University of Arkansas at Little Rock Law Review

Volume 42 | Issue 2 Article 2

2020

Document Review: You're Doing It Wrong Cognitive Psychology and the Attorney's Mental Plate

Robert Keeling

DOCUMENT REVIEW: YOU'RE DOING IT WRONG COGNITIVE PSYCHOLOGY AND THE ATTORNEY'S MENTAL PLATE

Robert Keeling[*]

ABSTRACT[**]

The review of documents in response to discovery requests is more expensive, burdensome, and time-consuming than ever before. This is largely due to a drastic increase in the volume of electronically stored information. In larger cases, document sets contain hundreds of thousands (if not millions) of documents. A manual, "eyes-on" review is often required for many, if not all, of these documents. Many attorneys believe that they can make this manual review more efficient by reviewing each document only once. Reviewers are thus instructed to kill two birds with one stone by performing the multifaceted factual review of documents (i.e., identifying key factual evidence) alongside the typical privilege review and the production review (i.e., determining whether a document is responsive to a discovery request).

This article challenges this common method of document review, arguing that it decreases review efficiency, and thus, increases cost of the review. *First*, studies show that production review may not be particularly accurate on its own. Research demonstrates that different reviewers may categorize the same document differently due to human error, fatigue, or subjective considerations. It is therefore inadvisable to distract reviewers examining a document for responsiveness by having them complete additional complex tasks such as reviewing for key factual evidence. *Second*, combining the factual, privilege, and production reviews subjects the reviewing attorneys to cognitive limitations. Applying principles from studies in cognitive psychology to the document review context makes clear that the common approach to document review is often flawed. These studies show that overloading an attorney's working memory by asking the attorney to

---

[*] Robert Keeling is a partner at Sidley Austin LLP, an experienced litigator whose practice includes a special focus on electronic discovery matters, and co-chair of the firm's eDiscovery Task Force. The author wishes to thank Matthew Letten and Christopher Joyce of Sidley Austin LLP for their assistance in writing this Article.

[**] This article has been prepared for informational purposes only and does not constitute legal advice. This information is not intended to create, and the receipt of it does not constitute, a lawyer-client relationship. Readers should not act upon this without seeking advice from professional advisers. The content herein does not reflect the views of Sidley Austin LLP.

simultaneously complete a series of complicated tasks negatively impacts the efficacy and quality of a document review. This may lead directly to an increase in the client's document review costs.

This article proposes a solution for attorneys looking to address the issues that come with inefficient review: splitting document review into distinct phases. By doing this, attorneys can alleviate some of the cognitive burdens unique to document review that attorneys face. This, in turn, reduces both the frequency of attorney error and the cost of the review. Splitting the review also carries certain benefits beyond those gained from lessening the burden on reviewing attorneys. Reviewing a document for responsiveness is a binary decision that can be made quickly. In making this decision, the attorney is exposed to, and gains familiarity with, the documents in a short time. This familiarity with the documents leaves the attorney better equipped to tackle the more challenging effort to identify key evidence.

TABLE OF CONTENTS:

I. INTRODUCTION 259
II. THE COST OF DOCUMENT REVIEW 261
III. CONSOLIDATED DOCUMENT REVIEWS: PROMOTING EFFICIENCY OR ERROR? 264
A. The Production Review Is Not a Wholly Accurate Process on Its Own 265
B. Limits on Cognition 268
C. Multitasking 271
D. Implications for Document Review. 272

IV. A SIMPLE SOLUTION: SPLITTING THE REVIEW 274
A. A Phased Approach to Document Review 275
B. Real World Evidence on the Impact of Multitasking on Review Performance 277
V. CONCLUSION 278

I. INTRODUCTION

Reviewing documents in response to a discovery request or a government subpoena can be difficult, time consuming, and expensive. One study indicates that document review is responsible for almost three-quarters of the entire cost of discovery.[1] These costs are driven by the volume of documents being searched and the resource-intensive nature of most reviews. Notwithstanding the use of modern techniques such as technology-assisted review, attorneys (typically junior associates or contract attorneys) predominantly undertake a manual, "eyes-on" review of a large number of documents potentially eligible for production. This review requires that the attorneys read the document and make various decisions about the document's relevancy, importance, and possible privileged contents. Attorneys review one document at a time, but are often tasked with reviewing hundreds of thousands,[2] which only serves to compound both the time and the monetary costs.

This document review process typically serves several key functions in litigation. First, between parties, it facilitates the exchange of documents relevant to the claims and defenses of that particular matter. A party served with discovery requests or subpoenas must determine what documents are substantively relevant, or "responsive," to the request. This process is often described as coding for relevancy or responsiveness, but I will refer to it as the "production review."[3]

Second, each party must evaluate which of the responsive documents may be safeguarded from disclosure or subject to other evidentiary protections. This determination typically involves assessing very different issues than the ones necessary for determining responsiveness, such as attorney-

---

1. NICHOLAS M. PACE & LAURA ZAKARAS, RAND INST. FOR CIVIL JUSTICE, WHERE THE MONEY GOES: UNDERSTANDING LITIGANT EXPENDITURES FOR PRODUCING ELECTRONIC DISCOVERY, xv, 42 (2012), http://www.rand.org/content/dam/rand/pubs/monographs/2012/RAND_MG1208.pdf.

2. *See In re* Target Corp. Customer Data Sec. Breach Litig., MDL No. 14-2522 (PAM/JJK), 2015 WL 13652715, at *7 (D. Minn. May 5, 2015) (attorneys reviewed 230,000 documents).

3. *See, e.g.*, Fed. Deposit Ins. Corp. for Valley Bank v. Crowe Horwath LLP, No. 1:17-CV-04384 (MMR), 2018 WL 3105987, at *2 (N.D. Ill. June 25, 2018) ("[T]he FDIC-R must conduct a relevance review and produce all relevant documents not already produced . . . ."); Gen. Elec. Co. v. United States, 119 F. Supp. 3d 17, 19 (D. Conn. 2015) ("The United States does not contest the right of GE to conduct a privilege review, and it is only logical for GE counsel to conduct a responsiveness review of the documents before conducting a privilege review."); Progressive Cas. Ins. Co. v. Delaney, No. 2:11-CV-00678 (LRH), 2014 WL 3563467, at *4 (D. Nev. July 18, 2014) ("The court-approved ESI protocol gave Progressive the option to perform a relevance review to make certain only relevant documents were produced.").

client privilege, confidentiality, or other such limitations on production. I will call this the "privilege review."

Third, document review helps parties understand the case and develop their overall legal strategies. Going into a review, a party will often identify key legal issues at stake and develop "issue codes" that can be used to organize relevant documents by topic and help senior attorneys prepare for settlement discussions, motions practice, and trial. Reviewing attorneys are also asked to identify "hot" or "key" documents that speak directly to the core issues in the case. This type of review is often called issue-coding, but I will refer to it as the "factual review."

In an attempt to make document review more efficient, attorneys often combine the production review, the privilege review, and the factual review into a single process or workflow. In other words, they ask reviewers to code for responsiveness, privilege, and key issues at the same time. The reasoning underlying this approach is intuitive: it would seem more efficient to have someone review a document only once. The reviewer can mark all applicable codes for that document, and the document will not need to be re-reviewed by other attorneys.[4]

The reality is not that simple. Combining the production review with the privilege and factual reviews does not always save time and money. The reason for this is straightforward: our brains can handle only a limited number of mental tasks at a time.

Combining reviews into a single workflow can thus actually slow down the pace of the overall review and inject a higher risk of error into the process. These added costs and errors can have a significant impact on the outcome of a case. For example, confusing different production requests or otherwise misremembering the requests can lead to mistakes in the production review. These mistakes could cut both ways, leading to the production of irrelevant documents that are not responsive to the requests and the erroneous withholding of relevant documents. Confusion when applying issue codes as part of the factual review may not directly impact the production but can still cause significant problems for the senior attorneys who rely on issue codes when crafting legal strategy, drafting motions, preparing or deposing witnesses, and conducting any other number of steps in preparation of trial. And mistakes in the privilege review can carry potentially case-altering consequences. The inadvertent production of documents containing attorney-client privileged communications or attorney work product might

4. This assumption that a document should never be reviewed more than once is flawed for many reasons. One major flaw, which is beyond the scope of this paper, is that a reviewer's knowledge grows as they read more documents, so decisions made at the beginning of a document review will be different, and potentially less accurate, than those made later.

reveal sensitive legal strategy and advice.[5] Even worse, in certain limited cases the production of privileged documents could trigger a waiver of the privilege, leading to the compelled disclosure of other privileged documents concerning the same subject matter.[6]

In this article, I take issue with the most common, conventional document review approach and suggest a simple solution to mitigate some of the problems associated with the conventional wisdom. Attorneys would do better to resist their intuition to combine multiple tasks into one review. Instead, for many document reviews, they should split the process into separate reviews: the production review first, then the privilege review, then the factual review, and then potentially others. Dividing the review in this fashion reduces the risk of cognitive strain on reviewers' brains, promotes faster decision-making, allows for more efficient and accurate work product, and decreases the number of hours spent overall and the likelihood of error. Additionally, the final review can focus on identifying and prioritizing the few documents that may be used as actual evidence, rather than the many documents that are merely responsive to discovery requests.

I begin with a discussion as to why document review must become a more efficient process, identifying some of the problems that appear inherent in document review.[7] I then survey and apply lessons from the study of cognitive psychology and multitasking to document review.[8] In doing so, I hope to demonstrate how conventional approaches to document review actually exacerbate the problems researchers have already identified. Finally, I close with a brief discussion as to why splitting the review can mitigate some of the problems of the most common approach and may even hold longer-term benefits.[9]

II. THE COST OF DOCUMENT REVIEW

The introduction of electronically stored information, or "ESI," has resulted in a massive increase of data that is now eligible for preservation,

5. *See, e.g.*, Starr Int'l Co. v. United States, 121 Fed. Cl. 428, 434–35 (2015), *aff'd in part, vacated in part by* 856 F.3d 953 (Fed. Cir. 2017) (showing that the government produced an email exchange with outside counsel that stated that the government's legal basis for taking over AIG was on "thin ice").

6. *See* Amgen Inc. v. Hoechst Marion Roussell, Inc., 190 F.R.D. 287, 288 (D. Mass. 2000) (describing the waiver of the attorney-client privilege as "the misstep feared by all litigators"); Paul W. Grimm et al., *Federal Rule of Evidence 502: Has It Lived Up to Its Potential?*, 17 RICH. J.L. & TECH. 8, 1 (2011) ("Nothing causes litigators greater anxiety than the possibility of doing, or failing to do, something . . . that waives attorney-client privilege or work-product protection.").

7. *See infra* Part II.

8. *See infra* Part III.

9. *See infra* Part IV.

collection, and review in the typical litigation.[10] Naturally, companies have witnessed a corresponding increase in the cost of producing such data, even though the number of documents that are actually used in litigation remains comparatively small.[11] Consider the production procedures for an average case that Microsoft detailed in a 2011 submission arguing for proposed changes to the current rules governing e-discovery.[12]

In its letter to the Chair of the Advisory Committee on Rules, Microsoft went step-by-step through each phase of the production process and explained how much data it was required to produce in response to a discovery request.[13] It then compared that number to the amount of produced material that was actually necessary to resolve the case. The ratio was shocking. Even after the collected data underwent a rigorous reduction process through "date ranges, search terms, de-duplication, and other data minimization processes," the relevant information amounted to "260 banker boxes of documents."[14] These 260 boxes were then manually reviewed by attorneys for privilege and responsiveness.[15] Upon review, only 22% (about 56 banker boxes) of the information was actually produced in the average matter.[16] Finally, of the 56 boxes produced (about 141,450 pages), an average of only 142 pages were used at trial.[17] Microsoft estimated that, "for each one-page trial exhibit, Microsoft produces an average of 1000 pages [and] manually reviews more than 4,500 pages."[18] Microsoft's own estimates are on par with the results of a survey undertaken by the Searle Center on Law, Regulation, and Economic Growth at Northwestern University, which found that "only 1 in 1000 pages produced in discovery is ever actually used as evidence to resolve the merits of the case."[19]

Although Microsoft did not release information concerning the cost of its production process, a study by the Rand Institute for Civil Justice provides an idea of how costs are distributed throughout the production process. The Rand Institute used a case-study method to collect data from eight different companies relating to production costs for fifty-seven cases in total.[20]

---

10. *See* Moore v. Publicis Groupe, 287 F.R.D. 182, 186 (S.D.N.Y. 2012) (over 3,000,000 emails collected for review).

11. PACE & ZAKARAS, *supra* note 1 at xiii.

12. *See* Letter from Microsoft Corporation to Judge David G. Campbell, Chair of the Advisory Committee on Civil Rules 5 (August 31, 2011) [hereinafter *Microsoft Letter*], https://www.uscourts.gov/sites/default/files/microsoft.pdf.

13. *See id.* at 3–5

14. *Id.* at 4.

15. *Id.*

16. *Id.*

17. *Id.* at 5.

18. *Microsoft Letter*, *supra* note 12, at 5.

19. *Id.* at 4 (emphasis omitted).

20. PACE & ZAKARAS, *supra* note 1, at xiii.

The study categorized the production process by task (collection, processing, and review), and then organized the costs associated with each task by source (internal, vendors, and outside counsel).[21] The study found that, "the major cost component in [its] cases was the review of documents for relevance, responsiveness, and privilege (typically about 73 percent)."[22] In one case, the cost of document production was $27 million.[23] The study further found that the cost "of outside counsel consumed about 70% of total e-discovery production costs," and attributed this large figure to the traditional farming out of document review to outside firms.[24]

Others have also weighed in on the increase of data available for review and the corresponding increase in the costs of complying with discovery requests. In *Moore v. Publicis Groupe*, Judge Andrew J. Peck of the United States District Court for the Southern District of New York noted that "linear manual review is simply too expensive" in light of the sheer volume of documents attorneys were expected to review.[25] In 2008, Federal Rule of Evidence 502 was amended in order to help "limit the prohibitive costs of privilege and work product review and retention."[26] And commenters, such as Judge Lee Rosenthal of the Southern District of Texas, have bluntly stated that "the sheer volume of electronically stored information makes privilege review expensive and time-consuming."[27]

---

21. *Id.* at xiv.

22. *Id.* The study suggests that for most of the cases it examined, reviewers reviewed the documents for responsiveness, and does not indicate that reviewers also undertook a factual review. *Id.* at 42. However, it is not difficult to infer that the costs found by the Rand Institute would be similar, if not higher, for cases where the production review and the factual review were combined into a single review.

23. *Id*. at 17.

24. *Id.* at xiv.

25. 287 F.R.D. 182, 190 (2012); *see also* Dynamo Holdings Ltd. P'ship v. Comm'r, 143 T.C. 183, 190 (2014) (approving use of predictive coding "to avoid the time and costs associated with traditional manual review"); Project on Predatory Lending of the Legal Servs. Ctr. of Harvard Law Sch. v. United States Dep't of Justice, 325 F. Supp. 3d 638, 654 (W.D. Pa. 2018) (searching for and producing documents from a potential pool of 20 million records is unduly burdensome); Nat'l Day Laborer Org. Network v. United States Immigration & Customs Enf't, No. 1:16-CV-00387 (PAE), 2017 WL 1494513, at *15 (S.D.N.Y. Apr. 19, 2017) (searching for and producing documents that could range from 436,000 to 1.3 million pages of records is unreasonably burdensome); Long v. Immigration & Customs Enf't, 149 F. Supp. 3d 39, 51 (D.D.C. 2015) (requiring search and redaction of five terabytes of information would be unduly burdensome).

26. FED. R. EVID. 502.

27. Lee H. Rosenthal, *A Few Thoughts on Electronic Discovery After December 1, 2006*, 116 YALE L.J. POCKET PART 167, 169 (2006); *see also* Sean Grammel, *Protecting Search Terms as Opinion Work Product: Applying the Work Product Doctrine to Electronic Discovery*, 161 U. PA. L. REV. 2063, 2065 (2018) ("Costs have skyrocketed due to the enormous volume of data being stored, a trend likely to continue.") Judge Rosenthal's comments ring even truer today, as it is now estimated that the "average" e-Discovery matter deals with

All of this makes clear that a significant portion of time and money are spent reviewing documents for relevancy and responsiveness, where most of the produced documents will never be used in defending or prosecuting the relevant claims. These conclusions, and the recognition that "eyes-on" manual review will never truly go away, suggest that attorneys need to become more efficient in determining which documents are responsive, both to combat the trend of overproduction and to reduce the costs sunk into the review process. Yet, the profession's current method for document review often does precisely the opposite.

III. CONSOLIDATED DOCUMENT REVIEWS: PROMOTING EFFICIENCY OR ERROR?

Document review tends to be fairly standardized. The reviewing attorney's first responsibility is to review the collected documents for responsiveness (the production review). In addition, the reviewing attorney may be charged with reviewing the documents for any safeguards from disclosure (the privilege review) and categorizing the documents by issue code to assist the more senior attorneys in building the case (the factual review). In theory, by applying a predefined issue code to each responsive document during the review, attorneys in the future will be able to quickly gather every document that is relevant to a particular issue or witness. To save time and reduce costs, assigning attorneys typically try to instruct junior attorneys or contract attorneys to perform the factual review alongside the initial production review and the privilege review. The intuition is that combining the three tasks into a single review makes the overall process more efficient.

But there are strong reasons to doubt this intuition. First, the law recognizes that discovery is an imperfect process that is prone to errors.[28] Sec-

---

approximately 100 gigabytes of data, or 6.5 million pages of Microsoft Word documents. *See* Kristin Kolasinski, *E-Discovery Fact Week Day Four: Examining E-Discovery Data Volumes*, EXTERRO (July 26, 2018), https://www.exterro.com/blog/e-discovery-fact-week-day-four-examining-e-discovery-data-volumes/.

28. There is a substantial body of case law addressing the issue of document review errors that lead to privileged documents being produced. *See* Victor Stanley, Inc. v. Creative Pipe, Inc., 250 F.R.D. 251, 257 (D. Md. 2008) (holding that party waived privilege after mistakenly disclosing 165 documents due to failure to use adequate keyword search terms); *see also* Am. Capital Homes, Inc. v. Greenwich, No. 2:09-CV-00622 (JCC), 2010 WL 11561400, at *3 (W.D. Wash. Aug. 3, 2010) (noting that "the only proper way to test the reliability of a keyword is to sample the documents so as to determine whether the search was over or under-inclusive") (internal citations omitted); United States v. Brewington, No. 1:15-CR-00073 (PAB), 2018 WL 1046804, at *3–4 (D. Col. Feb. 26, 2018) (noting that searching for the names of individuals in the email address field was reasonably calculated to prevent disclosure of privileged emails); Cole's Wexford Hotel, Inc. v. UPMC & Highmark, Inc., No. 2:10-CV-01609 (JFC), 2016 WL 462856, at *2 (W.D. Pa. Feb. 8, 2016) (noting that counsel took reasonable steps to protect the privilege, including searching for in-house and outside

ond, studies suggest that reviewers have enough trouble accurately making the binary decision between responsive and non-responsive, without adding other reviews into the mix.[29] Finally, there is a robust body of science in the study of memory and cognition demonstrating that humans can tolerate only a finite number of items on their mental plates.[30] These cognitive limitations can become overwhelmed in the face of too much information, causing even simple assignments to take much longer and to be more susceptible to error. Applying these findings to the document review process suggests that combining too many tasks within a single review may not make the review more efficient. Instead, it may prolong the review, and exacerbate the risk of error for any of the individual tasks to be performed.

A. The Production Review Is Not a Wholly Accurate Process on Its Own

Recent studies demonstrate that different reviewers will disagree over a document's responsiveness, although, as I will discuss, these studies may overstate the extent of that disagreement. For instance, in a study conducted by Herbert Roitblat, Anne Kershaw, and Patrick Oot, an initial team of reviewers found that 9.8% of the reviewed documents were responsive to a request.[31] Two subsequent teams of reviewers, A and B, found 24.2% and 28.76%, respectively, of those same documents to be responsive to the same request.[32] As to the documents that were judged to be responsive, team A was in agreement with the initial team only 28% of the time.[33] Similarly, the agreement rate between team B and the initial team was only 27.3%.[34] The agreement rate between teams A and B fared only marginally better at 43.8%.[35] A separate study, conducted by the Discovery of Electronically

---

counsel names in the full text of documents). The breadth of this case law affirms that document review is an imperfect process, that it is only getting worse with the constant increase of reviewable data, and that reviewers are prone to error when overwhelmed.

29. *See, e.g.*, Herbert L. Roitblat, Anne Kershaw & Patrick Oot, *Document Categorization in Legal Electronic Discovery: Computer Classification vs. Manual Review*, 61 J. AM. SOC'Y INFO. SCI. & TECH. 70, 75 (2010).

30. *See, e.g.*, George Miller, *The Magical Number Seven, Plus or Minus Two: Some Limits on Our Capacity for Processing Information*, 63 PSYCHOL. REV. 81, 82, 92–93 (1956).

31. Roitblat, Kershaw & Oot, *supra* note 29, at 75.

32. *Id.* at 77.

33. PACE & ZAKARAS, *supra* note 1, at 56. The authors of this study calculated these percentages, also called the proportion of specific agreement, by taking "twice the number of instances in which both teams agreed that a document was relevant, divided by the sum of (1) twice the number of agreed relevant documents, (2) the number of documents judged as relevant by team A but not team B, and (3) the number of documents judged as relevant by team B but not team A." *Id.* at 56 n.44.

34. *Id.* at 56.

35. *Id.*

Stored Information (DESI) workshop, produced similar results.[36] In this study, seven different teams of reviewers were given similar instructions and training prior to the document review assignment.[37] Despite the uniformity in treatment, the study found that "the seven teams differed significantly on the percentage of [documents] to be responsive, ranging from a low of 23.1 percent to a high of 54.2 percent."[38] Additionally, "[t]he overall agreement between various pairs of reviewing teams ranged from 65.5 percent to 84.9 percent."[39]

While these studies certainly support the common-sense intuition that manual review is not error free, their results very likely overstate the degree to which reviewers will reach different results as part of a well-managed document review.[40] Notably, the studies do not do enough to account for variability in reviewer reliability and the steps that can be taken to improve the quality of a manual review.[41] Some reviewers and review teams will perform better than others and this performance can be refined through proper review management. Over the course of a manual review, for example, more experienced reviewers might perform a second-level analysis of documents in order to provide real-time feedback.[42] This second-level review can be used to identify reoccurring mistakes or common areas of misunderstanding that can be addressed with the entire review team. It can also identify particular reviewers who should receive additional training or, in certain circumstances, be terminated from the review team. The studies discussed above included some initial training for reviewers, but did not implement the types of mid-review training and feedback that can improve reviewer reliability and performance.[43] Moreover, the studies did not attempt to establish a "gold standard" for relevance judgments that could be used to benchmark the review.[44] Given these shortcomings, attorneys should not be

36. Thomas I. Barnett & Svetlana Godjevac, *Faster, Better, Cheaper Legal Document Review, Pipe Dream or Reality?,* ICAIL 2011/DESI IV: WORKSHOP ON SETTING STANDARDS FOR SEARCHING ELECTRONICALLY STORED INFORMATION IN DISCOVERY PROCEEDINGS (2011), http://www.umiacs.umd.edu/~oard/desi4/proceedings.pdf.

37. PACE & ZAKARAS, *supra* note 1, at 57.

38. *Id.*; *see also* Barnett & Godjevac, *supra* note 36, at 5.

39. PACE & ZAKARAS, *supra* note 1, at 57 (footnote omitted)*; see also* Barnett & Godjevac, *supra* note 36, at 8.

40. *See* William Webber, *Re-examining the Effectiveness of Manual Review*, SPECIAL INT. GROUP INFO. RETRIEVAL (Ass'n Computing Machinery, Beijing, China), July 28, 2011; William C. Dimm, Predictive Coding: Theory & Practice 7–10 (Dec. 8, 2015) (unpublished manuscript) (on file with author).

41. Webber, *supra* note 40, at 4.

42. *Id.* at 4–6.

43. *See, e.g.*, Barnett & Godjevac, *supra* note 36, at 5.

44. Dimm, *supra* note 40, at 8.

so quick to dismiss the reliability of manual review, even in an age of technology-assisted review.

Setting those concerns aside, the question remains, what is driving the variability among determinations of responsiveness found in these studies? Roitblat, Kershaw, and Oot suggest two potential causes for the substantial disagreement among human reviewers. The first may be attributable to "random factors . . . unrelated to the material being judged or to any stable trait of the judges."[45] Reviewers' minds wander, and they become fatigued or distracted. Had the reviewers paid closer attention to each document, it is likely they would have categorized the documents differently.[46] The second source of disagreement may be more systematic, resulting from "the interaction between the content of the documents and stable properties of the reviewers, and to individual differences among reviewers."[47] In plain language, subjective considerations shape every review, and those considerations will influence the reviewers' calls. For example, reviewers' knowledge of the issues involved, background from prior reviews, their aggressiveness or cautiousness regarding the withholding of documents, and their assumptions regarding the risks involved all feed into a unique mental model of what constitutes a responsive document.[48] No two reviewers will approach the same document with the exact same background set of views and assumptions. Thus, the variability among calls does not demonstrate error, per se, but shows that responsiveness is less a binary determination and lies more along a spectrum.

Others have researched rates of disagreement in document reviews to investigate this very possibility. One study evaluated whether inconsistency among calls for responsiveness between a review team and a Topic Authority—a senior attorney who is familiar with the case and adjudicates any disagreements on responsiveness[49]—was a matter of reasonable differences of opinion or human error.[50] The authors found that most of the disputed documents were clearly responsive or clearly unresponsive as defined by the production request and the coding guidelines.[51] Therefore, the authors attributed a majority of the disagreements between reviewers to human error by one party or the other. [52]

---

45. Roitblat, Kershaw & Oot, *supra* note 29, at 77.
46. *Id.*
47. *Id.*
48. *Id.* (citation omitted).
49. Martha Grossman & Gordon V. Cormack, *Inconsistent Responsiveness Determination in Document Review: Difference of Opinion or Human Error?*, 32 PACE L. REV. 267, 272 (2012).
50. *Id*. at 276.
51. *Id.* at 284–85.
52. *Id.*

Given the errors that are seemingly unavoidable in the production review for responsiveness, it would be ill-advised to insert additional tasks into the process. Reviewing for factual importance and privilege in addition to reviewing for responsiveness will magnify the problems inherent in the production review because the factual review presents the reviewers with more options to consider for each document. Not only do reviewers have to decide whether a document is responsive, they must also determine if the document is privileged and which issues in the case are implicated. Thus, there are many more opportunities for distraction and subjective considerations to influence the reviewers' decisions. More importantly, there is an increased opportunity for human error. To reduce the rates of disagreement that must be evaluated as potential error, attorneys should consider minimizing the number of decisions a reviewing attorney must make during each review.

B. Limits on Cognition

Since the 1950s, researchers in cognitive psychology have understood that there are strict limits on the amount of information an individual can process at any one time.[53] As the amount of information exceeds those limits, the individual is more likely to make mistakes when transmitting or recalling the information.[54] The space in which we absorb information and then retain it for immediate accessibility is called *working memory*.[55]

Working memory is important for "mental tasks, such as language comprehension (for example, retaining ideas from early in a sentence to be combined with ideas later on), problem solving . . . and planning."[56] It "implicat[es] the simultaneous temporary storage and processing of any given information," and "is defined by the ability to retain information during short periods of time while performing a concurrent (and interfering) processing."[57] Working memory also helps "to maintain task relevant information in the face of competing irrelevant information."[58]

---

53. Miller, *supra* note 30, at 82, 92–93.

54. *Id.*

55. *See* Alessandra S. Souza, Laura Rerko & Klaus Oberauer, *Unloading and Reloading Working Memory: Attending to One Item Frees Capacity*, 40 J. EXPERIMENTAL PSYCHOL.: HUM. PERCEPTION & PERFORMANCE 3, 1237, 1237 (2014).

56. Nelson Cowan, *The Magical Mystery Four: How is Working Memory Capacity Limited, and Why?*, 19 CURRENT DIRECTIONS IN PSYCHOL. SCI. 51, 51 (2010) [hereinafter Cowan, *Magical Mystery*].

57. Roberto Colom et al., *Intelligence, Working Memory, and Multitasking Performance*, 38 INTELLIGENCE 543, 543–44 (2010).

58. Evan F. Risko et al., *Everyday Attention: Mind Wandering and Computer Use During Lectures*, 68 COMPUTERS & EDUC. 281 (2013).

There are clear advantages to leveraging working memory in the context of large-scale document review. Working memory enables reviewing attorneys to store key information about a particular document in a document set. That information can then be quickly and accurately recalled when the reviewer is deciding about another document in the set. Working memory is thus an invaluable cognitive tool for reviewing attorneys. A reviewer could, for example, use information about the relevance of a prior document to make a difficult call about the relevance of another document. Or a reviewer could make a connection between two documents and recognize the importance of the document to a particular issue in the litigation.

Working memory, however, is not unlimited. The limits identified early on by researchers demonstrate that our working memory "can hold only small amounts of information immediately accessible at the same time."[59] Many researchers suggest that average adults possess a working memory capacity capable of handling around three to five items.[60] Pure capacity limits, however, are difficult to observe because the brain has multiple mechanisms for retaining information that allow an individual to increase their recall.[61]

A different model of working memory suggests that working memory's limitations boil down to "resource-sharing."[62] Resource-sharing can be observed under dual-task conditions where subjects are asked to perform two tasks simultaneously.[63] Under these conditions, "[i]f two tasks require the same resource, then performance should deteriorate when those two tasks are performed simultaneously compared to single-task performance or to a situation in which two tasks relying primarily on different modules are performed simultaneously."[64] Research in this area has generally demonstrated that "deficits are typically observed when two tasks are performed simulta-

59. Klaus Oberauer & Elke B. Lange, *Interference in Verbal Working Memory: Distinguishing Similarity-based Confusion, Feature Overwriting, and Feature Migration*, 58 J. OF MEMORY & LANGUAGE 730, 730 (2008).

60. Cowan, *Magical Mystery*, *supra* note 56, at 52; Nelson Cowan, *The Magical Number 4 in Short-term Memory: A Reconsideration of Mental Storage Capacity*, 24 BEHAV. & BRAIN SCI. 87, 88 (2000) [hereinafter Cowan, *Magical Number*]; WILLIAM C. DIMM, PREDICTIVE CODING: THEORY & PRACTICE 7–10 (DEC. 8, 2015) (unpublished manuscript) (on file with author).

61. Cowan, *Magical Number*, *supra* note 60, at 88; Cowan, *Magical Mystery*, *supra* note 56, at 51.

62. Candice C. Morey et al., *Asymmetric Cross-Domain Interference Between Two Working Memory Tasks: Implications for Models of Working Memory*, 69 J. MEMORY & LANGUAGE 324, 325 (2013).

63. *Id.*

64. *Id.*

neously, regardless of the domain of the memoranda involved, but deficits observed with two tasks from the same domain are usually larger."[65]

A third model of working memory focuses on the decay of information held in working memory through interference. For instance, there can be interference "between items to be held in working memory simultaneously, and interference between memory items and representations involved in a concurrent processing task."[66] Examples of interference include: (1) item confusion, where an individual confuses one item with other items available in working memory, to the extent that the items are similar in nature or are otherwise triggered by similar cues;[67] (2) feature migration, where items held in working memory are composed of features and "interference arises from the interaction between the features," causing them to "migrate from one item to another, thereby creating illusory conjunctions;"[68] and (3) feature overwriting, which posits that "the more items are held in working memory simultaneously, and the more they overlap, the more each item's representation is degraded, thereby posing a limit to the capacity of working memory."[69]

Although each of these models provides a different explanation for why working memory is limited, researchers generally agree that working memory is limited or constrained in some manner.[70] My purpose in discussing these different theories is not to take a position on the merits of one theory or another. Rather, the point is to understand that an individual can handle only so much information on his or her mental plate, and that these limitations have very real implications for document review. Moreover, limitations on working memory also have significant consequences for an individual's ability to multitask, something reviewers (and attorneys, more generally) are expected to do every day. When assigning attorneys ask reviewers to undertake a review for responsiveness, privilege, and factual importance during a single pass at the documents, they are essentially asking reviewers to complete multiple tasks at the same time. Given the strict limitations on working memory and the close relationship between working memory and multitasking, it will come as no surprise that multitasking may also negatively affect the efficacy of a document review.

---

65. *Id.*
66. Oberauer & Lange, *supra* note 59, at 730–31.
67. *Id.* at 731.
68. *Id*.
69. *Id.*
70. Cowan, *Magical Number*, *supra* note 60, at 88 (describing commonly held views on working memory). Redick et al. offer a similar definition of multitasking: "(a) performing multiple tasks, (b) consciously shifting from one task to another, and (c) performing the component tasks over a relatively short time span." Thomas S. Redick et al., *Cognitive Predictors of a Common Multitasking Ability: Contributions from Working Memory, Attention Control, and Fluid Intelligence*, 145 J. EXPERIMENTAL PSYCHOL. GEN. 1473, 1474 (2016).

C. Multitasking

Although multitasking is increasingly common in the classroom and the workplace, research has shown that individuals do not actually perform multiple tasks simultaneously; rather, individuals rapidly switch their attention back and forth among tasks, giving off the impression that they are doing multiple things at once.[71] Indeed, a common definition for multitasking is "the cognitive ability to perform multiple task goals in the same time period by engaging in frequent switches between individual tasks."[72] Working memory capacity has been shown to be closely connected with the ability to multitask.[73] In fact, working memory is "more highly correlated with multitasking than intelligence."[74] One study even found that "intelligence does not predict multitasking once its correlation with [working memory capacity] is controlled for."[75] This may be the case "because tasks tapping [working memory capacity] and multitasking situations share coping with two or more competing concurrent cognitive requirements."[76]

The consequence of this rapid switching is that even simple tasks, when combined, have built-in costs. For example, consider a situation in which individuals are given two different tasks and a cue determines which task the individual is to perform. When the cue is given, the individual must gather all the information they have about the triggered task before they are able to respond.[77] This information gathering takes time, more time than if individuals undertook each task one at a time.[78]

One often-cited article suggests that individuals can lose up to 40% of their productivity when they attempt to multitask.[79] Brain imaging demonstrates that "when a single area of the brain . . . has to do two things at once . . . there is less brain activation than occurs" when the area of the brain

---

71. Jon Hamilton, *Think You're Multitasking? Think Again*, NPR (Oct. 2, 2008), https://www.npr.org/templates/story/story.php?storyId=95256794; *see also* Talk of the Nation, *The Myth of Multitasking*, NPR (May 10, 2013), https://www.npr.org/2013/05/10/182861382/the-myth-of-multitasking.

72. Colom et al., *supra* note 57, at 543 (internal quotation marks omitted).

73. Redick et al., *supra* note 70, at 1476 ("WM accounts for significant multitasking variance").

74. Colom et al., *supra* note 57, at 544*; see also* Redick et al., *supra* note 70, at 1475–76.

75. Colom et al., *supra* note 57, at 548.

76. *Id.* at 544.

77. Hamilton, *supra* note 71.

78. Susan Weinschenk, *The True Cost of Multi-Tasking*, PSYCHOL. TODAY (Sept. 18, 2012), http://www.psychologytoday.com/blog/brain-wise/201209/the-true-cost-multi-tasking.

79. *Id.*; *see also Multitasking: Switching Costs*, AM. PSYCHOL. ASS'N (Mar. 20, 2006), http://www.apa.org/research/action/multitask.aspx.

has to do only one thing,[80] suggesting an overall decrease in productivity when multitasking. Furthermore, when individuals multitask, they increase the risk of error for both tasks than if they completed the tasks individually.[81] These costs are compounded if the tasks are complex,[82] which is often the case with tasks involving the review of dense or unfamiliar documents.

There are various theories describing why multitasking inhibits performance, many of which mirror those on working memory. But there appears to be no agreement as to which theory is most promising.[83] Undertaking multiple tasks at one time is generally thought to cause interference among the various tasks, thus reducing performance.[84] Other theories are built upon a resource-sharing perspective, which "proposes that cognitive resources are limited and information processing is dependent on their availability."[85] Thus, when an individual multitasks, the demand for cognitive resources overwhelms the actual supply, jeopardizing information processing and task performance.[86] Again, the takeaway for my purposes is simply that multitasking, despite popular belief or the expectations of employers and clients, is not a habit that enhances productivity.[87]

D. Implications for Document Review

The research discussed above has practical implications for everyday legal tasks, including document review. Typically, when reviewers receive a document review assignment, they are given a set of instructions for conducting the review. These instructions may include exemplar or training documents. Exemplar documents provide examples of what is responsive to a particular request or what might be relevant to various issue codes. They are meant to be used by reviewers as a guide against which to compare the

80. Sandra Blakeslee, *Car Calls May Leave Brain Short Handed*, N.Y. TIMES (July 31, 2001), https://www.nytimes.com/2001/07/31/science/car-calls-may-leave-brain-short-handed.html; *see also* LAUREL J. FELT & MICHAEL B. ROBB, TECHNOLOGY ADDICTION: CONCERN, CONTROVERSY, AND FINDING BALANCE, COMMON SENSE MEDIA 1 (2016), https://www.commonsensemedia.org/sites/default/files/uploads/research/csm_2016_technology_addiction_research_brief_1.pdf

81. *See* Blakeslee, *supra* note 80.

82. *Id.*

83. Colom et al., *supra* note 57, at 543.

84. *Id.*; *see also* Hamilton, *supra* note 71.

85. Jatin Srivastava, *Media Multitasking Performance: Role of Message Relevance and Formatting Cues in Online Environments*, 29 COMPUTERS IN HUM. BEHAV. 888, 890 (2013); *see also* Victor Mittelstädt & Jeff Miller, *Separating Limits on Preparation Versus Online Processing in Multitasking Paradigms: Evidence for Resource Models*, 43 J. EXPERIMENTAL PSYCHOL.: HUM. PERCEPTION & PERFORMANCE 89, 89 (2017).

86. Srivastava, *supra* note 85, at 890.

87. Mittelstädt & Miller, *supra* note 85, at 89 ("It is well known, however, that performance often suffers when people perform multiple tasks at the same time.").

rest of the documents they review. As part of the instructions, the reviewers are typically made aware of the specific production requests and the applicable issue codes. The number of issue codes will depend on the matter and the preferences of the attorney who designed the review, but the number of issue codes may range from one to as many as fifteen or twenty. To successfully complete their assignment, reviewers must remain mindful of the specific production requests and issue codes, while maintaining the examples provided by the exemplar documents in an easily accessible state for comparison. Making matters worse, issue codes and production requests are frequently similar or overlapping in nature. When this happens, there is a high potential for interference among the various items that increases the likelihood of mistakes in all aspects of the review.[88]

As reviewers struggle to maintain these various items within a reachable state of their working memory, they are also expected to read and process documents, the content of which may be dense in substance, lengthy, or wholly unfamiliar, and then compare that information to the items they are attempting to recall. This reduces the total amount of brain-power available for each task, and thereby diminishes productivity.[89] One researcher believed that reading a single essay could stretch an individual's working memory capacity past its breaking point.[90] Yet, what assigning attorneys are asking reviewers to do—analyze documents typically at a rate of thirty to fifty documents per hour, but sometimes at a rate of one hundred documents per hour[91]—is far more taxing on reviewers' cognitive capabilities. As a result, reviewers may need to re-read documents or refer back to the training materials to make a decision, causing significant delay.

Ultimately, reviewers who find themselves assigned to conducting a factual review alongside the initial production review and the privilege review also find themselves in the unfortunate position of multitasking. While reading a single document, the reviewer is required to make several different and, at times, difficult determinations at once. As a result, the reviewer will rapidly shift from thinking about responsiveness to thinking about factual importance or privilege. And as the reviewer's attention shifts from one task

88. *See* Santiago Pelegrina et al., *Similarity-Based Interference in a Working Memory Numerical Updating Task Age-Related Differences Between Younger and Older Adults*, 59 EXPERIMENTAL PSYCHOL. 183, 183 (2012) ("When the information in a list is similar . . . items are prone to be confused at the moment of recall . . . .").

89. *See* AM. PSYCHOL. ASS'N, *Multitasking: Switching Costs*, *supra* note 79; *see also Multitasking Limits and Predictors*, AM. PSYCHOL. ASS'N, (Jan. 26, 2017), https://www.apa.org/pubs/highlights/peeps/issue-86.aspx ("[M]ultitasking performance is constrained by available resources.").

90. Cowan, *Magical Mystery*, *supra* note 56, at 52.

91. PACE & ZAKARAS, *supra* note 1, at xvi, 50.

to another, time is lost and the risk of error is increased.[92] Error arising from human fault is common enough during an initial production review on its own,[93] but will be made worse by adding the factual and privilege reviews as additional, simultaneous tasks, especially if the factual issues are complex or unfamiliar. Further, those who multitask may "have more trouble tuning out distractions than people who focus on one task at a time," which invites the possibility of further time-lag and error.[94]

Demanding too many tasks of a reviewer during a single review can also often defeat a critical goal of manual review. Manual review remains an essential part of document review even after the use of technology-assisted review, in part because attorneys must become familiar with the documents in order to build a case. Junior associates charged with document review are often also charged with being the keepers of the facts and are relied upon by senior attorneys to call attention to those facts where necessary. To fulfill this duty, associates must actually remember what they read during the review and be able to draw connections between the documents. But to the extent that document review becomes a multitasking venture that surpasses the reviewer's cognitive limitations, the reviewer's ability to process and remember the information they read long-term will be limited. While the brain can obscure pure capacity limits through alternative mechanisms for storing and recalling information, some tasks, because of their length or complexity, may prevent these mechanisms from doing their job.[95] Therefore, if the documents and issues are complex, it is possible that reviewers will be unable to process information into their long-term memory, thereby obstructing a key purpose of manual review. To obtain a level of familiarity with the documents that assigning attorneys often expect, reviewers will need to re-read the documents, thereby increasing the time spent on the review.

IV. A SIMPLE SOLUTION: SPLITTING THE REVIEW

Despite advances in document-review technology, most matters still require some type of manual review. While predictive coding and other forms of technology-assisted review have been effective at reducing the costs associated with discovery,[96] these technological tools are not used in many—if not most—reviews. Even when used, technology-assisted review

92. *See* AM. PSYCHOL. ASS'N, *Multitasking: Switching Costs*, *supra* note 79; s*ee also* Erez Reuveni, *Copyright, Neuroscience, and Creativity*, 64 ALA. L. REV. 735, 766 (2013).

93. *Supra* Part III.A.

94. Kendra Cherry, *Multitasking: The Cognitive Costs of Multitasking*, VERYWELLMIND (Mar. 16, 2019), https://www.verywellmind.com/multitasking-2795003.

95. Cowan, *Magical Mystery*, *supra* note 56, at 51–52.

96. PACE & ZAKARAS, *supra* note 1, at 60.

has not eliminated the need for manual "eyes-on" review so much as it has reduced the number of documents subjected to it. For example, with predictive coding, human reviewers code a subset of documents that a machine learning algorithm uses to create a predictive model that can analyze other documents and classify them as relevant or non-relevant.[97] Depending on the protocol being used, human reviewers might continue to review documents that are fed into the algorithm to improve the predictive model.[98] In other words, the manual review of documents is here to stay and attorneys should not overlook ways in which the traditional approach to manual review might be improved. In this section, I propose an approach to document review that incorporates research on cognitive limitations by limiting the number of cognitive tasks a reviewer is asked to perform. I also offer some real-world evidence on the impact of cognitive limitations on review performance from my own practice that suggests my proposed approach will produce better review results.

A. A Phased Approach to Document Review

The traditional approach to manual review—assignment of the production review, factual review, and privilege review in a single pass—attempts to achieve efficiency by streamlining the review. It combines different reviews so that reviewers can complete three tasks while reviewing each document just once. But as I have discussed, structuring the review this way severely strains, and can overcome, the reviewer's mental capacity for storing and processing information. When the reviewers become overwhelmed with information, the review itself not only takes much longer, and therefore costs more money, but is also more error-prone, reducing the overall quality of the work despite the higher price. Splitting the review into distinct phases, first the production review, then the privilege review, and finally, the factual review, reduces those risks and carries certain benefits. At a minimum, it would be advisable to separate the factual review from the production and privilege reviews.[99]

---

97. *See* Maura R. Grossman & Gordon V. Cormack, *The Grossman-Cormack Glossary of Technology-Assisted Review*, 7 FED. CTS. L. REV. 1, 26 (2013).

98. *See* Maura R. Grossman & Gordon V. Cormack, *Technology-Assisted Review in E-Discovery Can Be More Effective And More Efficient Than Exhaustive Manual Review*, 17 RICH. J. L. & TECH. 11, at 39–43 (2011) (describing different protocols for predictive coding).

99. At the same time, assigning attorneys would do well to be thoughtful about the number and types of issue codes they create. Splitting the review, but then instructing the reviewer to assign for 20 different issue codes during the later factual review may generate some of the same problems that plague reviewers when the three processes are combined in the first instance.

Splitting the review in this fashion can take advantage of the fact that reviewers generally find the production and privilege reviews easier and faster than the factual review. Reviewers prefer the binary, yes or no, questions at stake in the production and privilege reviews, as opposed to questions that pose multiple and alternative possibilities, like those presented during the factual review.[100] Because humans "cannot keep in mind and switch back and forth among three or more alternatives,"[101] the factual review is typically a more cognitively complex task, especially when reviewers are seeing the material for the first time. The production and privilege reviews are thus less likely to inhibit the brain's alternative mechanisms for storing and processing information long-term, increasing a reviewer's opportunity to bypass pure capacity limits.[102] Similarly, reducing the number of moving parts a reviewer must keep track of reduces the chance for interference among those moving parts, as well as the demand on limited cognitive resources. In these circumstances, reviewers are more likely to process and retain information long-term when calls for responsiveness or privilege are made on their own, rather than alongside calls for factual relevancy or importance.

Splitting the review leaves a reviewer better equipped to tackle the more challenging factual review later, especially where the documents implicate multiple issues, the issues are complex, or the issues appear very similar in content, requiring a nuanced eye to discern among them. Anyone who has reviewed documents before will attest that as time progresses and you become more familiar with the documents and issues in the case, the quality of your decision-making often improves. This is especially true for reviewers who at the beginning of the review are unfamiliar with the substantive content of the documents. An email reviewed early on may not make sense, and therefore may not appear relevant to any particular issue, until considered in connection with an email reviewed later in the process.

Splitting the review also allows reviewers to truly focus on the actual task before them. For example, the goal of the production review should not be confused with, and, indeed, is entirely separate from, the goal of the factual review. The goal of a production review is simply to ensure that available documents that are relevant to a production request (and not privileged) are produced. During a production review the reviewer need only match documents to requests. The factual review, on the other hand, is designed to locate the documents that are critical for trial. During the factual review,

100. Katherine Harmon, *Motivated Multitasking: How the Brain Keeps Tabs on Two Tasks at Once*, SCI. AM. (April 15, 2010), http://www.scientificamerican.com/article.cfm?id=multitasking-two-tasks.

101. *Id.* (internal quotation omitted).

102. Redick et al., *supra* note 70, at 1475 (complexity impacted the relationship between cognitive abilities, like working memory, and multitasking).

reviewers are more concerned with whether the document tends to help or harm the client's position in the litigation. By performing the processes simultaneously, reviewers are likely to conflate these distinct goals and code technically relevant documents as factually important. This creates additional time-intensive work on the back-end, when attorneys must actively search for the important documents. For these same reasons, issue-coding nominally responsive documents serves to clutter key legal issues, potentially obscuring factually important documents.

B. Real World Evidence on the Impact of Multitasking on Review Performance

One potential objection to splitting a document review into distinct phases—the production review, the privilege review, and the factual review—is that it will be more efficient to have a reviewer look at a document just once. As I have argued in this Article, that intuition about the efficiency of a combined document review is inconsistent with what we know about cognitive limitations and multitasking. The more tasks we ask a reviewer to perform as part of a single pass review, the greater the risk of error.

Real world evidence supports this view. Consider the following data from two large-scale document reviews. Both reviews involved a substantial volume of documents and the client ended up producing millions of pages of materials in response to these requests. As part of the production review in both cases, first-level reviewers were asked to confirm whether the document was responsive to various discovery requests.

The two reviews differed in the number of tags the reviewers were asked to apply. In Review A, the reviewers were asked to choose between twelve different tags to a document: (1) four tags for responsiveness (Responsive, Non Responsive, Foreign Language, and Technical Issue); (2) one tag for privilege; and (3) seven different issue tags. In Review B, the reviewers had eight different tags to potentially apply: (1) the same four tags for responsiveness (Responsive, Non Responsive, Foreign Language, and Technical Issue); (2) one tag for privilege; and (3) only three different issue tags.

We can use data on the number of documents where the responsiveness designation was changed as part of a second- or third-level review as an indication of review accuracy for both reviews. The more documents where the first-level review coding was "overturned" means that, all else being equal, the first-level reviewers were making more mistakes. The "overturn rate" was 11.6% in Review A and 8.4% in Review B, respectively. The below table sets out the percentage of documents where the responsiveness designation was changed:

Table 1. *Overturn is an indicator of mistakes in the first review.*

Review	Number of Tags	Overturn Rate
Review A	12	11.6%
Review B	8	8.4%

The data from these two reviews suggests that the reviewers in Review B (who had fewer tags to consider), were more accurate than the reviewers in Review A (who had more tags to potentially apply). This correlation between the number of tags and the overturn rate further indicates that the reviewers who were asked to perform more tasks and consider whether a document satisfied additional issue tags did a worse job at deciding whether a document was responsive. In other words, the more reviewers were asked to multitask, the more documents they coded incorrectly. This is precisely what you would expect to see in light of the research on cognitive limitations and multitasking that I have discussed in this Article.

That said, I do not mean to imply that there is enough evidence to suggest a causal relationship in this data between the number of document tags and review performance. There could be any number of additional reasons that the overturn rate in Review B was lower, including the guidance and training provided to the reviewers, the experiences of the reviewers, the subject matter of the review, and the method used for selecting documents for second- and third-level review. But at a minimum, this evidence is consistent with insights from psychology about cognitive limitations and multitasking.

V. CONCLUSION

Given the costs associated with manual review, it is no wonder that assigning attorneys are anxious to reduce the time and expense involved with manual review and make the process more efficient overall. But combining reviews for production with reviews for factual importance and privilege is a step away from, not toward, that end. Rather than shortening the amount of time spent on document review, consolidating the different reviews into a single process places a significant amount of strain on reviewers' mental faculties. This often causes the review to take longer and be more susceptible to error. Separating the review into distinct components is less likely to bombard a reviewer with too much information at once, making it easier for reviewers to process the information, which will ultimately lead to more decisive and accurate decisions and a level of familiarity with the information that cannot be achieved in a multi-task situation.