ALFA International
THE GLOBAL LEGAL NETWORK

2026 Transportation Seminar

April 28-30, 2026

A GHOST FROM THE DEPTHS

How Rare Outlier Verdicts Haunt Our Claim Evaluations

Brian T. Coolidge
Moderator
MAYER LLP
Houston, Texas
bcoolidge@mayerllp.com

Julie J. Busch
JOHNSON & BELL
Chicago, Illinois
buschj@jbltd.com

INTRODUCTION: COGNITIVE BIAS IS A PROBLEM

It is the morning after a verdict comes down in St. Louis — a **$462 million** judgment against a trailer manufacturer because an intoxicated driver crashed into the back of a truck.[i] The phones start ringing. Claim files get pulled. Reserve committees convene. Suddenly, every trucking case on the docket looks a little more like that one. Every plaintiff with a back injury is mentally compared to a catastrophic crash victim. Settlement demands that once seemed absurd start to seem merely optimistic.

This is not a failure of diligence. It is a failure of cognition — and it happens to smart, experienced people every single time a headline verdict hits the newswires. The human brain is not wired for base-rate analysis. It is wired for vivid stories. And in the litigation world, a $462 million trucking verdict is about as vivid as it gets.

A central question is whether claims are being evaluated on their merits, or whether we are haunted by the last scary verdict we read. It's worth considering how cognitive biases affect us. Cognitive shortcuts such as the availability heuristic, recency bias, anchoring, and the curse of knowledge, can distort claim evaluations, causing litigators and claims professionals to overweight rare, headline-grabbing outlier verdicts and misprice the overwhelming majority of claims as a result.

The remedy is not to ignore the headline verdicts. The trend toward exceptionally large jury awards is real, increasing, and warrants serious attention. The remedy is to understand exactly why those verdicts happened, what made them outliers, and how to build a disciplined evaluation process that uses that information appropriately, rather than letting it haunt the analysis.

Part II of this paper provides a framework for understanding cognitive bias in the context of legal decision-making. Part III examines the availability heuristic and its distorting effect on claim evaluation. Part IV addresses recency bias, anchoring, and the curse of knowledge as a compounding set of distortions. Part V situates the nuclear verdict phenomenon in statistical context and examines what the actual data shows. Part VI engages the strongest counterarguments. Part VII offers a practical debiasing framework for claims professionals and litigators. Part VIII closes with a synthesis of these points and their practical implications for claim evaluation.

THE ARCHITECTURE OF HUMAN JUDGMENT: COGNITIVE BIAS IN LEGAL DECISION-MAKING

The Brain Is Not a Calculator

As legal and claims professionals, we like to think of ourselves as rational actors. We are trained in logic, evidence, and reasoned argument. We spend careers distinguishing good facts from bad ones. But the psychological research of the past fifty years has established something humbling: under conditions of complexity, uncertainty, and time pressure — conditions that describe virtually every claims file ever opened — the human brain routinely substitutes intuitive shortcuts for careful analysis.

Psychologists Daniel Kahneman and Amos Tversky, whose work earned Kahneman the Nobel Prize in Economics in 2002, described two modes of cognition.[ii] "System 1" thinking is fast, automatic, and intuitive — it runs in the background, pattern-matches against experience, and produces answers that feel right. "System 2" thinking is slow, deliberate, and effortful — it is the kind of thinking that requires concentration. The uncomfortable truth is

that System 1 drives far more of our decisions than we realize, including decisions we believe we are making through careful System 2 analysis.

In the context of claim evaluation, this distinction matters enormously. Claims professionals who review a new trucking file are not operating in a vacuum. They bring every verdict they have heard about, every deposition they have sat through, every catastrophic case they have handled. System 1 is pattern-matching the new file against an entire experiential library. When System 1 finds a recent match — say, the publicity following a $462 million verdict — it flags the new file accordingly. System 2 then has the difficult job of correcting for that flag. And as the research shows, System 2 often fails to fully do so.

What Is a Cognitive Bias?

In psychology and cognitive science, cognitive biases are systematic patterns of deviation from rational judgment.[iii] They arise from the mental shortcuts — heuristics — that the brain uses to navigate complex environments efficiently. Crucially, these shortcuts are not signs of stupidity or laziness. They evolved because they work reasonably well most of the time. The problem is that they fail in predictable, systematic ways.

The biases most relevant to claim evaluation cluster around estimation, decision-making under uncertainty, and memory. They include the availability heuristic (we judge probability based on what comes to mind most easily), recency bias (we weight recent events disproportionately), anchoring bias (we rely too heavily on the first piece of information we encounter), the curse of knowledge (once we know something, we cannot imagine not knowing it), and a supporting cast of related distortions. Each of these will be examined in detail below.

Why Legal Decision-Makers Are Especially Vulnerable

Cognitive biases affect everyone, but the legal and claims environment creates conditions that amplify their impact. Consider the typical claims review: the professional is evaluating a case with genuine uncertainty about outcome, is under time pressure, is managing a caseload of many similar files simultaneously, and is often working from summaries rather than primary materials. These conditions are precisely the ones in which System 1 thinking dominates and System 2 correction is most difficult.

Add to this the adversarial structure of litigation. Plaintiff attorneys are not naive about psychology. "Reptile theory" — a plaintiff litigation strategy explicitly designed to trigger primal, emotional responses in jurors — has been widely discussed in the literature and at seminars like this one for nearly two decades. The demand letter that opens with a three-paragraph description of the most catastrophic possible outcome is not an accident. The mediation presentation that leads with the most recent high-value verdict in the jurisdiction is not coincidental. These are deliberate attempts to exploit the very cognitive shortcuts described in this paper.

The defense side is not immune either. A 2024 survey found that eight in ten corporate counsel reported increasing concern over the growth of exceptionally large jury awards.[iv] That concern is legitimate. But concern, when it is not grounded in accurate probability assessment, can itself become a bias — one that drives systematic overpayment across a portfolio of claims.

THE AVAILABILITY HEURISTIC — WHAT COMES TO MIND COMES TO MARKET

Defining the Heuristic

The availability heuristic is the tendency to overestimate the likelihood of events that are more easily recalled from memory — events that are vivid, recent, unusual, or emotionally charged. It was first formally described by Tversky and Kahneman in 1973 and has been replicated across hundreds of studies in the decades since.[v] The classic demonstration: ask people whether more words in the English language begin with the letter 'K' or have 'K' as their third letter. Most people answer that words beginning with K are more common, because they are much easier to recall. In fact, words with K as the third letter are far more numerous — but they are harder to call to mind.

The implications for legal decision-making are direct. When a claims professionals ask "what are the chances this trucking case results in a large verdict?", they are not performing a statistical analysis of a jurisdiction's verdict history. They are, to a significant degree, asking "how easily can I recall an example of a large verdict in a case like this one?" The more easily such an example can be recalled, whether because it was recent, unusually large, or widely covered in industry publications and legal reporting, the more likely a claims professional is to overestimate its probability.

How Outlier Verdicts Become "Available"

The media and publication ecosystem surrounding nuclear verdicts is uniquely designed to maximize availability. Consider what happens when a jury returns a $100 million verdict against a trucking company. Within hours, it appears on legal news services. Within days, it is the subject of client alerts, risk management bulletins, and insurance industry publications. Within weeks, it has been discussed at claims conferences, referenced in mediator opening statements, and cited in demand letters across the country.

Marathon Strategies' 2024 nuclear verdict study found that attorney advertising — which now exceeds $2.4 billion annually — aggressively markets the largest verdicts as evidence of systemic risk.[vi] Plaintiff attorneys use these verdicts both to attract new clients and to anchor the expectations of defense counsel and claims professionals in pending matters. The verdict becomes a tool precisely because of its availability-enhancing properties.

The mechanism can be described succinctly: *outlier jury verdicts create strong availability bias, causing the industry to confuse what is most probable with what is most memorable*. That observation bears emphasis because it captures the core problem addressed here. Probability and memorability are not the same. But under the availability heuristic, the brain treats them as if they are.

What the Data Actually Shows

Let's put the nuclear verdict phenomenon in statistical context. A recent analysis of verdict data found 135 lawsuits against corporate defendants resulted in verdicts with awards exceeding $10 million in 2024.[vii] This represents a 52% increase over 2023, and it generated appropriately significant attention.

But let's think about what 135 nuclear verdicts means in context. The National Center for State Courts tracks tens of thousands of tort trials annually across the country, with the understanding that actual filings number in the millions. Of the cases that are filed, the vast majority settle before trial. Of those that go to trial, defense wins are

common. Motor vehicle and trucking cases historically have defense win rates often exceeding 50% in many jurisdictions. And of the plaintiff wins, only a small fraction would be considered in the "nuclear range".

In the trucking sector specifically, eight of those 135 verdicts totaled approximately $790 million.[viii] Across an industry that logs billions of miles annually, that handles hundreds of thousands of accident claims each year, eight nuclear verdicts is small. The trucking sector in 2024 had 8 cases reach nuclear verdict status out of a universe of claims that dwarfs that number by several orders of magnitude.

The top ten trucking and transportation verdicts from August 2024 through February 2026 ranged from $38 million to $462 million.[ix] These are genuinely extraordinary numbers. But they are, by definition, the top ten. They are not representative. They are the outliers. The availability heuristic causes us to reason from these outliers as if they were representative — as if the next trucking case has a meaningful probability of landing in this range. For most cases, that probability is vanishingly small.

The Problem in Practice: How Availability Distorts Evaluation

Here is how the distortion plays out in a real claims context. A plaintiff's counsel files a trucking negligence case involving a rear-end collision. The plaintiff sustained a herniated disc requiring surgery and is claiming six figures in medical expenses and lost wages. Liability is reasonably clear. The case has moderate value, perhaps somewhere in the low to mid-six figures based on comparable verdicts in the jurisdiction.

But the claims professional handling the file just read about the $462 million Wabash verdict. Her System 1 is alert: this is also a trucking case; it also involves a rear-end collision; it also has a seriously injured plaintiff. Her System 2 knows that the cases are dramatically different — one involved a defective product, punitive damages, and catastrophic quadriplegic injuries; the other involves a routine negligence claim. But System 2's corrections are incomplete. The availability of the Wabash verdict has inflated her estimate of what this case "could" be, and that inflation drives up the settlement authority she requests, the reserve she sets, and the value she communicates to defense counsel.

Plaintiff's counsel, of course, knows exactly what she is doing. The demand letter leads with the nuclear verdict. The mediation presentation features a slide deck of the top five trucking verdicts in the state. The anchor is set before the first substantive conversation about the case begins.

The result is systematic overpayment, not on every case, but across the portfolio. And because no one gets a phone call when a case settles for less than it was reserved, the feedback loop that might correct the overvaluation never fires. The ghost stays in the room.

Recency Bias, Anchoring, and the Curse of Knowledge — A Triple Threat

Recency Bias: The Last Verdict You Read Is the One You Remember

Recency bias is the tendency to weight recent events more heavily than older ones in making predictions and judgments.[x] It is related to — but distinct from — the availability heuristic. While the availability heuristic broadly concerns ease of recall, recency bias specifically privileges temporally recent information over older information of equal or greater relevance. Together, they create a feedback loop: recent verdicts are both more available (easier to recall) and more heavily weighted (because they are recent), making them doubly distorting.

The practical implications for claims evaluation are significant. The 2024 verdict year was by virtually every metric a historic outlier — 135 nuclear verdicts, $31.3 billion in total value, a 52% increase over 2023. These numbers are genuinely alarming. But they reflect conditions in a specific year, driven in part by extraordinary cases (three verdicts against a single beverage company totaling over $8 billion; a $4.7 billion verdict against the NFL that was subsequently overturned) that artificially inflated the totals.

Claims professionals who set reserves based on the 2024 nuclear verdict climate are not making neutral statistical judgments. They are anchoring the evaluation to an extraordinary recent year and projecting those conditions forward. This is a form of recency bias operating at the portfolio level. The actual base rate of nuclear verdicts, while genuinely rising, is nowhere near what a 2024-only snapshot suggests.[xi]

Moreover, the recency effect is asymmetric. A claims professional who just handled a catastrophic trucking case will carry that experience into subsequent evaluations more heavily than a string of routine resolutions. Defense wins do not generate the same cognitive imprint as plaintiff wins. The bad outcomes are stickier, which means the evaluation baseline drifts upward over time through accumulation of availability-heavy memories.

Anchoring Bias: The First Number in the Room

Anchoring bias, sometimes called focalism, is the tendency to rely too heavily on the first piece of information encountered when making quantitative estimates.[xii] In classic anchoring experiments, participants asked to estimate an unknown quantity after being shown a random number (even one they know to be random) will give estimates that are meaningfully influenced by that number. The anchor need not be relevant to be powerful.

In litigation, anchors are often carefully crafted and rarely random. The plaintiff attorney's demand letter is an anchor. The opening number at mediation is an anchor. And increasingly, the verdict from a different jurisdiction that opened the demand letter is an anchor.

Research shows that anchoring can distort valuations by 20-30% or more, even among experts and even when subjects are aware of the anchoring effect.[xiii] Awareness of a bias does not neutralize it. This is the uncomfortable truth about most cognitive biases: knowing they exist is a necessary but insufficient condition for overcoming them. A claims professional who is fully aware that a plaintiff's attorney just anchored with a $300 million trucking verdict is still affected by that anchor, even as he or she consciously discounts it.

The anchoring problem is compounded in the trucking and transportation context by the behavior of plaintiff attorneys who have become highly sophisticated in its use. The practice of beginning demand letters, mediation briefs, and opening statements with the most alarming comparable verdict in the jurisdiction — or sometimes just the country — is now standard plaintiff strategy. It is not coincidence that a demand letter for a moderate-value trucking case opens with two pages about the $82 million Ohio trucking verdict or the $44 million Dallas pileup case. The anchor is being set, deliberately, before the first substantive communication about the actual case.

The Curse of Knowledge: You Can't Unknow What You Know

The curse of knowledge is a cognitive bias that occurs when, in trying to communicate with others, we unknowingly assume that they have the background knowledge necessary to understand.[xiv] More broadly, it refers to the difficulty, once information is known, of accurately imagining how matters appear to those who do not share that knowledge. The term was introduced by economists Colin Camerer, George Loewenstein, and Martin Weber in a 1989 study that demonstrated the phenomenon in information-asymmetric trading contexts.

In claims evaluation, the curse of knowledge operates in several distinct and damaging ways. First, consider the claims professional who has spent weeks immersed in the facts of a bad case — reviewing accident reconstruction reports, medical records, witness statements, and deposition testimony. They know that the driver was fatigued, that the company's safety program had documented deficiencies, and that the plaintiff's medical records reflect a catastrophic spinal injury with a compelling prognosis. That knowledge cannot be set aside.

When evaluating how a jury would respond, the analysis is conducted from the perspective of someone who has absorbed the equivalent of weeks of briefing. The task, however, is to estimate how twelve jurors, drawn from the general population and exposed to a curated presentation of the evidence over a limited trial period, will process the case. Any estimate of how those jurors will think and respond is necessarily shaped, and often distorted, by the depth of information already known.

This produces a specific and systematic error: the curse of knowledge leads experienced claims professionals and litigators to overestimate how much juries will understand and care about the most damaging facts. The experienced litigator who has spent months preparing a case has internalized its complexity and understands which facts are truly consequential. A jury encountering the evidence for the first time, however, may not connect those points in the same way. The curse of knowledge can inflate estimates of a plaintiff's likelihood of success in the most challenging cases and, paradoxically, may also lead to underestimation in simpler cases where the underlying narrative is less compelling.

Second, the curse of knowledge affects how evaluations are communicated to clients and supervisors. A claims handler who fully understands why a case presents significant risk may struggle to convey that risk to someone who has not reviewed the file, because it is difficult to identify and articulate the gaps in knowledge that would make the risk apparent to a fresh set of eyes. This asymmetry can lead to both under-reserving, when the handler knows the case is defensible but cannot explain that position as clearly as the risks, and over-reserving, when a vivid, damaging narrative is easier to internalize than to communicate in a balanced way to a reviewing supervisor.

THE REAL VERDICT LANDSCAPE — PUTTING OUTLIERS IN STATISTICAL CONTEXT

Understanding the Nuclear Verdict Trend

Let's be clear about what the nuclear verdict data actually shows, because conflating the trend line with the baseline risk on any individual case is itself a form of statistical error.[xv]

A recent report documents a genuine and significant trend. In 2024, 135 corporate verdicts exceeding $10 million were identified, totaling $31.3 billion, representing a 52% increase in number and a 116% increase in total value over 2023. Verdicts exceeding $100 million increased from 27 to 49, and the median award rose to $51 million. These are not aberrations confined to a single year; they reflect a pattern that has been building since courts resumed normal operations following pandemic-related closures and case backlogs began to move.

The underlying forces driving that increase are real and warrant acknowledgment. Corporate mistrust, social pessimism, erosion of tort reform, and public desensitization to large numbers are contributing factors. Jury pool demographics are also shifting, with Millennial and Gen Z jurors, who tend to be more pro-plaintiff and more skeptical of corporate defendants, comprising a growing share of the venire. Attorney advertising, now exceeding $2.4 billion annually, is designed both to recruit plaintiffs and to shape juror perceptions. These are structural developments, not isolated anomalies.

The critical analytical point, however, is that the trend toward exceptionally large jury awards reflects what is occurring at the tail of the distribution. It captures the cases that become headline news. It does not describe, and is often misapplied to, the broad middle of the claims distribution where most cases actually reside.

The Trucking Sector: Eight Cases in Context

The trucking sector is particularly instructive because it is one of the most discussed and feared nuclear verdict environments.[xvi] A recent compilation identified eight trucking verdicts exceeding $10 million in 2024, totaling approximately $790 million. That is genuinely significant. But let's think about what those eight cases looked like.

The two largest trucking-related verdicts in recent memory share a striking characteristic: they involved defective product claims with allegations of concealed knowledge spanning decades. The $462 million Wabash verdict involved allegations that Wabash failed to meet government standards for rear underride bumpers, with the trailer's design meeting only the minimum federal requirements of 2004 — more than twenty years before the verdict. The $160 million Daimler/Western Star verdict involved allegations that the defendants failed to increase cab strength for nearly thirty years despite documented knowledge of rollover dangers. These were not ordinary negligence cases. They were cases with massive punitive damage components rooted in allegations of deliberate indifference to known dangers over long periods of time.

The top ten trucking and transportation verdicts from the last eighteen months — ranging from $38 million to $462 million — all share specific characteristics that distinguish them from the routine claims portfolio.[xvii] They involve catastrophic injuries (death, quadriplegia, traumatic brain injury), clear or admitted liability in most cases, significant aggravating conduct (intoxicated driver who was not properly screened, defective design known to the manufacturer, failure of basic safety protocols), or some combination of the three. They are not simply "bad" trucking cases. They are worst-case trucking cases.

The Post-Verdict Reality

A critical and frequently overlooked point in discussions of nuclear verdicts is that the reported verdict number is not the collected amount. The news cycle captures the jury's verdict. It rarely follows the case through post-trial motions, remittitur, appeal, and ultimate resolution.

Consider several examples from the record. The $2.25 billion Roundup verdict against Bayer was "later slashed by a judge to $400 million."[xviii] The $4.7 billion NFL Sunday Ticket verdict was thrown out by the trial judge.[xix] The Wabash National verdict's reported "post-verdict reduction" is noted in the very source that reported it.[xx] The $847 million Verizon patent verdict was overturned and then settled for an undisclosed amount.[xxi]

This pattern is consistent with the academic literature on exceptionally large jury awards. Such verdicts are reduced at materially higher rates than ordinary awards, through both remittitur and appellate review. The headline figure, the one that becomes most readily available in memory, is the pre-reduction number. The amount ultimately collected is lower, often substantially so.

This creates a systematic availability distortion in a specific direction: the numbers we remember are bigger than the numbers that were actually paid. Claims professionals setting reserves based on "comparable verdicts" are, without realizing it, often working from a dataset that overstates collection amounts.

Geographic Concentration: Location Changes Everything

Another critical contextual factor is geography. The 2024 nuclear verdict data was highly concentrated in specific jurisdictions. The top five courts by verdict total — Clark County (Nevada), the Central District of California, Philadelphia County (Pennsylvania), New York County (New York), and the Eastern District of Texas — accounted for a disproportionate share of the total $31.3 billion.[xxii] The Eastern District of Texas alone generated six massive verdicts totaling over $1.8 billion, driven largely by patent litigation.

For trucking cases specifically, the risk of exceptionally large jury awards is heavily concentrated in plaintiff-friendly state court venues. St. Louis, Philadelphia, and certain Texas district courts have appeared repeatedly in reported high-value verdict data. This is not because trucking cases are inherently more dangerous in those locations. It reflects forum selection where rules permit, the development of venue-specific track records that attract aggressive plaintiff counsel, and differences in jury pool characteristics relative to the national baseline.

A trucking case in rural Wyoming does not carry the same risk profile as a trucking case in St. Louis County. Treating them identically because both involve trucking is a categorical error that availability bias and anchoring can reinforce. The headline verdict from a plaintiff-friendly venue can influence evaluation of a case in a materially different forum even when the underlying conditions that drove the larger award are not present.

COUNTERARGUMENTS

Large Verdicts Reflect Real Risk and Ignoring Them Is Dangerous

The strongest argument on the other side is that exceptionally large jury awards are not a distortion to be corrected, but a signal to be heeded. That argument has genuine force.

The trend is real. It is not merely a statistical artifact or the product of media amplification. The year-over-year increase in nuclear awards, the rising median amounts, and their expansion across jurisdictions reflect actual jury behavior in courtrooms. Millennial and Gen Z jurors tend to be more receptive to plaintiff arguments. Anti-corporate sentiment has increased. Attorney advertising continues to shape juror perceptions. Social inflation is driving liability costs at rates that exceed general economic inflation.

Empirical research supports this view. The Swiss Re Institute reported that social inflation in U.S. liability claims rose by 5.4% annually from 2017 to 2022, compared to 3.7% general economic inflation over the same period, and characterized the trend as distinctly pronounced in the United States.[xxiii] The Orrick jury research study found that anti-corporate sentiment has nearly doubled among jurors and that today's jurors are more likely to act on those feelings by delivering their own sense of justice.[xxiv]

From this perspective, a claims professional who accounts for the risk of exceptionally large jury awards is not acting irrationally but responding to a measurable trend. A defense lawyer who believes the facts strongly favor the defense may still underestimate how a contemporary jury will interpret those facts. Ignoring this environment is not analytical rigor; it is a failure to account for material changes in jury behavior.

Response

Acknowledging the trend toward exceptionally large jury awards and applying it accurately are two different things, and the distinction matters. The point is not to disregard the trend, but to apply it with precision, focusing on cases that share the risk profile associated with such outcomes rather than extending it across an entire sector.

A 52% year-over-year increase in nuclear awards means that the probability, starting from a low baseline, has increased by 52%. It does not mean that any particular case carries a 52% likelihood of such an outcome. The base rate matters. The characteristics of the specific case matter. The jurisdiction matters. The conduct of the defendant matters. The severity of the injury matters. Applying the trend without these contextual filters is the error, not recognition of the trend itself.

Put differently: yes, the ghost is real. But that does not mean every house is haunted.

Conservative Evaluations Are Prudent Risk Management

A related counterargument comes from within the risk management community itself: that conservative claim evaluation is simply good risk management, and that systematic debiasing toward lower valuations exposes clients to catastrophic outcomes. No one, the argument goes, ever got fired for settling a case. But they do get fired for taking cases to trial and losing badly.

This argument has real operational force. Risk managers are responsible for managing portfolios of risk, not just individual claims. From that perspective, carrying higher reserves can be prudent because it accounts for uncertainty. If the tail risk is $50 million, reserving at $2 million based solely on a most-likely defense outcome does not reflect disciplined risk management. It reflects optimism rather than analysis. The asymmetry of outcomes supports some degree of conservatism. A $10 million overpayment in settlement may be absorbed within the portfolio, while a $100 million adverse verdict may not be absorbed.

Response

This argument conflates accurate risk assessment with conservative risk assessment, and the distinction is consequential. Effective risk management depends on accurate probability estimates, not inflated ones. An insurer that systematically overpays claims by 30% across a portfolio of 10,000 trucking claims annually is committing hundreds of millions of dollars above defensible valuations. Those costs are ultimately borne through premiums, meaning transportation companies subsidize an evaluation process influenced by availability bias.

The asymmetry argument is also often overstated. When a case genuinely presents meaningful tail risk based on its specific facts, the appropriate response is not a premium-driven settlement untethered from the record. It is a disciplined, case-specific assessment grounded in the actual distribution of outcomes. That analysis may support a higher settlement than the defense would prefer, but it should be driven by the file itself, not by recent, high-profile verdicts.

A consistently conservative evaluation culture also produces broader systemic effects. When claims professionals and defense counsel regularly assess cases above their reasonable range, plaintiff attorneys adjust accordingly. High anchors become more effective. Opening demands increase. The reference point for valuation shifts upward. Over time, this feedback loop reinforces itself. Inflated evaluations beget inflated demands, which beget further inflated evaluations. Breaking that loop requires accurate evaluation, not conservative evaluation.

Restoring Rational Evaluation: A Practical Framework for Debiasing

Knowing about cognitive biases is not sufficient to correct for them. Research consistently shows that awareness of bias reduces its impact, but it does not eliminate it. Instead a systematic process is required, a set of habits and practices that build in the analytical safeguards that unaided intuition fails to provide. The following framework is offered not as a complete answer, but as a practical starting point for claims professionals and litigators who want to evaluate their cases more accurately.

Impose a Reference Class Before a Case-Specific Evaluation

Before turning to the specific facts of a new file, and certainly before reviewing a demand, establish how comparable cases resolve. The relevant reference class consists of matters with similar characteristics: same jurisdiction, same general claim type, comparable injury severity, and a similar liability profile. What do those cases typically resolve for? What percentage proceed to trial? What percentage result in plaintiff verdicts? What does the distribution of outcomes look like?

This is what behavioral economists call “reference class forecasting,” a method developed to counteract optimism and availability bias in prediction. It begins with the outside view, grounding the analysis in how comparable cases have performed, before turning to the particulars of the file. The specific facts then adjust that baseline up or down, rather than establishing the anchor from the outset.

Use Structured Probability Estimation

Rather than asking "what is this case worth?", ask a series of more specific questions whose answers can be combined to produce a more disciplined estimate. For example:

- What is the probability this case goes to trial? (Most cases do not.)
- If it does go to trial, what is the probability of a plaintiff verdict? (Varies by jurisdiction and facts, but defense wins regularly.)
- If the plaintiff wins, what is the distribution of verdict amounts? (Consider percentiles: what is the 25th percentile outcome? The median? The 75th? The 90th? The 99th?)
- What is the probability of a nuclear-range verdict specifically? (Very low in most cases, even in plaintiff-friendly jurisdictions, absent specific aggravating factors.)

Multiplying these probability estimates produces a more disciplined expected value calculation than intuition alone allows. It also requires explicit engagement with the base-rate question: how likely is an exceptionally large jury award in cases like this one, as opposed to this particular case? That is precisely where availability bias most distorts the analysis.

Separate the Evaluation from the Latest Headline

Implement a simple procedural rule: when beginning a significant claim evaluation, do not start by reviewing recent verdict reports from the jurisdiction. Start by pulling the base-rate data — the distribution of outcomes in comparable cases over a meaningful time period. Let that establish the analytical foundation. Then review the recent high-profile verdicts as contextual data that may modify the baseline, rather than as the starting point that establishes it.

This is not about ignoring large verdicts. It is about sequencing the analysis so that the vivid, recent cases are weighed against a proper statistical backdrop rather than setting the frame before any statistical context exists.

Run the Pre-Mortem Analysis

A pre-mortem analysis asks decision-makers to imagine that the decision has already been made and has failed spectacularly, and then to explain what happened. In the claims context, the exercise is straightforward. Imagine it is two years from now, the case proceeded to trial, and the plaintiff obtained a $75 million verdict. What happened? Which facts drove the outcome? Which decisions in the evaluation process contributed to it?

The value of the pre-mortem lies in its ability to counter overconfidence that develops with familiarity. As a file becomes more familiar, confidence in the evaluation tends to increase, and the depth of knowledge can create an illusion of control. The pre-mortem requires articulation of the specific scenario in which the evaluation proves materially wrong, which is a more rigorous exercise than simply assessing confidence levels.

Equally important is to conduct the inverse analysis. Assume it is two years later, the case was settled for $3 million based on an inflated evaluation influenced by high-profile verdicts, and the actual trial would have resulted in a defense verdict. What drove that outcome? Which availability effects contributed to the overpayment? This forces the same discipline on the high side of the distribution as the pre-mortem forces on the low side.

Calibrate for Jurisdiction-Specific Data

Availability bias is particularly distorting when a verdict from a different jurisdiction dominates the mental landscape. A $462 million St. Louis verdict should not meaningfully affect the evaluation of a trucking case in Houston, Texas — different jury pool, different tort law, different local legal culture. Yet without deliberate calibration, it will.

The practical corrective is to maintain jurisdiction-specific verdict databases and to require that comparable verdict analysis in any significant evaluation be drawn primarily from the jurisdiction where the case will be tried. National verdict trends are relevant context. They should not be the primary comparables for a case that will be tried in a jurisdiction with a very different verdict history.

Build in Structured Second Opinions

One of the most effective debiasing techniques is structured dissent: deliberately incorporating a perspective that has not been exposed to the primary evaluation and tasking that reviewer with identifying reasons the current assessment may be too high. This approach counteracts confirmation bias, which drives the search for information that supports existing conclusions, and the curse of knowledge, which obscures how familiarity with the file can itself distort judgment.

This need not be a time-intensive process on routine matters. But for significant reserves and meaningful settlement authority requests, a second opinion from someone who has not previously reviewed the file, and who is specifically asked to play devil's advocate on the high side, is an inexpensive safeguard against systematic overvaluation.

CONCLUSION

The ghost of the last large verdict you read about is in the room every time you pick up a new file. It does not announce itself. It does not distinguish between the case with the defective product, the egregious conduct, and the catastrophically injured plaintiff — and the case with clear but unremarkable liability and a manageable injury. It treats every trucking case as a potential $462 million Wabash, every transportation claim as a potential $329 million Tesla. And because the brain is not naturally equipped to perform base-rate analysis under time pressure and uncertainty, the ghost is surprisingly hard to exorcise.

The distorting effect of that ghost, which operates through the availability heuristic, recency bias, anchoring, and the curse of knowledge, is one of the most significant and underappreciated sources of error in claim evaluation. National verdict trends are real, important, and deserve serious attention. But attention to the trends is not the same as accurate evaluation of individual cases. The tendency to conflate what is most memorable with what is most probable is the core cognitive error, and it operates in both directions: toward overvaluation of cases that happen to resemble a recent nuclear verdict, and toward undervaluation of cases that do not.

The practical tools discussed in Part VII — reference class forecasting, structured probability estimation, jurisdiction-specific calibration, pre-mortem analysis, and structured second opinions — are not complex. None of them require additional budget or technology. They require commitment to slowing down the evaluation process at key moments, inserting the kind of deliberate System 2 analysis that System 1 is inclined to bypass, and building into the institutional claims handling process the safeguards that unaided individual judgment consistently fails to provide.

Good evaluation is not about ignoring risk. It is about measuring it accurately. The landscape of exceptionally large jury awards warrants serious attention in cases that present the characteristics associated with those outcomes. In the remainder of the portfolio, which accounts for the vast majority of claims, the risk should be recognized, understood, and then confined to its proper place in the analysis.

The most important question you can ask about your next file is not: "what is the worst verdict I have ever heard of in a case like this?" It is: "given the actual facts of this case, the actual jurisdiction, and the actual distribution of outcomes in comparable matters, what is the probable distribution of outcomes here, and what does that distribution tell me about how to evaluate and resolve this claim?"

Answer that question accurately, and the ghost loses its power.

[i] Perkins, et al. v. Wabash National Corp., Circuit Court, City of St. Louis (22nd Judicial Circuit), St. Louis, Missouri (Sept. 2024). The jury award of $462,000,000 included $450,000,000 in punitive damages. Post-verdict reduction has been reported.
[ii] Amos Tversky & Daniel Kahneman, Availability: A Heuristic for Judging Frequency and Probability, 5 COGNITIVE PSYCHOL. 207 (1973); DANIEL KAHNEMAN, THINKING, FAST AND SLOW (2011).
[iii] For a great list of these biases, see, List of Cognitive Biases, WIKIPEDIA (last visited Mar. 17, 2026), https://en.wikipedia.org/wiki/List_of_cognitive_biases. See also Marlene Dumas et al., Biases in Human Decision Making, in PERSPECTIVES ON PSYCHOLOGICAL SCIENCE (2020).
[iv] Annual Litigation Trends Survey, Norton Rose Fulbright (Jan. 2025).
[v] Tversky & Kahneman, supra note 2, at 207–08.
[vi] MARATHON STRATEGIES, CORPORATE VERDICTS GO THERMONUCLEAR: 2025 EDITION 10 (2025) [hereinafter MARATHON REPORT].
[vii] *Id.* at 2.
[viii] *Id.* at 19.
[ix] *Hilary Remick, Magna Legal Services, Confidential Client Memorandum: Top Jury Verdicts — Trucking/Transportation (Last 18 Months) (Feb. 24, 2026) (on file with authors)*.
[x] See generally RICHARD H. THALER & CASS R. SUNSTEIN, NUDGE: IMPROVING DECISIONS ABOUT HEALTH, WEALTH, AND HAPPINESS 36–37 (2008) (discussing recency and availability effects).
[xi] MARATHON REPORT, supra note 6, at 5.
[xii] List of Cognitive Biases, supra note 3 (anchoring bias entry).
[xiii] DAN ARIELY, PREDICTABLY IRRATIONAL: THE HIDDEN FORCES THAT SHAPE OUR DECISIONS 26–31 (2008).
[xiv] Colin Camerer, George Loewenstein & Martin Weber, The Curse of Knowledge in Economic Settings: An Experimental Analysis, 97 J. POL. ECON. 1232 (1989).
[xv] MARATHON REPORT, supra note 6, at 2, 5.
[xvi] Id. at 15.
[xvii] Remick Memorandum, supra note 9.
[xviii] MARATHON REPORT, supra note 6, at 12.
[xix] Id. at 11.
[xx] Remick Memorandum, supra note 9 (noting 'post-verdict reduction reported' for the Wabash verdict).
[xxi] MARATHON REPORT, supra note 6, at 14 (noting that the Verizon $847 million verdict 'was later thrown out, and in December, the parties agreed to an undisclosed settlement').
[xxii] Id. at 23.
[xxiii] Swiss Re Institute, Social Inflation: Litigation Costs Drive Claims Inflation (Sept. 7, 2024), https://www.swissre.com/institute/research/sigma-research/sigma-2024-04-social-inflation.html.
[xxiv] MARATHON REPORT, supra note 6, at 9 (citing Orrick, Juror Attitudes in a Polarized Society).