Docket

DOCKET / RESEARCH & PRACTICAL GUIDANCE

Decision Hygiene for Deal Teams: Less Noise and Bias in IC

> Build a more reviewable acquisition pipeline with Docket. Our deal-origination software and managed research help private equity teams screen companies against a mandate and understand the evidence behind each finding.

Published by Docket · 2026-10-29 · investment committee process · cognitive bias private equity · decision hygiene · noise in decision making

Original article: https://docket.capital/articles/investment-committee-decision-hygiene-bias



In this article

1. Bias versus noise in deal decisions
2. Known biases that hit screening
3. Independent assessments before discussion
4. Structure the rubric
5. Reference classes and the outside view
6. Pre-mortems
7. Memo and meeting mechanics
8. A one-page decision-hygiene checklist

Bias versus noise in deal decisions

An investment committee process can go wrong in two different ways, and the fixes differ.

Bias is error in a consistent direction. If every reviewer on a deal team overweights the first revenue figure they saw, the team's judgments lean the same way, and averaging more opinions does not cancel the lean.

Noise is scatter. The publisher's description of *Noise* by Daniel Kahneman, Olivier Sibony and Cass R. Sunstein defines it as "variability in judgments that should be identical," with examples such as two doctors in the same city giving different diagnoses to identical patients, or the same person deciding differently "depending on whether it is morning or afternoon" (Hachette Book Group).

A 2016 *Harvard Business Review* article by Kahneman, Andrew M. Rosenfield, Linnea Gandhi and Tom Blaser opens with a business case. At a global financial services firm the authors worked with, a longtime customer accidentally submitted the same application file to two offices; the employees were supposed to follow the same guidelines, yet the offices "returned very different quotes," and the customer took the business to a competitor (HBR, 2016). HBR's published description of the article explains how organizations can perform a "noise audit": members of a professional unit evaluate a common set of cases, and the degree to which their assessments vary provides the measure of noise (HBR Store, product R1610B).

Translate that to origination. If two associates read the same target's filings and one scores it a strong fit while the other passes, that is noise: the target did not change, the reader did. If both associates consistently rate founder-led businesses higher than the evidence supports, the firm has bias. A screening process should be designed to detect both, and the two require different tools:

Problem	What it looks like in screening	Primary tool
Noise between people	Two reviewers score the same target far apart	Independent scoring on a shared rubric, then agreement checks
Noise within a person	One reviewer scores similar targets differently on different days	Anchored scale descriptions, rescoring a sample blind
Directional bias	The whole team leans the same way on a class of targets	Reference-class checks, pre-mortems, structured dissent
Process bias	The committee hears the sponsor's framing first and argues from it	Written scores before discussion, fixed speaking order

Known biases that hit screening

Anchoring. Amos Tversky and Daniel Kahneman's 1974 paper in *Science* described three heuristics people use when judging under uncertainty: representativeness, availability, and "adjustment from an anchor, which is usually employed in numerical prediction when a relevant value is available." The abstract states that these heuristics "are highly economical and usually effective, but they lead to systematic and predictable errors" (Tversky and Kahneman, 1974). In screening, the relevant value is often a revenue estimate from a data vendor, an asking multiple in a teaser, or the first number in the associate's memo. Later numbers get read as adjustments to it.

Optimistic forecasts. Dan Lovallo and Daniel Kahneman's 2003 *HBR* article opens with Oxford Health Plans, which began building a claims and payments system in 1992; when the company disclosed on October 27, 1997 that its system and accounts were in disarray, its stock price "dropped 63%, destroying more than $3 billion in shareholder value in a single day" (Lovallo and Kahneman, 2003). The article's summary attributes overly optimistic investment forecasts to "a combination of cognitive biases (including anchoring and competitor neglect) and organizational pressures" (Lovallo and Kahneman, 2003). A deal team that has spent weeks on a target is under exactly that kind of pressure.

Confirmation and loss aversion. The 2011 *HBR* article by Kahneman, Lovallo and Sibony gives short working definitions: confirmation bias "leads people to ignore evidence that contradicts their preconceived notions," anchoring causes people "to weigh one piece of information too heavily," and loss aversion "makes them too cautious." The same opening adds that, in the authors' experience, awareness of biases "has done little to improve the quality of business decisions" (Kahneman, Lovallo and Sibony, 2011).

That last point is the case for process over training. Telling a committee about anchoring does not stop anchoring. Changing the order in which information arrives can.

Where these show up in a typical origination funnel:

Stage	Likely mechanism	Example trigger
List building	Availability	Targets the team has heard of get researched first and more thoroughly
First screen	Anchoring	A vendor revenue estimate sets the frame for every later size judgment
Memo drafting	Confirmation	The author searches for support for a thesis already chosen
IC discussion	Anchoring, conformity	The deal lead's opening view frames the debate
Pass decisions	Loss aversion	Reluctance to pass on a target after sunk research time, or excess caution on an unfamiliar sector

Independent assessments before discussion

The single most practical step is to collect judgments before people talk. Once a senior partner says "I like this one," every later score is partly a response to that statement.

A workable protocol:

1. **Freeze the evidence packet.** Every reviewer gets the same retained pages, filings and notes for each target. If one reviewer has extra information, disagreement may reflect unequal evidence rather than inconsistent judgment.
2. **Score alone, in writing.** Each reviewer scores each criterion on the shared screening rubric and writes one sentence citing the evidence behind each score. No overall score yet.
3. **Submit before seeing others.** Scores become visible only after everyone has submitted.
4. **Compare by criterion, not by total.** Look for criteria where reviewers differ by two or more points. Those are either ambiguous rubric language or genuinely borderline evidence.

5. **Discuss the gaps, then rescore.** Keep the original scores. The record of disagreement is how you improve the rubric.
6. **Track agreement over time.** Measure inter-rater reliability on a recurring sample. This is the noise audit idea from the 2016 *HBR* article, applied to screening: a common set of cases, several reviewers, and a measure of how much they vary (HBR, 2016).

Illustrative example (fictional). Three reviewers score three fictional targets on a 1 to 5 scale before any discussion.

Criterion	Reviewer	Company A	Company B	Company C
Recurring revenue evidence	R1	4	2	3
	R2	4	2	5
	R3	4	3	2
Owner transition signals	R1	3	4	2
	R2	3	4	2
	R3	2	4	2
Customer concentration	R1	3	3	4
	R2	3	4	4
	R3	3	3	4

Company C's recurring revenue scores span 2 to 5. That spread is the signal. In the fictional discussion, R2 counted a one-time project list as recurring work, while R3 required a stated service contract. The fix is not a vote; it is a rubric definition of "recurring" that names the evidence required.

If the committee had discussed Company C first and heard R2's 5 aloud, the spread would probably never have surfaced.

Structure the rubric

Independent scoring only helps if the rubric gives reviewers the same target to aim at. Several drafting rules reduce noise directly:

- **Define each criterion by observable evidence.** "Strong recurring revenue" invites opinion. "The company publishes service agreements, maintenance plans or subscription pricing on its own site or in filings" can be checked.
- **Anchor every scale point.** Write what a 1, 3 and 5 look like for each criterion, with a short example. Unanchored scales are where within-person noise lives.
- **Name the allowed sources.** Company-published pages, filings and registry records behave differently from vendor estimates. State which count.

- **Score criteria before any overall judgment.** An overall impression formed first becomes an anchor for the components. Compute or record the overall view last.
- **Separate "no" from "unknown."** A missing fact is not a negative fact. Give "unknown" its own value so it is not silently scored as a 1 or a 3.
- **Version the rubric.** When a definition changes after a disagreement, record which version governed each score.

Our page on screening criteria covers building and weighting the scored dimensions.

Reference classes and the outside view

Lovallo and Kahneman's remedy for optimistic forecasts is what the *HBR* summary calls the "outside view," also known as reference-class forecasting. It "completely ignores the details of the project at hand"; managers instead examine the experiences of a class of similar projects, lay out a rough distribution of outcomes for that class, and then position the current project in that distribution (Lovallo and Kahneman, 2003).

For a deal team, the reference class should come from the firm's own records, not from a market statistic quoted in a pitch book. Prompts to add to the IC memo template:

- **What is our class?** For example, all add-on targets in this sector the firm took to LOI in the past several years, or all founder-owned targets above a size band the team has screened.
- **What happened to that class?** Of those, how many closed, how many broke in diligence, and for what stated reasons? Use the firm's own deal log.
- **Where does this target sit, and why?** The memo author must argue for a position relative to the class, citing specific evidence, before presenting the inside-view case.
- **What would move it?** List the evidence that would push the target above or below the class median.

Two cautions. A small firm's history may hold too few comparable deals to form a meaningful class; say so rather than inventing a base rate. And the outside view can itself become an anchor, so present it as the first estimate and require inside-view adjustments to be justified in writing.

Pre-mortems

Gary Klein's 2007 *HBR* article describes the technique: "In a premortem, team members assume that the project they are planning has just failed... and then generate plausible reasons for its demise." The article's framing is about dissent: those with reservations "may speak freely at the outset, so that the project can be improved rather than autopsied," and making it safe for knowledgeable dissenters to speak up improves a project's chances of success (Klein, 2007).

An IC version, run before the final vote:

1. **Set the scene.** The chair states: "It is two years after close. This investment has failed badly."
2. **Write alone.** Each member writes, silently and independently, the most plausible reasons for that failure. Several minutes is enough.

3. **Collect in rotation.** Go around the room, one reason per person per round, starting with the most junior member and ending with the deal lead.
4. **Sort.** Group reasons into those the team has evidence on, those it could test before signing, and those it cannot test.
5. **Assign.** Each testable reason gets an owner and a diligence item. Untestable reasons go into the decision log as accepted risks.

Illustrative example (fictional). A committee runs a premortem on a fictional regional services platform. Reasons include: the founder's informal customer relationships do not transfer; the largest customer's contract renewal was assumed rather than confirmed; the add-on pipeline is thinner than the thesis needs. The second reason is testable before signing, so it becomes a confirmatory diligence item with a named owner. The first is partly testable through reference calls. The third is logged as a known risk with a review date.

Memo and meeting mechanics

Much of an investment committee process's decision hygiene is sequencing and record-keeping. The 2011 *HBR* article makes the general case: with important decisions, executives should conduct a careful review "not only of the content of recommendations but of the recommendation process," and the authors offer a 12-question checklist for doing so (Kahneman, Lovallo and Sibony, 2011). Practical mechanics that follow from that idea:

- **Circulate scores and memo in writing first.** Members read and score before the meeting.
- **Speaking order.** Junior members first, deal lead and senior partners last. The aim is to hear independent views before the most influential one.
- **Devil's advocate, with limits.** A rotating assigned skeptic can surface objections, but an assigned role is easy to discount as theater. Pair it with written dissent: anyone who disagrees with the outcome records why, in one paragraph, in the decision log.
- **Separate the evidence from the recommendation.** The memo should state what is known (with sources), what is unknown, and what is inferred, in that order. See our evidence standards for how we separate those categories.
- **Decision log.** Record the decision, the vote, the key assumptions, the premortem risks, the dissent and a date to revisit. Without a log, the reference class for future deals does not exist.

When to trust intuition. Experienced partners often have strong instincts about targets. Kahneman and Gary Klein's 2009 paper in *American Psychologist* addresses when that is warranted. Its abstract concludes that "evaluating the likely quality of an intuitive judgment requires an assessment of the predictability of the environment in which the judgment is made and of the individual's opportunity to learn the regularities of that environment," and that "subjective experience is not a reliable indicator of judgment accuracy" (Kahneman and Klein, 2009).

Applied to deal teams, that suggests two questions before deferring to a partner's gut call:

Question	Points toward trusting intuition	Points toward caution
How predictable is this kind of judgment?	Recurring, well-defined calls (for example, whether a website describes a service business at all)	Long-horizon outcomes (for example, whether a platform will compound through add-ons)
Has this person had feedback on similar calls?	Many similar deals with known outcomes, reviewed afterward	Few comparable deals, or outcomes never reviewed

Confidence alone is not on the list, because the paper's abstract says subjective experience does not indicate accuracy.

A one-page decision-hygiene checklist

Print this and attach it to the IC memo template.

Before screening

- [ ] Rubric criteria defined by observable evidence, with anchored scale points
- [ ] Allowed sources named per criterion; "unknown" is a separate value
- [ ] Rubric version recorded

During screening

- [ ] Every reviewer scores from the same frozen evidence packet
- [ ] Scores submitted independently, before anyone sees others' scores
- [ ] Criteria scored before any overall judgment
- [ ] Disagreements of two or more points flagged by criterion; originals preserved

Memo

- [ ] Reference class from the firm's own deal log, or a statement that none is large enough
- [ ] Target positioned against the class before inside-view adjustments
- [ ] Known, unknown and inferred facts separated, each with sources

IC meeting

- [ ] Written scores circulated before the meeting
- [ ] Junior members speak first, deal lead last
- [ ] Premortem run, with reasons sorted into tested, testable and untestable
- [ ] Written dissent recorded

After

- [ ] Decision log entry: decision, assumptions, risks, dissent, review date
- [ ] Recurring agreement check on a sample of screened targets, with results tracked over time

Limits. None of the sources cited here studies private equity screening or fund returns, and this article makes no claim that these practices improve any firm's returns. What these practices can show directly is narrower and measurable: whether reviewers agree more often, whether disagreements trace to fixable rubric language, and whether the committee's records support learning from past decisions.

Docket's Scout agent answers a client's screening criteria as typed values from a fixed list, each quoting a retained page verbatim.

Source links

Links cited in this article, in order of first appearance. Full addresses are provided for readers using extracted text.

1. https://www.hachettebookgroup.com/titles/daniel-kahneman/noise/9780316451390
2. https://hbr.org/2016/10/noise
3. https://store.hbr.org/product/noise-how-to-overcome-the-high-hidden-cost-of-inconsistent-decision-making/R1610B
4. https://doi.org/10.1126/science.185.4157.1124
5. https://hbr.org/2003/07/delusions-of-success-how-optimism-undermines-executives-decisions
6. https://hbr.org/2011/06/the-big-idea-before-you-make-that-big-decision
7. https://hbr.org/2007/09/performing-a-project-premortem
8. https://doi.org/10.1037/a0016755

THE PUBLISHER

About Docket

Docket provides deal-origination research software and managed research for private equity firms. We help investment teams investigate acquisition targets using structured screening criteria, retained sources and reviewable company evidence. Teams can work through a self-serve platform or use managed research, depending on how they want research delivered.

Research against a defined mandate

A useful target list needs more than company names. Docket focuses on the questions that determine whether a company fits an investment mandate, including the evidence needed to support or qualify each answer. Our research approach makes the connection between screening criteria, source material and conclusions visible to the team reviewing the work.

Triage, Scout and Audit

Docket's three named research agents perform complementary tasks. Triage resolves company identity and screens fit. Scout collects sourced answers against the mandate. Audit checks retained evidence, addresses contradictions and leaves unsupported answers visibly unresolved. This structure helps reviewers distinguish established findings from missing information and questions requiring further investigation.

Evidence that supports investment-team judgment

Our research library covers market mapping, screening criteria, private-company data, succession and ownership, source evaluation and evidence standards. These resources explain the methods and limitations behind origination research. Findings support a team's judgment; they do not establish that a company is for sale or guarantee a transaction or investment outcome.

Work with Docket

Visit Docket to explore the platform, managed research and the current contact path. Read about screening criteria, evidence standards, AI research agents and market mapping.

Public examples are illustrative unless explicitly identified otherwise. Research preparation and authorized outreach are separate activities; confidential target lists and customer outcomes should never be inferred from an educational example.

Website: https://docket.capital

This article is educational. Verify consequential medical, legal, financial, regulatory and technical decisions with appropriate professionals and the cited primary sources. Company services are described separately from the article's research findings.