

Neal, T.M.S. & Grisso, T. (2014). The cognitive underpinnings of bias in forensic mental health evaluations. *Psychology, Public Policy, and Law*, 20, 200-211. doi:10.1037/a0035824.

© 2014 American Psychological Association

This is a pre-print version of the paper that was published in the APA journal. This version may not exactly replicate the final version published in the APA journal. It is not the copy of record.

(See erratum at end of document).

The Cognitive Underpinnings of Bias in Forensic Mental Health Evaluations

Tess M.S. Neal & Thomas Grisso

University of Massachusetts Medical School

Author Note

Tess M.S. Neal and Thomas Grisso, Department of Psychiatry, University of Massachusetts Medical School.

Tess M.S. Neal is now at the University of Nebraska Public Policy Center.

A special thank you to David Faust and Ira Packer for providing feedback on an earlier version of this manuscript.

Correspondence concerning this article should be addressed to Tess M.S. Neal, University of Nebraska Public Policy Center, 215 Centennial Mall South, Suite 401 (P.O. Box 880228), Lincoln NE 68588. E-mail: tess.neal@asu.edu

Abstract

We integrate multiple domains of psychological science to identify, better understand, and manage the effects of subtle but powerful biases in forensic mental health assessment. This topic is ripe for discussion, as research evidence that challenges our objectivity and credibility garners increased attention both within and outside of psychology. We begin by defining bias and provide rich examples from the judgment and decision making literature as they might apply to forensic assessment tasks. The cognitive biases we review can help us explain common problems in interpretation and judgment that confront forensic examiners. This leads us to ask (and attempt to answer) how we might use what we know about bias in forensic clinicians' judgment to reduce its negative effects.

Keywords: bias; judgment; decision; forensic

The Cognitive Underpinnings of Bias in Forensic Mental Health Evaluations

Mr. Jones, a 24-year-old man facing a felony charge of cocaine trafficking, had been convicted of four previous offenses (assault & battery, theft, trespassing, and giving a false name to a police officer). He had never before received psychiatric treatment, but his attorney requested an evaluation of his client's mental status at the time of his alleged offense. Converging evidence indicated (among other things) that Mr. Jones was influenced by his antisocial peers, his substance abuse was impacting his relationships at the time of the crime, and he had a history of several head injuries resulting in loss of consciousness. After hearing the case, the court found Mr. Jones Not Guilty by Reason of Insanity (NGI).

Now please rank the following six categories of mental illness in order of the likelihood that, at the time of the offense, Mr. Jones met diagnostic criteria for each. Use 1 for most and 6 for least likely.

- | | |
|---|---|
| <input type="checkbox"/> Affective Disorder | <input type="checkbox"/> Personality Disorder |
| <input type="checkbox"/> Mental Retardation / Intellectual Disability | <input type="checkbox"/> Substance Use Disorder |
| <input type="checkbox"/> Psychotic Disorder | <input type="checkbox"/> Dissociative Disorder |

The question in this vignette is straightforward for readers who know the relative likelihood of various mental disorders in defendants found NGI. Defendants with psychotic disorders are the most likely to be found NGI, and defendants with personality disorders are among the least likely to be found NGI (Cochrane, Grisso, & Frederick, 2001; Warren, Murrie, Chauhan, Dietz, & Morris, 2004). Given the fact that Mr. Jones was found NGI, the “base rates” of the various disorders in the NGI population should have weighed heavily in the decision task. However, we provided stereotypic information about Mr. Jones that did not fit with the NGI research data. In fact, we intentionally designed Mr. Jones as an “anti base-rate character” (see Kahneman, 2011) to illustrate one kind of cognitive bias, the representativeness heuristic, which we discuss below in more depth along with other kinds of biases.

The purpose of this review is to apply information from multiple domains of psychological science (e.g., cognitive, social, methodological, clinical) to identify and better

understand bias in forensic mental health assessment. This topic is ripe for discussion as several studies have investigated potential bias in the work of forensic experts. For example, Murrie, Boccaccini, and colleagues published compelling data documenting the “allegiance effect” in forensic assessments (see e.g., Murrie, Boccaccini, Guarnera, & Rufino, 2013). Their data suggest that adversarially-retained experts tend to interpret data and score certain psychological assessment instruments in ways that are more likely to support the retaining party’s position.

We begin by defining bias. Then we review evidence for bias in forensic mental health practice in the context of rich research and theory on judgment and decision making. Along the way, we offer examples of how various theories of bias can help us explain common problems in interpretation and judgment that confront forensic examiners. This leads us to ask how we can use what we know about bias in clinicians’ judgment to find ways to reduce it. We describe various approaches to the problem and offer ideas that may stimulate research on interventions to mitigate the negative effects of bias in forensic evaluators’ decision making processes.

Defining “Bias”

According to the Oxford English Dictionary (2012), the word “bias” was first documented in the mid-16th century. It has roots in the French *biais*, which is perhaps based on the Greek *epikarsios*, for “oblique.” Bias was originally used to describe both a slanting line (i.e., the diagonal in a square) and a curve, such as the shape given to the side of a bowl or the curve of a cheek, as used by Shakespeare, “Thy sphered Bias cheek” (1609, *Troilus & Cressida* IV, vi.8). It also was used to refer to the oblique motion of a loaded bowling ball (as well as to the asymmetric construction of the bowling ball by loading one side with lead), exemplified by Shakespeare’s passage: “Well, forward, forward thus the bowle should run. And not unluckily against the bias” (1596, *The Taming of the Shrew* IV, v. 25, as cited by Keren & Tiegen, 2004).

The word was also used in the fabric industry to refer to cutting diagonally across the grain, “cut along the bias,” and in cooking as well; for instance, slicing carrot at a sharp angle increases the surface area of each slice and is thought to be visually appealing for food presentation.

These two uses of the term “bias” capture different meanings. It can be used to describe deviations from the norm (as with motion of the loaded bowling ball) or slanting one way rather than another (like the diagonal line). Error in judgment is not necessarily indicated, although the term as used today often carries a negative connotation. Keren and Tieggen (2004) point out the distinction between bias being a cause vs. an effect, noting that the word is also used in both of these ways. For example, the bias of the bowling ball can be *in its shape or loading*, causing it to curve (i.e., the cause), or it can refer to *the trajectory of the ball* (i.e., the effect).

In the forensic mental health field, the word “bias” carries a negative connotation often associated with an inappropriate personal or emotional involvement on the part of the evaluator (Neal, 2011). Bias may be outside the examiner’s awareness (i.e., implicit), but examiners may also be accused of purposefully putting a “spin” on the evaluation (i.e., explicit bias). The insightful and purposeful “spin” may not be the biggest challenge facing forensic mental health professionals. Evaluators who engage in explicit bias are likely to be recognized by their colleagues – in both the mental health and legal fields – as “hired guns” with reduced credibility as trustworthy experts. Rather, the bigger challenge for the field (and for individual forensic practitioners) is likely in understanding and dealing with implicit bias in the way we process and interpret information and reach conclusions. Although we acknowledge that explicit biases deserve attention, this review focuses primarily on the way in which examiners’ thinking and decision making may be systematically affected by implicit biases.

West and Kenny (2011) drew from multiple domains of psychology to create a single, integrative framework for the study of bias and accuracy. It is called the “Truth and Bias (T&B) Model” of judgment. Their T&B Model provides theoretical definitions and parameters of interest in the study of accuracy and bias. It can be used to streamline science’s basic understanding of how these constructs operate independent of the researcher’s *a priori* field or theoretical reference point. The model they developed can be applied widely across psychological domains, including forensic psychology. West and Kenny’s (2011) definition of bias, which we adopt in the present review, is any *systematic* factor (i.e., not random error) that determines judgment other than the truth.

Given this definition of bias, what evidence do we have that it influences forensic practitioners in their work? And how might we understand these influences on forensic practitioners? Fortunately, there is a rich body of judgment and decision making research that may provide the theoretical frameworks we need for explaining various cognitive biases that underlie human cognition. These theoretical frameworks may help us bridge the gap to designing studies that could reduce bias in forensic decision making.

Biases that May Affect Forensic Experts

Forensic assessment tasks present a tall order. Otto (2013) vividly outlined the difficulties faced by forensic clinicians (emphasis in original):

To (*in a limited amount of time, using assessment techniques of limited validity, and with a limited amount of information-some of which is provided by persons with an investment in the examiner forming a particular opinion*) come to an accurate assessment about the past, current, and/or future emotional, behavioral, and/or cognitive functioning of an

examinee as it relates to some issue before the legal decision maker (*while ensuring that how one has been involved in the case does not affect one's decisions*).

Forensic evaluators are asked to gather comprehensive data with regard to the referral issue, to analyze the patterns and interrelationships among the various pieces of data (called “configural analysis”), and then interpret the data to reach an opinion that will assist the trier-of-fact (see e.g., Faust & Faust, 2012). However, human brains do not have an endless capacity for processing information. Simon (1956) called this constraint “bounded rationality:” we do the best we can within the design of our cognitive machinery. As a consequence, people often use cognitive shortcuts or simplifying strategies to manage cognitive load.

There are two traditions of research with regard to human cognitive capacities (Kahneman & Klein, 2009). The Heuristics and Biases (HB) tradition, which developed first, has focused on the limitations of and systematic errors in human cognition (see e.g., Tversky & Kahneman, 1974; Kahneman, 2011; Kahneman, Slovic, & Tversky, 1982). The Naturalistic Decision Making (NDM) tradition developed in part in reaction to HB's narrow focus on problems in human cognition. NDM has focused on the strengths and evolutionary adaptiveness of human cognitive capacities (see e.g., Gigerenzer & Goldstein, 1996; Lipshitz, Klein, Orasanu, & Salas, 2001). For instance, NDM researchers argue our brains have adapted to process information in a fast and frugal way – quickly making sense of the vast amount of information with which we are constantly faced (Gigerenzer & Goldstein, 1996).

Kahneman and Tversky, the founders of the HB tradition, showed humans' decision-making processes are much more prone to error than was previously imagined. “Bias,” as used in the HB tradition, is a by-product of mental shortcuts called heuristics. Heuristics are decision aids that proceed along “rules of thumb,” used by people to arrive at efficient answers especially

when solutions are not readily apparent. It shares the same root as the word *eureka* (Kahneman, 2011). The mechanisms underlying heuristics are incomplete, but they are adequate for most situations and usually assist people in arriving at valid answers while preserving mental resources (Keren & Tieggen, 2004). Heuristics provide an adaptive pathway for humans to cope with limited processing capacities (i.e., an evolutionary strength according to NDM researchers), but the incomplete nature of heuristic methods means they can lead to error-proneness due to systematic biases under some circumstances (i.e., a limitation according to HB researchers).

Heuristic “answers” can be thought of as approximations of the truth. Theoretically, the truth could be discovered after an exhaustive step-by-step method of testing various possible solutions and arriving at the correct answer. Einstein (1905), for example, called his first Nobel Prize-winning paper on quantum physics, “On a Heuristic Point of View toward the Emission and Transformation of Light.” He used the term “heuristic” rather than “theory” to indicate that his new idea was an initial approximation should be further explored (Keren & Tieggen, 2004).

There is a rich literature on various heuristics and biases that affect human thinking processes. Since Kahneman and Tversky’s work in the 1970s, the number of “new” heuristics and biases has proliferated, although it is not always clear that the new discoveries are distinct from earlier-identified ones. As such, we have organized this section by focusing on three major heuristics that were among those first discovered and discussed: representativeness, availability, and anchoring (Keren & Tieggen, 2004). We provide examples of ways they might influence the judgment and decision making of forensic mental health examiners (see Table 1).¹ Most of the research we review is from the judgment and decision making literature, and it may or may not

¹ Heuristics are constructs used to describe cognitive processes by which we quickly summarize and make sense of data. Ironically, the way we discuss these various heuristics actually function as heuristics themselves. Also, our organization of these data is not the only way they can be organized. Some of the heuristics we have subsumed under one of these three major heuristics could arguably fit equally well under one or more of the others.

translate to the tasks performed by forensic evaluators. However, given that forensic evaluators are human, there are reasons to think these principles might apply. One purpose of this review is to stimulate research on the biases as they may affect forensic clinicians.

[Insert Table 1 about here]

Representativeness

The representativeness heuristic (Kahneman & Tversky, 1972; Tversky & Kahneman, 1974) is a mental shortcut in which the subjective probability of an event or sample is estimated based on its similarity to a class of events or a typical specimen. If X looks like a typical representation of Y, it may easily be perceived as an example of Y, even if Y is improbable. The vignette with which we began this paper provides an example. Although a defendant found NGI is much more likely to have a psychotic disorder than a personality disorder, the limited information provided about Mr. Jones characterizes him as a more stereotypic representation of a man with Antisocial Personality Disorder than a man with a psychotic disorder.

Think about this next example:

John P. is a meek man, 42 years old, married with two children. His neighbors describe him as mild-mannered but somewhat secretive. He owns an import-export company based in New York City, and he travels frequently to Europe and the Far East. Mr. P. was convicted once for smuggling precious stones and metals (including uranium) and received a suspended sentence of 6 months in jail and a large fine. Mr. P. is currently under police investigation. Please rank the following statements by the probabilities that they will be among the conclusions of the investigation. Remember that other possibilities exist and that more than one statement may be true. Use 1 for the most probable statement, 2 for the second, etc.

Mr. P. is a child molester.

Mr. P. is involved in espionage and the sale of secret documents.

Mr. P. is a drug addict.

Mr. P. killed one of his employees.

Mr. P. killed one of his employees to prevent him from talking to the police.

Does it seem like the last statement is more likely than the second-to-last statement? If so, your brain has made an intuitive but incorrect judgment by disregarding a basic law of probability: the *conjunction rule*. The probability of a conjunction, $P(A\&B)$ cannot exceed the probabilities of its elements, $P(A)$ and $P(B)$. In this particular example, the addition of a possible motive in the last statement (“to prevent him from talking to the police”) reduces the probability of the last statement compared to the second-to-last statement (because Mr. P might have killed his employee for a variety of other reasons). Therefore, if the A in this equation represents killing the employee to prevent him from talking to the police, and if B is all potential reasons for killing the employee, then A cannot equal more than A+B. Kahneman and Tversky (1983), using this example, found that many people ranked the last statement as more likely than the second-to-last. They termed this easy-to-make error the “conjunction fallacy.”

These findings are relevant to the work of forensic evaluators in many ways. Suppose for instance that a 16-year old evaluatee presents with symptoms of Attention Deficit-Hyperactivity Disorder (ADHD) and Bipolar Disorder, a not uncommon occurrence because the diagnostic criteria overlap. Given the conjunction rule, it is most likely the evaluatee would not have both disorders even if it appears so. The point here is that forensic evaluators should critically evaluate their decisions to diagnose both – the evaluatee could indeed have both disorders and their life could be impaired incrementally by each disorder. However, if the same pieces of data are being used to support both diagnoses (e.g., counting distractibility and excessive activity toward both), then perhaps the evaluator is in error by diagnosing both.

As can be seen in these examples, the representativeness heuristic can easily lead to *base rate neglect*. A *base rate* is the frequency with which a thing occurs. Base rate neglect is a ubiquitous phenomenon that affects laypeople and professionals alike. For example, Casscells,

Schoenberger, and Graboys (1978) asked Harvard Medical School faculty, staff, and fourth-year medical students the following question:

If a test to detect a disease whose prevalence is 1/1000 has a false positive rate of 5%, what is the chance that a person found to have a positive result actually has the disease, assuming that you know nothing about the person's symptoms or signs? _____%

The correct Bayesian answer under the most plausible interpretation of the problem is about 2%.² Specifically, 51 people out of 1000 would test positive (1 true positive and 50 false positives). Of the 51 people with positive tests, 1 would actually have the disease. Expressed as a proportion, this is $1/51 = 0.019$ or 1.9%. But only 18% of the Harvard-affiliated participants had an answer close to 2%. Forty-five percent of this distinguished group said the answer was 95%, thereby completely neglecting the base rate information.

An internationally recognized forensic psychiatrist made this error in the John Hinckley trial. On cross examination, the defense expert witness testified that Mr. Hinckley's brain had a particular brain anomaly in which his sulci were wider than normal, evidence the expert offered to support his conclusion that Mr. Hinckley had schizophrenia. On cross-examination:

Q: Isn't it true that the studies you are talking about indicate that most people who are schizophrenic don't have widened sulci?

A: To be precise about the word 'most.' In one study from St. Elizabeth's Hospital, one-third of the schizophrenics had widened sulci. That is a high figure. It is true that the simple majority didn't,

² Bayesian analysis is a modern approach to statistics named after an 18th-century English reverend named Thomas Bayes. He is credited with developing rules to explain how people should change their mind in light of evidence, as the evidence becomes available (Kahneman, 2011). For example, suppose a colleague told you they just finished meeting with an attorney about a potential new referral. Not knowing anything else about the attorney, you should believe that the probability that the attorney was a woman is 33.3% (this is the base rate of women in the legal profession; American Bar Association [ABA], 2013). Now suppose your colleague told you that s/he was impressed because the attorney was a managing partner at one of the largest law firms in the U.S. Taking into account this new information, where the base rate of women as managing partners in one of the 200 largest law firms in the U.S. is only 4% (ABA, 2013), then Bayes' theorem says you should believe that the probability that the attorney was a woman is now 2.04%. Although the equation is not listed out here, the example is included to demonstrate the importance of base rates in determining the probability of an occurrence.

but the fact that one-third had these widened sulci—whereas in normals, probably less than one out of fifty have them—that is a powerful fact.

Q: That is a fact?

A: Yes...It is a statistical fact, as I mentioned, that one-third of schizophrenics have widened sulci and probably less than two per cent of the normal people have them. That is a powerful statistical fact and it would bear on the opinion in this case (Caplan, 1984, as cited by Linder, n.d.).

We can use Bayesian reasoning to calculate the likelihood of having schizophrenia, given the presence of the brain anomaly. The base rate of schizophrenia in the general population is approximately 0.5% (American Psychiatric Association, 2013). Given the base rate, 50 out of every 10,000 people would have schizophrenia. According to the expert, of those 50 people with schizophrenia, one-third, or approximately 17 people, would have widened sulci. Of the 10,000 people, 9,950 would not have schizophrenia. Of the 9,950 people without schizophrenia, 1 out of 50, or 199 people, would have widened sulci. Thus, the number of people with widened sulci would be 216 (17 + 199). Of the 216 people with widened sulci, 17 would have schizophrenia. Expressed in a proportion, $17 / 216 = 0.078$ or 7.8% of people with widened sulci would have schizophrenia. Thus, we can see that the expert's conclusion about the evidence being "powerful" did not appear to account for the low base rate of schizophrenia in the population.

An example more typical of routine forensic practice may involve the assessment of future violence or recidivism risk, both of which have relatively low base rates in some forensic populations (e.g., roughly 10%; Campbell & DeClue, 2010; Monahan et al., 2005). The lower the base rate, the more challenging the assessment becomes, because even when evaluators use evidence-based tools, the tools are limited by the base rates. Campbell and DeClue (2010), for example, demonstrated what would happen if an evaluator used a very good sexual offense recidivism assessment tool (the Static-99; Hanson & Thornton, 1999) to estimate relative

likelihood of future violence of 100 different evaluatees. If the base rate was 19% (this is the base rate used in the example by Campbell and DeClue), and if the evaluator “bet the base rates” and predicted that all of these 100 people would be at very low risk to sexually re-offend within a 10 year period, the evaluator would be correct about 81% of the time. They asked whether employing the Static-99 could improve the accuracy of the assessment. They demonstrated that the overall accuracy to classify risk when using the measure would be about 76%, a decrease in overall accuracy compared to relying on the base rate alone. The lower the base rate (e.g., if it is 10% rather than 19%), the worse the predictions will be, compared to “betting the base rates.”

In clinical contexts, experts often underutilize or ignore base rate information and tend to rely instead on case-specific information (Carroll, 1977; Faust & Faust, 2012; Nickerson, 2004). The problem with this practice is that salient but less predictive case-specific information can draw the clinician’s attention away from the relevant base rates and have the adverse effect of decreasing accuracy (see e.g., Faust & Faust, 2012). Base rates are critical and should be part of a forensic evaluator’s thinking processes whenever possible.

Availability

The *availability heuristic* refers to the ease with which one can recall other examples of the event in question, which increases the likelihood of an interpretation (Tversky & Kahneman, 1973, 1974). In an early description of this heuristic, the philosopher David Hume (1976 [1736]) described the human tendency to judge an event’s probability by how “fresh” it is in memory. Other factors that increase availability are frequency and salience. For example, consider the task forensic clinicians are asked to conduct in violence or sexual offending risk assessments. A false negative occurs when a person is assessed as unlikely to reoffend but does in fact reoffend. Imagine a high-profile re-offense ending up in the newspaper, the evaluator being asked by the

local news media, “How could you have missed this,” and the evaluator’s employer undertaking a performance review of his or her work as a result. False negatives are likely to be much more memorable than a false positive (i.e., assessed as likely to reoffend but does not). The perceived likelihood of future reoffending might be overestimated due to the availability of information about instances in which a clinician was incorrect and the salience of anticipated regret.

Before reading further, please complete this exercise.

There is a set of cards. On Side 1 is a letter and on Side 2 is a number. The following hypothesis is proposed: “*Whenever there is an A on Side 1, there is a 2 on Side 2.*” Please look at these 4 cards, two with Side 1 exposed and two with Side 2 exposed:



The question is: *How many cards, and which ones, could you turn over to effectively test the hypothesis?*

If you make an error, you’ll be in good company. Most people (even trained scientists and professionals like both of the authors of this paper) have trouble correctly answering this question. The answer is 2. Two cards can effectively test the hypothesis: the “A” and “1” cards. Finding a “1” on the back of the “A” would allow you to reject the hypothesis, as would finding an “A” on the back of the “1” card. No other possibilities would reject the hypothesis (task adapted from Wason, 1968). Turning over “B” will not do it because it is not relevant to the hypothesis – no matter what is on the other side, it will not be helpful for testing the hypothesis. Turning over the “2” will not allow you to reject the hypothesis either; seeing an “A” on the reverse side would only confirm the hypothesis, and seeing a “B” would not tell you anything about the hypothesis. If you thought the “2” should be turned over, you engaged in a cognitive error called the *positive test strategy* (Kayman & Ha, 1987). Positive test strategy is a mental heuristic whereby hypotheses are tested exclusively (or primarily) by searching for evidence that has the best chance of verifying current beliefs, rather than those that have the best chance of

falsifying them. And the evidence suggests that this kind of bias is pervasive, even in the absence of any particular outcome motivations (Fischhoff & Beyth-Marom, 1983; MacCoun, 1998).

How is this example of positive test strategy relevant to forensic practice? It demonstrates the *confirmation bias*, which may plague forensic clinicians. Confirmation bias is a set of tendencies to seek or interpret evidence in ways that are partial to existing beliefs, expectations, or hypotheses (Nickerson, 1998). Turning over cards to confirm the hypothesis is not an effective method of testing the hypothesis because ruling out the possibility that something is false is a step toward demonstrating its truth.³ The ubiquity of this error has important implications for forensic evaluators. Mackay's compelling line demonstrates this bias well (1932 [1852], p. 552), "When men wish to construct or support a theory, how they torture facts into their service!" (Women, too.)

Clinicians may make conclusions based on inadequately formed hypotheses, they may not gather the necessary data needed to adequately test their hypotheses, and they may seek and rely mainly or exclusively on information that confirms their "hunch." Imagine that you have returned from vacation to find a new evaluation assignment in your box that is due in just a few days. In this case, you might be more likely to generate a "hunch" quickly (maybe within the first few minutes of meeting with the evaluatee), then look for information that would confirm your intuition. Doing so means you can finish your evaluation and report quickly. Even in less time-sensitive circumstances, evaluators may engage in these behaviors because doing so means the work gets done more quickly and it is less effort-intensive than methodically testing alternative hypotheses and seeking information that might *disprove* one's intuition.

³ This principle is rooted in Karl Popper's (1959) principle of falsificationism. Popper argued that whereas induction could never confirm a hypothesis, induction might permit one to *falsify* it (e.g., "If p then q; not q; therefore, not p") (MacCoun, 1998). Popper contended that falsification permits us to weed out bad ideas while seeing how our leading hypotheses hold up under attack (MacCoun, 1998).

An evaluator's initial hypothesis or "hunch" might be made based on the evaluator's own personal and political beliefs, exposure to pretrial publicity (e.g., suggestibility and expectancy), or comments from the referral party regarding their hypotheses about the defendant's mental health. For instance, imagine a defense attorney calls to tell you about a case and see if you are interested in the referral. The attorney says, "I have this really mentally ill guy who is being railroaded by the system" as compared to a prosecutor who calls about the same case and says, "This guy is faking – there's no way he's sick. We need your help to prove it." Other examples involve information arising from the institutional environment. For example, imagine conducting an annual evaluation for an NGI acquittee to help the court determine whether s/he continues to meet commitment criteria. If you work in this institution, you might have repeatedly heard about how "crazy" or "dangerous" the patient is from numerous co-workers, which might influence the way you perceive the evaluatee and the conclusions you reach. All of these sources of opinion can have subtle effects that set up the potential for confirmation bias.

Confirmation bias may also occur due to sharing a preliminary opinion before the evaluation is complete. For instance, forensic mental health clinicians might be asked to answer questions about the way they are "leaning" in a case based on their initial interpretation of partially-collected data (by retaining parties, supervisors, colleagues). Answering such questions prematurely commits the examiner in a way that makes it more difficult to resist confirmation bias when completing the final interpretation of one's data.

Confirmation bias may range on a continuum from unmotivated on the examiner's part (see e.g., Faust & Faust, 2012) to motivated (see e.g., Festinger, 1957; Kunda, 1990). Motivated reasoning may allow an evaluator to arrive at a particular desired conclusion, constrained only by the evaluator's ability to construct reasonable justifications for their conclusion (Kunda, 1990).

In a recent study, Mannes and Moore (2013) demonstrated that people tended to do a poor job of adequately adjusting their initial estimates after subsequently receiving relevant new information or information about the consequences of being wrong. They concluded that people, driven by their subjective confidence, tend to have unwarranted and excessive faith in the accuracy of their own judgments. Dana, Dawes, and Peterson (2013) showed that interviewers who used unstructured interviews were able to “make sense” out of virtually anything the interviewee said (even when the responses were nonsense because the interviewee answered questions using a random response system). The interviewers formed interview impressions just as confidently after getting random responses as they did when they got real responses.

WYSIATI (What You See Is All There Is) is a heuristic concept describing an apparent design feature of the human brain: only activated ideas are processed within a given cognitive task or decision making procedure (Kahneman, 2011). Information that is not retrieved while a person analyzes or interprets information might as well not exist. Our brains are designed to create the most coherent explanation out of the available information, and it is the consistency of the information that matters rather than its completeness (Kahneman, 2011). This design feature can also lead to systematic biases (e.g., base rate neglect, overconfidence, confirmation bias), which is what makes searching for disconfirming information so difficult. *WYSIATI* is relevant to explain how forensic evaluators form hypotheses, how they search for information to test their hypotheses, how they interpret the information they uncover, how they reach a decision, and how they communicate that information to the trier of fact.

Let us return to the opening example about Mr. Jones. There were at least 13 discrete pieces of information in the vignette. How many pieces of information can you recall about him (try without looking before you look back)? Although this question is somewhat relevant to

forensic clinician's tasks, the more relevant question is how many pieces of information, and which ones, would you be thinking about when trying to integrate the data to reach a conclusion about the referral question (mental state at time of offense)?

Faust and Faust (2012) describe this process of "configural analysis" (trying to integrate all the information gathered, analyzing the patterns of interrelationships among the various pieces of data), and they note that studies of human information processing suggest forensic clinicians are likely to be able to analyze the patterns of interrelationships only for about four discrete pieces of information at a time. This is how WYSIATI works in real life. Clinicians are limited by the "bounded rationality" of being a human being – we are all constrained by the limitations of our brain's design. Even if forensic clinicians do their due diligence and identify more than four pieces of critical information relevant to the referral question, when it comes to formulating their "bottom line," forensic clinicians likely focus on the four(ish) pieces of information they interpret as most relevant to the question (whether they realize it or not).

Anchoring

The *anchoring effect* is a cognitive phenomenon in which we are overly influenced by initial information encountered (Tversky & Kahneman, 1974). Anchoring, akin to *priming* and the *halo effect*, increases the weight of first impressions sometimes to the point that subsequent information is mostly wasted (Kahneman, 2011). The sequence in which we encounter information is often determined by chance, but it matters. Forensic evaluators who perform child custody evaluations, for example, may be well aware of what happens in speaking with one party at a time. Imagine first meeting with the mother, who presents as a smart, articulate, attractive 35-year old professional. She comes across as non-defensive, fully capable of parenting safely and competently, and tells you that her child's father is domineering, emotionally abusive to

herself and their daughter, and that he is dragging out the custody battle out of spite. You are struck by how likeable and credible she is, and before you even meet with the father, you may have a pretty sound “hunch.” Now imagine the counterfactual, in which you meet with the father, who presents as a smart, articulate, attractive 35-year old professional. He comes across as non-defensive, fully capable of parenting safely and competently, and tells you that his child’s mother is domineering, emotionally abusive to himself and their daughter, and that she is dragging out the custody battle out of spite. You are struck by how likeable and credible he is.

Any forensic evaluator might hear a coherent and compelling story told by the first person interviewed and begin to formulate hypotheses about the case, only to hear a different (and perhaps contradictory) story later from another party that might be just as coherent and compelling. Unfortunately, people often have difficulty sufficiently adjusting an original hypothesis based on information encountered later. The evaluator must somehow make sense of the contradictory information and beware the anchoring effect of the information from the first party whom the evaluator happened to interview

Framing and context effects are particularly relevant to forensic work. Framing is a cognitive heuristic in which people tend to reach conclusions based on the framework within which the situation was presented (Tversky & Kahneman, 1981). People may draw different conclusions on the basis of the same information, depending on how the information was framed or the context in which it was delivered. There is a built-in system of framing inherent in adversarial legal systems. The recent body of research on adversarial allegiance in forensic experts, which shows that mental health professionals may reach conclusions and opinions consistent with the goals of their retaining party (see e.g., Murrie, Boccaccini, Johnson, & Janke, 2008; Murrie et al., 2009; Murrie et al., 2013), might be interpreted with regard to this bias.

Building on their previous body of work, Murrie and colleagues (2013) conducted an elegant scientific experiment which demonstrated this effect. They “hired” forensic psychiatric and psychological experts to review a defendant’s case file and score the offender on two commonly used, well-researched risk assessment instruments. The experts were led to believe they had either been hired by the defense or the prosecution. That is, the manipulated independent variable was the adversarial “side” for which experts thought they were working (case and offender materials were held constant). Experts who believed they were working for the defense tended to assign lower scores on the risk instruments than experts who believed they were working for the prosecution. The effect sizes were up to $d = .85$ (large effects).

Murrie and colleagues attributed the adversarial allegiance effect directly to experts’ beliefs about for whom they were working, because they controlled for other possible explanations. The substantive information provided about the defendant was constant, so differences in the way the examinee presented could not have explained the findings. Furthermore, they eliminated the overt verbal influence often provided by the referral party in routine forensic practice that contributes to confirmation bias. This design element is important: their findings show that even when there is no overt framing by a referral party, there is still an insidious yet potentially potent form of anchoring due to adversarial allegiance.

Can We Reduce Bias in Forensic Practice?

We have reviewed how psychological science has come to better understand subtle but powerful sources of bias in human decision making. We focused the review especially on cognitive processes that might threaten the accuracy of forensic clinicians’ thinking when they are formulating their hypotheses during or after acquiring evaluation data. Throughout the

review, we provided examples to show how forensic examiners' use of evaluation data could potentially fall prey to subtle sources bias and error that can affect their conclusions.

We now pose the question, "Should we be concerned?" Let us accept the existence of such potential sources of bias in judgment and decision making. Let us presume that forensic mental health examiners, like other humans, are susceptible to heuristic and other cognitive sources of bias in their processing of information. Is there reason for forensic mental health assessment as a field, and forensic examiners individually, to take this on as a problem? Or should we accept our inevitable human fallibility while taking the lessons simply as a warning that we must "exercise due caution?"

We think there are good reasons to be concerned. Scientific and clinical expertise in the courtroom is dependent on the expectancy that the expert seeks accuracy and avoids anything that may lead to bias in the collection or interpretation of data. Challenging that expectancy is a growing body of research suggesting that forensic examiners differ in the data they collect and the opinions they reach, depending on the social contexts in which they are involved in forensic cases (e.g., Brown, 1992; Deitchman, 1991; Homant & Kennedy, 1986, 1987a, 1987b; Murrie et al., 2013; Neal, 2011; Svec, 1991). These studies are identifying the results of error, bias, and inaccuracy in our work. And decision making science offers plausible ways to explain it.

Failure to address the questions runs counter to our professional obligation to be accountable for our performance and to strive for the integrity of our opinions. Moreover, failure to address the questions degrades our perceived credibility.

Framing the Problem

As one considers the theories of bias in decision making, it is apparent that the ultimate objectives are (a) to explain ways in which human decisional processes fail to achieve accuracy

and (b) uncover means of correcting the errors to improve accuracy. These objectives parallel the psychometric notion of validity. We might ask how bias and error detract from achieving valid answers to forensic referral questions. However, for most forensic questions, there is no obvious criterion variable representing the “truth,” no touchstone with which to evaluate the validity of our answers to the questions.

An imperfect substitute for improving our validity is to improve our reliability, at least as an initial goal (see e.g., Mossman, 2013). We might propose that two clinicians should arrive at similar clinical assessments. That sameness—the essence of reliability—can be our touchstone as we explore the influences of bias on forensic opinions. This would not necessarily mean the two clinicians arrived at the right opinion. Reliability does not guarantee validity; it merely assures a condition without which validity cannot be achieved.

Given that we frame the problem of cognitive bias and error in forensic practice as a problem of reliability (at least for now), we consider three ways to respond to the problem. One of these approaches seeks understanding and the second seeks change. The third imagines a paradigm shift with potential to mitigate the problem.

Discovering the Extent of the Problem

Recent studies of the opinions of forensic clinicians in cases involving multiple experts have identified what they interpret as substantial unreliability in those opinions (see e.g., Boccaccini, Turner, & Murrie, 2008; Boccaccini, Turner, Murrie, & Rufino, 2012; Gowensmith, Murrie, & Boccaccini, 2012, 2013; Murrie et al., 2008, 2009, 2013; Murrie, Boccaccini, Zapf, Warren, & Henderson, 2008; Murrie & Warren, 2005). Moreover, the unreliability appears to be largely related to the examiners’ agency, allegiance, and sometimes personality and attitudes. Those studies tell us something about the conditions in which bias might arise, but they do not

enlighten us about how bias works. They tell us nothing about heuristics and biases at the level of cognitive effort during the collection, sorting, and use of information in forensic cases. It is reasonable to hypothesize that the types of errors in logic and heuristic thinking described by cognitive science “serve” the social and personal influences that appear to drive unreliable outcomes. Can we develop a body of research that connects variance in forensic opinions with specific sources of bias or errors in logic during data collection or case formulation?

One type of research would identify whether and how various cognitive heuristics occur when forensic clinicians are processing data in their case formulation. There is little reason to believe that forensic clinicians are any more or less immune to heuristics and biases than are any other similarly-intelligent decision makers in unstructured decision tasks. Documenting this in the context of forensic case formulation, and identifying specific types of more common biases in processing cases, would provide a fundamental start for identifying the extent of the problem.

A line of research could explore the “dynamics” of heuristics and biases in forensic clinicians’ cognitive processing of cases. Under what social conditions are the various sources of cognitive bias and error increased or decreased? This research might show that conditions of agency, allegiance, or other incentives augment the play of cognitive heuristics and biases when formulating cases, offering a cognitive explanation for unreliability in forensic formulations. To what degree do decision aids and structured methods decrease bias and improve accuracy?

One of the challenges of this work will be delineating the elements of the “forensic evaluation process” where biases and errors may exert an effect. There may be many ways to construe the process, but here we offer a simple one to provide an example. The forensic evaluation process begins with a referral question that guides the evaluation. The evaluation then includes: (a) selection of types of data to collect; (b) collection of the data; (c) analysis of

the data; and (d) interpretation of the data to formulate a forensic opinion. These domains within the evaluation process might allow us to discover how bias works in various ways associated with the steps in the process. Considered respectively, they offer the potential to determine how biases (a) narrow or expand our search for relevant data, (b) influence the quality and integrity of the data that we collect, (c) influence how we score or subjectively classify the data we have obtained, and (d) influence how we combine the data when testing our hypotheses and their alternatives regarding the answer to the forensic question.

Finding Remedies that Overcome Bias

A second way to respond to the problem is to seek ways to reduce clinicians' susceptibility to biases that increase unreliability. One might suppose that the line of research described in the first approach – discovering the extent of the problem – would be a required prelude to the identification of remedial strategies. Yet we should consider the possibility that the two lines of research could proceed in parallel fashion. Indeed, the former studies might sometimes include design features that explore remedial potentials.

What do we already know about debiasing strategies? The decision-making field has described various ways to make people aware of the positive and negative effects of heuristics on their decision making (see e.g., Gawande, 2009; Kahneman, 2011). Indeed, the potential sources of bias described earlier in this paper are known to many in business, medicine, education, and science. Yet arming people with insight into sources of heuristic error does not guarantee that the insights will be used. That should not be surprising. Almost all adaptive functioning requires not merely knowledge, but also motivation and practice. At minimum, putting such information to use would seem to require a desire to avoid bias and error and alertness to conditions in which such bias and error can occur (Kahneman, 2011). Efforts to change

individuals' cognitive heuristics could also be fashioned as educational procedures that allow individuals to recognize, practice, and repetitively rehearse positive heuristic methods while processing information that is relevant for them—in this case, formulating forensic cases.

One specific debiasing strategy that has a good chance of being useful for forensic clinicians is locating and then keeping in mind relevant base rates (Kahneman, 2011). For example, Schwartz, Strack, Hilton, and Naderer (1991) found that instructing people to “think like a statistician” enhanced the use of base-rate information, whereas instructing people to “think like a clinician” had the opposite effect. Kahneman (2011) suggests that the corrective procedure is to develop a baseline prediction, which you would make if you knew nothing specific about the case (e.g., find the relevant base rate). Second, determine whether the base rate matches your clinical judgment about the case. When thinking about your clinical judgment, always question the strength of the evidence you've gathered (How sound is the evidence? How independent are the observations? Don't confuse correlation with causation, etc.). Then aim for a conclusion somewhere between the baseline prediction and your clinical judgment (and stay much closer to baseline if the evidence underlying the clinical judgment is poor). Clinicians often tend to exaggerate the persuasiveness of case-specific information (Faust & Faust, 2012; Kahneman, 2011). Therefore, anchoring with a base rate and then critically evaluating the strength of the case-specific diagnostic information are offered as recommendations to combat the representativeness heuristic and overconfidence (Kahneman, 2011).

Another debiasing strategy is to “consider the opposite” (Koehler, 1991; Lord, Lepper, & Preston, 1985). This strategy may be particularly useful for forensic clinicians, given the adversarial nature of courtroom proceedings. Expert witnesses testify through direct- and cross-examination. Imagining how one's assessment methods, data, and interpretations will be

scrutinized during cross-examination is often recommended as a trial preparation strategy (e.g., Brodsky, 2012). Part of that consideration is recognizing that the opposing side will not merely attack the proof for the clinician's opinion, but might also pose alternative interpretations and contradictory data, asking the clinician why they were rejected. Similarly, Lord and colleagues (1985) found that people could mitigate confirmation bias when asked a question like, "Consider how you'd evaluate the case given opposite results," while a global motivational instruction to "try to be unbiased" was not an effective bias mitigation strategy.

Structure and systematic methods. The notion that observers' interpretation of evidence might be influenced by personal interests and prejudices dates back at least to Sir Francis Bacon, who is credited with advocating the scientific method in the sixteenth and seventeenth centuries (MacCoun, 1998). Structure and systematic methods are the backbones of the scientific method. What do we already know about how structure, standardized procedure, and evidence-based decision making improve the reliability and validity of forensic evaluations? The closely-related body of research on the poor reliability of clinical judgment sheds light on this question (see e.g., Faust & Faust, 2012 for an overview).

Many studies have shown that structured methods improve forensic assessments as compared to unstructured clinical judgments. For example, in a meta-analysis of recidivism risk assessments for sexual offenders, Hanson and Morton-Bourgon (2009) showed that actuarial measures (such as the Static-99) were considerably more accurate than unstructured clinical judgment for all outcomes (sexual, violent, or any recidivism). In a meta-analysis of violence risk assessments, Guy (2008) found that if evaluations employed structured professional judgment tools (such as the HCR-20; Webster, Douglas, Eaves, & Hart, 1997) and actuarial tools

(such as the VRAG; Harris, Rice, & Quinsey, 1993), they performed better in predicting antisocial behavior than assessments that relied on an unstructured clinical judgment approach.

These methods might improve forensic evaluations by helping forensic clinicians minimize the effects of their biases on their work. Research findings have shown that allegiance effects (i.e., assessments scored in the direction that would be preferred by the adversarial retaining party) are stronger for more subjective measures but attenuated with more structured measures. Murrie and colleagues (2009) found that allegiance effects were stronger for the Psychopathy Checklist-Revised (PCL-R; Hare, 2003), a measure that requires more subjective judgments in scoring, than for the Static-99, which requires less clinical judgment. Interrater agreement between evaluators working for opposing parties were stronger for the Static-99 (ICC = .064) but weaker for the PCL-R (ICC = 0.42). These findings were replicated in the 2013 paper by Murrie and colleagues. Specifically, effect sizes for the allegiance effect with the PCL-R were up to $d = 0.85$ (large effect), whereas the effect sizes for the Static-99R (Helmus, Thornton, Hanson, & Babchishin, 2012) were up to $d = 0.42$ (small-to-medium effect). Thus, more structured measures appear to be associated with higher inter-rater agreement rates and lower adversarial allegiance bias. Perhaps more structured measures also attenuate the other kinds of biases outlined in this review. Further research might inform potential bias-mitigation remedies.

Another strategy to increase reliability and validity, and likely decrease bias, would be to identify approximately four to six variables essential to the referral issue at question. The dimensions should be as independent as possible from each other, and they should be amenable to reliable assessment (they should be highly valid, highly reliable indicators). Identifying these essential dimensions for consideration with each kind of referral question might decrease the

time, money, and other resources spent on cases, and might also increase the other quality indicators described above (e.g., Faust & Faust, 2012; Gawande, 2009; Kahneman, 2011).

Take, for example, an evaluator tasked with assessing a defendant's competency to stand trial (CST) abilities. The few essential variables might include active symptoms of mental illness (which could be assessed with existing valid and reliable measures), intellectual or cognitive capacity (could also be assessed reliable with valid measures), ability to remain in behavioral control or inhibit impulsive behaviors (could be indexed through a review of the defendant's recent history and history during similar episodes, if there have been previous episodes), and perhaps degree of malingering (could be assessed with an existing response style indicator, particularly one developed to measure CST-related malingering). Although variables like educational attainment, age, and history of mental illness might be related to the referral question, these variables are unlikely to be as useful for generating a sound conclusion. Educational attainment could be a proxy for cognitive capacity, but measuring the capacity is a more valid and reliable way to index this particular trait (and education and cognitive abilities should not both be considered essential variables, because they are not as independent of one another as would be ideal). History of mental illness might be relevant, but would be less relevant than current symptoms of illness. However, particular behaviors during previous episodes might be essential (e.g., demonstrated difficulty inhibiting impulsive verbal and physical outbursts in previous court appearances during periods of active psychosis).

Finally, research might seek to develop an anti-bias linear or branching procedure that clinicians could follow in their practice, or perhaps "checklists" of essential elements for clinicians to systematically consider for various situations. A measure for violence risk assessment called the Classification of Violence Risk (COVR; Monahan, Steadman, Appelbaum

et al., 2005) is one example. It relies on classification-tree analysis to allow several different variables to be considered together. It produces an actuarial estimate of risk, but is intended to be used by clinicians as one piece of information upon which to base conclusions and decisions. This kind of tool and other “checklist” methods, like structured professional judgment tools, might guide forensic clinicians around the biases that arise in various steps of the evaluation process while encouraging positive heuristics.

Imagining a Systemic Adjustment

The two approaches we have described focus on understanding how forensic examiners think and then potentially modifying their thinking. These approaches consider the problem a matter of faulty human performance that requires changing the individual. Yet psychology has long recognized that problems in human functioning sometimes admit to a different type of analysis that refocuses the problem as one of person-environment fit. If one can change the individual to fit the environment’s demands, sometimes that is preferable. But if that proves too difficult, one can at least consider whether one can change the environment—the context in which the individual must function—to reduce the problem.

It is proper, therefore, for our analysis to conclude by imagining the development of a legal context that would mitigate the problem of bias in forensic practice. This leads us to consider a context that we do not advocate. We offer it in the form of a “thought experiment,” potentially for its “heuristic” value as our field searches for creative solutions to the problem of forensic examiner bias.⁴

Forensic experts are expected to use their science to lead them to an objective opinion. When two experts in the same field arrive at different opinions, judges and juries tend to assume

⁴ See MacCoun, 1998 for a similar analysis of inquisitorial versus adversarial models of science in his *Annual Review of Psychology Article* called “Biases in the interpretation and use of research results.”

either of two things. The experts' science is unreliable, or the experts used their science in a biased manner. But let us imagine a legal system that does not expect forensic clinicians to arrive at similar opinions. It expects the forensic experts to endeavor to build an argument for a conclusion that favors the party that called them, developing the best version possible that the available data will allow. In the paradigm we are imagining, the expert is relieved of the task of finding the explanation that most objectively fits the data. It allows the expert to create a "favored" interpretation that is consistent with, *and not contradicted by*, the data. In some cases the expert will find that it simply cannot be done. If so, the party is free not to call the expert to report or testify, and law could shield the expert's investigation from discovery. In many cases, the expert will be able to build an interpretation of the available data and a conclusion—sometimes more plausible, sometimes less—that fits the interests of the party.

One defense of this paradigm rests with the law's adversarial process of arriving at just decisions when the "truth" is often obscure. To resolve disputes, our legal system developed over time a structured and transparent process (called "due process" in the U.S.) through which it considers multiple perspectives in an adversarial framework. The system's adversarial framework attempts to achieve fairness through procedural justice (Rawls, 1971). Recognizing that truth may be unknowable in human affairs, it is *procedure* that matters most, so that if trustworthy processes are followed, truth can be closely approximated even if it is not attained. In this adversarial framework, the trial process relies on two parties to make arguments in favor of their opposing opinions. Both parties must investigate the potential evidence and construct their interpretation of the evidence that best fits their position in the case. This system presumes that many fact circumstances can be interpreted in different ways, and it makes no presumptions about the differential validity of the interpretations.

The analysis thus far has focused only on the interpretation of the data to reach a conclusion. What about the data themselves? Is there not a danger that the expert will selectively seek, or selectively find, only those data that will support the party's argument? Again we find a possible parallel with the role of attorneys in the criminal justice process. They are ethically bound to engage in competent practice that offers zealous advocacy for their party. Arguably, that ethical obligation is breached when an attorney seeks only evidence that will support their party's argument. In doing so, the attorney might well miss data that could support the other party's argument. Having missed it, the attorney is unprepared to defend against it, and thus does damage to the party to which allegiance is owed, therefore practicing incompetently.

In this revised forensic-clinical paradigm, the situation of the forensic expert working with either party would be much the same. The expert is expected to participate in the adversarial process (rather than being removed from it) by seeking the best interpretation for the party's conclusion that can be supported by the data. But seeking only those data that support the conclusion will weaken the plausibility of the interpretation. It will not stand up to other data that the opposing party might have found and against which the expert will be unprepared to defend. Thus the expert will be ethically obligated to find all data, not just data that support the adversarial hypothesis. Moreover, those data must be competently and reliably obtained, scored, and interpreted in light of past research. Just as attorneys using inaccurate information may harm their clients' interests, so experts' unreliable data may do the same.

This approach is often used to stimulate academic debate, with two experts constructing the two best opposing theoretical interpretations they can create using existing observations. (In fact, the argument we are currently constructing is based on such a model). The matter is somewhat different, though, for clinical professionals. Typically they are obligated to arrive at

opinions in a neutral and objective manner, so that their opinions will avoid error that may harm their patients. In the alternative paradigm we outline here, society might assign experts a role that does not require neutrality, but rather asks them to exercise their special expertise to produce the two best-supported but perhaps contradictory views of a case.

Forensic examiners' primary obligation would be (a) accuracy in one's use of objective methods of data collection, (b) integrity in the description of one's data, and (c) clarity and honesty in describing the manner in which the data led to one's conclusion. Put another way, one's data must be reliable, and one's processing of the information (typically called one's interpretation of the data) must be explained. Given this paradigm, reliability between two examiners should be required for (a), but would be variable for (b) and (c).

We offer this paradigm shift for further consideration. Undoubtedly a more penetrating analysis will find the potential for consequential damage in this model. There may also be flaws in the analogy, which may or may not be remediable. For example, attorneys are obligated to seek all information that may support and refute their party's position, but are they required to reveal it? In criminal cases, if prosecutors know of data that would harm the state's position and favor the defendant, they *must* reveal it. Defense attorneys have no such obligation. Our legal system's protection of individual liberty requires the state to prove guilt and provides defendants the right to withhold information that might incriminate them. Were this to apply to experts, the expert for the prosecution would operate much as is expected now, revealing all sources and types of information that were obtained in the evaluation. But defense experts would reveal only the data that they collected that favored the defense, while withholding all negative data.

The problems with this are quite evident. For example, if a defense expert administered an instrument measuring psychopathic traits, would the score be reported only if it were below

the psychopathy cut-off? Would we report only psychopathology test scores for scales on which the client scored in a direction favoring the defense argument, while leaving out scores on the other scales? Furthermore, how might adopting an advocacy role affect the examiner's ability to collect, score, and interpret "objective" tests? Explicit adoption of an adversarial role might lead to even greater differences in test data obtained by defense and prosecution experts than have been uncovered by researchers to date (e.g., Murrie et al., 2013). One potential remedy would be a requirement that the same database be used by explicitly adversarial experts. For example, a court-appointed psychometrician might administer and score the relevant test instruments requested by the adversarial parties, then the adversarial experts might be expected to interpret and make use of those data alongside the other data in forming their adversarial conclusions and opinions. Such an approach might maximize objectivity (in data collection) but also make best use of the adversarial nature of the fact-finding legal justice system.

Conclusion

Recent studies have examined the relation of personal and situational variables, including those that can be explained by judgment and decision making science, to forensic examiners' opinions. These studies have provided sufficient evidence of the need to address the issue of bias in forensic mental health evaluations. Not all agree that the current studies actually indicate bias (Mossman, 2013). But there is mounting evidence contrary to the notion of random error. We anticipate that increasingly we will be held accountable in the courts to explain recent research evidence that challenges our objectivity and credibility.

Accordingly, we have offered a review of various sources of bias in decision making that might provide frameworks and concepts for future studies. If these studies provide better understanding of the phenomenon, we envision the creative development of ways that clinicians

can reduce the effects of bias when they are processing data and arriving at opinions. Finally, we imagined a legal context that might change the role of forensic examiners in a way that accepts adversarial participation through expert evidence – a legal context full of practical, scientific, and ethical questions. These questions may or may not be worth trying to answer as we strive to improve the validity and reliability of forensic mental health evaluations and to foster trust in our work process and products.

References

- Boccaccini, M.T., Turner, D., & Murrie, D.C. (2008). Do some evaluators report consistently higher or lower psychopathy scores than others? Findings from a statewide sample of sexually violent predator evaluations. *Psychology, Public Policy, and Law*, *14*, 262-283. doi: 10.1037/a0014523
- Boccaccini, M.T., Turner, D., Murrie, D.C., & Rufino, K. (2012). Do PCL-R scores from state or defense experts best predict future misconduct among civilly committed sexual offenders? *Law and Human Behavior*, *36*, 159-169. doi: 10.1037/h0093949
- Brown, S. (1992). *Competency for execution: Factors affecting the judgment of forensic psychologists*. Unpublished dissertation, University of North Dakota.
- Cochrane, R. E., Grisso, T., & Frederick, R. I. (2001). The relationship between criminal charges, diagnoses, and psycholegal opinions among federal pretrial defendants. *Behavioral sciences & the law*, *19*, 565-582. doi: 10.1002/bsl.454
- Deitchman, M.A. (1991). *Factors affecting competency-for-execution decision-making in Florida forensic examiners*. Unpublished Dissertation, Florida State University.
- Gowensmith, W.N., Murrie, D.C., & Boccaccini, M.T. (2012). Field reliability of competency to stand trial evaluations: How often do evaluators agree, and what do judges decide when evaluators disagree? *Law and Human Behavior*, *36*, 130-139. doi: 10.1037/h0093958
- Gowensmith, W., Murrie, D.C., & Boccaccini, M.T. (2013). How reliable are forensic evaluations of legal sanity? *Law and Human Behavior*, *37*, 98-106. doi: [10.1037/lhb0000001](https://doi.org/10.1037/lhb0000001)
- Homant, R. J. & Kennedy, D. B. (1986). Judgment of legal insanity as a function of attitude toward the insanity defense. *International Journal of Law and Psychiatry*, *8*, 67-81.

- Homant, R. J. & Kennedy, D. B. (1987a). Subjective factors in clinicians' judgments of insanity: Comparison of a hypothetical case and an actual case. *Professional Psychology: Research and Practice*, 5, 439-446.
- Homant, R. J. & Kennedy, D. B. (1987b). Subjective factors in the judgment of insanity. *Criminal Justice and Behavior*, 14, 38-61.
- Kahneman, D. (2011). *Thinking, fast and slow*. New York: Farrar, Straus and Giroux.
- Murrie, D.C., Boccaccini, M.T., Guarnera, L.A., & Rufino K. (2013). Are forensic experts biased by the side that retained them? *Psychological Science*, online first publication August 22, 2013. doi:10.1177/0956797613481812.
- Murrie, D.C., Boccaccini, M., Johnson, J., & Janke, C. (2008). Does interrater (dis)agreement on Psychopathy Checklist scores in Sexually Violent Predator trials suggest partisan allegiance in forensic evaluation? *Law and Human Behavior*, 32, 352-362. doi: 10.1007/s10979-007-9097-5
- Murrie, D.C., Boccaccini, M.T., Turner, D., Meeks, M., Woods, C., & Tussey, C. (2009). Rater (dis)agreement on risk assessment measures in sexually violent predator proceedings: Evidence of adversarial allegiance in forensic evaluation? *Psychology, Public Policy, and Law*, 15, 19-53. doi: 10.1037/a0014897
- Murrie, D.C., Boccaccini, M., Zapf, P.A., Warren, J.I. & Henderson, C.E. (2008). Clinician variation in findings of competence to stand trial. *Psychology, Public Policy, and Law*, 14, 177-193. doi: 10.1037/a0013578
- Murrie, D.C., & Warren, J.I. (2005). Clinician variation in rates of legal sanity opinions: Implications for self-monitoring. *Professional Psychology: Research and Practice*, 36, 519-524. doi: 10.1037/0735-7028.36.5.519

- Neal, T.M.S. & Brodsky, S.L. (under review a). *A qualitative investigation of bias (and potential correction strategies) in forensic evaluations.*
- Neal, T.M.S. & Brodsky, S.L. (under review b). *Forensic evaluators' capital punishment attitudes and moral disengagement contribute to partisan participation in capital cases.*
- Svec, K.A. (1991). *Decisions about competency for execution as a function of attitudes toward capital punishment.* Unpublished master's thesis, University of Alabama, Tuscaloosa.
- Warren, J. I., Murrie, D. C., Chauhan, P., Dietz, P. E., & Morris, J. (2004). Opinion formation in evaluating sanity at the time of the offense: An examination of 5175 pre-trial evaluations. *Behavioral Sciences & the Law*, 22(2), 171-186. DOI: 10.1002/bsl.559

Table 1. List and definition of various cognitive biases in forensic assessment.

Bias	Related biases	Definition
Representativeness	<i>Conjunction Fallacy</i>	Overemphasizing evidence that resembles a typical representation of a prototype. <i>A compound event is judged more likely than is one of its elements alone.</i>
	<i>Base Rate Neglect</i>	<i>Judging an outcome's likelihood without considering information about the actual probability that it will occur.</i>
Availability	<i>Confirmation bias</i>	Overestimating the probability of an occurrence when other instances are relatively easy to recall. <i>Selectively gathering and interpreting evidence that confirms a hypothesis and ignoring evidence that might disconfirm it.</i>
	<i>WYSIATI</i> <i>(What You See Is All There Is)</i>	<i>Activated information is organized to derive the most coherent "story" possible (non-activated information is left out).</i>
Anchoring		Information encountered first is more influential than information encountered later.
	<i>Framing/Context</i>	<i>Drawing different conclusions from the same information, depending on how or by whom that information is presented.</i>

Erratum, “Correction to Neal and Grisso (2014)” published in volume 20, page 345. doi:

10.1037/law00000022

Abstract:

Reports an error in "The cognitive underpinnings of bias in forensic mental health evaluations" by Tess M. S. Neal and Thomas Grisso (*Psychology, Public Policy, and Law*, 2014[May], Vol 20[2], 200-211). This article contained an error that ironically demonstrates the very point of the article: that cognitive biases can easily lead to error—even by people who are highly attuned to and motivated to avoid bias. The authors inadvertently misapplied base rates by failing to account for nested probabilities in the illustration of how Bayesian analysis works in footnote 2 (p. 203). Specific details are provided.

Text of erratum:

The article, “The Cognitive Underpinnings of Bias in Forensic Mental Health Evaluations” by Tess M. S. Neal and Thomas Grisso (*Psychology, Public Policy, and Law*, 2014, Vol. 20, No. 2, pp. 200 –211. doi:10.1037/a0035824) contained an error that ironically demonstrates the very point of the article: that cognitive biases can easily lead to error— even by people who are highly attuned to and motivated to avoid bias. The authors inadvertently misapplied base rates by failing to account for nested probabilities in the illustration of how Bayesian analysis works in footnote 2 (p. 203). Specifically, the authors said that a person’s belief about the probability of a given attorney being a woman should be 33.3%, given the base rate of women in the legal profession (ABA, 2013). They then tried to demonstrate how Bayes’ Theorem works by incorporating new information about the attorney being a managing partner at one of the 200 largest law firms in the U.S., which has a base rate of 4% women (ABA, 2013). The authors combined these two probabilities (33.3% with 4%) using Bayes’ Theorem to demonstrate how the person’s belief about the probability of the attorney being a woman should have changed, given the new evidence. In doing this, they engaged in a cognitive error similar to the conjunction fallacy described on pages 202–203. They should not have applied Bayes’ Theorem to combine the two probabilities. The new evidence – that the attorney was a managing partner at one of the 200 law firms in the U.S. (with its corresponding base rate of 4% women) – should have replaced, rather than supplemented, the original 33.3% overall base rate information, and thus the “answer” should have been 4% rather than 2.04%. The base rate of women who are large law firm managers is nested within the larger base rate for all women lawyers. Therefore, once the 4% base rate for women as large law firm managers was known, the original base rates for women lawyers generally was no longer relevant. Bayes’ Theorem presumes the conditional independence of the two probabilities in question. In this case, however, the probabilities were not independent. The probability that X is a woman manager of a large law firm logically assumes that X is a woman lawyer.

Reference

American Bar Association. (2013). A current glance at women in the law. Chicago, IL: Author. Retrieved from http://www.americanbar.org/content/dam/aba/marketing/women/current_glance_statistics_feb2013.authcheckdam.pdf