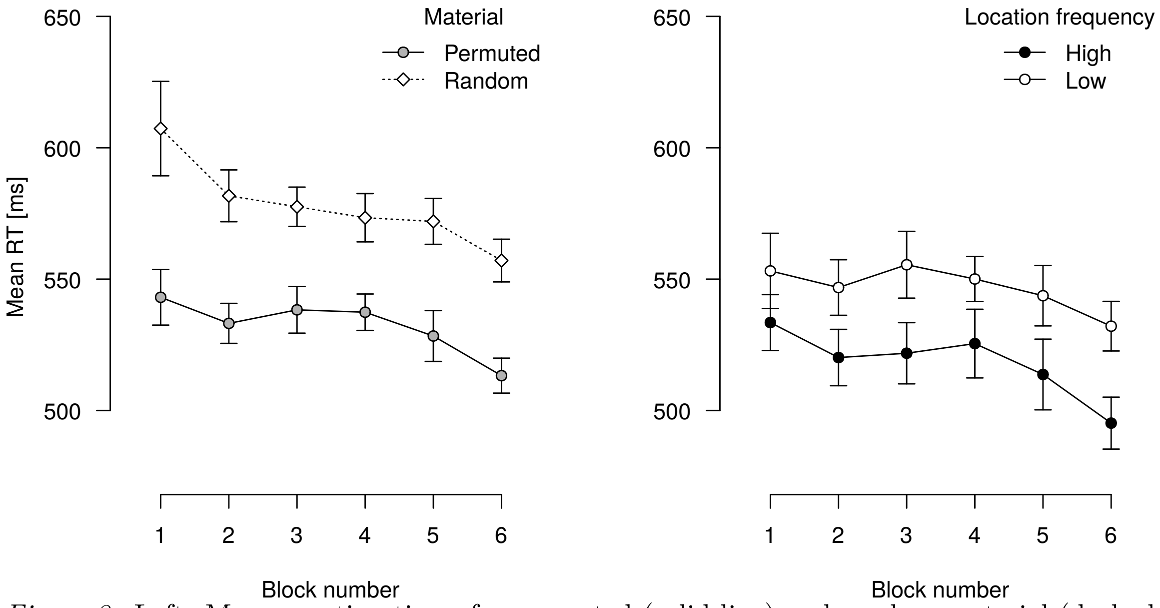

Figure 3. Left: Mean reaction times for permuted (solid line) and random material (dashed line). Right: Mean reaction times for permuted material, split by high-frequency (filled circles) vs. low-frequency locations (open circles). Error bars represent 95% within-subjects CIs.

SRTT Training phase

The mean RTs obtained over training blocks are shown in the left panel of Figure 3, separately for permuted and random material (in this and all following training-phase analyses, we excluded RTs of the first trial of each block and of trials that resulted in an error). Analysing RTs using a 2 (*material*: permuted vs. random) × 6 (*block number*) ANOVA revealed a significant main effect of *material*, $F(1, 117) = 5.62$, $MSE = 66,693.00$, $p = .019$, $\hat{\eta}^2_G = .042$, with slower responses for random material than for permuted material, and a significant main effect of *block number*, , reflecting practice effects. The factors *material* and *block number* did not interact, , indicating that any knowledge was acquired already during the first block.

For the group with permuted material, some locations were presented twice within each pseudo-sequence. The right panel of Figure 3 shows the mean RTs for these two types of stimuli. A 2 (*frequency*: high vs. low) × 6 (*block number*) ANOVA revealed a main effect of *frequency*, $F(1, 55) = 19.31$, $MSE = 7,105.13$, $p < .001$, $\hat{\eta}^2_G = .030$; high-frequency responses were faster than low-frequency responses. The main effect of *block number* was also significant, , but not the interaction, . These findings suggest that participants in the

[2]For all our analyses, we used R (Version 3.6.1; R Core Team, 2018) and the R-packages *ibdreg* (Version 0.2.5; Sinnwell & Schaid, 2013), *knitr* (Version 1.24; Xie, 2015), *papaja* (Version 0.1.0.9842; Aust & Barth, 2018), and *rmarkdown* (Version 1.15; Allaire et al., 2018).

permuted condition acquired some form of simple response-location frequency knowledge that benefitted their SRTT performance, and especially so for the more frequent responses.

Generation task

In the free generation condition, after an initial sequence of three cue trials, participants freely generated 120 consecutive response locations. In the cued generation condition, in each of 120 trials, three to five stimulus locations (taken from learning materials) were presented as cues and participants had to respond with the corresponding key. For each of the 120 trials, a response *triplet* consisted of the previous two locations as well as the location of the current response (in the cued generation condition, the response triplet consisted of the last two cue locations and the current response location; in the free generation condition, the response triplet consisted of the previous two response locations as well as the current response location; for the first two trials of each generation block, the locations of the corresponding cue trials were used). We calculated the proportion of triplets that were consistent with training sequences (see Appendix A).[3] Figure 4 depicts the pattern of correctly generated triplets as a function of generation task, material, and PD instruction.

Correctly generated second order transitions. The proportions of correctly generated triplets were analysed using a 3 (*material*: permuted vs. random vs. no-learning) $\times$ 2 (*generation task*: free vs. cued) $\times$ 2 (*PD instruction*: inclusion vs. exclusion) $\times$ 2 (*order*: inclusion first vs. exclusion first) ANOVA. It revealed a main effect of *PD instruction*, $F(1, 169) = 4.75$, $MSE = 0.00$, $p = .031$, $\hat{\eta}^2_G = .010$, more correct triplets were generated during inclusion than during exclusion blocks. This $I > E$ pattern suggests the presence of explicit knowledge, despite the absence of sequence information in the training material.

The ANOVA also revealed an interaction of *material* $\times$ *generation task*, $F(2, 169) = 5.88$, $MSE = 0.01$, $p = .003$, $\hat{\eta}^2_G = .042$. Analysing only free generation using a 3 (*material*: permuted vs. random vs. no-learning) $\times$ 2 (*PD instruction*: inclusion vs. exclusion) $\times$ 2 (*order*: inclusion first vs. exclusion first) ANOVA revealed a main effect of *material*, $F(2, 83) = 6.72$, $MSE = 0.01$, $p = .002$, $\hat{\eta}^2_G = .076$ (and again the main effect of *PD instruction*, $F(1, 83) = 4.13$, $MSE = 0.01$, $p = .045$, $\hat{\eta}^2_G = .024$). Tukey's HSDs revealed a significant difference between *permuted* and *random*, $p = .001$, the difference between *permuted* and *no-learning* group trended to be significant, $p = .083$, *random* and *no-learning* groups did not differ from each other, $p = .257$. Analysing only cued generation revealed no

[3] Because we only used pseudo-random sequences in this study, we defined that, given the last two locations, the location that was most frequently presented after these two locations during training as being the "correct" response. For the *no-learning* condition, correct responses were computed on the basis of sequences identical to those used in the *random* group.

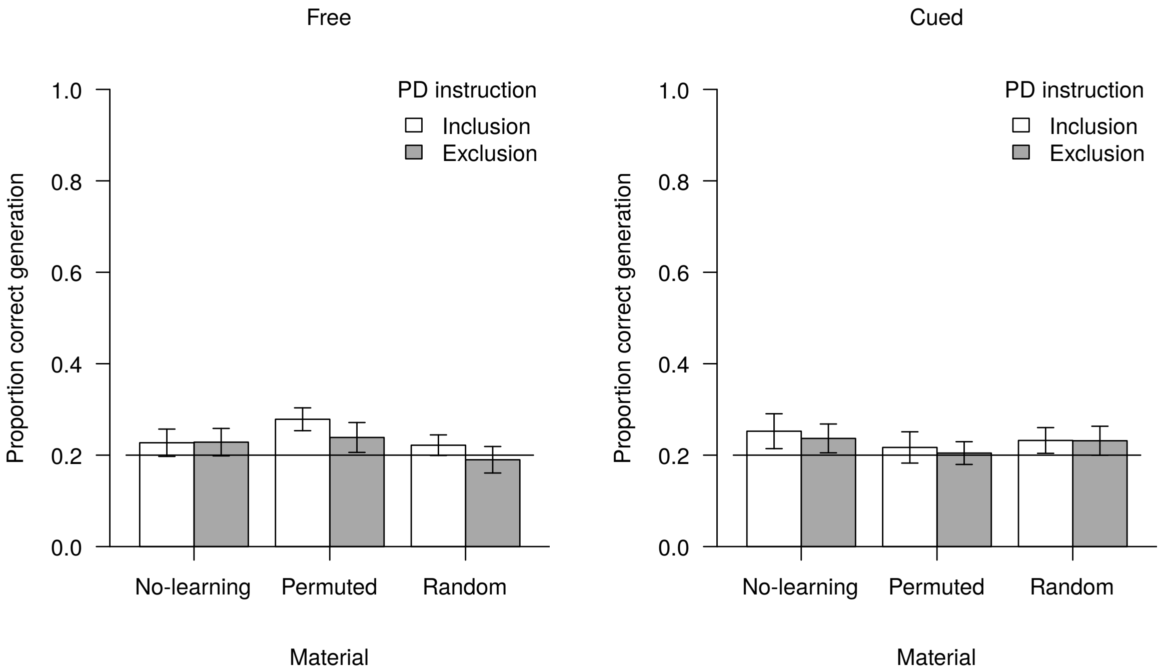

Figure 4. Proportions of correctly generated triplets during PD generation task (chance level = .2). Error bars represent 95% CIs.

effects, suggesting that the free generation task was more sensitive in picking up learning effects on generation performance.

To summarize, there was a main effect of PD instruction as well as an effect of material in the free-generation data; no effects were obtained in the cued-generation data. In the process-dissociation logic, the effect of PD instruction suggests the presence of explicit knowledge. In addition, the effect of material suggests the presence of implicit knowledge in the permuted condition. Next, we will apply two different PD analysis stategies to these data to investigate more formally whether, given the assumptions underlying the PD approach can be upheld, the pattern is indeed indicative of the presence of explicit and implicit knowledge.

Ordinal PD. We first applied the ordinal PD approach (Hirshman, 2004). It proposes a model-free analysis strategy of PD data by deriving four critical data patterns of inclusion and exclusion performance changes across two conditions that imply differences in automatic and controlled processes between the conditions. For instance, it can be concluded that both the automatic and the controlled processes are stronger in one condition if inclusion performance is greater in that condition but exclusion performance is comparable in both conditions. Also, it can be concluded that the automatic process is greater in one condition if both inclusion and exclusion performance are greater in that condition.

We screened the free-generation data for these data patterns by separately comparing inclusion and exclusion performance across conditions of the material factor. Inclusion performance was affected by material, $F(2, 83) = 5.42$, $MSE = 0.01$, $p = .006$, $\hat{\eta}^2_G = .116$: it was greater in the *permuted* than in both the *no-learning* group, $F(1, 55) = 6.70$, $MSE = 0.01$, $p = .012$, $\hat{\eta}^2_G = .109$, and the *random* group, $F(1, 53) = 11.65$, $MSE = 0.00$, $p = .001$, $\hat{\eta}^2_G = .180$; in the latter two, performance was comparable, $F(1, 58) = 0.08$, $MSE = 0.01$, $p = .777$, $\hat{\eta}^2_G = .001$. Exclusion performance was also affected by material, $F(2, 83) = 2.93$, $MSE = 0.01$, $p = .059$, $\hat{\eta}^2_G = .066$; specifically, it was greater in the *permuted* than the *random* condition, $F(1, 53) = 5.26$, $MSE = 0.01$, $p = .026$, $\hat{\eta}^2_G = .090$, but did not differ between the other two conditions (*no-learning* vs. *permuted*: $F(1, 55) = 0.24$, $MSE = 0.01$, $p = .628$, $\hat{\eta}^2_G = .004$; *no-learning* vs. *random*: $F(1, 58) = 3.53$, $MSE = 0.01$, $p = .065$, $\hat{\eta}^2_G = .057$).

Comparing the *permuted* and *no-learning* conditions, performance was greater in the *permuted* group under inclusion instructions but was identical in both groups under exclusion instructions. This data pattern implies increased levels of both the automatic and the controlled process in the permuted as compared to the no-learning condition.

Comparing the *permuted* and *random* conditions, performance was greater in the *permuted* group under both inclusion and exclusion instructions. This data pattern implies an increase in the automatic process in the permuted as compared to the random condition but no effect on controlled process.

These results suggest that participants acquire some form of implicit knowledge from permuted material that they can use to produce above-chance levels of correct triplets in the free generation task. They also suggest that the permuted and random materials allow participants to acquire some form of explicit knowledge which they can use to perform better under inclusion than under exclusion conditions.

PD equations. The same interpretation was suggested when the PD equations were used to obtain quantitative estimates of explicit and implicit knowledge. Based on the opposition logic, correct performance under inclusion instructions can arise due to controlled processes (C) or – should these fail – due to automatic processes (A), $p(correct|inclusion) = C + (1 - C) * A$. Under exclusion instructions, the controlled process leads to an incorrect response, and correct performance is based only on automatic processes in the absence of controlled processing, $p(correct|exclusion) = (1 - C) * A$. Automatic and controlled parameters could differ across experimental conditions, resulting in a total of 12 parameters (i.e., separate sets of A and C for two generation task variants, multiplied by three levels of the material factor). Parameters were estimated using the HMMTree software and the MPTinR package (Singmann & Kellen, 2013; Stahl & Klauer, 2007). Parameter estimates are given in Appendix B.

Hypotheses are tested by imposing restrictions on parameters and evaluating whether these restrictions significantly harm the model's goodness of fit, in which case the associated hypothesis cannot be upheld.[4]

To investigate the evidence for any form of learning, we analysed whether there was evidence for explicit knowledge by testing whether controlled parameters could be restricted to zero. This was not possible without substantially harming goodness of fit, $\Delta G^2_{(df=6)} = 28.19, p < .001$. This is consistent with the above findings of an effect of PD instruction, as well as the conclusions drawn from the ordinal PD approach about the presence of explicit knowledge. As above, this effect was strongest in the free generation condition $\Delta G^2_{(df=3)} = 23.55, p < .001$; it was clearly present in both the permuted ($\Delta G^2_{(df=1)} = 13.95, p < .001$) and the random material ($\Delta G^2_{(df=1)} = 9.60, p = .001$). In the cued-generation condition, there was a much weaker effect which was only marginally significant, $\Delta G^2_{(df=3)} = 4.64, p = .074$, and restricted to the permuted material ($\Delta G^2_{(df=1)} = 2.03, p = .077$) but absent for the random material ($\Delta G^2_{(df=1)} = 0.00, p > .999$).

Next we analysed whether there was evidence for implicit knowledge, first by testing whether a restriction of A parameters to an a-priori chance level of .2 can be maintained. This was clearly not the case, $\Delta G^2_{(df=6)} = 172.68, p < .001$. Even if only the no-learning condition was considered, the A parameters were greater than the a-priori chance level, $\Delta G^2_{(df=2)} = 74.26, p < .001$, suggesting that participants had acquired substantial implicit knowledge. The A parameters could also not be set equal across materials, $\Delta G^2_{(df=4)} = 53.14, p < .001$, confirming the above findings that the amount of acquired implicit knowledge differed across conditions. This was the case in the free-generation condition, $\Delta G^2_{(df=2)} = 35.92, p < .001$, as well as in the cued-generation condition, $\Delta G^2_{(df=2)} = 17.22, p < .001$. However, the patterns differed across conditions: In the free-generation task, implicit knowledge estimates were ordered *permuted* > *no-learning* > *random*; in the cued-generation task, the order was *no-learning* = *random* > *permuted*.

Taken together, when the PD equations were used to obtain estimates of explicit and implicit knowledge from the present data, results suggest that participants acquired both implicit and explicit knowledge in the control conditions. Whereas explicit knowledge was detected only in the free-generation task, implicit knowledge was detected in both tasks (but patterns differed between tasks).

[4]Restricting parameters to be equal across *generation task* (free vs. cued) harmed goodness of fit, $\Delta G^2_{(df=6)} = 68.14, p < .001$; therefore, the amount of controlled and/or automatic processes must be assumed to differ across tasks. Similarly, parameters could not be equated across materials (no-learning, permuted, random), $\Delta G^2_{(df=8)} = 81.59, p < .001$, implying effects of material.

Interim summary

In three different control conditions, we computed the proportion of generated responses that matched the learning materials to test for any effect of implicit or explicit knowledge acquired from the learning phase. We obtained converging evidence from three different approaches: (1) In ANOVAs, the proportions of correctly generated responses differed as a function of material (this effect was restricted to the free generation task), as well as of PD instruction. (2) The ordinal PD approach, when applied to the free generation data, yielded evidence for greater implicit knowledge in the permuted than in the random and no-learning groups. It also yielded evidence for explicit knowledge in the permuted condition (and, by implication, in the random condition). (3) The PD equations yielded estimates of the controlled process, reflecting explicit knowledge, that were significantly different from zero for the free-generation data. They also yielded estimates of the automatic process, reflecting implicit knowledge, that were above chance levels for 4 out of 6 experimental conditions.

Taken together, despite the fact that the material in the learning phase did not contain any sequence information, the PD approach using generation tasks yielded 'evidence' for both implicit and explicit knowledge. In the following sections, we will try to account for these findings in terms of sequence-unrelated frequency properties, response tendencies, and cueing artifacts.

Sequence-unrelated frequency properties

First, we analysed the proportion of reversals as well as the proportion of high- versus low-frequency locations (as manipulated in the permuted condition) that participants generated in both versions of the task.

Reversals. Figure 5 shows participants' proportions of reversals (e.g., 1-3-1) generated during inclusion and exclusion (where triplets that contained a repetition were excluded from analyses). By chance, a reversal would be generated in 1 out of 5 cases.

We conducted a 3 (*material*: permuted, random, no-learning) × 2 (*generation task*: free vs. cued) × 2 (*PD instruction*: inclusion vs. exclusion) × 2 (*order*: inclusion first vs. exclusion first) ANOVA of the proportion of reversals. The ANOVA revealed a main effect of *PD instruction*, $F(1, 169) = 22.14$, $MSE = 0.01$, $p < .001$, $\hat{\eta}^2_G = .057$: More reversals were generated under exclusion than under inclusion instruction.

The ANOVA also revealed a main effect of *material* $F(2, 169) = 3.75$, $MSE = 0.01$, $p = .025$, $\hat{\eta}^2_G = .024$, that was qualified by an interaction of *material* and *generation task*, : The effect of *material* was restricted to the free-generation condition, $F(2, 83) = 4.84$, $MSE = 0.01$,

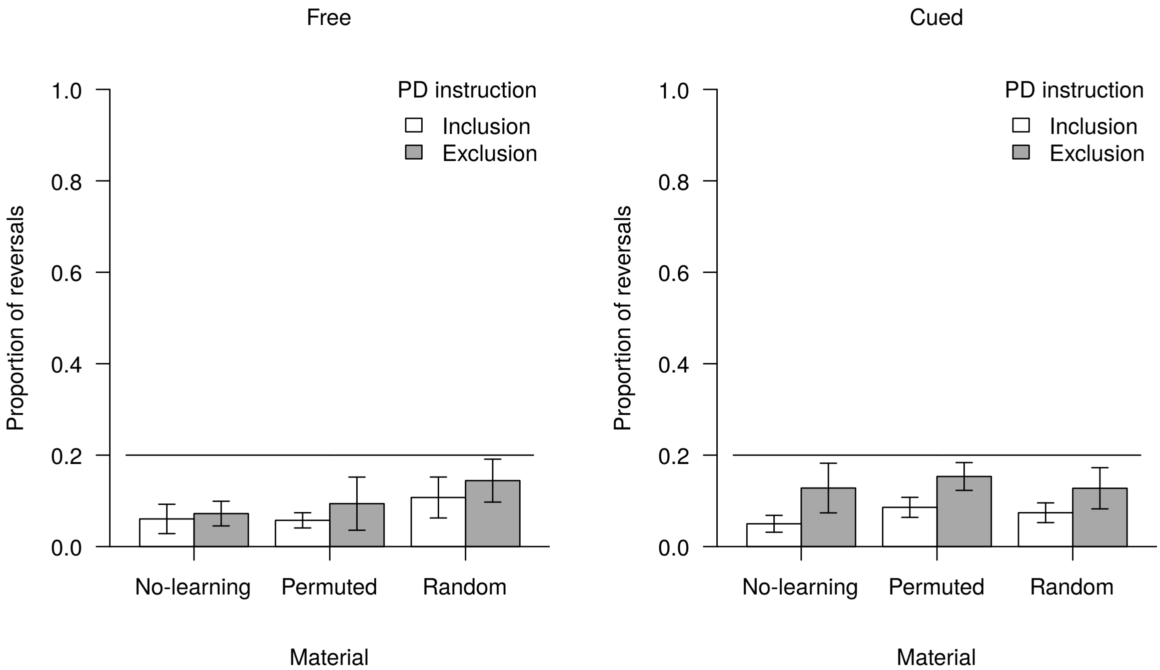

Figure 5. Proportions of reversals produced in generation task (chance level = .2). Error bars represent 95% CIs.

$p = .010$, $\hat{\eta}^2_G = .063$; it was not found in the cued-generation condition, $F(2, 86) = 1.70$, $MSE = 0.01$, $p = .189$, $\hat{\eta}^2_G = .019$. We further analysed the effect of material in the free-generation condition using Tukey's HSDs: The *random* group generated more reversals than the other groups (*no-learning* group, $p = .012$, *permuted* group, $p = .053$), which did not differ from each other, $p = .897$.

Finally, and perhaps most importantly, as apparent from Figure 5 the proportion of reversals was below chance for all materials. This effect was most prominent in the *no-learning* group which had the smallest proportion of reversals. The overall below-chance generation proportions of reversals can therefore not be interpreted as an effect of training.

It is more likely that they reflect a response bias that participants bring into the lab (e.g., Boyer et al., 2005). If reversals represent a regular pattern, according to participants' subjective theory of randomness, they should tend to avoid generating such regularities when attempting to produce a random sequence. The finding that more reversals were generated under exclusion conditions than under inclusion conditions would then reflect participants' attempt to generate a non-random sequence under exclusion conditions that is most dissimilar to the random sequence from the training phase. This strategy may underlie the $I > E$ pattern of correct responses obtained above that would suggest the presence of explicit knowledge: In cases where reversals occurred at chance levels in the sequence, the tendency to avoid generating reversals would lead to an underestimation of (implicit and/or

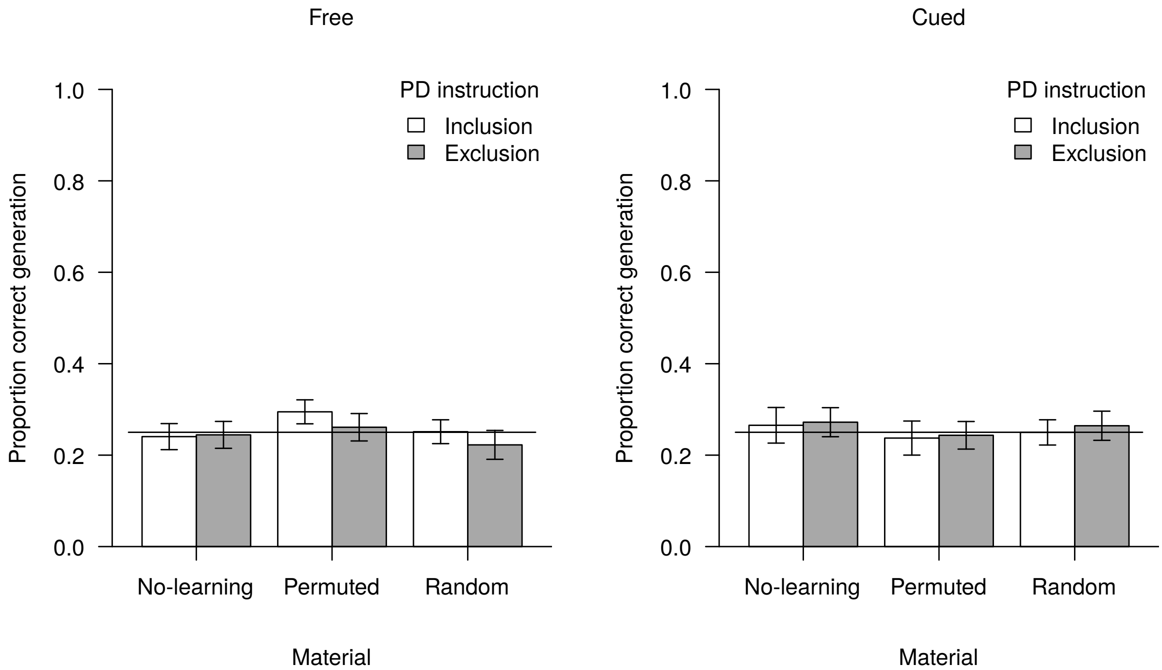

Figure 6. Proportions of correctly generated triplets during PD generation task after excluding reversals (chance level = .25). Error bars represent 95% CIs.

explicit) knowledge. In the present case, in which no reversals were encountered during learning, suppressing reversals leads to overestimated implicit knowledge. If this tendency is stronger under inclusion instructions than under exclusion instructions, as we have found here, it may erroneously suggest that participants have acquired explicit knowledge. Thus, the proportion of reversals that differed across PD instructions may be responsible for the above findings.

Correct generation performance when excluding reversals. Reversal generation proportions, whether reflective of extra-experimental response biases or not, may have distorted the above PD analyses and erroneously suggested the presence of implicit and/or explicit knowledge. We therefore repeated the above analyses after excluding reversals, to investigate whether the above patterns of $I > E$ and $E > B$ hold also for the remaining types of transitions. Figure 6 shows the proportion of correct generation responses after removal of reversals.

After excluding reversals from analyses, an ANOVA revealed only a significant interaction of *material* × *generation task*, $F(2, 169) = 4.39$, $MSE = 0.01$, $p = .014$, $\hat{\eta}^2_G = .031$, and a trend towards a *generation task* by *PD instruction* interaction, $F(1, 169) = 3.46$, $MSE = 0.01$, $p = .065$, $\hat{\eta}^2_G = .008$. Crucially, the main effect of *PD instruction* was no longer significant, $F(1, 169) = 0.42$, $MSE = 0.01$, $p = .519$, $\hat{\eta}^2_G = .001$.

As above, the effect of *material* was limited to the free-generation task, $F(2, 83) = 4.53$, $MSE = 0.01$, $p = .014$, $\hat{\eta}^2_G = .053$, it was absent from the cued-generation task, $F(2, 86) = 1.24$, $MSE = 0.01$, $p = .296$, $\hat{\eta}^2_G = .019$. Tukey's HSDs replicated the above finding that the *permuted* group generated more correct responses than the *random* group, $p = .001$, and tended to generate more correct responses than the *no-learning* group, $p = .083$. The *random* and *no-learning* groups did not differ from each other, $p = .257$.

Excluding reversals eliminated the main effect of PD instruction, an effect that is typically interpreted as evidence for explicit knowledge. This shows that conclusions drawn from applications of PD to the SRTT can be distorted by response-tendency artifacts: Participants appeared to possess explicit and controllable knowledge about the material but in fact merely avoided generating reversals, especially under inclusion instructions.

High-frequency locations. Next, we investigated effects of the response location frequency manipulation in the *permuted* condition. The above differences in correct generation performance between *permuted* and the other two groups might be explained by the fact that the *permuted* group was able to acquire knowledge about the unequally distributed location frequencies: The *permuted* group was trained on 8-response sequences constructed from six response locations, with two of the locations doubled. For a given participant, the two selected locations remained constant throughout the training phase and were therefore practiced more frequently than the other four response locations. This was reflected in faster responses at high-frequency compared to low-frequency response locations (see above, Figure 3).

In the absence of learning, one would expect one third (i.e., two out of six) of generated responses to be high-frequency locations. If participants in the permuted condition acquired and used knowledge to generate more high-frequency locations, this would increase their chances of producing a correct triplet. This is because approximately two thirds of the possible triplets end in a high-frequency location (in the *permuted* learning material, on average, a high-frequency location was the correct response for approximately 22 out of 30 transitions). Thus, the probability of generating a correct triplet may have been inflated by a tendency to generate an above-chance proportion of high-frequency locations.

The proportion with which high-frequency locations were generated is illustrated in Figure 7. We analyzed these proportions by way of a 2 (*generation task*: free vs. cued) × 2 (*PD instruction*: inclusion vs. exclusion) × 2 (*order*: inclusion first vs. exclusion first) ANOVA. Consistent with the above finding that participants responded faster to high-frequency locations during training, we found evidence that participants learned to match the response frequencies during generation: responses encountered more frequently during the SRTT were also generated more frequently. Yet, this tendency was affected by the type of *generation*

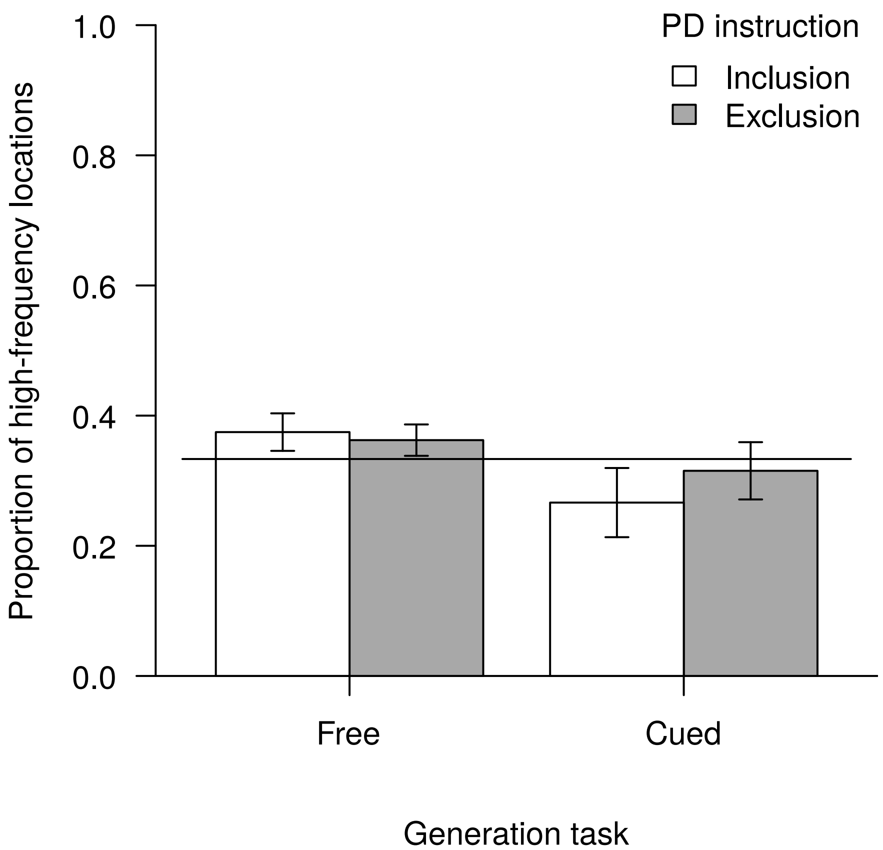

Figure 7. Proportion of high-frequency locations during generation in the permuted material condition (chance level = .33). Error bars represent 95% CIs.

task, $F(1, 52) = 13.97$, $MSE = 0.01$, $p < .001$, $\hat{\eta}^2_G = .144$: more high-frequency locations were generated in the free generation than in the cued generation test. Additionally, there was a trend towards an interaction of *generation task* and *PD instruction*, $F(1, 52) = 3.04$, $MSE = 0.01$, $p = .087$, $\hat{\eta}^2_G = .021$.

For free generation, the proportion of high-frequency locations generated was above the 1/3 chance level in both the inclusion blocks, $t(26) = 2.95$, $p = .007$, $d = 0.57$, and in the exclusion blocks, $t(26) = 2.47$, $p = .020$, $d = 0.48$. The proportion was *below* 1/3 for cued generation in the inclusion blocks, $t(28) = \text{-}2.58$, $p = .015$, $d = \text{-}0.48$, bot not in the exclusion blocks, $t(28) = \text{-}0.84$, $p = .407$, $d = \text{-}0.16$. The proportion of high-frequency locations in cued generation may be influenced by cueing artifacts as illustrated below.

In free generation, there was a tendency to generate high-frequency response *locations* at above-chance levels. As suggested above, this tendency might explain the above-chance generation of correct *triplets* in the permuted condition that remained after excluding reversals.

Effect of elevated high-frequency location generation rates on correct generation performance. To provide support for this interpretation, we investigated whether an above-chance generation rate as observed in the permuted condition is necessary and sufficient to account for this pattern. First, we corrected for the trend to generate more high-frequency locations by equally weighing correct generations for each response location.

For the *permuted* group we calculated the proportions of correctly generated triplets (excluding repetitions and reversals) separately for triplets that were completed by a high-frequency location and those completed by a low-frequency location. We then calculated the weighted mean correct performance, with one third as a weight for the two high-frequency locations and two thirds as the weight for low-frequency locations. Analyses of the equally-weighted free-generation data no longer revealed any significant effects. Crucially, the effect of material was no longer significant, $F(2, 83) = 1.44$, $MSE = 0.01$, $p = .243$, $\hat{\eta}^2_G = .017$. Perhaps trivially so, all significant findings disappeared after we eliminated the greater weight of the high-frequency locations in determining the proportion of correct triplets. This suggests that the higher proportion of high-frequency locations was necessary to produce the above effect.

To test whether the only slightly above-chance generation of high-frequency responses was sufficient to explain the above pattern of correctly generated triplets, we simulated data for each participant based only on their average proportions of high-frequency and low-frequency responses.[5]

For a participant's simulated dataset, we then determined the proportions of correctly generated triplets as above (i.e., based on the material in the learning phase, and after excluding reversals). We conducted the same analysis as above on the simulated data for the free-generation task. Importantly, the simulated data replicated the effect of material, $F(2, 87) = 49.64$, $MSE = 0.00$, $p < .001$, $\hat{\eta}^2_G = .447$: Tukey's HSDs revealed significant differences between *permuted* and the other two groups, both ps $< .001$, and no difference between *random* and *no-learning* group, $p = .986$.

The above results demonstrate that the slightly elevated generation rate of high-frequency responses was necessary and sufficient to produce the performance advantage in the permuted-material condition. They illustrate how sequence-unrelated frequency properties of the training material may affect generation performance, and thereby, PD estimates of underlying processes.

[5] All other information from their empirical distribution of response location frequencies as well as any order information that may be present in participants' generation responses was discarded: A high-frequency (low-frequency) response location was simulated with the probability given by the participant's rate of actually generating a high-frequency (low-frequency) response in the generation task; within the set of high-frequency (low-frequency) locations, each was generated with equal probability. Datasets were simulated by drawing responses from this distribution (with replacement, but with the constraints that no direct repetitions were allowed).

Effects of generation task

Whereas the proportion of reversals as well as high-frequency locations affected performance in free generation, similar influences were not observed on cued-generation data. This suggests that the cued generation task may be less sensitive to subtle effects of learning. The cued generation task has been criticized because the cues may contain information about the sequence that may affect generation performance and distort estimates of learning. Here we investigated potential effects of cues on generation responses that may occur even in control conditions and in the absence of informational influence.

As cues, participants were presented with brief segments of response locations taken from the learning phase. It is known that recent responses may be less likely to be generated (e.g., Boyer et al., 2005). This bias may affect generation performance selectively in the *permuted* condition. Because the permuted material contained high-frequency as well as low-frequency response locations, the same was true for the cues presented to participants in the permuted condition during the cued-generation task. As a consequence, the bias to avoid recent locations would apply more strongly to high-frequency locations. This could account for the above finding that high-frequency locations were generated at below-chance levels in the cued generation task, and for the suppressed levels of correct performance in the permuted condition in that task.

For the cued generation task, Figure 8 shows the proportion of response locations that were also presented as a cue on their respective trial. With three to five cues presented on each trial (and direct repetitions prohibited and excluded from analyses), chance level of generating a location that had just been presented equalled $3/5 = .6$.

A 3 (*material*: permuted, random, no-learning) $\times$ 2 (*PD instruction*: inclusion vs. exclusion) $\times$ 2 (*block order*: inclusion first vs. exclusion first) ANOVA revealed a main effect of *material*, $F(2, 86) = 3.41$, $MSE = 0.03$, $p = .038$, $\hat{\eta}^2_G = .042$. Tukey's HSDs revealed a significant difference between the *permuted* and *no-learning* groups, $p = .026$, all other ps $> .24$. The ANOVA also revealed a main effect of *PD instruction*, $F(1, 86) = 10.24$, $MSE = 0.03$, $p = .002$, $\hat{\eta}^2_G = .051$: more repetitions of cue locations were generated under exclusion conditions. This effect was qualified by an interaction *PD instruction* $\times$ *block order*, $F(1, 86) = 6.06$, $MSE = 0.03$, $p = .016$, $\hat{\eta}^2_G = .031$, which indicates that the effect of PD instruction differed between participants who worked under inclusion instructions in the first generation block and under exclusion instructions in the second block and those who received the PD instructions in the reverse order. To further explore this interaction, we analysed first and second blocks using separate ANOVAs, thus turning the *PD instruction* factor into a between-subjects factor. In the first block, only the main effect of *material* turned out to

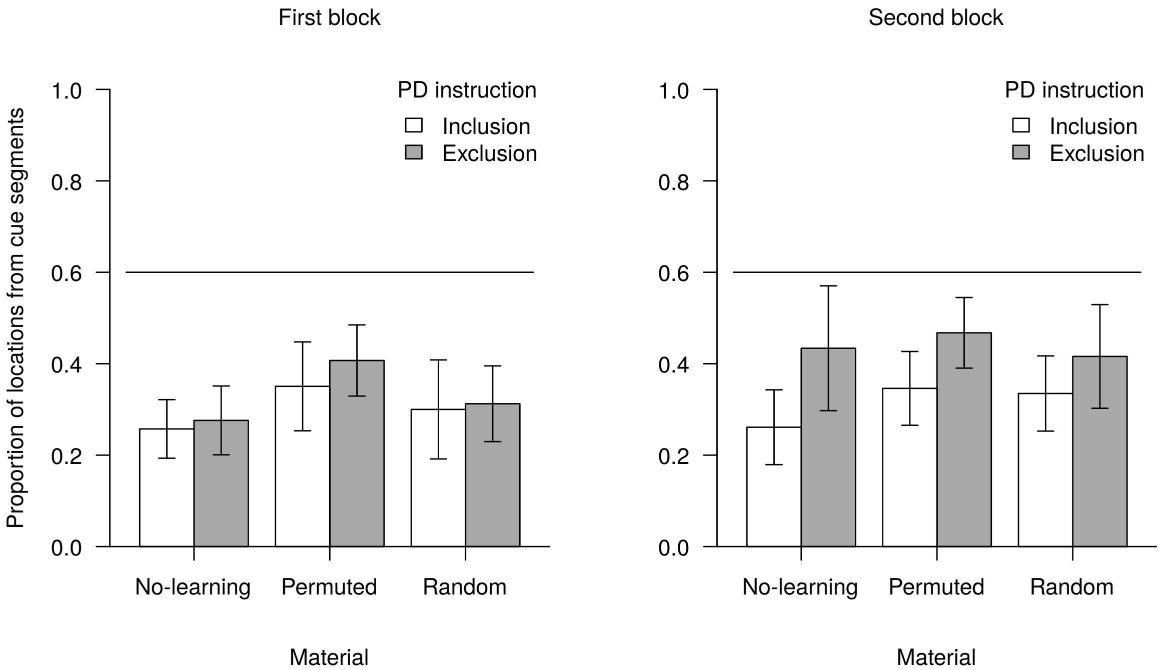

Figure 8. Proportions of locations generated in the cued generation task that were presented as a cue on their respective trial, separately for the first and the second generation block (chance level = .6). Error bars represent 95% CIs.

be significant, $F(2, 86) = 3.67$, $MSE = 0.03$, $p = .030$, $\hat{\eta}^2_G = .079$. In the second block, only the main effect of *PD instruction* was significant, $F(1, 86) = 10.91$, $MSE = 0.03$, $p = .001$, $\hat{\eta}^2_G = .113$.[6]

These results suggest three conclusions: First, the fact that the tendency to avoid generating cued locations was also present in the no-learning condition suggests that it was not acquired during the SRTT training but either reflects pre-experimental response tendencies, or, alternatively, it may reflect a tendency of the cued generation task to bias generation performance against repeating the cue locations.

Second, the cue-avoidance bias was influenced by the type of learning material in the first but not the second block. This suggests that, while the cued generation task is sensitive to effects of learning during the first block, this is no longer the case during the second block. This finding further supports the notion that the cued generation format may affect generation performance, and that over time this influence becomes stronger than the effects of learning.

[6]This data pattern remained even after excluding reversals: Analysing only the first block, only the main effect of *material* trended to be significant, $F(2, 86) = 3.67$, $MSE = 0.03$, $p = .030$, $\hat{\eta}^2_G = .079$. Analysing only the second block revealed a main effect of *PD instruction*, $F(1, 86) = 10.91$, $MSE = 0.03$, $p = .001$, $\hat{\eta}^2_G = .113$.

Third, cue-avoidance bias was stronger under inclusion than under exclusion instructions in the second block of the generation task. This suggests that participants can acquire response strategies during the cued-generation task that help them produce different outcomes under inclusion and exclusion instructions.

Whether or not this tendency leads to artificially elevated levels of correct responses largely depends on the sequence material that is used. Common sequences (e.g., Fu et al., 2010, 2008; Wilkinson & Shanks, 2004) contain only one reversal within a four-position sequence consisting of 12 triplets. Typically, only two locations are presented as cues on each trial to avoid informative influences, and direct repetitions are not allowed. In this case, the strategy to avoid generating a cue location would lead to a correct-performance rate of approximately 46% (i.e., a .5 chance to generate a correct response for 11 out of 12 triplets and a zero chance for the single reversal among the 12 triplets). This is rather high relative to a chance-level baseline of .33 (c.f., Fu et al., 2010).

To summarise, in the cued generation task, participants avoid generating a response that had been part of the cue sequence. This finding was independent of the type of training material, and it occurred even for the no-learning grooup, suggesting that it is a training-independent response tendency (i.e., a bias against generating locations present in the cue segment).

The findings further suggest that, in the permuted condition, this response strategy interacted with cue properties to influence generation performance. Here, cue properties vary as a function of the frequency manipulation in the permuted condition: Frequent locations are more likely to be included in the cue, and are therefore more likely to be subject to the avoidance bias; this can explain the reduction, in the *permuted* condition, of high-frequency responses in the cued as compared to the free generation task.

Finally, the cue-avoidance bias differed across PD instructions: cued responses were generated less frequently under inclusion than under exclusion instructions. Interestingly, this difference emerged only in the second generation block, suggesting that the cued generation task allowed participants to acquire a strategy of selectively generating more cue locations under exclusion instructions: If participants perceived the learning material as random, they will aim at generating a random sequence under inclusion instructions (i.e., generate a sequence similar to that in the learning phase); under exclusion instructions, they will aim at generating a sequence that does not conform to their subjective notion of randomness (i.e., generate a sequence that is as dissimilar as possible to that in the learning phase). When generating subjectively random sequences, participants typically produce more alternations and fewer repetitions than would be expected by chance (e.g., Boyer et al., 2005). By attempting to deviate from this subjectively experienced randomness they then generate more cue locations.

Taken together, the cued generation task was not only less sensitive to learning effects. It also appears to induce – or interact unfavorably with – a bias to avoid recently generated response locations. Both the below-chance generation of cue locations as well as its modulation by PD instruction may distort findings and conclusions regarding the presence or absence of explicit and implicit sequence knowledge.

Discussion

The present findings extend previous knowledge about the influence of pre-experimental response biases, simple frequency information, task format, and their interaction on performance in the generation task. First, it is known that zero-order frequency information may affect generation performance (Reed & Johnson, 1994); here, we show that this effect depends on the type of generation task. Second, it is known that, in cued generation, cues that carry information about the sequence may affect generation performance (Fu et al., 2010). We present first evidence that even in the absence of any sequence information, cues may affect generation performance by way of their simple frequency properties and in interaction with response tendencies. Third, it is known that participants may be biased to avoid generating recent response locations (Boyer et al., 2005). Here, we demonstrated that this bias may interact with properties of the learning material (zero-order frequencies, proportion of reversals) and task format, and that it can differentially affect inclusion and exclusion performance. These findings have important implications for the validity of the process-dissociation procedure as applied to the generation task.

Process-dissociation results

The present study realized three different control conditions, using different types of pseudo-random materials, in which participants could not learn any second-order regularity. Implicit and explicit knowledge was assessed with the PD procedure, using two different generation tasks. Despite the irregular learning materials, and independent of the type of analysis, the results of a comparison between performance in the permuted and random conditions in the free-generation task suggested the presence of implicit knowledge. Similarly, independent of the type of analysis, the results also suggested the presence of explicit knowledge under some conditions (i.e., due to the asymmetric generation of reversals under inclusion versus exclusion instructions). Finally, the results erroneously underestimated knowledge in some conditions: Although participants learned to distinguish between high- and low-frequency locations, they expressed this knowledge in the free-generation condition but did not express this knowledge in the cued-generation condition. These findings illustrate that unwanted influences can affect generation performance and may thereby artificially inflate or mask estimates of implicit and explicit knowledge.

Taken together, the findings suggest that there are several problems with the PD approach in its application to investigating the processes underlying sequence learning when effects of extra-experimental influences on performance cannot be excluded. In those cases, baseline performance may differ in inclusion and exclusion instructions and/or across experimental conditions. Under those circumstances, the ordinal PD approach is no longer valid (Hirshman, 2004).[7] In other words, given the possible influence of response tendencies and their interaction with properties of the material, we cannot draw conclusions about the relative contribution of implicit and explicit knowledge by simply comparing inclusion and exclusion performance, be it directly (e.g., $I > E$), across experimental conditions (e.g., $E_A > E_B$), or with a baseline (e.g., $E > B$). Instead, it is necessary to quantify the unwanted extraneous influence and separate it from the experimental effects of interest. This can be done by extending the basic PD design and model.

An extended PD model

In applications of the PD, exclusion performance is sometimes compared to an a-priori chance level baseline. The present study has shown that the proportion of correctly generated triplets may deviate from such an a-priori chance-level baseline, and may even differ across inclusion and exclusion conditions, in the absence of both implicit and explicit sequence knowledge. In other words, we found that response tendencies, alone or in interaction with properties of the learning material, may affect generation performance, may cause deviations from a-priori chance baselines, and may distort the PD model's estimates of implicit and explicit knowledge.

To accommodate this potential confound, first, it is necessary to use an empirical baseline. One could extend the PD design and model, as applied to the generation task, by adding a control condition along with separate response-tendency or nuisance parameters, as has been done by Buchner and colleagues for the recollection task (Buchner et al., 1998, 1997b). In the control condition, that is, in the absence of sequence knowledge, generation performance will then reflect sequence-unspecific properties of the material as well as response tendencies.

[7]The ordinal PD approach assumes that baseline performance does not differ across conditions or tasks (see Hirshman, 2004, p.559, section "Comparison of the Current Approach and the Process-Dissociation Procedure" and Footnote 8). Such baseline differences may arise from different guessing or response bias, and they invalidate the conclusions that can be drawn from the comparison of performance in conditions. This is because the performance differences between conditions or tasks that are used to draw conclusions about differences in automatic and/or controlled processes may instead reflect baseline differences due to guessing or response bias.

In the experimental condition, these same influences may affect generation performance over and above the effects of learning higher-order sequence information. The nuisance parameter can then be equated across the control and experimental conditions, capture the unwanted influences, and separate them from the effects of implicit and explicit knowledge. In this model, any differences in performance between the control and experimental conditions will be reflected in estimates of implicit and/or explicit knowledge. Second, the problem that response tendencies may differentially affect inclusion and exclusion performance can be solved by using as separate baselines the performance of a control group in the inclusion and exclusion conditions, and by introducing one nuisance parameter for each condition. Such an extended model allows for the possibility that response tendencies affect generation performance over and above explicit and implicit knowledge. Similar extensions for capturing extraneous influences have been proposed and demonstrated as necessary and useful in other domains to address similar response-tendency issues (e.g., pre-experimental familiarity, Buchner et al., 1995; right-hand bias; Stahl & Degner, 2007).

The extended model's equations deviate from the original equations (see introduction) by including nuisance parameters that account for baseline performance (along the lines suggested by Buchner et al., 1995, 1997b). Because baseline performance may be different under inclusion and exclusion instructions, two different parameters $R_{inclusion}$ and $R_{exclusion}$ are added to the model. The equations of the experimental conditions of the extended model are: $p(correct|inclusion, experimental) = C + (1-C)*A + (1-C)(1-A)*R_{inclusion}$, and $p(correct|exclusion, experimental) = (1-C)*A + (1-C)(1-A)*R_{exclusion}$. For the empirical control or baseline condition, it would be assumed that only extra-experimental or nuisance factors affect generation performance: $p(correct|inclusion, control) = R_{inclusion}$, and $p(correct|exclusion, control) = R_{exclusion}$.

To illustrate, we applied this model to the permuted and random conditions in the free generation task. Recall that both traditional approaches suggested the presence of explicit and implicit knowledge in the permuted group, as well as the presence of explicit knowledge in the random group. The extended model allows us to quantify the effect of training with permuted material over a training phase with random material on controlled and automatic processes. We used the random group as the empirical control condition; this is equivalent to the assumption that whatever affects performance in the random group can be subsumed as nuisance factors. Estimates of $R_{inclusion}$ and $R_{exclusion}$ therefore reflect generation performance in the random group, $R_{inclusion} = .251$ and $R_{exclusion} = .228$. Differences between the random and permuted condition are then reflected in the parameters for controlled and automatic processes. Results of the extended model suggests that training with permuted material affected the automatic process, $A = .047, \Delta G^2_{(df=1)} = 16.44, p < .001$, but did not affect the controlled process, $C = .001, \Delta G^2_{(df=1)} = 0.004, p = .95$. The effect on the automatic process – implicit knowledge – reflects the location frequency effect, and it is consistent with the finding that participants responded faster to high-frequency than to low-frequency locations during SRTT training. The extended model indicates the absence of an effect on explicit knowledge. The nuisance parameters capture the difference between inclusion and exclusion performance in both the random and permuted group that was interpreted as evidence for explicit knowledge in both the ordinal PD approach as well as in the original model equations.

Note that the levels of generation performance in the control condition can serve as valid baselines – and the controlled and automatic parameters can be valid estimates – only to the degree that the processes captured by the nuisance parameters are independent from explicit and implicit knowledge processes. In other words, the extended model requires additional assumptions, namely that the processes underlying the parameters C and A are independent from the nuisance processes (e.g., guessing, response bias) that are reflected in the new parameters $R_{inclusion}$ and $R_{exclusion}$. These assumptions are as of yet untested, and future research is needed to supply empirical evidence as to whether these new independence assumptions are met.

Free versus cued generation

In free generation, participants sometimes produce relatively sparse data, for instance, by repeatedly generating a small subset of all possible transitions, thereby generating missing data for the remaining transitions. With the cued generation task, researchers can control the cues and thereby ensure that comparable (or at least considerable) numbers of observations are obtained for each transition. This is an advantage especially when different types of transitions and their properties are to be compared in a within-subject approach.

However, the present findings support previous warnings that the cued generation task must be treated with caution because the choice of cues may influence generation performance (Fu et al., 2010; Jiménez & Vázquez, 2005). In addition to previous findings regarding informative cues, the present study provides an additional argument against the cued generation task: The tendency to avoid generating locations that were presented as cues may systematically bias generation performance, and may erroneously suggest the presence of both implicit and explicit knowledge.

In light of these problems, we currently do not see an empirical way of obtaining an equal number of observations for each transition in the generation task. This could imply that attempts of modeling individual transitions as a within-participant factor will face additional challenges. One way of avoiding artifacts resulting from a limited and biased selection of generated transitions is to use adequate weighting procedures. For instance, researchers could compute correct-generation rates for each of the cells in the transition matrix and estimate participants' mean proportion of correctly generated transitions as a weighted average.

Limitations and open questions

We briefly sketch open research questions posed by the additional independence assumptions, discuss the current development of hierarchical multinomial modeling techniques, and note the potential influence of additional unknown factors affecting performance in the generation task.

On the independence assumptions underlying the extended model. As mentioned above, the extended PD model assumes that response tendencies as modelled by $R_{inclusion}$ and $R_{exclusion}$ are independent of controlled and automatic processes. If this assumption is violated, parameter estimates are no longer valid measures of the underlying psychological processes. This possiblity deserves to be taken seriously; in other domains, empirical violations of independence between controlled and automatic processes have been obtained (e.g., Rouder et al., 2008). The open question is whether response tendencies can be affected by the same properties of the learning material that may also affect implicit and explicit learning. There are some findings in the literature that at least suggest such an interaction may well be possible. For instance, participants tended to show lag effects (reflecting negative recency) in a permuted condition but not in a random condition (Boyer et al., 2005). Here, the tendency to avoid generating recent locations interacted with properties of the learning material. Perhaps more critically, the present results suggest that in the permuted condition, the bias against generating cued locations was reduced whereas implicit knowledge was enhanced (i.e., participants learned about the frequency with which certain response locations occurred during training). If the response bias is reduced in a condition with greater implicit knowledge, then the response bias estimates (i.e., the R parameters) obtained from the control condition (where implicit knowledge is lower or entirely absent) will tend to overestimate the response bias in that condition. The present finding may reflect effects of other factors, but if such an interaction pattern between response tendency and implicit knowledge can be substantiated, it would constitute evidence against the independence assumption. In turn, this would require further elaborating and refining the extended PD paradigm and model.

Data aggregation and hierarchical modeling. In our application of the PD model, we aggregated the data across participants, as this has been the standard procedure in common applications of PD and our aim was to illustrate potential problems with such applications.

Our results demonstrate that the PD procedure, as commonly applied, yields erroneous conclusions supporting the presence of implicit and/or explicit knowledge, even in the absence of such knowledge. Note that such erroneous conclusions about the presence of an influence on parameters are less likely when, instead of aggregating, individual data are analyzed using hierarchical models. This is because data aggregation assumes parameter homogeneity across participants, which is likely to be violated in most cases. As a consequence of aggregation in the presence of heterogeneity, confidence intervals can be underestimated, and nominal statistical error levels can be violated by statistical tests.[8]

This problem does not arise if hierarchical modeling techniques are used (e.g., Rouder et al., 2008; Klauer, 2006, 2010) that account for the heterogeneity across participants (and items) in their estimates of the parameters' variability. Hierarchical models often draw a more realistic picture of the variability of parameter estimates. In all cases in which parameters must be assumed to vary across participants, they should be preferred to the traditional approach of data aggregation.[9]

Characteristics of response tendencies. In the present study, we obtained evidence that participants may use two types of response tendencies: First, they generated fewer reversals than expected by chance. Second, in the cued generation task, they avoided generating responses that were part of the cue segments. It is unclear whether these patterns reflect the same bias (e.g., based on subjective notions of randomness) or whether they reflect different response tendencies, and whether they reflect consciously applied strategies or rather more implicit trends. In addition, results obtained with the extended model suggested that, even in the absence of both implicit and explicit knowledge and after controlling for the identified response tendencies, there are other unknown factors that allow participants to perform better than chance, and even better under inclusion than under exclusion conditions. Future research should address these questions in order to better understand their potential for affecting generation performance.

[8] We applied the hierarchical Bayesian PD model proposed by Rouder et al. (2008) to the present data in order to account for person and item variability. The resulting parameter estimates are available in the supplemental material that can be obtained from https://github.com/methexp/pdl1. The results corroborated the findings obtained with the traditional analyses reported above (i.e., $C > 0$, $A > .2$, and the ordering of A estimates across conditions).

[9] The use and adoption of hierarchical modeling approaches is currently limited by the availability of general-purpose software (but see, e.g., Matzke, Dolan, Batchelder, & Wagenmakers, 2013; Stahl & Klauer, 2007).

Conclusion

Broadly speaking, the process-dissociation approach is to create two (or more) experimental conditions within a given experimental paradigm so that there is overlap with regard to most but not all of the psychological processes that are relevant for performance. Combined with measurement models such as the PD equations or more complex multinomial models, this general approach has been extremely valuable and has been fruitful in a wide variety of research areas and experimental paradigms (Erdfelder et al., 2009; Yonelinas & Jacoby, 2012). Whereas a simple and elegant design is generally preferable, the present findings suggest that the application of process-dissociation methodology to the generation task may require more differentiation of experimental design and measurement model.

Chapter III

Assumptions of the Process-Dissociation Procedure are Violated in Implicit Sequence Learning

In implicit sequence learning, a process-dissociation (PD) approach has been proposed to dissociate implicit and explicit learning processes. Applied to the popular generation task, participants perform two different task versions: inclusion instructions require generating the transitions that form the learned sequence; exclusion instructions require generating transitions other than those of the learned sequence. Whereas accurate performance under inclusion may be based on either implicit or explicit knowledge, avoiding to generate learned transitions requires controllable explicit sequence knowledge. The PD approach yields separate estimates of explicit and implicit knowledge that are derived from the same task; it therefore avoids many problems of previous measurement approaches. However, the PD approach rests on the critical assumption that the implicit and explicit processes are invariant across inclusion and exclusion conditions. We tested whether the invariance assumptions hold for the PD generation task. Across three studies using first-order as well as second-order regularities, invariance of the controlled process was found to be violated. In particular, despite extensive amounts of practice, explicit knowledge was not exhaustively expressed in the exclusion condition. We discuss the implications of these findings for the use of process-dissociation in assessing implicit knowledge.

Riding a bicycle is an easy task, but most of us will be hard-pressed to describe in detail the coordinated movements necessary for pedaling, keeping direction, and maintaining balance. Capturing this intuition, theories of human learning commonly distinguish two types of knowledge: Explicit learning that is accompanied by awareness of its contents, and implicit learning that operates independently of awareness (Shanks & St. John, 1994).

Such implicit learning has been demonstrated using the Serial Reaction Time Task (SRTT, Nissen & Bullemer, 1987), which has participants respond to stimuli presented at four horizontal screen locations by pressing the key that corresponds to the stimulus location. Unbeknownst to participants, the stimulus locations follow a regular sequence. With practice, participants learn to respond faster on trials with regular stimulus-location transitions than on irregular transitions. Critically, on a subsequent task, participants are often not able to express explicit knowledge about the sequential structure (Cohen et al., 1990; Nissen & Bullemer, 1987; Willingham, Nissen, & Bullemer, 1989).

There has been a long-lasting debate whether or not this effect is evidence for implicit learning, a question entwined with methodological considerations of how to properly measure and separate the contributions of supposedly implicit and/or explicit learning systems to this task (for a recent review, see Timmermans & Cleeremans, 2015). One of the most promising methods has been the process-dissociation (PD) approach as applied to the free generation task (Destrebecqz & Cleeremans, 2001); yet, its validity rests on a set of previously untested assumptions. The present study assesses two of the crucial assumptions on which this method is based.

Measuring implicit knowledge in the SRTT

In order to conclude that the learning effect in the SRTT (i.e., an RT advantage for regular transitions) is based on implicit knowledge, dissociations from subsequent assessments of explicit knowledge are typically sought. They depend on the assumptions that the explicit task is as *sensitive* to explicit sequence knowledge as the SRTT (the absence of an explicit effect may otherwise be due to lower reliability); and that it is also an *exhaustive and exclusive* measure of explicit knowledge, such that performance on the explicit task reflects all explicit but *no* implicit knowledge (Reingold & Merikle, 1990; Shanks & St. John, 1994). Multiple explicit-knowledge assessment tasks have been proposed, including verbal reports (i.e., recall of the sequence), recognition, prediction, and generation tests. Yet, while dissociations from RT advantages in the SRTT have been demonstrated in some studies, these tests have also been criticized for not meeting the above criteria, or the reported dissociations did not replicate (Shanks & Perruchet, 2002).

Contrary to the reported dissociations, studies utilizing recognition tests typically found substantial *associations* of the RT advantage in the SRTT with explicit knowledge (Perruchet & Amorim, 1992; Perruchet, Bigand, & Benoit-Gonin, 1997; Perruchet & Gallego, 1993). It has been argued that these associations were found because the subsequently used recognition task might not be exclusive to explicit but might also be driven by fluency-based processes. To test this alternative explanation, Buchner and colleagues (Buchner et al., 1997a, 1998) used the process-dissociation approach to disentangle (explicit) recollection and (implicit) fluency in the recognition task, finding that recognition is in fact driven by both processes. Still, Shanks and Johnstone (1999) argued that fluency-based recognition judgments cannot be equated with implicit knowledge, leading them to conclude that there was no conclusive evidence for implicit learning in the SRTT literature.

Given the interpretative problems of the recognition task, Destrebecqz and Cleeremans (2001) introduced the process-dissociation approach to the free-generation task, a measure that was considered to be the most sensitive to sequence learning (Perruchet & Amorim, 1992). Participants were instructed, after finishing the SRTT, to generate a sequence that is either similar (in the inclusion condition) or dissimilar (in the exclusion condition) to that encountered during the SRTT. If participants can generate a similar sequence under the inclusion instruction, they can be said to have acquired knowledge about the sequence; yet, this knowledge may reflect both implicit and explicit knowledge because both may be used to re-generate the learned sequence. However, only explicit knowledge is assumed to be under participants' control: When asked to generate a sequence that is dissimilar to the learned sequence – that is, to *exclude* their explicit knowledge – participants can avoid generating similar transitions only *if their sequence knowledge is explicit.* If, instead, their sequence knowledge is implicit, they would still generate a sequence *similar* to the learned sequence despite being instructed to do the opposite.

To selectively impair explicit knowledge, Destrebecqz and Cleeremans (2001) manipulated the (presence versus absence of a) response-stimulus interval (RSI), speculating that a certain minimal amount of preparation time would be necessary to acquire explicit knowledge during the SRTT. In both an RSI and a no-RSI condition, performance in the free-generation task was above a chance baseline, corroborating previous findings that the generation task is sensitive to sequence knowledge. Critically, in the no-RSI condition, performance under inclusion (I) was similar to performance under exclusion (E) instructions (i.e., $I = E$), suggesting that participants had no control over their sequence knowledge, and that the sequence knowledge in the no-RSI condition was fully implicit. (In addition, exclusion performance was above baseline, i.e., $E > B$, indicating that participants in the no-RSI condition were not able to withhold generating parts of the sequence they previously had implicitly learned.) Conversely, in the RSI condition, a robust inclusion-exclusion performance difference (i.e., $I > E$) indicated that participants were able to control their sequence knowledge, suggesting that this knowledge is explicit.

Assumptions underlying the PD approach

These conclusions about the presence versus absence of explicit knowledge, based on comparisons of inclusion and exclusion performance, depend on two assumptions: First, explicit knowledge must be assumed to be fully controllable (otherwise, the lack of an inclusion-exclusion difference cannot be interpreted as the absence of explicit knowledge but may instead reflect uncontrollable explicit knowledge). Put differently, conclusions drawn from the PD approach are limited to controllable explicit knowledge and do not extend to knowledge that may be explicit but not controllable (in the sense that it may be used to affect the similarity of the generated sequence with the learned sequence). This is unproblematic as long as the PD approach is used to investigate theories that hold controllability as a central tenet of explicit knowledge. Second, comparisons between inclusion and exclusion task performance are only meaningful if both tasks are indeed comparable measures of sequence knowledge.[10] In other words, the processes underlying free-generation performance are assumed to be *invariant to the inclusion versus exclusion instructions.* This assumption is critical for the validity of the PD approach, but it has so far not been tested directly.

The PD generation task has been used repeatedly to investigate sequence learning, but results were typically less clear-cut than those of the initial studies. First, most studies found $I > E$, suggesting the presence of at least some amount of controllable (explicit) knowledge even under no-RSI conditions (Wilkinson & Shanks, 2004). The debate focused on the evidence for residual implicit knowledge under exclusion instructions: Some studies replicated the $E > B$ finding of Destrebecqz and Cleeremans (2001), and concluded that SRTT learning is driven by implicit knowledge (e.g., Destrebecqz & Cleeremans, 2003; Fu et al., 2008; Haider et al., 2011); other studies found only $E = B$, a pattern interpreted as evidence that only explicit knowledge is acquired during the SRTT (e.g., Destrebecqz, 2004; Norman et al., 2006; Shanks et al., 2005; Wilkinson & Shanks, 2004). Wilkinson and Shanks (2004) failed to replicate the $E > B$ finding and speculated that this may come about because participants attempt to refrain from generating regular sequences under exclusion by resorting to various perseverative response strategies (i.e., by repeatedly generating regular-looking runs such as 1−2−3−4). If participants indeed use different strategies under the inclusion and exclusion instructions, this may violate the invariance assumption.

Moreover, in the presence of explicit knowledge, conclusions about the presence or absence of implicit knowledge, based on comparing exclusion performance with a baseline (i.e., $E > B$ vs. $E = B$), depend on additional assumptions regarding the interplay between both types of knowledge. If both types of knowledge may be involved, additional assumptions must be met if one aims at comparing inclusion and exclusion performance across two experimental conditions in order to draw conclusions about the relative contributions of explicit and implicit of knowledge; the *ordinal* PD approach formulates such a set of assumptions (Hirshman, 2004). Further assumptions are required for a *parametric* PD measurement model that can provide quantitative estimates of the underlying latent cognitive processes, for instance if the relative magnitude of the effect of a manipulation on explicit versus implicit knowledge is the quantity of interest. We next discuss critical assumptions underlying these two candidate methodological frameworks for the PD paradigm.

The ordinal PD approach. Analyzing their data by comparing inclusion and exclusion performance with a baseline, Destrebecqz and Cleeremans (2001) adopted an analysis

[10] For instance, if inclusion and exclusion performance differ in their sensitivity to implicit knowledge, this might lead to an artificial $I > E$ finding suggestive of the presence of explicit knowledge.

strategy that has been later formalized — with modifications — by Hirshman (2004) as the *ordinal-PD* approach. Instead of providing quantitative estimates of implicit and explicit knowledge, the ordinal-PD approach identified specific patterns of results that allow for ordinal comparisons between two experimental conditions (i.e., conclusions about increasing or decreasing amounts of explicit and/or implicit knowledge). In the light of the then-ongoing controversy about the PD method, this has been critically acclaimed as a way around the strong assumptions underlying the original (parametric) PD (Curran, 2001).

However, even this approach is based on assumptions that might be violated in a specific application: First, it is assumed that baseline performance is the same under both inclusion and exclusion instructions – an assumption that may be violated in sequence learning (Stahl, Barth, & Haider, 2015). Perhaps more critically, the second basic assumption of the ordinal-PD approach holds that both inclusion and exclusion performance are a monotonically increasing function of implicit knowledge; and that inclusion performance monotonically increases but *exclusion performance monotonically decreases as a function of explicit knowledge.* The exclusion strategies suggested by Wilkinson and Shanks (2004) would, however, imply that explicit knowledge does not necessarily inform exclusion performance: If participants adopted a perseverative response strategy (instead of engaging in an effortful search for their explicit knowledge, and attempting to implement this knowledge into a motor pattern consistent with the exclusion instructions), they would still be able to suppress their exclusion performance to baseline; however, they would not be able to suppress their exclusion performance *below* baseline (i.e. $E < B$).[11] The present Experiment 1 provides a first empirical test of this basic assumption of the ordinal-PD approach as applied to the free-generation task.[12]

The parametric PD model. The parametric PD model provides quantitative estimates of the underlying processes but relies on stronger assumptions. This section introduces the parametric PD model and its assumptions and then discusses its relation to the ordinal PD.

The PD model can be formalized as a set of equations describing inclusion (I) and exclusion (E) performance as a function of the probabilities of controlled process, C, reflecting explicit

[11] Moreover, if response strategies are informed by fragmentary knowledge about the regularity, such fragmentary knowledge might influence exclusion performance in any direction, depending on whether or not the chosen strategy is consistent with what the researcher considers to be successful exclusion. This effect might even outweigh performance changes due to the available explicit knowledge.

[12] Even though Destrebecqz and Cleeremans (2001) deviated from the ordinal PD as put forward by Hirshman (2004), their conclusions still rest on the assumptions of the ordinal PD specified here. Moreover, in order to interpret $I - E$ differences, they implicitly assume that the *same* strictly monotonic function links automatic and controlled processes with both inclusion and exclusion (whereas Hirshman allowed inclusion and exclusion performance to be linked by different functions). This additional assumption remains untested, yet.

knowledge, and the automatic process, A, reflecting implicit knowledge, as follows:

$$I = C + (1 - C) * A$$

and

$$E = (1 - C) * A$$

These equations reflect the notions that (1) regular transitions generated under inclusion can arise from either the controlled process (with probability C) or, given that it fails (with probability $1 - C$), from the automatic process A; and (2) regular transitions generated under exclusion are solely due to the automatic process in the absence of the influence of the controlled process, $(1 - C) * A$. Solving these equations for C and A (or using parameter estimation techniques for multinomial models) yields estimates of the contributions of the controlled and automatic process.

The validity of the PD method and model has been the target of debate since its introduction by Jacoby (1991; see, e.g., Buchner et al., 1995; Curran and Hintzman, 1995; Graf and Komatsu, 1994). This is because the PD approach is not a theory-free measurement tool but rests on a set of strong and possibly problematic assumptions. First and obviously, it assumes the existence of two qualitatively different—controlled and automatic—processes, and it aims to measure the magnitude of their respective contributions. It is, however, not well-suited for comparing single- and dual-process models: To illustrate, Ratcliff, Van Zandt, and McKoon (1995) found that data generated from a single-process model could produce a data pattern that, when analyzed using the PD approach, appears to support the existence – and differential contributions – of two qualitatively distinct processes. This implies that empirical dissociations between the controlled and automatic estimates do not necessarily imply the existence of two qualitatively different underlying processes.

Second, it is assumed that both processes operate independently; that is, on each trial, both the explicit and the implicit process attempt to produce a candidate response in parallel, without influencing each other.[13] In particular, the response proposed by the automatic process is assumed to be uninfluenced by whether the controlled process proposes the same or a different candidate response. Relatedly, the model assumes that independence holds across persons and items; when data are aggregated over (potentially heterogeneous) participants and items, a violation can lead to biases in parameter estimates. There has been considerable debate about the independence assumption in applications of the PD to episodic memory

[13] As an alternative to independence, a redundancy relation has been proposed such that the implicit process always operates, whereas the explicit process operates only in a subset of cases (Joordens & Merikle, 1993). An empirical comparison of the independence and redundancy assumptions has, however, supported independence (Joordens, Wilson, Spalek, & Paré, 2010).

paradigms (Curran & Hintzman, 1995, 1997; Hintzman & Curran, 1997; Jacoby & Shrout, 1997). Evidence suggests that aggregation independence may often be violated; hierarchical extension of the PD model have been proposed to address this problem (Rouder et al., 2008).

Third, and most important for the present study, it is assumed that both the controlled and automatic processes are *invariant* across the inclusion and exclusion instructions. This is reflected in the PD equations by the use of a single parameter C instead of separate parameters for inclusion and exclusion; in other words, the PD equations represent a simplified model that incorporates the invariance assumption $C = C_{Inclusion} = C_{Exclusion}$. Similarly, the PD equations include only a single parameter A, reflecting the simplifying assumption that the automatic process is invariant across inclusion and exclusion, $A = A_{Inclusion} = A_{Exclusion}$. If the PD instruction affects those cognitive processes, the PD equations do no longer yield valid estimates. Recently, the invariance assumption was indeed found to be violated for the controlled process across three different paradigms (Klauer et al., 2015).[14]

To summarize, Wilkinson and Shanks (2004) speculated that participants might use perseverative response strategies *especially in the exclusion condition* of the PD generation task; as a consequence, explicit knowledge would be less likely to affect exclusion performance. In terms of the parametric PD model, this would translate into an invariance violation of the controlled process with $C_I > C_E$. If the probability of controlled processes in exclusion C_E is negligible small, or if the invariance violation increases with increasing explicit knowledge, it cannot be assumed that explicit knowledge reliably decreases with explicit knowledge; thus, in terms of the ordinal-PD approach, an invariance violation of this kind would translate into a violation of the monotonicity assumption. In contrast, if neither is the case (e.g., if the invariance violation remains constant across different levels of explicit knowledge), the monotonicity assumption may hold despite an invariance violation. It is therefore important to test both the monotonicity and the invariance assumptions.

Overview of the present studies

The present study aimed at testing, in the free-generation task, the assumptions underlying both the ordinal- and the parametric-PD methods. For this purpose, it was necessary to extend the traditional PD design by manipulating both explicit knowledge (in Experiments 1-3) and implicit knowledge (in Experiments 2 & 3).

We manipulated *explicit* knowledge by explicitly informing participants, after the SRTT

[14]This assumption has not been tested earlier because the PD equations represent a saturated model: With two data points (i.e., the proportion of correct responses under inclusion and exclusion conditions), only two parameters (i.e., C and A) can be estimated. An extension of the design is needed to allow for estimating separate parameters $C_{Inclusion}$ and $C_{Exclusion}$, and/or $A_{Inclusion}$ and $A_{Exclusion}$.

training phase, about a subset of the regular transitions (e.g., 1 out of 6) of the sequence. By presenting information about the transitions *after training* we ensured that the manipulation did not affect the amount of sequence knowledge acquired during training (i.e., we made sure that participants did not use that information during the SRTT to strategically search for more regular transitions). We manipulated *implicit* knowledge by varying the amount of regularity present in the SRTT training sequence. For this purpose, we used materials with a mere probabilistic regularity; such materials are typically assumed to produce robust implicit knowledge but no explicit knowledge (Jiménez & Méndez, 1999; Jiménez et al., 1996).

In a test of the monotonicity assumption underlying the ordinal PD approach, Experiment 1 explored the speculation that explicit knowledge remains underutilized in exclusion. To foreshadow, we found that this was indeed the case and that the monotonicity assumption was violated. Results suggested that this is because invariance of the controlled process is violated.

In Experiments 2 and 3, we directly tested the invariance assumptions of the parametric PD model, closely following the methodology used by Klauer et al. (2015): We fit an extended process-dissociation model $\mathcal{M}_1$ that allowed for testing the invariance assumption of both the controlled and the automatic process. The model provided us with separate estimates for these processes for both inclusion and exclusion tasks; and we used the differences between these estimates to test the invariance assumption. This model relies on the auxiliary assumptions that each experimental manipulation selectively influenced only one of both processes; these assumptions are tested by goodness-of-fit tests proposed by Klauer (2010). Moreover, in order to justify the auxiliary assumptions, we specified a standard process-dissociation model $\mathcal{M}_2$ that does not enforce the auxiliary assumptions but enforces the invariance assumption; model comparison techniques (DIC; Spiegelhalter, Best, Carlin, & Van Der Linde, 2002) were then used to compare model $\mathcal{M}_1$ and model $\mathcal{M}_2$. If model $\mathcal{M}_1$ is favored over model $\mathcal{M}_2$, this can be taken as evidence in favor of our auxiliary assumptions over the invariance assumption. Finally, instead of aggregating data, we used hierarchical Bayesian extensions of all models (cf., Klauer, 2010; Rouder & Lu, 2005; Rouder et al., 2008).[15]

Experiment 1

A critical assumption of the ordinal-PD approach is that explicit knowledge monotonically increases inclusion performance *and* monotonically decreases exclusion performance. If (contrary to this assumption) explicit knowledge does not affect exclusion performance at

[15] This modeling approach controls for interindividual differences and circumvents aggregation artifacts.

all, the ordinal PD approach may technically still be applied; however, the results would be misleading if a difference in explicit (but not implicit) knowledge between two conditions led to a difference in inclusion but not in exclusion performance. In this case, the ordinal PD would suggest that the two conditions differ in explicit *and implicit* knowledge (Hirshman, 2004, Data Pattern I). For the ordinal PD approach to yield valid results, exclusion performance must decline with increasing explicit knowledge, and even fall *below baseline* when explicit knowledge is sufficiently strong so as to counter the influence of implicit knowledge (this would yield Hirshman's Data Pattern IV, which indicates an increase in explicit knowledge only). Therefore, a critical empirical test for the ordinal-PD approach is whether (and under which conditions) participants are able to use explicit knowledge to suppress generation below baseline levels under exclusion conditions. Our primary goal of Experiment 1 was to test this assumption; therefore, while keeping implicit sequence constant at moderate levels across conditions, we manipulated *explicit* knowledge by revealing parts of the sequence (i.e., explicit knowledge about 0, 1, or 2 transitions) to participants *after* finishing the SRTT.

A secondary goal of Experiment 1 was to manipulate the amount of *practice* participants had *with including and excluding their explicit knowledge*: If, in a sequence learning study, participants acquired explicit knowledge about a transition during the SRTT, they are likely to encounter the same transition again during the remainder of the SRTT several times: These additional exposures to the transition amount to an opportunity to practice *including* the explicit knowledge (e.g., intentionally implementing it into a motor pattern). This practice might be essential for the validity of the subsequent PD generation task. Therefore, one might wonder if the explicit sequence knowledge that is acquired during learning is comparable to explicit knowledge via instruction as implemented in our studies. To ensure that the effects of our explicit-knowledge manipulation were comparable with those of acquired knowledge, we amended the generation task by short *generation-practice blocks* involving inclusion/exclusion instructions that preceded the main inclusion/exclusion blocks. To explore the effects of practice, we manipulated whether a transition was revealed *prior to* or *after* these practice blocks: Some transitions were revealed *prior to* practice, and participants were instructed to implement a motor pattern including (or excluding) the revealed transition. Other transitions were revealed only *after* practice; for these transitions there was no opportunity to practice.

In five between-subjects conditions, we manipulated both the level of explicit sequence knowledge (i.e., the number of revealed transitions) and–nested within that factor–the amount of generation practice (i.e., whether a transition was revealed prior to or after practice blocks):

1. In the *Control* group, no explicit knowledge was revealed to participants.

2. In the *No-Practice* group, one transition was revealed immediately before the first main generation block, but *after* the practice blocks that preceded the first generation block. To avoid carry-over of practice effects from the first generation block, a different non-practiced transition was revealed after the second set of practice blocks and immediately preceding the second main generation block.

3. In the *Unspecific-Practice* group, one transition was revealed to participants *after* practice, immediately before each main generation block (as in the *No-Practice* group). In the third practice block before the exclusion task, participants were asked to inhibit a specific response location (i.e., they were asked *not* to use the 5^{th} location/N key).

4. In the *Practice* group, one transition was revealed to participants immediately *before* the practice blocks. Participants were encouraged to include (exclude) the revealed transition during generation-practice and in the main generation block.

5. In the *Transfer* group, information about two transitions was revealed; one of them was non-practiced (as in the No-Practice group), the other one practiced (as in the Practice group). The practiced transition was revealed before the first practice blocks. After these practice blocks, the second (non-practiced) transition was revealed immediately before the first generation block. The practiced transition was again named before participants worked on the practice blocks of the second generation phase. After these practice blocks, a second non-practiced transition was revealed immediately before the second generation block.

This design allowed us to assess overall generation performance for different levels of explicit knowledge. The monotonicity assumption states that with increasing levels of explicit knowledge, the proportion of regular transitions generated under inclusion should increase, while under exclusion it should decrease.

Beyond overall performance, we also analyzed data from specific transitions to test whether explicit knowledge may not be exhaustively expressed during the generation task (and, in particular, under exclusion instructions), and whether the level of its expression depends on generation practice. Our design allowed us to assess generation performance for three main transition types; (1) *non-revealed* transitions, (2) transitions that were revealed but remained *non-practiced*, and (3) transitions that were revealed and *practiced*.[16]

A comparison of *non-revealed* with (revealed but) *non-practiced* transitions allows us to

[16] In the second generation block of the No-Practice, Unspecific-Practice, and Transfer groups, a fourth transition type can be distinguished: Transitions that were revealed but non-practiced before the first generation block. Because participants included (or excluded) these transitions in the previous (i.e., the first) generation block, performance on these transitions should be more similar to practiced than to non-practiced transitions in the second block.

assess the degree to which participants can spontaneously (i.e., without practice) make use of their explicit knowledge in the generation task. Comparing *non-practiced* with *practiced* transitions allowed us to assess whether transition-specific generation practice could increase the use of explicit knowledge. We also compared whether performance for revealed but *non-practiced* transitions differs between the No-Practice and Transfer groups, as would be expected if the effect of specific practice transfers to non-practiced explicit knowledge. Finally, we explored whether unspecific inhibition practice affects performance for both revealed but *non-practiced* and/or *non-revealed* transitions.

In sum, we hypothesized that possessing explicit knowledge may not be sufficient for its expression in the generation task. Specifically, (a) explicit knowledge without practice (*No-Practice* group) may fail to lead to below-baseline exclusion performance, and (b) this may also hold for non-practiced transitions for participants who practiced another transition (*Transfer* group). If the exclusion task is not sensitive to manipulations of explicit knowledge, the ordinal PD would yield erroneous conclusions; it would also suggest that the invariance assumption for explicit knowledge of the parametric PD might also be violated. We had no clear hypothesis regarding the unspecific response-inhibition practice, but wanted to explore whether it would be as effective as transition-specific exclusion practice in improving the validity of the generation task as a measure of explicit knowledge.

Method

Design. The study realized a 5 (*Explicit-knowledge-and-Practice*: Control, No-Practice, Unspecific-Practice, Practice, Transfer) × 2 (*PD instruction*: inclusion vs. exclusion) × 2 (*Block order*: inclusion first vs. exclusion first) design with repeated measures on the *PD instruction* factor.

Participants. One hundred and forty-seven participants (113 women) aged between 17 and 55 years (M = 23.7 years) completed the study. Most were undergraduates from Heinrich-Heine-Universität Düsseldorf. Participants were randomly assigned to experimental conditions. They received either course credit or 3.50 Euro for their participation.[17]

Materials. A *probabilistic* sequence was generated from the first-order conditional sequence 2 – 6 – 5 – 3 – 4 – 1. With a probability of .6, a stimulus location was followed by the next location from this sequence; otherwise, another stimulus location was randomly chosen from a uniform distribution. There were no direct repetitions of response locations.

[17]The present research used procedures that are exempt from mandatory formal ethical approval under the ethical guidelines of the Deutsche Gesellschaft für Psychologie.

Procedure. The experiment consisted of three consecutive parts: Participants first worked on an SRTT (the *acquisition task*), followed by a *generation task* and, finally, a debriefing phase. In the acquisition task, all participants performed an SRTT consisting of eight blocks with 144 trials each (for a total of 1,152 responses). SRTT and generation task were run on 17" CRT monitors (with a screen resolution of 1,024 px × 768 px). The viewing distance was approximately 60 cm. A horizontal sequence of six white squares (56 px) was presented on a gray screen. The distance between squares was 112 px. Each screen location corresponded to a key on a QWERTZ keyboard (from left to right Y, X, C, B, N, M). Participants had to respond whenever a square's color changed from white to red by pressing the corresponding key. They were instructed to place the left ring-, middle- and index fingers on the keys Y, X and C. The right index-, middle- and ring fingers were to be placed on keys B, N and M. There was no time limit for responses in the learning phase (nor in the generation phase). A warning beep indicated an incorrect response. The response-stimulus interval (RSI) was 250 ms; there were no pauses within a single learning block.

Following the SRTT phase, participants were told that stimulus locations during the SRTT followed an underlying sequential structure (but were not informed about the exact sequence). They were then asked to try to generate a short sequence of six locations that followed this structure.

The generation task followed, consisting of two *main generation blocks* of 120 responses that were each preceded by three *generation-practice blocks* of twelve responses. Before entering practice blocks, one transition was revealed to participants in the *Practice* and *Transfer* groups. After practice blocks, another transition was revealed to participants in the *No-Practice*, *Unspecific-Practice*, and *Transfer* groups. Participants were told to memorize those transitions and to use their knowledge in all following tasks.

The main inclusion block was preceded by three practice blocks that were all performed under inclusion instructions. The main exclusion block was preceded by two practice blocks that were performed under inclusion instructions and a third practice block that was performed under exclusion instructions. The first two practice blocks (those that always involved inclusion instructions) were aimed at allowing participants to integrate their acquired sequence knowledge with just-revealed sequence information. The third block (that involved either inclusion or exclusion instructions, depending on the instructions of the subsequent main generation block) was aimed at allowing participants to familiarize with inclusion/exclusion instructions.[18]

In both main generation and generation-practice blocks, under inclusion (exclusion) instructions, participants were told to generate a sequence as similar (dissimilar) as possible to the sequence from the acquisition task. Participants were instructed to follow their intuition if they had no explicit knowledge about the underlying sequence. Participants who had received information about transitions were instructed to include (exclude) the revealed transitions. Question marks appeared at all locations and participants' key presses were reflected by the corresponding square's color changing to red. Direct repetitions were explicitly discouraged and were followed by a warning beep.

Upon completing the computerized task, participants were asked to complete a questionnaire containing the following items (translated from German): (1) "One of the tasks mentioned a sequence in which the squares lit up during the first part of the study. In one of the experimental conditions, the squares did indeed follow a specific sequence. Do you think you were in this condition or not?", (2) "How confident are you (in %)?", and (3) "Can you describe the sequence in detail?". Subsequently, participants were asked to indicate, for each of the six response keys, the next key in the sequence on a printed keyboard layout and to indicate how confident they were in this decision. Finally, participants were thanked and debriefed.

Data analysis. All analyses were performed using the R software[19] and Stan (Carpenter et al., 2016). For analyses of reaction times during the acquisition task, we excluded the first trial of each block as well as trials with errors, trials succeeding an error, reactions faster than 50 ms and those slower than 1,000 ms. For analyses of error rates during the acquisition task, we excluded the first trial of each block.

Generation task analyses were conducted with the first trial of each block as well as any response repetitions excluded. During the generation task, participants generated 120 keypresses. We coded these data as 119 first-order conditional transitions (e.g., a 4-key sequence $1-2-3-4$ was coded as the three transitions $1-2$, $2-3$, and $3-4$); we then computed the frequency of transitions that were consistent (i.e., part of) or inconsistent with (i.e., not part of) the training sequence. This scoring procedure follows the one used

[18]In Experiment 1, we held constant the number of generation-practice blocks involving inclusion/exclusion instructions. We based the number of presented practice blocks on our observations in Experiment 2 (that was conducted earlier): In Experiment 2, participants worked on practice blocks until they had consistently included/excluded a revealed transition. Prior to main inclusion blocks, participants in Experiment 2 needed $M = 2.98$, $Md = 2$ inclusion practice blocks. Prior to main exclusion blocks, participants in Experiment 2 needed $M = 1.26$, $Md = 1$ exclusion and $M = 1.52$, $Md = 1$ inclusion practice blocks. This suggests that the choice of 3 practice blocks — 3 under inclusion instructions (for the inclusion block), or 2 under inclusion and 1 under exclusion instructions (for the exclusion block) — should be sufficient for the majority of participants.

[19]We used R (Version 3.6.1; R Core Team, 2018) and the R-packages *afex* (Version 0.24.1; Singmann, Bolker, Westfall, & Aust, 2018), and *papaja* (Version 0.1.0.9842; Aust & Barth, 2018).

in the studies of Destrebecqz and Cleeremans (2001) and Wilkinson and Shanks (2004).[20] Response repetitions were excluded from analyses, as these were explicitly discouraged in the instructions. For repeated-measures ANOVAs, Greenhouse-Geisser-corrected degrees of freedom are reported.

Results

We first analyzed the performance data from the SRTT to determine whether sequence knowledge had been acquired during the task. Next, we analyzed generation task performance using an ordinal-PD approach (full descriptive statistics and additional model-based analyses are reported in Appendices A and C). Finally, to test our predictions regarding the different effects of practice, we analyzed generation performance for transitions about which explicit knowledge had been revealed.

Acquisition task. If participants acquired knowledge about the regularity underlying the sequence of key presses, we expect a performance advantage for regular over irregular transitions, reflected in reduced RT and/or error rate. If this advantage is due to learning, it is expected to increase over SRTT blocks.

Reaction times.

Figure 9 shows reaction times during acquisition. We conducted an 8 (*Block number*) $\times$ 2 (*FOC transition status*: regular vs. irregular) repeated-measures ANOVA. There was a main effect of *block number*, $F(4.05, 591.75) = 80.42$, $MSE = 1,658.05$, $p < .001$, $\hat{\eta}^2_G = .048$, with RT decreasing over blocks. There also was a main effect of *FOC transitions status*, $F(1, 146) = 716.67$, $MSE = 982.05$, $p < .001$, $\hat{\eta}^2_G = .062$, reflecting faster responses to regular than to irregular transitions. The interaction of *block number* and *FOC transition status* was also significant, $F(6.39, 933.34) = 45.89$, $MSE = 257.06$, $p < .001$, $\hat{\eta}^2_G = .007$, reflecting the finding that the RT advantage for regular transitions increased over blocks, which indicated successful sequence learning.

Error rates.

Figure 10 shows error rates during acquisition. The pattern of findings was similar to that obtained for RT. We conducted an 8 (*Block number*) $\times$ 2 (*FOC transition status*:

[20] This scoring procedure ignores sequential dependencies inherent in the free-generation data. For instance, the frequency with which a specific location is generated determines how often a transition starting from this location can be generated, and thereby, how well the knowledge available about this transition can be estimated: To illustrate, if the starting point of a transition is never generated, it is not possible to learn anything about the knowledge participants may have acquired about this transition. We believe this is not a serious threat to the present analysis because participants generated the locations at comparable rates. Still, other types of dependencies may yet turn out to be more problematic, and future research should consider modeling entire generation sequences instead of individual transitions.

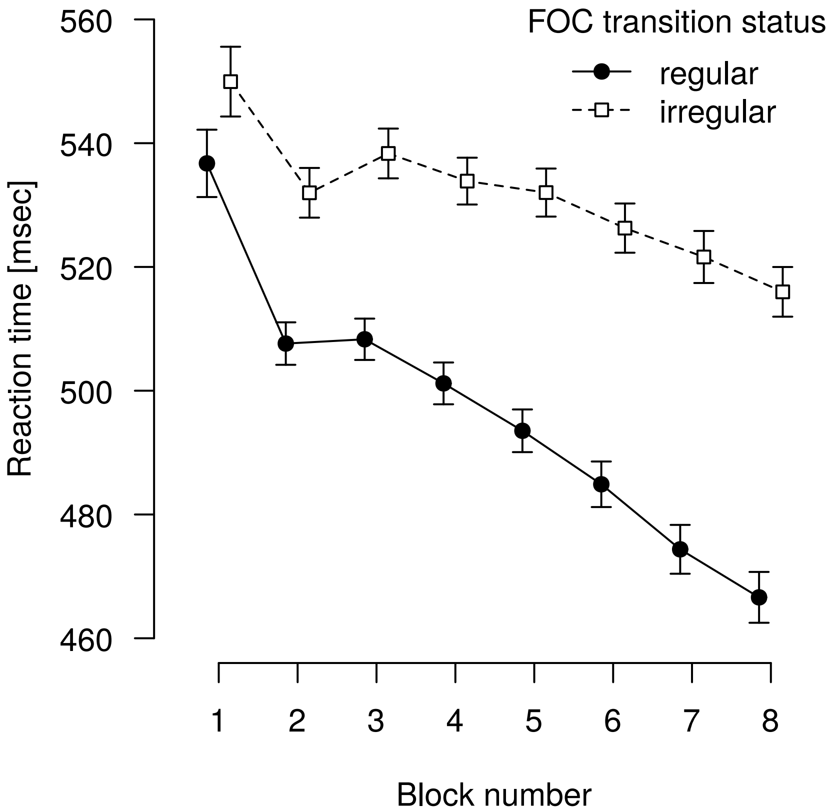

Figure 9. RTs during acquisition phase of Experiment 1, split by *FOC transition status*. Error bars represent 95% within-subjects confidence intervals.

regular vs. irregular) repeated-measures ANOVA that revealed a main effect of *block number*, $F(6.29, 917.65) = 8.35$, $MSE = 9.42$, $p < .001$, $\hat{\eta}^2_G = .015$, reflecting increasing error rates over blocks; and a main effect of *FOC transition status*, $F(1, 146) = 188.88$, $MSE = 11.92$, $p < .001$, $\hat{\eta}^2_G = .066$, reflecting an accuracy advantage (i.e., lower error rates) for regular transitions. The interaction of *block number* and *FOC transition status* was also significant, $F(6.53, 953.88) = 7.36$, $MSE = 7.09$, $p < .001$, $\hat{\eta}^2_G = .011$, reflecting an increase of the accuracy advantage for regular (as compared to irregular) transitions over blocks, indicating successful sequence learning.

Generation task. We first analyzed generation performance by applying standard ANOVA techniques to the proportions of regular transitions generated in inclusion and exclusion blocks. We then analyzed generation performance for those transitions that were revealed to participants, testing our hypotheses about the effects of practice on generation performance.

Overall generation performance.

Figure 11 shows the overall generation performance. We conducted a 5 (*Condition*: Control vs. No-Practice vs. Unspecific-Practice vs. Practice vs. Transfer) × 2 (*Order*: Inclusion first

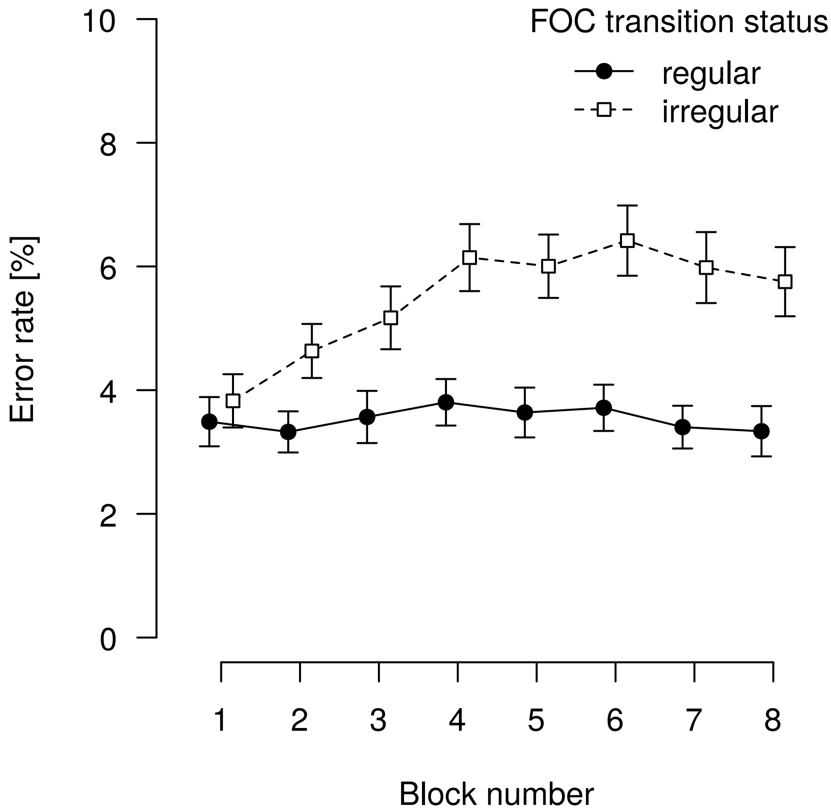

Figure 10. Error rates during acquisition phase of Experiment 1, split by *FOC transition status*. Error bars represent 95% within-subjects confidence intervals.

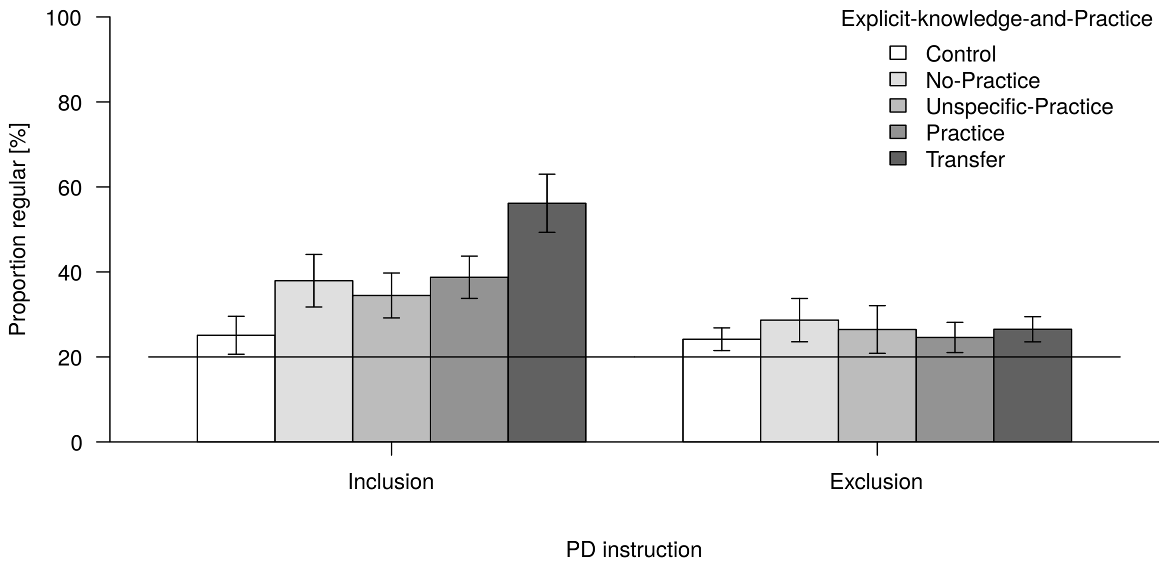

Figure 11. Generation performance in Experiment 1, excluding repetitions. Error bars represent 95% confidence intervals. Horizontal lines represent chance baseline.

vs. Exclusion first) $\times$ 2 (*PD instruction*: Inclusion vs. Exclusion) ANOVA that revealed a main effect of *PD instruction*, $F(1, 137) = 64.03$, $MSE = 176.93$, $p < .001$, $\hat{\eta}^2_G = .199$, participants generated more regular transitions in inclusion than exclusion blocks; and a main effect of *Explicit-knowledge-and-Practice*, $F(4, 137) = 13.81$, $MSE = 155.01$, $p < .001$, $\hat{\eta}^2_G = .158$, indicating a clear influence of our manipulation of explicit knowledge and on generation performance. Moreover, the interaction of *Explicit-knowledge-and-Practice* and *PD instruction* reached significance, $F(4, 137) = 9.63$, $MSE = 176.93$, $p < .001$, $\hat{\eta}^2_G = .130$, indicating that the effect of *Explicit-knowledge-and-Practice* is qualified by *PD instruction*. The interaction of *PD instruction* and *block order* also reached significance, $F(1, 137) = 10.89$, $MSE = 176.93$, $p = .001$, $\hat{\eta}^2_G = .041$. To disentangle these interactions, we analyzed inclusion and exclusion performance, separately.

Analyzing inclusion blocks, a 5 (*Condition*: Control vs. No-Practice vs. Unspecific-Practice vs. Practice vs. Transfer) $\times$ 2 (*Order*: Inclusion first vs. Exclusion first) ANOVA revealed a significant main effect of *Explicit-knowledge-and-Practice*, $F(4, 137) = 17.74$, $MSE = 211.85$, $p < .001$, $\hat{\eta}^2_G = .341$, indicating that inclusion performance increased with the number of revealed transitions; and a main effect of *block order*, $F(1, 137) = 9.95$, $MSE = 211.85$, $p = .002$, $\hat{\eta}^2_G = .068$: participants generated more regular transitions if inclusion followed exclusion; the interaction of *Explicit-knowledge-and-Practice* and *block order* did not reach significance, $F(4, 137) = 0.52$, $MSE = 211.85$, $p = .723$, $\hat{\eta}^2_G = .015$.

Analyzing exclusion blocks, a 5 (*Condition*: Control vs. No-Practice vs. Unspecific-Practice vs. Practice vs. Transfer) $\times$ 2 (*Order*: Inclusion first vs. Exclusion first) ANOVA revealed *no* significant effects on exclusion performance (all ps $\geq .143$). Specifically, increasing levels of explicit knowledge did *not* reduce the level of regular transitions generated under exclusion instructions.

Generation performance for revealed transitions.

To explore effects of practice, we analyzed generation performance for only those transitions about which explicit knowledge was revealed (see Figure 12). A 5 (*Condition*: Control vs. No-Practice vs. Unspecific-Practice vs. Practice vs. Transfer) $\times$ 2 (*Order*: Inclusion first vs. Exclusion first) $\times$ 2 (*PD instruction*: Inclusion vs. Exclusion) ANOVA revealed a nonsignificant main effect of *Explicit-knowledge-and-Practice*, $F(3, 110) = 2.00$, $MSE = 660.29$, $p = .119$, $\hat{\eta}^2_G = .028$, but a significant main effect of *PD instruction*, $F(1, 110) = 243.88$, $MSE = 575.67$, $p < .001$, $\hat{\eta}^2_G = .508$, and their significant interaction, $F(3, 110) = 5.59$, $MSE = 575.67$, $p = .001$, $\hat{\eta}^2_G = .066$. The main effect of *PD instruction* reflects the clear influence of the PD instruction on the expression of explicit knowledge depicted in Figure 12. It is present in all practice conditions but modulated by amount of knowledge and type of practice (i.e., greater effects given specific practice): The effect was greatest in

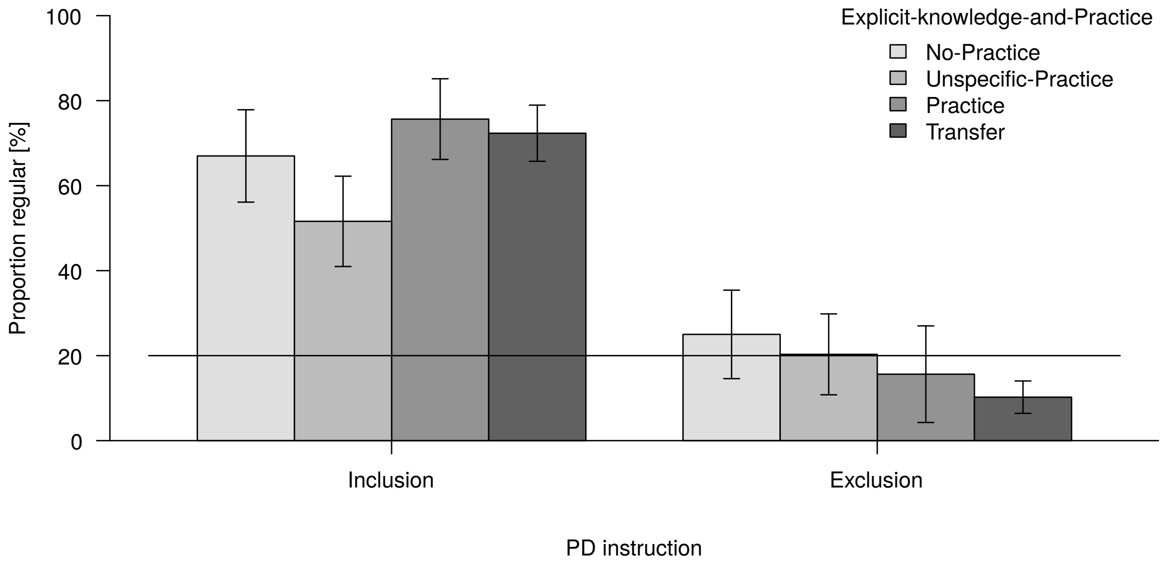

Figure 12. Generation performance for revealed transitions in Experiment 1. Error bars represent 95% confidence intervals. Horizontal lines represent chance baseline.

the *Transfer* group, $t(29) = -14.84$, $p < .001$, $d = -2.71$; somewhat smaller in the *Practice* group, $t(28) = -9.79$, $p < .001$, $d = -1.82$; it was still smaller and comparable in the *No-practice* group, $t(28) = -5.25$, $p < .001$, $d = -0.97$, and the *Unspecific-practice* group, $t(29) = -5.13$, $p < .001$, $d = -0.94$.

We investigated this issue more closely in two sets of follow-up analyses. Whereas the above findings suggest that practice improves the degree to which explicit knowledge is expressed in the generation task, it does not elucidate the mechanism by which this occurs. One mechanism by which practice may improve performance is by boosting the proportion of regular transitions in inclusion blocks.

Inclusion performance for revealed transitions in the *No-Practice* and *Practice* groups was analyzed as a function of practice (practiced vs. non-practiced) . Results showed no effect of practice on inclusion performance, $F(1, 56) = 0.21$, $MSE = 696.48$, $p = .652$, $\hat{\eta}^2_G = .004$. Similarly, when we compared inclusion performance for practiced vs. non-practiced transitions in the *Transfer* group, there was no effect of practice, $F(1, 29) = 1.19$, $MSE = 365.77$, $p = .285$, $\hat{\eta}^2_G = .014$. We conclude that practice did not affect inclusion performance for revealed transitions.

Next, we analyzed whether practice improves suppressing the regular transition in the exclusion task. We speculated that, without training, participants might not be able to suppress their generation of regular transitions below baseline level in the exclusion task. We compared generation performance for the revealed transitions between the *No-Practice* and *Practice* groups (see Figure 13, left panel). The expected below-baseline performance was not

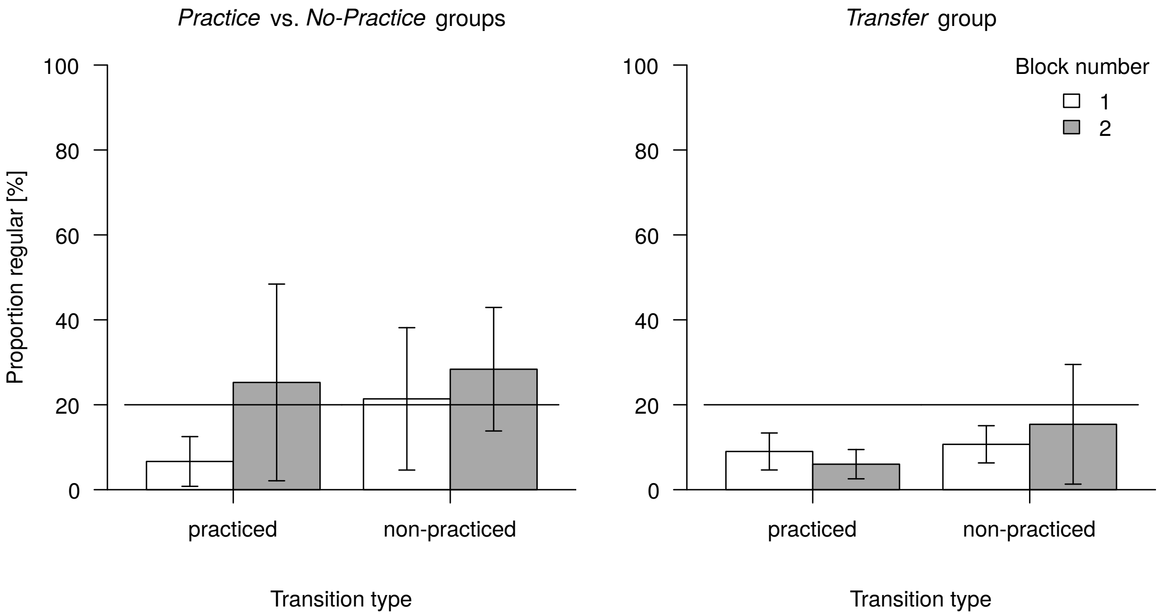

Figure 13. Exclusion performance for revealed transitions in Experiment 1. Left: Between-subjects comparison between *Practice* and *No-Practice* groups. Right: Within-subjects comparison in *Transfer* group. Horizontal lines represent chance baseline.

found when aggregating across both blocks: Whereas the direction of effects was as expected, there was no deviation from chance, neither for the practice condition, $t(28) = -0.79$, $p = .219$, $d = -0.15$, nor for the no-practice condition, $t(28) = 1.60$, $p = .940$, $d = 0.30$. However, the pattern was present in the first block: Below-chance performance was found for the practice condition, $t(14) = -4.89$, $p < .001$, $d = -1.26$, but not for the no-practice condition, $t(13) = 0.18$, $p = .569$, $d = 0.05$.

To more directly establish a practice effect, we next turned to the data from the *Transfer* group for a within-subjects comparison of practiced and non-practiced transitions. In doing so, we also addressed the transfer hypothesis: If training on one transition transfers to other transitions, we should find below-chance performance also for non-practiced transitions in the *Transfer* group. This was confirmed: Generation performance was below baseline both for practiced, $t(29) = -9.60$, $p < .001$, $d = -1.75$ and non-practiced transitions, $t(29) = -2.04$, $p = .025$, $d = -0.37$, indicating transfer of exclusion practice from practiced to non-practiced transitions (see Figure 13, right panel).[21]

[21] Analyzing only the first block revealed the same pattern of results: Generation performance was below chance for practiced, $t(14) = -5.42$, $p < .001$, $d = -1.40$, as well as for non-practiced transitions, $t(14) = -4.56$, $p < .001$, $d = -1.18$.

Discussion

Participants in Experiment 1 acquired knowledge about the sequence, as expressed in RT and accuracy advantages for regular transitions that increased over SRTT blocks. Participants received different amounts of instructed explicit knowledge, and they were able to express this knowledge in the inclusion task, as revealed by a main effect of *Explicit-knowledge-and-Practice* on inclusion performance. Conversely, participants were not able to express their knowledge in the exclusion task, as there was no effect of our explicit knowledge manipulation on exclusion performance. This finding violates the monotonicity assumption.

Analyzing exclusion performance of only revealed transitions, performance differed across groups (i.e., practice conditions), suggesting that explicit knowledge was indeed expressed under exclusion instructions, and that specific exclusion practice was beneficial to implementing these instructions. However, even with practice, inclusion performance did not reach ceiling and exclusion performance did not reach floor levels, indicating that participants were not able to exhaustively express their explicit knowledge of these transitions in the generation task. This pattern of results is also in line with Wilkinson and Shanks's (2004) speculation that participants adopt perseverative response strategies especially under exclusion instructions; these might be mildly informed by strong explicit knowledge (e.g., in our *Transfer* group).

Importantly, the results showed no effect of practice on inclusion performance of revealed transitions. Such an effect would be expected if explicit knowledge revealed to participants after the end of the SRTT differed from explicit knowledge acquired by participants during the SRTT (e.g., because in the latter case, during the remainder of the SRTT participants would have repeated opportunities to practice including their explicit knowledge by intentionally implementing it into a motor pattern). The absence of this effect corroborates the validity of the present explicit-knowledge manipulation.

Furthermore, even if (after repeated opportunity to practice) participants were able to refrain from generating some of the revealed transitions, this was not consistently reflected in below-baseline overall generation performance. It can thus be concluded that increasing amounts of explicit knowledge do not necessarily lead to fewer regular transitions being generated; the monotonicity assumption of the ordinal PD is thus violated. As a consequence, if the ordinal PD were applied to such data, a change in only explicit knowledge between two conditions would thus be misinterpreted as reflecting changes in both implicit *and* explicit knowledge.

In sum, Experiment 1 showed that, first, increasing amounts of explicit knowledge were not reflected in decreasing levels of exclusion performance, showing that the monotonicity assumption underlying the ordinal PD approach is violated. Second, explicit knowledge can nevertheless be used under exclusion instructions to decrease performance to below-baseline levels (if not exhaustively, and only under specific practice conditions); thus, we can reject the hypothesis that explicit knowledge does not affect exclusion performance at all.

Third, the usage of explicit knowledge in the generation task was higher under inclusion than exclusion, suggesting a violation of invariance (i.e., $C_I > C_E$). Experiments 2 and 3 more directly tested this assumption.

Experiment 2

Experiment 2 applied the parametric PD model and tested the invariance assumption for automatic and controlled processes using materials with first-order conditional regularities. We implemented two different levels of implicit knowledge by presenting either random or probabilistic sequences to participants during the SRTT. Orthogonally, we implemented two different levels of explicit knowledge by experimentally inducing such knowledge: After the SRTT, we informed one half of participants about one of the six transitions in the regular sequence.

Method

Design. The study realized a 2 (*Material*: random vs. probabilistic) × 2 (*Explicit knowledge*: no transition revealed vs. one transition revealed) × 2 (*PD instruction*: inclusion vs. exclusion) × 2 (*Block order*: inclusion first vs. exclusion first) design with repeated measures on the *PD instruction* factor.

Participants. One hundred and twenty-one participants (87 women) aged between 17 and 51 years ($M = 23.7$ years) completed the study. Most were undergraduates from University of Cologne. Participants were randomly assigned to experimental conditions. They received either course credit or 3.50 Euro for their participation.

Materials. We used two different types of material:

- A *random* sequence was randomly generated for each participant anew by drawing with replacement from a uniform distribution of six response locations.
- A *probabilistic* sequence was generated similar to the sequence in Experiment 1.

In both materials there were no direct repetitions of response locations. In the random group, there was no 'regular' sequence, and transition frequencies varied across persons. To compute the dependent variable in the generation task (i.e., the proportion of rule-adhering or regular transitions), we used the generating sequence for participants who worked on *probabilistic* material; for participants who worked on *random* material, we determined an individual criterion for each participant. In order to calculate the individual criteria, we first generated all possible sequences that follow the constraints that they are 6-item-sequences that do not contain repetitions and contain all six response locations. Then, for each participant, we calculated how many of the transitions that were presented during the acquisition phase followed each of those 120 non-redundant 6-item-sequences. We then chose, for each participant anew, the sequence that most frequently adhered to the transitions presented during acquisition phase and took this 6-item-sequence to calculate the dependent variable during the generation phase. Given probabilistic materials, this scoring leads to the same results as using the generating sequence as a criterion. For the group that was instructed about a regular transition, the *criterion sequence* always contained the revealed transition.

Procedure. During an SRTT consisting of eight blocks with 144 trials each (for a total of 1,152 responses), participants were trained on either random or probabilistic sequences. After the SRTT, participants were informed about the underlying sequential structure of stimulus locations and asked to generate a short sequence of six key presses that followed this (unspecified) structure.

The generation task followed, with counterbalanced order of inclusion versus exclusion blocks. Prior to the inclusion task, two generation-practice blocks involved inclusion instructions; prior to the exclusion task, the first generation-practice block was performed under inclusion instructions and the second generation-practice block was performed under exclusion instructions. If participants who were explicitly informed about one transition failed to include (exclude) the revealed transition in practice blocks, they were informed that they did something wrong; the already revealed transition was again presented and two additional practice blocks had to be performed (if a participant failed to include the transition during the first practice block, they were immediately presented with the sequence knowledge, again). This procedure was repeated until the revealed transition was successfully included (excluded) in two consecutive practice blocks (in contrast to Experiment 1, where the number of practice blocks was held constant). Upon completing the computerized task, participants answered the same questionnaire as in Experiment 1.

Data analysis. For analyses of reaction times during the acquisition task, we excluded the first trial of each block as well as trials with errors, trials succeeding an error, reactions faster than 50 ms and those slower than 1,000 ms. For analyses of error rates during the acquisition task, we excluded the first trial of each block.

Generation task analyses were conducted with the first trial of a block as well as any response repetitions excluded. For the model-based analyses, we used hierarchical Bayesian extensions of the process-dissociation model (Klauer, 2010; Rouder & Lu, 2005; Rouder et al., 2008). We estimated model $\mathcal{M}_1$ that extended the traditional process-dissociation model by allowing for a violation of the invariance assumption: Controlled and automatic processes were allowed to vary as a function of instruction (inclusion vs. exclusion). The first-level equations of this model were given by:

$$\begin{aligned} I_{ij} &= C_{ijm} + (1 - C_{ijm})A_{ijm}, \quad && m = 1 \\ E_{ij} &= (1 - C_{ijm})A_{ijm}, && m = 2 \end{aligned}$$

where i indexes participants, j indexes transition type (i.e., revealed: $j = 1$; nonrevealed: $j = 2$), and m indexes the *PD instruction* condition (inclusion: $m = 1$; exclusion: $m = 2$).

Parameters C_{ijm} and A_{ijm} are probabilities in the range between zero and one; following previous work (e.g. Albert & Chib, 1993; Klauer et al., 2015; Rouder et al., 2008), we used a probit function to link these probabilities to the second-level parameters as follows:

$$C_{ijm} = \begin{cases} \Phi(\mu^{(C)}_{km} + \delta^{(C)}_{im}) & \text{if } j = 1 \text{ (item has been revealed)} \\ 0 & \text{if } j = 2 \text{ (item has not been revealed)} \end{cases}$$

and

$$A_{ijm} = \Phi(\mu^{(A)}_{jkm} + \delta^{(A)}_{ijm})$$

where Φ denotes the standard normal cumulative distribution function, $\mu^{(C)}_{km}$ is the fixed effect of material k (that participant i worked on during the SRTT) and *PD instruction* condition m on controlled processes. $\delta^{(C)}_{im}$ is the ith participant's deviation from his or her group's mean.

Accordingly, $\mu^{(A)}_{jkm}$ is the fixed effect of transition type j, material k, and *PD instruction* condition m on automatic processes, and $\delta^{(A)}_{ijm}$ is the ith participant's deviation from the corresponding mean. Priors on parameters are given in the Appendix D.

This specification imposes two auxiliary assumptions to the model: First, it is assumed that controlled processes C are set to zero for nonrevealed transitions (i.e., $C = 0$ for $j = 2$), in other words, we assumed that no explicit knowledge has been acquired during the SRT phase. Second, it is assumed that automatic processes A do not vary as a function of the between-subjects manipulation of explicit knowledge l (i.e., $\mu^{(A)}_{l=1} = \mu^{(A)}_{l=2}$). These assumptions allowed us to relax and test the invariance assumption by obtaining separate estimates of both C and A for the inclusion and exclusion conditions (note that a *full* model relaxing all three assumptions cannot be estimated).

To assess goodness of fit, we used posterior predictive model checks as proposed by Klauer (2010): Statistic T_{A1} summarizes how well the model describes the individual category counts for the eight categories (revealed vs. nonrevealed $\times$ regular vs. nonregular $\times$ inclusion vs. exclusion). Statistic T_{B1} summarizes how well the model describes the covariations in the data across participants.

Additionally, we also estimated a model $\mathcal{M}_2$ that does not impose the auxiliary assumptions but enforces the invariance assumptions (i.e., parameters were not allowed to vary as a function of PD instruction condition m):

$$I_{ij} = C_{ij} + (1 - C_{ij})A_{ij}$$
$$E_{ij} = (1 - C_{ij})A_{ij}$$

The second-level equations of model $\mathcal{M}_2$ are then given by:

$$C_{ij} = \Phi(\mu_{jkl}^{(C)} + \delta_{ij}^{(C)})$$

and

$$A_{ij} = \Phi(\mu_{jkl}^{(A)} + \delta_{ij}^{(A)})$$

where i indexes participants, j indexes transition type, k indexes the learning material that participant i worked on during the SRTT, and l indexes the manipulation of explicit knowledge (i.e., whether or not a transition has been revealed to participant i). Note that, given this model specification, separate parameters are estimated for each between-subjects condition kl and each transition type j, while the invariance assumption is maintained (i.e., there is no index m for *PD instruction* in the model equations).

These two models were compared using the deviance information criterion DIC (Spiegelhalter et al., 2002; Spiegelhalter, Best, Carlin, & van der Linde, 2014); the DIC is an extension of AIC for Bayesian hierarchical models, and differences of 10 are considered to imply *strong* evidence in favor of the model with the lower DIC value (Klauer et al., 2015). Therefore, if model $\mathcal{M}_1$ outperforms model $\mathcal{M}_2$, it can be concluded that the auxiliary assumptions are less problematic than the invariance assumptions.

Furthermore, model $\mathcal{M}_1$ yields separate estimates of controlled and automatic processes for both inclusion and exclusion. The invariance assumption can be targeted directly by calculating the posterior differences $A_I - A_E$ and $C_I - C_E$: If the posterior distributions of these differences include zero, it can be concluded that the respective invariance assumption holds; if the posterior does not contain zero, it can be concluded that the respective invariance

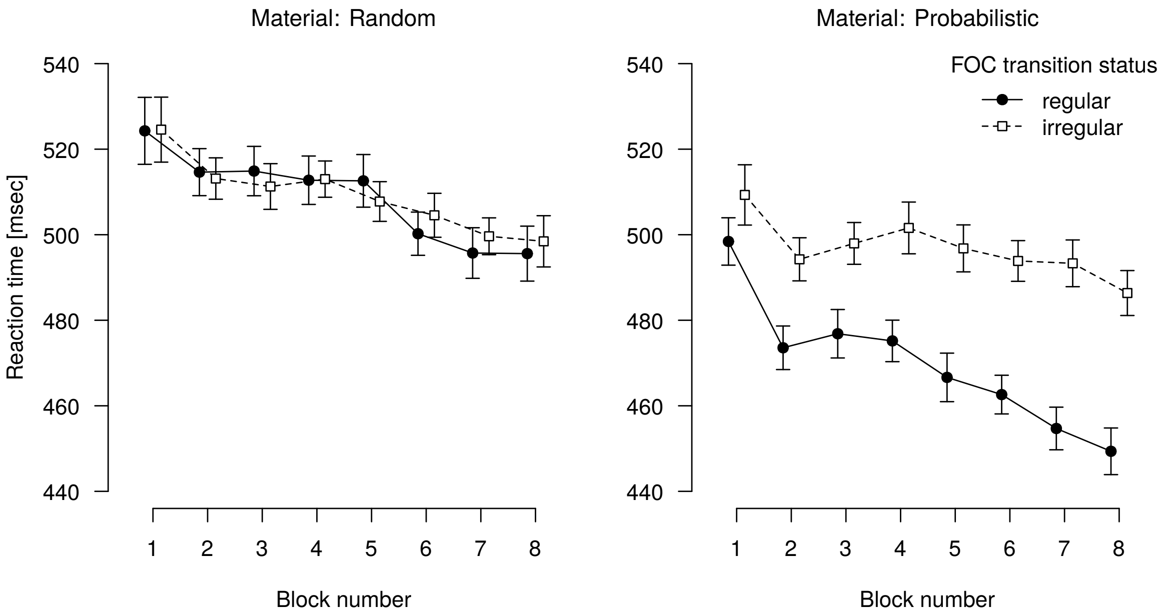

Figure 14. RTs during acquisition phase of Experiment 2, split by *material* and *FOC transition status*. Error bars represent 95% within-subjects confidence intervals.

assumption is violated.

Results

We first analyzed the performance data from the SRTT to determine whether sequence knowledge had been acquired during the task. Next, we analyzed generation task performance using hierarchical PD models (descriptive statistics and ordinal-PD analyses are reported in Appendices A and B).

Acquisition task. If participants acquired knowledge about the (probabilistic) regularity underlying the sequence of key presses, we expect a performance advantage for regular over irregular transitions, reflected in reduced RT and/or error rate. If this advantage is due to learning, it is expected to increase over SRTT blocks.

Reaction times.

Figure 14 shows reaction times during the SRTT. We conducted a 2 (*Material*: Random vs. Probabilistic) $\times$ 8 (*Block number*) $\times$ 2 (*FOC transition status*: regular vs. irregular) ANOVA that revealed a main effect of *material*, $F(1, 119) = 8.11$, $MSE = 39,617.25$, $p = .005$, $\hat{\eta}^2_G = .055$; a main effect of *block number* $F(4.89, 582.06) = 33.35$, $MSE = 1,032.91$, $p < .001$, $\hat{\eta}^2_G = .029$; a main effect of *FOC transition status*, $F(1, 119) = 125.46$, $MSE = 714.88$, $p < .001$, $\hat{\eta}^2_G = .016$; an interaction of *material* and *FOC transition status*, $F(1, 119) = 121.57$, $MSE = 714.88$, $p < .001$, $\hat{\eta}^2_G = .015$; an interaction of *block number*

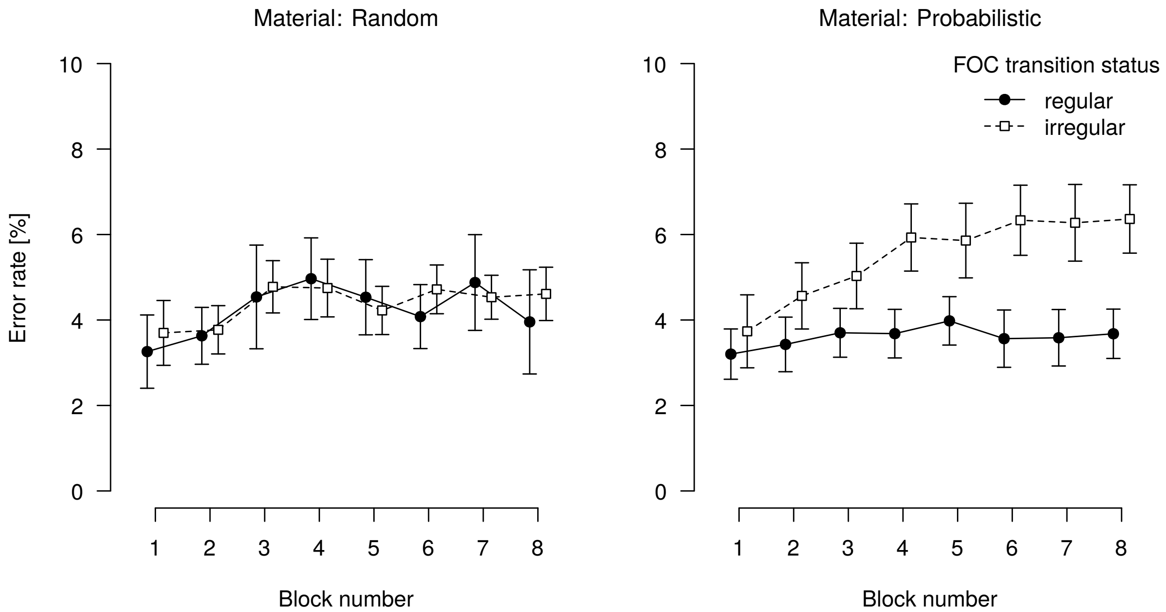

Figure 15. Error rates during acquisition phase of Experiment 2, split by *material* and *FOC transition status*. Error bars represent 95% within-subjects confidence intervals.

and *FOC transition status*, $F(6.32, 752.52) = 10.68$, $MSE = 197.96$, $p < .001$, $\hat{\eta}^2_G = .002$; and a three-way interaction between *material*, *block number*, and *FOC transition status*, $F(6.32, 752.52) = 5.70$, $MSE = 197.96$, $p < .001$, $\hat{\eta}^2_G = .001$.

Separate ANOVAs for each *material* condition yielded, for random material, only a significant main effect of *block number*, $F(4.38, 258.47) = 13.09$, $MSE = 1,276.78$, $p < .001$, $\hat{\eta}^2_G = .026$, with RTs decreasing over blocks (all other Fs < 1). For probabilistic material, in contrast, we obtained main effects of *block number*, $F(5.07, 304.28) = 22.09$, $MSE = 891.30$, $p < .001$, $\hat{\eta}^2_G = .035$; and of *transition status*, $F(1, 60) = 182.32$, $MSE = 976.60$, $p < .001$, $\hat{\eta}^2_G = .061$ (i.e. responses to regular transitions were faster than those for irregular transitions); importantly, we also obtained an interaction of *block number* and *transition status*, $F(5.93, 356.02) = 15.83$, $MSE = 194.03$, $p < .001$, $\hat{\eta}^2_G = .007$, showing that the RT difference between regular and irregular transitions increased over blocks, indicating learning of the regularities inherent in the probabilistic material.

Error rates.

Figure 15 shows error rates during acquisition. We conducted a 2 (*Material*: Random vs. Probabilistic) $\times$ 8 (*Block number*) $\times$ 2 (*FOC transition status*: regular vs. irregular) ANOVA that revealed a main effect of *block number*, $F(5.83, 693.83) = 6.06$, $MSE = 11.83$, $p < .001$, $\hat{\eta}^2_G = .016$, indicating that error rates increased over blocks, and a main effect of *FOC transition status*, $F(1, 119) = 38.19$, $MSE = 13.49$, $p < .001$, $\hat{\eta}^2_G = .019$, indicating that error rates were higher for nonregular transitions. The interaction of *material* and *FOC* *transition status* was also significant, $F(1, 119) = 27.61$, $MSE = 13.49$, $p < .001$, $\hat{\eta}^2_G = .014$, reflecting the finding that the effect of the latter factor was limited to the probabilistic material. The three-way interaction of *material*, *block number*, and *FOC transition status* was however not significant, $F(6.55, 778.97) = 1.84$, $MSE = 7.94$, $p = .082$, $\hat{\eta}^2_G = .004$.

To disentangle these interactions, we analyzed both *material* groups separately. As for RT, an ANOVA for the random material group revealed only a main effect of *block number*, $F(4.94, 291.45) = 2.50$, $MSE = 16.03$, $p = .031$, $\hat{\eta}^2_G = .013$ (all other Fs < 1). The probabilistic material group showed a main effect of *block number* $F(5.73, 343.65) = 4.63$, $MSE = 10.29$, $p < .001$, $\hat{\eta}^2_G = .022$, and a main effect of *FOC transition status*, $F(1, 60) = 62.50$, $MSE = 14.23$, $p < .001$, $\hat{\eta}^2_G = .070$. Importantly, the interaction of *block number* and *FOC transition status* was significant, $F(5.9, 353.81) = 3.23$, $MSE = 7.85$, $p = .004$, $\hat{\eta}^2_G = .012$, indicating that the difference in error rates between regular and irregular transitions increased across blocks, consistent with the learning effect obtained for reaction times.

Generation task. In a second step, we investigated how learned knowledge was expressed in the generation task. We analyzed generation performance by fitting two hierarchical models, $\mathcal{M}_1$ and $\mathcal{M}_2$. $\mathcal{M}_1$ allows the automatic and controlled processes to vary between inclusion and exclusion, but it assumes that participants acquired only implicit knowledge during the SRTT, and that revealing explicit knowledge after the SRTT did not affect implicit knowledge. $\mathcal{M}_2$ is a hierarchical extension of the classical PD model that enforces the invariance assumption. We computed model fit statistics to test whether each model could account for the means, T_{A1}, and covariances, T_{B1}, of the observed frequencies. We compared both models using the DIC statistic that provides a combined assessment of parsimony and goodness of fit and penalizes models for unnecessary complexity. Parameter estimates from model $\mathcal{M}_1$ were used to address the invariance assumptions, directly.

The model checks for model $\mathcal{M}_1$ were satisfactory,

$$T_{A1}^{observed} = 491.06, T_{A1}^{expected} = 469.94, p = .369,$$

$$T_{B1}^{observed} = 9.05, T_{B1}^{expected} = 6.95, p = .366.$$

In contrast, the model checks for model $\mathcal{M}_2$ revealed significant deviations of the model's predictions from the data,

$$T_{A1}^{observed} = 1,092.06, T_{A1}^{expected} = 473.88, p = .002,$$

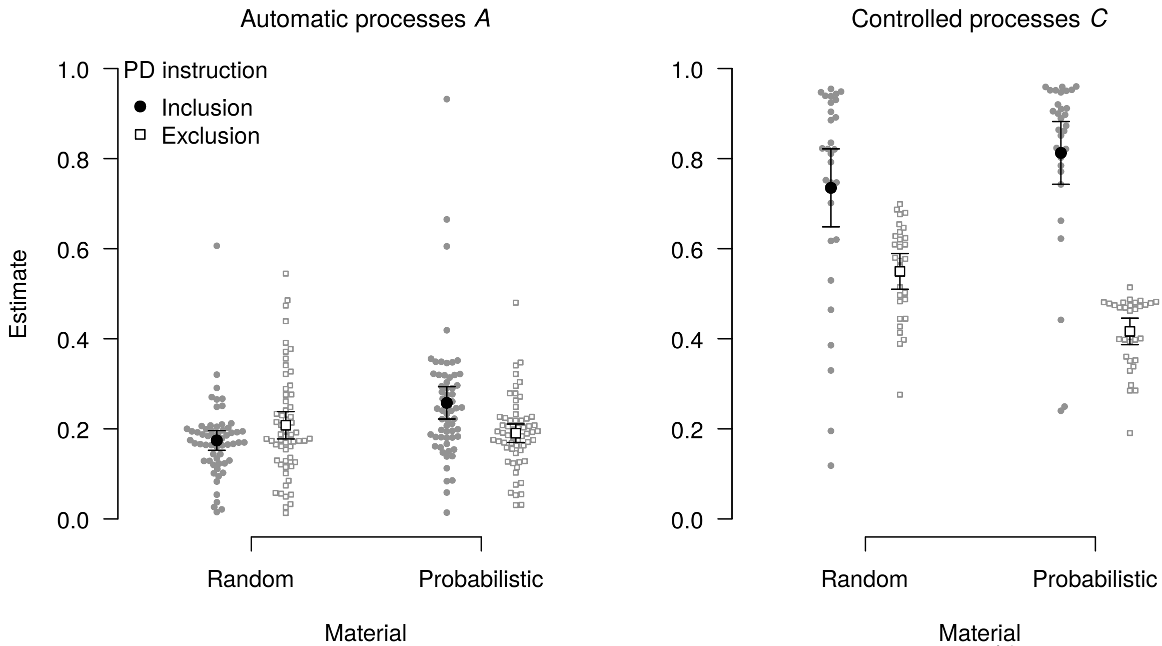

Figure 16. Parameter estimates from Experiment 2. Error bars represent 95% confidence intervals.

$$T_{B1}^{observed} = 190.05, T_{B1}^{expected} = 6.93, p < .001.$$

Model $\mathcal{M}_1$ attained a DIC value of 25,293.45 and clearly outperformed model $\mathcal{M}_2$ that attained a DIC value of 25,891.74, $\Delta\text{DIC}_{\mathcal{M}_1-\mathcal{M}_2} = -598.29$. This implies that the auxiliary assumptions we introduced to $\mathcal{M}_1$ were much less problematic than the invariance assumption. Moreover, the standard PD model enforcing the invariance assumption was not able to account for the data.

Figure 16 shows the parameter estimates obtained from model $\mathcal{M}_1$. Figure 17 shows that the invariance assumption for the automatic process was violated with $A_I > A_E$, 95% CI [.01, .03], and a Bayesian $p < .001$ ($p = .360$ for revealed transitions). The invariance assumption for the controlled process was also violated with $C_I > C_E$, 95% CI [.08, .54], and a Bayesian $p = .003$.

Robustness checks. Next we assessed whether these findings were sensitive to the assumptions of our models. Despite the fact that the auxiliary assumptions could be upheld in model comparison, and that the incorporating model was well able to account for the data, it may nevertheless still be the case that violations have biased parameter estimates. Specifically, if participants had in fact acquired explicit knowledge about nonrevealed transitions during learning, they may have used this knowledge to generate more regular transitions

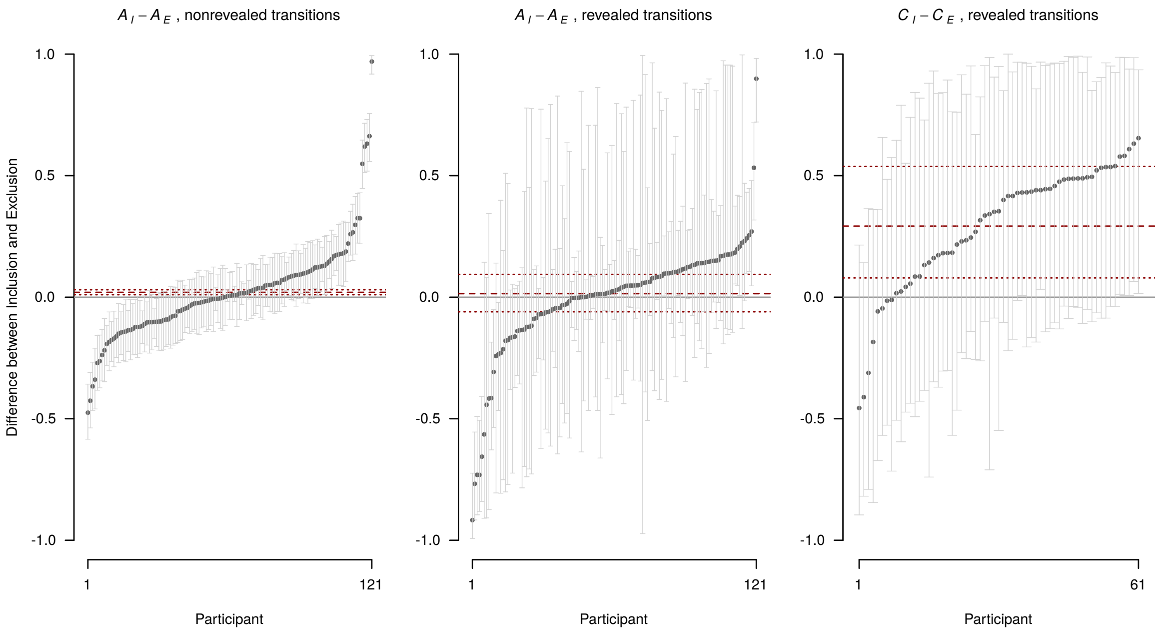

Figure 17. Posterior differences between $A_I - A_E$ and $C_I - C_E$ in Experiment 2, plotted for each participant (gray dots) with 95% credible intervals. Dashed lines represent the posterior means of the differences between mean parameter estimates. Dotted lines represent 95% credible intervals.

under inclusion than exclusion. Because of our assumption that $C = 0$ for nonrevealed transitions, this performance difference would have been reflected in greater estimates of implicit knowledge under inclusion than exclusion, and might account for the observed $A_I > A_E$ pattern.

To assess this possibility, we used the questionnaire data to exclude any transitions that participants reported in their explicit description of the sequence (while keeping the revealed transitions); if the acquired explicit knowledge was indeed the cause of the invariance violation, excluding the transitions for which knowledge was reported should make the violation disappear. To the contrary, excluding all correctly reported transitions (9.04% of cases) did not affect the pattern of results.[22]

We also tested the invariance assumption using a different model $\mathcal{M}_{1R}$ that extended $\mathcal{M}_1$ by relaxing the assumption that $C = 0$ for nonrevealed transitions (see Appendix C for details). The invariance violation for the controlled process, $C_I > C_E$, replicated in the absence of the assumption $C = 0$, demonstrating its robustness. However, the small invariance violation for the automatic process was no longer evident in $\mathcal{M}_{1R}$.

[22]Of the reported (nonrevealed) transitions, only approximately 25.47% were indeed regular transitions. After excluding *all* reported transitions regardless of whether they reflect correct knowledge or not (27.55% of cases), the invariance violation was descriptively unchanged but no longer statistically significant, Bayesian $p = .221$.

Discussion

Based on the SRTT results, we can conclude that participants acquired sequence knowledge during learning. In addition, explicit knowledge about one of the six transitions had a clear effect on generation performance for that transition.

The extended process-dissociation model $\mathcal{M}_1$ revealed a violation of the invariance assumptions for both the controlled process (i.e., $C_I > C_E$) and the automatic process (i.e., $A_I > A_E$). Model $\mathcal{M}_1$ rested on two auxiliary assumptions: It was assumed that controlled processes were not affected by learning material (i.e., no explicit knowledge was acquired from the SRTT), and that automatic processes were not affected by the manipulation of explicit knowledge (i.e., revealing a transition). Both assumptions found support in the current data as they did not harm model fit. Moreover, model comparison by the DIC showed that model $\mathcal{M}_1$ was a better account of the data than the standard process-dissociation model $\mathcal{M}_2$ that did not impose these assumptions but instead imposed the invariance assumption.

Invariance of the automatic process was significantly violated, but the magnitude of the violation was small, and it disappeared entirely under a relaxed model ($\mathcal{M}_{1R}$; see Appendix C). Given the small magnitude, and its lack of robustness to the modeling assumptions, the invariance violation of A appears to be no serious threat to the validity of the PD at this point.

In contrast, invariance of the controlled process was consistently found to be violated and the violation was large in magnitude: Confirming the speculation that explicit knowledge is not exhaustively used in exclusion, explicit knowledge was used to a greater degree under inclusion than exclusion.

Experiment 3

The main goal of Experiment 3 was to replicate the previous findings and extend them to second-order conditional (SOC) material.

A secondary goal was to explore whether different amounts of implicit knowledge are acquired with *mixed* versus *pure* SOC material. Previous studies of the SRTT using a PD generation task have employed 12-item-sequences of four response locations (e.g., SOC1 = 3−4−2−3−1−2−1−4−3−2−4−1; SOC2 = 3−4−1−2−4−3−1−4−2−1−3−2, Destrebecqz & Cleeremans, 2001; Wilkinson & Shanks, 2004). Analyzing these sequences more closely, it is evident that they did not only contain second order information (i.e., the last two locations predict the next location), but they also incorporate lower-order information: First, direct repetitions never occur; and reversals occur below chance (i.e., 1/12, whereas chance level would equal 1/3 given that repetitions are prohibited). Second, the last location of a triplet L_3 is not independent of the first location L_1 (e.g., for SOC1, $p(L_3 = 2|L_1 = 3) = 2/3$). In other words, in two out of three cases, the third location of a triplet can be predicted by the first location of a triplet alone. It is plausible that participants are able to learn this lower-order information, and that learning effects may not (only) be based on second-order information (cf., Koch & Hoffmann, 2000a; Reed & Johnson, 1994). To investigate this possibility, Experiment 3 implemented two types of probabilistic material: A *mixed SOC* material that incorporated both second-order and first-order types of information, and another *pure SOC* material that only followed a second-order regularity.

Method

Design. The study realized a 3 (*Material*: random, mixed SOC, pure SOC) $\times$ 2 (*Explicit knowledge*: no transition revealed vs. two transitions revealed) $\times$ 2 (*PD instruction*: inclusion vs. exclusion) $\times$ 2 (*Block order*: inclusion first vs. exclusion first) design with repeated measures on the *PD instruction* factor.

Participants. One hundred and seventy-nine participants (120 women) aged between 18 and 58 years ($M = 22.8$ years) completed the study. Most were undergraduates from Heinrich-Heine-Universität Düsseldorf. Data from 8 participants were excluded from generation task analyses because they had received erroneous exclusion instructions. Participants were randomly assigned to experimental conditions. They received either course credit or 3.50 Euro for their participation.

Materials. We implemented three different types of material:

- A *random* sequence was randomly generated for each participant anew by drawing with replacement from a uniform distribution of six response locations.
- A *mixed SOC* sequence incorporated two types of information: First, the third location of a triplet was conditional upon the first two locations. Second, within such regular triplets, given a fixed first-position location, there was one highly probable third-position location and two somewhat less probable third-position locations; the other three response locations never occurred for this first-position location.
- A *pure SOC* sequence followed only the second-order regularity.

In both probabilistic materials (*mixed* and *pure* SOC), 87.5% of trials adhered to the second-order regularity, which was individually and randomly selected for each participant anew. In all conditions, the material adhered to the following (additional) restrictions: (1) there were no direct repetitions of response locations, and (2) there were no response location reversals (i.e., 1-2-1). To compute the dependent variable in the generation task (i.e., the number of rule-adhering triplets), for both *probabilistic* groups, we used the second-order sequence that was used to generate each participant's materials. For the *random* group, there is no 'regular' sequence and we again computed an individual criterion sequence for each participant. For convenience, we did not generate all possible second-order sequences for these participants (as we did for first-order materials in Experiment 1), but chose to use individual criterion sequences that were randomly generated similar to the *pure SOC* material.

Procedure. The experimental procedure closely followed that of Experiment 1: In the acquisition task, participants performed a SRTT consisting of eight blocks with 180 trials each (for a total of 1,440 responses). The response-stimulus interval (RSI) was 0 ms. Following the SRTT phase, participants were told that stimulus locations during the SRTT followed some underlying sequential structure. They were then asked to try to generate a short sequence of thirty locations that followed this structure.

The generation task followed, with inclusion versus exclusion block order counterbalanced. We fixed the number of generation-practice blocks that preceded both inclusion and exclusion task: Prior to the inclusion task, three practice blocks involved inclusion instructions; prior to the exclusion task, the first and second practice block involved inclusion instructions, and the third involved exclusion instructions. Before working on practice blocks, two transitions were revealed to one half of the participants.

Upon completing the computerized task, participants were asked to complete a questionnaire containing the following items: (1) "Did you notice anything special working on the task? Please mention anything that comes to your mind.", (2) "One of the tasks mentioned a sequence in which the squares lit up during the first part of the study. In one of the experimental conditions, the squares did indeed follow a specific sequence. Do you think you were in this condition or not?", (3) "How confident are you (in %)?", (4) "Can you describe the sequence in detail?". Subsequently, participants were asked to indicate, for ten first-order transitions, the next three keys in the sequence on a printed keyboard layout. The first-order transitions were individually selected for each participant so that each participant had the chance to express full explicit knowledge about the second-order regularity.

Data analysis. For analyses of reaction times during the acquisition task, we excluded the first two trials of each block because the first two locations cannot be predicted, as well as error trials, trials succeeding an error, reactions faster than 50 ms and slower than 1,000 ms. For analyses of error rates during the acquisition task, we excluded the first two trials of each block.

Generation task analyses were conducted with the first two trials of a block as well as any

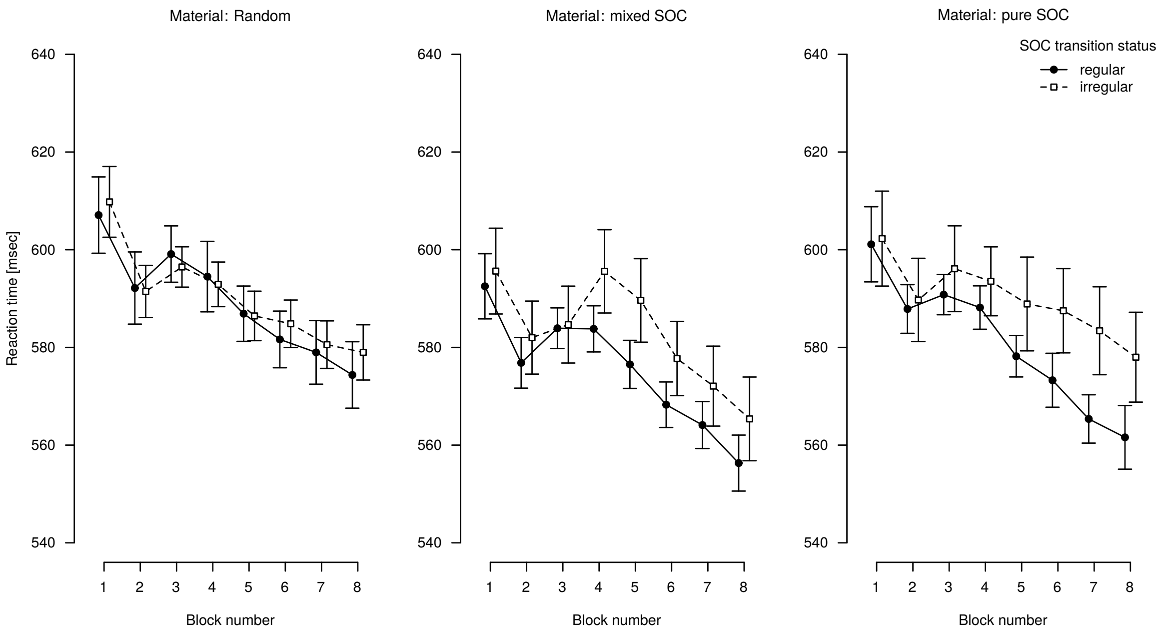

Figure 18. RTs during acquisition phase of Experiment 3, split by *material* and *SOC transition status*. Error bars represent 95% within-subjects confidence intervals.

response repetitions and reversals excluded. Model-based analyses were conducted with models $\mathcal{M}_1$ and $\mathcal{M}_2$ analogous to those used in Experiment 2 (see Appendix D for details).

Results

We first analyzed reaction times and error rates during the SRTT to determine whether sequence knowledge had been acquired during the task. Next, we analyzed generation task performance using hierarchical PD models (descriptive statistics and ordinal-PD analyses are reported in Appendices A and B).

Acquisition task. If participants acquired sequence knowledge from probabilistic materials, we expect a performance advantage for regular over irregular transitions, reflected in reduced RT and/or error rate. If this advantage is due to learning, it is expected to increase over SRTT blocks. If participants are able to learn lower-order information that is only present in *mixed SOC* material, the advantage is expected to be greater in *mixed SOC* material compared to *pure SOC*. If participants are able to learn second-order information, a performance advantage is to be expected not only in *mixed SOC* but also in *pure SOC* material.

Reaction times.

Figure 18 shows reaction times during acquisition. We conducted a 3 (*Material*: random vs. pure SOC vs. mixed SOC) $\times$ 2 (*Transition status*: regular vs. irregular SOC) $\times$ 8 (*Block number*) ANOVA with repeated measures on the last two factors that revealed a main effect of *block number*, $F(4.46, 780.51) = 41.53$, $MSE = 1,515.93$, $p < .001$, $\hat{\eta}^2_G = .020$, reflecting decreasing RT over blocks; a main effect of *transition status*, $F(1, 175) = 40.02$, $MSE = 582.10$, $p < .001$, $\hat{\eta}^2_G = .002$, reflecting an RT advantage for regular transitions; and an interaction of *block number* and *transition status*, $F(6.39, 1118.42) = 2.81$, $MSE = 439.60$, $p = .009$, $\hat{\eta}^2_G = .001$, reflecting the finding that the RT advantage for regular transitions increased over block (i.e., the sequence learning effect). We also found an interaction of *material* and *transition status*, $F(2, 175) = 7.40$, $MSE = 582.10$, $p = .001$, $\hat{\eta}^2_G = .001$, reflecting the finding that the effect of *transition status* was absent in the random material group, $F(1, 58) = 0.44$, $MSE = 380.19$, $p = .510$, $\hat{\eta}^2_G = .000$; trivially, no sequence knowledge was learned from random material.

The three-way interaction was not significant, $F(12.78, 1118.42) = 0.92$, $MSE = 439.60$, $p = .535$, $\hat{\eta}^2_G = .000$, suggesting that the sequence-learning effect did not differ across material groups. We conducted separate analyses to probe for sequence-learning effects in each material condition. Analyzing only the random material group revealed only a main effect of *block number*, $F(3.82, 221.55) = 15.74$, $MSE = 1,484.04$, $p < .001$, $\hat{\eta}^2_G = .020$ (all other ps $> .05$). In the *pure SOC* group, in contrast, a main effect of *block number*, $F(3.96, 229.51) = 12.04$, $MSE = 2,038.65$, $p < .001$, $\hat{\eta}^2_G = .019$, was accompanied by a main effect of *transition status*, $F(1, 58) = 28.48$, $MSE = 637.73$, $p < .001$, $\hat{\eta}^2_G = .004$, and an interaction of both factors, $F(6.03, 349.61) = 2.47$, $MSE = 530.13$, $p = .023$, $\hat{\eta}^2_G = .002$, reflecting a sequence learning effect on RT.

In the *mixed SOC* group, we obtained only main effects of *block number*, $F(4.91, 289.7) = 15.95$, $MSE = 1,334.22$, $p < .001$, $\hat{\eta}^2_G = .024$, and of *transition status*, $F(1, 59) = 18.83$, $MSE = 725.90$, $p < .001$, $\hat{\eta}^2_G = .003$, but the interaction of *block number* and *transition status* was not significant, $F(5.74, 338.77) = 1.15$, $MSE = 571.40$, $p = .331$, $\hat{\eta}^2_G = .001$. This is despite the fact that the effect of transition status is also likely to be a result of sequence learning, and it is of similar magnitude to that obtained in the pure SOC group. The notion that both learning effects are similar was also supported by a joint analysis of the pure SOC and mixed SOC groups: The two-way interaction between block number and transition status was significant, $F(6.14, 718.32) = 2.75$, $MSE = 527.50$, $p = .011$, $\hat{\eta}^2_G = .001$, but the three-way-interaction of *material*, *block number*, and *transition status* was not significant, $F(6.14, 718.32) = 0.87$, $MSE = 527.50$, $p = .521$, $\hat{\eta}^2_G = .000$. Taken together, we interpret these findings to show that the learning effect in the mixed SOC group was comparable to that observed in the pure SOC group but too small to reach significance in a separate analysis.

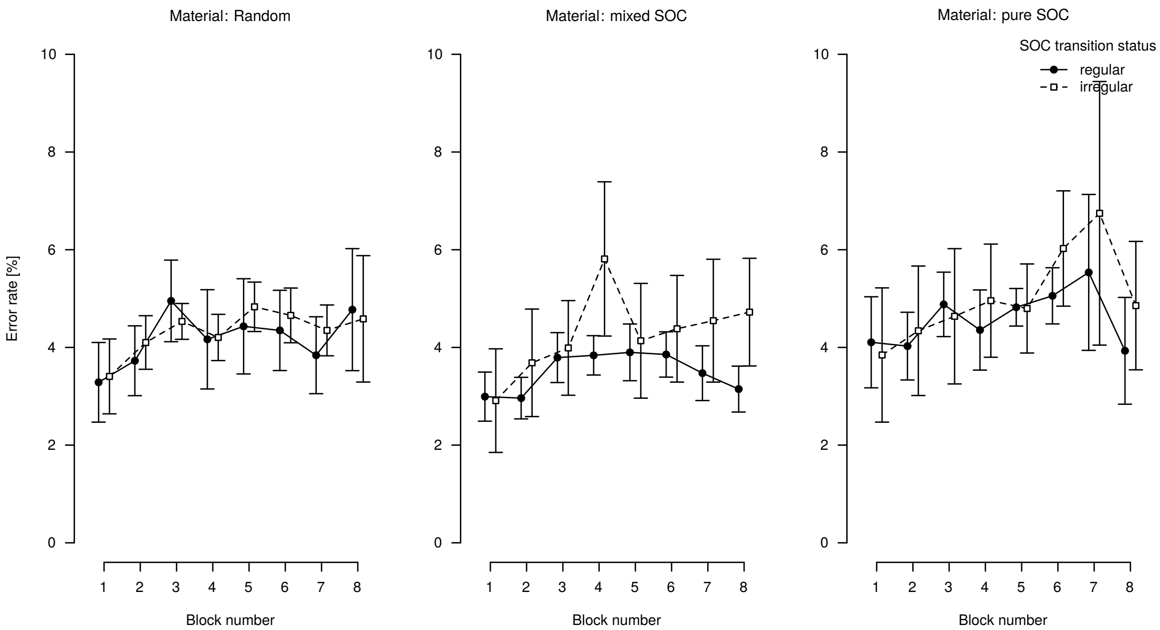

Figure 19. Error rates during acquisition phase of Experiment 3, split by *material* and *SOC transition status*. Error bars represent 95% within-subjects confidence intervals.

Error rates.

Figure 19 shows error rates during acquisition. We conducted a 3 (*Material*: Random vs. mixed SOC vs. pure SOC) $\times$ 8 (*Block number*) $\times$ 2 (*SOC transition status*: regular vs. irregular) ANOVA with repeated measures on the last two factors that revealed a main effect of *block number*, $F(3.66, 644.87) = 3.78$, $MSE = 39.10$, $p = .006$, $\hat{\eta}^2_G = .008$, reflecting increasing error rates over blocks, and a main effect of *transition status*, $F(1, 176) = 16.14$, $MSE = 9.08$, $p < .001$, $\hat{\eta}^2_G = .002$, reflecting an accuracy advantage for regular transitions. The interaction of *material* and *transition status* was not significant, $F(2, 176) = 2.66$, $MSE = 9.08$, $p = .073$, $\hat{\eta}^2_G = .001$,

Separate analyses yielded no significant effects in the random material group (all ps $> .05$). Importantly, an effect of *transition status* was clearly absent from the random material group, $F(1, 58) = 0.62$, $MSE = 7.68$, $p = .433$, $\hat{\eta}^2_G = .000$. In the *mixed SOC* group, a main effect of *block number* was found, $F(5.66, 334.01) = 2.96$, $MSE = 15.46$, $p = .009$, $\hat{\eta}^2_G = .017$, along with a main effect of *transition status*, $F(1, 59) = 12.88$, $MSE = 11.29$, $p = .001$, $\hat{\eta}^2_G = .009$, reflecting higher error rates for irregular than for regular transitions. Finally, in the *pure SOC* group, block number did not affect error rates, $F(1.87, 110.6) = 1.72$, $MSE = 133.60$, $p = .185$, $\hat{\eta}^2_G = .011$; but a main effect of *transition status* was also found, $F(1, 59) = 5.55$, $MSE = 8.24$, $p = .022$, $\hat{\eta}^2_G = .001$, reflecting higher error rates for irregular than regular transitions.

Taken together, error rates mirror RTs in that they also reflect a performance advantage for regular transitions in the mixed and pure SOC groups that was not evident in the random control group. Deviating from the RT result pattern, this advantage did not reliably increase across blocks.

Generation task. We analyzed generation performance by fitting the two hierarchical models $\mathcal{M}_1$ and $\mathcal{M}_2$ that we introduced above to the data from Experiment 3. For both models, we computed model fit statistics to assess whether each model could account for the data; we then compared both models using the DIC. Parameter estimates from model $\mathcal{M}_1$ were then used to address the invariance assumptions directly.

The model checks for model $\mathcal{M}_1$ were satisfactory,

$$T_{A1}^{observed} = 692.77, T_{A1}^{expected} = 653.45, p = .291,$$

$$T_{B1}^{observed} = 8.44, T_{B1}^{expected} = 6.04, p = .292.$$

In contrast, the model checks for model $\mathcal{M}_2$ revealed significant deviations of the model's predictions from the data,

$$T_{A1}^{observed} = 1,077.52, T_{A1}^{expected} = 652.79, p = .003,$$

$$T_{B1}^{observed} = 49.97, T_{B1}^{expected} = 6.06, p < .001.$$

Model $\mathcal{M}_1$ attained a DIC value of 38,907.43 and outperformed model $\mathcal{M}_2$ that attained a DIC value of 39,210.66, $\Delta\text{DIC}_{\mathcal{M}_1-\mathcal{M}_2} = -303.23$. This implies that our auxiliary assumptions were less problematic than the invariance assumption. Moreover, the standard PD model enforcing the invariance assumption was not able to account for the data.

Figure 20 shows the parameter estimates obtained from model $\mathcal{M}_1$. Figure 21 shows that the invariance assumption for controlled processes was again violated with $C_I > C_E$, 95% CI [.27, .63], Bayesian $p < .001$. The invariance violation was also obtained with model $\mathcal{M}_{1R}$, showing that it is robust to the specific modeling assumptions (see Appendix C). In contrast to the results of Experiment 2, the invariance assumption for automatic processes was not violated but could be upheld, 95% CI [-.01, .01], Bayesian $p = .638$ for non-revealed transitions and 95% CI [-.10, .05], $p = .763$ for revealed transitions.

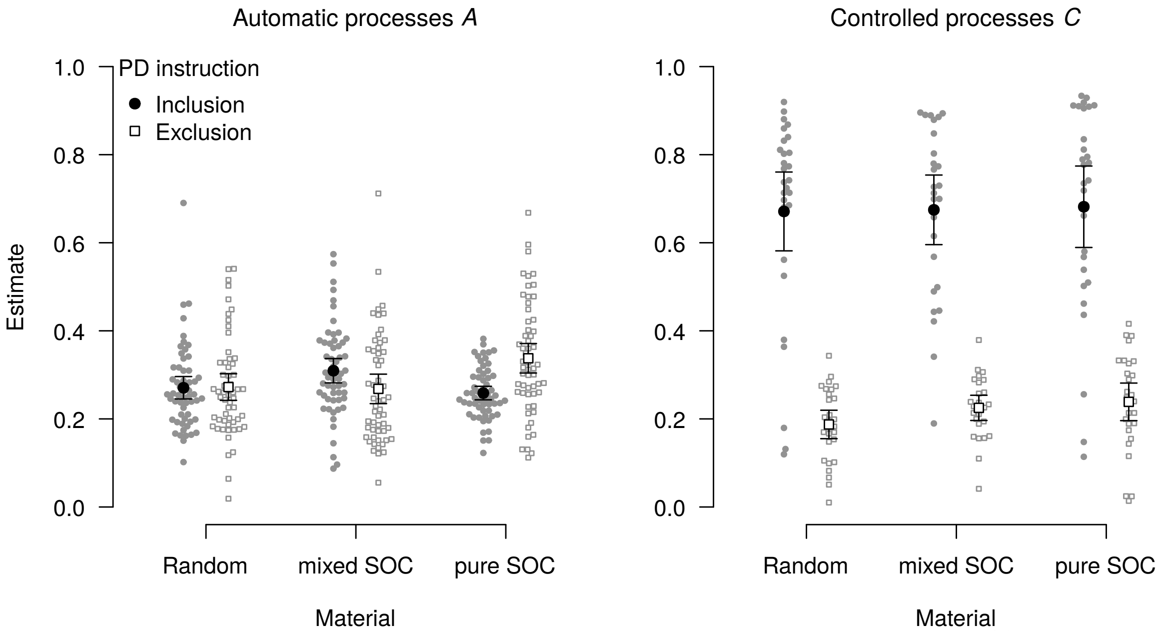

Figure 20. Parameter estimates from Experiment 3. Error bars represent 95% confidence intervals.

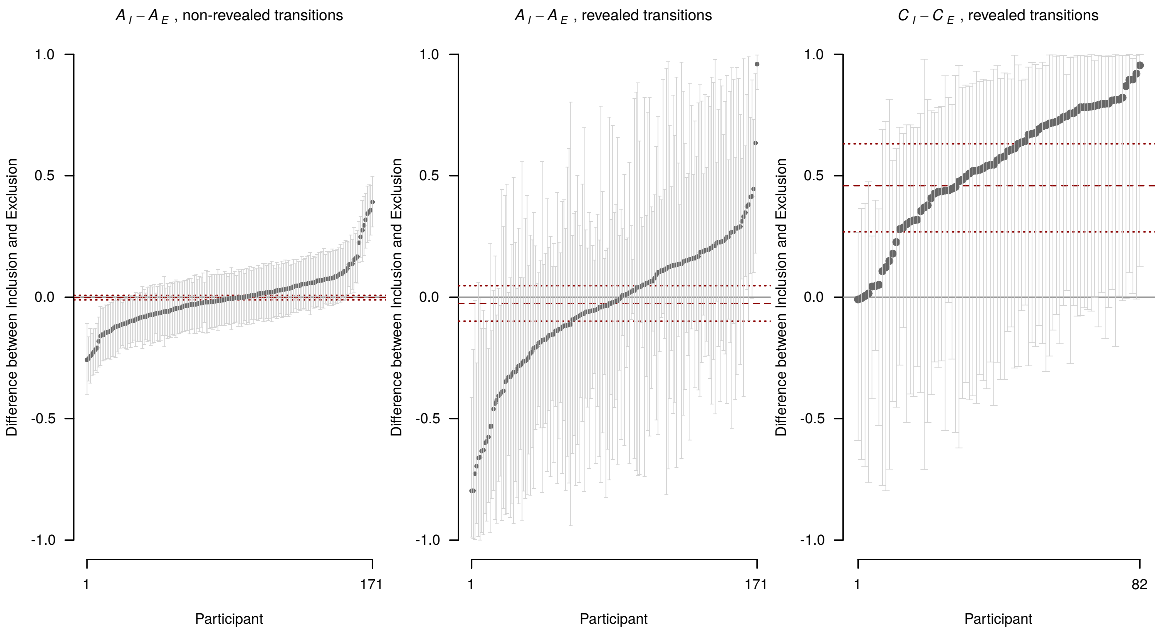

Figure 21. Posterior differences $A_I - A_E$ and $C_I - C_E$ in Experiment 3, plotted for each participant (gray dots) with 95% credible intervals. Dashed lines represent the posterior means of the differences between mean parameter estimates. Dotted lines represent 95% credible intervals.

Discussion

Based on the SRTT results, we can conclude that participants acquired some (albeit weak) sequence knowledge during learning. In addition, generation performance was clearly affected by instructed explicit knowledge, as revealed by the clearly above-zero estimates of the C parameters for revealed transitions.

An extended process-dissociation model $\mathcal{M}_1$ revealed a violation of the invariance assumption for controlled processes with $C_I > C_E$. The invariance assumption for automatic processes could be upheld. Model $\mathcal{M}_1$ rested on two auxiliary assumptions: It was assumed that controlled processes were not affected by learning material, and that automatic processes were not affected by the manipulation of explicit knowledge. Both assumptions found support in the current data as they did not harm model fit. Moreover, model selection strongly favored model $\mathcal{M}_1$ over a standard process-dissociation model $\mathcal{M}_2$ that did not impose these assumptions.

Regarding our secondary goal to explore whether different amounts of sequence knowledge are acquired from mixed versus pure second-order conditional material, we did not find evidence for a difference between these two types of material in the SRTT. This may well be due to the overall low levels of acquired sequence knowledge in the present study. Clearly, the present data are not strong enough to rule out such differences; this question requires further study.

General Discussion

Summary of main findings

The process-dissociation approach as applied to sequence learning assumes either (1) that automatic processes monotonically increase both inclusion and exclusion performance, while controlled processes increase inclusion but decrease exclusion performance (if the ordinal approach is used), or (2) that the controlled and automatic process are invariant under inclusion and exclusion instructions (if the parametric model is used).

In three sequence-learning experiments, we tested whether the monotonicity and invariance assumptions hold in the generation task. The results show a consistent pattern.

Monotonicity assumption. Increases in explicit knowledge across conditions consistently increased overall *inclusion* performance, but were insufficient to reliably decrease overall *exclusion* performance. Participants were largely unable to use their explicit knowledge to suppress the proportion of regular transitions generated in the exclusion task to levels below baseline. Below-baseline generation levels for revealed transitions were robustly found only for material with a first-order regularity, and only in participants who had explicit knowledge about (at least) two transitions and engaged in generation-task practice specific to a given to-be-excluded transition (Exp. 1, Transfer condition). In these participants, there was even some evidence that below-chance exclusion performance transferred to non-practiced explicit knowledge. However, transition-specific practice was (necessary but) not sufficient for successful exclusion: Whereas participants without such practice (i.e., the No-Practice and Unspecific-Practice conditions of Exp.1) failed to reach below-chance levels, participants with practice also failed to attain below-chance levels under exclusion instructions if they worked on the inclusion task first (i.e., Exp. 1, Practice condition). Taken together, these results confirm Wilkinson and Shanks's (2004) speculation that inclusion and exclusion strategies may differ and that explicit knowledge is not exhaustively expressed in the generation task's exclusion condition, to the effect that increasing explicit knowledge does not result in decreased generation of regular transitions under exclusion.

Invariance of the controlled process. The finding that explicit knowledge was less likely to affect exclusion performance also suggests a violation of invariance. Experiments 2 and 3 showed that, indeed, the invariance assumption for explicit knowledge was consistently violated, in first-order as well as second-order material, and despite extensive opportunity for practice. In all cases, explicit knowledge was expressed to a greater degree under inclusion than under exclusion instructions: Participants succeeded in generating the revealed transition under inclusion conditions, but failed to consistently refrain from generating that transition under exclusion conditions; specifically, under exclusion conditions, participants typically generated the revealed transition at chance levels, instead of suppressing its generation altogether as instructed.

Limitations and open questions

Before turning to the implications of the present findings, we discuss potential limitations and identify open questions.

The invariance violation of the automatic process may reflect learned explicit knowledge. In Experiment 2 that used first-order conditional material we found evidence suggesting a violation of the invariance assumption for implicit knowledge; no such evidence was however found for the second-order conditional material used in Experiment 3. If interpreted in a standard PD framework, the inclusion-exclusion performance difference resulting from this violation may lead to erroneous conclusions about the presence of explicit knowledge (if such knowledge is indeed absent), or to overestimation of the contribution of explicit knowledge. We believe these findings of an inclusion-exclusion difference in estimates of the automatic parameter should be interpreted with some caution, for at least three reasons. First, the finding was inconsistent across studies, and there are multiple possible causes of this inconsistency: The lack of a violation in Experiment 3 may be due to specific properties of the material, or it may be due to the fact that sequence knowledge levels in that study were too low for differences in its expression to be measurable.

Second, the violation was relatively small (i.e., the $A_I - A_E$ difference ranged between .01 and .03 in Exp.2; and between .00 and .03 in Exp.1, see Appendix C). In the absence of controlled influences, this would be equivalent to a difference between inclusion and exclusion performance of approximately 2 percentage points — an effect barely noticeable under typical conditions.

Third, it is unclear whether the observed invariance violation of parameter A reflects implicit knowledge at all. Note that the parameter for the automatic process captures the sum of all non-controlled influences on generation performance. In particular, it might reflect guessing strategies, and these may differ under inclusion versus exclusion conditions (Stahl et al., 2015). In other words, the above effect may reflect a violation of invariance of guessing or response strategies instead of a violation of invariance of the automatic expression of implicit knowledge. Taken together, we interpret the finding as too weak to conclude that the invariance assumption is violated also for the automatic process.

Instead of being due to guessing, the inclusion-exclusion difference in estimates of the automatic parameter may be due to explicit knowledge acquired during learning. Such an effect, if present at all, is likely to be small given that (1) the material was probabilistic and therefore difficult to learn explicitly; (2) the model incorporating the assumption that no learned explicit knowledge was learned fitted the data well; and (3) the results were unchanged when we excluded the data from transitions that participants (correctly) reproduced during debriefing. However, we cannot exclude the possibility that small amounts of explicit knowledge, obtained during the SRTT phase, may have distorted our model's parameter estimates. This interpretation could also account for the lack of such an effect in Experiment 3 given that explicit knowledge was less likely to be learned from the more complex second-order conditional material used in that study. If this were true, then any differences between inclusion and exclusion that were attributed by the model to an invariance violation of the implicit process may in fact have been a consequence of residual explicit knowledge that was not reflected in our debriefing questionnaire (perhaps due to participants' conservative reporting criteria).

To further address this possibility, we conducted additional model analyses for Experiments 2 and 3 (reported in Appendix C) that aimed at estimating the amount of this residual explicit knowledge; we still found a violation of invariance for the automatic process, but of different direction — a finding that we consider to be an artifact of the auxiliary modeling assumptions. This limitation is another reason for caution in interpreting the above finding as evidence for a violation of invariance of the automatic process. Note that it does not limit the interpretation of our main finding of the invariance violation of the controlled process, which was robust against changes in auxiliary assumptions.

The evidence for sequence learning was weak for SOC material in Experiment 3. As expected, second-order conditional material (Experiment 3) was more difficult to learn than first-order conditional material (Experiments 1 & 2). This was reflected here in the finding that (despite a 20% increase in learning trials) there was only weak evidence for sequence learning in Experiment 3. Specifically, responses to regular transitions were clearly faster and more accurate for both variants of the SOC materials, but the interaction between regularity and training block, which is critical for unambiguously interpreting a performance advantage for regular transitions as an effect of learning, was not significant. Clearly, an even larger amount of SRTT training should be realized in future studies using SOC materials. Yet, it is unlikely that the observed RT advantage for regular transitions has any other causes than learning, given that it was absent from the random condition, and that the effect could not be attributed to properties of specific transitions because regularity of a transition was randomized for each participant anew. Nevertheless, because evidence for (implicit) sequence learning was not beyond doubt, it is not warranted to interpret the modeling results as stringent tests of the invariance assumption for the automatic process.

Explicit knowledge learned via instruction may be qualitatively different from acquired explicit knowledge. The present study manipulated explicit knowledge via instruction. Although it is an established method (e.g., Liefooghe, Wenke, & De Houwer, 2012) that has yielded important insights in other domains, one might argue that explicit knowledge acquired via instruction is somehow qualitatively different from explicit knowledge acquired during SRTT training, and that therefore the present results do not speak to the question of interest regarding the invariance of the expression of acquired knowledge. We believe our manipulation to be valid for the following reasons. First, the instructed explicit knowledge communicated the same proposition about the sequence that participants would have acquired during SRTT training (i.e., that a specific location was regularly followed by another location). Second, we took precautions to avoid any inconsistency or conflict with learned sequence knowledge: Transitions that were revealed to participants were part of the regular sequence and therefore compatible with acquired (implicit or explicit) sequence knowledge. Third, we allowed participants to integrate instructed and acquired knowledge during the practice blocks before the generation task.

Given that the instructed and acquired propositions are identical, we would argue that qualitative differences between acquired and instructed knowledge are likely to involve non-propositional forms of knowledge; such non-propositional knowledge is typically considered to be implicit. Indeed, it is likely that strong implicit knowledge is a precondition for acquiring explicit knowledge (Cleeremans & Jiménez, 2002; Haider & Frensch, 2009): Instructed and acquired explicit knowledge are therefore likely to differ in the degree to which they are correlated with implicit knowledge. If participants are better able to control acquired than instructed explicit knowledge, this would then be due, paradoxically, to the presence of acquired implicit knowledge. Finally, even if that was the case, note that this would not salvage the PD method because a strong correlation between explicit and implicit knowledge would violate the independence assumption, thereby posing another problem for its validity.

Boundary conditions of the process-dissociation approach may be violated. Jacoby and colleagues (Jacoby et al., 1997; Jacoby, Toth, & Yonelinas, 1993; Toth, Reingold, & Jacoby, 1994) emphasized the importance of avoiding floor effects when applying the process-dissociation approach. In the present studies, floor effects may be present if participants succeed in avoiding to generating any regular transitions under exclusion instructions. In such cases, parameter estimates of controlled and automatic processes may be biased and might have artifactually produced an invariance violation for controlled processes (i.e., $C_I > C_E$). We would argue this is not the case in our data for the following reasons: First, generation performance in the present studies fails to show evidence for floor effects (i.e., both low overall levels as well as reduced variability): Regarding overall levels, mean performance in most conditions deviated no more than $\pm 1 SD$ from chance baseline (i.e., 20% in Experiments 1 and 2, and 25% in Exp.3). Regarding variability, the cells with the lowest overall performance (i.e., with the greatest risk of floor effects) showed variability comparable to that exhibited by the other conditions. While we found below-chance exclusion performance for revealed transitions in some conditions of our experiments (i.e., successful exclusion of the revealed transitions), the invariance violation for controlled processes was replicated not only in these conditions, but across all conditions that involved non-zero explicit knowledge. Importantly, it also replicated in Experiment 3 that realized a higher baseline level of 25%: In that study, revealed transitions were generated at rates of 25-30% under exclusion conditions, rates that are clearly unconspicuous of reflecting floor effects.

Second, while Jacoby and colleagues warned about floor effects, they also described the mechanism by which zero counts pose a threat to the validity of the estimation procedure, and proposed means to deal with this problem: If individual participants' data are analyzed separately, floor effects might be accompanied by inflated levels of zero counts in the exclusion condition (i.e., perfect exclusion) for some participants. The original PD equations would then lead to an estimate of $A = 0$ for this participant; therefore, averaging over individual A parameters would lead to an underestimation of the *automatic* parameter. As a means to circumvent this estimation problem, Toth et al. (1994) proposed complete pooling (i.e., aggregate analysis) of data. In our studies, we followed and extended this recommendation by utilizing a Bayesian hierarchical multinomial model for estimating parameters; while allowing for individual differences between parameter estimates, a participant's parameter estimate is not only affected by the data that are directly linked to this estimate, but also by the higher-level distribution for this parameter; the influence of outlier values (such as zero counts provided by some individuals) on the parameter estimates is thus minimized. Another type of estimation bias may arise if our data contained otherwise inflated levels of zero counts; this would suggest higher levels of control ability and lead to an *over*estimation of C parameters for exclusion conditions. Such inflated levels of zero counts would therefore work against our main invariance-violation finding that explicit knowledge remains underutilized under exclusion conditions (i.e., $C_I > C_E$).

Finally, our conclusion that explicit knowledge remains underutilized under exclusion is not only based on analyses of the parametric PD model (which are susceptible to biased parameter estimates due to floor effects) but was consistently corroborated by the results from ordinal-PD analyses that did not depend on estimates based on the parametric PD equations: Across all three studies, we consistently found a violation of the monotonicity assumption in the sense that explicit knowledge does not reliably decrease the proportion of regular transitions in exclusion conditions (see Appendices A and B).

Implications

We will first discuss implications for the PD approach before we suggest ways to improve measurement of sequence knowledge using the generation task. We conclude with a few broader implications.

Validity of the PD method. The present findings show that participants fail to exhaustively suppress generating regular transitions under exclusion instructions; this finding has repercussions for both the ordinal- and parametric-PD approaches.

In the ordinal approach, given a single experimental condition, it is concluded that implicit knowledge is present if exclusion performance is above a (chance or empirical) baseline; and it is concluded that explicit knowledge is present if inclusion performance exceeds exclusion performance. These conclusions depend on the assumption that a monotonically increasing controlled process should lead to a monotonic increase of inclusion performance and at the same time a monotonic decrease of exclusion performance. The present study shows, however, that exclusion performance cannot be assumed to reliably decrease with increasing explicit knowledge. This implies that the assumptions underlying the ordinal-PD approach are violated for the generation task as applied to sequence learning. In addition, we have previously shown that another assumption of ordinal PD, namely that baseline performance is identical in the inclusion and exclusion tasks, is also violated at least in some cases (Stahl et al., 2015). Given that these two fundamental assumptions are violated, the analysis approach adopted in the SRTT literature is also compromised.

The controlled process was found to operate less effectively under exclusion than inclusion instructions; in terms of the parametric PD model, invariance for the controlled process was violated with $C_I > C_E$. A model that nevertheless incorporates the invariance assumption will likely fail to adequately account for the data, and will yield distorted estimates of the automatic and controlled process. To illustrate, assume that the true values of the parameters are $C_{Inclusion} = .8, C_{Exclusion} = .4$, and $A_{Inclusion} = A_{Exclusion} = .25$. This yields the following generation proportions of regular transitions $I = .8 + (1 - .8) * .25 = 0.85$ and $E = (1 - .4) * .25 = 0.15$. When fitting a traditional PD model enforcing the invariance assumption $C = C_{Inclusion} = C_{Exclusion}$ to these data, we get $C = .7$ that lies somewhere between the true values of C, and $A = .5$ which is a vast overestimation of the true A. Importantly, note that if the true value of $A = .25$ represents chance level, applications of the traditional PD method might lead to the erroneous conclusion that implicit knowledge had been learned even if such knowledge was in fact entirely absent. In addition, if we are interested in the amount of explicit knowledge learned from the SRTT training phase, it might be argued that the higher estimate obtained from the inclusion condition might be a more valid estimate of learned explicit knowledge; the inability to express this knowledge under exclusion may be of secondary interest. By this argument, applying the traditional PD method also yields an underestimation of explicit knowledge.

We therefore recommend against using the PD method unless separate estimates of $C_{Inclusion}$ and $C_{Exclusion}$ can be obtained, for example as we have done in the present study. To do so, an extension of the standard design is necessary; for instance, in the present study we implemented two levels of an explicit-knowledge factor across which we equated the A parameters; this allowed us to estimate separate C parameters for inclusion and exclusion. Note that this strategy may not be broadly applicable in typical SRTT studies because of the strong correlation between (acquired) C and A; the assumption that the level of implicit knowledge is constant across two different levels of explicit knowledge will be warranted only in special cases such as realized in the present studies (e.g., if explicit knowledge is revealed).

Generation task as a measure of sequence knowledge. The generation task has been proposed as a useful and sensitive measure of implicit knowledge (Jiménez et al., 1996; Perruchet & Amorim, 1992). Its sensitivity may be called into question by the finding that RT effects obtained during the SRTT were often greater than implicit-knowledge effects in the generation task. In part, this may be attributed to the greater reliability of the RT measure, as it relies on aggregation across a larger number of trials than does the generation task. Another possible reason is that the generation task's sensitivity as a measure of implicit knowledge may be lower than previously thought. For instance, previous findings of implicit knowledge using the generation task may have been overestimates of implicit knowledge due to a violation of invariance for the controlled process with $C_I > C_E$. Note that most studies used much easier-to-learn materials (with four instead of six locations); it is thus plausible that participants acquired more explicit knowledge than they did in our experiments, and that the overestimation bias was more severe in those studies.

Another possible reason for overestimating implicit knowledge is that the regularities in the sequences implemented in previous research were such that the probability of reversals (e.g., 1-2-1) was below chance. Given that participants spontaneously tend to generate reversals at below-chance levels, this implies that they instead generate other regular transitions at slightly above-chance levels even in the absence of any true sequence knowledge (Stahl et al., 2015). As a consequence of this reversal-avoidance bias, implicit knowledge might be overestimated if one uses chance baselines as a reference. This problem has been discussed before (Destrebecqz & Cleeremans, 2003; Reed & Johnson, 1994; Shanks & Johnstone, 1999), and was solved by comparing performance on the training sequence with performance on a transfer sequence containing a similarly low proportion of reversals. This implies, however, that the PD approach does not provide a measure of the absolute level of implicit or explicit knowledge; instead, by relying on a comparison of performance across two sequences, it yields a difference measure that is associated with reduced reliability. In addition, the reversal-avoidance bias may not only mimic implicit knowledge; it may also mimic (or mask) explicit knowledge if it interacted with the inclusion-exclusion instructions, perhaps via different response strategies or criteria adopted under inclusion versus exclusion instructions.

Conclusion and Outlook

In light of the present findings suggesting limited validity of the PD generation task, what can we conclude about explicit and implicit sequence knowledge from its previous applications? Clearly, the violation of basic assumptions implies that PD results cannot be unambiguously interpreted: Unless we have a better understanding of the processes that drive generation performance, and the degree to which they operate under inclusion versus exclusion instructions, comparisons between inclusion and exclusion performance do not support conclusions about implicit and explicit knowledge. This also implies that a reanalysis of previous findings (which is beyond the scope of the present article) would probably provide limited insight. In this section we therefore take a different approach: We initially accept the conclusions reported in the literature about the contribution of implicit and explicit knowledge at face value; consider the implications of these conclusions about the presence of distortions arising from the invariance violation; and then discuss how the initial conclusion should be corrected in light of these distortions. To recap, the invariance violation results in overestimation of implicit knowledge and underestimation of explicit knowledge. These distortions differentially affect the three patterns of results found in the literature (i.e., evidence for only implicit knowledge, for only explicit knowledge, or both).

The first pattern, evidence for implicit but no explicit knowledge, was found in only two studies (no-RSI condition, Destrebecqz & Cleeremans, 2001; and Exp.3, 6-blocks condition, Fu et al., 2008). In these studies, however, explicit knowledge may nevertheless have been acquired; the observed lack of significant evidence for explicit knowledge may instead reflect the underestimation bias resulting from the invariance violation, perhaps combined with relatively low statistical power (with $N = 12$ and $N = 24$ in the respective conditions).

Other attempts to replicate this finding were unsuccessful and instead produced the second, opposite, pattern — evidence for explicit but no implicit knowledge (e.g., Wilkinson & Shanks, 2004). In this case, the evidence for explicit knowledge suggests that the distortions due to the invariance violation apply: Obtaining evidence for explicit knowledge despite the underestimation bias implies that explicit knowledge was likely present. Obtaining no evidence for implicit knowledge despite the likely presence of an overestimation bias supports the absence of implicit knowledge (or, alternatively, it may reflect lack of statistical power).

The third pattern—evidence for both explicit and implicit knowledge—was reported in several studies (e.g., Destrebecqz & Cleeremans, 2001, 2003; Jiménez, Vaquero, & Lupiáñez, 2006). The evidence for explicit knowledge suggests that the distortions resulting from the invariance violation may have compromised the results: Again, the evidence for explicit knowledge obtained despite the underestimation bias should probably be assumed to be reliable; however, the evidence for implicit knowledge may be an artifact of the overestimation bias and should be interpreted with caution.

Taken together, when considering the limitations discovered in our studies, the PD approach to using the generation task as a measure of implicit and explicit sequence knowledge in the SRTT has so far yielded few reliable conclusions. If anything, results support the presence of explicit knowledge and call into question the interpretation of PD results as indicative of implicit knowledge.

It might be possible to devise a version of the generation task that allows for the separation of automatic and controlled processes but does not depend on exclusion of explicit knowledge and does not induce different response criteria. For example, D'Angelo, Milliken, Jiménez, and Lupiáñez (2013) implemented such a generation task variant in artificial grammar learning in which two different inclusion instructions were compared: After learning about two different grammars, participants were asked, in the first (second) inclusion block to generate exemplars from the first (second) grammar. Under certain assumptions, performance differences between blocks can be interpreted as evidence for explicit controllable knowledge. Exclusion failure and different criteria presumably do not matter in this task: Participants were not instructed to exclude explicit knowledge, and it is plausible that the similarity of instructions for both generation tasks also induced comparable response criteria. As another example, in the domain of recognition memory, the PD procedure can be replaced by a source-memory task in which, instead of including versus excluding items from one of two study lists (A and B), participants are asked to indicate the source of the word (list A or list B; Buchner et al., 1997a; Steffens, Buchner, Martensen, & Erdfelder, 2000; Yu & Bellezza, 2000). Perhaps with a similar modification, an improved generation task may prove a useful measure of sequence knowledge. Future research should also consider using alternative methods of assessing implicit and explicit knowledge (for a recent overview, see Timmermans & Cleeremans, 2015).

One of the great benefits of multinomial models such as the PD model is that they are flexibly adaptable measurement models for studying latent cognitive processes using a wide variety of experimental paradigms (Erdfelder et al., 2009). To validate a new model, it is common to assess its goodness of fit, and to empirically demonstrate that its parameters can be selectively manipulated and interpreted psychologically (i.e., parameter estimates reflect targeted experimental manipulations in the predicted manner; Batchelder & Riefer, 1999). In many cases, however, simplifying assumptions need to be made; for instance, latent processes are equated across two or more experimental conditions (e.g., a single controlled process C is assumed to operate under inclusion and exclusion conditions). Whenever such assumptions of invariance are made, we propose that they should also be tested empirically as part of the model-validation effort when a new model is proposed, before it is used to investigate substantive issues (for an example, see Brainerd, Reyna, & Mojardin, 1999).

Chapter IV

Cognitive Processes in Implicit and Explicit Sequence Learning: A Diffusion-Model Analysis

Implicit sequence learning is frequently considered to be mediated by the formation of simple associations between stimulus features, response features, or both. In recent years, this view has been criticized as being overly simplistic, and more complex representations have been proposed: Schumacher and Hazeltine (2016) proposed that stimuli, responses, and task features are represented together in a hierarchically organized *task file*. Eberhardt, Esser, and Haider (2017) proposed that features of stimuli and response are represented together in *abstract feature codes*. Both of these newer accounts have in common that they assume a common representational coding of both stimulus and response features. To investigate the issue, it is helpful to consider the cognitive processes involved in performing the SRTT: While stimulus encoding and response execution are possibly mediated by simple associations between either stimuli or responses, response selection necessarily relies on information about both. Therefore, an involvement of response selection in implicit learning provided evidence for a common coding of stimulus and response features.

In two SRTT experiments, we analyzed response times and accuracy with a drift-diffusion model. We found that implicit sequence learning was expressed by multiple processes, involving stimulus encoding, response selection, and response execution. Importantly, we found a mediating role of response-selection processes in the expression of sequence knowledge, indicating that the representations acquired in implicit sequence learning necessarily contain information about both stimuli and responses. Explicit sequence knowledge resulted in anticipatory responding that overruled other measurable effects of learning. Implications for theories of implicit and explicit sequence learning are discussed.

Implicit learning has been demonstrated using the Serial Reaction Time Task (SRTT, Nissen & Bullemer, 1987), which has participants respond to stimuli presented at four horizontal screen locations by pressing the key that corresponds to the stimulus location. Unbeknownst to participants, the stimulus locations follow a regular sequence. With practice, participants learn to respond faster on trials with regular stimulus-location transitions than on irregular transitions. Critically, on subsequent tasks aimed at assessing the amount of acquired explicit sequence knowledge, participants are often not able to express explicit knowledge about the sequential structure. (Cohen et al., 1990; Nissen & Bullemer, 1987; Willingham et al., 1989).

Early accounts try to explain learning by the formation of simple associations

An unresolved issue in the SRTT literature is *what* is learned in implicit sequence learning, i.e., what type of information is represented in the cognitive system that allows for performance benefits in the SRTT (for a review, see Abrahamse et al., 2010). Early accounts assumed that it is a single type of representation that is formed during learning, interpreting evidence for one type of representation as evidence against the other. However, over almost three decades of research on this topic, at least three different types of representations attained strong empirical support in the literature: *Response-based learning* (i.e., R–R learning) has been the dominant model of implicit sequence learning, and found considerable support in the literature (e.g., Deroost & Soetens, 2006a; Nattkemper & Prinz, 1997; Willingham, 1999). It assumes that learning is the result of the formation of direct associations between features of consecutive responses. However, also purely *perceptual learning* (i.e., S–S learning) has been observed; it refers to the formation of associations between consecutive stimulus features (e.g., Clegg, 2005; Howard, Mutter, & Howard, 1992; Mayr, 1996; Song, Howard, & Howard, 2008). A third option, *response-effect learning* (i.e., R–S learning), refers to the formation of associations between consecutive *compounds* of a response and a subsequent stimulus (Hoffmann, Sebald, & Stöcker, 2001; Stöcker, Sebald, & Hoffmann, 2003; Ziessler & Nattkemper, 2001).

A fourth option is *learning at the response-selection stage*, which received only equivocal empirical support in the SRTT literature. It assumes that *stimulus-response* (S–R) associations are formed between consecutive stimulus-response pairs. According to this view, performing the SRTT basically involves *selecting responses* from a set of task-relevant S–R pairs; multiple S–R pairs are concurrently maintained in memory across multiple trials, where contingencies between activated pairs allow for the formation of associations between these pairs. Important for the present study, it is assumed that response selection can only be facilitated by representations that contain both stimulus and response features, which typically implies that they contain both visual and motor components.

Processing stages were considered to be mappable to specific associations

These different types of associations have been frequently considered to be relatable to the different processing stages (Donders, 1969; cf. Sternberg, 1969) that are involved in performing the SRTT (Abrahamse et al., 2010; Schwarb & Schumacher, 2012). Consider a participant performing a single trial of the SRTT: it is necessary to detect and encode a stimulus, select a corresponding response, and execute the response.

Depending on what type of information is acquired and represented in the cognitive system, different processing stages should be affected by learning. Response-based learning, by preactivating response representations via R–R associations, is considered to facilitate response execution on regular trials, while it potentially results in conflicts between actually chosen and preactivated responses. Perceptual learning (by preactivating stimulus representations via S–S associations) or response-effect learning (by preactivating stimulus representations through bidirectional R–S associations) facilitates stimulus encoding on regular trials (and, in the case of a sequence of stimulus locations, faster detection), but results in slower stimulus encoding on nonregular trials. Learning at the response-selection stage, by preactivating associated S–R pairs, results in faster response selection on regular trials, and possibly slower response selection on nonregular trials.

Criticism of associationism

Common features of these accounts are (1) that they assume the formation of relatively simple associations that preactivate a stimulus or a response in a feedforward fashion, and (2) that the main distinguishing feature of different types of associations is whether they contain stimulus features, response features, or both. Arguing that there exists evidence for at least three of these types of representations (R–R, S–S, and R–S learning), Abrahamse et al. (2010) integrated these accounts into a multi-level account of implicit sequence learning that assumes that each of these three types of representation may be formed, a proposal that also fits nicely with the multiple-systems model of sequence learning proposed by Keele et al. (2003).

However, in recent years, alternatives to the formation of such simple associations have been proposed. Hazeltine and Schumacher (2016; see also, Schumacher and Hazeltine, 2016) argued that simple associations are not rich enough to explain the patterns of transfer and flexibility that are observed in sequence learning. Schumacher and Hazeltine (2016) proposed a different type of representation which they call a *task file*: a task file consists of a set of hierarchical associations between stimulus features, response features, current task goals, and drives. A central tenet of this account is that task files span across multiple representational levels, including both stimulus and response features. The formed associations are considered to be bidirectional and to become increasingly abstract. More directly targeted to implicit learning, Eberhardt, Esser, and Haider (2017; see also, Haider, Esser, and Eberhardt, 2018; Esser and Haider, n.d.) proposed to abandon the idea that representations in implicit learning can be conceptualized as simple associations between stimulus or response features. Instead, implicit learning might be represented by *abstract feature codes* that span across stimulus, response , and task features, an idea borrowed from the *Theory of Event Coding* (*TEC*, Hommel, Müsseler, Aschersleben, & Prinz, 2001). Both proposals have in common that features of stimuli and responses are always coded together in a common representational code.

Taking these recent developments into account, it is evident that identifying learning at the response-selection stage with simple S–R associations is overly simplistic: It is conceivable that either a task file or an abstract feature code might subserve response-selection processes. Still, asserting that both stimulus and response features are necessarily represented together to subserve response selection, finding evidence for a mediating role of response selection in implicit sequence learning may be considered evidential for a common coding. Studies investigating the representational basis of implicit sequence learning typically employed an opposite logic: Evidence against simple S–R associations has been interpreted as evidence against a mediating role of response selection in implicit sequence learning. Bearing in mind that this equating is possibly unwarranted, the next section will briefly review the findings related to response-selection processes and S–R associations in implicit sequence learning.

Response-selection learning was equated with simple S–R associations

Studies that have been considered as being informative regarding the role of response selection in implicit sequence learning may be subdivided into two general approaches. The first approach investigated whether or not S–R associations underlie implicit learning; this is typically achieved by changing the mapping of stimuli to responses *between a training and a transfer task*. Asserting that learning at the response-selection stage is driven by S–R associations, it is then possible to infer on a possible mediating role of response selection processes. The second approach manipulated the difficulty of response selection more directly, typically by using more or less demanding stimulus-response mappings or by creating dual-task situations that arguably lead to interference at a central response-selection stage; if implicit learning is moderated by these manipulations, it is also possible to infer on a mediating role of response selection. We will briefly review the results from these two lines of research, adding a few studies that used different approaches.

Early evidence favoring a role of S–R associations in implicit sequence learning was provided by Willingham et al. (1989). In their Experiment 3, participants responded to stimulus colors that were mapped to spatially arranged responses; for one group of participants, stimuli were presented centrally on the screen and stimulus colors (and, hence, response locations) followed a regular sequence. For another group of participants, stimulus colors (and response locations) were randomly presented; however, the locations of stimuli—a task-irrelevant stimulus feature—followed the same underlying regularity that stimulus colors followed in the other group. On a subsequent *transfer task*, participants had to respond to stimulus locations instead of color, and either stimulus or response locations followed the sequence participants were already exposed to during training. In contrast to their predictions, they found no transfer of sequence knowledge in both groups, and speculated that neither sequences of stimuli nor sequences of responses, but *condition-action rules* (i.e., S–R rules) had been learned, which could not be used after switching from stimulus colors to stimulus locations in the transfer task.

Willingham, Nissen, and Bullemer (1999, Experiment 3) directly investigated the role of S–R learning – however, their findings led them to the conclusion that S–R learning does not mediate implicit sequence learning: They trained participants on an SRTT with an incompatible S–R mapping (response locations were shifted one position to the right). During a subsequent testing phase, participants were switched to a compatible S–R mapping, and either the sequence of stimuli or the sequence of responses remained intact. They found transfer of sequence knowledge if the response sequence, but not if the stimulus sequence remained intact. They concluded that only response-based learning had happened; because stimulus-response pairs changed from training to transfer, they also interpreted this finding as evidence against learning of S–R associations. However, Schwarb and Schumacher (2010) argued that this finding might still be consistent with the view that S–R associations were learned: They argued that shifting responses one position to the right is a rather simple transformation that has to be applied to S–R associations, and such a simple transformation of the acquired S–R associations is sufficient to still use S–R knowledge during transfer. Consistent with this reasoning, they replicated the original transfer effect of Willingham (1999) using the original incompatible S–R mapping that allowed for a simple transformation; importantly, using another incompatible S–R mapping, that would have required a far more complex transformation, eliminated transfer of sequence knowledge.

Studies that followed the second approach of manipulating the difficulty of response selection during the SRTT were at least in part motivated by a finding of Koch and Hoffmann (2000b): In a study investigating the role of spatial information in sequence learning, they orthogonally manipulated stimulus- and response sequences, which required to use either compatible or incompatible stimulus-response (S–R) mappings. In their Experiment 3, they found that incompatible S–R mappings had a beneficial effect on sequence learning. Deroost and Soetens (2006b) argued that such manipulations of S–R compatibility have a selective influence on response-selection processes: more demanding S–R mappings lead to a more controlled response-selection process that benefits learning (cf., Kornblum, Hasbroucq, & Osman, 1990). In their study, they indeed found better learning of a sequence when participants performed the SRTT with an incompatible compared to a compatible S–R mapping. In contrast, Hoffmann and Koch (1997) found that manipulating S–R compatibility (high: four locations, low: four symbols in one location) only affected general RT levels, but not learning scores. Moreover, Kinder, Rolfs, and Kliegl (2008) minimized the necessary processing at the response-selection stage by using saccadic eye movements as responses to stimulus locations. Assuming that saccadic eye movements are highly overlearned responses to stimulus locations, there is not much to learn from a response-selection point of view; in contrast, Kinder et al. (2008) found robust sequence learning even under such conditions, and therefore concluded that learning at the response-selection stage is not necessary for sequence learning.

Some studies that investigated the role of response selection used different approaches: Schwarb and Schumacher (2009) compared the brain areas involved in sequence learning and spatial response selection, finding that both rely on many of the same brain areas. Moreover, Schumacher and Schwarb (2009) found that, under dual-task conditions, sequence learning was disrupted if a secondary task involved selecting a response in parallel; in contrast, sequence learning remained intact if the secondary task did not require selecting a response, or if participants had enough time to serially perform both tasks.

Goschke (1998) demonstrated simultaneous learning of a stimulus and a response sequence. Asserting that parallel processing of two independent sequences at one central processing stage would likely lead to interference, this finding indicates that stimuli and response are at least in part processed independently. Goschke (1998) concluded that learning at the response-selection stage cannot be the only mechanism underlying implicit sequence learning.

Response selection might mediate explicit, but not implicit sequence learning

A possible way to reconcile these contradictory findings was proposed by Abrahamse et al. (2010), who argued that learning at the response-selection stage may be limited to *explicit*, but not *implicit* learning. Evidence pointing into this direction was reported by Abrahamse (2010), who failed to replicate the learning advantage for incompatible S–R mappings reported by Deroost and Soetens (2006b). In contrast to the original study that employed a hybrid first-order conditional sequence, he used a probabilistic second-order conditional sequence, stimulus material that is typically considered to produce robust implicit knowledge but no explicit knowledge (Jiménez & Méndez, 1999; Jiménez et al., 1996). The difference between both studies could thus be explained by an explicit learning mechanism at work in the Deroost and Soetens (2006b) study, but not in the Abrahamse (2010) study.

Abrahamse et al. (2010) argued that this idea converges with findings from studies examining the role of response-effect learning in sequence learning: Response-effect learning is considered to be represented by bidirectional R–S associations (Koch, Keller, & Prinz, 2004), and therefore also involves representations that contain both stimuli and responses. In these studies, explicit sequence knowledge is typically very high, and participants who learn response effects typically show the highest learning scores (Ziessler & Nattkemper, 2001; cf. Abrahamse et al., 2010). Recently, Esser and Haider (2018) found that response effects benefited the acquisition of explicit sequence knowledge. Accordingly, these findings might be interpreted as another indication that representations containing both stimuli and responses benefit the acquisition of explicit sequence knowledge.

In contrast to this reasoning, other authors argued that response-selection processes play a role in *implicit* sequence learning, but are overcome as soon as participants acquire explicit sequence knowledge: In Experiment 2 of Koch (2007), participants responded to symbols (instead of stimulus locations) that were presented either at compatible or incompatible screen locations. Stimulus locations produced a Simon effect (i.e., slower responses for incompatible compared to compatible locations), that diminished over the course of sequence learning in sequenced compared to random blocks. Moreover, participants who were post-hoc classified as having acquired explicit sequence knowledge showed a reduced Simon effect compared to participants who had not acquired explicit sequence knowledge. In his Experiment 3, instructing the sequence to participants prior to the SRTT almost completely eliminated the Simon effect for sequenced materials, while it reappeared in a random block at the end of training. Koch (2007) explained these findings as evidence that the acquisition (or instruction) of explicit sequence knowledge caused a shift from stimulus-based to motor-based action control — he concluded that “forming an explicit representation of the sequence can reduce the influence of stimulus information on response selection”. Similar results were reported by Tubau and López-Moliner (2004), who also observed a reduced Simon effect in post-hoc classified explicit learners. In line with this reasoning, Haider et al. (2011) found that participants who acquired explicit sequence knowledge showed an abrupt increase of response speed during the SRTT; after such an increase in response speed, participants also showed a reduced Stroop-congruency effect. They interpreted these findings as an indication that the acquisition of explicit sequence knowledge induced a switch from stimulus-driven to top-down processing.

Problems of assessing the role of response-selection processes. Studies manipulating S–R compatibility to investigate the role of response selection in implicit sequence learning all suffer from the same epistemological problem: manipulating S–R compatibility might change the mode of processing during the SRTT, emphasizing response-selection processes, and thereby bolstering the influence of these. Therefore, experimental designs manipulating S–R compatibility might not be suitable to study the processes that are *typically* involved in a standard SRTT (Clegg, 2005). Dual-task designs may suffer from a similar problem, as it has been noted that participants tend to try to integrate both tasks, thereby changing the involved processes and/or representations (the *task integration hypothesis*, Schmidtke & Heuer, 1997; see also, Röttger, Haider, Zhao, & Gaschler, 2017).

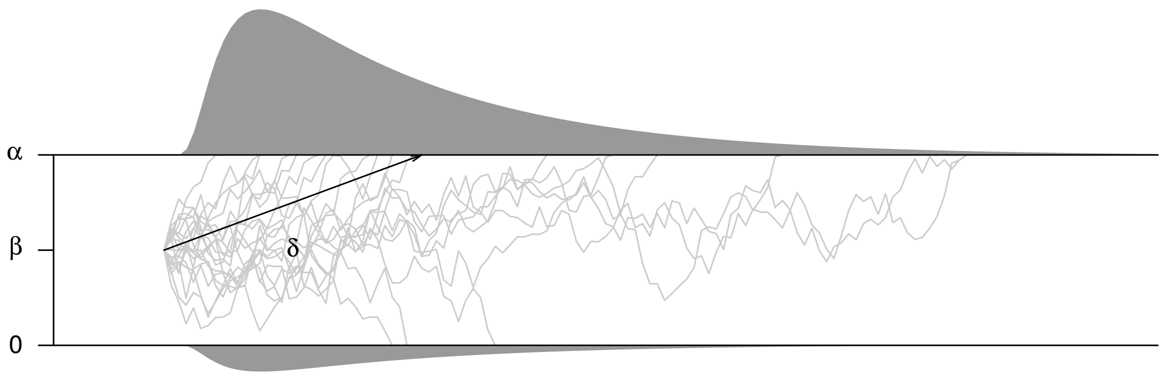

Figure 22. The diffusion model. On each trial, the decision process (depicted as grey lines) begins at a starting point that is determined by parameter β. The spread of the thresholds is determined by parameter α. Evidence is accumulated in a random-walk fashion. Whenever one of the two thresholds is reached, a decision is made. The average rate of evidence accumulation is determined by parameter δ. The decision process is preceded by stimulus encoding and succeeded by response execution, the duration of both processes is captured by nondecision time τ. This basic diffusion model may be extended by response-execution bias ξ and intertrial variabilites for each of the core parameters.

Voss et al. (2013b) addressed a similar problem in semantic and categorical priming: A common strategy to investigate the underlying processes of both phenomena has been to manipulate the task that has to be performed on the targets (e.g., pronunciation instead of lexical-decision task). However, such manipulations of the task that are aimed at controlling or eliminating the influence of specific processes might not only affect the processes they are targeted at, but might also change the processes that are elicited by the stimuli. Therefore, Voss et al. (2013b) proposed a new way of dissociating the effects of stimulus encoding, response selection, and response execution by combining experimental manipulations with a diffusion-model analysis. They found that semantic priming was largely driven by response-selection processes, while categorical priming effects were driven by response competition (i.e., response conflict or facilitation) at the response-execution stage.

In the present study, we apply this rationale to the SRTT. In the next section, we will introduce the diffusion model, and show how it can be applied to the SRTT.

The Diffusion Model

The diffusion model (Ratcliff, 1978; Ratcliff & McKoon, 2008; see also Wagenmakers, 2009; Voss et al., 2013a) is a model of the cognitive processes involved in simple decision tasks, that has been applied to a wide array of applications in cognitive psychology. It disentangles the decision process (i.e., response selection) from other nondecision processes like stimulus encoding and response execution.

Figure 22 illustrates the diffusion model with its core parameters. The diffusion model assumes that during the decision process, evidence continuously accumulates until one of two thresholds is reached. The average *rate of evidence accumulation* is the drift rate δ, which is driven by the quality of information available from a stimulus. *Decision caution* is captured by boundary separation α, the spread of the two thresholds that are reached by the decision process. *Decision tendencies* are captured by bias parameter β, which is the starting point of the decision process relative to both thresholds. Nondecisional processes, such as *stimulus encoding and response execution*, are captured by nondecision time τ. Voss, Voss, and Klauer (2010) extended the diffusion with an additional parameter, response-execution bias ξ, that captures differences in response-execution times for responses that correspond to the upper or the lower threshold: It can, for example, capture differences in execution time for left or right hands (if one response is mapped to the left hand, the other response to the right hand). Moreover, in a variety of experimental paradigms, response competition (i.e., response facilitation and/or conflict), resulting in different response-execution times for both responses, are discussed as possible explanations for response-time effects; if such differences are present in an experiment but are not accounted for, the default diffusion model is unable to capture the differences adequately. Instead, incorporating response-execution bias into the model allows for testing contrasting theoretical accounts of a phenomenon. Voss et al. (2013b) used response-execution bias ξ to dissociate response conflict/response facilitation from effects on response-selection, measured by changes in drift rate.

Applying the diffusion model to the SRTT

Importantly, nondecision time τ and response-execution bias ξ are conceptually distinct from response-selection processes, which are mapped onto the diffusion process. If sequence learning is indeed mediated by response-selection processes, it should be indicated by changes in the parameters of the diffusion process (i.e., drift rate, bias, or boundary separation); if, instead, sequence learning relies on noncentral processing stages (stimulus encoding or response execution), learning should be indicated by changes in nondecision time τ and/or response-execution bias ξ. If implicit sequence learning is response-based, responses that adhere to the sequence regularity should be facilitated by reactivation of the rule-adhering response; if a response is selected that does not adhere to the regularity, conflict with the pre-activated (rule-adhering) response might occur. If implicit sequence learning is stimulus-based, stimuli that adhere to the regularity should be encoded more easily; if a stimulus is presented that does not adhere to the regularity, it should be encoded less easily.

An important difference between the priming studies of Voss et al. (2013b) and a typical SRTT is that, in an SRTT, there are usually more than two response options. Therefore, it is necessary to use *accuracy coding*, with correct responses mapped to the upper threshold, and error responses mapped to the lower threshold of the diffusion process. This allows us to further distinguish the effects of regular vs. nonregular *stimuli* and of regular vs. nonregular *responses*: If a *stimulus* is presented in a regular location, an upper-threshold (i.e., correct) *response* is also regular, leading to response facilitation; a lower-threshold (i.e., error) response is not regular, leading to response conflict. If, instead, a stimulus is presented in a nonregular location, an upper-threshold (correct) response is also nonregular, thereby possibly leading to response conflict; a lower-threshold (error) response that follows the regularity would result in response facilitation, a lower-threshold response that does not follow the regularity would result in response conflict. That is, conflict or facilitation at the response level depends on whether the response that is chosen by the decision process is regular or not. In contrast, conflict or facilitation at the level of stimulus encoding does not rely on the response that is chosen by the decision process, but only depends on the stimulus. Therefore, if accuracy coding with more than two response options is used, conflict/facilitation trials vs. upper-/lower-threshold trials are deconfounded, and it is possible to disentangle effects of stimulus encoding and response execution. As a result, an effect of *stimulus regularity* indicates conflict/facilitation at stimulus encoding, while an effect of *response regularity* indicates conflict/facilitation at response execution. These factors are implemented by nondecision time being mapped onto *stimulus regularity*, and response-execution bias being mapped onto *response regularity*. Note, however, that both factors are highly correlated if participants commit relatively few errors; therefore, it may be difficult to distinguish between the effects of these two factors, empirically.

Overview of the present studies. In the present study, we applied a diffusion model analysis to the data of two SRTT experiments: In Experiment 1 (published in Barth, Stahl, & Haider, 2018), we aimed at testing the hypothesis that learning at the response-selection stage mediates implicit sequence learning. In addition, we aimed to assess possible additional influences of stimulus-based and response-based learning, allowing for the possibility that multiple processing stages are influenced by learning of the sequence. To exclude the possibility that the data were contaminated by explicit sequence knowledge, we assessed participants' sequence awareness after the SRTT using a post-experimental interview. However, given the criticism of such measures of sequence awareness, it may be imprudent to conclude that participants did not acquire explicit sequence knowledge.

Therefore, in Experiment 2, we applied the same diffusion-model analysis to data from another SRTT experiment where participants showed robust explicit sequence knowledge. If the analyses of Experiments 1 and 2 showed qualitatively different patterns of results, it may be warranted to conclude that participants had also acquired qualitatively different types of sequence knowledge. However, if the patterns of results only differ by effect size, that would be an indication that the acquired sequence knowledge is qualitatively the same. Moreover, if the acquisition of explicit sequence knowledge is indeed accompanied by a shift from stimulus-based to plan-based action control as suggested by Koch (2007) and Haider et al. (2011), this should reduce or eliminate effects on response selection.

Experiment 1

In Experiment 1, participants worked on an SRTT that employed a probabilistic sequence of stimulus locations, stimulus material that is typically considered to generate robust implicit, but no explicit knowledge. If learning at the stage of response selection mediates this type of learning, this should be indicated by higher drift rates for regular compared to nonregular transitions. Learning at noncentral processing stages should be indicated by differences in nondecision time (for learning at the stage of stimulus encoding), or effects on response-execution bias (for response competition attributable to learning).

Method

Design. The study realized an 8 (*Block number*) $\times$ 2 (*FOC transition status*: regular vs. nonregular) design with repeated measures on both factors.

Participants. One hundred and forty-seven participants (113 women) aged between 17 and 55 years (M = 23.7 years) completed the study. Most were undergraduates from Heinrich-Heine-Universität Düsseldorf. They received either course credit or 3.50 Euro for their participation.[23]

Materials. A *probabilistic* sequence was generated from the first-order conditional (*FOC*) sequence $2-6-5-3-4-1$. With a probability of .6, a stimulus location was followed by the next location from this sequence; otherwise, another stimulus location was randomly chosen from a uniform distribution. There were no direct repetitions of response locations.

Procedure. The experiment consisted of three consecutive parts: Participants first worked on an SRTT, followed by a generation task and, finally, a debriefing phase. Participants performed an SRTT consisting of eight blocks with 144 trials each (for a total of 1,152 responses). SRTT and generation task were run on 17" TFT monitors (with a screen resolution of 1,024 px $\times$ 768 px). The viewing distance was approximately 60 cm. A horizontal sequence of six white squares (56 px) was presented on a gray screen. The distance between squares was 112 px. Each screen location corresponded to a key on a

[23]The present research used procedures that are exempt from mandatory formal ethical approval under the ethical guidelines of the Deutsche Gesellschaft für Psychologie.

QWERTZ keyboard (from left to right Y, X, C, B, N, M). Participants had to respond whenever a square's color changed from white to red by pressing the corresponding key. They were instructed to place the left ring-, middle- and index fingers on the keys Y, X and C. The right index-, middle- and ring fingers were to be placed on keys B, N and M. There was no time limit for responses in the SRTT (nor in the generation task). A warning beep indicated an incorrect response. The response-stimulus interval (RSI) was 250 ms; there were no pauses within a single learning block.

Following the SRTT, participants were told that stimulus locations had followed an underlying sequential structure (but were not informed about the exact sequence). They were then asked to try to generate a short sequence of six locations that followed this structure. The generation task followed, consisting of two generation blocks involving either inclusion or exclusion instructions. Results from the generation task will not be analyzed here, but see Barth, Stahl, and Haider (2018, Exp. 1), for details. Importantly, some participants received additional explicit sequence knowledge during this phase of the study. Upon completing the computerized task, participants were asked to complete a questionnaire containing the following items (translated from German): (1) "One of the tasks mentioned a sequence in which the squares lit up during the first part of the study. In one of the experimental conditions, the squares did indeed follow a specific sequence. Do you think you were in this condition or not?", (2) "How confident are you (in %)?", and (3) "Can you describe the sequence in detail?". Subsequently, participants were asked to indicate, for each of the six response keys, the next key in the sequence on a printed keyboard layout and to indicate how confident they were in this decision. Finally, participants were thanked and debriefed.

Data analysis. All analyses were performed using the R software[24] and JAGS (Plummer, 2015). For repeated-measures ANOVAs, Greenhouse-Geisser-corrected degrees of freedom are reported.

For the diffusion model analyses, reaction times t_n and accuracy-coded responses y_n (correct responses were coded as upper threshold responses, incorrect responses were coded as lower threshold response) on trial n were modeled as a function of a Wiener process,

$$(t_n, y_n) \sim \mathit{Wiener}\left(\alpha_{ijk}, \beta_{ijk}, \delta_n, \tau_n\right)$$

where boundary separation α and decision bias β varied as a function of participant i, block number j, and transition status k. Drift rate δ_n and nondecision time τ_n were estimated trialwise.

[24]We used R (Version 3.6.1; R Core Team, 2018) and the R-packages *afex* (Version 0.24.1; Singmann et al., 2018), *papaja* (Version 0.1.0.9842; Aust & Barth, 2018), and *runjags* (Version 2.0.4.4; Denwood, 2016).

Boundary separation α of participant i for block j and FOC transition status k was modeled by

$$\alpha_{ijk} \sim \mathcal{N}_{\mathcal{I}(\prime,\infty)}\left(\mu_{jk}^{(\alpha)}, \sigma^{(\alpha)}\right)$$

Decision bias β of participant i for block j and FOC transition status k was modeled by

$$\Phi^{-1}\left(\beta_{ijk}\right) \sim \mathcal{N}\left(\mu_{jk}^{(\beta)}, \sigma^{(\beta)}\right)$$

where Φ^{-1} denotes the inverse of the standard normal cumulative distribution function (i.e., the probit).

Drift rate was estimated for each trial n, where

$$\delta_n \sim \mathcal{N}\left(\nu_{ijk}, \eta_{ijk}\right)$$

The intertrial variability of drift rate η was modeled by

$$\eta_{ijk} \sim \Gamma\left(\mu_{jk}^{(\eta)}, \sigma^{(\eta)}\right)$$

The average drift rate for participant i for block j and FOC transition status k was modeled by

$$\nu_{ijk} \sim \mathcal{N}\left(\mu_{jk}^{(\nu)}, \sigma^{(\nu)}\right)$$

Nondecision time was estimated for each trial n,

$$\tau_n \sim \mathcal{N}_{\mathcal{I}(.1,.9)}\left(\theta_{ijk} + b\xi_{ij}, \chi_i\right)$$

, where $b = 1$ for response-conflict trials and $b = -1$ for response-facilitation trials. Note that, given this implementation, for each participants i within each block j, four different means of nondecision times are estimated: One for regular stimuli that were followed by a regular response (i.e., encoding *and* response facilitation), one for regular stimuli that were followed by a nonregular response (i.e., encoding facilitation, but response conflict), one for nonregular stimuli followed by a nonregular response (i.e., encoding "conflict" and response conflict), and one for nonregular stimuli followed by a regular response (i.e., *encoding* "conflict", but *response* facilitation).

The intertrial variability of nondecision time χ was given by

$$\chi_i \sim \mathcal{N}_{\mathcal{I}(0,\infty)}\left(\mu^{(\chi)}, \sigma^{(\chi)}\right)$$

The stimulus-specific nondecision time for participant i, block j, and *FOC transition status*

k was given by

$$\theta_{ijk} \sim \mathcal{N}\left(\mu_{jk}^{(\theta)}, \sigma^{(\theta)}\right)$$

Response conflict ξ was modeled by

$$\xi_{ij} \sim \mathcal{N}\left(\mu_{j}^{(\xi)}, \sigma^{(\xi)}\right)$$

The main effect of *FOC transition status* is then given by $\zeta^{(\alpha)} = \frac{1}{J}\sum_{j=1}^{J}\Delta\mu_j^{(\alpha)}$.

To assess an interaction of *block number* and *FOC transition status*, we calculated linear trends for the blockwise differences of parameter means

$$\psi^{(\alpha)} = \left(\sum_{j=1}^{J} c_j\right)^{-1} \sum_{j=1}^{J} c_j (\mu_{j,\mathrm{regular}}^{(\alpha)} - \mu_{j,\mathrm{nonregular}}^{(\alpha)})$$

where $c = (-7, -5, -3, -1, 1, 3, 5, 7)$.

For post-hoc comparisons, we also computed, within each of the learning blocks, the posterior differences between the estimated condition means for regular compared to nonregular transitions. If the posterior of these difference parameters does not contain zero, this can be interpreted in favor of an effect of *FOC transition status* on the respective parameter.

Results

We first analyzed reaction times and error rates to determine whether sequence learning occurred. We then assessed whether sequence knowledge remained implicit by analyzing the forced-choice data from the postexperimental interview. Finally, we analyzed reaction times and responses with a hierarchical Bayesian diffusion model.

Reaction times. Figure 23 shows reaction times from Experiment 1. We conducted a 8 (*Block number*) × 2 (*FOC transition status*: regular vs. nonregular) ANOVA that revealed a main effect of *block number*, $F(4.09, 597.64) = 63.23$, $MSE = 2,655.75$, $p < .001$, $\hat{\eta}^2_G = .043$, RTs decreased over blocks; a main effect of *FOC transition status*, $F(1, 146) = 629.78$, $MSE = 1,220.90$, $p < .001$, $\hat{\eta}^2_G = .047$, RTs were shorter for regular compared to nonregular transitions; and an interaction of *block number* and *FOC transition status*, $F(6.48, 945.44) = 32.57$, $MSE = 345.99$, $p < .001$, $\hat{\eta}^2_G = .005$, indicating learning of the sequence.

Error rates. Figure 23 shows error rates from Experiment 1. We conducted a 8 (*Block number*) × 2 (*FOC transition status*: regular vs. nonregular) ANOVA that revealed a main effect of *block number*, $F(6.12, 894.24) = 10.17$, $MSE = 8.47$, $p < .001$, $\hat{\eta}^2_G = .018$, error

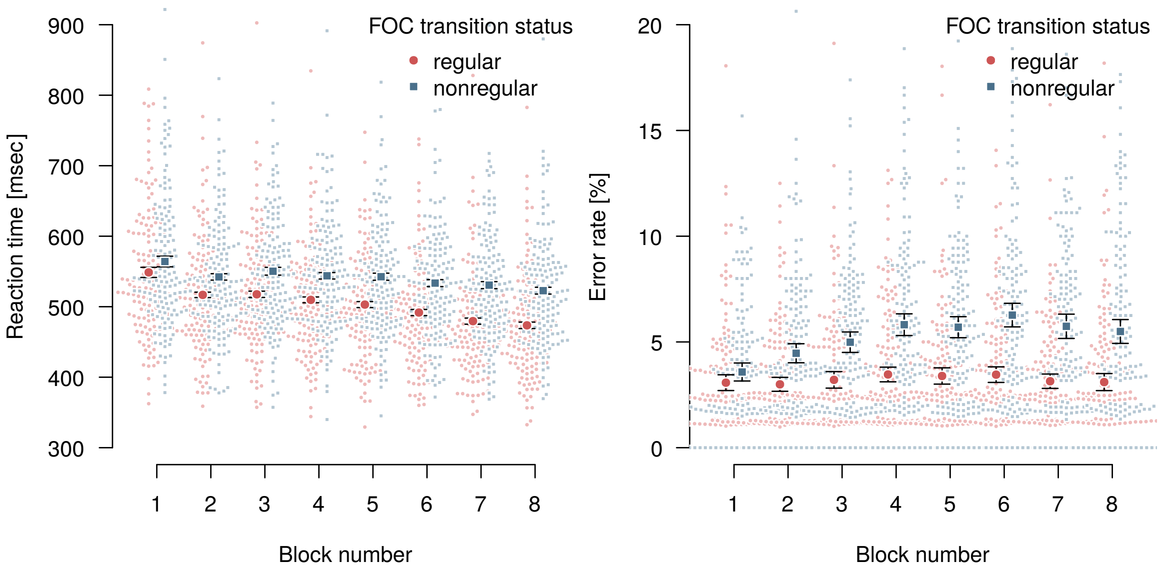

Figure 23. Left panel: RTs from Experiment 1, split by *block number* and *FOC transition status*. Right panel: Error rates from Experiment 1, split by *block number* and *FOC transition status*. Error bars represent 95% within-subjects confidence intervals.

rates increased over blocks; a main effect of *FOC transition status*, $F(1, 146) = 189.18$, $MSE = 12.78$, $p < .001$, $\hat{\eta}^2_G = .078$, more errors were committed for nonregular compared to regular transitions; and an interaction of *block number* and *FOC transition status*, $F(6.63, 967.32) = 6.18$, $MSE = 7.10$, $p < .001$, $\hat{\eta}^2_G = .010$, indicating learning of the sequence.

Explicit sequence knowledge. To assess whether participants acquired explicit sequence knowledge, we analyzed data from the postexperimental interview only of those participants who did not receive further information on the sequence during the generation task: Chance level for this task may be considered to be .25, because there were no direct repetitions in the stimulus materials, all participants were discouraged to generate repetitions, and participants tend to generate reversals at below-chance levels (Stahl et al., 2015). Participants chose the correct location at chance levels, $M = 0.26$, 95% CI $[0.20, \infty]$, $t(28) = 0.34$, $p = .369$ (one-sided). Therefore, we conclude that sequence knowledge remained largely implicit in this experiment.

Diffusion-model analysis. Figure 24 shows parameter estimates from the diffusion model.

Boundary separation α was unaffected by *FOC transition status*, neither the main-effect contrast, $\zeta^{(\alpha)} = 0.03$, 95%HDI$[0.00, 0.05]$, $p = .044$, nor the interaction, $\psi^{(\alpha)} = 0.00$, 95%HDI$[0.00, 0.01]$, $p = .430$, differed from zero. Decision bias β was unaffected by *FOC transition status*, $\zeta^{(\beta)} = 0.01$, 95%HDI$[0.00, 0.03]$, $p = .140$, $\psi^{(\beta)} = 0.00$, 95%HDI$[0.00, 0.00]$,

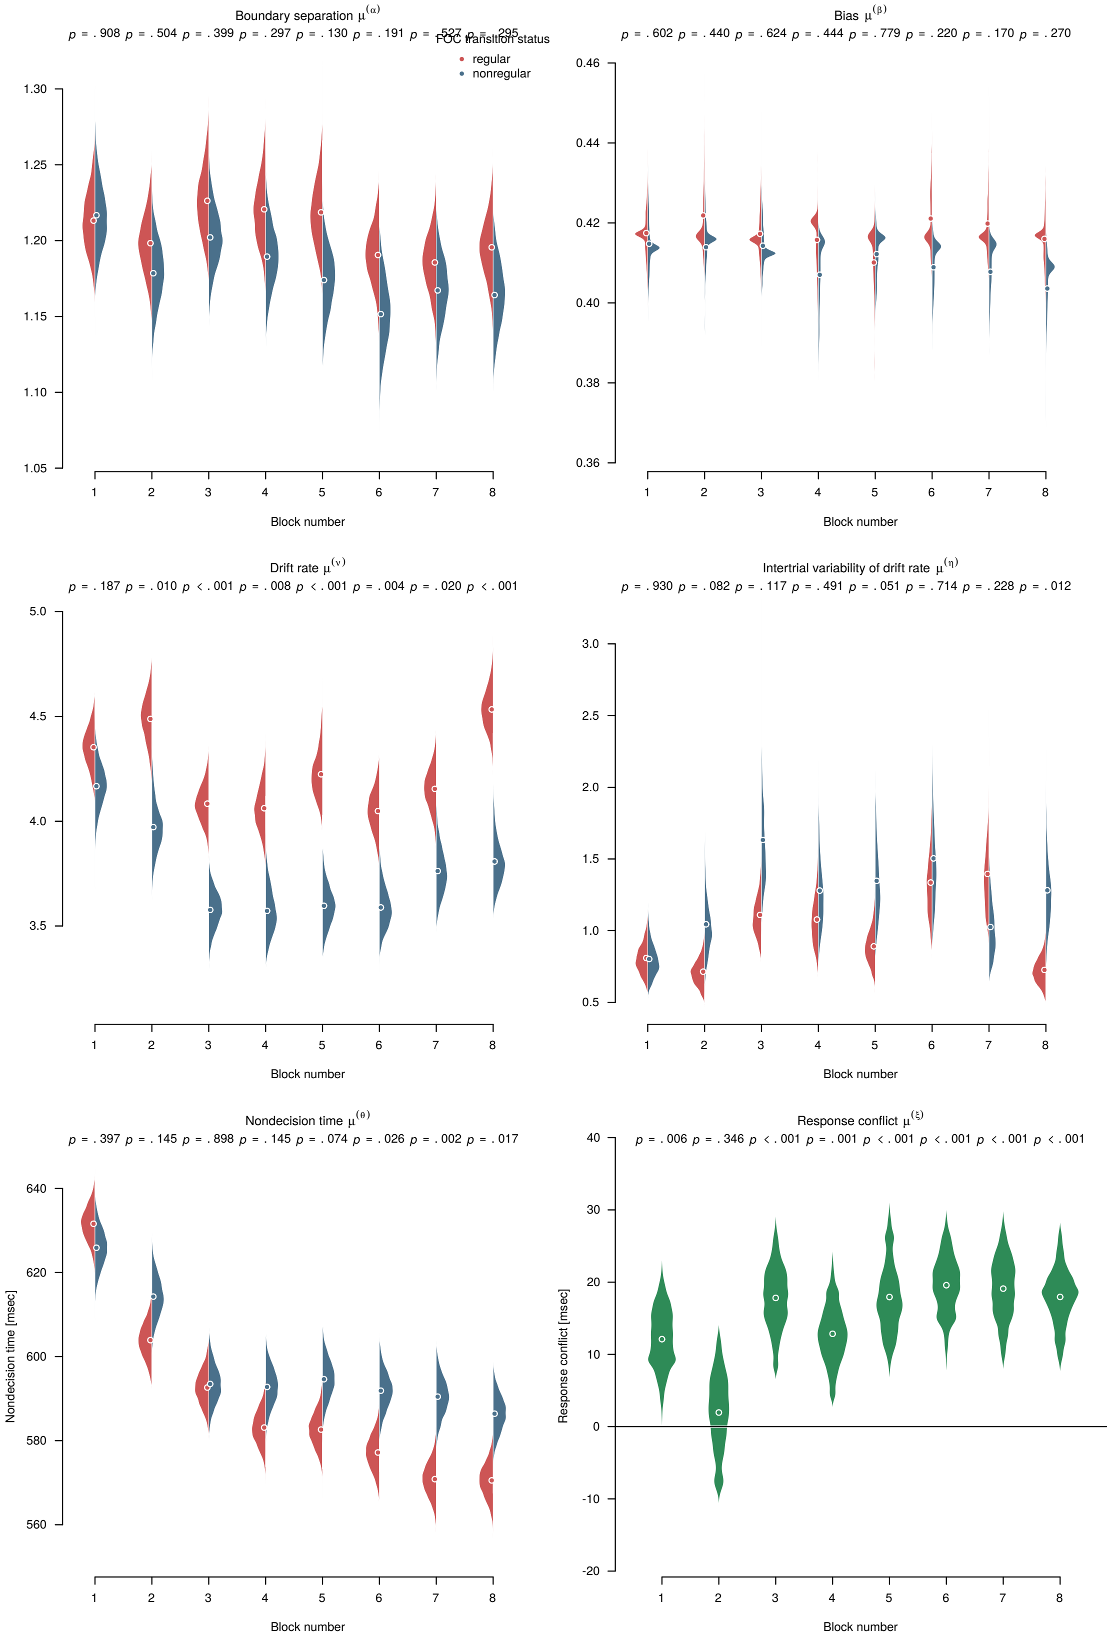

Figure 24. Parameter estimates from Experiment 1. Dots represent posterior means, densities represent posterior densities. Bayesian p values indicate whether the posterior differences of mean parameters in each block contain zero.

$p = .312$.

Drift rate ν varied as a function of *FOC transition status*, $\zeta^{(\nu)} = 0.49$, 95%HDI$[0.34, 0.62]$, $p < .001$, but did not enter an interaction with *block number*, $\psi^{(\nu)} = 0.02$, 95%HDI$[0.00, 0.04]$, $p = .109$. Post-hoc analyses of blockwise differences showed increased drift rates for regular compared to nonregular transitions for all except the first learning block;

Nondecision time θ was lower for regular compared to nonregular transitions, $\zeta^{(\theta)} = -9.67$, 95%HDI$[-14.49, -5.15]$, $p < .001$, indicating faster stimulus encoding for regular transitions, and also entered an interaction with *block number*, $\psi^{(\theta)} = -1.44$, 95%HDI$[-2.41, -0.38]$, $p = .006$, indicating that this effect increased over learning blocks.

Response competition ξ was greater than zero, $\zeta^{(\xi)} = 14.91$, 95%HDI$[11.22, 18.30]$, $p < .001$, but did not enter an interaction, $\psi^{(\xi)} = 0.82$, 95%HDI$[0.15, 1.40]$, $p = .020$. Analyzed separately for each block, response competition was always above zero descriptively, with one-sided Bayesian p values varying around conventional levels such as .05.

Discussion

In Experiment 1, participants worked on an SRTT where stimuli followed a probabilistic 6-item sequence, stimulus material that is considered to produce implicit, bot not explicit sequence knowledge. Separate analyses of response times and error rates provided evidence for robust sequence learning. In a postexperimental interview, participants were not able to recollect any explicit sequence knowledge above chance, indicating that the acquired sequence knowledge remained implicit.

Analyzing response times and response identities jointly with a diffusion model, we found that sequence learning was accompanied by changes in drift rate, nondecision time, and response-execution bias, while boundary separation and decision bias remained unaffected. Higher drift rates for regular compared to nonregular transitions indicate that response-selection processes partially mediated sequence learning. These differences in drift rate were already present during the second SRTT block. Above-zero response-execution bias indicates that facilitation and/or conflict at the response-execution stage also mediated performance effects in the present experiment. This effect also appeared early in training, and (with the exception of the second block) was relatively constant across the whole training. Differences in nondecision time between regular and nonregular stimuli increased over training. The interpretation of this effect is not clear-cut: it might either indicate facilitation or conflict at the processing-stage of stimulus encoding, or might be a spillover effect from response competition (because we used effect coding for response-execution bias $\mu^{(\xi)}$ and thus assumed that response facilitation and conflict are of equal size).

To summarize, Experiment 1 indicated that implicit sequence learning is mediated by both response-selection processes, but also noncentral processes (stimulus encoding and/or response execution). Both findings converge with the idea that it is not a single type of association that is formed during implicit sequence learning, but that the acquired representations contain information about both stimuli and responses.

Experiment 2

The aim of Experiment 2 was to explore qualitative changes in patterns of diffusion-model parameters that are associated with the acquisition of explicit, compared to implicit, sequence knowledge. For this purpose, we reanalyzed a data set that was originally collected for the same project as Experiment 1; it remained unpublished, because participants had acquired significant amounts of explicit sequence knowledge during the SRTT.

Method

Design. The study realized an 8 (*Block number*) × 2 (*FOC transition status*: regular vs. nonregular) × 3 (*material*: sequenced, permuted, random) design with repeated measures on the first two factors. Only results from the sequenced group will be reported here.

Participants. One hundred and forty-six participants (105 women) aged between 17 and 48 years ($M = 24$ years) completed the study. Most were undergraduates from University of Cologne. Participants were randomly assigned to experimental conditions. They received either course credit or 3.50 Euro for their participation.

Materials and Procedure. The experimental procedure was identical to the procedure of Experiment 1, except for the following changes: Participants responded either to a *random* sequence of stimulus locations, or to a mixed-*deterministic* sequence, where *runs* of 15 or 22 stimulus locations followed a deterministic 6-item first-order conditional sequence (individually generated for each participant), that was interspersed with runs of pseudorandom transitions. Similar to Experiment 1, approximately 60% of transitions followed the deterministic sequence. There were neither direct repetitions nor reversals present in the stimulus material. Stimuli were presented on 17" CRT monitors running at 100 Hz.

Results

We first analyzed reaction times and error rates to determine whether sequence learning occurred. We then assessed whether participants showed explicit sequence knowledge in the

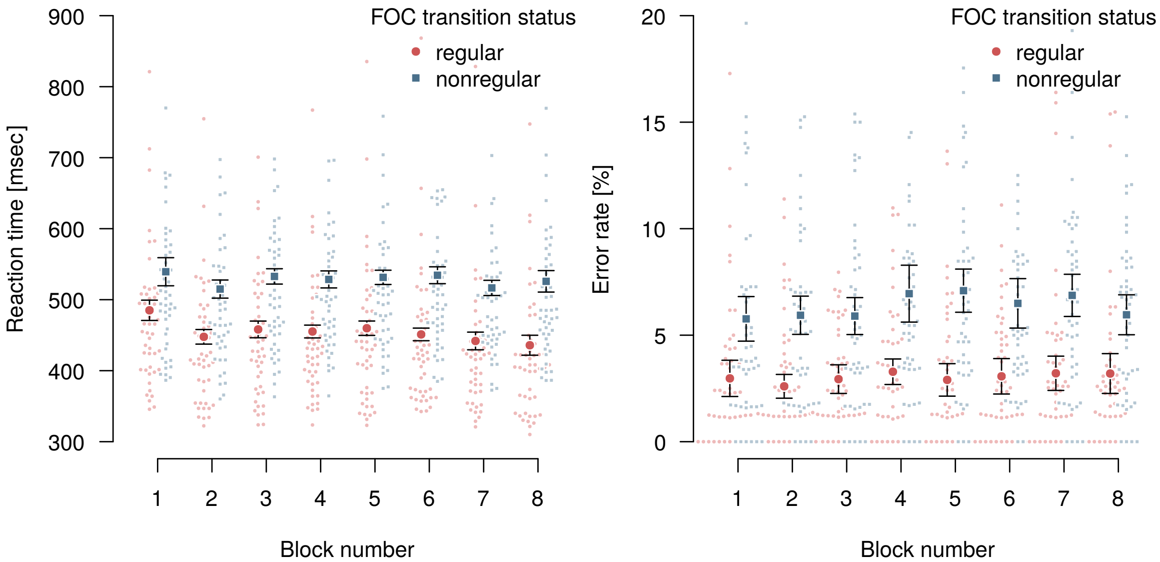

Figure 25. Left panel: RTs from Experiment 2, split by *block number* and *FOC transition status*. Right panel: Error rates from Experiment 2, split by *block number*, *FOC transition status*. Error bars represent 95% within-subjects confidence intervals.

forced-choice data of the postexperimental interview. Finally, we analyzed reaction times and responses with a hierarchical Bayesian diffusion model.

Reaction times. Figure 25 shows reaction times from Experiment 2. An 8 (*Block number*) $\times$ 2 (*FOC transition status*: regular vs. nonregular) ANOVA revealed a main effect of *block number*, $F(4.04, 194.09) = 4.82$, $MSE = 4,171.31$, $p = .001$, $\hat{\eta}^2_G = .006$, RTs decreased over blocks; a main effect of *FOC transition status*, $F(1, 48) = 144.32$, $MSE = 7,368.31$, $p < .001$, $\hat{\eta}^2_G = .073$ RTs were shorter for regular compared to nonregular transitions; and an interaction of *block number* and *FOC transition status*, $F(6.19, 297.21) = 5.13$, $MSE = 603.89$, $p < .001$, $\hat{\eta}^2_G = .001$, indicating learning of the sequence.

Error rates. Figure 25 shows error rates from Experiment 2. We conducted an 8 (*Block number*) $\times$ 2 (*FOC transition status*: regular vs. nonregular) ANOVA that revealed a main effect of *FOC transition status*, $F(1, 48) = 84.71$, $MSE = 25.94$, $p < .001$, $\hat{\eta}^2_G = .143$, participants committed less errors on regular compared to nonregular transitions; and an interaction of *block number* and *FOC transition status*, $F(5.55, 266.16) = 0.74$, $MSE = 10.19$, $p = .604$, $\hat{\eta}^2_G = .003$, the effect of *FOC transition status* increased over blocks, indicating learning of the sequence. The main effect of *block number* was not significant, $F(5.11, 245.16) = 1.10$, $MSE = 13.45$, $p = .359$, $\hat{\eta}^2_G = .006$.

Explicit knowledge. To assess whether participants had acquired explicit sequence knowledge during the SRTT, we analyzed data from the postexperimental interview only of those participants who did not receive further information on the sequence during the generation task: Chance level for this task may be considered to be .25, because there were neither direct repetitions nor reversals in the stimulus materials; all participants were discouraged to generate repetitions, and participants tend to generate reversals at below-chance levels (Stahl et al., 2015). Participants chose the correct location above chance levels, $M = 0.44$, 95% CI $[0.25, \infty]$, $t(15) = 1.76$, $p = .049$. Therefore, we conclude that participants acquired at least some amounts of explicit sequence knowledge.

Diffusion-model analysis. Figure 26 shows parameter estimates from Experiment 2.

We calculated contrasts ζ for main effects and ψ for interactions as in Experiment 1. Contrasts ζ indicated no main effects on boundary separation α, intertrial variability of drift rate η, and nondecision time θ (all Bayesian ps $\geq .103$).

We found a main effect on drift rate, $\zeta^{(\nu)} = 0.16$, 95%HDI$[0.01, 0.31]$, $p = .027$, indicating faster response selection for regular transitions; however, post-hoc analyses of blockwise differences revealed that this main effect was largely driven by an elevated drift rate for regular transition during the first learning block; afterwards, the difference in drift rate disappeared. We found no linear trend for the differences '$\psi^{(\nu)} = -0.02$, 95%HDI$[-0.05, 0.01]$, $p = .119$.

A similar pattern was found for response-execution bias ξ: Overall, response-execution bias was above zero, $\zeta^{(\xi)} = 5.75$, 95%HDI$[1.41, 11.82]$, $p = .016$. This main effect was qualified by an interaction of *FOC transition status* and *block number*, $\psi^{(\xi)} = -2.61$, 95%HDI$[-3.77, -1.38]$, $p < .001$. Inspecting the blockwise effects, it is obvious that main effect and interaction are largely driven by a pronounced effect in the first block that disappeared over training. As of the second block, response-execution bias was not different from zero.

In contrast, we found a main effect on decision bias β, $\zeta^{(\beta)} = 0.13$, 95%HDI$[0.11, 0.15]$, $p < .001$, and also an interaction with *block number*, $\psi^{(\beta)} = 0.01$, 95%HDI$[0.01, 0.01]$, $p < .001$. Post-hoc comparisons revealed that participants were strongly biased towards regular responses as soon as during the second learning block (all Bayesian ps $< .001$).

Discussion

In Experiment 2, participants worked on an SRTT where stimuli followed a mixed-deterministic 6-item sequence, stimulus material that is considered to more likely produce explicit sequence knowledge, Separate analyses of response times and error rates provided evidence for sequence learning. In a postexperimental interview, participants were able to recollect explicit sequence knowledge above chance, indicating that participants acquired some explicit knowledge about the underlying sequence.

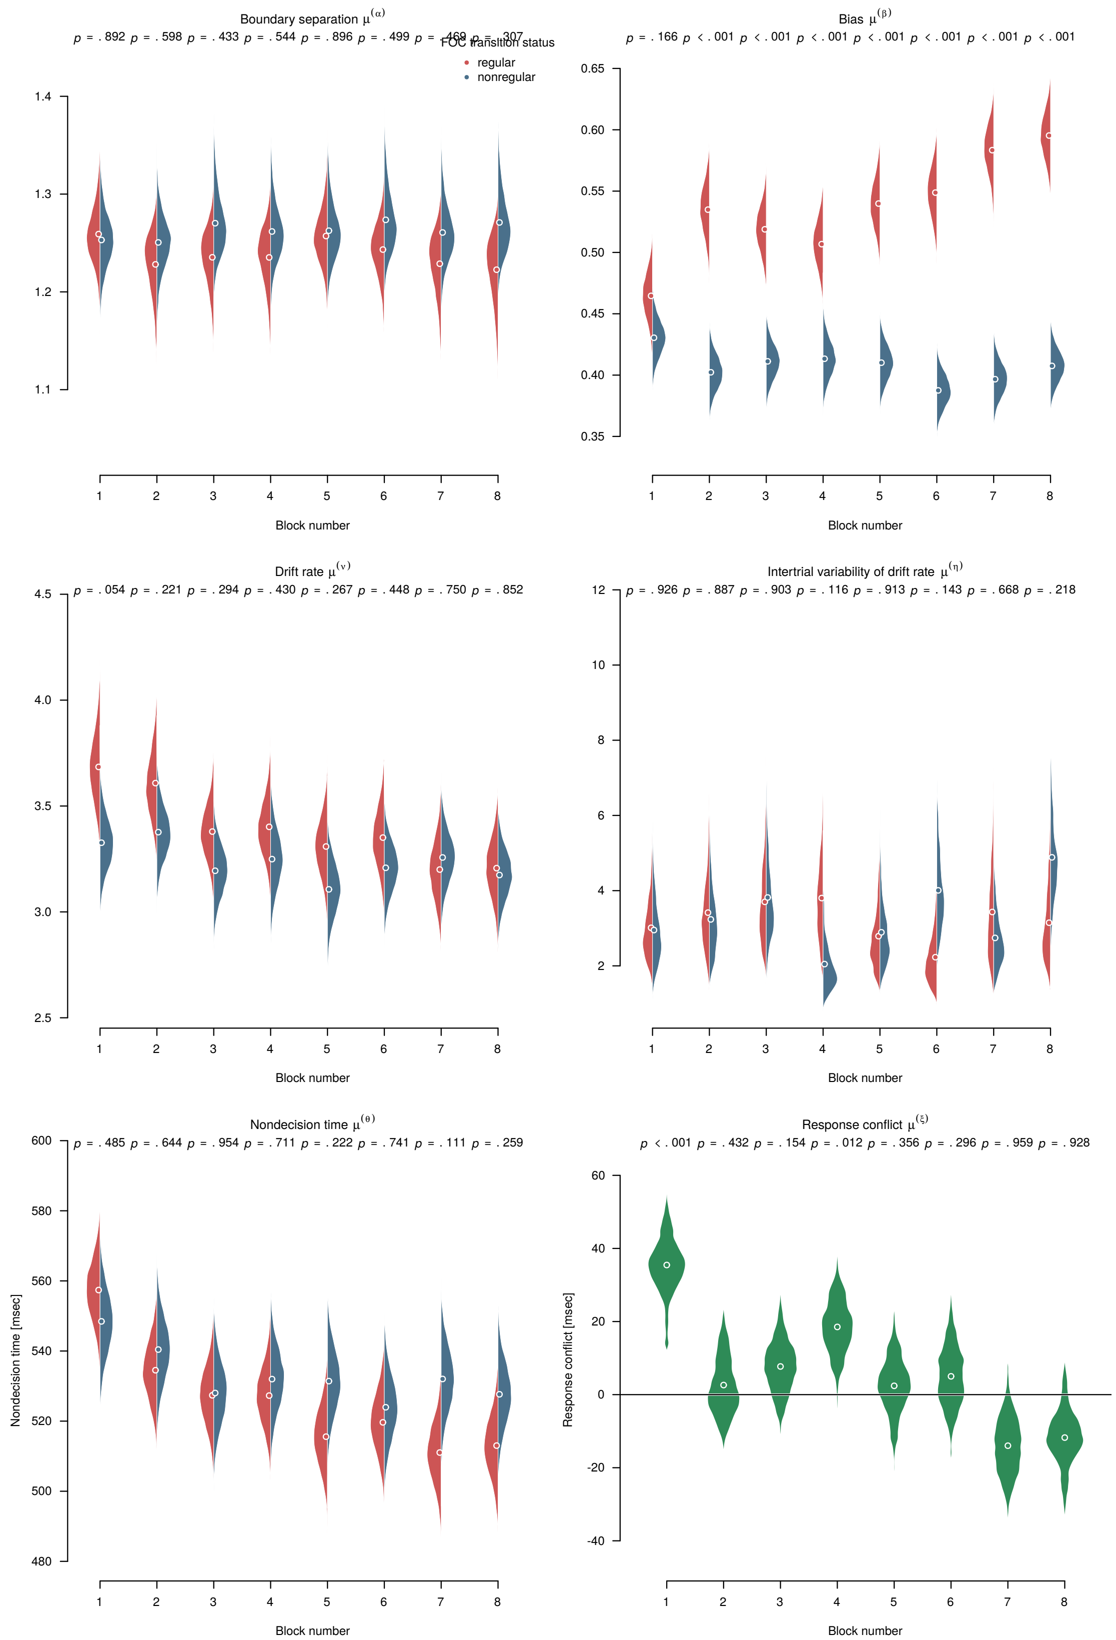

Figure 26. Parameter estimates from Experiment 2. Dots represent posterior means, densities represent posterior densities. Bayesian p values indicate whether the posterior differences of mean parameters in each block contain zero.

Analyzing response times and response identities jointly with a diffusion model, we found that learning was expressed by an increasing decision bias towards regular responses. Moreover, effects on nondecision time and drift rate that we observed in Experiment 1 were not observed, or disappeared early in training. Both findings may indicate a shift from stimulus-based to anticipatory responding, as suggested by Koch (2007). It is plausible that, in the presence of an explicit sequence representation and predictable stimuli, the response-selection (i.e., diffusion) process starts even before the respective stimulus is presented. Such an earlier starting point would likely translate into a decision bias as we observed in this study.

In addition to effects on decision bias, in the first learning block of Experiment 2, the diffusion model indicated a strong response conflict of approximately 40 ms. While we did not predict this result, it is in line with theories that assume that explicit sequence knowledge emerges from conflicts between expected and actually performed behavior (the *Unexpected-Event Hypothesis*, Haider & Frensch, 2009; Rünger & Frensch, 2008). The observed response conflict disappeared after the fist block; moreover, learning at the response-selection stage (as indicated by an effect on drift rate) could be observed in the first block, but also disappeared thereafter. Both findings might be explained by Koch (2007)'s idea that explicit sequence knowledge overrules the expression of implicit sequence learning, which would otherwise be expressed by changes in drift rate and response competition.

General Discussion

We reanalyzed data from two SRTT experiments with a diffusion model. In both experiments, performance data (reaction times and error rates) indicated robust learning of the sequences. In Experiment 1, where participants worked on a probabilistic sequence, sequence knowledge remained largely implicit, as indicated by at-chance generation performance in a postexperimental interview. In Experiment 2, where participants worked on a mixed-deterministic sequence, participants acquired at least partial explicit sequence knowledge, as indicated by above-chance generation performance in a postexperimental interview.

The diffusion-model analysis indicated that the expression of implicit sequence learning was mediated by faster response selection, response competition and faster stimulus encoding. This finding converges with the idea that both stimulus and response information is acquired implicit sequence learning; crucially, a mediating role of response selection indicates that stimulus and response features are not represented independently, but in a common representational code.

In Experiment 2, explicit sequence knowledge was largely expressed by a decision bias towards regular responses. Effects on stimulus encoding and response execution were absent or disappeared during the first few learning blocks. Before turning to the theoretical implications of our findings, we first discuss potential limitations to our study.

Limitations

The diffusion model has been successfully applied to a wide variety of performance tasks in psychology, and the validity of its parameters has been extensively tested in many paradigms using experimental manipulations (e.g., Voss, Rothermund, & Voss, 2004; Boywitt & Rummel, 2012); moreover, its parameters could be successfully related to other behavioral (e.g., eye tracking) and neurophysiological measures (EEG, fMRI, single-cell recordings) (for a review, see Ratcliff et al., 2016).

Its success as a model of speeded response tasks makes it an excellent candidate model for measuring SRTT performance and to disentangle decision processes from nondecisional processes involved. The present study provides a first application of the diffusion model to the SRTT; in order to show that its interpretation is valid in this paradigm, it needs to be applied more frequently, and to be combined with experimental manipulations or related to other measures. Alternative architectures (e.g., linear ballistic accumulators, Donkin et al., 2009) are possible and should be considered; however, given the extant evidence in favor of the diffusion model's validity, and the availability of fitting procedures, we deem it to be a good starting point for further investigation.

Performance data might not indicate differences in learning, but in the expression of knowledge. Frensch, Lin, and Buchner (1998) discussed the fundamental problem that performance data from the SRTT cannot be interpreted directly as evidence for or against effects on learning per se: It is always possible that a manipulation only affects the *expression*, but not the *acquisition*, of sequence knowledge. This problem also applies to the diffusion-model analyses presented in this study, and further work is needed to disentangle the acquisition of sequence knowledge from its expression.

Qualitative differences might be based on predictability of next stimulus, not implicit/explicit distinction. In Experiment 2, we used a deterministic instead of a probabilistic sequence, and expected higher degrees of explicit sequence knowledge if participants worked on deterministic materials. However, it is possible that qualitative differences between the patterns of results in Experiments 1 and 1 cannot be attributed to the distinction of implicit vs. explicit sequence knowledge: instead, the probabilistic materials that we used in Experiment 1 might not have encouraged participants to anticipate the next stimulus or response, because only 60% of trials were indeed regular. In contrast, the mixed-deterministic materials in Experiment 2 might have encouraged a different strategy:

While the rate of regular trials was approximately the same, arranging trials in runs of multiple consecutive regular or nonregular trials provides participants with some additional information: If participants are able to detect differences between runs of regular and nonregular trials (e.g., by detecting changes of fluency), they might exploit that additional information, and extract knowledge about "fluent" trials (e.g., "1–2") and "non-fluent" trials (e.g., "totally random"). If participants acquired such knowledge, this might have encouraged anticipating consecutive stimuli and/or responses with regular runs. Haider et al. (2005) already provided evidence that participants only apply such a strategy if the sequence reliably predicts future events.

Processing stages might be organized in parallel. It has been argued that the processing stages involved in sequence learning might be organized in parallel (Hazeltine & Schumacher, 2016; Schumacher & Hazeltine, 2016; Verwey, 2003). If this is indeed the case, the diffusion model might be unfit to describe these processes, especially if information is shared or exchanged between processes. However, such a misfit between model assumptions and the processes involved in task performance are likely to be detectable by further model-validation efforts.

For sequentially presented stimuli, processing might start even before a stimulus is presented. The diffusion model was originally developed for two-choice speeded responses to single stimuli. Applying the diffusion model to sequentially presented stimuli, it is possible that the diffusion process starts even before the stimulus on a given trial is presented: Evidence accumulation could start, for instance, after the response of the last trial was executed. This would result in a bias towards the rule-adhering response option (cf., Mulder, Wagenmakers, Ratcliff, Boekel, & Forstmann, 2012), but also in a change of drift rate as soon as the stimulus of the current trial is encoded and starts informing the diffusion process. This change of drift rate cannot be measured given the present implementation as a single-step decision process; if one was interested in the drift rate *before* a stimulus was presented, this would require an additional modification of the model with multiple contiguous diffusion processes. However, this problem is not unique to our study, but also applies to the study by Voss et al. (2013b), where primes and targets had a stimulus-onset asynchrony (SOA) of 250 ms.

Implications

Early accounts of implicit sequence learning assumed that it is a single type of simple association that is acquired in implicit learning; however, after more than 30 years of research on the topic that provided equivocal or contradictory results, sometimes providing evidence in favor of a specific type of representation, sometimes against such a representation, it is evident that such a position is untenable. Abrahamse et al. (2010) tried to reconcile these findings by proposing a multi-level account of implicit sequence learning, assuming that at least three types of associations may be formed; what type of representation may depend on task set and attentional processes. Opposing views have been formulated by Hazeltine and Schumacher (2016) and Eberhardt et al. (2017); both accounts assume that the representations that are acquired in implicit sequence learning are more complex than simple associations, and importantly, always contain information about both stimuli and responses.

The present study investigated the processing stages that are involved in the expression of implicit sequence learning, and found a mediating role of both peripheral (stimulus encoding and response execution) and central (response-selection) processes. Because response selection necessarily depends on both stimulus and response information, this finding provides evidence that the representations acquired in implicit learning contain both stimulus and response features. The experimental designs employed here do not allow us to further distinguish between S–R associations, R–S associations, task files, or abstract feature codes. However, combined with experimental manipulations, the diffusion model proposed here may be used in future research to distinguish between these opposing views, which will be elaborated below. Importantly, this is the first study to demonstrate a mediating role of response selection that did not manipulate the difficulty of response selection by changing S–R mappings or adding a secondary task to the SRTT — both manipulations that haven been criticized for possibly changing the processes involved in performing the SRTT.

In addition to an involvement of response selection, we also found effects of learning on stimulus encoding and response competition. At least two possible explanations for these effects are plausible: These processes might simply be able to access the same representations that underlie response-selection effects; alternatively, it is possible that these effects are indeed mediated by simple R–R and S–S associations. Such independent (i.e., encapsulated) learning mechanisms are compatible with both Abrahamse et al. (2010)'s multiple-level view, and also the multiple-systems model proposed by Keele et al. (2003).

Explicit sequence knowledge was largely expressed by a decision bias towards regular responses. This finding might be explained by a change from stimulus-based to plan-based action control, as it has been suggested by Koch (2007) and Haider et al. (2011). Consider a participant who acquired explicit sequence knowledge and also an intuition that some consecutive trials are highly predictable, while other consecutive trials are not. It is plausible that under such conditions, participants already prepare a decision for the regular response option.

Interpreting the diffusion-model parameters verbatim, our results indicate that it is not a *response* that is prepared, which would be indicated by response-competition parameter ξ; instead, a *decision* is prepared by adjusting the starting point of the decision process towards the expected threshold. Therefore, this effect is best interpreted as a switch to *plan-based*, not *motor-based*, action control.

Importantly, in Experiment 2, we did not find an involvement of drift rate, response execution bias, or nondecision time in the expression of sequence knowledge. Still, this is in line with Koch (2007), who suggested that the acquisition of explicit sequence knowledge overrules the expression of implicit sequence knowledge. It is not clear whether participants did not acquire implicit knowledge or did not express it.

In our study, we did not find evidence for Abrahamse et al. (2010)'s speculation that it is possibly explicit, but not implicit, sequence learning that is mediated by learning at the response-selection stage. In contrast, we found effects on response selection (indicated by differences in drift rate) in Experiment 1, where participants acquired implicit, but no robust explicit sequence knowledge. In Experiment 2, where we found evidence for participants having acquired at least some explicit sequence knowledge, the effect on drift rate disappeared after the first learning block.

A secondary finding of Experiment 2 was that an initially large effect of response competition was apparent in the first training block, but disappeared over the course of learning. While we did not anticipate this effect, we interpret this finding as evidence in favor of the *Unexpected-Event Hypothesis* (Haider & Frensch, 2009; Rünger & Frensch, 2008), which states that conflicts between predicted and actually performed behavior trigger the search for regularity in stimulus materials, and thereby bolsters the acquisition of explicit sequence knowledge.

Outlook

The diffusion model introduced here may be combined with experimental manipulations to further investigate the representational basis of implicit learning.

For instance, Goschke and Bolte (2012) showed concurrent learning of both a stimulus and an uncorrelated response sequence, a finding that was interpreted as evidence against an involvement of central processes in implicit learning, because it is plausible that such uncorrelated sequences might cause interference at central processing stages. Combining such concurrent learning with a diffusion-model analysis would be informative for at least two reasons: First, it could be tested whether uncorrelated sequences indeed eliminate or hamper the effect of learning at the response-selection stage. Second, if the parameters of stimulus encoding (i.e., nondecision time θ) selectively varied with stimulus regularity, and parameters of response competition (i.e., response-execution bias ξ) selectively varied with response regularity, this provided evidence for encapsulated learning at these noncentral processing stages. If, instead, these noncentral processing stages varied as a function of both sequences, this provided evidence for commonly coded representations that are also accessed by these noncentral processes.

Another interesting finding was provided by Gaschler, Frensch, Cohen, and Wenke (2012), who showed that, even under exactly the same learning conditions, the type of representation that is acquired in an SRTT depends on what type of information participants are instructed to respond to. If the representations that underlie learning effects at noncentral processing stages are encapsulated and independent of instructed task set (as implied by multiple-systems model by Keele et al., 2003), stimulus encoding and response competition should be independent of instructed task set; if both processes were affected by task set, this would again provide evidence for a common representation.

Conclusion

Recent theoretical work on implicit sequence learning proposes that stimulus and response features are represented in a common representational code. To investigate the issue, it is helpful to consider the cognitive processes involves in performing the SRTT: While stimulus encoding and response-execution may be mediated by separate representations of stimulus and response features, learning at the response-selection stage necessarily relies on representations containing both stimulus and response features. Using a diffusion-model analysis, we were able to disentangle the contributions of response selection, stimulus encoding, and response competition in the expression of sequence learning without manipulating S–R mappings or introducing a secondary task, both manipulations that have been criticized for their unpredictable effects on task representation and the processes involved.

We found that the expression of implicit sequence learning is partially mediated by response-selection processes, indicating that the representations acquired in implicit sequence learning contain information about both stimuli and responses. In addition, we found indications of implicit sequence learning at the processing stages of stimulus encoding and response execution — whether these effects rely on the same representations as response selection remains an open question.

The acquisition of explicit sequence knowledge was accompanied by indications of anticipatory responding that eliminated observable effects on other possible processes. Whether or not sequence learning at other processing stages does not occur in the presence of explicit sequence knowledge or is simply not expressed in the performance data remains an unresolved issue.

Diffusion models have been successfully applied to a variety of cognitive tasks, and furthered the understanding of the processes that are involved in these tasks. The present study introduced the diffusion model to the SRTT; combined with experimental manipulations, it may prove a useful tool to disentangle the possibly multiple processes at work in sequence learning, and to test new predictions from the theoretical frameworks discussed.

Chapter V

General Discussion

Using measurement models of task performance, the present studies investigated two questions in the study of sequence learning: Is there evidence for sequence learning that proceeds without awareness, and which processes mediate the expression of implicit and explicit sequence knowledge in the SRTT?

In Chapter II and III, we tested the assumptions of the process-dissociation (PD) approach as applied to the generation task. We found that the assumptions of equal baselines for inclusion and exclusion task (Chapter II), the monotonicity assumption, and the invariance assumption were violated (Chapter III). Each of these violations has the potential of overestimating implicit, and underestimating explicit sequence knowledge.

In Chapter IV, we applied a diffusion model to task performance in the SRTT. We found that the expression of implicit sequence learning was mediated by stimulus-encoding and response-execution, but also response-selection processes. The latter finding is indicative for a common coding of stimulus and response features in implicit sequence learning. In contrast, explicit sequence knowledge was expressed by a decision bias towards the regular response option, indicating that the acquisition of explicit sequence knowledge caused a shift from stimulus-based to plan-based action control.

Implications for theories of sequence learning

Ever since its introduction by Nissen and Bullemer (1987), sequence learning in the SRTT has been considered a key demonstration of learning that can occur incidentally and without awareness. Claims about the implicit nature of the acquired knowledge were largely founded on a simple logic of dissociation that contrasted performance gains in the SRTT with subsequent measures of awareness. However, the logic of dissociation has been heavily criticized, and, for the host of its problematic assumptions, is now considered insufficient for providing evidence for learning that occurs in the absence of awareness.

Destrebecqz and Cleeremans (2001) introduced the process-dissociation procedure to the generation task, and provided evidence for implicit sequence learning without relying on the assumptions of the dissociation logic. Therefore, studies that applied the PD approach to the generation task have been considered the most evidential result in favor of implicit sequence learning. However, the PD approach comes with its own set of critical assumptions. The experiments presented in Chapters II and III tested these assumptions, finding that they are violated in applications to sequence learning. Taking these violations of basic assumptions into account, the evidence for implicit learning has to be reevaluated: If anything, the results reported here support the presence of explicit knowledge and call into question the presence of implicit knowledge. Given the problematic assumptions of other approaches relying on the logic of dissociation, the field of implicit learning is still lacking unequivocal evidence for purely implicit learning.

In Chapter IV, we found that the expression of implicit sequence knowledge is mediated by both noncentral processes (stimulus encoding and response execution), but also central response-selection processes. Early accounts of implicit sequence learning assumed that it is mediated by the formation of a single type of simple associations between consecutive stimulus features, response features, or both. An involvement of response-selection processes in the expression of sequence is incompatible with accounts that assume that only R–R associations, or only S–S associations are formed during implicit sequence learning. Instead, it converges with recent ideas of Schumacher and Hazeltine (2016; see also, Hazeltine and Schumacher, 2016) and Eberhardt et al. (2017; see also, Esser and Haider, n.d.; Haider et al., 2018) who assume that implicit sequence learning is mediated by representations that contain information about both stimulus and response features.

An involvement of noncentral processes in the expression of implicit sequence learning raises the question whether these processes rely on the same commonly coded representations, or depend on different representations that contain information of either stimulus features or response features. Such independent coding would be in line with Keele et al. (2003)'s model that assumes that implicit learning may proceed in highly encapsulated learning modules. The acquisition of explicit sequence knowledge resulted in a shift from stimulus-based to plan-based action control, converging with earlier findings (e.g., Koch, 2007; Tubau, Hommel, & López-Moliner, 2007).

Methodological implications

In Chapters II and III, we found that the assumptions of the PD approach as applied to sequence learning are violated. For the parametric PD model, we found that the invariance assumption for controlled processes was violated with $C_{\text{Inclusion}} > C_{\text{Exclusion}}$, resulting in an overestimation of implicit, and an underestimation of explicit sequence knowledge. In addition, the monotonicity assumption of the ordinal-PD approach was also found to be violated, where an increase in explicit knowledge did not necessarily result in a reduced number of regular transitions generated under exclusion instructions, (i.e., $C_1 > C_2$ did not imply $E(C_1) < E(C_2)$). Both violations have to the potential to cause an overestimation of implicit sequence knowledge, and an underestimation of explicit sequence knowledge.

Our results indicating limited validity of the PD approach and its assumptions might lead one to conclude that measurement models such as the parametric PD model—because of the host of their possibly violated assumptions—are more problematic than standard statistical techniques such as linear models, and the adoption of standard techniques is to be preferred. We consider this conclusion to be false, for the following reasons: In their original study, Destrebecqz and Cleeremans (2001) adopted an analysis strategy that seemingly avoided the problematic assumptions of the parametric PD approach that were already discussed, for instance the independence assumption, by using linear models (i.e., ANOVA) to analyze their data. This step has been applauded even by outspoken critics of the parametric PD model (Curran, 2001). However, the entirety of the assumptions underlying the analysis approach were never made explicit, and not put at a critical test. A possible explanation for this is the high familiarity of psychologists with analysis strategies such as ANOVA, leading to an illusion of depth of understanding (Ylikoski, 2009). In contrast, the introduction of PD model immediately provoked a rich discussion on its assumptions and studies investigating the identified issues (i.e., the independence assumption). This was only possible because the PD model is relatively simple and precise, with the model being expressible in two mathematical equations.

A common research strategy in the sequence learning literature has been to study implicit and explicit learning in isolation, typically by realizing some experimental conditions that make the emergence of explicit knowledge unlikely; the assumption that learning remained implicit for some experimental conditions or a subgroup of participants was then justified by assessing sequence awareness with a separate measure (e.g., verbal report, recognition, or process dissociation). Taking into account both the criticism of other measures of awareness and our findings of limited validity of the PD approach, it is conceivable that in at least some of these experiments, participants may have acquired explicit sequence knowledge that remained undetected. If explicit sequence knowledge also affected SRTT performance, which we found in Chapter IV and was demonstrated by many others (e.g., Haider et al., 2011), the conclusions of these experiments regarding properties of implicit learning may be contaminated by more or less high amounts of acquired explicit knowledge. In order to circumvent erroneous conclusions about characteristics of implicit learning, it is therefore necessary to take into account the possibility that explicit sequence knowledge emerges in the course of training – characteristics of explicit sequence knowledge could otherwise be ascribed to implicit sequence knowledge.

In Chapter IV, we applied a diffusion-model analysis to the SRTT, and found that multiple processes are involved in the expression of implicit and explicit sequence knowledge. Moreover, we were able to test the prediction that stimulus and response features share a common representational basis in implicit sequence learning. Therefore, we believe that future research should more frequently apply such analyses of SRTT performance to further our understanding of both implicit and explicit sequence learning. In the next section, we will provide an outlook on possible research strategies that encompass both implicit and explicit sequence learning.

Outlook

Repetition priming and recognition have long been viewed as relying on distinct learning mechanisms. Berry, Shanks, Speekenbrink, and Henson (2012) used formal models to derive quantitative predictions from a single-system and two multiple-systems theories for both phenomena. By doing so, they were able to provide compelling evidence against one multiple-systems theory, and found that the single-system theory outperformed the remaining multiple-systems theory quantitatively, which was attributable to its relative parsimony. Such a modeling effort may also prove helpful in comparing the single-system and the multiple-systems view in sequence learning. Starns, Ratcliff, and McKoon (2012) criticized that Berry et al. (2012) used an ad-hoc link function between memory output and performance (response times and accuracy); and proposed to use a sequential-sampling model such as the diffusion model to link memory and performance. Taking into account our results from Chapter IV that multiple processes are sensitive to the regularity, it is obvious that using an ad-hoc link function would not be sufficient to capture the complex performance changes observed in the SRTT. Instead, the diffusion model developed in Chapter IV may be used as a link between a model representation of single- vs. multiple systems views and performance data. For the single-systems view, candidate models implementing such a sequential-sampling approach are the exemplar-based random-walk model (EBRW, Cohen & Nosofsky, 2003; Nosofsky & Palmeri, 1997; Palmeri, 1997) or the predictive temporal context model (pTCM, Shankar, Jagadisan, & Howard, 2009).

Haider et al. (2011) investigated whether explicit sequence knowledge in the SRTT emerges in a gradual or an all-or-none fashion using an online measure of response times: Arguing that the acquisition of explicit sequence knowledge should be accompanied by a rapid (i.e., all-or-none) decrease in response times (i.e., an *RT drop*), they calculated, for each participant and item, whether such RT drops occurred using a tailored algorithm. If a participant exhibited RT drops for four (out of six possible) transitions, they were classified as having shown RT drops. After the SRTT, participants who were classified as having shown an RT drop exhibited significantly more explicit sequence knowledge as indicated by performance on a process-dissociation generation task and post-decision wagering (Persaud et al., 2007). Haider et al. (2011) concluded that the acquisition of explicit sequence knowledge is accompanied by RT drops. RT drops could be successfully related to neuropsychological

(EEG and fMRI) data (Rose, Haider, & Büchel, 2010; Wessel, Haider, & Rose, 2012) – importantly, showing that RT drops are temporally preceded by changes in EEG data. The diffusion model developed in Chapter IV opens up new possibilities for scrutinizing and expanding upon this finding: First, a model assuming that sequence awareness is an all-or-none phenomenon can be compared quantitatively and qualitatively to a competitor model assuming a gradual emergence of awareness. Second, it could be tested whether RT drops reflect the same phenomenon such as the anticipatory responding that we observed in the diffusion-model analysis in Chapter IV (i.e., whether RT drops are correlated with changes in decision bias).

Conclusion

Embracing the notion that probably no task is process pure, and almost always multiple processes are involved in performing a task, the present studies utilized measurement models to disentangle the contributions of multiple processes to task performance. In the first two studies, we scrutinized the assumptions underlying the PD approach as applied to sequence learning, and found that its assumptions are typically violated. Taking these violations into account, earlier studies that used the PD approach and found indications of implicit sequence learning have to be reevaluated; given the limited validity of the PD approach, current results do not provide firm evidence for implicit sequence learning.

In a third study, we disentangled the processes involved in the expression of implicit and explicit sequence knowledge in the SRTT. We found that the expression of implicit sequence knowledge is mediated by noncentral processes (stimulus encoding and response execution), but also central response selection. The finding of an involvement of response selection indicates that the representations acquired in implicit sequence learning contain information about both stimulus and response features. Explicit sequence learning was expressed by a decision bias towards the regular response option, indicating that the acquisition of explicit sequence knowledge caused a shift from stimulus-based to plan-based action control.

The prefacing quote by Reingold and Merikle (1990, p. 20) stated that “the lack of definitional and conceptual clarity in the study of the unconscious stems from the implicit or explicit association of certain tasks with characteristics of observers or rememberers such as intentionality or phenomenal awareness”. Measurement models of task performance such as those utilized in the present studies allow the researcher to refrain from identifying tasks with processes by disentangling the contributions of multiple processes involved in task performance. Therefore, their application will hopefully contribute to a further clarification of definitions and concepts in implicit and explicit sequence learning.

References

Abrahamse, E. L. (2010). *Serial Action and Perception.* University of Twente, Enschede. Retrieved from https://ris.utwente.nl/ws/portalfiles/portal/6082584

Abrahamse, E. L., Jiménez, L., Verwey, W. B., & Clegg, B. A. (2010). Representing Serial Action and Perception. *Psychonomic Bulletin & Review*, *17*(5), 603–623. https://doi.org/10.3758/PBR.17.5.603

Albert, J. H., & Chib, S. (1993). Bayesian Analysis of Binary and Polychotomous Response Data. *Journal of the American Statistical Association*, *88*(422), 669–679. https://doi.org/10.1080/01621459.1993.10476321

Allaire, J., Xie, Y., McPherson, J., Luraschi, J., Ushey, K., Atkins, A., … Chang, W. (2018). *Rmarkdown: Dynamic documents for r.* Retrieved from https://CRAN.R-project.org/package=rmarkdown

Anzulewicz, A., Asanowicz, D., Windey, B., Paulewicz, B., Wierzchoń, M., & Cleeremans, A. (2015). Does level of processing affect the transition from unconscious to conscious perception? *Consciousness and Cognition*, *36*, 1–11.

Augusto, L. M. (2016). Lost in dissociation: The main paradigms in unconscious cognition. *Consciousness and Cognition*, *42*, 293–310. https://doi.org/10.1016/j.concog.2016.04.004

Augusto, L. M. (2018). Transitions Versus Dissociations: A Paradigm Shift in Unconscious Cognition. *Axiomathes*, *28*(3), 269–291. https://doi.org/10.1007/s10516-017-9366-y

Aust, F., & Barth, M. (2018). *papaja: Create APA manuscripts with R Markdown.* Retrieved from https://github.com/crsh/papaja

Barth, M., Stahl, C., & Haider, H. (2018). Assumptions of the process-dissociation procedure are violated in implicit sequence learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition.* https://doi.org/10.1037/xlm0000614

Batchelder, W. H., & Riefer, D. M. (1999). Theoretical and empirical review of multinomial process tree modeling. *Psychonomic Bulletin & Review*, *6*(1), 57–86. https://doi.org/10.3758/BF03210812

Bayne, T., Hohwy, J., & Owen, A. M. (2016). Are There Levels of Consciousness? *Trends in Cognitive Sciences*, *20*(6), 405–413. https://doi.org/10.1016/j.tics.2016.03.009

Berry, C. J., Shanks, D. R., Speekenbrink, M., & Henson, R. N. (2012). Models of recognition, repetition priming, and fluency: Exploring a new framework. *Psychological Review*, *119*(1), 40. Retrieved from http://psycnet.apa.org/psycarticles/2011-24233-001

Block, N. (1995). On a confusion about a function of consciousness. *Behavioral and Brain Sciences*, *18*(2), 227–247. https://doi.org/10.1017/S0140525X00038188

Block, N. (2007). Consciousness, accessibility, and the mesh between psychology and neuroscience. *Behavioral and Brain Sciences*, *30*(5-6). https://doi.org/10.1017/S0140525X07002786

Boyer, M., Destrebecqz, A., & Cleeremans, A. (2005). Processing abstract sequence structure: Learning without knowing, or knowing without learning? *Psychological Research*, *69*(5-6), 383–398. https://doi.org/10.1007/s00426-004-0207-4

Boywitt, C. D., & Rummel, J. (2012). A diffusion model analysis of task interference effects in prospective memory. *Memory & Cognition*, *40*(1), 70–82. https://doi.org/10.3758/s13421-011-0128-6

Brainerd, C. J., Reyna, V. F., & Mojardin, A. H. (1999). Conjoint recognition. *Psychological Review*, *106*(1), 160–179. https://doi.org/10.1037/0033-295X.106.1.160

Buchner, A., Erdfelder, E., Steffens, M. C., & Martensen, H. (1997a). The Nature of Memory Processes Underlying Recognition Judgments in the Process Dissociation Procedure. *Memory & Cognition*, *25*(4), 508–517. https://doi.org/10.3758/BF03201126

Buchner, A., Erdfelder, E., & Vaterrodt-Plünnecke, B. (1995). Toward Unbiased Measurement of Conscious and Unconscious Memory Processes within the Process Dissociation Framework. *Journal of Experimental Psychology: General*, *124*(2), 137–160. https://doi.org/10.1037/0096-3445.124.2.137

Buchner, A., Steffens, M. C., & Rothkegel, R. (1998). On the Role of Fragmentary Knowledge in a Sequence Learning Task. *The Quarterly Journal of Experimental Psychology Section A*, *51*(2), 251–281. https://doi.org/10.1080/713755757

Buchner, A., Steffens, M. C., Rothkegel, R., & Erdfelder, E. (1997b). A multinomial model to assess fluency and recollection in a sequence learning task. *The Quarterly Journal of Experimental Psychology Section A*, *50*(3), 631–663. https://doi.org/10.1080/713755723

Carpenter, B., Gelman, A., Hoffman, M., Lee, D., Goodrich, B., Betancourt, M., … Riddell, A. (2016). Stan: A probabilistic programming language. *Journal of Statistical Software*, *20*, 1–37. https://doi.org/10.18637/jss.v076.i01

Clark, H. H. (1973). The language-as-fixed-effect fallacy: A critique of language statistics in psychological research. *Journal of Verbal Learning and Verbal Behavior*, *12*(4), 335–359.

Cleeremans, A. (2011). The radical plasticity thesis: How the brain learns to be conscious. *Consciousness Research*, *2*, 86. https://doi.org/10.3389/fpsyg.2011.00086

Cleeremans, A., & Jiménez, L. (2002). Implicit learning and consciousness: A graded, dynamic perspective. *Implicit Learning and Consciousness*, 1–40.

Cleeremans, A., & McClelland, J. L. (1991). Learning the structure of event sequences. *Journal of Experimental Psychology: General*, *120*(3), 235. Retrieved from http://psycnet.apa.org/psycinfo/1992-00389-001

Cleeremans, A., Timmermans, B., & Pasquali, A. (2007). Consciousness and Metarepresentation: A Computational Sketch. *Neural Networks*, *20*(9), 1032–1039. https://doi.org/10.1016/j.neunet.2007.09.011

Clegg, B. A. (2005). Stimulus-specific sequence representation in serial reaction time tasks. *Quarterly Journal of Experimental Psychology: Section A*, *58*(6), 1087–1101. https://doi.org/10.1080/02724980443000485

Cohen, A., Ivry, R. I., & Keele, S. W. (1990). Attention and structure in sequence learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *16*(1), 17–30. https://doi.org/10.1037/0278-7393.16.1.17

Cohen, A. L., & Nosofsky, R. M. (2003). An Extension of the Exemplar-Based Random-Walk Model to Separable-Dimension Stimuli. *Journal of Mathematical Psychology*, *47*(2), 150–165. https://doi.org/10.1016/S0022-2496(02)00031-7

Curran, T. (2001). Implicit Learning Revealed by the Method of Opposition. *Trends in Cognitive Sciences*, *5*(12), 503–504. https://doi.org/10.1016/S1364-6613(00)01791-5

Curran, T., & Hintzman, D. L. (1995). Violations of the Independence Assumption in Process Dissociation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *21*(3), 531–547. https://doi.org/10.1037/0278-7393.21.3.531

Curran, T., & Hintzman, D. L. (1997). Consequences and Causes of Correlations in Process Dissociation. *Journal of Experimental Psychology. Learning, Memory & Cognition*, *23*(2), 496. https://doi.org/10.1037/0278-7393.23.2.496

D'Angelo, M. C., Milliken, B., Jiménez, L., & Lupiáñez, J. (2013). Implementing flexibility in automaticity: Evidence from context-specific implicit sequence learning. *Consciousness and Cognition*, *22*(1), 64–81. https://doi.org/10.1016/j.concog.2012.11.002

Denwood, M. J. (2016). runjags: An R package providing interface utilities, model templates, parallel computing methods and additional distributions for MCMC models in JAGS. *Journal of Statistical Software*, *71*(9), 1–25. https://doi.org/10.18637/jss.v071.i09

Deroost, N., & Soetens, E. (2006a). Perceptual or motor learning in SRT tasks with complex sequence structures. *Psychological Research Psychologische Forschung*, *70*(2), 88–102. https://doi.org/10.1007/s00426-004-0196-3

Deroost, N., & Soetens, E. (2006b). The role of response selection in sequence learning. *Quarterly Journal of Experimental Psychology*, *59*(3), 449–456.

Destrebecqz, A. (2004). The Effect of Explicit Knowledge on Sequence Learning: A Graded Account. *Psychologica Belgica*, *44*(4), 217. https://doi.org/10.5334/pb-44-4-217

Destrebecqz, A., & Cleeremans, A. (2001). Can Sequence Learning Be Implicit? New Evidence with the Process Dissociation Procedure. *Psychonomic Bulletin & Review*, *8*(2), 343–350. https://doi.org/10.3758/BF03196171

Destrebecqz, A., & Cleeremans, A. (2003). Temporal effects in sequence learning. In *Attention and implicit learning* (pp. 181–213). Amsterdam, Netherlands: John Benjamins Publishing Company. https://doi.org/10.1075/aicr.48.11des

Destrebecqz, A., Franco, A., Bertels, J., & Gaillard, V. (2015). Direct and indirect measures of statistical learning. *Behavioral Methods in Consciousness Research*, 107–120.

Dienes, Z. (2015). How Bayesian statistics are needed to determine whether mental states are unconscious. In *Behavioral methods in consciousness research* (pp. 199–220). Oxford, UK: Oxford University Press.

Dienes, Z., & Perner, J. (1999). A Theory of Implicit and Explicit Knowledge. *Behavioral and Brain Sciences*, *22*(05), 735–808.

Dienes, Z., & Seth, A. (2010). Gambling on the Unconscious: A Comparison of Wagering and Confidence Ratings as Measures of Awareness in an Artificial Grammar Task. *Consciousness and Cognition*, *19*(2), 674–681. https://doi.org/10.1016/j.concog.2009.09.009

Donders, F. C. (1969). On the speed of mental processes. *Acta Psychologica*, *30*, 412–431. https://doi.org/10.1016/0001-6918(69)90065-1

Donkin, C., Averell, L., Brown, S., & Heathcote, A. (2009). Getting more from accuracy and response time data: Methods for fitting the linear ballistic accumulator. *Behavior Research Methods*, *41*(4), 1095–1110. https://doi.org/10.3758/BRM.41.4.1095

Eberhardt, K., Esser, S., & Haider, H. (2017). Abstract feature codes: The building blocks of the implicit learning system. *Journal of Experimental Psychology: Human Perception and Performance*, *43*(7), 1275–1290. https://doi.org/10.1037/xhp0000380

Efron, B., & Morris, C. (1977). Stein's paradox in statistics. *Scientific American*, *236*(5), 119–127.

Erdfelder, E., Auer, T.-S., Hilbig, B. E., Aßfalg, A., Moshagen, M., & Nadarevic, L. (2009). Multinomial Processing Tree Models: A review of the literature. *Zeitschrift Für Psychologie / Journal of Psychology*, *217*(3), 108–124. https://doi.org/10.1027/0044-3409.217.3.108

Esser, S., & Haider, H. (2017). The Emergence of Explicit Knowledge in a Serial Reaction Time Task: The Role of Experienced Fluency and Strength of Representation. *Frontiers in Psychology*, *8*. https://doi.org/10.3389/fpsyg.2017.00502

Esser, S., & Haider, H. (2018). Action–effects enhance explicit sequential learning. *Psychological Research*, *82*(6), 1113–1129. https://doi.org/10.1007/s00426-017-0883-5

Esser, S., & Haider, H. (n.d.). How the Task-Set Influences Implicit and Explicit Learning.

Fazekas, P., & Overgaard, M. (2016). Multidimensional Models of Degrees and Levels of Consciousness. *Trends in Cognitive Sciences*, *20*(10), 715–716. https://doi.org/10.1016/j.tics.2016.06.011

Frensch, P. A., Haider, H., Rünger, D., Neugebauer, U., Voigt, S., & Werg, J. (2003). The Route From Implicit Learning to Verbal Expression of What has Been Learned. In L. Jimenez (Ed.), *Attention and Implicit Learning* (p. 335). John Benjamins.

Frensch, P. A., Lin, J., & Buchner, A. (1998). Learning versus behavioral expression of the learned: The effects of a secondary tone-counting task on implicit learning in the serial reaction task. *Psychological Research*, *61*(2), 83–98.

Fu, Q., Dienes, Z., & Fu, X. (2010). Can Unconscious Knowledge Allow Control in Sequence Learning? *Consciousness and Cognition*, *19*(1), 462–474. https://doi.org/10.1016/j.concog.2009.10.001

Fu, Q., Fu, X., & Dienes, Z. (2008). Implicit Sequence Learning and Conscious Awareness. *Consciousness and Cognition*, *17*(1), 185–202. https://doi.org/10.1016/j.concog.2007.01.007

Gaschler, R., Frensch, P. A., Cohen, A., & Wenke, D. (2012). Implicit sequence learning based on instructed task set. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *38*(5), 1389.

Goschke, T. (1998). Implicit learning of perceptual and motor sequences: Evidence for independent learning systems. In *Handbook of implicit learning* (pp. 401–444). Thousand Oaks, CA, US: Sage Publications, Inc.

Goschke, T., & Bolte, A. (2012). On the modularity of implicit sequence learning: Independent acquisition of spatial, symbolic, and manual sequences. *Cognitive Psychology*, *65*(2), 284–320. https://doi.org/10.1016/j.cogpsych.2012.04.002

Graf, P., & Komatsu, S.-I. (1994). Process Dissociation Procedure: Handle with Caution! *European Journal of Cognitive Psychology*, *6*(2), 113–129. https://doi.org/10.1080/09541449408520139

Haider, H., Eichler, A., & Lange, T. (2011). An Old Problem: How Can We Distinguish between Conscious and Unconscious Knowledge Acquired in an Implicit Learning Task? *Consciousness and Cognition*, *20*(3), 658–672. https://doi.org/10.1016/j.concog.2010.10.021

Haider, H., Esser, S., & Eberhardt, K. (2018). Feature codes in implicit sequence learning: Perceived stimulus locations transfer to motor response locations. *Psychological Research*, 1–12. https://doi.org/10.1007/s00426-018-0980-0

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). Retrieved from http://psycnet.apa.org/journals/xlm/28/2/392/

Haider, H., & Frensch, P. A. (2005). The generation of conscious awareness in an incidental learning situation. *Psychological Research*, *69*(5), 399–411. https://doi.org/10.1007/s00426-004-0209-2

Haider, H., & Frensch, P. A. (2009). Conflicts between expected and actually performed behavior lead to verbal report of incidentally acquired sequential knowledge. *Psychological Research*, *73*(6), 817–834. https://doi.org/10.1007/s00426-008-0199-6

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. https://doi.org/10.1016/j.concog.2004.12.002

Hazeltine, E., & Schumacher, E. H. (2016). Understanding Central Processes: The Case against Simple Stimulus-Response Associations and for Complex Task Representation. In B. H. Ross (Ed.), *Psychology of Learning and Motivation* (Vol. 64, pp. 195–245). Academic Press. https://doi.org/10.1016/bs.plm.2015.09.006

Heathcote, A., Brown, S., & Mewhort, D. J. K. (2000). The power law repealed: The case for an exponential law of practice. *Psychonomic Bulletin & Review*, *7*(2), 185–207.

Heck, D. W., & Erdfelder, E. (2016). Extending multinomial processing tree models to measure the relative speed of cognitive processes. *Psychonomic Bulletin & Review*, 1–26. https://doi.org/10.3758/s13423-016-1025-6

Hintzman, D. L., & Curran, T. (1997). More than One Way to Violate Independence: Reply to Jacoby and Shrout (1997). *Journal of Experimental Psychology. Learning, Memory & Cognition*, *23*(2), 511. https://doi.org/10.1037/0278-7393.23.2.511

Hirshman, E. (1998). On the Logic of Testing the Independence Assumption in the Process-Dissociation Procedure. *Memory & Cognition*, *26*(5), 857–859. https://doi.org/10.3758/BF03201168

Hirshman, E. (2004). Ordinal Process Dissociation and the Measurement of Automatic and Controlled Processes. *Psychological Review*, *111*(2), 553–560. https://doi.org/10.1037/0033-295X.111.2.553

Hoffmann, J., & Koch, I. (1997). Stimulus-response compatibility and sequential learning in the serial reaction time task. *Psychological Research*, *60*(1-2), 87–97. https://doi.org/10.1007/BF00419682

Hoffmann, J., Sebald, A., & Stöcker, C. (2001). Irrelevant response effects improve serial learning in serial reaction time tasks. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *27*(2), 470–482. https://doi.org/10.1037/0278-7393.27.2.470

Hommel, B., Müsseler, J., Aschersleben, G., & Prinz, W. (2001). The Theory of Event Coding (TEC): A framework for perception and action planning. *Behavioral and Brain Sciences*, *24*(05), 849–878. https://doi.org/10.1017/S0140525X01000103

Howard, J. H., Mutter, S. A., & Howard, D. V. (1992). Serial pattern learning by event observation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *18*(5), 1029.

Jacoby, L. L. (1991). A Process Dissociation Framework: Separating Automatic from Intentional Uses of Memory. *Journal of Memory and Language*, *30*(5), 513–541. https://doi.org/10.1016/0749-596X(91)90025-F

Jacoby, L. L., Begg, I. M., & Toth, J. P. (1997). In defense of functional independence: Violations of assumptions underlying the... *Journal of Experimental Psychology. Learning, Memory & Cognition*, *23*(2), 484. Retrieved from http://search.ebscohost.com/login.aspx?direct=true&db=bth&AN=9706175953&site=ehost-live

Jacoby, L. L., & Shrout, P. E. (1997). Toward a psychometric analysis of violations of the independence assumption in process dissociation. *Journal of Experimental Psychology. Learning, Memory & Cognition*, *23*(2), 505. https://doi.org/10.1037/0278-7393.23.2.505

Jacoby, L. L., Toth, J. P., & Yonelinas, A. P. (1993). Separating conscious and unconscious influences of memory: Measuring recollection. *Journal of Experimental Psychology: General*, *122*(2), 139–154. https://doi.org/10.1037/0096-3445.122.2.139

Jiménez, L., & Méndez, C. (1999). Which attention is needed for implicit sequence learning? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *25*(1), 236–259. https://doi.org/10.1037/0278-7393.25.1.236

Jiménez, L., Méndez, C., & Cleeremans, A. (1996). Comparing direct and indirect measures of sequence learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *22*(4), 948. https://doi.org/10.1037/0278-7393.22.4.948

Jiménez, L., Vaquero, J. M. M., & Lupiáñez, J. (2006). Qualitative differences between implicit and explicit sequence learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *32*(3), 475–490. https://doi.org/10.1037/0278-7393.32.3.475

Jiménez, L., & Vázquez, G. A. (2005). Sequence Learning under Dual-Task Conditions: Alternatives to a Resource-Based Account. *Psychological Research*, *69*(5-6), 352–368. https://doi.org/10.1007/s00426-004-0210-9

Joordens, S., & Merikle, P. M. (1993). Independence or redundancy? Two models of conscious and unconscious influences. *Journal of Experimental Psychology: General*, *122*(4), 462–467. https://doi.org/10.1037/0096-3445.122.4.462

Joordens, S., Wilson, D. E., Spalek, T. M., & Paré, D. E. (2010). Turning the Process-Dissociation Procedure inside-out: A New Technique for Understanding the Relation between Conscious and Unconscious Influences. *Consciousness and Cognition*, *19*(1), 270–280. https://doi.org/10.1016/j.concog.2009.09.011

Katahira, K. (2016). How Hierarchical Models Improve Point Estimates of Model Parameters at the Individual Level. *Journal of Mathematical Psychology*, *73*, 37–58. https://doi.org/10.1016/j.jmp.2016.03.007

Keele, S. W., Ivry, R., Mayr, U., Hazeltine, E., & Heuer, H. (2003). The cognitive and neural architecture of sequence representation. *Psychological Review*, *110*(2), 316–339. https://doi.org/10.1037/0033-295X.110.2.316

Kinder, A., Rolfs, M., & Kliegl, R. (2008). Short Article: Sequence Learning at Optimal

Stimulus–Response Mapping: Evidence from a Serial Reaction Time Task. *Quarterly Journal of Experimental Psychology*, *61*(2), 203–209. https://doi.org/10.1080/17470210701557555

Klauer, K. C. (2006). Hierarchical Multinomial Processing Tree Models: A Latent-Class Approach. *Psychometrika*, *71*(1), 7–31. https://doi.org/10.1007/s11336-004-1188-3

Klauer, K. C. (2010). Hierarchical Multinomial Processing Tree Models: A Latent-Trait Approach. *Psychometrika*, *75*(1), 70–98. https://doi.org/10.1007/s11336-009-9141-0

Klauer, K. C., Dittrich, K., Scholtes, C., & Voss, A. (2015). The Invariance Assumption in Process-Dissociation Models: An Evaluation Across Three Domains. *Journal of Experimental Psychology: General*, *144*(1), 198–221. https://doi.org/10.1037/xge0000044

Klauer, K. C., & Kellen, D. (2018). RT-MPTs: Process models for response-time distributions based on multinomial processing trees with applications to recognition memory. *Journal of Mathematical Psychology*, *82*, 111–130. https://doi.org/10.1016/j.jmp.2017.12.003

Knapp, B. R., & Batchelder, W. H. (2004). Representing Parametric Order Constraints in Multi-Trial Applications of Multinomial Processing Tree Models. *Journal of Mathematical Psychology*, *48*(4), 215–229. https://doi.org/10.1016/j.jmp.2004.03.002

Koch, I. (2007). Anticipatory response control in motor sequence learning: Evidence from stimulus–response compatibility. *Human Movement Science*, *26*(2), 257–274. https://doi.org/10.1016/j.humov.2007.01.004

Koch, I., & Hoffmann, J. (2000a). Patterns, Chunks, and Hierarchies in Serial Reaction-Time Tasks. *Psychological Research*, *63*(1), 22–35. https://doi.org/10.1007/PL00008165

Koch, I., & Hoffmann, J. (2000b). The role of stimulus-based and response-based spatial information in sequence learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *26*(4), 863–882. https://doi.org/10.1037/0278-7393.26.4.863

Koch, I., Keller, P., & Prinz, W. (2004). The ideomotor approach to action control: Implications for skilled performance. *International Journal of Sport and Exercise Psychology*, *2*(4), 362–375.

Kornblum, S., Hasbroucq, T., & Osman, A. (1990). Dimensional overlap: Cognitive basis for stimulus-response compatibility–a model and taxonomy. *Psychological Review*, *97*(2), 253.

Lewandowski, D., Kurowicka, D., & Joe, H. (2009). Generating random correlation matrices based on vines and extended onion method. *Journal of Multivariate Analysis*, *100*(9), 1989–2001. https://doi.org/10.1016/j.jmva.2009.04.008

Liefooghe, B., Wenke, D., & De Houwer, J. (2012). Instruction-based task-rule congruency effects. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *38*(5), 1325–1335. https://doi.org/10.1037/a0028148

Logan, G. D. (1988). Toward an instance theory of automatization. *Psychological Review*, *95*(4), 492.

Logan, G. D. (1990). Repetition priming and automaticity: Common underlying mechanisms? *Cognitive Psychology*, *22*(1), 1–35. https://doi.org/10.1016/0010-0285(90)90002-L

Logan, G. D. (1992). Shapes of reaction-time distributions and shapes of learning curves: A test of the instance theory of automaticity. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *18*(5), 883.

Matzke, D., Dolan, C. V., Batchelder, W. H., & Wagenmakers, E.-J. (2013). Bayesian Estimation of Multinomial Processing Tree Models with Heterogeneity in Participants and Items. *Psychometrika*, *80*(1), 205–235. https://doi.org/10.1007/s11336-013-9374-9

Mayr, U. (1996). Spatial attention and implicit sequence learning: Evidence... *Journal of Experimental Psychology. Learning, Memory & Cognition*, *22*(2), 350. Retrieved from http://search.ebscohost.com/login.aspx?direct=true&db=bth&AN=9604196042&site=ehost-live

Morey, R. D., Rouder, J. N., & Speckman, P. L. (2008). A Statistical Model for Discriminating between Subliminal and Near-Liminal Performance. *Journal of Mathematical Psychology*, *52*(1), 21–36. https://doi.org/10.1016/j.jmp.2007.09.007

Mulder, M. J., Wagenmakers, E.-J., Ratcliff, R., Boekel, W., & Forstmann, B. U. (2012). Bias in the Brain: A Diffusion Model Analysis of Prior Probability and Potential Payoff. *Journal of Neuroscience*, *32*(7), 2335–2343. https://doi.org/10.1523/JNEUROSCI.4156-11.2012

Nattkemper, D., & Prinz, W. (1997). Stimulus and response anticipation in a serial reaction task. *Psychological Research*, *60*(1-2), 98–112. https://doi.org/10.1007/BF00419683

Newell, B. R., & Shanks, D. R. (2014). Unconscious influences on decision making: A critical review. *Behavioral and Brain Sciences*, *37*(1), 1–19. https://doi.org/10.1017/S0140525X12003214

Nissen, M. J., & Bullemer, P. (1987). Attentional Requirements of Learning: Evidence from Performance Measures. *Cognitive Psychology*, *19*(1), 1–32. https://doi.org/10.1016/

0010-0285(87)90002-8

Norman, E., Price, M. C., & Duff, S. C. (2006). Fringe consciousness in sequence learning: The influence of individual differences. *Consciousness and Cognition*, *15*(4), 723–760. https://doi.org/10.1016/j.concog.2005.06.003

Nosofsky, R. M., & Palmeri, T. J. (1997). An Exemplar-Based Random Walk Model of Speeded Classification. *Psychological Review*, *104*(2), 266–300. https://doi.org/10.1037/0033-295X.104.2.266

Palmeri, T. J. (1997). Exemplar similarity and the development of automaticity. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *23*(2), 324–354. https://doi.org/10.1037/0278-7393.23.2.324

Pasquali, A., Timmermans, B., & Cleeremans, A. (2010). Know thyself: Metacognitive networks and measures of consciousness. *Cognition*, *117*(2), 182–190.

Perruchet, P., & Amorim, M.-A. (1992). Conscious Knowledge and Changes in Performance in Sequence Learning: Evidence against Dissociation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *18*(4), 785. https://doi.org/10.1037/0278-7393.18.4.785

Perruchet, P., Bigand, E., & Benoit-Gonin, F. (1997). The emergence of explicit knowledge during the early phase of learning in sequential reaction time tasks. *Psychological Research*, *60*(1-2), 4–13. https://doi.org/10.1007/BF00419676

Perruchet, P., & Gallego, J. (1993). Association between conscious knowledge and performance in normal subjects: Reply to Cohen and Curran (1993) and Willingham, Greeley, and Bardone (1993). https://doi.org/10.1037/0278-7393.19.6.1438

Persaud, N., McLeod, P., & Cowey, A. (2007). Post-Decision Wagering Objectively Measures Awareness. *Nature Neuroscience*, *10*(2), 257–261. https://doi.org/10.1038/nn1840

Plummer, M. (2015). JAGS Version 4.0.0 user manual. International Agency for Research on Cancer. Retrieved from http://ufpr.dl.sourceforge.net/project/mcmc-jags/Manuals/4.x/jags_user_manual.pdf

Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, *85*(2), 59.

Ratcliff, R., & McKoon, G. (2008). The Diffusion Decision Model: Theory and Data for Two-Choice Decision Tasks. *Neural Computation*, *20*(4), 873–922. https://doi.org/10.1162/neco.2008.12-06-420

Ratcliff, R., Smith, P. L., Brown, S. D., & McKoon, G. (2016). Diffusion Decision Model:

Current Issues and History. *Trends in Cognitive Sciences*, *20*(4), 260–281. https://doi.org/10.1016/j.tics.2016.01.007

Ratcliff, R., Van Zandt, T., & McKoon, G. (1995). Process dissociation, single-process theories, and recognition memory. *Journal of Experimental Psychology: General*, *124*(4), 352–374. https://doi.org/10.1037/0096-3445.124.4.352

R Core Team. (2018). *R: A language and environment for statistical computing.* Vienna, Austria: R Foundation for Statistical Computing. Retrieved from https://www.R-project.org/

Reed, J., & Johnson, P. (1994). Assessing Implicit Learning with Indirect Tests: Determining What Is Learned about Sequence Structure. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *20*(3), 585–594. https://doi.org/10.1037/0278-7393.20.3.585

Reingold, E. M. (2004). Unconscious Perception and the Classic Dissociation Paradigm: A New Angle? *Perception & Psychophysics*, *66*(5), 882–887. https://doi.org/10.3758/BF03194981

Reingold, E. M., & Merikle, P. M. (1988). Using direct and indirect measures to study perception without awareness. *Perception & Psychophysics*, *44*(6), 563–575. https://doi.org/10.3758/BF03207490

Reingold, E. M., & Merikle, P. M. (1990). On the Inter-relatedness of Theory and Measurement in the Study of Unconscious Processes. *Mind & Language*, *5*(1), 9–28. https://doi.org/10.1111/j.1468-0017.1990.tb00150.x

Rose, M., Haider, H., & Büchel, C. (2010). The Emergence of Explicit Memory during Learning. *Cerebral Cortex*, *20*(12), 2787–2797. https://doi.org/10.1093/cercor/bhq025

Rosenthal, D. M. (1990). *A theory of consciousness.* Zentrum für interdisziplinäre Forschung.

Rouder, J. N., & Lu, J. (2005). An introduction to Bayesian hierarchical models with an application in the theory of signal detection. *Psychonomic Bulletin & Review*, *12*(4), 573–604. https://doi.org/10.3758/BF03196750

Rouder, J. N., Lu, J., Morey, R. D., Sun, D., & Speckman, P. L. (2008). A Hierarchical Process-Dissociation Model. *Journal of Experimental Psychology: General*, *137*(2), 370–389. https://doi.org/10.1037/0096-3445.137.2.370

Rouder, J. N., Lu, J., Speckman, P., Sun, D., & Jiang, Y. (2005). A Hierarchical Model for Estimating Response Time Distributions. *Psychonomic Bulletin & Review*, *12*(2), 195–223. Retrieved from http://link.springer.com/article/10.3758/BF03257252

Rouder, J. N., Morey, R. D., & Pratte, M. S. (2013). Bayesian Hierarchical Models. In *Measurement and Methodology* (Vol. 1, p. 82). Cambridge, MA, US: Cambridge University Press.

Rouder, J. N., Morey, R. D., Speckman, P. L., & Pratte, M. S. (2007). Detecting chance: A solution to the null sensitivity problem in subliminal priming. *Psychonomic Bulletin & Review*, *14*(4), 597–605. https://doi.org/10.3758/BF03196808

Rouder, J. N., Morey, R., & Wagenmakers, E.-J. (2016). The Interplay between Subjectivity, Statistical Practice, and Psychological Science. *Collabra*, *2*(1). https://doi.org/10.1525/collabra.28

Rowland, L. A., & Shanks, D. R. (2006). Attention modulates the learning of multiple contingencies. *Psychonomic Bulletin & Review*, *13*(4), 643–648. https://doi.org/10.3758/BF03193975

Röttger, E., Haider, H., Zhao, F., & Gaschler, R. (2017). Implicit sequence learning despite multitasking: The role of across-task predictability. *Psychological Research*. https://doi.org/10.1007/s00426-017-0920-4

Rünger, D., & Frensch, P. A. (2008). How Incidental Sequence Learning Creates Reportable Knowledge: The Role of Unexpected Events. *Journal of Experimental Psychology. Learning, Memory & Cognition*, *34*(5), 1011–1026. https://doi.org/10.1037/a0012942

Schmidtke, V., & Heuer, H. (1997). Task integration as a factor in secondary-task effects on sequence learning. *Psychological Research*, *60*(1-2), 53–71. https://doi.org/10.1007/BF00419680

Schumacher, E. H., & Hazeltine, E. (2016). Hierarchical Task Representation : Task Files and Response Selection. *Current Directions in Psychological Science*, *25*(6), 449–454. https://doi.org/10.1177/0963721416665085

Schumacher, E. H., & Schwarb, H. (2009). Parallel response selection disrupts sequence learning under dual-task conditions. *Journal of Experimental Psychology: General*, *138*(2), 270.

Schwarb, H., & Schumacher, E. H. (2009). Neural evidence of a role for spatial response selection in the learning of spatial sequences. *Brain Research*, *1247*, 114–125. https://doi.org/10.1016/j.brainres.2008.09.097

Schwarb, H., & Schumacher, E. H. (2010). Implicit sequence learning is represented by stimulus-response rules. *Memory & Cognition*, *38*(6), 677–688.

Schwarb, H., & Schumacher, E. H. (2012). Generalized lessons about sequence learning from the study of the serial reaction time task. *Advances in Cognitive Psychology*, *8*(2), 165–178. https://doi.org/10.2478/v10053-008-0113-1

Shankar, K. H., Jagadisan, U. K., & Howard, M. W. (2009). Sequential learning using temporal context. *Journal of Mathematical Psychology*, *53*(6), 474–485. Retrieved from http://www.sciencedirect.com/science/article/pii/S0022249609000984

Shanks, D. R. (2005). Implicit Learning. In K. Lamberts & R. Goldstone (Eds.), *Handbook of Cognition*. London, UK: Sage.

Shanks, D. R., & Johnstone, T. (1999). Evaluating the Relationship between Explicit and Implicit Knowledge in a Sequential Reaction Time Task. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *25*(6), 1435–1451. https://doi.org/10.1037/0278-7393.25.6.1435

Shanks, D. R., & Perruchet, P. (2002). Dissociation between Priming and Recognition in the Expression of Sequential Knowledge. *Psychonomic Bulletin & Review*, *9*(2), 362–367. https://doi.org/10.3758/BF03196294

Shanks, D. R., Rowland, L. A., & Ranger, M. S. (2005). Attentional Load and Implicit Sequence Learning. *Psychological Research*, *69*(5-6), 369–382. https://doi.org/10.1007/s00426-004-0211-8

Shanks, D. R., & St. John, M. F. (1994). Characteristics of Dissociable Human Learning Systems. *Behavioral and Brain Sciences*, *17*(3), 367–395. https://doi.org/10.1017/S0140525X00035032

Singmann, H., Bolker, B., Westfall, J., & Aust, F. (2018). *Afex: Analysis of factorial experiments*. Retrieved from https://CRAN.R-project.org/package=afex

Singmann, H., & Kellen, D. (2013). MPTinR: analysis of multinomial processing tree models in r. *Behavior Research Methods*, *45*(2), 560–575. https://doi.org/10.3758/s13428-012-0259-0

Sinnwell, J. P., & Schaid, D. J. (2013). *Ibdreg: Regression methods for ibd linkage with covariates*. Retrieved from https://CRAN.R-project.org/package=ibdreg

Smith, J. B., & Batchelder, W. H. (2010). Beta-MPT: Multinomial processing tree models for addressing individual differences. *Journal of Mathematical Psychology*, *54*(1), 167–183. https://doi.org/10.1016/j.jmp.2009.06.007

Song, S., Howard, J. H., & Howard, D. V. (2008). Perceptual sequence learning in a serial reaction time task. *Experimental Brain Research*, *189*(2), 145–158. https:

//doi.org/10.1007/s00221-008-1411-z

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. https://doi.org/10.1111/1467-9868.00353

Spiegelhalter, D. J., Best, N. G., Carlin, B. P., & van der Linde, A. (2014). The deviance information criterion: 12 years on. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, *76*(3), 485–493. https://doi.org/10.1111/rssb.12062

Stadler, M. A. (1989). On Learning Complex Procedural Knowledge. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *15*(6), 1061–1069. https://doi.org/10.1037/0278-7393.15.6.1061

Stadler, M. A. (1992). Statistical Structure and Implicit Serial Learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *18*(2), 318–327. https://doi.org/10.1037/0278-7393.18.2.318

Stahl, C., Barth, M., & Haider, H. (2015). Distorted Estimates of Implicit and Explicit Learning in Applications of the Process-Dissociation Procedure to the SRT Task. *Consciousness and Cognition*, *37*, 27–43. https://doi.org/10.1016/j.concog.2015.08.003

Stahl, C., & Degner, J. (2007). Assessing Automatic Activation of Valence: A Multinomial Model of EAST Performance. *Experimental Psychology*, *54*(2), 99–112. https://doi.org/10.1027/1618-3169.54.2.99

Stahl, C., & Klauer, K. C. (2007). HMMTree: A Computer Program for Latent-Class Hierarchical Multinomial Processing Tree Models. *Behavior Research Methods*, *39*(2), 267–273. https://doi.org/10.3758/BF03193157

Starns, J. J., Ratcliff, R., & McKoon, G. (2012). Modeling Single versus Multiple Systems in Implicit and Explicit Memory. *Trends in Cognitive Sciences*, *16*(4), 195–196. https://doi.org/10.1016/j.tics.2012.02.005

Steffens, M. C., Buchner, A., Martensen, H., & Erdfelder, E. (2000). Further evidence on the similarity of memory processes in the process dissociation procedure and in source monitoring. *Memory & Cognition*, *28*(7), 1152–1164. https://doi.org/10.3758/BF03211816

Sternberg, S. (1969). The discovery of processing stages: Extensions of Donders' method. *Acta Psychologica*, *30*, 276–315.

Stöcker, C., Sebald, A., & Hoffmann, J. (2003). The Influence of Response—Effect Compatibility in a Serial Reaction Time Task. *The Quarterly Journal of Experimental Psychology Section A*, *56*(4), 1–19. https://doi.org/10.1080/02724980244000585

Thomas, M. S. C., & McClelland, J. L. (2008). Connectionist Models of Cognition. In R. Sun (Ed.), *The Cambridge Handbook of Computational Psychology* (pp. 23–58). Cambridge: Cambridge University Press. https://doi.org/10.1017/CBO9780511816772.005

Timmermans, B., & Cleeremans, A. (2015). How can we measure awareness? An overview of current methods. In M. Overgaard & M. Overgaard (Eds.), *Behavioral Methods in Consciousness Research* (pp. 21–46). Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199688890.003.0003

Timmermans, B., Schilbach, L., Pasquali, A., & Cleeremans, A. (2012). Higher order thoughts in action: Consciousness as an unconscious re-description process. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *367*(1594), 1412–1423. https://doi.org/10.1098/rstb.2011.0421

Toth, J. P., Reingold, E. M., & Jacoby, L. L. (1994). Toward a redefinition of implicit memory: Process dissociations following elaborative processing and self-generation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *20*(2), 290–303. https://doi.org/10.1037/0278-7393.20.2.290

Tubau, E., Hommel, B., & López-Moliner, J. (2007). Modes of Executive Control in Sequence Learning: From Stimulus-Based to Plan-Based Control. *Journal of Experimental Psychology: General*, *136*(1), 43–63. https://doi.org/10.1037/0096-3445.136.1.43

Tubau, E., & López-Moliner, J. (2004). Spatial interference and response control in sequence learning: The role of explicit knowledge. *Psychological Research*, *68*(1), 55–63. https://doi.org/10.1007/s00426-003-0139-4

Verwey, W. B. (2003). Processing modes and parallel processors in producing familiar keying sequences. *Psychological Research*, *67*(2), 106–122. https://doi.org/10.1007/s00426-002-0120-7

Voss, A., Nagler, M., & Lerche, V. (2013a). Diffusion Models in Experimental Psychology. *Experimental Psychology*, *60*(6), 385–402. https://doi.org/10.1027/1618-3169/a000218

Voss, A., Rothermund, K., Gast, A., & Wentura, D. (2013b). Cognitive processes in associative and categorical priming: A diffusion model analysis. *Journal of Experimental Psychology: General*, *142*(2), 536–559. https://doi.org/10.1037/a0029459

Voss, A., Rothermund, K., & Voss, J. (2004). Interpreting the parameters of the diffusion model: An empirical validation. *Memory & Cognition*, *32*(7), 1206–1220. https://doi.org/10.3758/BF03196893

Voss, A., Voss, J., & Klauer, K. C. (2010). Separating response-execution bias from decision bias: Arguments for an additional parameter in Ratcliff's diffusion model. *British Journal of Mathematical and Statistical Psychology*, *63*(3), 539–555. https://doi.org/10.1348/000711009X477581

Wagenmakers, E.-J. (2009). Methodological and empirical developments for the Ratcliff diffusion model of response times and accuracy. *European Journal of Cognitive Psychology*, *21*(5), 641–671. https://doi.org/10.1080/09541440802205067

Wagenmakers, E.-J., Lodewyckx, T., Kuriyal, H., & Grasman, R. (2010). Bayesian hypothesis testing for psychologists: A tutorial on the Savage–Dickey method. *Cognitive Psychology*, *60*(3), 158–189. https://doi.org/10.1016/j.cogpsych.2009.12.001

Wessel, J. R., Haider, H., & Rose, M. (2012). The transition from implicit to explicit representations in incidental learning situations: More evidence from high-frequency EEG coupling. *Experimental Brain Research*, *217*(1), 153–162. https://doi.org/10.1007/s00221-011-2982-7

Wickens, T. D., & Keppel, G. (1983). On the choice of design and of test statistic in the analysis of experiments with sampled materials. *Journal of Verbal Learning and Verbal Behavior*, *22*(3), 296–309. https://doi.org/10.1016/S0022-5371(83)90208-6

Wilkinson, L., & Shanks, D. R. (2004). Intentional Control and Implicit Sequence Learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *30*(2), 354–369. https://doi.org/10.1037/0278-7393.30.2.354

Willingham, D. B. (1999). Implicit motor sequence learning is not purely perceptual. *Memory & Cognition*, *27*(3), 561–572.

Willingham, D. B., Nissen, M. J., & Bullemer, P. (1989). On the Development of Procedural Knowledge. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *15*(6), 1047–1060. https://doi.org/10.1037/0278-7393.15.6.1047

Windey, B., & Cleeremans, A. (2015). Consciousness as a graded and an all-or-none phenomenon: A conceptual analysis. *Consciousness and Cognition*, *35*, 185–191.

Wixted, J. T., & Mickes, L. (2010). A continuous dual-process model of remember/know judgments. *Psychological Review*, *117*(4), 1025.

Xie, Y. (2015). *Dynamic documents with R and knitr* (2nd ed.). Boca Raton, Florida:

Chapman; Hall/CRC. Retrieved from https://yihui.name/knitr/

Ylikoski, P. (2009). The Illusion of Depth of Understanding in Science. In H. D. Regt, S. Leonelli, & K. Eigner (Eds.), *Scientific Understanding: Philosophical Perspectives* (pp. 100–119). University of Pittsburgh Press.

Yonelinas, A. P., & Jacoby, L. L. (2012). The Process-Dissociation Approach Two Decades Later: Convergence, Boundary Conditions, and New Directions. *Memory & Cognition*, *40*(5), 663–680. https://doi.org/10.3758/s13421-012-0205-5

Yu, J., & Bellezza, F. S. (2000). Process dissociation as source monitoring. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *26*(6), 1518. Retrieved from http://psycnet.apa.org/journals/xlm/26/6/1518/

Ziessler, M., & Nattkemper, D. (2001). Learning of event sequences is based on response-effect learning: Further evidence from a serial reaction task. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *27*(3), 595–613. https://doi.org/10.1037/0278-7393.27.3.595

Appendices

Appendix A
Performance Data and Parameter Estimates

This appendix provides the raw generation performance and additional model analyses for the experiment reported in Chapter II.

Table A1
Proportion of correctly generated triples

Material	Generation task	Order	Inclusion *M*	Inclusion 95% CI	Exclusion *M*	Exclusion 95% CI
No-learning	Free	Inclusion first	.23	[.20, .27]	.21	[.18, .25]
		Exclusion first	.22	[.17, .27]	.24	[.19, .30]
	Cued	Inclusion first	.22	[.17, .27]	.22	[.17, .28]
		Exclusion first	.28	[.22, .34]	.25	[.21, .28]
Permuted	Free	Inclusion first	.29	[.26, .32]	.23	[.18, .28]
		Exclusion first	.27	[.22, .31]	.25	[.20, .29]
	Cued	Inclusion first	.24	[.20, .28]	.21	[.17, .25]
		Exclusion first	.19	[.13, .25]	.20	[.16, .23]
Random	Free	Inclusion first	.23	[.20, .25]	.20	[.15, .26]
		Exclusion first	.21	[.18, .25]	.18	[.15, .21]
	Cued	Inclusion first	.23	[.19, .28]	.24	[.18, .29]
		Exclusion first	.23	[.20, .26]	.23	[.19, .26]

Note. CI = Confidence intervals

Table A2
Estimates of model parameters before or after reversals were removed

Material	Generation task	Full dataset *A*	Full dataset *C*	After excluding reversals *A*	After excluding reversals *C*
No-learning	Free	.23 [.22, .24]	.00 [-.02, .02]	.24 [.23, .26]	.00 [-.02, .02]
	Cued	.24 [.23, .25]	.02 [.00, .04]	.27 [.26, .28]	.00 [-.02, .02]
Permuted	Free	.25 [.24, .26]	.04 [.02, .06]	.27 [.26, .28]	.03 [.01, .06]
	Cued	.21 [.20, .22]	.01 [-.01, .03]	.24 [.23, .25]	.00 [-.02, .02]
Random	Free	.20 [.19, .21]	.03 [.01, .05]	.23 [.22, .24]	.02 [.00, .05]
	Cued	.23 [.22, .24]	.00 [-.02, .02]	.26 [.25, .27]	.00 [-.02, .02]

Note. 95% confidence intervals are in parentheses.

Hierarchical-model analysis

We analysed our data using a modified version of Rouder, Lu, Morey, Sun, & Speckman (2008)'s three-level hierarchical process-dissociation model.

The first level is the process-dissociation model:

Table A3
Category counts

Material	Generation task	Inclusion Correct	Inclusion Incorrect	Exclusion Correct	Exclusion Incorrect
Full dataset					
No-learning	Free	851	2894	851	2875
	Cued	886	2626	825	2675
Permuted	Free	881	2285	755	2427
	Cued	723	2584	673	2621
Random	Free	777	2731	669	2824
	Cued	879	2916	888	2914
After excluding reversals					
No-learning	Free	851	2671	851	2605
	Cued	886	2452	825	2233
Permuted	Free	881	2101	755	2127
	Cued	723	2306	673	2112
Random	Free	777	2360	669	2324
	Cued	879	2642	888	2434

$$I_{ijk} = C_{ijk} + (1 - C_{ijk})A_{ijk}$$

and

$$E_{ijk} = (1 - C_{ijk})A_{ijk}$$

where i and j index participants and items, and k indexes the experimental condition. The parameters A and C represent probabilities that range between zero and one; they are transformed via a probit link to the reals, where a and c denote the transformed parameters:

$A_{ijk} = \Phi(a_{ijk})$ and $C_{ijk} = \Phi(c_{ijk})$

The second level is a main effects models on transformed parameters a and c:

$$c_{ijk} = \alpha_i^{(c)} + \beta_j^{(c)} + \mu_k^{(c)}$$

and

$$a_{ijk} = \alpha_i^{(a)} + \beta_j^{(a)} + \mu_k^{(a)}$$

where α denotes participant effects, β denotes item effects, and μ denotes condition effects that lead to conscious or unconscious contributions to task performance.

Participant and item effects are modeled as draws from bivariate normals whose covariance matrices were estimated from the data:

$$\left(\frac{\alpha_i^{(c)}}{\alpha_i^{(a)}}\right) \sim N_2(0, \textstyle\sum_\alpha), i = 1, \cdots, I.$$

and

$$\left(\frac{\beta_i^{(c)}}{\beta_i^{(a)}}\right) \sim N_2(0, \textstyle\sum_\beta), j = 1, \cdots, J.$$

This model was estimated within a Bayesian modeling framework using MCMC sampling. For further detail, refer to Rouder et al. (2008).

Results. For each group, we sampled three chains of 50,000 iterations, discarding the first 20,000 as burn-in. Mixing was monitored by $\hat{R}$ which was below 1.02. Table A4 shows estimates of the posterior distribution of the grand-mean parameters μ_k of the model. Table A5 shows the estimates equivalent to C and A from traditional analyses. As can be seen, the results corroborated the findings obtained with the traditional analyses reported above (i.e., $C > 0$, $A > .2$, and the ordering of A estimates across conditions).

Table A4

Parameter estimates from the hierarchical process-dissociation model. Parameters $\mu_k^{(c)}$ and $\mu_k^{(a)}$ refer to estimates of the grand mean.

		Full dataset				Reversals excluded			
		$\mu_k^{(a)}$		$\mu_k^{(c)}$		$\mu_k^{(a)}$		$\mu_k^{(c)}$	
No-learning	Free	-0.82	-0.96, -0.67	-6.11	[-8.61, -4.18]	-0.75	-0.90, -0.61	-8.58	-12.95, -5.51
	Cued	-0.82	-1.02, -0.62	-5.13	[-7.38, -3.44]	-0.74	-0.94, -0.55	-5.88	-7.97, -4.11
Permuted	Free	-0.71	-0.87, -0.55	-6.30	[-9.92, -3.70]	-0.64	-0.79, -0.48	-7.48	-10.98, -4.61
	Cued	-0.88	-1.04, -0.71	-6.62	[-10.59, -3.62]	-0.77	-0.94, -0.61	-7.14	-9.56, -4.64
Random	Free	-0.88	-1.04, -0.73	-6.40	[-12.02, -3.66]	-0.78	-0.94, -0.62	-6.88	-11.43, -3.61
	Cued	-0.78	-0.95, -0.61	-4.06	[-5.79, -2.69]	-0.68	-0.85, -0.52	-4.39	-6.12, -2.86

Note. 95% credible intervals are in parentheses.

Table A5

Parameter estimates from the hierarchical PD model. Parameters A and C denote the Bayesian equivalent to parameter estimates obtained from classical analyses.

		Full dataset				Reversals excluded			
		A		*C*		*A*		*C*	
No-learning	Free	.21	[.21, .22]	.03	[.03, .04]	.23	[.22, .24]	.03	[.03, .04]
	Cued	.23	[.22, .24]	.04	[.03, .05]	.26	[.24, .27]	.04	[.03, .05]
Permuted	Free	.25	[.23, .26]	.03	[.03, .04]	.27	[.26, .28]	.03	[.02, .03]
	Cued	.20	[.19, .21]	.04	[.03, .05]	.23	[.22, .24]	.04	[.03, .04]
Random	Free	.20	[.19, .21]	.02	[.02, .03]	.22	[.21, .23]	.04	[.03, .04]
	Cued	.22	[.21, .23]	.03	[.02, .03]	.25	[.24, .26]	.03	[.02, .04]

Note. 95% credible intervals are in parentheses.

Appendix B
Generation performance

This appendix provides the raw generation performance for all experiments in Chapter III in tables B1, B2, and B3.

Table B1
Mean percentage of regular transitions generated in Experiment 1, excluding repetitions. Standard deviations are given in parentheses.

Condition	Inclusion	Exclusion
Full dataset		
Control	25.10 (11.74)	24.17 (7.02)
No-Practice	37.94 (16.26)	28.66 (13.39)
Unspecific-Practice	34.46 (14.14)	26.46 (15.02)
Practice	38.74 (13.08)	24.59 (9.34)
Transfer	56.16 (18.32)	26.51 (7.93)
Nonrevealed transitions		
Control	25.10 (11.74)	24.17 (7.02)
No-Practice	29.20 (18.56)	31.90 (14.01)
Unspecific-Practice	30.38 (15.48)	29.34 (14.06)
Practice	29.63 (14.62)	26.81 (11.35)
Transfer	45.68 (24.66)	43.95 (17.03)
Revealed, but nonpracticed transitions		
No-Practice	47.64 (39.71)	24.65 (31.82)
Unspecific-Practice	33.91 (32.58)	20.07 (26.96)
Transfer	59.65 (33.59)	16.72 (22.52)
Revealed-and-practiced transitions		
Practice	75.65 (24.96)	15.63 (29.87)
Transfer	79.51 (21.81)	7.50 (7.13)

Table B2
Mean percentage of regular transitions generated in Experiment 2, excluding repetitions. Standard deviations are given in parentheses.

	Random		Probabilistic	
Condition	Inclusion	Exclusion	Inclusion	Exclusion
Full dataset				
No transition revealed	17.06 (8.64)	18.94 (10.99)	25.80 (19.20)	23.37 (10.16)
One transition revealed	30.00 (14.91)	15.26 (10.44)	41.56 (15.60)	22.38 (11.58)
Nonrevealed transitions				
No transition revealed	17.06 (8.64)	18.94 (10.99)	25.80 (19.20)	23.37 (10.16)
One transition revealed	18.46 (17.67)	16.80 (11.47)	31.29 (17.49)	25.82 (14.26)
Revealed transitions				
One transition revealed	79.37 (24.65)	8.74 (11.51)	86.75 (20.28)	6.77 (12.20)

Table B3

Mean percentage of regular transitions generated in Experiment 3, excluding repetitions and reversals. Standard deviations are given in parentheses.

	Random		Mixed SOC		Pure SOC	
Condition	Inclusion	Exclusion	Inclusion	Exclusion	Inclusion	Exclusion
Full dataset						
No transition revealed	26.57 (9.25)	23.16 (9.58)	28.43 (11.07)	25.11 (9.51)	28.85 (13.03)	25.98 (9.13)
Two transitions revealed	34.27 (8.75)	25.82 (6.00)	38.87 (10.47)	27.02 (10.85)	34.26 (8.86)	29.54 (8.93)
Nonrevealed transitions						
No transition revealed	26.57 (9.25)	23.16 (9.58)	28.43 (11.07)	25.11 (9.51)	28.85 (13.03)	25.98 (9.13)
Two transitions revealed	19.54 (8.86)	24.46 (7.59)	27.05 (11.64)	27.19 (8.94)	22.01 (9.61)	27.93 (7.63)
Revealed transitions						
Two transitions revealed	78.73 (28.18)	28.90 (31.97)	80.02 (21.62)	24.59 (27.24)	78.64 (26.53)	29.92 (31.57)

Appendix C
Additional ordinal-PD analyses

This appendix provides results of additional ordinal-PD analyses for Experiments 2 and 3 in Chapter III.

Experiment 2

Figure C1 shows the overall generation performance. We conducted a 2 (*Material*: Random vs. Probabilistic) $\times$ 2 (*Condition*: No transition revealed vs. One transition revealed) $\times$ 2 (*Order*: Inclusion first vs. Exclusion first) $\times$ 2 (*PD instruction*: Inclusion vs. Exclusion) ANOVA that revealed a main effect of *PD instruction*, $F(1, 113) = 28.43$, $MSE = 156.22$, $p < .001$, $\hat{\eta}^2_G = .109$, participants generated more regular transitions in inclusion than exclusion blocks; and a main effect of *explicit knowledge*, $F(1, 113) = 13.00$, $MSE = 164.96$, $p < .001$, $\hat{\eta}^2_G = .056$, indicating a clear influence of the explicit knowledge manipulation on generation performance. Moreover, we found a main effect of *material*, $F(1, 113) = 22.95$, $MSE = 164.96$, $p < .001$, $\hat{\eta}^2_G = .094$, participants generated more regular transitions if they had worked on regular material during the SRTT; the effect of *block order* also trended to be significant, $F(1, 113) = 3.57$, $MSE = 164.96$, $p = .062$, $\hat{\eta}^2_G = .016$, participants generated slightly more regular transitions if inclusion followed exclusion. These main effects were qualified by two-way interactions of *explicit knowledge* and *block order*, $F(1, 113) = 10.31$, $MSE = 164.96$, $p = .002$, $\hat{\eta}^2_G = .045$; and of *explicit knowledge* and *PD instruction*, $F(1, 113) = 26.64$, $MSE = 156.22$, $p < .001$, $\hat{\eta}^2_G = .103$; moreover, the four-way interaction of *material*, *explicit knowledge*, *block order*, and *PD instruction* was also found to be significant, $F(1, 113) = 5.42$, $MSE = 156.22$, $p = .022$, $\hat{\eta}^2_G = .023$. To disentangle these interactions, we analyzed inclusion and exclusion performance, separately.

Inclusion.

Analyzing the number of regular transitions generated in inclusion blocks, a 2 (*Material*: Random vs. Probabilistic) $\times$ 2 (*Condition*: No transition revealed vs. One transition revealed) $\times$ 2 (*Order*: Inclusion first vs. Exclusion first) ANOVA revealed a main effect of *material*, $F(1, 113) = 14.72$, $MSE = 207.66$, $p < .001$, $\hat{\eta}^2_G = .115$, participants generated more regular transitions if they had worked on probabilistic materials; and a main effect of *explicit knowledge*, $F(1, 113) = 29.57$, $MSE = 207.66$, $p < .001$, $\hat{\eta}^2_G = .207$, indicating a clear influence of our explicit-knowledge manipulation on inclusion performance. This effect was qualified by a significant interaction of *explicit knowledge* and *block order*, $F(1, 113) = 9.64$, $MSE = 207.66$, $p = .002$, $\hat{\eta}^2_G = .079$, indicating that participants used their explicit sequence knowledge more extensively if inclusion followed exclusion (i.e., after we had represented the

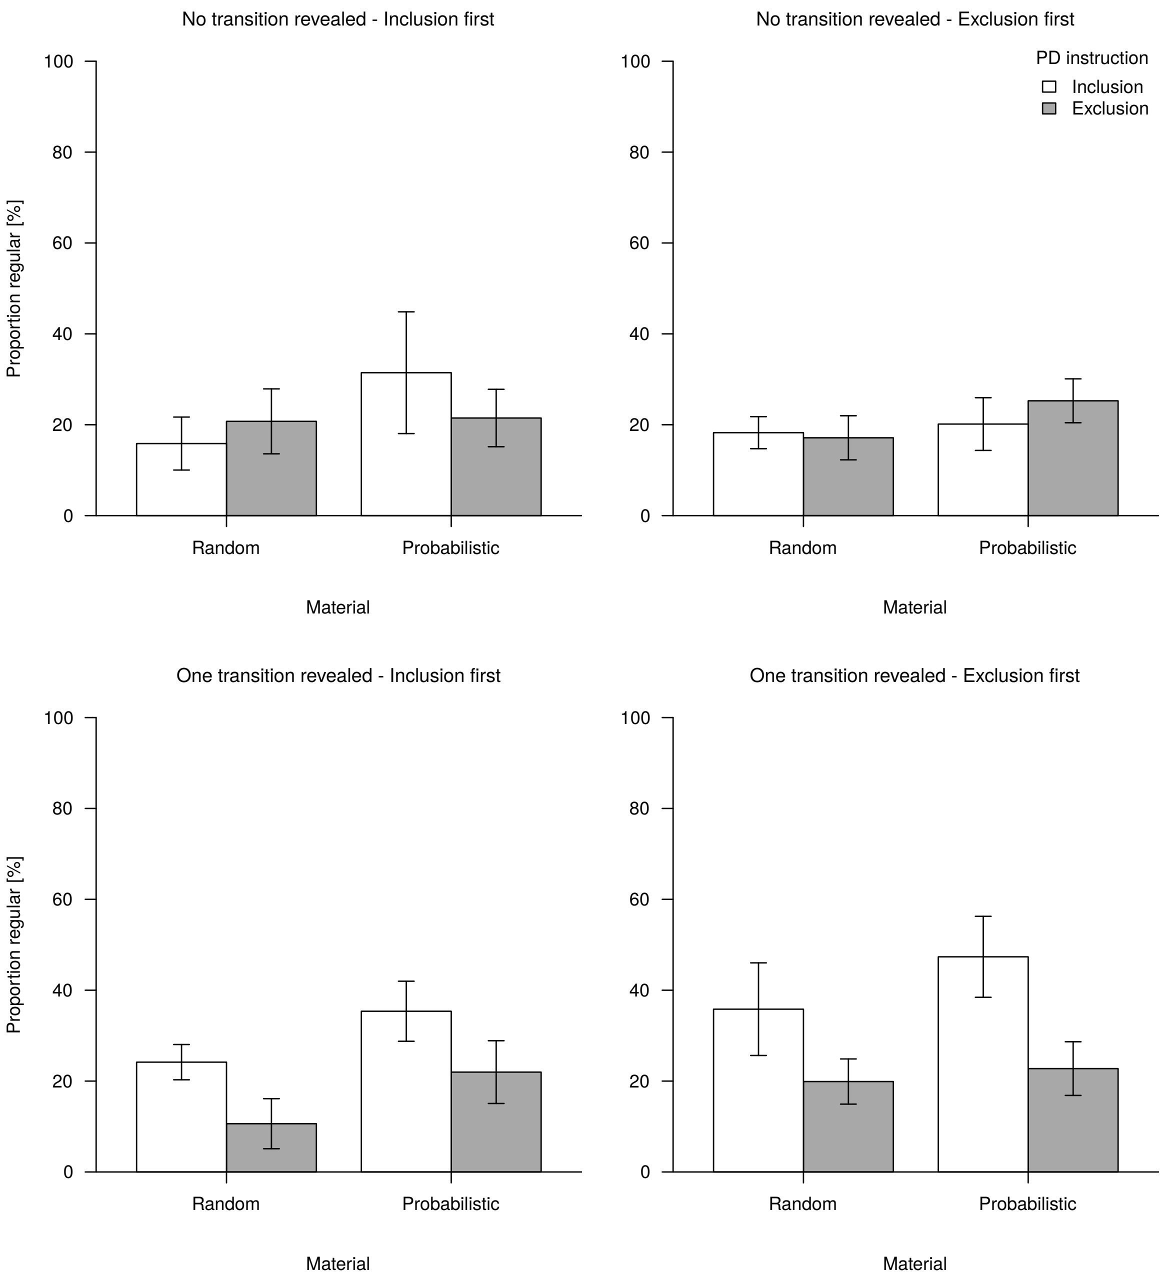

Figure C1. Mean proportion of correct FOCs during the generation task of Experiment 2, excluding repetitions. Error bars represent 95% confidence intervals.

transition a second time).

Exclusion.

Analyzing the number of regular transitions generated in exclusion blocks, a 2 (*Material*: Random vs. Probabilistic) $\times$ 2 (*Condition*: No transition revealed vs. One transition revealed) $\times$ 2 (*Order*: Inclusion first vs. Exclusion first) ANOVA revealed a main effect of *material* $F(1, 113) = 8.87$, $MSE = 113.52$, $p = .004$, $\hat{\eta}^2_G = .073$, participants generated more regular transitions if they had worked on probabilistic materials during the SRTT. We also found a significant three-way interaction of *material*, *explicit knowledge*, and *block order*, $F(1, 113) = 4.21$, $MSE = 113.52$, $p = .042$, $\hat{\eta}^2_G = .036$: Exclusion performance was below baseline only if exclusion followed inclusion *and* participants had worked on random material during the SRTT (i.e., they only had knowledge about one single transition of the sequence and had maximum practice in including/excluding this transition) – that is, if participants had no sequence knowledge but the single transition that we had revealed to them and they had already used this knowledge during the inclusion block, they were able to generate less regular transitions than baseline during the following exclusion block. The monotonicity assumption of the ordinal-PD approach is thus not violated in this single cell of the design. It is, hoewever, violated if exclusion preceded inclusion, or if participants had worked on probabilistic materials.

Experiment 3

Figure C2 shows the overall generation performance. A 3 (*Material*: Random vs. mixed SOC vs. pure SOC) $\times$ 2 (*Condition*: No transition revealed vs. Two transitions revealed) $\times$ 2 (*Order*: Inclusion first vs. Exclusion first) $\times$ 2 (*PD instruction*: Inclusion vs. Exclusion) ANOVA revealed a main effect of *PD instruction*, $F(1, 159) = 30.61$, $MSE = 94.53$, $p < .001$, $\hat{\eta}^2_G = .087$, participants generated more regular transitions in inclusion than exclusion blocks; and a main effect of *explicit knowledge*, $F(1, 159) = 25.01$, $MSE = 97.20$, $p < .001$, $\hat{\eta}^2_G = .074$, indicating a clear influence of the explicit knowledge manipulation on generation performance. Moreover, the interaction of *explicit knowledge* and *PD instruction* reached significance, $F(1, 159) = 6.18$, $MSE = 94.53$, $p = .014$, $\hat{\eta}^2_G = .019$, indicating that the effect of *explicit knowledge* is qualified by *PD instruction.* The interaction of *PD instruction* and *block order* almost reached significance, $F(1, 159) = 3.04$, $MSE = 94.53$, $p = .083$, $\hat{\eta}^2_G = .009$. To disentangle these interactions, we analyzed inclusion and exclusion performance, separately.

Inclusion.

Analyzing the number of regular transitions generated in inclusion blocks, a 3 (*Material*: Random vs. mixed SOC vs. pure SOC) $\times$ 2 (*Condition*: No transition revealed vs. Two

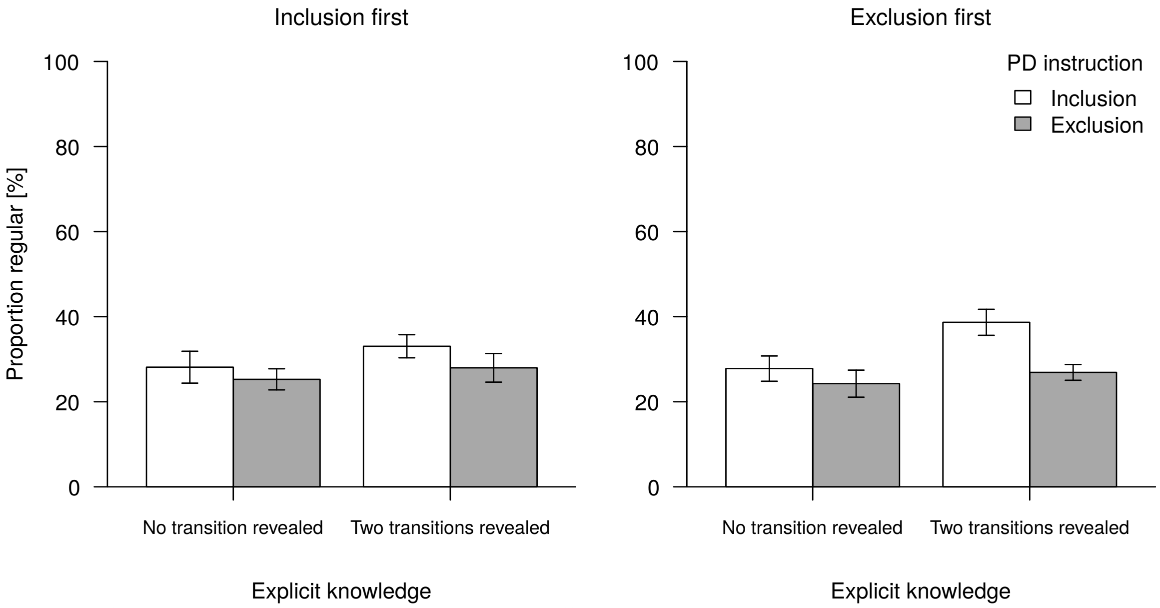

Figure C2. Mean proportion of correct SOCs during the generation task of Experiment 3, excluding repetitions and reversals. Error bars represent 95% confidence intervals.

transitions revealed) $\times$ 2 (*Order*: Inclusion first vs. Exclusion first) ANOVA revealed a significant main effect of *explicit knowledge*, $F(1,159) = 25.27$, $MSE = 106.81$, $p < .001$, $\hat{\eta}^2_G = .137$, indicating that our manipulation of explicit knowledge influenced inclusion performance. The main effect of *block order* trended to be significant, $F(1,159) = 2.84$, $MSE = 106.81$, $p = .094$, $\hat{\eta}^2_G = .018$, which was qualified by an almost significant interaction of *explicit knowledge* and *block order*, $F(1,159) = 3.70$, $MSE = 106.81$, $p = .056$, $\hat{\eta}^2_G = .023$. This pattern indicated that more regular transitions were generated if participants had received explicit knowledge about two transitions and inclusion followed exclusion, i.e. the explicit knowledge had been presented a second time (once prior to exclusion, once prior to inclusion).

Exclusion.

Analyzing the number of regular transitions generated in exclusion blocks, a 3 (*Material*: Random vs. mixed SOC vs. pure SOC) $\times$ 2 (*Condition*: No transition revealed vs. Two transitions revealed) $\times$ 2 (*Order*: Inclusion first vs. Exclusion first) ANOVA revealed only an almost significant main effect of *explicit knowledge*, $F(1,159) = 3.72$, $MSE = 84.92$, $p = .056$, $\hat{\eta}^2_G = .023$; revealing explicit knowledge about the sequence slightly *increased* the proportion of regular transitions generated. This pattern, again, violates the core assumption of the ordinal-PD approach that increasing amounts of explicit knowledge monotonically decrease the proportion of regular transitions in exclusion blocks. Moreover, it also shows that increasing explicit knowledge might produce a data pattern that is typically interpreted as evidence for increasing amounts of implicit knowldge.

Appendix D
Additional model analyses

This appendix provides results of additional model analyses not included in the main text of Chapter III.

Experiment 1, model $\mathcal{M}_1$

In Experiment 1, we fitted model $\mathcal{M}_1$ and used posterior analyses to evaluate the invariance assumption. We adapted the equations from Experiment 2 to the design of Experiment 1 (which did not contain experimental groups with random material). In order to accommodate for the more complex design, we used a model specification that allowed for participant and item (i.e., transition) effects and their interactions by estimating fixed effects for each transition type plus individual participants' deviations from these effects. The model equations of model $\mathcal{M}_1$ are given by:

$$C_{ijm} = \begin{cases} \Phi(\mu^{(C)}_{jlm} + \delta^{(C)}_{ijm}) & \text{if } j \epsilon 1, 2 \text{ (item has been revealed \& practiced, revealed \& non-practiced)} \\ 0 & \text{if } j = 3 \text{ (item has not been revealed)} \end{cases}$$

and

$$A_{imt} = \Phi(\mu^{(A)}_{mt} + \delta^{(A)}_{imt})$$

where $\mu^{(C)}_{jlm}$ is the fixed effect of transition type j (non-revealed, revealed & practiced, revealed & non-practiced) in condition l and *PD instruction* condition m on controlled processes, and $\delta^{(C)}_{ijm}$ is the ith participant's deviation from the corresponding mean. Accordingly, $\mu^{(A)}_{mt}$ is the fixed effect of *PD instruction* condition m and transition t on automatic processes, and $\delta^{(A)}_{imt}$ is the ith participant's deviation from the corresponding mean.

Model $\mathcal{M}_1$ imposes two auxiliary assumptions: First, it assumed that no explicit knowledge has been acquired during the SRT phase (i.e., $C = 0$ for non-revealed transitions). Second, it assumed that revealing sequence knowledge did not affect automatic processes (i.e., A does not vary as a function of the between-subjects manipulation of explicit knowledge, index l). Both auxiliary assumptions were tested by posterior predictive checks. In addition to reporting T_{A1} and T_{B1} as in Experiments 2 and 3, we calculated additional model check statistic T_{A2}, which summarizes how well the model describes the item-wise category counts (aggregated over participants), and T_{A3}, which summarizes how well the model describes the category counts per participant-item combination; finally, the additional statistic T_{B2} summarizes how well the model describes the variances and covariances introduced by items. We also calculated the posterior differences $C_I - C_E$ and $A_I - A_E$ to more directly test the invariance assumption.

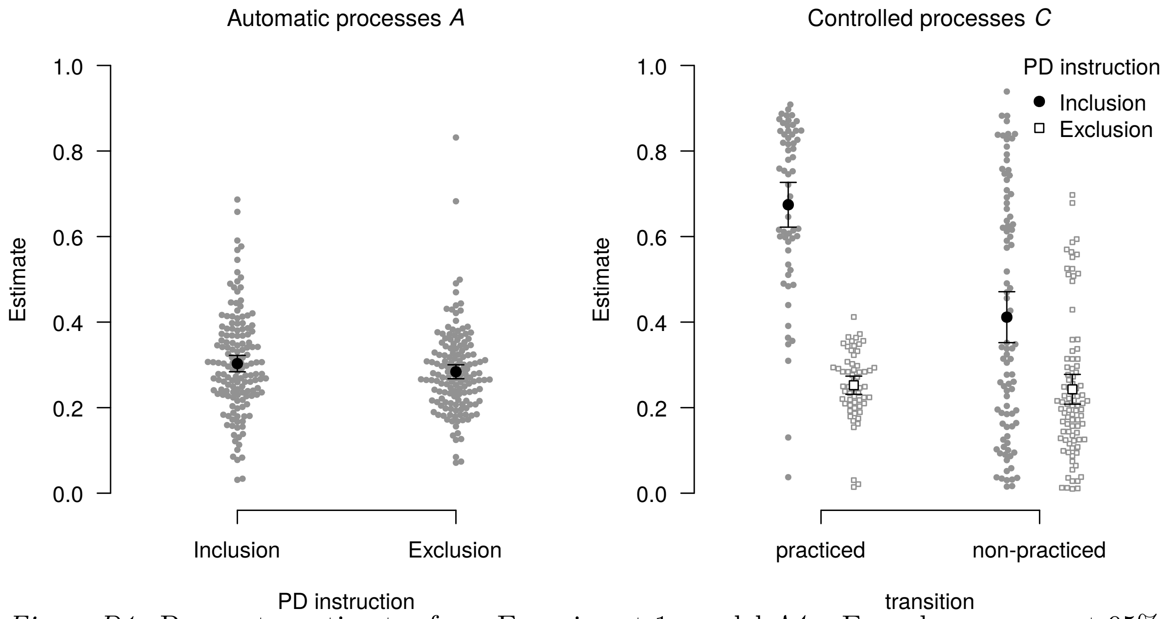

Figure D1. Parameter estimates from Experiment 1, model $\mathcal{M}_1$. Error bars represent 95% confidence intervals.

Results. We analyzed generation performance by fitting $\mathcal{M}_1$ and computed model fit statistics to assess whether each model can account for the data. Parameter estimates from model $\mathcal{M}_1$ were used to address the invariance assumptions, directly. The first trial of a block as well as any response repetitions were excluded from all generation task analyses.

The model checks for model $\mathcal{M}_1$ were satisfactory,

$$T_{A1}^{observed} = 35.97, T_{A1}^{expected} = 33.96, p = .322,$$

$$T_{A2}^{observed} = 0.05, T_{A2}^{expected} = 0.05, p = .480,$$

$$T_{A3}^{observed} = 1,763.79, T_{A3}^{expected} = 1,720.63, p = .372,$$

$$T_{B1}^{observed} = 5.31, T_{B1}^{expected} = 4.62, p = .457,$$

$$T_{B2}^{observed} = 3,852.65, T_{B2}^{expected} = 3,393.90, p = .464.$$

Figure D1 shows the parameter estimates obtained from model $\mathcal{M}_1$; while estimates of the automatic process were only slightly above chance in both *PD instruction* conditions, estimates of the controlled process differ strongly between *PD instruction* conditions.

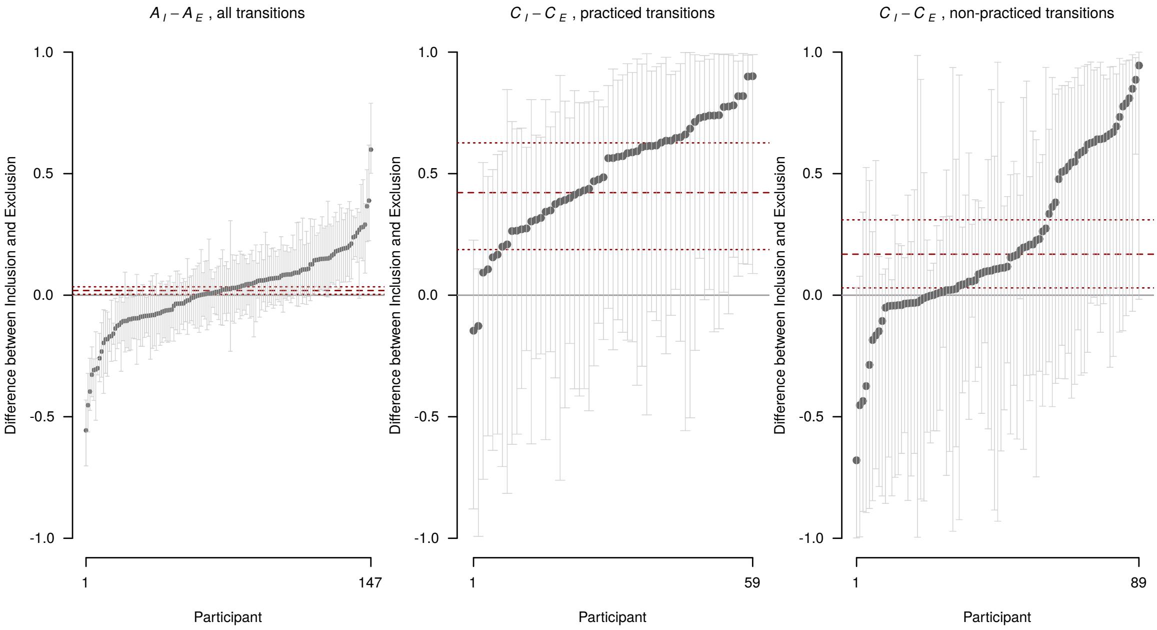

Figure D2. Posterior differences between $A_I - A_E$ and $C_I - C_E$ in Experiment 1, plotted for each participant (gray dots) with 95% credible intervals. Dashed lines represent the posterior means of the differences between mean parameter estimates. Dotted lines represent 95% credible intervals.

Figure D2 shows that the invariance assumption for automatic processes was violated with $A_I > A_E$, 95% CI [.00, .03], and Bayesian $p = .008$. For revealed and practiced transitions, the invariance assumption was violated with $C_I > C_E$, 95% CI [.19, .63] and a Bayesian $p = .001$. For revealed but non-practiced transitions, the invariance assumption was violated with $C_I > C_E$, 95% CI [.03, .31] and a Bayesian $p = .005$.

Experiment 2, model $\mathcal{M}_{1R}$

To test whether our results are robust against changes in auxiliary assumptions, we fitted another model $\mathcal{M}_{1R}$ with different auxiliary assumptions. Specifically, we dropped the assumption that $C = 0$ for nonrevealed transitions and instead estimated explicit-knowledge parameters for all transitions. Instead, we imposed ordinal restrictions (Knapp & Batchelder, 2004) as follows: In model $\mathcal{M}_{1R}$, it is assumed that C parameters are greater under inclusion than exclusion. We also fitted a parallel model with the reversed assumption, but estimation of this model failed to converge.

The second-level equations of model $\mathcal{M}_{1R}$ are given by:

$$\begin{aligned} C_{ij1} = C_{ij,Inclusion} &= \Phi(\mu^{(C)}_{jk,Inclusion} + \delta^{(C)}_{ij,Inclusion}) \\ C_{ij2} = C_{ij,Exclusion} &= \Phi(\mu^{(C)}_{jk,Exclusion} + \delta^{(C)}_{ij,Exclusion}) * C_{ij,Inclusion} \end{aligned}$$

and

$$A_{ijm} = \Phi(\mu^{(A)}_{jkm} + \delta^{(A)}_{ijm})$$

$\mu^{(C)}_{jkm}$ is the fixed effect of material k (that participant i worked on during the SRTT), *transition type* j ($j = 1$ if a transition has actually been revealed, $j = 2$ if not), and *PD instruction* condition m on controlled processes. $\delta^{(C)}_{ijm}$ is the ith participant's deviation from the respective group mean. For participants who did not receive explicit knowledge about a single transition, we assumed that all $\mu^{(C)}_{jk,Inclusion} = \mu^{(C)}_{k,Inclusion}$ and $\mu^{(C)}_{jk,Exclusion} = \mu^{(C)}_{k,Exclusion}$, i.e. we assumed that the grand mean of explicit knowledge did not vary as a function of the transition that *would* have been revealed if participants *were* in another condition. Accordingly, $\mu^{(A)}_{jkm}$ is the fixed effect of transition type j ($j = 1$ for the transition that was or *would* have been revealed, i.e. transition 2–6, $j = 2$ for all other transitions), material k, and *PD instruction* condition m on automatic processes, and $\delta^{(A)}_{ijm}$ is the ith participant's deviation from the corresponding mean.

Note that this specification imposes two auxiliary assumptions to the model: First, it is assumed that

$$\forall ij(C_{ij,Inclusion} \geq C_{ij,Exclusion})$$

Second, it is assumed that automatic processes A do not vary as a function of the between-subjects manipulation of explicit knowledge l (both assumptions were necessary so that the model was identified; an alternative model imposing an order constraint $C_I < C_E$ was also not identified).

Results. The model checks for model $\mathcal{M}_{1R}$ were satisfactory,

$$T^{observed}_{A1} = 484.60, T^{expected}_{A1} = 470.11, p = .409,$$

$$T^{observed}_{B1} = 9.13, T^{expected}_{B1} = 6.88, p = .358.$$

and attained a DIC value of 25,294.53, a value comparable to our extended model $\mathcal{M}_1$ presented in the main text and clearly outperforming $\mathcal{M}_2$. This again implies that our auxiliary assumptions introduced to $\mathcal{M}_{1R}$ were much less problematic than the invariance assumption.

Figure D3 shows the parameter estimates obtained from model $\mathcal{M}_{1R}$. The pattern of results mostly replicates the estimates from model $\mathcal{M}_1$. The main difference was that C parameters were slightly greater than zero for nonrevealed transitions (these were set to zero for model $\mathcal{M}_1$). This may suggest that some explicit knowledge may have been acquired during the learning phase. Alternatively, it may also reflect a technical issue with the

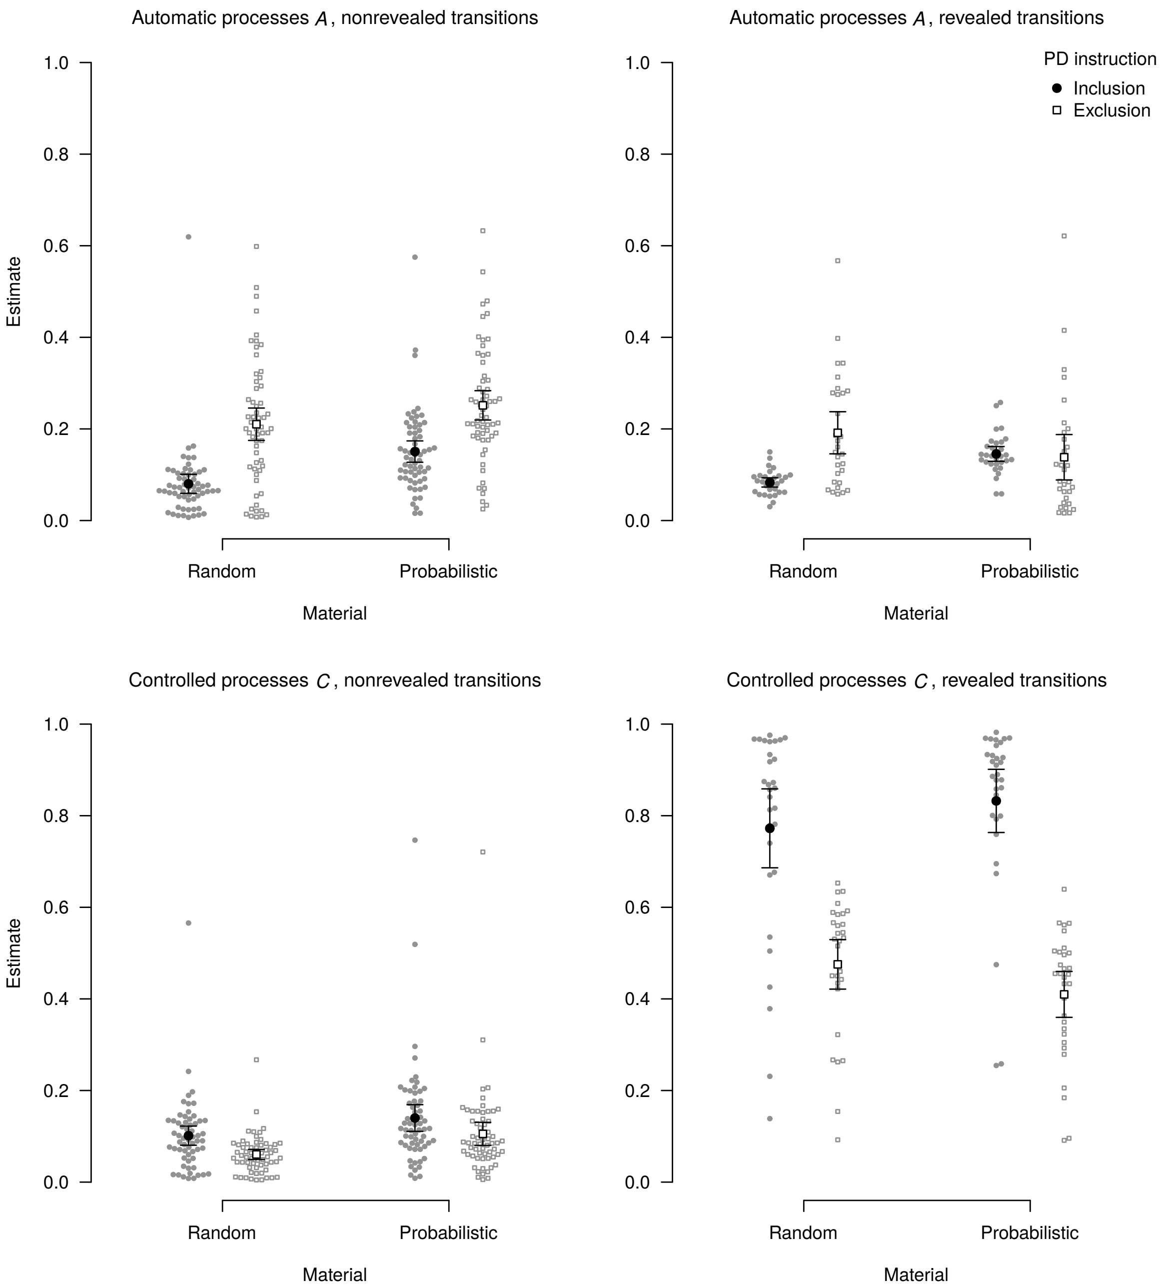

Figure D3. Parameter estimates from Experiment 2, model $\mathcal{M}_{1R}$. Error bars represent 95% confidence intervals.

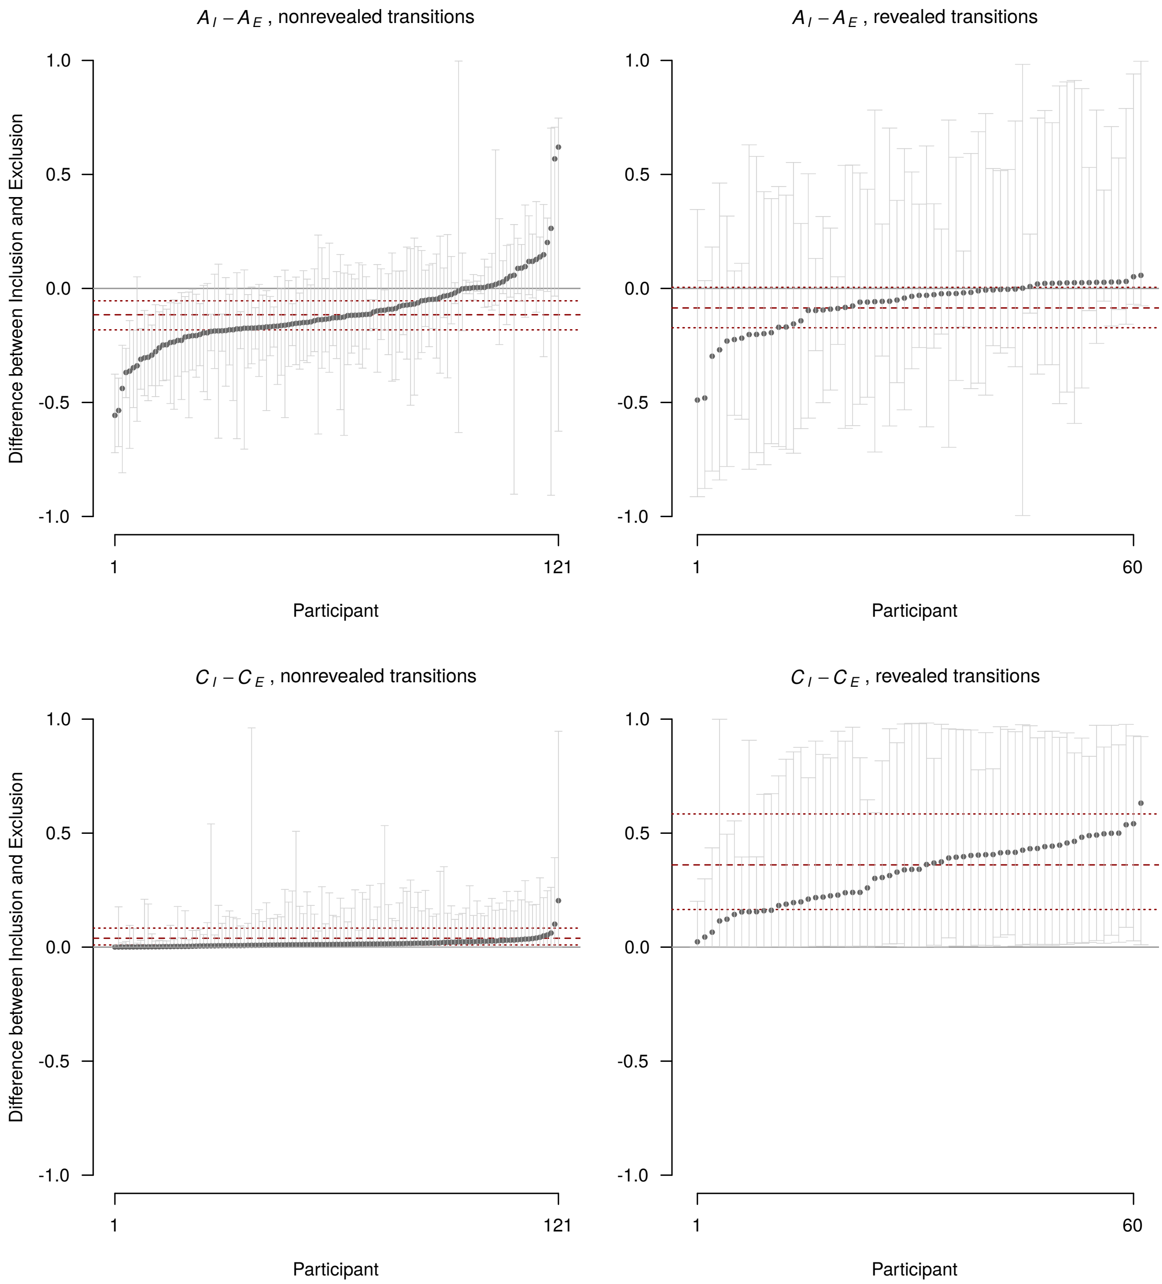

Figure D4. Posterior differences between $A_I - A_E$ and $C_I - C_E$ in Experiment 2, model $\mathcal{M}_{1R}$, plotted for each participant (gray dots) with 95% credible intervals. Dashed lines represent the posterior means of the differences between mean parameter estimates. Dotted lines represent 95% credible intervals.

present family of models that biases estimates away from zero: Specifically, for nonrevealed transitions, the inclusion-exclusion difference in C estimates should vary around zero, with half below zero and half above zero; the auxiliary assumption however forces all of them to be positive, which biases the corresponding C parameters. Either way, the effect is not substantial, as suggested by the finding that model $\mathcal{M}_1$, which assumes $C = 0$, achieved an equally good fit. The $C > 0$ estimates also have a tradeoff effect on A parameters, with lower estimates under inclusion and slightly higher estimates under exclusion. This biasing effect eliminated (for revealed transitions) or even inverted (for nonrevealed transitions) the invariance-violation effect found in $\mathcal{M}_1$.

Figure D4 shows the posterior differences obtained from model $\mathcal{M}_{1R}$. Most importantly, the pattern of results shows that the invariance violation for controlled processes C for revealed transitions (i.e., whenever substantial explicit knowledge is present) is robust to the change in auxiliary assumptions.

Experiment 3, model $\mathcal{M}_{1R}$

For the data of Experiment 3, we additionally fitted model $\mathcal{M}_{1R}$ analogous to $\mathcal{M}_{1R}$ of Experiment 2.

Results. The model checks for model $\mathcal{M}_{1R}$ were satisfactory,

$$T_{A1}^{observed} = 689.87, T_{A1}^{expected} = 657.24, p = .314,$$

$$T_{B1}^{observed} = 8.94, T_{B1}^{expected} = 6.02, p = .263.$$

and attained a DIC value of 38,881.68, a value somewhat smaller than the DIC of our extended model $\mathcal{M}_1$ presented in the main text and clearly outperforming $\mathcal{M}_2$. This again implies that our auxiliary assumptions introduced to $\mathcal{M}_{1R}$ were much less problematic than the invariance assumption.

Figure D5 shows the parameter estimates obtained from model $\mathcal{M}_{1R}$. The pattern of results mostly replicates the estimates from model $\mathcal{M}_1$; with parameters for controlled processes C being estimated close to zero for nonrevealed transitions.

Figure D6 shows the posterior differences obtained from model $\mathcal{M}_{1R}$. The pattern of results again demonstrates robustness of the invariance violation for controlled processes C for revealed transitions (i.e., whenever substantial explicit knowledge was present). There was again some indication of an invariance violation for automatic processes A; however, the effect was very small and depended on the specific modeling assumptions.

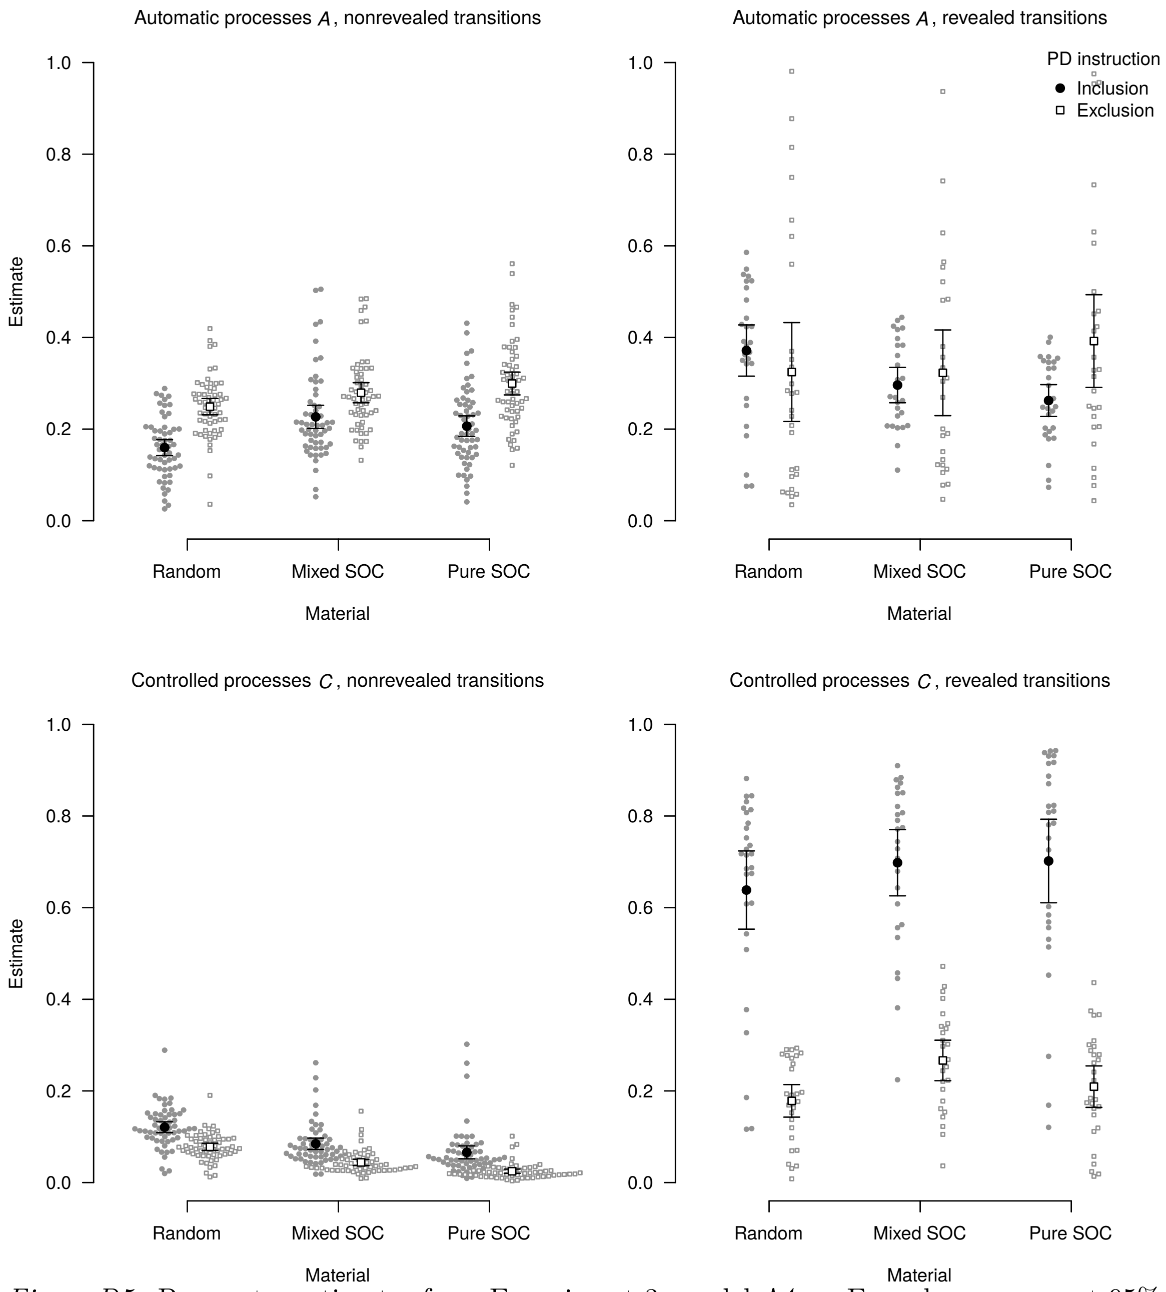

Figure D5. Parameter estimates from Experiment 3, model $\mathcal{M}_{1R}$. Error bars represent 95% confidence intervals.

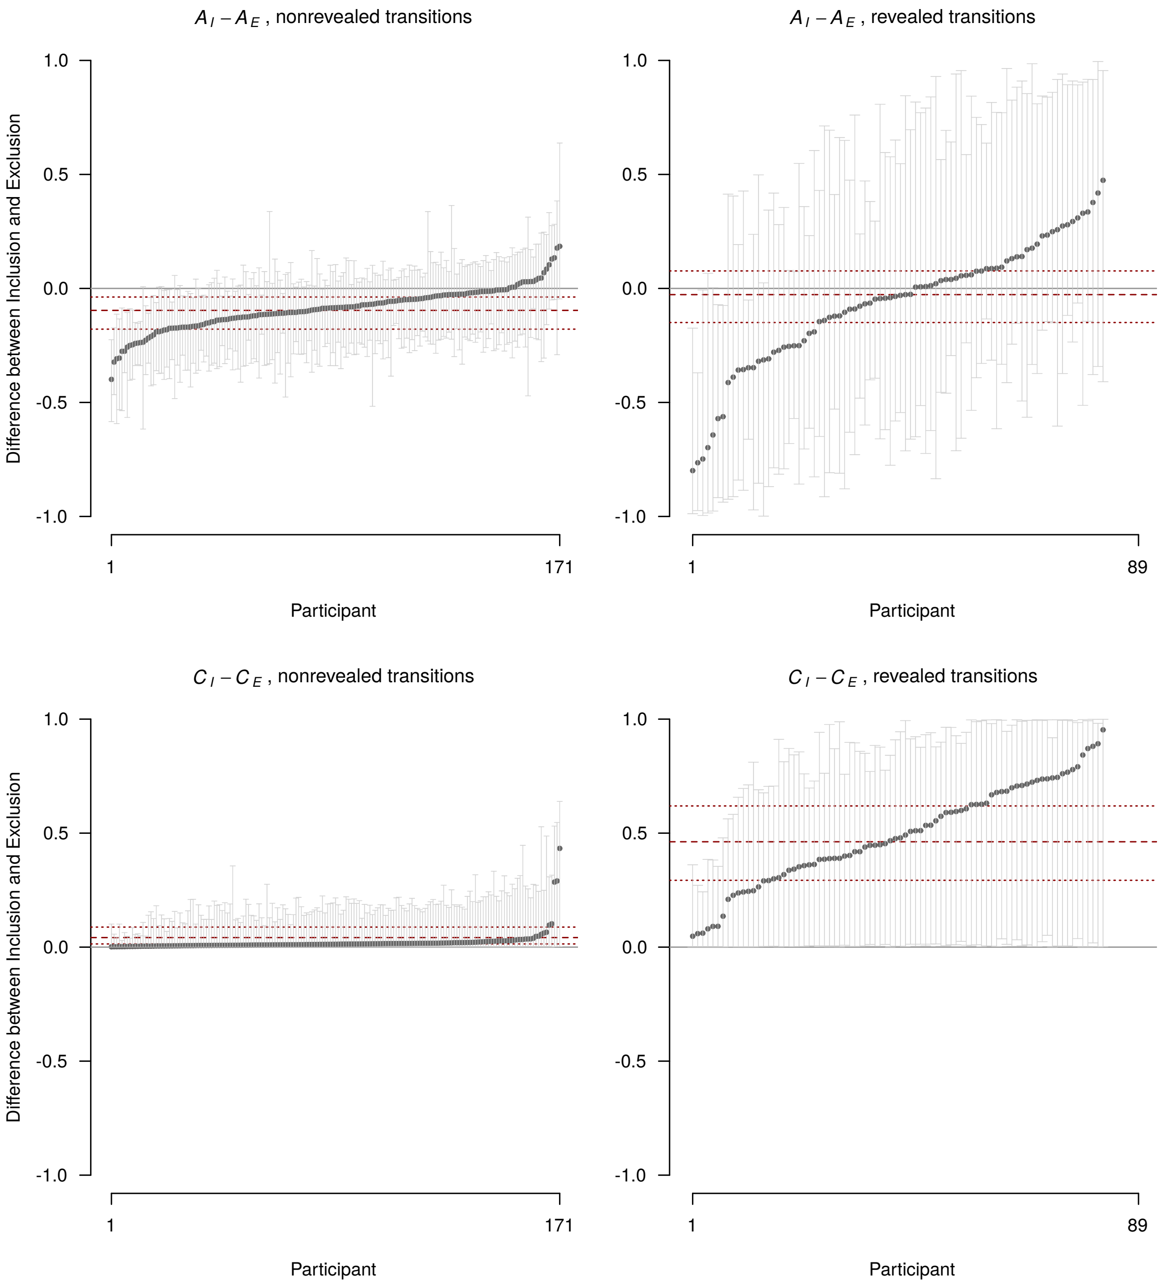

Figure D6. Posterior differences between $A_I - A_E$ and $C_I - C_E$ in Experiment 3, model $\mathcal{M}_{1R}$, plotted for each participant (gray dots) with 95% credible intervals. Dashed lines represent the posterior means of the differences between mean parameter estimates. Dotted lines represent 95% credible intervals.

Appendix E
Specification of priors

This section provides a complete specification of the models and priors used. Code (**R**/***Stan***) is available at https://github.com/methexp/pdl2.

Experiment 1, model $\mathcal{M}_1$

Priors on fixed effects were

$$\mu^{(C)}_{jlm} \sim N(0,1), j = \{1,2\}; l = \{1,2\}; m = \{1,2\}$$
$$\mu^{(A)}_{mt} \sim N(0,1), t = \{1,...,6\}; m = \{1,2\}$$

where j indexes *transition type* (revealed & practiced vs. revealed & non-practiced), l indexes practice condition (Control, No-practice, Unspecific-practice, Practice, Transfer), t indexes specific items (i.e., transitions), and m indexes *PD instruction* (inclusion vs. exclusion). Participant effects $\delta^{(A)}_{imt}$ and $\delta^{(C)}_{ijm}$ can be written as vectors $\boldsymbol{\delta}_i$. For participants in the *Control* group, these were modeled by

$$\boldsymbol{\delta}_i \sim N_{12}(0, \Sigma_l), i = 1, ..., I$$

For participants in the *No-Practice*, *Unspecific-Practice*, and *Practice* groups,

$$\boldsymbol{\delta}_i \sim N_{14}(0, \Sigma_l), i = 1, ..., I$$

For participants in the *Transfer* group

$$\boldsymbol{\delta}_i \sim N_{16}(0, \Sigma_l), i = 1, ..., I$$

The covariance matrices Σ_l were modeled separately and independently for each between-subjects condition. Priors on these matrices were as described below for Experiment 2.

Experiment 2, model $\mathcal{M}_1$

Priors on fixed effects were

$$\mu^{(C)}_{km} \sim N(0,1), k = \{1,2\}; m = \{1,2\}$$
$$\mu^{(A)}_{jkm} \sim N(0,1), j = \{1,2\}; k = \{1,2\}; m = \{1,2\}$$

where j indexes transition type (revealed vs. non-revealed), k indexes learning material presented during the SRTT (random vs. probabilistic), and m indexes *PD instruction* condition (inclusion vs. exclusion). For participants who did not receive explicit knowledge about a single transition, we assumed that all $C_{ijkm} = 0$. Therefore, participant effects are only required for automatic processes ($\delta^{(A)}_{ijkm}$). In participants who received explicit knowledge about one transition, two additional participant effects were needed to model controlled processes for revealed transitions ($\delta^{(C)}_{ikm}$). We thus provide the specification of participant effects for these two groups of participants separately.

Participants who did not receive explicit knowledge about one transition.

For participants who did not receive explicit knowledge about one transition, participant effects $\delta^{(A)}_{ijm}$ can be written as vectors $\boldsymbol{\delta}_i$ that were modeled as draws from a multivariate normal

$$\boldsymbol{\delta}_i \sim N_4(0, \Sigma_{kl}), i = 1, ..., I$$

where k indexes the learning material that was presented to participant i and l indexes his or her level of the explicit-knowledge factor. The covariance matrices Σ_{kl} were obtained from the standard deviations of participant effects $\boldsymbol{\sigma}_{kl}$ and correlation matrices Ω_{kl}

$$\Sigma_{kl} = Diag(\boldsymbol{\sigma}_{kl})\ \Omega_{kl}\ Diag(\boldsymbol{\sigma}_{kl}), k = \{1, 2\}, l = \{1, 2\}$$

Each element σ_{klp} of the vectors of standard deviations $\boldsymbol{\sigma}_{kl}$ was drawn from independent half-normal prior distributions.

$$\sigma_{klp} \sim N(0, 1)_{\mathcal{I}(0,\infty)}, k = \{1, 2\}, l = \{1, 2\}$$

For the correlation matrices Ω_k, we used LKJ priors with a scaling factor of 1 (Lewandowski, Kurowicka, & Joe, 2009):

$$\Omega_{kl} \sim LKJcorr(\nu = 1), k = \{1, 2\}, l = \{1, 2\}$$

Participants who received explicit knowledge about one transition.

For participants who received explicit knowledge about one transition, participant effects $\delta^{(A)}_{ijm}$ and $\delta^{(C)}_{im}$ can be written as vectors $\boldsymbol{\delta}_i$ that were modeled as draws from a multivariate normal

$$\boldsymbol{\delta}_i \sim N_6(0, \Sigma_{kl}), i = 1, ..., I$$

where k indexes the learning material that was presented to participant i and l indexes his or her level of the explicit-knowledge factor. The covariance matrices $\Sigma_k l$ were specified as above, with the only exception that six instead of four parameters were required.

Experiment 2, model $\mathcal{M}_2$

Priors on fixed effects were

$$\mu_{jkl}^{(C)} \sim N(0,1), j=\{1,2\}; k=\{1,2\}; l=\{1,2\}$$
$$\mu_{jkl}^{(A)} \sim N(0,1), j=\{1,2\}; k=\{1,2\}; l=\{1,2\}$$

Participant effects $\delta_{ij}^{(A)}$ and $\delta_{ij}^{(C)}$ can be written as vectors $\boldsymbol{\delta}_i$ that were modeled by

$$\boldsymbol{\delta}_i \sim N_4(0, \Sigma_{kl}), i=1,...,I$$

Priors for the covariance matrix Σ_{kl} were specified as above.

Experiment 2, model $\mathcal{M}_{1R}$

Priors on fixed effects were

$$\mu_{jkm}^{(C)} \sim N(0,1), j=\{1,2\}; k=\{1,2\}; m=\{1,2\}$$
$$\mu_{jkm}^{(A)} \sim N(0,1), j=\{1,2\}; k=\{1,2\}; m=\{1,2\}$$

where j indexes transition type (revealed vs. non-revealed), k indexes learning material presented during the SRTT (random vs. probabilistic), and m indexes *PD instruction* condition (inclusion vs. exclusion). Participant effects $\delta_{ijm}^{(A)}$ and $\delta_{ijm}^{(C)}$ can be written as vectors $\boldsymbol{\delta}_i$ that were modeled as draws from a multivariate normal

$$\boldsymbol{\delta}_i \sim N_8(0, \Sigma_{kl}), i=1,...,I$$

where k indexes the learning material that was presented to participant i and l indexes his or her level of the explicit-knowledge factor. Priors for the covariance matrix Σ_{kl} were specified as above.

Experiment 3, models $\mathcal{M}_1$, $\mathcal{M}_2$, and $\mathcal{M}_{1R}$

For the model-based analyses, we used models $\mathcal{M}_1$, $\mathcal{M}_2$, and $\mathcal{M}_{1R}$ analogous to those used in Experiment 2.