The International Journal of Indian Psychology
ISSN 2348-5396 (Online) | ISSN: 2349-3429 (Print)
Volume 14, Issue 3, July- September, 2026
DIP: 18.01.056.20261403, DOI: 10.25215/1403.056
https://www.ijip.in

Review Paper

Cognitive Biases and Decision Noise in Hiring Decisions: A Literature Review

Ashna Sehgal[1], Rajisha Sharma[1]*

ABSTRACT

Hiring decisions determine who enters an organization, yet the process by which recruiters and managers evaluate candidates is rarely as objective as organizations intend it to be. This literature review synthesizes research on the cognitive biases and decision noise that distort hiring judgments, drawing on foundational work in judgment and decision-making (Tversky & Kahneman, 1974; Kahneman, Sibony, & Sunstein, 2021) as well as more recent empirical and organizational research. The review defines cognitive bias and decision noise, surveys the most consequential biases identified in the selection literature—including the halo and horn effects, confirmation bias, similarity or affinity bias, anchoring, availability, representativeness, stereotyping, and status-quo bias—and examines the theoretical frameworks used to explain them, namely Dual Process Theory, the Heuristics and Biases framework, Bounded Rationality, and Noise theory. It then considers empirical evidence of bias in real hiring contexts, including field experiments on racial discrimination in callbacks (Bertrand & Mullainathan, 2004) and cultural matching in elite firms (Rivera, 2012), before turning to the promises and risks of algorithmic hiring tools. The review concludes by summarizing evidence-based mitigation strategies—structured interviews, decision hygiene, blind screening, and governed use of artificial intelligence—and argues that addressing both bias and noise is essential to fair and effective personnel selection.

Keywords: *Cognitive Biases, Decision Noise, Hiring Decisions, Review*

Hiring is one of the most consequential decisions an organization makes. Who joins a team shapes productivity, culture, turnover, and, over time, competitiveness itself; human resource managers are sometimes described as second only to the chief executive in the scale of their influence over organizational outcomes (Salehzadeh & Ziaeian, 2024). The premise behind most selection systems is that hiring rests on a fairly straightforward match between what a candidate can do and what a role requires. Decades of research in psychology, behavioral economics, and human resource management complicate that premise considerably. Two distinct sources of error creep into hiring judgments: cognitive bias, the systematic and directional distortion that arises when people lean on mental shortcuts rather than exhaustive analysis (Tversky & Kahneman, 1974), and decision noise, the messier problem of unwanted variability—different evaluators, or the same

[1]University of Delhi
*Corresponding Author

Received: July 12, 2026; Revision Received: July 26, 2026; Accepted: July 31, 2026

evaluator on different days, reaching different verdicts about an identical candidate for reasons that have nothing to do with merit (Kahneman, Sibony, & Sunstein, 2021).

This review draws together classic judgment-and-decision-making research with newer work on algorithmic hiring, tracing how bias and noise show up across the stages of recruitment and selection—resume screening, interviewing, reference checking, and the final offer. It begins by defining cognitive bias and surveying the categories of bias most often documented in the selection literature, then turns to decision noise as a related but distinct source of error. From there it works through four theoretical lenses—Dual Process Theory, the Heuristics and Biases framework, Bounded Rationality, and Noise theory—that help explain why these errors keep recurring even among trained, well-intentioned evaluators. The later sections look at empirical evidence from real hiring settings, including the rapidly expanding use of artificial intelligence in recruitment, before closing with a review of the strategies—structured interviews, decision hygiene, blind screening, governed AI use—that the literature suggests actually reduce bias and noise, rather than simply raising awareness of them.

The Meaning of Cognitive Bias in Hiring

Cognitive bias, at its simplest, is a systematic error in thinking that pushes judgment predictably in one direction rather than another (Tversky & Kahneman, 1974). Evaluators rarely judge candidates purely on demonstrated skills and qualifications; personal impressions, prior beliefs, and mental shortcuts creep in largely outside conscious awareness (Thomas & Reimann, 2022). Interviews and other selection methods are, by their nature, exercises in human judgment, so they are inherently subjective—different interviewers can sit across from the same candidate and walk away with markedly different conclusions (Schmitt & Kim, 2007). It is this subjectivity that leaves the door open for bias in the first place. And the door tends to open wider under pressure: bias grows more likely as decisions become more complex and time grows shorter (Thomas & Reimann, 2022), which describes most real-world hiring, where recruiters work through large candidate pools against tight deadlines.

None of this requires bad intentions. Nickerson (1998) drew a useful line between deliberate case-building, where someone knowingly seeks evidence for a conclusion they already favor, and unwitting selectivity, where information processing gets skewed without any awareness that it's happening. Most workplace bias sits in the second category. Recruiters tend to genuinely believe they are being fair and objective, even as the shortcuts they aren't aware of are already shaping the outcome (Berthet, 2022).

Types of Cognitive Biases in Hiring

Halo and Horn Effects

A single positive trait can quietly recolour everything else an evaluator notices about a candidate—this is the halo effect (Nisbett & Wilson, 1977). Nisbett and Wilson's original demonstration remains one of the clearest illustrations of it: undergraduates watched one of two videotaped interviews with the same instructor, who came across as either warm or cold, and those who saw the warm version rated his appearance, mannerisms, and even his accent more favourably than those who saw the cold version—despite watching an identical person. Translated into a hiring context, a candidate who speaks confidently, or who lists a prestigious university, is often assumed competent and hardworking well before there is any direct evidence for either (Dessler, 2020, as cited in the assignment literature). The horn effect is the mirror image: one negative cue, arriving late or seeming nervous, can drag down an evaluator's reading of everything else the candidate did (Harvard University, as cited in student coursework, 2023). What both effects share is a kind of evaluative shortcut—a part standing in for the whole.

Confirmation Bias

Interviewers rarely walk into a conversation as blank slates. Confirmation bias—the tendency to seek out, interpret, and remember information that confirms what one already believes while discounting what doesn't (Nickerson, 1998)—often starts forming within the first few minutes of an interview, sometimes before a substantive question has even been asked. Nickerson (1998) called it one of the most pervasive and consequential reasoning errors in psychology, contributing to belief perseverance long after contradicting evidence has appeared. In an interview, the mechanics are fairly predictable: the interviewer forms an early hypothesis about a candidate, then unconsciously steers the rest of the conversation toward confirming it, discounting or reinterpreting whatever answers would complicate the story. What ends up scored is less a measure of the candidate than a record of the interviewer's first few minutes of instinct.

Similarity (Affinity) Bias and Cultural Matching

People tend to like people who remind them of themselves, and hiring is no exception. Similarity or affinity bias captures interviewers' tendency to favour candidates who share their background, education, interests, or way of speaking (Thomas & Reimann, 2022). Rivera (2012) documented this at scale in a study of hiring at elite investment banks, law firms, and consulting firms, drawing on 120 interviews with employers and direct observation of a hiring committee. Employers in that study routinely prioritized "cultural fit"—shared leisure pursuits, self-presentation styles, social background—over pure productivity signals, in effect treating hiring as an exercise in cultural matching rather than skills sorting. The short-term payoff is interpersonal comfort; the longer-term cost is a narrower organization, and qualified candidates overlooked simply for not looking or sounding like the people already in the room.

First Impression Bias

First impressions form within seconds, long before an interview has properly begun, and they carry disproportionate weight against everything that follows over the next thirty minutes or more. A candidate's posture, vocal confidence, or general polish often shapes an interviewer's overall judgment more than the actual content of their answers does—even though these surface cues are, on their own, weak predictors of how someone will perform on the job.

Stereotyping, In-Group Bias, and Status-Quo Bias

Stereotyping bias enters when a recruiter judges a candidate against assumed traits of a demographic group—gender, age, race, nationality—rather than the individual sitting in front of them (Thomas & Reimann, 2022). A close relative, in-group bias, shows up as interviewers rating candidates who share their own age, gender, ethnicity, or educational background more favourably (Thomas & Reimann, 2022). Status-quo bias then compounds both: because familiar options feel safer, hiring panels gravitate toward candidates who resemble people already in similar roles, quietly reproducing the existing workforce rather than assessing each new applicant on their own terms (Wharton Executive Education, 2025). Samuelson and Zeckhauser's original account of the phenomenon traced this preference for the familiar to loss aversion and discomfort with change rather than to any rational weighing of risk—which is part of why hiring panels resist new evaluation criteria even when the evidence says they would help.

Anchoring Bias

An early number or impression can lodge itself in an evaluator's mind and barely budge afterward, regardless of what comes next—this is anchoring bias (Tversky & Kahneman, 1974). A candidate's stated salary expectation, an unusually strong or weak opening answer, or an early interview score can all function as anchors, pulling the final rating or offer toward them, because evaluators tend to adjust only slightly from a reference point rather than reassessing the evidence from scratch.

Availability and Representativeness Heuristics

Two related heuristics distort how likely or typical something seems. The availability heuristic leads people to judge how common or risky something is by how easily an example comes to mind (Tversky & Kahneman, 1974); a single memorable negative anecdote about a candidate from an unfamiliar background can make a manager overestimate the risk that background carries, even when the anecdote is not representative of anything broader. The representativeness heuristic works a bit differently: evaluators size up a candidate against a mental prototype of the "ideal employee" and judge fit by resemblance rather than by the full evidence in front of them, which means candidates bringing different but equally valuable experience can be undervalued simply for not matching the mold.

Contrast Effect and Order Effects

Judgment can also depend on what came just before it. The contrast effect means a candidate gets judged against whoever preceded them rather than on independent merit—an average candidate can look strong right after a weak one, or weak right after an outstanding one. Order effects more broadly go back at least to Farr (1973), who found recency effects—later information carrying disproportionate weight—among 140 personnel interviewers, a result later replicated under laboratory conditions by Highhouse and Gallo (1997, as cited in Thomas & Reimann, 2022). The peak-end rule adds one more wrinkle: interviewers tend to remember an interaction by its most intense moment and its ending rather than by the interview as a whole (Thomas & Reimann, 2022), so a strong closing answer can outweigh an otherwise mediocre performance.

Overconfidence Bias

Many interviewers trust their own read on a candidate far more than the evidence warrants. Overconfidence bias describes this excessive faith in one's ability to judge capability from a short conversation, despite a substantial body of research showing that unstructured intuitive judgment predicts job performance poorly (Highhouse & Brooks, 2023). The practical consequence is that structured or data-driven methods get less traction than they should, simply because evaluators are already convinced their gut instinct is good enough.

Gender, Racial, and Age Bias

Beyond these general-purpose heuristics, several biases attach specifically to demographic categories, and here the evidence is unusually well documented. Gender bias reflects stereotyped assumptions about which roles men and women are suited for; consistent with Social Role Theory (Eagly, 1987), recruiters of either gender may unconsciously favor men for leadership or technical positions and women for supportive or nurturing ones, reinforcing occupational segregation even among evaluators who genuinely intend to be impartial.

Racial and ethnic bias has been demonstrated with unusual rigor through field experimentation. In a widely cited audit study, Bertrand and Mullainathan (2004) sent nearly five thousand fictitious resumes to newspaper job listings in Boston and Chicago, randomly assigning applicants distinctively White-sounding or African American-sounding names while holding resume quality fixed. Resumes with White-sounding names received about 50% more callbacks, and the gap held steady across occupation, industry, and employer size—evidence that discrimination can operate before a candidate is ever interviewed. The Stereotype Content Model (Fiske, Cuddy, Glick, & Xu, 2002) offers one explanation: people evaluate social groups along perceived warmth and competence, and when a group is stereotyped as lower on either dimension for a given role, recruiters may downgrade its members without registering that they are doing so.

Age bias follows a related but distinct logic. Older candidates are often stereotyped as less adaptable or less comfortable with technology than younger ones, even though the age stereotype content model finds that older workers are typically seen as high in competence but low in perceived warmth (Kite & Johnson, 1988)—a stereotype about likability, not ability, that nonetheless gets read as a verdict on capability.

Stereotype Threat

A different mechanism operates through the candidate rather than the interviewer. Stereotype threat, first identified by Steele and Aronson (1995), describes how simply being aware that one's group is negatively stereotyped on some dimension can impair performance on that dimension—independent of anything the interviewer says or does. Whysall (2018) frames this as a second, largely invisible channel through which bias enters selection: candidates from stigmatized groups may underperform in interviews or assessment centers not because they are less capable, but because the situation itself primes anxiety about confirming a stereotype. The research Whysall reviews suggest the effect bites hardest for people who care most about the domain being evaluated—meaning the most motivated and talented candidates from a stereotyped group can be paradoxically the most vulnerable to it. Because the threat originates in the environment rather than in any individual evaluator's head, addressing it means attending to situational cues—who sits on an interview panel, what imagery appears in recruitment materials—rather than relying on interviewer training alone.

Manifestation of Bias Across the Hiring Process

Bias does not confine itself to the interview room; it surfaces at every stage of recruitment and selection. Resume screening is often the first point of contact, and cues as simple as a candidate's name, university, or estimated age can quietly signal suitability—applicants with minority-associated names receive fewer callbacks (Bertrand & Mullainathan, 2004), and older candidates get screened out on unfounded assumptions about adaptability. "Cultural fit," evaluated a stage later, frequently functions as affinity bias wearing a legitimate-sounding label, steering recruiters toward candidates whose backgrounds mirror the existing team (Rivera, 2012). Inside the interview itself, gender and age stereotypes shape impression formation—assertiveness gets assumed in male candidates, technological fluency in younger ones—and the tone and difficulty of the questions asked can shift depending on how much an interviewer subconsciously favours someone, quietly unbalancing each candidate's opportunity to show what they can do. And when the final call rests on unstructured, gut-level impressions rather than documented evidence, none of this tends to get caught (Highhouse & Brooks, 2023).

Bias in Promotion and Reward Decisions

Bias does not vanish once a candidate is hired; most of the literature above concerns the point of entry, but the same distortions shadow people through their careers. If anything, Whysall (2018) argues, promotion and reward decisions are more exposed to implicit bias than hiring decisions, precisely because they tend to rest on vaguer, less formalized criteria. Ruderman and Ohlott's (1994) review of promotion decisions across three Fortune 500 companies found that formally collected performance data often played a surprisingly small role in who actually got promoted; the real process was intuitive and subjective, built around how familiar and comfortable decision-makers already felt with the candidate. Attribution bias sharpens the effect further: evaluators tend to credit an in-group member's success to stable internal qualities like talent or effort, while crediting an out-group member's success to external, unstable causes like luck (Whysall, 2018).

Actual compensation data shows how this discretion turns into measurable disparities. Castilla (2008), studying a large U.S. service-sector employer, found no significant race or gender gap in the performance evaluations managers assigned—yet a significant white-male advantage in the size of the salary increases that followed from otherwise identical evaluations. He traced the gap to a lack of accountability and transparency at the point where an evaluation gets converted into a dollar figure, a stage of the process largely invisible to outside scrutiny. Executive compensation shows a comparable pattern: Kulich, Trojanowski, Ryan, Haslam, and Renneboog (2011), examining a matched sample of male and female UK executive directors, found that bonuses awarded to men were both larger on average and more sensitive to fluctuations in company performance than bonuses awarded to women—strong performance by men, in other words, gets noticed and rewarded more readily. Panel composition matters too: Powell and Butterfield (2002) found that the effects of applicant race and gender on promotion to top management showed up when promotion panels were homogeneous (all white and male) but disappeared when panels were demographically mixed, echoing a broader pattern in which diverse panels curb biases that homogeneous ones let through.

Decision Noise in Hiring

If bias pulls judgment systematically in one direction, noise scatters it unpredictably in all directions. Kahneman, Sibony, and Sunstein (2021) popularized this distinction, arguing that hiring outcomes can vary substantially even once systematic bias has been minimized, simply because of factors that have nothing to do with a candidate's actual qualifications. They break the problem into three components. Level noise is the tendency of some evaluators to run consistently stricter or more lenient than others, regardless of who they're assessing. Pattern noise is subtler: different evaluators weigh the same set of candidate attributes differently, so one interviewer prizes communication skills above all else while another leans on technical knowledge. Occasion noise, the most fleeting of the three, captures how much a judgment that should in principle be stable can shift with an evaluator's mood, fatigue, or simply the time of day.

Awareness of the problem is not new—Highhouse and Brooks (2023) traced it back nearly a century in personnel selection research—but noise has historically drawn far less attention than bias, despite often doing just as much damage to decision quality. Drawing on decades of selection research, they found noise can be reliably reduced by breaking judgments down into specific components, agreeing on consistent standards for each one, and aggregating independent ratings mechanically rather than leaning on a single holistic impression. Their review also found that intuition-based judgment consistently underperforms structured, mechanical combinations of ratings, even when both approaches start from identical information (Highhouse & Brooks, 2023).

Evidence from Organizational Noise Audits

Kahneman, Rosenfield, Gandhi, and Blaser (2016), writing several years before the fuller treatment in Kahneman, Sibony, and Sunstein (2021), illustrated the scale of the noise problem with a memorable anecdote: a longtime customer at a financial services firm accidentally submitted the same application file to two different offices, which were supposed to apply the same guidelines, and received two very different quotes—leading the customer to take the business elsewhere. Curious how widespread this kind of inconsistency was, the authors conducted formal noise audits inside two financial services organizations, asking dozens of professionals in each to independently evaluate a shared set of realistic cases without being told that consistency, rather than the substance of their judgment, was under study. Executives in both firms predicted beforehand that disagreement between professionals evaluating the same case would fall somewhere between 5% and 10%, a level they considered normal for "matters of judgment." The measured noise index instead ranged from 34% to 62% in one organization and 46% to 70% in the other, and, strikingly, professionals with five or more years of experience were no more consistent than their junior colleagues. The authors traced this blind spot to a simple psychological mechanism: experienced professionals tend to be highly confident in their own judgment and hold their colleagues in equally high regard, which together lead them to assume, incorrectly, that a colleague reviewing the same case would reach much the same conclusion.

The Interplay of Bias and Noise: Unintentional Errors in Hiring

In practice, bias and noise rarely operate in isolation—they compound each other. Berthet (2022) found that decision-makers under time pressure lean more heavily on heuristics, which drives up both the systematic distortion of bias and the random scatter of noise at once, since rushed judgments end up simultaneously less consistent and more susceptible to irrelevant cues like a candidate's perceived similarity to the interviewer. One practical result: two equally qualified candidates can land very different outcomes—one hired, one rejected—not because of any real difference in ability, but because of which interviewer they happened to meet, what mood that interviewer was in, or where in the schedule their interview fell (Berthet, 2022).

None of this is cost-free. Poor-quality hires driven by biased or noisy judgment drag down team performance and are expensive to fix through additional recruitment, onboarding, and training (Highhouse & Brooks, 2023). Unstructured interviews, which are especially exposed to both bias and noise, show low predictive validity for future job performance in meta-analytic evidence, while structured interviews and objective tests do substantially better (McDaniel, Whetzel, Schmidt, & Maurer, 1994; Schmidt & Hunter, 1998). And beyond the immediate hiring decision, processes that feel inconsistent or unfair can erode candidates' and employees' trust in the organization and dent its reputation as an employer (Kahneman et al., 2021).

THEORETICAL FRAMEWORKS

Four theoretical traditions in psychology and behavioral economics provide a coherent explanation for why bias and noise arise in hiring decisions.

Dual Process Theory
Kahneman (2011) and others built Dual Process Theory around a distinction between two modes of thinking: System 1, fast, automatic, and intuitive, and System 2, slow, deliberate, and analytical. Rapid impressions of candidates—based on appearance, accent, or a vague sense of "fit"—typically come from System 1, and they feel confident precisely because they're fast, even though that speed is what makes them prone to the shortcuts behind halo effects and stereotyping. Structured hiring procedures exist, in effect, to force System 2 online and override the default pull toward intuition.

The Heuristics and Biases Framework
Tversky and Kahneman (1974), later elaborated in Kahneman, Slovic, and Tversky (1982), catalogued the specific mental shortcuts behind the Heuristics and Biases framework—anchoring, availability, and representativeness among them—each producing predictable errors under uncertainty. Hiring decisions are usually made with incomplete information and against a clock, which makes them especially fertile ground for these heuristics.

Bounded Rationality
Decision-makers, in Simon's (1955) account, are limited by finite cognitive capacity, incomplete information, and time pressure, and so tend to "satisfice"—settling for the first acceptable option rather than searching exhaustively for the best one. Bounded rationality is a useful lens on why recruiters lean on simplified proxies like institutional prestige or a polished resume format instead of conducting a full evaluation of every candidate's capability (Simon, 1955, 1991).

Noise Theory
Noise theory takes the study of judgment error somewhere the other three frameworks don't go. Kahneman, Sibony, and Sunstein (2021) formally theorized noise as distinct from bias and proposed "decision hygiene" practices—structured procedures, independent judgments, aggregation rules—as the primary fix. Where the other frameworks explain directional distortion, this one explains inconsistency across judgments, and the two problems call for different remedies.

Applied together, these frameworks show that hiring is not a purely rational organizational process but a psychologically complex activity shaped by cognitive limitations, intuitive thinking, and contextual variability. This integrated understanding underlies the practical recommendations discussed in the following section.

Artificial Intelligence in Recruitment: Boon or a Bane
Artificial intelligence gets pitched as a fix for human bias in hiring: algorithms can, in principle, apply identical criteria to every candidate. The empirical record complicates that promise. The most widely cited cautionary tale is Amazon's experimental AI recruiting tool, developed starting in 2014 to score resumes on a one-to-five scale (Dastin, 2018). Trained on ten years of historical resumes submitted to a male-dominated technology workforce, the system learned to penalize resumes containing terms associated with women—membership in women's organizations, attendance at women's colleges—while favouring language patterns more common in men's resumes. Amazon eventually shut the project down after concluding it couldn't reliably guarantee the tool's neutrality (Dastin, 2018). The broader lesson travels well beyond this one case: an algorithm trained on biased historical data can replicate, and even amplify, the very biases it was built to remove.

The Amazon case is not an isolated data point. More recent work, though it remains an emerging and unevenly consolidated literature, points in a similar direction: AI hiring systems appear to inherit, and can obscure rather than expose, the discriminatory patterns already present in their training data, which is part of why careful auditing matters more than any single technical safeguard. There is also preliminary evidence that algorithmic decision-making introduces its own form of noise—models behaving inconsistently across cases that look superficially similar—suggesting automation may relocate variability rather than remove it. Some scholars go further, cautioning that algorithmic tools on their own cannot touch the deeper structural inequalities built into organizational hiring practices; on that view, a better model is necessary but nowhere near sufficient without broader institutional change.

The picture gets more nuanced once human-AI collaboration enters the frame. Pairing algorithmic recommendations with human oversight can improve decision quality—but the evidence suggests this only holds when recruiters actually scrutinize the AI's output rather than accepting it uncritically; automatic deference to algorithmic suggestions risks reinforcing discrimination rather than reducing it. One relevant experiment found that a reliable AI tool by itself was not enough to improve diversity outcomes; the gains instead traced back to AI-generated explanations tied explicitly to diversity considerations, clearly stated organizational diversity protocols, and decision-makers actually willing to reflect critically on what the tool recommended—not to how sophisticated the algorithm was or how experienced the decision-maker happened to be. Taken together, this body of work suggests AI can be part of debiased, low-noise hiring, but only inside a governance structure that actively monitors for embedded bias and keeps meaningful human judgment in the loop rather than replacing it wholesale (Highhouse & Brooks, 2023).

Evidence-Based Strategies for Reducing Bias and Noise

The literature converges on a consistent set of practices for reducing both bias and noise in hiring, all of which share the underlying logic of replacing global, intuitive judgment with structured, decomposed, and independently aggregated evaluation.

Structured Interviews

Structured interviews—every candidate asked the same set of predetermined questions, scored against a consistent rubric—are among the best-validated interventions in the selection literature. McDaniel, Whetzel, Schmidt, and Maurer (1994), synthesizing 245 validity coefficients drawn from more than 86,000 individuals, found structured interviews substantially more predictive of job performance than unstructured ones, a result Schmidt and Hunter (1998) corroborated in their synthesis of 85 years of personnel selection research, which identified structured interviews as among the strongest available predictors of future performance.

Validated Testing and Work Samples

Standardized cognitive ability tests, work-sample exercises, and structured assessments provide measurable, comparable data about candidates that reduce reliance on subjective, holistic impressions (Schmidt & Hunter, 1998). Because these instruments assess the same competencies in the same way for every applicant, they are inherently more resistant to both directional bias and random noise than free-form conversation.

Panel Interviews and Independent, Aggregated Ratings

Using multiple interviewers who record their evaluations independently before discussing them helps average out individual idiosyncrasies in judgment. Highhouse and Brooks (2023) emphasized that mechanically aggregating independent ratings—rather than allowing a group to reach a single consensus impression through discussion, which can amplify the most persuasive voice in the room—is one of the most effective noise-reduction techniques available.

Joint Versus Separate Evaluation

How candidates are compared, and not only what information evaluators see, also shapes how much bias enters a decision. Bohnet, van Geen, and Bazerman (2012) found that assessors evaluating candidates one at a time, in isolation, tended to rely more heavily on group stereotypes than assessors who evaluated the same candidates jointly, side by side. Joint evaluation appeared to redirect attention toward evidence of future job performance and away from category-based impressions, making it not only less prone to bias but, the authors noted, more cost-effective than a long series of separate, sequential evaluations. This finding complements the case for blind screening: where blind review removes identity cues before a judgment is formed, joint or comparative review changes the structure of the judgment itself, and the two strategies can be combined.

Clear Job Criteria and Blind Screening

Defining precise, job-relevant criteria in advance helps interviewers focus on competencies that predict performance rather than personal preferences. Removing identifying information such as name, age, or photograph from initial resume screening (blind recruitment) can further reduce the influence of demographic stereotypes before a candidate is even seen or heard.

A frequently cited natural experiment offers some of the clearest evidence that blind evaluation changes outcomes rather than merely intentions. Goldin and Rouse (2000) studied the effect of introducing screens that concealed a musician's identity during preliminary rounds of symphony orchestra auditions in the United States. The switch to blind auditions increased the probability that a woman would advance out of preliminary rounds by roughly 50%, and the authors estimate that it explains between 30% and 55% of the increase in the share of women among newly hired orchestra musicians since 1970. Because the only change was whether the evaluator could see the candidate, the result is difficult to attribute to anything other than the removal of an identity cue that had previously been shaping judgment, and it offers a template for what blind resume or portfolio review might plausibly achieve in other hiring contexts.

Training and Debiasing

Training interviewers to recognize their own susceptibility to bias raises awareness, which is a necessary first step, but the literature is clear that awareness alone is usually insufficient to change behavior. Fasolo, Heard, and Scopelliti's review of roughly 100 studies on cognitive bias mitigation in organizations distinguished between debiasing, which targets the individual's thought process through training, warnings, and feedback, and choice architecture, which restructures the decision environment itself—for example, through defaults or clearer information presentation—so that better choices require less individual effort. Their analysis found both approaches effective under different conditions but rarely compared directly, and recommended combining structural tools such as mandatory rubrics and forced-choice evaluation items with debiasing training for the strongest effect.

Governed Use of Artificial Intelligence

Algorithmic tools can help standardize evaluation, but only when accompanied by active governance: auditing training data for embedded historical bias, monitoring outputs for disparate impact, maintaining transparency about how scores are generated, and preserving meaningful human oversight rather than automatic deference to algorithmic recommendations (Dastin, 2018; Highhouse & Brooks, 2023).

Algorithms and Reasoned Rules

A common assumption is that building an algorithm to replace human judgment requires large volumes of historical outcome data—for hiring, this would mean years of records linking resume characteristics to eventual job performance. Kahneman, Rosenfield, Gandhi, and Blaser (2016) show that this assumption is largely unnecessary. They describe a class of simple predictive formulas they call "reasoned rules," built by selecting a small number of variables that are obviously related to the outcome of interest, converting each candidate's values on those variables into standardized scores, and averaging the scores with equal weights—no historical outcome data or statistical modeling is required, only a modest set of current cases and a defensible list of relevant criteria. The authors report that such rules have performed comparably to more sophisticated statistical models across a range of applied settings, including personnel selection, because the weights assigned by complex models are frequently unstable and contribute little additional predictive power beyond a sensible equal-weighting scheme. For hiring specifically, this suggests that even organizations without the data infrastructure to build a full statistical hiring model can still construct a low-cost, noise-free scoring rule from a handful of job-relevant criteria, so long as human decision-makers retain the authority to override the rule in clear-cut exceptional cases and remain responsible for translating its output into a final hiring decision.

CONCLUSION

The literature reviewed here demonstrates that hiring decisions, however well-intentioned, remain vulnerable to two distinct but interacting sources of error: cognitive bias and decision noise. Biases such as the halo and horn effects, confirmation bias, similarity and affinity bias, anchoring, availability, representativeness, stereotyping, and status-quo preference systematically distort how resumes are screened, how interviews are conducted, and how candidates are ultimately judged (Nisbett & Wilson, 1977; Nickerson, 1998; Rivera, 2012; Thomas & Reimann, 2022). At the same time, decision noise introduces unpredictable inconsistency into these same judgments, so that equally qualified candidates can receive very different outcomes depending on who evaluates them and under what circumstances (Kahneman, Sibony, & Sunstein, 2021; Highhouse & Brooks, 2023).

Four theoretical frameworks—Dual Process Theory, the Heuristics and Biases framework, Bounded Rationality, and Noise theory—together explain why these errors persist even among trained, well-meaning professionals: human judgment is inherently limited, prone to intuitive shortcuts, and inconsistent across time and context. The growing use of artificial intelligence in recruitment offers a partial remedy but also new risks, since algorithms trained on biased historical data can reproduce and even amplify existing patterns of discrimination (Dastin, 2018).

Fortunately, the literature also points toward concrete, well-validated remedies. Structured interviews, validated testing, independently aggregated panel ratings, blind screening, clear job criteria, and carefully governed algorithmic tools have each been shown to reduce bias, noise, or both, and in combination they offer organizations a realistic path toward fairer and more accurate hiring (McDaniel et al., 1994; Schmidt & Hunter, 1998; Highhouse & Brooks, 2023).

As with any narrative synthesis, these conclusions are only as strong as the underlying studies allow. Several of the effects reviewed here rest on a single audit study or laboratory demonstration rather than large-scale field replication, and their generalizability across industries, seniority levels, and national labor markets remains an open empirical question; readers should treat the strength of individual findings as varying accordingly rather than uniformly settled. Addressing cognitive bias and decision noise in hiring is nonetheless not merely a technical refinement but an ethical and strategic necessity: it is central to building workplaces where opportunity is distributed according to merit rather than the accident of who happens to be interviewing, on what day, and in what mood.

REFERENCES

Berthet, V. (2022). The impact of cognitive biases on professionals' decision-making: A review of four occupational areas. *Frontiers in Psychology, 12*, Article 802439. https://doi.org/10.3389/fpsyg.2021.802439

Bertrand, M., & Mullainathan, S. (2004). Are Emily and Greg more employable than Lakisha and Jamal? A field experiment on labor market discrimination. *American Economic Review, 94*(4), 991–1013. https://doi.org/10.1257/0002828042002561

Bohnet, I., van Geen, A., & Bazerman, M. H. (2012). *When performance trumps gender bias: Joint versus separate evaluation* (Working Paper No. 12-083). Harvard Business School.

Castilla, E. J. (2008). Gender, race, and meritocracy in organizational careers. *American Journal of Sociology, 113*(6), 1479–1526. https://doi.org/10.1086/588738

Dastin, J. (2018, October 10). Amazon scraps secret AI recruiting tool that showed bias against women. *Reuters*. https://www.reuters.com/article/us-amazon-com-jobs-automation-insight-idUSKCN1MK08G

Dhingra, M. (2018). Factors affecting hiring decision: Employers' perspective in National Capital Region. *Journal of Strategic Human Resource Management, 7*(3), 15–21.

Eagly, A. H. (1987). *Sex differences in social behavior: A social-role interpretation.* Lawrence Erlbaum Associates.

Fasolo, B., Heard, C., & Scopelliti, I. (2024). Mitigating cognitive bias to improve organizational decisions: An integrative review, framework, and research agenda. *Journal of Management.* Advance online publication.

Fiske, S. T., Cuddy, A. J. C., Glick, P., & Xu, J. (2002). A model of (often mixed) stereotype content: Competence and warmth respectively follow from perceived status and competition. *Journal of Personality and Social Psychology, 82*(6), 878–902. https://doi.org/10.1037/0022-3514.82.6.878

Goldin, C., & Rouse, C. (2000). Orchestrating impartiality: The impact of "blind" auditions on female musicians. *American Economic Review, 90*(4), 715–741. https://doi.org/10.1257/aer.90.4.715

Highhouse, S., & Brooks, M. E. (2023). Improving workplace judgments by reducing noise: Lessons learned from a century of selection research. *Annual Review of Organizational Psychology and Organizational Behavior, 10*, 519–533. https://doi.org/10.1146/annurev-orgpsych-120920-050708

In V. Caven & S. Nachmias (Eds.), *Hidden inequalities in the workplace: A guide to the current challenges, issues and business solutions* (pp. 215–243). Palgrave Macmillan. https://doi.org/10.1007/978-3-319-5

Kahneman, D. (2011). *Thinking, fast and slow.* Farrar, Straus and Giroux.

Kahneman, D., Rosenfield, A. M., Gandhi, L., & Blaser, T. (2016). Noise: How to overcome the high, hidden cost of inconsistent decision making. *Harvard Business Review, 94*(10), 38–46.

Kahneman, D., Sibony, O., & Sunstein, C. R. (2021). *Noise: A flaw in human judgment.* Little, Brown Spark.

Kahneman, D., Slovic, P., & Tversky, A. (Eds.). (1982). *Judgment under uncertainty: Heuristics and biases.* Cambridge University Press.

Kite, M. E., & Johnson, B. T. (1988). Attitudes toward older and younger adults: A meta-analysis. *Psychology and Aging, 3*(3), 233–244. https://doi.org/10.1037/0882-7974.3.3.233

Kulich, C., Trojanowski, G., Ryan, M. K., Haslam, S. A., & Renneboog, L. D. R. (2011). Who gets the carrot and who gets the stick? Evidence of gender disparities in executive remuneration. *Strategic Management Journal, 32*(3), 301–321. https://doi .org/10.1002/smj.878

Mandal, A., & Shrivastava, R. (2023). The hiring process and its significance. *IRE Journals, 7*(1), 256–257.

McDaniel, M. A., Whetzel, D. L., Schmidt, F. L., & Maurer, S. D. (1994). The validity of employment interviews: A comprehensive review and meta-analysis. *Journal of Applied Psychology, 79*(4), 599–616. https://doi.org/10.1037/0021-9010.79.4.599

Nickerson, R. S. (1998). Confirmation bias: A ubiquitous phenomenon in many guises. *Review of General Psychology, 2*(2), 175–220. https://doi.org/10.1037/1089-2680.2.2 .175

Nisbett, R. E., & Wilson, T. D. (1977). The halo effect: Evidence for unconscious alteration of judgments. *Journal of Personality and Social Psychology, 35*(4), 250–256. https:// doi.org/10.1037/0022-3514.35.4.250

Powell, G. N., & Butterfield, D. A. (2002). Exploring the influence of decision makers' race and gender on actual promotions to top management. *Personnel Psychology, 55*(2), 397–428. https://doi.org/10.1111/j.1744-6570.2002.tb00116.x

Rajyalaxmi, M. (2023). Human resource: Recruitment and selection process. *International Journal of Novel Research and Development, 8*(10), c118–c119.

Rivera, L. A. (2012). Hiring as cultural matching: The case of elite professional service firms. *American Sociological Review, 77*(6), 999–1022. https://doi.org/10.1177/000 3122412463213

Ruderman, M. N., & Ohlott, P. J. (1994). *The realities of management promotion.* Center for Creative Leadership.

Salehzadeh, R., & Ziaeian, M. (2024). Decision making in human resource management: A systematic review of the applications of analytic hierarchy process. *Frontiers in Psychology, 15*, Article 1400772. https://doi.org/10.3389/fpsyg.2024.1400772

Schmidt, F. L., & Hunter, J. E. (1998). The validity and utility of selection methods in personnel psychology: Practical and theoretical implications of 85 years of research findings. *Psychological Bulletin, 124*(2), 262–274. https://doi.org/10.1037/0033-2909.124.2.262

Schmitt, N., & Kim, B. (2007). Selection decision-making. In P. Boxall, J. Purcell, & P. Wright (Eds.), *The Oxford handbook of human resource management* (pp. 300–323). Oxford University Press.

Simon, H. A. (1955). A behavioral model of rational choice. *Quarterly Journal of Economics, 69*(1), 99–118. https://doi.org/10.2307/1884852

Simon, H. A. (1991). Bounded rationality and organizational learning. *Organization Science, 2*(1), 125–134. https://doi.org/10.1287/orsc.2.1.125

Steele, C. M., & Aronson, J. (1995). Stereotype threat and the intellectual test performance of African Americans. *Journal of Personality and Social Psychology, 69*(5), 797–811. https://doi.org/10.1037/0022-3514.69.5.797

Thomas, O., & Reimann, O. (2022). The bias blind spot among HR employees in hiring decisions. *German Journal of Human Resource Management, 37*(1), 5–22. https://doi.org/10.1177/23970022221094523

Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science, 185*(4157), 1124–1131. https://doi.org/10.1126/science.185.4157.1124

Wharton Executive Education. (2025, May 21). *What is status quo bias and how does it affect the workplace?* University of Pennsylvania. https://executiveeducation.wharton.upenn.edu/thought-leadership/wharton-online-insights/status-quo-bias/

Whysall, Z. (2018). Cognitive biases in recruitment, selection, and promotion: The risk of subconscious discrimination.

Acknowledgment
The author(s) appreciates all those who participated in the study and helped to facilitate the research process.

Conflict of Interest
The author(s) declared no conflict of interest.

How to cite this article: Sehgal, A. & Sharma, R. (2026). Cognitive Biases and Decision Noise in Hiring Decisions: A Literature Review. *International Journal of Indian Psychology*, *14*(3)*,* 667-680. DIP:18.01.056.20261403, DOI:10.25215/1403.056