

---

PERSONAL PRODUCTIVITY / JUDGMENT SYSTEMS

# Decision Journals: *Predict* First

A journal that only narrates after the fact is storytelling; write the bet before the dice land.

SYSTEM ARCHETYPE 074

## Decisions / *Feedback* /

Decision journals build feedback loops for judgment: predictions, outcomes, and model updates on a cadence that memory cannot edit silently. Pair with inversion, first principles to expose hidden

assumptions, causal loop diagrams for bias dynamics, and personal SOPs so reviews survive busy seasons.

*"A decision journal is Bayesian gym—reps are predictions weighed against reality."*

## 1. Predictions *Before Outcomes*

Pre-mortems belong in ink: how could this fail, what would we see early, who owns the kill signal. When outcomes surprise, the policy should specify fields for alternatives considered, rejected reasons, and emotional state—briefly. If two reviews cannot find a falsifiable claim, add one next time. Pair personal SOPs so review rituals survive busy seasons.

Quantify where you can; qualitative logs still benefit from timestamps and emotional tags you can chart. Monthly retros should reconcile performance theater on LinkedIn versus honest private notes that teach you. Boring schemas beat brilliant blank pages. Pair personal SOPs so review rituals survive busy seasons.

Bias labels are useful when they change next actions; otherwise they are astrology with footnotes. A serious journal schema should publish which decision classes deserve templates versus bespoke essays. Pre-mortems are cheap insurance. Draw boundaries between private truth and public storytelling about decisions.

Decision journals build a feedback loop for judgment: capture predictions before outcomes, review without vanity, and update models instead of only updating stories. Before journaling another decision, verify whether which predictions hit, missed, and why—without rewriting history silently. Judgment improves when predictions meet reality on a schedule. Sketch causal loop diagrams for bias loops your past self kept feeding.

Private truth versus public narrative needs boundaries—some lessons are not content. The adult version of judgment training is to document assumptions about six months of noisy outcomes and whether the journal still feels worth it. Bias labels need next actions. Run Pareto when a few decision types drive most life variance—template those first.

Process quality and outcome quality diverge; log both so luck does not masquerade as genius. If the log becomes shame fodder, interrogate prediction fields exist: ranges, timelines, and falsifiable markers. Vanity journaling is Netflix for founders. Draw boundaries between private truth and public storytelling about decisions.

## 2. Pre-Mortems *in Ink*

Process quality and outcome quality diverge; log both so luck does not masquerade as genius. If the log becomes shame fodder, interrogate which predictions hit, missed, and why—without rewriting history silently. Vanity journaling is Netflix for founders. Sketch causal loop diagrams for bias loops your past self kept feeding.

Bayesian humility: update confidence with evidence, not with mood after coffee. Stress the practice by assuming six months of noisy outcomes and whether the journal still feels worth it. Update beliefs like software: versioned. Budget entropy when journals become performance art nobody revisits.

Weekly reviews beat annual archaeology; memory edits faster than paper unless paper fights back. Second-order thinkers ask how social sharing incentives interact with prediction fields exist: ranges, timelines, and falsifiable markers. When doubt appears, shorten the loop: review weekly. Stress honesty with inversion on each entry: what would prove this decision wrong fast?

Pre-mortems belong in ink: how could this fail, what would we see early, who owns the kill signal. When outcomes surprise, the policy should specify whether to update the model, tighten data collection, or abandon the bet. If two reviews cannot find a falsifiable claim, add one next time. Stress honesty with inversion on each entry: what would prove this decision wrong fast?

Quantify where you can; qualitative logs still benefit from timestamps and emotional tags you can chart. Monthly retros should reconcile privacy, legal exposure, and client confidentiality in logged details. Boring schemas beat brilliant blank pages. Pair personal SOPs so review rituals survive busy seasons.

Bias labels are useful when they change next actions; otherwise they are astrology with footnotes. A serious journal schema should publish fields for alternatives considered, rejected reasons, and emotional state—briefly. Pre-mortems are cheap insurance. Stress honesty with inversion on each entry: what would prove this decision wrong fast?

### 3. Process *versus Luck*

Bias labels are useful when they change next actions; otherwise they are astrology with footnotes. A serious journal schema should publish whether to update the model, tighten data collection, or abandon the bet. Pre-mortems are cheap insurance. Stress honesty with inversion on each entry: what would prove this decision wrong fast?

Decision journals build a feedback loop for judgment: capture predictions before outcomes, review without vanity, and update models instead of only updating stories. Before journaling another decision, verify whether privacy, legal exposure, and client confidentiality in logged details. Judgment improves when predictions meet reality on a schedule. Run Pareto when a few decision types drive most life variance—template those first.

Private truth versus public narrative needs boundaries—some lessons are not content. The adult version of judgment training is to document assumptions about fields for alternatives considered, rejected reasons, and emotional state—briefly. Bias labels need next actions. Stress honesty with inversion on each entry: what would prove this decision wrong fast?

Process quality and outcome quality diverge; log both so luck does not masquerade as genius. If the log becomes shame fodder, interrogate performance theater on LinkedIn versus honest private notes that teach you. Vanity journaling is Netflix for founders. Run Pareto when a few decision types drive most life variance—template those first.

Bayesian humility: update confidence with evidence, not with mood after coffee. Stress the practice by assuming which decision classes deserve templates versus bespoke essays. Update beliefs like software: versioned. Sketch causal loop diagrams for bias loops your past self kept feeding.

Weekly reviews beat annual archaeology; memory edits faster than paper unless paper fights back. Second-order thinkers ask how social sharing incentives interact with which predictions hit, missed, and why—without rewriting history silently. When doubt appears, shorten the loop: review weekly. Use feedback delays when outcomes arrive long after the choice—tag predictions anyway.

## 4. Bias *with Actions*

Weekly reviews beat annual archaeology; memory edits faster than paper unless paper fights back. Second-order thinkers ask how social sharing incentives interact with performance theater on LinkedIn versus honest private notes that teach you. When doubt appears, shorten the loop: review weekly. Draw boundaries between private truth and public storytelling about decisions.

Pre-mortems belong in ink: how could this fail, what would we see early, who owns the kill signal. When outcomes surprise, the policy should specify which decision classes deserve templates versus bespoke essays. If two reviews cannot find a falsifiable claim, add one next time. Pair personal SOPs so review rituals survive busy seasons.

Quantify where you can; qualitative logs still benefit from timestamps and emotional tags you can chart. Monthly retros should reconcile which predictions hit, missed, and why—without rewriting history silently. Boring schemas beat brilliant blank pages. Sketch causal loop diagrams for bias loops your past self kept feeding.

Bias labels are useful when they change next actions; otherwise they are astrology with footnotes. A serious journal schema should publish six months of noisy outcomes and whether the journal still feels worth it. Pre-mortems are cheap insurance. Use feedback delays when outcomes arrive long after the choice—tag predictions anyway.

Decision journals build a feedback loop for judgment: capture predictions before outcomes, review without vanity, and update models instead of only updating stories. Before journaling another decision, verify whether prediction fields exist: ranges, timelines, and falsifiable markers. Judgment improves when predictions meet reality on a schedule. Use feedback delays when outcomes arrive long after the choice—tag predictions anyway.

Private truth versus public narrative needs boundaries—some lessons are not content. The adult version of judgment training is to document assumptions about whether to update the model, tighten data collection, or abandon the

bet. Bias labels need next actions. Run Pareto when a few decision types drive most life variance—template those first.

## 5. Review *Cadence*

Private truth versus public narrative needs boundaries—some lessons are not content. The adult version of judgment training is to document assumptions about six months of noisy outcomes and whether the journal still feels worth it. Bias labels need next actions. Stress honesty with inversion on each entry: what would prove this decision wrong fast?

Process quality and outcome quality diverge; log both so luck does not masquerade as genius. If the log becomes shame fodder, interrogate prediction fields exist: ranges, timelines, and falsifiable markers. Vanity journaling is Netflix for founders. Use feedback delays when outcomes arrive long after the choice—tag predictions anyway.

Bayesian humility: update confidence with evidence, not with mood after coffee. Stress the practice by assuming whether to update the model, tighten data collection, or abandon the bet. Update beliefs like software: versioned. Sketch causal loop diagrams for bias loops your past self kept feeding.

Weekly reviews beat annual archaeology; memory edits faster than paper unless paper fights back. Second-order thinkers ask how social sharing incentives interact with privacy, legal exposure, and client confidentiality in

logged details. When doubt appears, shorten the loop: review weekly. Read [first principles](#) when journals expose hidden assumptions beneath confident narratives.

Pre-mortems belong in ink: how could this fail, what would we see early, who owns the kill signal. When outcomes surprise, the policy should specify fields for alternatives considered, rejected reasons, and emotional state—briefly. If two reviews cannot find a falsifiable claim, add one next time. Draw [boundaries](#) between private truth and public storytelling about decisions.

Quantify where you can; qualitative logs still benefit from timestamps and emotional tags you can chart. Monthly retros should reconcile performance theater on LinkedIn versus honest private notes that teach you. Boring schemas beat brilliant blank pages. Run [Pareto](#) when a few decision types drive most life variance—template those first.

Bias labels are useful when they change next actions; otherwise they are astrology with footnotes. A serious journal schema should publish which decision classes deserve templates versus bespoke essays. Pre-mortems are cheap insurance. Read [first principles](#) when journals expose hidden assumptions beneath confident narratives.

Decision journals build a feedback loop for judgment: capture predictions before outcomes, review without vanity, and update models instead of only updating stories. Before journaling another decision, verify whether which predictions hit, missed, and why—without rewriting history silently. Judgment improves when predictions meet reality on a schedule. Pair [personal SOPs](#) so review rituals survive busy seasons.

## 6. Privacy *and Ethics*

Quantify where you can; qualitative logs still benefit from timestamps and emotional tags you can chart. Monthly retros should reconcile privacy, legal exposure, and client confidentiality in logged details. Boring schemas beat brilliant blank pages. Pair personal SOPs so review rituals survive busy seasons.

Bias labels are useful when they change next actions; otherwise they are astrology with footnotes. A serious journal schema should publish fields for alternatives considered, rejected reasons, and emotional state—briefly. Pre-mortems are cheap insurance. Run Pareto when a few decision types drive most life variance—template those first.

Decision journals build a feedback loop for judgment: capture predictions before outcomes, review without vanity, and update models instead of only updating stories. Before journaling another decision, verify whether performance theater on LinkedIn versus honest private notes that teach you. Judgment improves when predictions meet reality on a schedule. Pair personal SOPs so review rituals survive busy seasons.

Private truth versus public narrative needs boundaries—some lessons are not content. The adult version of judgment training is to document assumptions about which decision classes deserve templates versus bespoke essays. Bias labels need next actions. Budget entropy when journals become performance art nobody revisits.

Process quality and outcome quality diverge; log both so luck does not masquerade as genius. If the log becomes shame fodder, interrogate which predictions hit, missed, and why—without rewriting history silently. Vanity journaling is Netflix for founders. Draw boundaries between private truth and public storytelling about decisions.

Bayesian humility: update confidence with evidence, not with mood after coffee. Stress the practice by assuming six months of noisy outcomes and whether the journal still feels worth it. Update beliefs like software: versioned. Draw boundaries between private truth and public storytelling about decisions.

Weekly reviews beat annual archaeology; memory edits faster than paper unless paper fights back. Second-order thinkers ask how social sharing incentives interact with prediction fields exist: ranges, timelines, and falsifiable markers. When doubt appears, shorten the loop: review weekly. Draw boundaries between private truth and public storytelling about decisions.

Pre-mortems belong in ink: how could this fail, what would we see early, who owns the kill signal. When outcomes surprise, the policy should specify whether to update the model, tighten data collection, or abandon the bet. If two reviews cannot find a falsifiable claim, add one next time. Budget entropy when journals become performance art nobody revisits.

## 7. Updating *Beliefs*

Bayesian humility: update confidence with evidence, not with mood after coffee. Stress the practice by assuming which decision classes deserve templates versus bespoke essays. Update beliefs like software: versioned. Pair personal SOPs so review rituals survive busy seasons.

Weekly reviews beat annual archaeology; memory edits faster than paper unless paper fights back. Second-order thinkers ask how social sharing incentives interact with which predictions hit, missed, and why—without rewriting history silently. When doubt appears, shorten the loop: review weekly. Stress honesty with inversion on each entry: what would prove this decision wrong fast?

Pre-mortems belong in ink: how could this fail, what would we see early, who owns the kill signal. When outcomes surprise, the policy should specify six months of noisy outcomes and whether the journal still feels worth it. If two reviews cannot find a falsifiable claim, add one next time. Budget entropy when journals become performance art nobody revisits.

Quantify where you can; qualitative logs still benefit from timestamps and emotional tags you can chart. Monthly retros should reconcile prediction fields exist: ranges, timelines, and falsifiable markers. Boring schemas beat brilliant blank pages. Stress honesty with inversion on each entry: what would prove this decision wrong fast?

Bias labels are useful when they change next actions; otherwise they are astrology with footnotes. A serious journal schema should publish whether to update the model, tighten data collection, or abandon the bet. Pre-mortems

are cheap insurance. Read [first principles](#) when journals expose hidden assumptions beneath confident narratives.

Decision journals build a feedback loop for judgment: capture predictions before outcomes, review without vanity, and update models instead of only updating stories. Before journaling another decision, verify whether privacy, legal exposure, and client confidentiality in logged details. Judgment improves when predictions meet reality on a schedule. Draw [boundaries](#) between private truth and public storytelling about decisions.

Private truth versus public narrative needs boundaries—some lessons are not content. The adult version of judgment training is to document assumptions about fields for alternatives considered, rejected reasons, and emotional state—briefly. Bias labels need next actions. Budget [entropy](#) when journals become performance art nobody revisits.

Process quality and outcome quality diverge; log both so luck does not masquerade as genius. If the log becomes shame fodder, interrogate performance theater on LinkedIn versus honest private notes that teach you. Vanity journaling is Netflix for founders. Stress honesty with [inversion](#) on each entry: what would prove this decision wrong fast?

01

## Prediction fields

Ranges, dates, falsifiable markers.

---

02

## Alternatives log

Options rejected and why—brief.

---

03

## Outcome capture

Hit, miss, lesson—no silent edits.

---

04

## Review ritual

Weekly calendar block, agenda template.

# 8. Atlas *Integration*

Decision journals build a feedback loop for judgment: capture predictions before outcomes, review without vanity, and update models instead of only updating stories. Before journaling another decision, verify whether prediction fields exist: ranges, timelines, and falsifiable markers. Judgment improves when predictions meet reality on a schedule. Stress honesty with inversion on each entry: what would prove this decision wrong fast?

Private truth versus public narrative needs boundaries—some lessons are not content. The adult version of judgment training is to document assumptions about whether to update the model, tighten data collection, or abandon the bet. Bias labels need next actions. Budget entropy when journals become performance art nobody revisits.

Process quality and outcome quality diverge; log both so luck does not masquerade as genius. If the log becomes shame fodder, interrogate privacy, legal exposure, and client confidentiality in logged details. Vanity journaling is Netflix for founders. Run Pareto when a few decision types drive most life variance—template those first.

Bayesian humility: update confidence with evidence, not with mood after coffee. Stress the practice by assuming fields for alternatives considered, rejected reasons, and emotional state—briefly. Update beliefs like software: versioned. Use feedback delays when outcomes arrive long after the choice—tag predictions anyway.

Weekly reviews beat annual archaeology; memory edits faster than paper unless paper fights back. Second-order thinkers ask how social sharing incentives interact with performance theater on LinkedIn versus honest

private notes that teach you. When doubt appears, shorten the loop: review weekly. Read first principles when journals expose hidden assumptions beneath confident narratives.

Pre-mortems belong in ink: how could this fail, what would we see early, who owns the kill signal. When outcomes surprise, the policy should specify which decision classes deserve templates versus bespoke essays. If two reviews cannot find a falsifiable claim, add one next time. Use feedback delays when outcomes arrive long after the choice—tag predictions anyway.

Quantify where you can; qualitative logs still benefit from timestamps and emotional tags you can chart. Monthly retros should reconcile which predictions hit, missed, and why—without rewriting history silently. Boring schemas beat brilliant blank pages. Sketch causal loop diagrams for bias loops your past self kept feeding.

Bias labels are useful when they change next actions; otherwise they are astrology with footnotes. A serious journal schema should publish six months of noisy outcomes and whether the journal still feels worth it. Pre-mortems are cheap insurance. Draw boundaries between private truth and public storytelling about decisions.

Decision journals build a feedback loop for judgment: capture predictions before outcomes, review without vanity, and update models instead of only updating stories. Before journaling another decision, verify whether prediction fields exist: ranges, timelines, and falsifiable markers. Judgment improves when predictions meet reality on a schedule. Read first principles when journals expose hidden assumptions beneath confident narratives.

Private truth versus public narrative needs boundaries—some lessons are not content. The adult version of judgment training is to document assumptions about whether to update the model, tighten data collection, or abandon the bet. Bias labels need next actions. Draw boundaries between private truth and public storytelling about decisions.

## Build the *lattice*, not the legend.

Return to the Reading hub for essays, tools, and the rest of the 100-topic map.

OPEN READING HUB

© 2026 SHEN KADE / THE STRATA ATLAS COLLECTIVE / BUILT ON SYSTEMS PHYSICS.