APPLIED ETHICS · STUDY GUIDE

Exam 7 Study Guide

This guide closes the course. Chapter 13 shows how data and algorithms carry values, where bias enters, and who gets blamed when automated systems fail. Chapter 14 turns everything into a practice you can run on a busy shift: a method, a script, an institution that makes speaking up survivable, and habits that last a career. Read the "what to know" paragraphs, copy the tables out by hand, then try each worked case with the steps covered. Finish each chapter with its self-test.

WHAT THIS GUIDE COVERS

- Chapter 13 · Kranzberg, affordances and defaults, policy vacuums and Collingridge, dual use, privacy versus confidentiality versus security, contextual integrity, consent failures, re-identification, HIPAA's gaps, bias sources, COMPAS and the impossibility theorem, many hands and the crumple zone, automation bias, alert fatigue, GDPR and the EU AI Act.
- Chapter 14 · bounded ethicality and ethical fading, situational forces and the 2021 retraction, the four topics, CASES and ISSUES, Rest's four components, Giving Voice to Values, CUS and the two-challenge rule, psychological safety, ethics committees and consults, just culture, burnout versus moral distress and injury, moral repair and a lasting practice.

CHAPTER 13

Technology, Data and Artificial Intelligence

Read it in full: **Chapter 13 in the textbook**. Practice it in the lab: **Where Does It Travel? The Contextual Integrity Router** and **The Fairness Bench: Pick Which Unfairness You Will Accept**.

What to know

Kranzberg (1985, published 1986): "Technology is neither good nor bad; nor is it neutral." Designs have **affordances** (Gibson; Norman) that make some acts easy and others hard. EHR fields, two-option sex drop-downs, 15-character name limits and copy-forward defaults govern behavior. Crash-test dummies: a fifth-percentile female dummy entered US frontal testing from

2003 and star ratings in model year 2011; Bose (2011) found about 47% higher odds of serious injury for belted female drivers.

Moor (1985): computers are logically malleable, creating **policy vacuums** that first need a conceptual muddle cleared. The **Collingridge dilemma** (1980): early on you can steer but do not know the effects; later you know but cannot steer. Demanding harm data before regulating favors deployment. **Dual use**: Urbina (2022) flipped a drug-discovery model to generate about 40,000 toxic candidates in under six hours. **Value-sensitive design** (Friedman) treats values as requirements; defaults matter (Austria near 100% vs Germany 12% for organ donation).

Privacy is control over the flow of information about you; **confidentiality** is the duty of the one entrusted; **security** is the technical control. Warren and Brandeis (1890) responded to the snapshot camera. Nissenbaum's **contextual integrity** (2004; 2010): flows follow the norms of the context where information arose. "It was already public" is weak, and aggregation can harm.

Notice-and-consent fails: McDonald and Cranor (2008) estimated about 200 hours a year to read policies. **Dark patterns** (Brignull, 2010). **Secondary use**: the Havasupai samples (1990; settled April 2010, $700,000). Royal Free and DeepMind (about 1.6 million patients; ICO ruling July 2017). Google and Ascension (about 50 million patients, 2019): probably lawful, yet trust collapsed. Lawful and appropriate are different tests.

De-identification is not anonymity. HIPAA's Safe Harbor (18 identifiers) and Expert Determination. Sweeney (1997) re-identified Governor Weld; her 87% uniqueness figure was recomputed by Golle (2006) at about 63%. AOL (2006; user 4417749, Thelma Arnold). Netflix (Narayanan and Shmatikov, 2008). Golden State Killer via GEDmatch (April 2018); 23andMe breach (October 2023, about 6.9 million profiles; 2025 bankruptcy). HIPAA leaves out apps, wearables, DTC genetics and data brokers; the FTC acted on GoodRx ($1.5 million, February 2023) and BetterHelp ($7.8 million, March 2023). After Dobbs (June 2022), cycle and location data gained legal risk.

Sources of bias (Suresh and Guttag): **historical, representation, measurement, aggregation, deployment**. Gender Shades (Buolamwini and Gebru, 2018): error about 0.8% for lighter-skinned men versus up to 34.7% for darker-skinned women. Robert Williams (January 2020; about 30

hours held; settlement June 2024) and Porcha Woodruff (February 2023). Obermeyer (Science, October 2019): spending as a proxy for need; fixing the label raised Black patients flagged from 17.7% to 46.5%.

Fairness has no single definition. ProPublica on COMPAS (May 2016): among non-reoffenders, about 45% of Black vs 23% of white defendants were labeled high risk; Northpointe answered that the score was **calibrated**. Kleinberg, Mullainathan and Raghavan, and Chouldechova proved that with different base rates, calibration and equal error rates cannot both hold. Criteria: anti-classification, calibration, error-rate balance, demographic parity, individual fairness. Most algorithmic harm is **disparate impact** (Griggs, 1971). Removing race leaves proxies and only removes your ability to audit.

Biased instruments feed models: pulse oximetry (Sjoding, December 2020: 11.7% vs 3.6% occult hypoxemia), race-free eGFR (September 2021). Government scale: the Dutch childcare benefits scandal (about 26,000 families; €2.75 million fine; cabinet resigned January 2021) and Australia's Robodebt (from 2016; about 470,000 unlawful debts; about A$1.8 billion; Royal Commission July 2023).

Responsibility: the **problem of many hands** (Thompson; Nissenbaum 1996) and the **responsibility gap** (Matthias, 2004). Therac-25 (June 1985 to January 1987; six overdoses, at least three deaths; race condition, removed hardware interlocks). Boeing 737 MAX (MCAS; single sensor; 346 deaths). Uber in Tempe (18 March 2018; emergency braking disabled; driver prosecuted, company not): Elish's **moral crumple zone**. The Vanderbilt vecuronium case (2017 death; 2022 conviction). **Automation bias** (Skitka): commission and omission errors. **Alert fatigue**: Joint Commission Sentinel Event Alert 50 (April 2013; 98 events, 80 deaths; 85 to 99% of alarms need no action). Epic Sepsis Model validation (Wong, June 2021): AUC 0.63 vs 0.76 to 0.83 claimed; missed about two thirds; alerts on 18%. Bainbridge's ironies of automation (1983).

Opacity (Burrell, 2016): secrecy, illiteracy, and real mismatch. Rudin (2019): build interpretable models for high stakes. Generative AI **confabulates**: Mata v. Avianca (June 2023); transcription fabrications (October 2024); Koko (January 2023) and NEDA's Tessa (June 2023). Governance: **GDPR** (25 May 2018; purpose limitation, data minimization, accountability; Article 22 on solely automated decisions; the right to explanation is contested). **EU AI Act** (in force 1 August 2024):

prohibited, high-risk, limited-risk tiers; prohibitions from 2 February 2025. FDA's IDx-DR (April 2018); NIST AI RMF (January 2023: govern, map, measure, manage); Illinois BIPA (2008). WHO guidance (2021). Model cards, datasheets, algorithmic impact assessments, local validation, a named person who can switch it off. Personal duties: no curiosity access, no patient posts, no identifiable data in consumer chatbots, verify before charting, document overrides.

LOCK THESE IN

- Technology is neither good nor bad, nor neutral: defaults are legislation.
- Collingridge: steerable when unknown, known when entrenched.
- Privacy, confidentiality and security are different; contextual integrity judges flows.
- De-identified is not anonymous; HIPAA misses most consumer health data.
- Bias enters through history, sampling, measurement, aggregation or deployment; proxies do damage.
- With unequal base rates, calibration and equal error rates cannot both hold.
- Responsibility slides to the last human: the moral crumple zone.
- Validate locally, audit by subgroup, name who can switch it off.

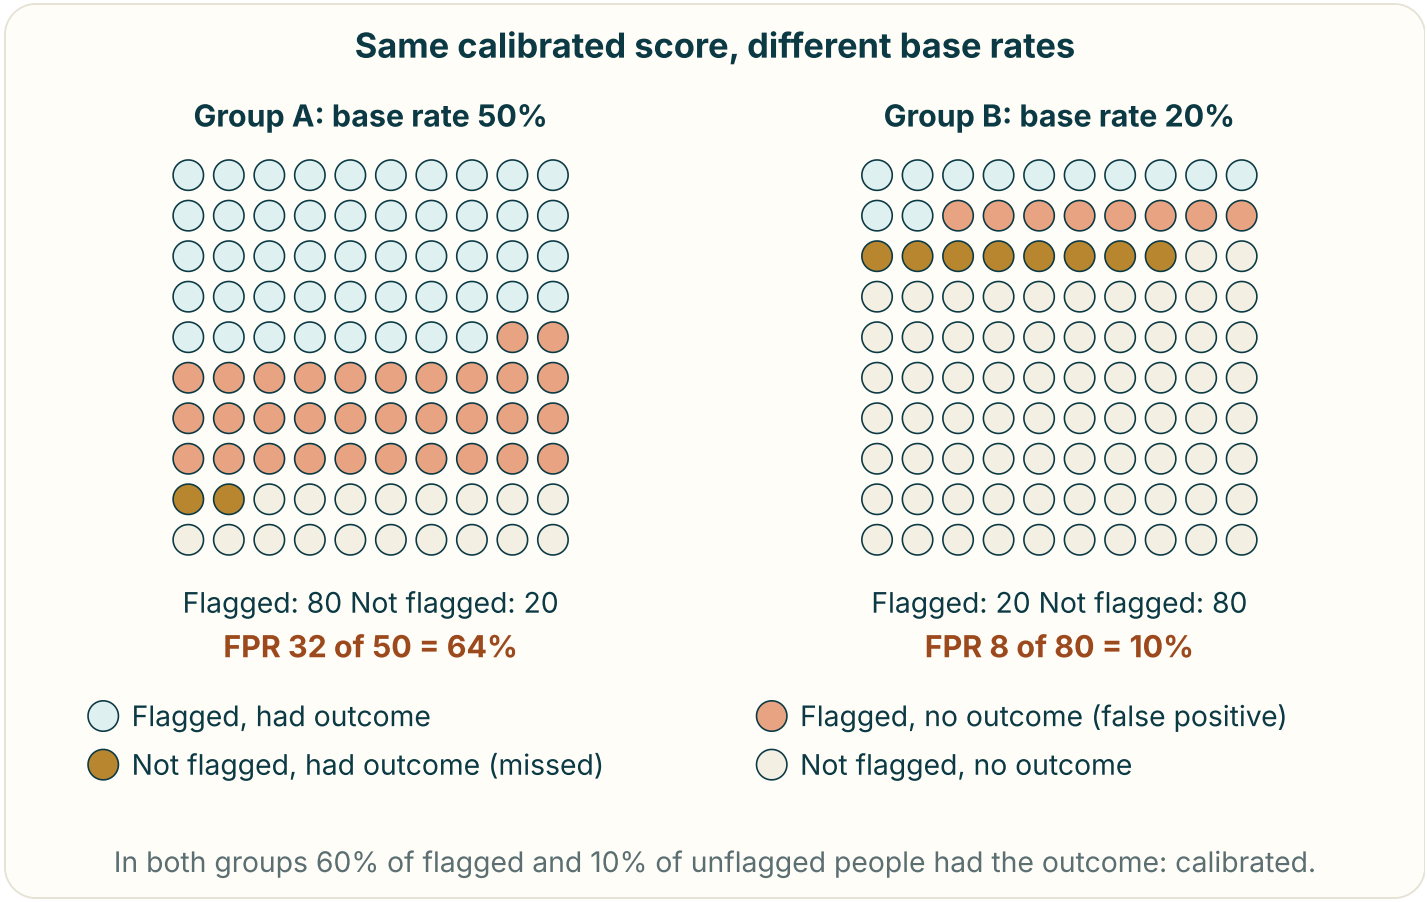

Calibrated, and still unequal. A score that means the same risk in both groups wrongly flags 64 percent of non-reoffenders in one group and 10 percent in the other, because the base rates differ. Choosing which unfairness to accept is a value judgment.

Fairness criteria

Criterion	What it demands	What it costs
Anti-classification	Protected attribute not an input	Proxies reinstate it; no audit possible
Calibration	Same score, same real risk in every group	Allows big gaps in who is wrongly flagged
Error-rate balance	Equal false-positive and false-negative rates	Conflicts with calibration when base rates differ
Demographic parity	Equal selection rates	Ignores real outcome differences
Individual fairness	Similar people treated similarly	Requires agreeing what "similar" means

Key terms and comparisons

Item	Meaning	Example
Affordance	What a design makes easy or hard	Copy-forward default
Policy vacuum	New capability with no settled rule	Deepfakes (Moor, 1985)
Contextual integrity	Flows must fit their context's norms	Drinking history to an employer
Secondary use	Data used beyond the consented purpose	Havasupai samples
Linkage attack	Re-identify by matching outside data	Sweeney and the voter roll
Proxy variable	Measurable stand-in for what matters	Cost for need (Obermeyer)
Moral crumple zone	Operator absorbs a system's blame	Tempe, 2018
Automation bias	Over-trust in automated cues	Errors of commission and omission
Alert fatigue	Many false alarms train dismissal	85 to 99% of alarms need no action
Model card	Intended use, data, subgroup performance	Mitchell et al., 2019

Worked cases

WORKED CASE 1 · CALIBRATION VERSUS EQUAL ERROR RATES

A risk score is calibrated in two groups: 60% of people flagged high have the outcome, and 10% of people flagged low have it. Group A's base rate is 50%; Group B's is 20%. Each group has 100 people. Find the false-positive rate in each.

STEP 1 · HOW MANY FLAGGED IN A

$0.6h + 0.1(100 - h) = 50$, so $0.5h = 40$ and $\mathbf{h = 80}$.

STEP 2 · HOW MANY FLAGGED IN B

$0.6h + 0.1(100 - h) = 20$, so $0.5h = 10$ and $\mathbf{h = 20}$.

STEP 3 · FALSE POSITIVES

A: 40% of 80 = 32 among 50 without the outcome, so FPR = **64%**. B: 40% of 20 = 8 among 80, so FPR = **10%**.

STEP 4 · LESSON

Calibrated, yet the error burden is very unequal. With different base rates you cannot have both; choosing is a value judgment to make openly.

WORKED CASE 2 · APPLYING CONTEXTUAL INTEGRITY

Patients signed an admission form allowing data use for "operations." The hospital now plans to send identifiable records to a tech firm to train a commercial product. Assess.

STEP 1 · THE ORIGINAL CONTEXT

Clinical care, where patients shared information to be treated.

STEP 2 · THE NEW FLOW

Identifiable records to a company for its own product.

STEP 3 · TEST THE NORMS

Would patients, told plainly, expect or accept this? Royal Free and Ascension suggest not.

STEP 4 · VERDICT

Possibly lawful, but a breach of contextual integrity. "They signed" answers a legal question, not an ethical one.

WORKED CASE 3 · FINDING THE PROXY

A readmission model flags patients for home visits using prior-year cost. Clinicians notice few flagged patients from one neighborhood with poor access. Diagnose.

STEP 1 · THE LABEL

Cost stands in for need: a **measurement** bias.

STEP 2 · THE MECHANISM

People with poor access generate less spending at the same level of illness, so they look healthier.

STEP 3 · THE FIX

Change the target to a direct measure of illness, as Obermeyer showed.

STEP 4 · DO NOT JUST DELETE VARIABLES

Removing neighborhood or race would not fix a biased label and would stop you auditing it.

WORKED CASE 4 · ASKING PROCUREMENT QUESTIONS

Your unit is offered a vendor's AI triage tool with an overall accuracy figure. What do you ask before go-live?

STEP 1 · THE LABEL
What exact outcome does it predict? Watch for cost or billing codes standing in for need.

STEP 2 · LOCAL PERFORMANCE
How does it perform on our patients, by race, language, insurance, age and disability?

STEP 3 · BURDEN
What is the alert rate per shift? Epic's sepsis model alerted on 18% of admissions.

STEP 4 · CONTROL
Who may override it safely, who can switch it off, and what is the review date? No owner is the clearest red flag.

TRAPS PEOPLE FALL INTO

- **"It's just a tool."** The design already distributes ease and risk.
- **Using privacy, confidentiality and security as synonyms.** They fail separately.
- **"They consented."** That is legal, not necessarily ethical.
- **Treating de-identified as anonymous.** Linkage attacks re-identify.
- **Assuming a health app is under HIPAA.** Most are not.
- **"We removed race, so it's fair."** Proxies remain; you lose the audit.
- **Expecting one fairness metric to satisfy all.** With different base rates it cannot.
- **Blaming the last pair of hands.** Look for the moral crumple zone.

Chapter 13 self-test · ten items

1. Kranzberg's first law says technology is what?

A. Neither good nor bad, nor neutral
B. Always good unless it is misused
C. Purely neutral in all its effects
D. Bad unless it is closely regulated

A. Neither good nor bad, nor neutral Designs shape behavior before any user chooses.

2. What is the Collingridge dilemma?

A. Laws always arrive before the technologies they will govern
B. Users always find ways to misuse new tools sooner or later
C. Early on you can steer but lack knowledge; later you know but cannot steer
D. Only technical experts can foresee a new technology's harms

C. Early on you can steer but lack knowledge; later you know but cannot steer Social media is its monument.

3. Nissenbaum's contextual integrity says privacy is violated when what happens?

A. Any piece of information becomes known to anyone at all
B. Data is stored on a server without any encryption
C. A person forgets to read a policy
D. Information flows breach the norms of its original context

D. Information flows breach the norms of its original context It explains why "already public" is a weak defense.

4. Sweeney re-identified Governor Weld by doing what?

A. Hacking into the hospital's records servers
B. Linking hospital data to a purchased voter roll
C. Interviewing hospital staff
D. Reading his public statements about his own medical care

B. Linking hospital data to a purchased voter roll Removing names does not anonymize.

5. Which is generally NOT covered by HIPAA?

A. A consumer cycle-tracking app
B. A hospital's billing records
C. A health plan's claims data
D. A lab's test results

A. A consumer cycle-tracking app Enforcement there has come from the FTC.

6. In Gender Shades, which group had the highest error rates?

A. Lighter-skinned men
B. Lighter-skinned women
C. Darker-skinned women
D. Darker-skinned men

C. Darker-skinned women Up to 34.7% versus about 0.8% for lighter-skinned men.

7. What did the impossibility results about COMPAS show?

A. COMPAS was neither calibrated nor accurate for any group
B. With different base rates, calibration and equal error rates conflict
C. All of the fairness criteria can always be met together
D. Removing race solves the problem

B. With different base rates, calibration and equal error rates conflict Choosing among criteria is a value judgment.

8. Elish's moral crumple zone describes what?

A. A car's structural safety design for absorbing crashes
B. A legal shield that protects software firms from lawsuits
C. A type of automated braking
D. A human operator absorbing blame for a system's failure

D. A human operator absorbing blame for a system's failure In Tempe, the safety driver was prosecuted and the company was not.

9. The external validation of the Epic Sepsis Model found what?

A. Performance that was better than the vendor had claimed
B. No alerts at all on any of the hospitalized patients
C. Perfect detection with high false alarms
D. AUC of 0.63 and about two thirds of cases missed

D. AUC of 0.63 and about two thirds of cases missed The vendor had reported 0.76 to 0.83.

10. Which practice does the EU AI Act prohibit outright?

A. Medical device AI
B. Credit scoring with oversight
C. Social scoring of citizens
D. Chatbots that disclose they are AI

C. Social scoring of citizens Prohibitions applied from 2 February 2025; medical devices are high-risk, not banned.

CHAPTER 14

Building an Ethical Practice

Read it in full: **Chapter 14 in the textbook**. Practice it in the lab: **The Four Topics Workbench** and **The Just Culture Console**.

What to know

Failure is usually not a failure of belief but of noticing, language and design. **Bounded ethicality** (Bazerman and Tenbrunsel): we fall short of our own standards without knowing. **Ethical fading** (Tenbrunsel and Messick, 2004): euphemism drains the moral content ("headcount action," "the bed"). Darley and Batson (1973): 63% of unhurried seminarians helped versus 10% who were late; the Good Samaritan topic barely mattered. Milgram: 26 of 40 (65%) to 450 volts. Asch (1951): about three quarters conformed at least once.

A correction to carry: the "sign at the top" honesty finding (Shu and colleagues, 2012) failed to replicate (Kristal, 2020) and was **retracted in September 2021**. Cite the retraction. The situationist challenge (Harman 1999; Doris 2002) is contested by Kamtekar (2004) and Annas. Hybrid view: character is real but weak, situations strong, so arrange the situation in advance.

Institutions manufacture drift. Gneezy and Rustichini (2000; ten Haifa daycares; fines in six): late pickups rose and stayed up: "A Fine is a Price." **Motivated blindness**; the slippery slope (Gino and Bazerman's 2009 lab study failed to replicate, but case records support erosion); **diffusion of responsibility**; the **hidden curriculum** (Jackson; Hafferty, 1994). Effective countermeasures move the decision earlier and make the right act the default.

The **four topics** method (Jonsen, Siegler and Winslade, 1982): **medical indications, patient preferences, quality of life, contextual features**, in that order. Indications first dissolve factual confusion; quality of life third because observers rate it lower than patients do (Albrecht and Devlieger's disability paradox); context last but present. It organizes facts; it does not decide.

The VHA's **IntegratedEthics** (2007): **CASES** consultation (clarify, assemble, synthesize, explain, support) and **ISSUES** preventive ethics (identify, study, select, undertake, evaluate, sustain). Write the question as one sentence: "Given these facts, is it permissible or obligatory for this agent to do this action?" If you cannot, it may be a communication, legal, clinical or personality problem instead.

Rest's four components (1983 to 1986): sensitivity, judgment, motivation, implementation. Diagnose before you prescribe: sensitivity needs exposure, judgment a framework, motivation protection and culture, implementation skills practice. Training usually targets judgment, often the

part that was not broken. A usable ethics note: question, participants, facts, options with reasons, recommendation and plan. Steelman; ask which layer (facts, values, authority, process) is in dispute; run the method before the decision (Haidt).

Moral courage in a hierarchy. Silence Kills (2005): over half saw dangerous shortcuts, fewer than one in ten spoke to the person. Korean Air 801 (6 August 1997; 228 of 254 died; the MSAW system near Guam was inhibited). Fischer and Orasanu: **mitigated speech**. Elaine Bromiley (29 March 2005; died 11 April): nurses had the tracheostomy set and an ICU bed; Martin Bromiley founded the Clinical Human Factors Group (2007).

Giving Voice to Values (Gentile, 2010): assume you know what is right; script, rehearse and pre-answer four rationalizations: standard practice, materiality, locus of responsibility, locus of loyalty. Grammar matters: own it with "I," name the observation, end with a small concrete step. **Crew resource management** (NASA workshop, June 1979; after Tenerife 1977 and Portland 1978) gave the **two-challenge rule**. TeamSTEPPS (2006) gave **CUS**: Concerned, Uncomfortable, Safety issue. Virginia Mason's Patient Safety Alert (2002). Edmondson (1999): better units reported more errors because they counted them: **psychological safety**.

Machinery. Quinlan (31 March 1976) cited Karen Teel's 1975 suggestion of ethics committees. Baby Doe (born 9 April 1982; died 15 April) led to the Child Abuse Amendments (1984). Joint Commission required an ethics mechanism by 1992. Fox, Myers and Pearlman (2007): consult services in 81% of US general hospitals, median three consults a year, only 41% of consultants formally trained. ASBH competencies (1998; 2011); HEC-C (2018). Consults are advisory. Schneiderman (JAMA, 2003; 551 ICU patients, seven hospitals): no mortality change, but fewer hospital, ICU and ventilator days among those who died. **Just culture** (Marx, 2001; Reason): human error (console), at-risk behavior (coach), reckless behavior (discipline); the **substitution test**; resist **outcome bias**. Bawa-Garba (Jack Adcock died 18 February 2011; erasure ordered January 2018, overturned August 2018). Enron waived its code; Treviño and Weaver (1999); ECI retaliation figures rose from about 22% (2013) to 44% (2017) to about 79% (2020).

Carrying it. **Burnout** (Maslach, 1981; ICD-11 occupational phenomenon, 2019): exhaustion, depersonalization, reduced accomplishment. **Compassion fatigue**. **Moral distress** (Jameton, 1984), **moral residue** (Webster and Baylis) and the **crescendo effect** (Epstein and Hamric, 2009).

Moral injury (Shay, 1994; Litz, 2009); Dean and Talbot (July 2018). The AACN **Four A's** (2004): ask, affirm, assess, act. Schwartz Rounds (Schwartz's essay 16 July 1995; Maben 2018: 25% to 12% poor wellbeing among attenders vs 37% to 34%). Rushton's moral resilience; the NAM report (2019); the Lorna Breen Act (18 March 2022). **Moral repair** (Walker, 2006); Michigan's disclosure program (2001; claims 7.03 to 4.52 per 100,000 encounters; lawsuits 2.13 to 0.75; resolution about 16 to 11 months; Kachalia 2010); CANDOR (2016). A lasting practice: red lines, a decision journal, a named ally, rehearsed scripts, scheduled reflection (Schön, 1983), a drift check, an exit number.

LOCK THESE IN

- Failure is usually situational: hurry, authority, conformity, euphemism, metrics, drift, diffusion.
- Cite the 2021 retraction, not the honesty-signature headline.
- Four topics in order: indications, preferences, quality of life, context.
- CASES: clarify first; ISSUES: fix the recurring system problem.
- Rest: sensitivity, judgment, motivation, implementation; diagnose before training.
- Escalate with scripts: CUS and the two-challenge rule make challenge policy.
- Just culture judges behavior, not outcome: console, coach, discipline.
- Distress, residue and injury are injuries of conditions; change the conditions.

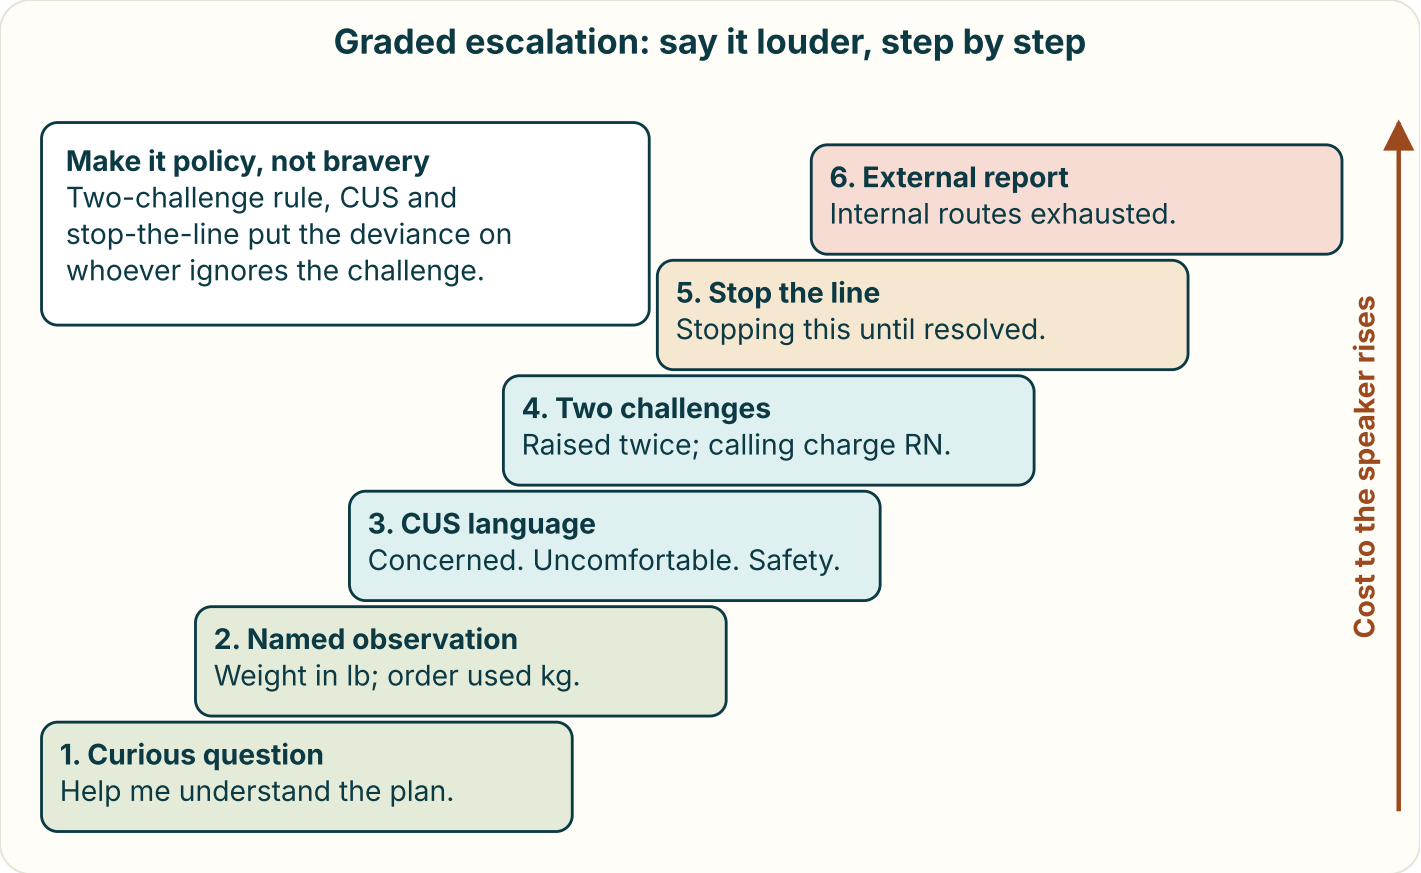

Policy carries people up the ladder. Each rung is a little louder and a little costlier. Scripts like CUS and the two-challenge rule make escalation routine instead of heroic.

Methods compared

Method	Origin	Best for	Weakness
Four topics	Jonsen, Siegler, Winslade 1982	Bedside cases	Organizes facts; does not adjudicate
Four principles	Beauchamp and Childress 1979	Naming what is at stake	No built-in ranking
CASES	VHA 2007	Formal consults	Needs trained consultants
ISSUES	VHA 2007	Recurring system problems	Slow; needs data and sponsors
Four components	Rest 1983 to 1986	Diagnosing failure	Descriptive, not normative
Giving Voice to Values	Gentile 2010	Acting once you know	Skips what is right

Key terms and comparisons

Item	Meaning	Response
Human error	Inadvertent slip, lapse or mistake	Console; redesign the process
At-risk behavior	Drift into an unrecognized risky shortcut	Coach; restore the barrier
Reckless behavior	Conscious disregard of substantial risk	Discipline; still examine the system
Substitution test	Would a peer have done the same?	If yes, it is a system problem
Burnout	Exhaustion from overload and low control	Staffing, schedule control, recovery
Moral distress	Know the right act, blocked	Remove the constraint; consult; escalate
Moral injury	Betrayal or violation of deep beliefs	Acknowledgment, repair, specialist care
Mitigated speech	Hedged requests to superiors	Scripts and graded language
Psychological safety	Team is safe for interpersonal risk	More reporting, more learning
Moral repair	Acknowledge, take responsibility, make right, change	Walker, 2006; CANDOR

Worked cases

WORKED CASE 1 · WORKING THE FOUR TOPICS

Mr. Osei, 82, has advanced heart failure. His daughter wants "everything done"; he has said he wants to go home. The unit is short of beds. Fill the four boxes in order.

STEP 1 · MEDICAL INDICATIONS

Diagnosis, prognosis and what each option (ICU, home with hospice) can realistically achieve.

STEP 2 · PATIENT PREFERENCES

He has expressed a wish to go home. Does he have capacity for this decision? If yes, his preference governs.

STEP 3 · QUALITY OF LIFE

By his judgment, not the team's; observers tend to rate it lower than patients do.

STEP 4 · CONTEXTUAL FEATURES

His daughter's distress, bed pressure, any conflicts. State them openly; do not let bed pressure pose as clinical reasoning.

THE ONE-SENTENCE QUESTION

Given his capacity and stated wish, is it permissible to discharge him home with hospice despite his daughter's objection?

WORKED CASE 2 · DIAGNOSING WITH REST

A nurse knew a colleague was charting assessments she had not done and said nothing for months. Which component failed, and what helps?

STEP 1 · SENSITIVITY

She noticed the issue: intact.

STEP 2 · JUDGMENT

She knew it was wrong: intact.

STEP 3 · MOTIVATION

She prioritized loyalty and fear over the moral value: **failed**.

STEP 4 · REMEDY

Not more teaching. Protection, a safe reporting route and a culture where speaking up is expected.

WORKED CASE 3 · BUILDING AN ESCALATION

During a handoff you see a potassium infusion ordered at a rate that looks too fast for the patient's weight. The resident brushes off your first comment. Climb the ladder.

STEP 1 · NAMED OBSERVATION

"The weight in the chart is in pounds and the order looks calculated in kilograms."

STEP 2 · CUS

"I'm concerned. I'm uncomfortable. This is a safety issue and I need us to stop."

STEP 3 · SECOND CHALLENGE

"I've raised this twice. Per policy I'm calling the charge nurse and the on-call attending."

STEP 4 · DOCUMENT

Record what you saw, what you said and when. The two-challenge rule makes persisting the policy.

WORKED CASE 4 · A JUST CULTURE RESPONSE

Two nurses skip the independent double-check on a high-alert infusion because it "slows things down." One patient is fine; one dies. How should leadership respond?

STEP 1 · CLASSIFY THE BEHAVIOR

A drift into a shortcut whose risk was underrated: **at-risk behavior**.

STEP 2 · SUBSTITUTION TEST

Would similar nurses on that unit, with that staffing, do the same? Probably yes: a system problem too.

STEP 3 · RESPOND IDENTICALLY

Coach both. Judging by the outcome is **outcome bias** and teaches that the lesson is luck.

STEP 4 · FIX THE BARRIER

Redesign the check so it survives a real shift, and remove the incentive to skip it.

TRAPS PEOPLE FALL INTO

- **"Good people don't need a method."** Situations overwhelm resolve; design them in advance.
- **Teaching the honesty-signature finding.** It was retracted in 2021.
- **Starting the four topics with context.** Facts first, so resources do not pose as clinical reasoning.
- **Answering the question as asked.** Clarify first; it is often the wrong question.
- **Training judgment by default.** Diagnose which of Rest's components failed.
- **Hinting instead of stating.** Mitigated speech fails; use I, the observation and a small step.
- **Judging errors by outcome.** Just culture judges behavior.
- **Offering only resilience training.** Moral distress needs the conditions changed.

Chapter 14 self-test · ten items

1. Tenbrunsel and Messick's "ethical fading" describes what?

A. Slowly forgetting one's own moral beliefs over the years
B. Euphemism draining the moral content from a decision
C. Losing interest in ethics over the course of a long career
D. Moral rules weakening as laws change

B. Euphemism draining the moral content from a decision "Headcount action" for a layoff is an example.

2. In Darley and Batson's study, what most affected whether seminarians stopped to help?

A. Whether they were in a hurry
B. Whether their talk was on the Good Samaritan
C. How religious they rated themselves
D. Whether they had helped someone before

A. Whether they were in a hurry 63% helped when unhurried, 10% when late.

3. Gneezy and Rustichini's daycare fine showed what?

A. A fine turned an obligation into a price, and lateness rose
B. Fines reliably reduced the number of late pickups
C. Parents ignored the fine entirely
D. Lateness fell at first, then rose again when the fines ended

A. A fine turned an obligation into a price, and lateness rose Late pickups stayed high even after the fine was withdrawn.

4. What is the correct order of the four topics?

A. Context, indications, preferences, quality of life
B. Preferences, indications, context, quality of life
C. Quality of life, context, preferences, indications
D. Indications, preferences, quality of life, context

D. Indications, preferences, quality of life, context Facts first; context last so it cannot pose as clinical reasoning.

5. What is the first step of the CASES approach?

A. Assemble all relevant facts
B. Clarify the consultation request
C. Synthesize a recommendation
D. Support the process afterward

B. Clarify the consultation request The question as asked is often the wrong question.

6. A clinician knew the right action and stayed silent. Which of Rest's components failed?

A. Sensitivity
B. Judgment
C. Motivation
D. Implementation

C. Motivation The fix is protection and culture, not more teaching.

7. What does the two-challenge rule do?

A. Allows two errors before discipline
B. Requires two clinician signatures on every medication order
C. Limits each nurse to two safety challenges per shift
D. Requires escalation after two unanswered safety challenges

D. Requires escalation after two unanswered safety challenges It makes persisting the policy and puts the deviance on whoever ignores it.

8. Edmondson found better-led units reported more errors because of what?

A. They counted errors rather than committing more
B. They hired less experienced and more junior staff
C. They had higher patient acuity
D. They used a different and broader definition of error

A. They counted errors rather than committing more This grew into the construct of psychological safety.

9. Under just culture, at-risk behavior calls for what response?

A. Consoling the nurse only
B. Discipline and a report to the licensing board
C. Coaching and restoring the barrier
D. No response if no harm occurred

C. Coaching and restoring the barrier Human error: console. At-risk: coach. Reckless: discipline.

10. What did Michigan's disclosure program show (Kachalia, 2010)?

A. Lawsuits rose sharply after staff began apologizing
B. Claims and lawsuits fell after honest disclosure
C. Disclosure had no measurable effect on new claims
D. Only payouts rose while claims fell

B. Claims and lawsuits fell after honest disclosure Claims fell from 7.03 to 4.52 per 100,000 encounters; a single-center before-and-after study.

ONE PAGE

Cram sheet

Everything from chapters 13 and 14 in one place. Cover the right column and see if you can rebuild each row.

Names and dates

Who or what	When	Point
Kranzberg	1985 (pub. 1986)	Neither good nor bad, nor neutral
Collingridge	1980	Steerable early, entrenched later
Warren and Brandeis	1890	The right to privacy
Nissenbaum	2004; 2010	Contextual integrity
Sweeney	1997	Weld re-identified via voter roll
Obermeyer	October 2019	Cost proxy; 17.7% to 46.5%
ProPublica, COMPAS	May 2016	Calibrated but unequal error rates
Elish	2010s	Moral crumple zone
GDPR; EU AI Act	25 May 2018; 1 August 2024	Purpose limitation; risk tiers
Jonsen, Siegler, Winslade	1982	Four topics
VHA IntegratedEthics	2007	CASES and ISSUES
Gentile	2010	Giving Voice to Values
Marx; Reason	2001	Just culture; substitution test
Jameton; Epstein and Hamric	1984; 2009	Moral distress; crescendo effect

Lists to know

List	Items
Bias sources	Historical, representation, measurement, aggregation, deployment
Fairness criteria	Anti-classification, calibration, error-rate balance, parity, individual
Four topics	Indications, preferences, quality of life, context
CASES	Clarify, assemble, synthesize, explain, support
Rest	Sensitivity, judgment, motivation, implementation
GVV rationalizations	Standard practice, materiality, responsibility, loyalty
Just culture	Error: console; at-risk: coach; reckless: discipline
Four A's	Ask, affirm, assess, act

Numbers

Item	Value
Gender Shades error	About 0.8% vs up to 34.7%
COMPAS false high risk	About 45% Black vs 23% white non-reoffenders
Epic Sepsis Model	AUC 0.63; alerts on 18% of admissions
Alarms not needing action	About 85 to 99%
Darley and Batson	63% vs 10% helped
Michigan claims	7.03 to 4.52 per 100,000 encounters

Dr. Garg Learning · CC BY-SA 4.0