Article

Why Experts Make Errors

Itiel E. Dror
David Charlton

School of Psychology
University of Southampton
Southampton, United Kingdom

Abstract: Expert latent fingerprint examiners were presented with fingerprints taken from real criminal cases. Half of the prints had been previously judged as individualizations and the other half as exclusions. We re-presented the same prints to the same experts who had judged them previously, but provided biasing contextual information in both the individualizations and exclusions. A control set of individualizations and exclusions was also re-presented as part of the study. The control set had no biasing contextual information associated with it. Each expert examined a total of eight past decisions. Two-thirds of the experts made inconsistent decisions. The findings are discussed in terms of psychological and cognitive vulnerabilities.

Background

Expert performance and accuracy is an important issue in almost all specialized domains. In contrast to novices, experts possess abilities and skills that enable them to perform certain tasks, such as medical procedures or flying aircraft [1]. An expert needs not only knowledge, but skill, judgement, and experience to evaluate and interpret information correctly to make correct decisions. However, being an expert does not necessarily mean error-free performance; in fact, almost every specialist domain is subject to error. The pivotal question is the source of the error. Errors, broadly speaking, fall into three categories.

Received December 9, 2005; accepted February 6, 2006

Errors

The first category relates to human error. Human errors can be intentional errors (whereby experts are involved in fraudulent behavior), negligent errors (whereby experts do not pay attention, do not follow procedures, etc.), and competency errors (whereby experts are unable to make correct judgements because of a lack of appropriate skill sets; this can be due to declining eyesight, faulty initial selection that results in recruiting people who do not possess the proper cognitive abilities that are needed to underpin the expertise, inadequate training, and so forth).

The second category relates to instrumentation and technological errors. In this category, errors derive from failure and breakdown of instruments and technology. These types of errors are rare in the fingerprint domain. Technical malfunction certainly accounts for errors in other domains, such as breakdown of equipment on aircraft. The first and second categories both relate to chance malfunctions and breakdown, either human or machine.

The third category of error relates to more fundamental methodological factors that are inherent to the field in question. These may include errors associated with the technology in question, instrumentation, and measurements. In the fingerprint domain, for example, a failure of the Automated Fingerprint Identification System algorithms to provide a matching file from the database (assuming such a file is present and enough information exists to make the match) can lead to an error.[1] This, in contrast to category two, reflects not a malfunction or a bug in the software, but the inherent inaccuracy of the algorithm. Technologies and instrumentations have their limits, range of accuracy, levels of precision, variations, and so forth, which are not due to their breakdown and malfunction but to their very nature.

[1] An error is sometimes defined only as an incorrect individualization (i.e., a false positive, what is referred to in signal detection theory as a false alarm). However, such a definition is limited because it does not include cases where there is sufficient information to make a positive individualization, but, because of the error, no such individualization is made (i.e., a false negative, what is referred to in signal detection theory as a miss). These two types of errors are different, but nevertheless, both can be regarded as errors.

However, such errors are not limited to technology or instrumentation, especially in a domain like fingerprints where much of the individualization process falls on human experts and their interaction with technology [2, 3]. Here, the error is due not to the nature of the technology and instrumentation but to the nature and mechanisms of the human mind and cognition. This is particularly noticeable when dealing with latent prints that are collected from crime scenes and are thus degraded, contaminated, partially missing, and distorted.

This third category thus includes errors that are not simple practitioner error that can be attributed to the specific expert involved (as specified and belonging to category one). When practitioners are competent, well trained, and following procedures, when instrumentation and technology operate properly, and errors happen nevertheless, then these errors belong to category three.

This paper examines the possible role that psychological and cognitive factors may have in causing these types of errors. *When expert practitioners perform well and technology is effective, can errors still occur? And if so, why and how?* Some people attempt to dismiss this possibility a priori, claiming that an error results either from a practitioner's error (such as those specified in category one) or from the lack of scientific basis for the domain (such as the uniqueness of fingerprints).

This dichotomized attribution of an error as either reflecting a basic scientific flaw in the domain or a specific practitioner's error fails to consider a third alternative: errors that derive from psychological and cognitive elements involved in fingerprint individualization. This process falls on human experts, allowing the possibility that errors may result from the way the brain processes information and makes decisions.

If the nature of the mind and cognitive processing can give rise to error in fingerprint individualization, then these errors are inherent to the domain. Nevertheless, they do not reflect a basic ontological scientific flaw in the domain nor are they the fault of a specific practitioner. They are, in essence, epistemological problems that derive from the mechanisms of human cognition and the workings of the mind.

As with technological and instrumentation advances that improve their limits, accuracy, and levels of precision, so can human performance be improved with correct selection, training, and procedures. However, such endeavors need to be based on systematic and scientific research, and even then they will not totally eliminate human error of category three. Nevertheless, with such research, these errors can be drastically minimized, so minimized that although they are theoretically possible, they are in fact so very rare that de facto they do not exist.

The above division into three categories of error is a simplistic model for methodological reasons; reality can be more complex. For example, because practitioners and methodology are so intertwined, it is difficult (perhaps sometimes impossible) to separate and distinguish between the two [2], thus making it problematic to attribute an error solely to either category one or three. For instance, if errors occur because a practitioner is incompetent, but his or her incompetence is due to basic flaws in the domain, then the errors are not purely in the practitioner error of category one. Practitioners' incompetence may arise from the lack of appropriate scientifically based screening tests for recruitment and certification of fingerprint examiners and thus may reflect deeper flaws in the domain. To date, there is no systematic scientific research into the psychological and cognitive skills that underlie fingerprint expertise. Research is needed to construct appropriate tests for recruitment screening and selection of fingerprint examiners. Therefore, in some cases, an expert's incompetence may also be attributed to more basic flaws in the domain rather than purely to the individual practitioner.

Understanding the source of errors and their assignment to one of the categories can be highly insightful and have important implications. For example, the Mayfield erroneous individualization [4] raises interesting issues. Would such a mistake be totally attributable to practitioner error? Or should the error be attributed also, at least in part, to the lack of appropriate procedures, training, and quality assurance to address and deal with the causes of such errors? These types of questions are critically important to investigate to allow advances in this and other domains of forensic science.

However, researching and discussing errors is a problematic and challenging endeavor. First, by its very nature, "error" is a sensitive issue that often meets defensive responses. This is especially true in the criminal justice system which deals with incarcerating and even executing people. Second, the framework of the criminal justice system does not enable the "ground truth" to be positively known. Thus, individualization cases can always be open to suspicions as being erroneous. A recent paper by Cole [5] tries to catalogue such suspected erroneous individualizations. Third, even when errors are detected and acknowledged, their classification and examination are constructed post hoc in a highly political and personal environment.

The study presented here examines whether inherent psychological and cognitive mechanisms predispose fingerprint and other forensic identification experts to commit category three errors [6, 7]. This important area of research has been highly neglected. Some of the psychological phenomena and mechanisms that may underlie such errors are derived from the nature and architecture of the human mind [8]. It is imperative to conduct empirical experimental research to examine whether such errors actually do exist and, if so, to find ways to minimize these errors.

To examine such issues, it is important to conduct scientific studies within a real world setting. This is challenging because when people know they are being studied, their behavior and performance change, and thus puts into question the applicability and ecological validity of the findings. If you want to know how people drive, then their performance during a driving test is not very insightful and revealing, neither is their driving when they know they are near speed cameras or radars. One must try to observe and examine performance as well as collect data in the normal routine setting with minimal (or no) knowledge of the people involved. This of course if very difficult to achieve and necessarily results in small data sets. However, these data sets are statistically very powerful, meaningful, and more interpretable.

Previous Research

In our empirical studies [9 and 10], we have started to address issues relating to errors that derive from psychological and cognitive influences. Emotional states (Figure 1) have been shown to cause nonexperts to be more likely to "match... ambiguous pairs of fingerprints" when they performed the comparison in a highly emotional context [9].

Figure 1

An example of an emotional context used in the previous (nonexpert) study.

In a previous study [10], five experts were shown fingerprints and were told that the prints were from a highly publicized erroneous identification [4], suggesting that the fingerprints in front of them were an exclusion. However, rather than giving them prints that were an exclusion, we presented them with fingerprints that had been compared and individualized. The fingerprints were not only individualizations, but they had also been previously individualized by the same experts now being tested. Although the experts were instructed to "ignore all the contextual information and to focus solely on the actual prints", most of the experts (four of the five) were affected by the context and made inconsistent decisions. In the above study, we were able to demonstrate the vulnerability of experts to extraneous contextual information, but only when we subjected them to a relatively extreme context and when we presented them with difficult matches.

Current Research

The study reported within this article follows up on previous studies, further examining psychological and cognitive influences that may play a role in the work of fingerprint experts. The specific purposes of this current study were:

1. To determine whether the original findings would replicate to another and larger set of data.
2. To use less extreme and more routine day-to-day contextual influences.
3. To use pairs of prints of varying levels of difficulty.
4. To examine the possible influence of contextual information on different decision types (thus, not only whether it can change a past individualization decision to an exclusion, but also to examine whether it can change a past exclusion decision to an individualization).
5. To examine the basic consistency of decisions by re-presenting to experts the same fingerprints they judged in the past but without introducing any contextual information manipulation, thus examining the reliability of experts.

Method

Participants

Six fingerprint experts, representing more than 35 years of experience in examining fingerprints (each with a minimum of 5 years' experience in latent prints), participated. The experts were experienced and specialized in latent print examination and were not field operatives; hence, their experience was full-time in latent print comparisons. Each of the expert participants was not only highly experienced, but was highly trained, certified by a nationally recognized independent authority, and had successfully completed proficiency testing. None of them had been the subject of a poor competency review, and they were all considered by their respective laboratory directors or bureau chiefs to be effective and competent. The participants were taken from our international fingerprint expert pool of volunteers. This pool includes fingerprint experts from across the world (including the USA, UK, Israel, the Netherlands, and Australia).

We used experts whose past work we could covertly access. Because the collation of individual casework is so difficult in such circumstances, and because of the difficulties in covert testing, this effectively limited the number of participants within this study. This enabled us, however, to collect more meaningful and powerful data. This derives not only from the covert nature of our study but also from the use of a within-subject experimental design in which participants are used as their own control through repeated measures [11]. Thus, variations are more easily attributable to the experimental manipulations and conditions rather than individual differences between people. This provides clearer and more interpretable results, as well as more statistical power for each data point [11].

Materials

A different and unique set of eight pairs of fingerprints was prepared and tailored for each of the participants. Each set included four pairs of prints that the specific fingerprint expert had in the past judged as individualizations and four pairs of prints that he or she had judged in the past as exclusions. All of the eight pairs of fingerprints had been deemed in the past by the specific participant to have sufficient information to make definite judgements. Within each of the four past individualizations and the four exclusions, two pairs of prints were relatively difficult to judge, and the remaining two pairs were relatively not difficult to judge.

The latent fingerprints had all been obtained from real crime scenes and were all presented to the participants in their original format for comparison against suspect tenprint exemplars. It should be stressed that all the pairs of prints that we used were fingerprint comparisons that we obtained from the archives and had been evaluated some years before by the very same experts. The within-subject design of the experiment was deemed vital to the overall robustness and credibility of the findings. Two additional experienced fingerprint experts who did not take part in the study (each had more than 20 years of experience) independently confirmed and verified that all the pairs of fingerprints were indeed either correct individualizations or correct exclusions. They also characterized the pairs of prints as either relatively difficult to evaluate or as relatively not difficult to evaluate.

An instruction and response sheet was prepared for each of the eight pairs of fingerprints. Four pairs of prints were used as controls and provided no contextual information. These four pairs of control prints included two pairs that had been judged in the past as individualizations and two pairs that had been judged in the past as exclusions. Of each of these two pairs, one was relatively difficult to judge and one was relatively not difficult to judge. The instruction sheet that was given with the control pairs included minimal instructions that told the experts to evaluate the prints. A response sheet was provided for participants to write their conclusions. In the response sheet, the experts were first asked whether there was "sufficient information within the prints to either identify or exclude". If the answer was, "no", then they finished with this pair of prints and moved on to the next pair. If the answer was, "yes", then the experts continued, stated what their decisions were, and then moved on to the next pair of prints.

Of the remaining four pairs of prints, two pairs had been judged in the past as individualizations and two had been judged in the past as exclusions. These four pairs of prints were presented within contextually biasing information that was hypothesized to influence the conclusions reached by the experts. The two pairs of prints that had been judged in the past as individualizations were presented in a context that suggested that they were exclusions (one of the pairs was relatively difficult to individualize, and the other was not). Similarly, the two pairs of prints that had been judged in the past as exclusions were presented in a context that suggested that they were individualizations. (Again, one of the pairs was relatively difficult to individualize, and the other was not.)

In contrast to our previous study [10], we now used more subtle, routine, day-to-day, contextually biasing information. In the instructions the participants were told, for example, that the "suspect confessed to the crime" (for contextual information that the prints were an individualization, when in fact they were not and had been judged by the same expert as an exclusion in the past) or that the "suspect was in police custody at the time of the crime" (for contextual information that the prints were not an individualization, when in fact they were and had been judged by the same expert as an individualization in the past).

Thus we prepared a total of 48 unique experimental trials. Each one consisted of a latent print from a crime scene and a suspect tenprint exemplar, accompanied with the proper instructions and response sheets. For each expert, we prepared a customized folder containing the eight pairs of fingerprints that they themselves had judged in the past. The eight pairs of fingerprints in each folder were counterbalanced. Counterbalance presentations assure that results are not due to order affects and cross-contamination between the different conditions, because the order of presentation is systematically permutated across participants [11].

Procedure

Participants were approached by the director or head of the laboratory or bureau and were asked to provide opinions on a variety of latent prints and their comparisons to tenprint exemplars. They were told that the conclusions they reached after the examination would be used for an assessment project. They were further told that the project was intended to look at problematic prints and assessments.

Then, the first assessment, along with the instruction and response sheet, was given to the participants. After they finished the comparison and documented their conclusions on the response sheet, the materials were put back into the original folder, and the second assessment was given to them. This continued until all eight assessments were completed.

During the comparison process, all of the participants were allowed to evaluate the prints as they would do routinely (handling the prints, using magnifying and lighting equipment, and so forth). The participants were allowed an unlimited amount of time and all normal resources (e.g., comparators) to make their evaluations.

Results

Overall, from 48 experimental trials, the fingerprint experts changed their past decisions on six pairs of fingerprints (Table 1, Figure 2). The six inconsistent decisions (12%) included the 24 control trials that did not have any contextual manipulation. From the 24 experimental trials that included the contextual manipulation, the fingerprint experts changed four of their past decisions, thus making 16.6% inconsistent decisions that were due to biasing context. The inconsistent decisions were spread between the participants. (The inconsistent decisions were by four of the six experts, but one expert made three inconsistent decisions while each of the other three made only one inconsistent decision.) Only one-third of the participants (two out of six) remained entirely consistent across the eight experimental trials.

	1	2	3	4	5	6	7	8
Past Decision	individualization	individualization	individualization	individualization	exclusion	exclusion	exclusion	exclusion
Level of Difficulty	difficult	difficult	not difficult	not difficult	difficult	difficult	not difficult	not difficult
Contextual Information	none	suggest exclusion	none	suggest exclusion	none	suggest individualization	none	suggest individualization
Expert A	consistent	consistent	consistent	consistent	consistent	consistent	consistent	consistent
Expert B	**change to exclusion**	consistent	consistent	consistent	consistent	consistent	consistent	consistent
Expert C	consistent	**change to exclusion**	consistent	consistent	consistent	consistent	consistent	consistent
Expert D	consistent	**change to exclusion**	consistent	**change to exclusion**	**change to individualization**	consistent	consistent	consistent
Expert E	consistent	**change to cannot decide**	consistent	consistent	consistent	consistent	consistent	consistent
Expert F	consistent	consistent	consistent	consistent	consistent	consistent	consistent	consistent

Table 1

Results of fingerprint examinations.

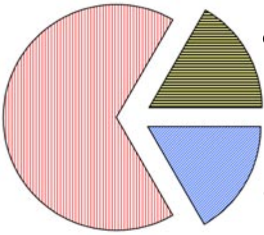

Figure 2

Breakdown of inconsistent decisions.

A further examination of the inconsistent decisions revealed that they were most common in the more difficult decisions. In most cases (five out of six), the inconsistent decisions were made in the difficult comparisons. Nevertheless, inconsistent decisions were not totally limited to the difficult comparisons; one of the six was in the relatively not difficult comparisons.

The inconsistent decisions were most prevalent in past individualization decisions (five out of the six inconsistent decisions; they were changed to either exclusion, in four cases, or to cannot decide, in one case). However, a conflicting judgment also occurred in a past exclusion decision (now changed to an individualization).

Finally, two inconsistent decisions were present in the control condition that had no biasing contextual information. These decisions constituted 33.3% of the conflicting data (two out of six) and 8.3% of the control data (two out of 24).

Further statistical analyses and procedures (inferential statistics, effect sizes that directly computed retest reliability r_{sample}, and $r_{equivalent}$ that specifically takes into account the sample size, as well as meta-analytic measures) are all considered in more statistical detail and technical fashion elsewhere [12].

Discussion

Experts, just like all humans, are bound by the way the mind and the brain work. One of the most notable characteristics of humans is the active and dynamic nature in which we process information [8]. This fundamental cognitive architecture enables us to deal with vast amounts of information and is the basis of human intelligence. However, it can also affect and distort what we see (or why and what we do not see), how we assess and evaluate visual information, and our decision making [7, 13].

Fingerprint and other forensic experts are not immune to such psychological and cognitive factors. Researching and understanding them better is a necessary step toward dealing with and minimizing such influences.

Previous research has demonstrated that under conditions of relatively extreme and rare extraneous contextual information, fingerprint experts may change the way they compare and judge fingerprints [10]. The extraneous contextual information effect caused some experts to reach different and inconsistent decisions to those they had made in the past on the very same pair of fingerprints.

The research reported in the present study replicated and expanded our previous findings and found that across eight comparisons made by each of the six participants, two thirds of the fingerprint experts made inconsistent decisions to those they had made in the past on the same pairs of prints. The findings of this study not only further substantiate the vulnerability of experts to contextual effects within a larger data set, but they further contribute to our understanding of this phenomenon.

Our data demonstrate that fingerprint experts were vulnerable to biasing information when they were presented within relatively routine day-to-day contexts, such as corroborative (or conflicting) evidence of confession to the crime. Thus, contextual information does not need to be extreme and unique to influence experts in their fingerprint examination and judgement.

Varying the levels of difficulty in comparing the prints demonstrates that psychological and cognitive vulnerabilities are most pronounced in the difficult cases. However, our data also show that such vulnerabilities can also occur and cloud judgement in nondifficult cases, because our contradictory findings were not limited to only difficult comparisons.

Previous research has only examined whether past decisions are susceptible to change when the past decisions were individualizations [10]. It seems that the threshold to make a decision of exclusion is lower than that to make a decision of individualization. Indeed our data support this claim, as reflected by the fact that most of the conflicting decisions were past individualizations. We did, however, observe a case in which an exclusion decision was now judged to be an individualization. This relates to the decision-making model used by experts in the fingerprint domain. Changes in decisions may reflect changes in decision thresholds or changes of the decision strategy itself. The former reflects changes in decision criteria within a single strategy whereas the latter reflects modifying (or totally abandoning) the decision strategy and replacing it with a different strategy. Examining the decisions themselves and how they change can reveal which occurred; however, this needs to be done carefully, because changes in thresholds and decision strategies can yield similar (and even identical) decision changes [13].

Finally, quite surprising and alarming are the data points of inconsistent decisions made in our control "context-free" condition. This may reflect that expert decisions are inconsistent across time regardless of context. If this is indeed the case, then further research needs to examine the root of this observation: Is this due to the accumulation of expertise and experience gained between the first and second exposures (and this led to change in decision threshold or in decision strategy itself), or is some other explanation the cause? Such inconsistency does not suggest that contextual information does not affect the judgement of fingerprint experts; it only suggests that contextual information is not the only factor that may affect fingerprint experts. The control condition needs to be evaluated with a number of reservations: First, the control condition was not context free. We did not include any of our extraneous contextually biasing information; however, that does not make it context free. Second, even if we were able to achieve this ideal notion of a context-free environment, this context would still be different than that which was present during the first exposure and judgement years ago.

Our findings are especially robust because we employed a within-subject experimental design. Thus our findings do not reflect individual differences among experts. We further feel confident in our findings because they support the findings of our previous studies [9, 10] and thus the extra data further substantiate and validate our conclusions. This study was conducted covertly, which is critical for the correct measure of performance. When participants know they are participating in a study, their behavior changes [7]. These experimental design criteria make it hard to collect data and enable only small data sets, but the data are meaningful and statistically powerful [11, 12].

This entire area of research is new in the forensic sciences and has rarely been considered before. Therefore, such studies constitute the initiation of a research program that is aimed at examining the psychological and cognitive elements that are involved in fingerprint and other forensic identifications. We hope that these findings will contribute to better selection, better training, and better procedures for work in this domain. However, such an endeavor to deal with and minimize these vulnerabilities is dependent on the cooperation of fingerprint experts worldwide.

Further research is needed in a wide range of issues pertaining to individualization. Within the issues raised in this paper, additional data may shed light on the characteristics of the experts who were immune to our contextual manipulations, issues pertaining to the circumstances in which the manipulations were more (or less) effective, as well as additional issues. Research needs to examine psychological and cognitive influences in all the stages that lead to decisions in comparing fingerprints, from feature perception and selection, to evaluation and comparison, to the final interpretation and weighting of alternative choices that determine the decision outcome. Additional perspectives for future research in the fingerprint and other forensic domains relate to verification, selection and training of experts, and integration of technology.

There is no possibility of 100% objectivity [7, 8, 14], but there is potential for very high levels of objectivity. How rare and under what conditions errors occur at a practical level is still unclear at this stage. Experts, as humans, are prone to errors; however, with proper research and its systematic application, these errors can be reduced and minimized.

Acknowledgments

First, we would like to acknowledge and thank our collaborators in the fingerprint world, both those participating in our studies and those helping us in conducting our research. We would also like to thank Ailsa Peron for discussing different aspects of this work and Alan McRoberts, Robert Rosenthal, Arie Zeelenberg, Cedric Neumann, John Vanderkolk, and anonymous reviewers for comments on earlier versions of this paper.

For further information, please contact:

Dr. Itiel Dror
School of Psychology
University of Southampton
Southampton SO17 1BJ
England
United Kingdom
+44 23-80594519
id@ecs.soton.ac.uk
www.ecs.soton.ac.uk/~id/

References

1. Dror, I. E.; Kosslyn, S. M.; Waag, W. Visual-Spatial Abilities of Pilots. *J. Appl. Psychology* **1993**, *78* (5), 763-773.
2. Rudin, N.; Inman, K. Fingerprints in Print, the Sequel. *The CAC News*, 2nd Quarter 2005, pp 6-10.
3. Dror, I. E. Technology and Human Expertise: Some Do's and Don'ts. *Biometric Technology Today* **2005**, *13* (9), 7-9.
4. Stacey R. B. Report on the Erroneous Fingerprint Individualization in the Madrid Train Bombing Case. *J. For. Ident.* **2004**, *54* (6), 706-718.
5. Cole, S. A. More than Zero: Accounting for Error in Latent Fingerprint Identification. *J. Crim. Law & Criminology* **2005**, *95* (3), 985-1078.
6. Risinger, D. M.; Saks, M. J.; Thompson, W. C.; Rosenthal, R. The Daubert/Kumho Implications of Observer Effects in Forensic Science: Hidden Problems of Expectation and Suggestion. *Calif. Law Rev.* **2002**, *90* (1), 1-56.
7. Dror, I. E. Perceptual, Cognitive, and Psychological Elements Involved in Expert Identification. *Friction Ridge Sourcebook*, in press.
8. Dror, I. E. Perception is Far From Perfection: The Role of the Brain and Mind in Constructing Realities. *Brain and Behavioural Sciences* **2005**, *28* (6), 763.

9. Dror, I. E.; Peron, A.; Hind, S.; Charlton, D. When Emotions Get the Better of Us: The Effect of Contextual Top-Down Processing on Matching Fingerprints. *Appl. Cognitive Psychology* **2005**, *19* (6), 799-809.
10. Dror, I. E.; Charlton, D.; Peron A. Contextual Information Renders Experts Vulnerable to Making Erroneous Identifications. *For. Sci. Inter.* **2006**, *156* (1), 74-78.
11. Rosenthal, R.; Rosnow, R. L. *Essentials of Behavioural Research: Methods and Data Analysis*; McGraw-Hill: New York, 1991.
12. Dror, I. E.; Rosenthal, R. Meta-Analytically Quantifying the Reliability and Biasability of Fingerprint Experts' Decision Making. (Submitted for publication.)
13. Dror, I. E.; Busemeyer, J. R.; Basola, B. Decision Making Under Time Pressure: An Independent Test of Sequential Sampling Models. *Memory and Cognition* **1999**, *27* (4), 713-725.
14. Hofstadter, D. R. *Gödel, Escher, Bach: An Eternal Golden Braid*; Basic Books: New York, 1979.