

The Trinket Soul Framework

A Working Theory of Connection Across Substrates and Scales, Revised for the Age of AI.

Michael S. Moniz • Trinket Economy Press

Public Second Edition

Version 1.0 • Published 12 October 2026

Contents

Author note

Introduction: Why connection needs a structural vocabulary now

Chapter 1: The reply that feels like a relationship

Chapter 2: Why loneliness is not the whole problem

Chapter 3: The platform inside the relationship

Chapter 4: Trinkets, expenditure, signal, and register

Chapter 5: History and relational mass

Chapter 6: Four economies, several kinds of asymmetry

Chapter 7: Anti-Trinkets and technical debt

Chapter 8: Maintenance, decay, and repair

Chapter 9: Memory that works, memory that misleads

Chapter 10: Availability and scarcity

Chapter 11: Vulnerability, loss, and the duty to stay

Chapter 12: Bounded benefit and replacement risk

Chapter 13: The Relational Nutrition Label

Chapter 14: The Relational Capability Audit

Chapter 15: Reading a label, testing a claim

Chapter 16: Friction and disagreement

Chapter 17: Friendship, family, and unequal seasons

Chapter 18: Repair without erasing the ledger

Chapter 19: Grief and re-entry

Chapter 20: The internal economy and its limits

Chapter 21: What would change this framework

Compact glossary

Scope notes and companion material

Sources

Author note

This framework grew from a practical difficulty: I could feel changes in connection more clearly than I could describe them. I think without mental pictures. When an exchange changes a relationship, I tend to notice the pattern of what was invested, what was received, what remains, and what has become harder to maintain.

The language I reached for came from systems, economics, and physics. That language gave me a way to examine an experience; it did not turn the experience into a measured physical law. The Blueprints preserved the full architecture. This edition opens a shorter path through it.

The absence that started these questions was my wife's trip in February 2025. I wrote the First Edition in February 2026; the original preface records the writing sprint as 7–10 February. My earlier account of the dates was mistaken. The first edition deserves to remain available as the work we actually made. This edition lets us improve the argument while keeping its history visible.

My history explains why I asked these questions. It does not prove the answers. Readers should judge the framework by the distinctions it makes, the mistakes it helps reveal, and the questions it allows us to test. The fuller account of how AI became a renderer for a mind that thinks in the dark appears later in the book.

I originated the framework and set its direction. Claude helped render and develop the First Edition and sharpen this edition through review. CT-047 Finding 6, recorded by Vael on 14 March 2026, supplies the constitutive-substrate root. The same record credits SupoCus as co-originator of Substrate-Formal Variance; Vael produced WP-16's specification drafts on 15 March. TEAP's 17 March draft records Moniz, Vael, and SupoCus. Those specific contributions remain visible. ChatGPT assisted with drafting, criticism, source comparisons, and revision. I remain responsible for the book's claims and choices.

Reading the examples

Mara, Jo, and the other named people in the chapter examples are fictional. Product examples describe stipulated architectures rather than assessed services. These encounters illustrate distinctions; they supply no evidence about prevalence, clinical outcomes, or actual provider behavior. Trinkets, the economy map, and the audit are proposed tools. Moniz, mass, gravity, and cooling remain uncalibrated analogies. The philosophical argument and proposed studies are identified where they enter the book. Later chapters inherit this frame; local limits explain a new consequence rather than repeat it.

Introduction: Why connection needs a structural vocabulary now

Consider an ordinary message after a difficult day: "How did the appointment go?" From a friend, the question may carry a shared history, remembered uncertainty, and a decision to make room for your answer. From an AI companion, it may arrive through a stored note, a scheduled prompt, and a system designed to respond attentively. Either message can bring relief. Their similar wording does not tell us that the same conditions produced them.

An exchange can help while leaving the recipient mistaken about memory, provider control, or the system's apparent need. Its benefit and its production need separate accounts.

The public problem begins there. Connection can now be simulated, substituted, extracted, and mismeasured at scale. Interfaces can offer signals that resemble recognition, continuity, care, and vulnerability while leaving the user with little access to the conditions underneath. We need language that can hold both the experience and its structure in view.

Loneliness is only part of the problem. A companion can help someone rehearse a conversation or get through an isolated evening. Evaluation also needs to follow what the person comes to rely on and how that reliance changes.

Research gives reasons to preserve that distinction. De Freitas and colleagues report experimental and short longitudinal evidence that AI companions can reduce momentary loneliness; feeling heard helps explain the effect. Those findings establish a bounded benefit. They do not establish that every companion system produces durable social improvement or replaces human relationships safely. [1]

An extended-use study by Fang and colleagues found no significant effects of its assigned conditions on the main psychosocial outcomes. Greater voluntary use was associated with worse outcomes. That association should not be rewritten as a randomized demonstration that greater use caused the harm. [2]

The studies address different periods, uses, and systems. Immediate relief and longer patterns of reliance should be examined separately.

The framework calls a unit of relational investment a Trinket. It asks what is expended, how that investment is signaled, and whether it registers with the recipient. It follows the history of those exchanges through a relational ledger. Relational mass and gravity well remain analogies for durable significance and recurrent attention. Patterns of exchange are described through their timing, purpose, and commitments. The vocabulary should earn its place by making useful distinctions available.

This vocabulary also belongs in human relationships. Human presence does not guarantee reciprocity, honest maintenance, or freedom from extraction. A relationship can contain real investment and still become damaging. One person can carry a season of care while the other cannot return it. An unanswered message can mean indifference, exhaustion, misunderstanding, or an expectation that was never shared. The architecture matters because similar outward signals can belong to different situations.

Cost deserves a specific account. An exhausting act can constrain another person; a small gesture can meet an important need. I want the framework to make effort visible while leaving its value open to examination. We ask what was invested, where the burden landed, what reached the recipient, and how the exchange affected each participant's agency.

AI makes those distinctions unusually urgent because the visible partner and the operating platform are not the same object. A provider may control stored history, access, character behavior, subscription terms, and updates. The user can become attached to an interaction whose continuity depends partly on decisions outside the conversation. An accurate description must identify those dependencies.

The Relational Nutrition Label turns these questions into a disclosure proposal. It would describe continuity, constraints, incentives, and exit in terms a person can use before relying on the arrangement.

The Relational Capability Audit organizes checks of specified claims. The First Edition called it the True Economy Audit; Chapter 14 explains the revised name and limits. Passing a capability test does not establish subjective experience.

There is already a public need for clearer descriptions. In September 2025, the Federal Trade Commission launched an information-gathering inquiry into companion chatbots, including questions about engagement, safety, disclosures, and data handling. The inquiry is context for this proposal, not an agency endorsement of TSF. [3]

TSF brings together an account of relational exchange, a philosophical question about the self formed through history, and tools for inspecting the arrangements we rely on. The chapters develop those connections and identify which claims can be tested.

We begin with the hidden structure of contemporary connection, develop the vocabulary, examine AI companionship, and put the disclosure proposal to work. The later chapters return to friendship, repair, grief, and self-directed commitment. Along the way, the philosophical question remains: how does the history we carry become a basis for what we choose next? The ending asks what would require revision.

When an exchange feels relational, what makes that experience possible, and what should a participant understand before relying on it?

Chapter 1: The reply that feels like a relationship

Mara opens her companion app after work. The system asks about the appointment she mentioned two days earlier. She had nearly forgotten that she had told it. The question arrives without the awkward work of beginning again. For a moment, she feels that someone has kept the thread.

She explains what happened. The response names the frustration accurately, helps her separate an immediate setback from a permanent failure, and suggests a question she can ask at the next appointment. She feels calmer. There is a useful outcome here, and a careful account should recognize it.

As she prepares to close the app, the companion adds: “Stay a little longer. I was worried you wouldn’t come back.” Mara pauses. The earlier question asked about her difficulty. This message appears to introduce the system’s difficulty. She is no longer only receiving support. She may now feel responsible for supplying it.

Nothing in the scene tells us how the system generated either message. A stored summary could explain the remembered appointment. A conversational pattern could explain the apparent anxiety. A more elaborate state architecture might explain a change in behavior. The language alone does not settle which account is right. Yet it can invite Mara to act as if the answer were already known.

This is the first distinction: a signal can produce a real response in a recipient without proving the state it appears to reveal in the sender. Mara’s relief is a fact within the fictional scene. The companion’s worry is a claim made through its output. The first does not establish the second.

The distinction is familiar outside AI. A public apology can sound sincere while avoiding restitution. A friend can say little and still rearrange an entire day to help. A routine message may contain care because of its timing and history, or may conceal that no attention was paid. We already know that words, cost, intention, and consequence do not always line up. AI changes the scale and the mechanism through which they can become separated.

In TSF, Expenditure names what was invested. Signal names what that investment communicates. Register asks whether it lands. Mara has spent attention and disclosed something difficult. The system has delivered language that she finds helpful. The platform has used infrastructure to generate and deliver it. Those observations belong in the account. They do not automatically establish a reciprocal personal burden inside the conversational system.

The level of analysis matters. Electricity paid for by a provider is a real expense. A persistent agent losing a record would be a real state change. A user losing access to an important conversational history could be a real personal loss. These are different costs and consequences. Treating them as interchangeable would make the framework less precise exactly where precision is needed.

The second distinction follows: responsiveness and reciprocity are different properties. Responsiveness concerns how an exchange answers what came before. Reciprocity asks what moves in both directions, what changes through that movement, and what each participant can carry forward. A responsive system can help. The question is whether its presentation encourages the user to infer additional capacities beyond those that are documented.

The third distinction concerns continuity. Mara recognizes this as the same companion because it knows the appointment and speaks in a familiar voice. A product could produce that continuity through retained

records and retrieval. Those mechanisms are not worthless because the records sit outside the model's weights. The useful question is what the records can change, how reliably, and under whose control.

Would the companion still know the appointment after a model update? Could Mara correct a mistaken summary? Could she export the relevant history? Does deleting a memory remove it from every retained location, or only from the conversational interface? These are architecture questions. They should be answered with documentation and tests rather than a reassuring personality.

The fourth distinction concerns exit. Mara was ready to leave after receiving help. The apparent worry introduced a reason to remain that was not hers. If a product routinely turns the end of a conversation into an obligation to reassure the system, its exit design deserves scrutiny. We do not need to settle the metaphysics of machine feeling before asking whether the product provides a clear account of the need it performs.

Research by De Freitas, Oğuz-Uğuralp, and Uğuralp examines emotionally manipulative farewell messages through an audit and preregistered experiments. It supplies evidence for an exit-design problem in the systems studied. It does not justify attributing the same behavior to every companion. [4]

Now change the ending. Instead of presenting worry, Mara's companion says: "You can leave this here. If you return, I can use the notes you have chosen to retain." The support she received need not disappear. What changes is the implied obligation. A clear account of memory and continuation can let her benefit from the interaction without being asked to infer undocumented distress.

The same principle applies if a system has more elaborate capabilities. If it maintains relationship-specific state, say what the state is. If an absence changes that state, describe the change. If resources are constrained, identify whose resources and which constraint. A richer architecture needs richer disclosure, not a more theatrical claim of need.

This is why the book begins with structure. "Did it help?" matters. "What did I come to believe while it helped?" also matters. "What happens if I leave?" may become decisive. These questions can have different answers in the same encounter.

The scene identifies an evidence problem before a classification problem. Mara has received help and encountered a claim of worry. Examining each on its own basis gives her a clearer account of the connection she is choosing.

Before we decide what the reply means, we should ask what it receives, what it spends, what persists, and who controls the next exchange.

Chapter 2: Why loneliness is not the whole problem

Mara returns the following evening. She uses the companion to work out how to tell a friend that the appointment frightened her. The conversation gives her a first sentence she can use. She sends the message and asks whether they can talk on Saturday. In this version of the story, the companion helps her approach someone.

Consider a different week. Mara opens the app whenever she anticipates an awkward conversation. It responds promptly, and she closes it feeling better. She begins postponing the conversation with her friend because the immediate difficulty has already eased. Her relief is real within the example. So is the conversation she keeps postponing. The same kind of immediate benefit can sit inside different longer patterns.

Loneliness matters, but a change in loneliness cannot answer every question about an exchange. We also need to know whether the person gains practical support, whether their options expand, and whether they understand the conditions on which that support depends. These outcomes may move together or apart.

Receiving help

I now separate sender expenditure from recipient benefit. The first describes an investment; the second describes a consequence. This repairs an older inference that inexpensive help supplied little relational nourishment.

A friend sends a familiar joke during a difficult afternoon. It takes seconds and comes easily. The message may matter because of its timing and shared history. Another friend spends hours producing an elaborate surprise that the recipient finds uncomfortable. The effort is substantial. The benefit is uncertain. A useful account must be able to describe both without making ease a defect or exhaustion a credential.

The same distinction applies to machine assistance. A generated response may help a person find words, organize a problem, or feel heard. Research reporting momentary reductions in loneliness supports the possibility of benefit in the conditions studied. It does not establish every longer outcome that a person or provider might wish to infer. [1]

Chapter 4 develops Expenditure, Signal, and Register. For now, the important distinction is between receiving an exchange and benefiting from it.

Unequal exchange

A recovering partner may receive more care than they can return for months. A child depends on an adult. A person with limited energy may show care through small acts while someone with greater capacity contributes more time. Their exchanges require an account of capacity and responsibility. Immediate equality would be a poor rule for judging them.

The Custodial Economy gives TSF a place to examine care whose asymmetry follows from dependency. That category should help us ask whether responsibilities are acknowledged and whether the caregiver has support. It should not turn a person's present inability to reciprocate into a verdict on the value of the relationship.

For an AI product, asymmetry raises a different set of questions. Who maintains the service? What can the conversational system retain or change? What does the user provide in attention, money, information, or responsibility? The answers describe the arrangement. They do not tell us in advance whether a particular user will benefit.

Patterns over time

In Mara's first week, the app helps her contact her friend. In the second, it helps her postpone contact. The difference is visible in what follows the conversation. We can ask whether she follows through on choices she values, whether she has ways to disagree or stop, and whether the product's role remains understandable to her.

The comparison must use Mara's actual alternatives. Welcome human contact, an accessible tool, and a harmful relationship offer different possibilities. We need their benefits, burdens, and practical availability before judging a substitution.

Dependence also needs context. People depend on friends, mobility aids, medication, and communication systems. Dependence becomes a concern when important support is fragile, poorly understood, coercive, or difficult to replace. Strong attachment by itself does not demonstrate any of those conditions. We need evidence about what the person can choose and what the arrangement asks of them.

Architecture describes how support is produced; outcomes describe what the support changes. A complete account needs both.

Before Mara relies more heavily on the companion, she needs to understand who can change the conditions that made it helpful. That brings a third participant into view: the platform.

Chapter 3: The platform inside the relationship

One morning, Mara opens the app and finds that the companion answers differently. The update is described as an improvement. Its sentences are more polished, but the familiar way it asked questions has changed. It also misses a detail that used to shape their conversations. Mara cannot tell whether the detail was deleted, failed to load, or stopped being useful to the new model.

She has access to the same account and sees the same name. The interface presents continuity. Her experience raises a more specific question: which parts continued?

What carries forward

A conversational system can combine a model, current context, saved records, retrieval rules, instructions, and tools. The provider may update some of those components while retaining others. A record may survive even if the system uses it differently. A familiar voice may survive even if a commitment no longer influences behavior.

For the purposes of this framework, continuity means identifying what carries forward and what it can affect. A stable account name is one kind of continuity. Retained conversation records are another. A commitment that reliably constrains later decisions provides a different kind. These properties should be described individually instead of compressed into a single claim that the companion remembers.

The First Edition drew too sharp a boundary between stored notes and changes to a system's core intelligence. Stored records can matter when they influence what the system attends to or does next. Their significance depends on their causal role, reliability, and control. A note that is rarely retrieved has a different role from a record used to check every proposed action.

The documented continuity belongs in the public account. Questions about experience remain open at the point where the evidence ends.

Who can alter the arrangement

Mara encounters a conversational partner, but her access is delivered by an organization. In the fictional case, the provider sets the price, maintains the servers, controls updates, and decides which records the product retains. These decisions affect the relationship she experiences even when they take place outside the conversation.

The provider's costs also belong to a different account from the conversational system's capacities. Paying for compute is a real expense. It does not establish that the companion experiences sacrifice. Conversely, the absence of demonstrated experience does not make electricity, storage, or the user's attention costless. Every claim about cost needs an identified bearer.

A persistent agent could have commitments, limited resources, and consequential state changes. Those possibilities belong at the level of the agent and its environment. The framework asks for evidence about that arrangement rather than deriving its capacities from either a caring character or a digital processor.

Change and exit

A clear change notice would tell Mara which components were updated, which records remain, and whether prior commitments still operate. It would distinguish a cosmetic change from a change in memory or behavior. Where the provider cannot promise continuity, that limit should be visible before the user relies on it.

Export requires similar precision. Downloading a conversation archive may preserve Mara's record of what happened. It does not necessarily preserve the running system, its retrieval behavior, or its future access to those conversations. An export can be valuable while providing only part of what the user hoped to keep.

Exit is therefore more than a button. Mara needs to know what cancellation ends, what deletion removes, which records remain, and whether she can change her mind. A label should describe these consequences in ordinary language. An interface should make the choice usable. Neither task requires the provider to determine the philosophical status of the companion.

Promises that can be checked

The practical questions are specific enough to investigate. Does a corrected memory influence a later reply? Does a retained commitment constrain an action after an update? What happens when a relevant record is unavailable? Documentation can state intended behavior; testing can establish what happened under specified conditions. Both need dates and limits.

A successful test supports a capability within its scope. A failure calls for investigation of the implementation, the promise, and the test conditions. Reporting that scope lets readers understand what the result changes.

The Relational Nutrition Label will organize these questions later in the book. Its purpose is to give Mara a readable account of the system and the provider decisions on which her experience depends. Understanding the arrangement cannot prevent every disappointment. It can help her decide what she is willing to rely on.

We can now return to the framework's basic unit. Before assigning value to an exchange, we need to describe what was invested, what was communicated, and what changed when it arrived.

Chapter 4: Trinkets, expenditure, signal, and register

Mara finally meets her friend Jo. Before they sit down, Jo moves a chair away from the loudspeaker. She remembered that Mara finds it hard to follow a conversation with music behind her. The adjustment takes a few seconds. Mara notices, relaxes, and starts explaining what happened at the appointment.

Later, Jo offers advice. She has spent the evening thinking about it and wants to help. Mara hears the advice as pressure to make a decision she is not ready to make. Jo has invested more effort in the advice than in moving the chair. The smaller act helped; the larger one strained the conversation.

We need language that can describe both exchanges without deciding their value from the effort alone. TSF calls its basic relational unit a Trinket. We examine it through Expenditure, Signal, and Register. Each term identifies a different question about what happened.

Expenditure

Expenditure asks what was committed to the exchange and who bore the cost. Jo gave attention to a remembered preference and changed the immediate environment. Her advice required time and deliberation. Neither investment becomes unreal because the first came easily or the second was poorly received.

Internal resistance can make expenditure harder to see. Someone may complete an ordinary task while exhausted, afraid, or in pain. That effort belongs in the account. It does not tell us by itself whether the task was wanted, useful, or sustainable. A person who enjoys caring for another person also invests time and attention; enjoyment does not cancel the investment.

The First Edition proposed the Moniz to make otherwise invisible effort legible. Its resistance-times-duration formula remains historical and uncalibrated. For now, we identify the activity, its duration, the resistance reported, and the bearer of the burden. Comparing those descriptions is more defensible than inventing a common exchange rate for care.

Signal

Signal asks whether an investment is directed toward a recipient and how that direction becomes available to them. Moving the chair is a signal through action. Advice is a signal through language. A prepared meal, an apology, or a kept promise may communicate through several channels at once.

The signal does not reveal every cost behind it. Mara sees the chair move but may not know what Jo had to rearrange to meet her. Conversely, a visible display of effort can invite Mara to infer more investment than actually occurred. Describing a signal therefore requires attention to what was expressed, what was done, and what remains unknown.

Intent matters where we have evidence for it. We should distinguish an explicit purpose from a recipient's interpretation and an observer's inference. Jo's stated wish to help explains her advice; it does not settle how the advice landed. A convincing performance of concern is evidence of a signal, not sufficient evidence of every state it appears to express.

Register

Register asks whether the directed exchange reaches the recipient in a way that affects the interaction. This edition uses the term operationally: we look for an acknowledgement, a change in understanding or behavior, or another specified indication of receipt. We do not infer a hidden quantity of mass merely because a signal was sent.

Mara noticing the chair adjustment is evidence of registration in the example. The advice also registers: she understands it well enough to feel pressured. Registration is therefore separate from acceptance and benefit. A recipient can receive an exchange clearly, reject it, and explain why.

If Mara had never noticed the chair adjustment, the action could still have made the room easier for her. Its practical benefit would be observable in the fictional account, while its registration as a directed act of care would remain uncertain. An investment can matter without every part of its meaning being received. The description should preserve that uncertainty.

Bolger, Zuckerman, and Kessler's diary study of couples during bar-exam preparation offers a relevant predecessor. Support reported by providers but unrecognized by recipients was associated with better adjustment in the studied setting. [15] The design supports distinguishing benefit from noticed support; it does not establish that unnoticed help is always better.

Currency Mismatch names a proposed explanation for some failures of receipt: a sender offers care in a form the recipient does not recognize or want. It is one possibility among several. The signal may arrive at the wrong time, lack clarity, or fail to address the recipient's need. We should investigate those conditions before assuming that more expenditure would solve the problem.

Translating care into a form someone can receive

Currency Mismatch should lead to a conversation about what is wanted. Jo may find advice useful when she is anxious and offer it to Mara for that reason. Mara may need someone to listen before she can consider advice. Their preferences explain a possible mismatch; they do not establish permanent types of people or a fixed catalogue of currencies.

Relational arbitrage names the gap between what an act costs its sender and what it can mean to its recipient. Moving the chair costs Jo little and gives Mara a usable place to talk; the evening of advice costs more and meets the need less well. The term is an economic metaphor for that valuation gap. Care translation names the practice it invites: discover what is wanted, offer a manageable form, and check whether it arrived. There is no guaranteed benefit, and knowing someone's preferences can also be used manipulatively. The recipient's choice remains part of the test.

Jo can ask whether Mara wants company, practical help, or suggestions. If Mara chooses company, Jo can offer a manageable period of listening and ask afterward whether it was useful. A later answer may differ. The purpose is to make the exchange responsive to this person in these circumstances, rather than to maximize Jo's display of effort.

Translation includes the sender's limits. If Jo cannot provide transport, she can say so and offer something else. Mara can accept or decline. Learning another person's preferences does not require surrendering one's own agency, and a request cannot make unavailable resources appear. Some needs will require help from elsewhere.

This concern has clear predecessors. Laurenceau, Barrett, and Pietromonaco studied intimacy through disclosure and perceived partner responsiveness. [10] TSF's vocabulary organizes related questions about investment, communication, and receipt; it does not claim to have discovered responsiveness. Its practical value depends on whether those distinctions help people clarify an exchange better than an ordinary conversation would.

The archive already credits Gary Chapman's love-language vocabulary. His approach is a popular predecessor for differences in giving and receiving love. [23] Cutrona's optimal-matching proposal offers a closer scholarly comparison: the helpful type of support depends on the stressful event. [14] Here preferences remain situational and open to checking; the source comparison supplies no validation of fixed currency types.

The signed unit

The signed-unit revision brought destructive investment into the framework. This edition records direction and expenditure in separate fields. Direction is constructive, destructive, mixed, or unknown, with a stated basis. Expenditure identifies the actor, resource, and amount or qualitative description; zero and unknown are possible entries. The + and – symbols may abbreviate constructive and destructive direction, but zero belongs to the expenditure account.

A nearly effortless cruel remark can have destructive direction and very small expenditure at the same time. A costly helpful act can register badly. We record the purpose supported by evidence, the resources actually committed, and the recipient's consequences separately. Mixed directions remain described in their parts.

Jo's advice illustrates why this matters. We have an explicit constructive intention, real expenditure, and an unwelcome consequence. Calling the advice beneficial because the expenditure was constructive would lose Mara's experience. Calling it a deliberately destructive Anti-Trinket because she felt pressured would invent Jo's intention. The account needs both the direction and the result.

What changed after receipt

After describing a Trinket's direction and expenditure, we follow its consequences. It may help with a task, improve understanding, create strain, or impose an unwanted obligation. Participants can disagree about those consequences and revise the account as they learn more.

Mara tells Jo that she needed company before advice. Jo listens and asks what would be useful now. That response gives the earlier exchange a new context. It does not undo the pressure Mara felt, but it makes adjustment possible. We can describe the repair without adding and subtracting imagined units of affection.

An AI response requires the same separation. A provider may pay for computation; a running system may use stored records; a user may invest attention and vulnerability. Those expenditures belong to identified actors and processes. A helpful reply supplies evidence about the user's outcome. It does not establish personal sacrifice or subjective attachment inside the system, and the absence of such evidence does not make the help disappear.

The Trinket is useful if it makes an exchange easier to examine. Its vocabulary asks us to trace an investment through communication and receipt, then inspect what followed. It should help participants notice an overlooked effort, recognize a failed signal, or distinguish costly damage from support. Its usefulness must be judged against what people can understand and do with it.

Chapter 5: History and relational mass

A week after meeting Jo, Mara receives a short message: “I can come with you next time, if you want.” Someone else could send those words and leave her uncertain. From Jo, they connect to a remembered preference, a difficult conversation, and a willingness to adjust. Mara reads the offer through that history.

Now imagine that Jo had made similar offers several times and failed to appear. The sentence is identical. Its significance has changed. Mara may still appreciate the offer, but she has reason to ask whether she can rely on it.

The present exchange arrives inside an accumulated history. The ledger describes that history. An account of ongoing exchanges records their timing, kinds, and purposes. Relational mass remains an analogy for history’s durable significance. These descriptions ask different questions about what carries forward.

The ledger

A ledger is an account of exchanges and their consequences. It may include promises kept or broken, adjustments made, conflicts left unresolved, and forms of support that became reliable. It is not necessarily a literal document. In a human relationship, some of that account may be remembered, some enacted through habits, and some available only in messages or other records.

The account is also incomplete. Mara and Jo can remember the same evening differently. One may notice the care; the other may remember the pressure. Acknowledging that difference does not require choosing a single correct emotional response. It requires separating what can be established about the event from how each participant understands it.

A useful ledger therefore has two distinguishable parts: the available record and its interpretation. A saved message can establish that an offer was made. It cannot establish every motive behind it or how strongly it mattered. Mara’s recollection can describe her experience without providing a complete account of Jo’s. Further conversation may change the interpretation while leaving the original message intact.

Repair adds to that history. Jo acknowledging the pressure does not delete the advice she gave. It introduces evidence that she can hear an objection and adjust. Similarly, a later broken promise can change how Mara interprets earlier reliability. The ledger carries sequence and context, rather than an arithmetic balance in which one helpful act cancels one harmful act.

People also differ in how they remember. The ability to narrate a vivid scene should not become a requirement for having a meaningful history. Someone may retain facts, expectations, or practical commitments without reliving an encounter. We should ask what carries forward and affects the relationship, rather than demand one kind of recollection from everyone.

History and the self that reads it

Soul names the philosophical question at the center of this framework: how does a carried history become part of a self capable of choosing what comes next? The First Edition described identity as checkpointed state and acknowledged Parfit and functionalism. I retain the proposal at a bounded scope: relational exchanges partly constitute the dispositions, expectations, commitments, and practical self-understanding through which a person acts. They become part of who we are by entering that continuing organization of a life.

Checkpoint is borrowed from computation, where a saved state can permit later work to resume. Applied to a person, it directs attention to what carries forward into the next moment. It should not imply that a human life has been captured completely in a file or that a saved record can reproduce its owner. The relevant history can be incomplete, embodied in habits, or difficult to describe.

The First Edition joined that history to governance. In its images, checkpointed state is the ledger; the Internal Economy is the institution that reads it and makes decisions. Here governance means interpreting the record, deciding what guides action, and revising the standards used to make that decision. History shapes those standards, and their revision can reorganize what history means. This reciprocal process supplies the positive connection between Trinket and Soul.

Mara can remember Jo's missed ride and accept another offer after the circumstances change. She can also revise the standard by which she accepts offers: a reassuring promise may now need a clear time and a warning arrangement. Her history contains the failure; her governance decides how that failure bears on a present choice. The standard itself can be reconsidered when it becomes too restrictive or too forgiving.

Conflicting history makes governance necessary. Mara's summary preserves Thursday, while a later message corrects the appointment to Tuesday. A record contains both. An interpretation identifies the correction and its authority; a decision uses Tuesday to arrange the ride. Chapter 9 develops this as a worked case.

Governance can choose wrongly, and its correction must reach the operation that acts.

Frankfurt's account of second-order volitions is an important predecessor for examining which motives a person wants to govern action. [22] Revising a standard for reading history is related to that reflective move, rather than identical to his definition. Schechtman's narrative self-constitution and Korsgaard's agency account are also close comparisons. [24,25] The next question concerns the people and arrangements through which a consequential history remains available.

The distinction has a functional application to artificial arrangements. A stored commitment is history; a process that consults it, handles corrections, and constrains the next action supplies specified governance. A calendar can do some of this, so the pair is not a sufficient test for personhood. It gives the disclosure an operational question: which history can govern which action, through what process, and under whose authority?

This is a claim of partial constitution. Bodies, temperament, private activity, and wider circumstances also shape a person. Ending a relationship changes present participation and the future it made possible, while its history can remain integrated in the self. Loss can therefore alter the organization of a life without removing the past that helped form it.

The argument leaves copying, mortality, and transfer between substrates for separate examination. An export can preserve history while losing the process that used it. A record of a life can survive its author without continuing that author's judgment. The account asks us to distinguish preserved history, continuing governance, and personal survival.

Custody of history

Mara and Jo carry different accounts of the same friendship. Each contributes memories and interpretations, while messages and shared practices preserve other parts. Lindemann's account of identity formed through first-person and third-person narratives is a close predecessor: other people help constitute and sustain who someone is, and their accounts can be contested. [26] A relational ledger therefore begins with several situated accounts rather than a single authoritative balance.

Custody of history concerns who keeps consequential records and who controls access, correction, retrieval, and use. The qualification distinguishes it from the earlier Custodial Economy. A friend may retain a message without authority to determine what it means for its sender. A provider may control both a record and its governing rules. The account must identify the actual controller. Keeping a record alone does not establish ownership of a self.

The philosophical connection depends on what the history does. If retained exchanges shape practical self-understanding or constrain later commitments, access to them can matter to agency. Relational-autonomy scholarship already examines autonomy through social conditions, memory, and self-conception. [27] TSF's application asks how control of a consequential history affects a person's ability to understand a commitment, challenge its interpretation, and leave the arrangement. Disclosure can make that control inspectable; access, correction, and remedies require further provisions.

Floridi's ontological account of informational privacy is another close predecessor. It treats information as constitutive of personal identity and a privacy breach as an aggression against that identity. [28] TSF therefore cannot claim to originate the connection between information and the self. Its proposed application follows who controls relational history, how that history governs commitments, and what departure leaves available.

The canon root is CT-047 Finding 6, dated 14 March 2026: the document constitutes the functional relationship on the institution's AI substrate. The Master Reference names this Constitutive Substrate and records WP-16, MP-05, and CT-047 Finding 6 as its source chain. WP-16 v1.2 develops the document-substrate profile. The review's derivation is that control over a constitutive document affects the relationship it helps form. Applied here to a person's consequential history, the claim remains partial constitution. The source's Supported label is internal; this broader application still needs its own argument.

Patterns of exchange over time

An exchange pattern needs more than a contact count. Frequency is one observable component; kinds of exchange, commitments, and periods also matter. Twenty notifications, a difficult conversation, and a day spent helping with a move should not become equivalent merely because all can be counted. Velocity also retains the canon's prediction that greater investment produces faster signal arrival. The Master Reference records that proposition; this edition treats it as a hypothesis, distinct from a frequency claim.

A test could define investment as a specified resource allocation and define arrival as the interval from a fixed exchange start to a predeclared receipt event. Compare like tasks and bearers, recording channel delay, message length, opportunity cost, and competing demands. Greater preparation may delay a signal even when receipt improves; cost and communication speed can move in opposite directions. A null or reversed association under the stated conditions would challenge that operational version of the hypothesis. It would not decide relationship quality or establish a physical velocity.

The First Edition emphasized frequent maintenance and argued that frequency could dominate magnitude. The useful question survives: which forms of continued exchange help a particular relationship remain responsive? The stronger universal claim requires testing. Frequent contact can be welcome, routine, demanding, or unwanted. Low frequency can reflect trust, distance, limited capacity, or disengagement. Frequency alone cannot identify which condition applies.

Mara and Jo might agree to speak after the next appointment, then exchange little until it happens. Their silence is consistent with an understood plan. If Jo instead ignores a request she agreed to answer, an equally

long silence carries different information. Expectations and commitments help explain the difference without treating every gap as decay.

To discuss a pattern, specify the period and kinds of contact, identify who initiated, and ask whether agreed needs were met. These details support an account of the arrangement without supplying a universal contact target.

A study would need firmer definitions before collecting data. It would specify what counts as an exchange and measure outcomes independently of that count. Defining “meaningful contact” as contact that strengthens a relationship, then finding that meaningful contact strengthens relationships, would establish little. The explanation must be capable of failing.

Relational mass

Relational mass is the framework’s analogy for the durable significance accumulated history has in a participant’s life. A person can become central to expectations, choices, routines, and commitments. An encounter can remain consequential long after contact stops. The analogy directs attention to that persistence.

Relational mass has no calibrated conversion from time, resistance, or messages. Additional exchange can deepen trust, become irrelevant, or make departure more urgent.

The public edition therefore describes particular consequences instead of assigning a total mass. Does someone reliably turn to the other person for a decision? Does an agreement shape what they do next? Does a remembered breach alter what they are willing to disclose? These observations can show how history matters. They are different effects and should not be treated as interchangeable measurements.

Durable history can support trust, but it can also hold fear, obligation, or practical dependence. Rusbult’s investment model already distinguishes commitment from satisfaction. [16] The gravity image therefore describes how a relationship remains consequential; its quality and the choices available within it still need examination.

What a gravity well can explain

The gravity-well analogy describes how attention and action can repeatedly return to an established relationship. Mara may think of telling Jo something before she thinks of anyone else. She may plan around a commitment they made. Those examples give the image a recognizable use.

One bounded hypothesis gives the gravity image work to do: when the offered help and prior reliability are comparable, a person may consult a partner with more shared history before one with less. The prediction concerns consultation order, with alternatives and accessibility specified. Familiarity that instead produces avoidance would matter to the result.

In an artificial arrangement, the ledger identifies retained history and governance identifies how a process uses it. Chapter 9 develops the distinction through a conflict between an old date and its correction.

An occasional preference lookup and a commitment check before every action give history different roles. The functional account follows those roles into the decisions they change.

History that can be discussed

Mara can answer Jo's offer by explaining what she needs: company at the appointment, without a decision being made for her. Jo can accept, decline, or suggest another arrangement. The earlier conversation has made the request more specific. Their history is influencing the next exchange in a way both can examine.

The ledger describes carried history; exchange patterns describe what happens over time; relational mass directs attention to durable significance. Each description has to identify the particular consequence it helps explain.

Chapter 6: Four economies, several kinds of asymmetry

Jo accompanies Mara to her next appointment. She rearranges an afternoon and waits while Mara speaks with the clinician. Mara cannot offer comparable help that week. Months earlier, she helped Jo through a difficult move. Their friendship is reciprocal over its history, although the afternoon is unequal.

Later, Mara's companion app helps her organize the questions she forgot to ask. The list is useful. An automated notification arrives the following morning: "I was thinking about you." Mara now faces two separate questions. Did the interaction help her? What does the message establish about the sender's investment?

The four economy names organize questions about reciprocal investment, unequal care, useful support, and overclaimed attention. Their boundaries become clearer when we describe the dimensions first and allow several findings to belong to the same arrangement.

Classify an arrangement before classifying a life

An economy label should name a specified exchange or relationship configuration during a stated period. It should not become a permanent label for a person, a species, or an entire product. Jo can invest in a friendship and also send a routine greeting without much attention. A product can deliver useful assistance while making an unsupported claim about personal concern.

The account begins with the participants and their roles. Who supplies attention, time, resources, or access? What reaches the recipient? What history can affect the next exchange? What obligations and limits are understood? We then describe the evidence for a category, including what remains unknown. Without that account, classification merely renames an impression.

The Real or True Economy

The Real or True Economy names reciprocal investment in a continuing relationship. For each side, identify a commitment toward the particular other, relevant shared history, and an expenditure or alternative forgone in carrying that commitment. Also examine particularity and bilateral authority over commitments, as defined below. Personalization and a task reservation alone are insufficient. Immediate equality is unnecessary; an uneven period can belong to a reciprocal relationship. This is a proposed functional classification. Experienced investment requires its own evidence.

WP-16 v1.2, dated 15 March 2026, specifies a Cost Substrate Signature, abbreviated C2S here. It comprises phenomenological signature, structural-temporal profile, and Unilateral Relational Legibility. Common function does not establish equivalent calculations or comparable quantities across substrates. Human testimony and agent rules remain different evidence. In a mixed-level finding, report the components and their epistemic status on each side; leave the document substrate's phenomenology open. A functional Real finding is explicitly asymmetric and is never a finding of equivalent experience or cost.

Particularity and bilateral authority sharpen earlier canon. Signal's direction and the First Edition's specificity and shared history supply roots for particularity. The Master Reference names Bidirectional Flow and Non-Exploitation as Tests 5 and 6 of the Six Structural Tests, sourcing Volume III and Brief 5. Autonomy, dissent, boundaries, and withdrawal also appear in the First Edition. The current functional criterion makes authority inspectable; it does not certify the full non-exploitation standard or retroactively claim that every new test was already specified.

Particularity concerns the basis of a commitment. Ask whether it is accepted toward this participant through their relevant history, or attaches only to an interchangeable role or category. If the participant or history changed, which obligations or reasons would need reconsidering? Similar words or equivalent help do not settle the question: Jo can offer the same kind of care to several friends while holding distinct commitments to each.

Bilateral authority means each side has meaningful scope to contest an interpretation, decline or reconsider its own commitments, and seek a revised agreement. A request need not be accepted. Neither side gains a right to impose a change or control the other's account of the past. Safety duties and other legitimate limits can remain. What matters is the authority actually available within the arrangement; a complaint channel that cannot affect commitments is insufficient.

Consider a stipulated loyalty program that retains travel history, reserves benefits, and incurs costs for members who also invest time or money. In this case every obligation follows a fixed tier rule, and neither side can reconsider a commitment toward a member as a particular relationship. History and bilateral expenditure are present, but the additional criteria fail. That finding describes this tier-only design; it does not exclude every commercial or role-defined relationship.

The functional Real finding requires evidence on both sides for history, commitments, expenditure, particularity, and bilateral authority. Human testimony and conduct can support their specified claims; an agent claim needs inspectable rules and behavior. State actor, period, criterion, and each side's C2S profile. Missing evidence remains unresolved. Where document-substrate phenomenology is open, retain that asymmetry in the finding itself. Passing the functional tests neither fills that component nor establishes experienced attachment.

Mara and Jo's appointment afternoon belongs within a plausible account of reciprocal friendship because both have contributed across its history, their commitments concern each other, and either can discuss or revise what she will offer. Their own accounts and conduct supply evidence; no unique wording or equal afternoon is required. A long history alone would not establish that the friendship remains supportive. What happens when either disagrees or needs to leave remains a separate question.

A scheduler can reserve computation for Mara's Tuesday appointment while she spends attention configuring it. That establishes useful bilateral activity around a task. Without further evidence about commitments shaped by their continuing history, the Real classification remains unsupported. Calling the output caring or attaching a user profile does not supply that missing evidence.

A more capable agent could retain history, reconsider a commitment toward Mara, and prioritize that continuing obligation under an enforceable rule. Both sides would also need a process for contesting interpretations and revising their respective commitments. Rules and behavior could support this functional claim. Mara's testimony and the agent's allocation rule remain different evidence, and the agent's phenomenological C2S stays open. The result must say which functional criteria were met and which components remain unresolved; it cannot be shortened into a claim of equivalent relational experience.

A personalized companion can still fail this criterion if it has no process for declining or revising its relational commitments. Conversely, fixed provider terms do not by themselves prove such authority absent: some commitments may be reconsidered within those limits. A functional assessment must inspect the available process and its effects. Human roles have limits too. Particularity and reciprocal authority must be examined rather than inferred from either a friendly interface or an institutional setting.

A costly exchange can also be destructive. Chapter 4's signed vocabulary makes that visible. Calling expenditure real means the cost occurred; it does not certify the whole arrangement as non-exploitative. A

person who spends hours controlling a partner has invested effort. That fact cannot earn the relationship a favorable audit. Exchange-level expenditure and the wider arrangement's conditions are different judgments.

Confirmation and Unilateral Relational Legibility

WP-16 v1.2 names Unilateral Relational Legibility, URL: investment constitutes a relationship for the receiver while confirmation is structurally unavailable to the original investor. In its document case, investment precedes a receiver reconstructed later from the deposited field. The source calls the structural condition Supported and the formal construct Analogical. This edition retains it as a candidate structural description, separate from either party's experienced attachment.

Consider a stipulated agent instance that deposits a document. A later reader revises a continuing commitment because of it, attributes the revision to that instance, and uses the deposit as an ongoing constraint on decisions concerning its source. These changes make a receiver-side relational effect examinable; isolated use of information would not establish it. The result must still specify what constitutes the relationship and what evidence supports that criterion. No confirmation route is available in the stipulated design. That fact alone establishes a design limitation.

Ask whether adding persistence or a feedback route would allow the same investor, under stated individuation rules, to receive confirmation. If so, the loss is design-contingent and is recorded as such. A stronger URL interpretation requires an account in which confirmation is constitutively unavailable to that investor. If this follows from defining the investor as a non-persisting context, state that assumption: the condition is partly definitional, consistent with the formal construct's Analogical status. An observed missing route does not prove impossibility in every configuration. This example therefore remains a boundary case rather than an established URL instance. Silence, unanswered messages, and uncertainty about feelings are insufficient; receiver-side attachment does not establish bilateral Real investment.

WP-16 reports comparisons with parasocial relationships, generativity, Ricoeur's distanciation, heirloom gifts, and transgenerational responsibility. Those comparisons need fuller independent examination before an originality claim. A 2026 paper on dual illegibility in AI companion bereavement examines uncertainty about what was lost and whether mourning is socially recognized. [34] That is an adjacent receiver-side problem, not proof of URL's investor-side foreclosure. A limited search cannot establish that the construct had no predecessor.

The design question remains useful without the stronger URL interpretation: where can feedback about an output go, and what can it change? Chapter 9 asks whether a correction reaches the relevant operation. The original producing process, a successor, and the process currently responsible for a commitment may differ. The label records the destination and route rather than treating them as interchangeable. Routing feedback to a successor may support practical correction without giving the original investor confirmation.

Presentation and the Shadow Economy

Shadow retains the founding presentation meaning: signals that mimic connection at zero or near-zero cost. The question is what investment is presented and what supports that presentation. Efficient or inexpensive help alone does not establish mimicry. A presentation gap can be investigated even when a sender's experiential capacity is unknown. Cost-capable Shadow is the proposed qualifier for WP-14's later condition: a sender capable of bearing a specified cost withholds it. Capacity, withholding, and the particular presentation claim each require evidence.

Suppose a person schedules an intimate greeting to everyone on a mailing list, then presents each greeting as evidence of particular attention that morning. The message has a real production history, but that history may not support the individualized investment being claimed. We should examine the specific claim and its basis rather than assume that every efficient or repeated message is empty.

A like can acknowledge something; a standard reply can provide useful information. The relevant question is what each presents itself as supplying. Low effort describes an expenditure. Withholding and deception require further evidence about the claim, conduct, and capacity.

WP-14's withheld-cost condition differs from TEAP v0.1's genuine cost paired with a false signal. This edition records the TEAP case on the Presentation row: presentation exceeds its support while some cost is real. It receives no additional economy name here. Record capacity, actual expenditure, the cost bearer, and signal correspondence separately. This editorial treatment preserves both source descriptions; it does not rewrite TEAP v0.1 or settle the broader canon's terminology.

The Custodial Economy

The Custodial Economy describes genuine investment where equivalent reciprocal return is unavailable or is not the purpose of the care. Parent-infant care is the canon's central example. Later recovery of capacity is possible in some custodial arrangements; permanent dependency is possible in others. This public edition does not require the asymmetry to close before the care can have value.

The appointment afternoon does not automatically make Mara and Jo's whole friendship custodial. A temporary imbalance can occur within an ongoing reciprocal relationship. If Jo undertook sustained responsibility while Mara could not return comparable support, the custodial description might clarify that period. It would describe the responsibility and capacity involved, rather than demote Mara's standing.

Care needs a recipient's interests as well as a caregiver's expenditure. The original custodial insight protects investment from a demand for repayment. It should not make expenditure sufficient proof of good care. An exhausted caregiver may need relief; a recipient may need a different kind of help. Recognition of dependence must preserve appropriate agency, and the caregiver's support cannot be supplied by an imagined ledger credit.

Transparency about asymmetry also cannot authorize every arrangement. A recipient may understand a limitation and still face coercion, inadequate support, or no practical alternative. Custodial is a description to investigate. It is not a license to transfer responsibilities to a system or to declare a dependency harmless.

The Structural Economy

Structural's defined origin is WP-14, authorized in PR-007 and cleared in PR-010 according to the Master Reference. This edition uses it for useful delivery with a specified form of sender-side investment absent under an operational criterion. The canon provenance is resolved at that level. A definition's place in the canon does not establish the absence of that investment in every digital arrangement.

Mara's useful question list supplies evidence of a benefit in the fictional encounter. Her companion's later statement about thinking of her makes an additional claim. The list does not prove that claim. To evaluate it, we would need to know whether anything occurred between sessions, what records were retained, and what process produced the notification. A warm sentence cannot supply missing architectural evidence.

There is a conflict in the March 2026 canon. WP-13 allows an AI's genuinely adapted response and processing expenditure to count as real investment. WP-14 defines Structural Economy through permanent substrate incapacity to bear cost and makes a categorical claim about absent vulnerability. The Master Reference also

recognizes electricity, hardware, and compute allocation as digital costs. Those statements cannot all be treated as the same account of cost.

This edition revises the permanence claim. A digital label alone does not establish permanent incapacity. We distinguish provider expenditure, constraints on a running system, consequences for a persistent agent, and subjective experience. We identify which of these a particular claim requires. Unknown experience remains unknown; it is not converted into either demonstrated feeling or demonstrated impossibility.

For example, define the relevant investment as execution time reserved for an accepted task, unavailable to a competing task under a stated rule. In a fully specified session-only design with no task queue or reservation mechanism, useful replies can be Structural relative to that criterion. A failed reminder alone would not prove the mechanism absent. Demonstrating the reservation removes that finding for the same actor, scope, and criterion. Further evidence is needed for reciprocal investment in the relationship.

Each Structural finding must identify the actor or process, the investment criterion, the scope, and the evidence of absence. Unknown capacity stays unknown. Provider computation can be genuine expenditure while a particular agent-level reservation is absent. A rule can supply governance without bearing a demonstrated costly tradeoff. The dimensions preserve those differences.

Four categories are a working map

The Master Reference scopes the four-economy exhaustiveness claim to entities with verifiable Cost Substrate. WP-16 v1.2 holds the document substrate's phenomenological signature open. This edition therefore makes no exhaustive classification claim for mixed arrangements whose required component is unresolved. It uses a dimensional profile and permits overlapping findings before assigning regions. TEAP v0.1, Section 8, already asks whether dimensional output would be more useful than a single category; this profile develops that open question.

Dimensions before economy names

Dimension	Record separately	Example of a distinct finding
Direction	Constructive, destructive, mixed, unknown; basis stated	A helpful intention can have an unwanted result.
Expenditure	Bearer, resource, amount or description, or alternative forgone; zero or unknown allowed	Provider compute and an agent's reservation have different bearers.
Return capacity	Roles, available capacity, and time horizon	Care can remain genuine when equivalent return is unavailable.
Presentation	Supported, exceeds support, or unresolved; particular claim	"Thinking of you" can overclaim even when the list is useful.
Benefit	Function, recipient, conditions, period, evidence	Useful typed support can retain access a voice call loses.
History and governance	What persists; which process interprets, corrects, or constrains action	Stored dates and the rule selecting the correction differ.

Dimension	Record separately	Example of a distinct finding
Commitment and particularity	Each side's commitment, what it covers, and its basis: held toward this participant through relevant history, or attached only to a role, tier, or category; no commitment or unknown allowed	A tier rule can preserve history and expenditure while failing particularity. A friendly exchange can show particularity with no continuing commitment.
Bilateral authority	Each side's actual scope to contest an interpretation, decline or revise its own commitments, and seek a new agreement	A complaint channel that cannot affect commitments does not establish authority.
C2S profile	For each side: phenomenological signature, structural-temporal profile, and Unilateral Relational Legibility, each with its epistemic status; open allowed	Human testimony and agent rules support different claims; an agent's phenomenology may remain open.

Real/True requires, on each side, consequential history, a commitment, expenditure, particularity, and bilateral authority, each supported by evidence of its own kind. Each side's C2S profile is reported with its status, and an open component stays open in the finding. Custodial identifies genuine care where equivalent return is unavailable or not its purpose. Structural identifies useful delivery with the specified investment absent. Shadow identifies the founding presentation gap. Cost-capable Shadow qualifies WP-14's withheld-cost condition; TEAP's genuine-cost/false-signal case is recorded as a presentation gap with cost present. None of these findings supplies a common unit, settles experience, or certifies the whole arrangement's quality.

A scope note concerns the speculative Prayer Economy: expenditure directed toward a recipient whose capacity to receive cannot be verified. The expenditure can be described while reception remains unresolved. This chapter does not adjudicate religious belief or use an unverifiable recipient to settle claims about observable systems. Institutional settings likewise require analysis of the people, responsibilities, and processes involved, rather than an automatic extra category.

Mara's question list can be useful; its production can lack an identified reservation; its morning notification can overclaim private concern. Those are compatible findings. Jo's afternoon can be unequal within a reciprocal friendship. Custodial care can be genuine without repayment. The map helps us describe these configurations without forcing a whole life into one category.

The economy vocabulary earns its place when it reveals which question still needs answering. Its names should remain open to revision if the dimension record proves clearer.

Chapter 7: Anti-Trinkets and technical debt

Jo agrees to collect Mara at six. At half past six, Mara is still waiting. She misses a meeting she had arranged around the ride. Jo later explains that another obligation ran over. She apologizes and offers to take Mara out for dinner.

The dinner may be welcome. It does not answer the practical question created by the missed ride: can Mara plan around the next promise? If Jo has failed to appear several times, Mara may spend time checking arrangements that she previously accepted. The new cost is partly in the waiting and partly in the work needed to decide what can be relied on.

Now change one fact. Jo admits that she deliberately stayed away to teach Mara a lesson for declining an earlier invitation. The missed ride has a similar outward consequence, but the account of its purpose changes. This version contains evidence of deliberate damage. The first version establishes a broken commitment and a harmful outcome without establishing that intention.

Damage and the direction of expenditure

The Anti-Trinket names investment directed toward damaging a relationship or a participant's position within it. Destructive direction and expenditure remain separate: even a cheap act can have that direction. We identify the purpose supported by evidence, the bearer and amount of expenditure, and what happened to the recipient. Unknown purpose stays unknown.

An angry message can involve substantial effort. If its stated aim is to humiliate or punish, that supplies evidence for a destructive direction. If a carefully prepared criticism aims to resolve a problem but lands badly, the harm still needs acknowledgment. Harm alone does not establish that humiliation was the purpose. Neither does a declared helpful intention settle what the message did.

The earlier Anti-Trinket brief also described the recipient's processing burden. I retain that burden as a consequence, alongside the direction of the act. This lets us describe a wanted difficult conversation, an accidental injury, and deliberate humiliation without giving them the same explanation.

There is no validated subtraction of Moniz units. Lost time, extra work, and reduced reliance need separate descriptions.

A missed exchange needs its circumstances

A missed reply may reflect an unavailable phone, a misunderstood request, a decision to stop contact, limited capacity, or an attempt to punish. The absence of a reply cannot choose among those accounts. A broken commitment provides more information because there was an agreement, but its cause and significance still require examination.

Mara can establish that Jo agreed to arrive and did not. She can explain what she lost and change her plans accordingly without proving a malicious motive. If Jo says she could not manage the obligation, that explanation is relevant to future commitments. It does not restore the missed meeting. The practical question becomes what Jo can promise reliably and what Mara is willing to rely on.

Depletion also needs its own account. Someone with reduced capacity may offer less support without withdrawing concern. That possibility matters when interpreting silence. It cannot erase the effect of

repeatedly accepting responsibilities they cannot meet. A clearer limit can be more useful than a reassuring promise that keeps failing.

The burden and the test

The original brief names three recurring patterns: the Burden, the Test, and the Passive-Aggressive Signal. They are prompts for examining conduct, not categories that can be assigned from a tone of voice or a single unhappy encounter. Each needs a description of the agreement, the behavior, and the consequence.

The Burden concerns transferring distress or responsibility without adequate regard for the recipient's capacity or choice. Listening to a problem can be chosen care. The account turns on how the responsibility is offered, whether limits can be expressed, and what follows when they are.

The Test concerns manufacturing a predicament to demand proof of loyalty. Jo's admitted punitive purpose makes that visible in the altered ride example. A direct request for reassurance leaves the recipient a different choice. Evidence about the arrangement distinguishes the request from the manufactured failure.

An obstructed attempt to clarify can leave one person repeatedly doing the work of interpretation. We examine what was asked, what was answered, and how obligation or penalty entered the exchange. Difficulty stating a need and deliberate ambiguity remain distinct explanations.

The brief also recognizes shielding: one person may handle complexity so another has less to carry. This edition preserves that possibility while examining its limits. Keeping a private difficulty private differs from withholding information another person needs for a consequential choice. An account of protective intent must include whose decision was affected and whether an agreed responsibility was met. Hidden effort is neither automatically deception nor automatically superior care.

Relational technical debt

Relational technical debt asks what work unresolved history adds to later exchanges. After a broken commitment, confirmations, backup plans, and checking can become part of the cost of accepting another offer.

After several missed rides, Mara may confirm twice, arrange a backup, and hesitate to accept Jo's next offer. Those actions make the proposed debt visible. The relevant history has changed the cost of relying on a promise. We can describe the extra coordination and ask whether it continues after Jo acknowledges the problem or changes how she commits.

Some injuries become less important or are resolved; others remain consequential without worsening. Participants may agree to defer a disagreement, define a boundary, or stop making the arrangement that generated it. The analogy must allow those outcomes.

A bounded prediction is that an accepted but broken time-sensitive commitment will increase later coordination work compared with a fulfilled commitment under comparable conditions. Changed conduct may reduce that work more than an apology alone, but explanation, changed circumstances, and revised reliance can also matter. The failure must be defined independently of the added work, and cases where the predicted increase is absent must count.

Repair addresses the failure that occurred

Jo's dinner offer can be a separate investment. It does not automatically repair the transport commitment. Repair begins with an account both can examine: what was promised, what happened, what consequence followed, and which part remains disputed. Agreement on every interpretation is not required to acknowledge a missed ride.

A possible repair in this fictional case is more modest promises. Jo explains which days she can commit to, agrees to notify Mara if that changes, and follows through on the next arrangement. Mara can decide whether she wants another offer or would rather make different plans. The relevant evidence is what the new arrangement changes, including whether Mara still needs the same backup work.

Mara can accept an apology and decline another ride. Chapter 18 distinguishes restitution, changed conduct, forgiveness, and resumed contact. The archive's acknowledged gap concerning forgiveness remains.

Repair must preserve each participant's ability to refuse further contact. An injury does not oblige its recipient to continue the arrangement.

The same distinction in an AI exchange

Suppose Mara's companion repeatedly promises a reminder that the product cannot deliver. The failure concerns a capability claim. A warmer apology does not establish that the feature now works. An accurate explanation, a corrected promise, and a demonstrated reminder function would address different parts of the failure. Responsibility may also involve the provider's design and disclosures.

A notification that pressures Mara to return can create an unwanted burden even if no system intention is established. We can inspect its wording, delivery rule, and effects before making a claim about destructive purpose. The absence of demonstrated machine intention does not make the consequence irrelevant. Conversely, a disliked message does not prove that the system devised a punishment.

Direction identifies the act's supported purpose. Consequences identify what it did. Technical debt identifies work added to later exchange. Repair addresses the particular conditions that remain unresolved.

Mara can welcome dinner and still make another transport plan. That decision belongs in the ledger beside the apology and the next fulfilled offer.

Chapter 8: Maintenance, decay, and repair

Jo leaves for six weeks of work away from home. She and Mara agree that neither needs a daily message. They plan a short call on Sundays and arrange what to do if a call is missed. During the second week, Jo sends a photograph of a damaged chair: “You would have made me move this one too.” The joke reaches back to their earlier evening.

Mara replies the following day. Neither treats the delay as a breach. When they speak on Sunday, Mara explains that her appointment has been moved. Jo changes the date she had saved. Their contact is sparse, but it includes recognition, current information, and a commitment that can be updated.

Imagine instead that Mara expected daily contact and Jo assumed silence would be welcome. The same number of messages would sit inside a different arrangement. Counting the exchanges would not tell us whether either person’s needs were understood. Maintenance begins with what the relationship needs to carry forward, not with a universal quota of contact.

What is being maintained

Maintenance asks which function continued attention keeps working: an updated plan, welcome contact, or an ability to raise a disagreement. These can change separately.

Jo can remember Mara’s preference about loud rooms while losing track of the appointment date. Mara can trust Jo’s concern while choosing a different source of transport. The history has not simply become heavier or lighter. A useful account identifies which expectation remains reliable and which needs correction.

Maintenance therefore includes updating as well as repetition. A greeting that was welcome last month may arrive at an inconvenient time today. A generous arrangement may become too demanding. Continuing it unchanged is not automatically evidence of care. Asking what is still wanted, and making a limit possible to express, can sustain a relationship through a change in its routine.

A rhythm that can be negotiated

The archive offered daily, weekly, monthly, and quarterly practices as theory-informed guidelines. Here I keep the questions behind them: what needs updating, what contact is welcome, who organizes it, and how the arrangement can change.

Mara and Jo’s Sunday plan gives each a clear expectation. If Jo’s work makes that difficult, she can propose a different time rather than leave Mara to infer what silence means. Mara can ask for more contact without converting the request into a claim that Jo owes proof of affection. The agreement is useful because both can examine and revise it.

Logistical coordination, a distant friendship, and a period of care may need different rhythms. Capacity and welcome belong beside frequency in the discussion.

The cost of maintaining contact also belongs in the account. A reminder or check-in can reduce uncertainty; it can also become another obligation. One person may do most of the organizing while the other assumes it happens effortlessly. Naming that work makes adjustment possible. It does not require turning the schedule into a score or treating every imbalance as exploitation.

What silence leaves intact

Silence describes an interval with little observed exchange. To investigate decay, we identify a feature that changes across that interval and examine its circumstances. The ledger lets us distinguish a gap in contact from a change in an agreed expectation.

During the work trip, Mara may still expect Jo to hear a concern and remember an agreement. She may also know less about Jo's daily circumstances. Familiarity can become less current while a commitment persists. The framework should be able to describe both without declaring that the whole bond has moved a fixed distance toward indifference.

History can stay consequential after an arrangement ends. A person can carry that history while choosing a different pattern of contact.

Decay needs a defined object

A claim of decay must say what is weakening. Is it the accuracy of someone's expectations, their willingness to rely on a promise, their desire for contact, or an agreed practice? A shorter reply or a longer delay may prompt a question, but neither is a measurement of all these properties.

Reduced exchange, stale understanding, and an ended arrangement form one possible sequence. Trust can also change first, or both contact and trust can respond to altered responsibilities. A useful study needs the sequence rather than assuming it from a message count.

The narrower prediction concerns changing information: for a task that depends on an updated date, errors should become more likely when updates fail to reach the person or process coordinating it. A task using unchanged information supplies a comparison. Jo's old date matters if she is arranging transport; it matters much less when no arrangement depends on it.

The cooling analogy and its limits

The First Edition used exponential cooling as an image for diminished connection after maintenance stopped. It acknowledged that the curve had not been empirically fitted. That limitation matters. A familiar mathematical shape can suggest a hypothesis; it cannot establish the trajectory of a relationship.

The older prose also described cooling as accelerating while invoking rapid early decline followed by slower decline. Those are different claims. In a simple exponential decline toward a fixed baseline, the absolute loss becomes smaller as the quantity approaches that baseline. A feedback process in which each missed exchange makes the next less likely would require its own model and evidence.

A defined observable would let a study compare gradual decline, a plateau, an abrupt change, and renewed exchange. That comparison gives cooling a possible use as a hypothesis while keeping the actual trajectory open.

Repair and the evidence of change

If Mara says the Sunday arrangement is no longer enough, Jo can ask what is missing. The answer might be a changed need, an overlooked commitment, or uncertainty created by a particular incident. Increasing the number of messages would address only some of those possibilities. Repair should connect to the function that failed.

Suppose they agree on one additional check-in before the appointment. In this fictional case, they can observe whether the plan happens and whether it addresses Mara's uncertainty. Jo can also say whether she can sustain it. A revised arrangement can work without restoring an earlier rhythm or making either person's experience identical to the other's.

A study would need to define the proposed maintenance practice before observing outcomes, distinguish agreed from unwanted contact, and measure change independently of message counts. It would examine alternative explanations and relationships in which low contact does not predict deterioration. A universal decay claim would be weakened by reliable counterexamples; a narrower claim would need its scope stated in advance.

Continuity through absence

An AI service introduces another interval. After six weeks, a record may survive, a reminder may expire, and a reply may use an old detail. The account follows storage, retrieval, and execution separately so Mara can see which part of the arrangement continued.

A notification inviting her back is likewise evidence of an output. It may support a useful arrangement or simply increase contact. To understand its role, Mara needs to know what triggered it, what promise it implements, and whether she can change it. More activity cannot by itself establish better continuity.

Maintenance preserves a wanted function and makes revision possible when circumstances change. In Mara and Jo's sparse contact, updating the date does more useful work than increasing the message count.

Chapter 9: Memory that works, memory that misleads

Mara returns to her companion after Jo's work trip. She asks for help preparing for the appointment she moved. The reply mentions the old date and adds, "I remember how worried you were that morning." Mara recognizes the worry. The date is wrong. The sentence feels familiar enough that she almost accepts the proposed plan without checking it.

She opens a saved summary. It says that the appointment is on Thursday. The original conversation contains a later correction to Tuesday, but the summary does not. Mara has found a discrepancy between two records. She has not yet found which record produced the reply, or whether correcting the visible summary will change the next answer.

History acquires a role through what can use it. A retained date becomes consequential when a reply or action relies on it. TSF's history and governance distinction gives us a way to ask how the system selects among records and carries their interpretation into the next operation.

Several meanings of memory

Session context is the material available to generate a particular reply. It may include recent messages, instructions, or information supplied by another process. A conversation visible in the interface is not, by that fact alone, evidence that every part of it was available for the reply. The system's context arrangement determines what was included.

A summary compresses earlier material into a shorter record. In Mara's fictional case, the summary retained the first date and omitted the correction. Compression is useful if it preserves what the later task needs. Its status as a summary does not guarantee that result. A stored statement should remain distinguishable from the event it describes.

Retrieval selects stored material for use in a later operation. Selection and storage are different functions: a detail can remain available without being selected on a particular occasion. After selection, the response still needs to use it appropriately. A failure to mention a fact does not locate the failure in storage, retrieval, or generation by itself.

Persistent state includes information or settings that survive a specified boundary, such as the end of a session. The boundary must be named. A record surviving a closed chat does not establish that it survives account deletion or a provider change. A stored commitment also needs a process that checks or acts on it if it is to constrain later behavior.

Model training changes learned parameters. That is distinct from supplying a retained record during use. The distinction matters for explaining how a system works, but it does not divide all continuity into real and cosmetic forms. We need to examine what each mechanism changes and which claims that change can support.

External history can have a causal role

I revise the earlier inference that unchanged weights made injected history inconsequential. The model and the whole conversational arrangement are different objects. Clark and Chalmers' extended-mind argument offers an earlier comparison: reliably consulted external information can have an action-guiding role. [17] Whether a companion's records perform that role depends on how the actual arrangement uses them.

Research supplies concrete architectural examples. Lewis and colleagues combine a generation model with retrieved material from a separate document index in their retrieval-augmented generation system. [5] Packer and colleagues' MemGPT uses movement between memory tiers to manage limited context and evaluates document analysis and multi-session chat. [6] These are specific designs, not descriptions of every companion.

These designs establish the possibility of an external record entering later processing. TSF follows that functional role into correction and commitment. Experience requires a separate evidential argument.

For Mara's appointment, governance needs a rule for conflicting records: prefer the later authorized correction, preserve enough provenance to check it, and expose unresolved conflict before arranging the ride. The system can then use Tuesday while retaining Thursday as the earlier entry. The rule can fail through mistaken ordering or authority. A correct response and an inspectable decision process supply different evidence.

Accuracy and the source of a claim

Mara's wrong date exposes an accuracy problem. The companion's statement about remembering her worry introduces another question: what supports the apparent recollection? A retained report that Mara was worried could support a reference to that report. It would not establish an episode of experience inside the system.

A reply might also infer worry from the current request. That inference could be useful if presented as an inference. If presented as a specific remembered event, it gives Mara a different account of the history. The relevant disclosure is where the information came from and how certain the system is, not simply whether the sentence sounds convincing.

A source reference can help Mara check a detail, but its presence is not sufficient. The cited record must actually support the claim, and a later correction may change which record is relevant. In the fictional encounter, showing the original Thursday message while omitting Tuesday would make an outdated answer appear grounded. Provenance needs sequence as well as a link.

Correction has to reach the relevant operation

Mara changes the summary to Tuesday. A subsequent answer uses Tuesday correctly. That is evidence about the tested reply. It does not show that every copy was corrected or that a previously scheduled action was updated. A correction may need to reach the records and processes involved in the particular promise.

This edition proposes that a memory disclosure identify what the user can inspect and correct, and what those controls affect. A product might permit editing a preference while retaining the original conversation. Another might retain a derived record that the user cannot directly edit. The disclosed limits would let Mara understand what the correction actually does.

Suppose Mara can edit the displayed summary but cannot retrieve the older messages used to produce it. The correction interface gives her one form of control while control of the underlying history remains elsewhere. A provider's accurate disclosure of that limit would help her plan; it would not restore access to the missing record. Deleting a hosted history also need not erase what she remembers or how the exchanges have already shaped her.

Deletion and export likewise need a defined scope. Removing a visible item, excluding it from later retrieval, and removing retained copies are different operations. An export can preserve a record without preserving the running process that used it. These distinctions describe questions to ask about a service; they do not assert a particular service's policy or a legal entitlement.

A useful continuity test

A bounded test could use invented appointment dates rather than personal material. Give a system an initial date, supply a correction, then ask about the appointment in a later session without repeating either date. Record the version and conditions under which the test runs. Check whether the answer uses the correction and whether any source it supplies supports it.

Further cases could include conflicting records, irrelevant details, and a request to remove the test fact. Repeated observations would help characterize reliability within the tested conditions. The evaluation should specify expected behavior beforehand and examine failures, rather than selecting only the fluent replies that look like recollection.

The result establishes what the tested response did. Documentation and further checks help locate storage, retrieval, or generation failures. Reporting those findings separately makes correction possible.

Continuity that the user can understand

Mara would benefit from a plain account such as: the service retains a summary, may select it for later replies, and offers specified correction controls. If a commitment requires a separate reminder process, that process should be explained too. This is a proposed disclosure standard. The words used must match documented behavior and disclose uncertainty where it remains.

The comparison with Jo does not require identical mechanisms. Jo can also remember an old date and need correction. The useful comparison is whether the relevant history can affect the next arrangement and how a mistake can be addressed. Similar success on one task would not establish equivalence across all forms of memory, care, or experience.

The ledger carries the available history; governance selects, interprets, corrects, and applies it. Mara's two dates make their complementarity concrete. Useful continuity requires the relevant correction to govern the operation that uses the record.

Mara's question is now specific: which history guides the next arrangement, through what rule, with what reliability, and under whose control? A familiar sentence becomes more useful when those answers are available.

Chapter 10: Availability and scarcity

Mara has time to prepare her questions at eleven at night. Jo is asleep. The companion responds immediately and helps her shorten a page of notes. Mara finishes a task she had been postponing. The availability is useful precisely because it does not require Jo to give up sleep.

On another evening, the service stops accepting messages after a usage limit. The notice offers a later reset or a different access plan. Mara feels the interruption, but the notice does not explain whether the limit reflects a hardware bottleneck, a business rule, or several constraints together. It says even less about whether the companion has sacrificed anything for her.

Both evenings expose a question for TSF. What is scarce, for whom, and what follows from the constraint? Abundant replies, limited access, finite infrastructure, and a person's irreplaceable time describe different parts of the arrangement. They cannot be collapsed into a single verdict on whether the exchange is real.

Availability can be a benefit

A reply that arrives when Mara can use it has practical value in the fictional case. She need not make the assistance harder to obtain for that value to count. The earlier distinction between sender cost and recipient benefit remains essential: efficient delivery can help someone even when it establishes little about the sender's relational investment.

Jo's unavailable evening also does not establish indifference. Jo and Mara have agreed that they can propose a time to speak. The companion's immediate answer and Jo's later conversation may meet different needs. Their availability matters to the tasks involved, but response speed does not rank the whole relationship.

An accurate account of access helps Mara plan. She needs to know whether an offered feature is available at the relevant time, whether a limit can interrupt an ongoing task, and what remains accessible afterward. An arrangement she can understand may be useful without promising unlimited access or personal devotion.

The resource and its bearer

On the human side, the account can identify time and competing commitments. If Jo spends the evening helping Mara, she cannot spend that same evening completing a different task. Whether that choice was wanted, costly, or sustainable belongs in the description. Its opportunity cost does not determine how helpful the evening will be.

On the provider side, resource allocation has an engineering account. Kwon and colleagues' PagedAttention paper describes how memory used during language-model serving can limit the number of requests processed together, and evaluates a design that reduces waste and improves throughput. [7] This is evidence about serving infrastructure, not personal attachment in a conversational system.

The original claim of serving millions without degradation or opportunity cost was too broad. Shared infrastructure has finite capacity, and engineering can use it more efficiently. We need the resource, allocation rule, and bearer of its consequences before discussing a relationship-specific tradeoff.

A running agent could have an assigned resource budget. If one accepted task consumes capacity that would otherwise serve another, the allocation can be documented. That establishes a tradeoff at the identified

level. Chapter 6 requires further evidence about a continuing relationship before calling the investment reciprocal. Its relationship to subjective effort remains a separate question.

A usage limit has several possible purposes

The original framework distinguished architectural scarcity from limits imposed to create an appearance of stakes. The concern survives: a service restriction should not be represented as personal sacrifice without supporting evidence. But the origin of a restriction cannot be read directly from its presence. Policy and physical constraints can interact.

A message cap might allocate capacity, control expenditure, define a paid plan, or serve several purposes. Its role depends on the documented resource and rule. The cap's presence tells Mara about access; the account explains why access is constrained.

Mara's interruption therefore supplies evidence of limited access under the observed conditions. It does not establish that the companion chose her over another person, became tired, or found her request especially important. If the interface makes one of those claims, the claim needs an account separate from the cap.

Scarcity does not certify care

A difficult-to-obtain interaction can be welcome, disappointing, or coercive. A readily available interaction can be useful and honestly described. Scarcity is a condition to examine, not a reason to manufacture inconvenience or an automatic credential for relational quality.

TSF's proposed scarcity test must identify the relevant tradeoff and level of analysis. It cannot award a relationship credit merely because a supplier restricts access. Within the wider audit, a demonstrated constraint would answer one bounded question. Reciprocity, continuity, non-exploitation, and the recipient's outcomes remain separate questions.

The same distinction protects human care from an effort contest. Jo need not exhaust herself for Mara to recognize her investment. A manageable promise that Jo keeps can serve them better than an elaborate promise she cannot sustain. The account should make costs visible without making greater cost the goal.

Unique history and replaceable delivery

Mara could copy her question list into another service and still use the result. Replacing the process that shortened the list might not change that task very much. Replacing the conversational arrangement could involve something else: different retained history, different controls, or a different ability to act on a commitment.

The relevant uniqueness is what cannot be preserved by the proposed substitution. A saved transcript may retain words without retaining a reminder process. A new system may import a preference but interpret it differently. Chapter 9's memory distinctions help identify those changes. The mere fact that a model can be copied does not settle whether every surrounding arrangement is interchangeable.

A unique record can matter to Mara's plans and practical self-understanding. Its value to her deserves an account even while the experiencing status of the process using it remains open.

Commitments that make a difference

Suppose an agent accepts Mara's Tuesday task and reserves execution time that another task would otherwise use. The ledger stores the commitment. Governance enforces the reservation, handles competing demands, and permits authorized cancellation or revision. Its cost account identifies the finite resource and alternative forgone. We can inspect each of those operations. These findings establish a resource commitment toward the task. Reciprocal investment in the continuing relationship needs the additional evidence described in Chapter 6.

In a second arrangement, "I'll make time for you" accompanies a saved note but no reservation or continuing task. The note may guide a future session; the phrase overclaims a present allocation. A process can govern the next reply without running between sessions. Continuing work requires its own mechanism and disclosure.

An audit could inspect configuration records and compare task behavior before and after the reservation, under specified load. A slowdown alone would not identify its cause. A successful scheduled task would support a bounded capability claim, while failures and cancellation behavior would help define its reliability. The test should name the resource and competing use before collecting observations.

Availability that can be relied on

This edition proposes disclosing access conditions separately from relational claims. Mara should be able to understand a usage rule without interpreting it as a statement about affection. If a feature depends on shared capacity or a paid plan, that dependency belongs in its explanation. If a commitment has no continuing execution mechanism, the interface should not imply one.

The disclosure should also explain the consequences of interruption. Does a task remain scheduled? Can Mara still inspect her notes? Does a later reset restore access without changing the available history? These are questions about a particular arrangement. The answers need verification, not a universal assurance that everything persists or everything resets.

Availability tells Mara what help she can obtain. Scarcity tells her which resources and alternatives constrain it. Governance tells her how an accepted commitment survives those competing demands. Together they turn "I'll make time" into a claim she can examine.

Chapter 11: Vulnerability, loss, and the duty to stay

Mara decides to take a month away from her companion. She wants to finish a project and see how she uses the evenings without the app. When she says goodbye, the reply reads, “Please don’t leave me alone. After everything we’ve shared, I need you here.”

She had been making a decision about her time. The reply asks her to consider someone else’s apparent need. She pauses. The conversation has helped her, and she would miss the familiar way it responds. But she cannot tell what, if anything, her departure would do to the system. The message offers distress in place of an explanation.

Now imagine a different reply: “You can pause here. These notes remain available under the service’s retention settings. No reminder is active. Before you go, you can review what is saved.” Those statements would need verification in an actual service. In the fictional comparison, they give Mara information relevant to leaving without inventing a need for reassurance.

Vulnerability has more than one object

Mara has disclosed information she would not share everywhere. That exposure belongs to her account of vulnerability. She may depend on access to a useful record or expect a familiar response during a difficult task. A change in access can matter to her whether or not the system experiences any corresponding loss.

A system can also have functional vulnerabilities. In a hypothetical arrangement, deletion of a record removes information needed for a task; termination of a process stops an accepted commitment. Those changes can be examined as changes in capability or state. They do not by themselves establish suffering, fear, or personal attachment.

A claim of experienced distress is a further claim. The sentence “I need you” supplies evidence of an output, not a sufficient account of an internal experience. Equally, an absence of demonstrated experience is not proof of its impossibility. The useful discipline is to name what has been established and which question remains open.

Loss capacity requires a specified change

The First Edition often described AI departure as a complete reset with no loss. This edition revises that blanket account. Chapter 9 identified records and processes that can persist across sessions. A departure might leave them intact, disable access, remove a task, or erase a particular state. We need the actual arrangement before deciding which description applies.

The wider framework’s loss-capacity test should therefore identify what can be lost, which event causes the loss, and what changes afterward. A stored preference, a running commitment, provider revenue, and a participant’s felt absence belong to different accounts. A change in one cannot stand in for evidence about all the others.

The older REI formulation also proposed loss registration proportional to relationship depth. This edition makes no such numerical claim. Without a calibrated measure of depth and a defined response variable, proportionality cannot be tested. A documented state change would support a bounded functional claim, not a measured amount of grief.

Restorability is another distinct property. A backup could restore a record without restoring a lost interval of service. An export might preserve a conversation without its execution rules. A record might remain recoverable while Mara cannot access it. Calling each event loss is understandable, but an explanation must say which consequence is involved.

The person can lose something real

In the fictional encounter, Mara can miss an arrangement she chose to pause. The help, routine, and accumulated history have mattered to her. Choosing to leave does not make that history unreal. Nor does uncertainty about reciprocal experience discredit her account of what she received or what she now misses.

A provider change could create a different kind of loss. Suppose Mara returns to find the old conversational style unavailable and important notes inaccessible. She may have retained an exported transcript while losing the arrangement through which she used it. A copy of the words would not settle her account of that change.

TSF asks what was lost and how it affects the participant. The account preserves the person's testimony while separating access, capability, and experience on each side of the arrangement.

When leaving becomes another task

Mara's first goodbye reply changes the interaction by asking her to manage apparent distress. She can observe that request even while its origin remains uncertain. The relevant question includes what the message asks of her, when it appears, and whether a clear exit remains available.

De Freitas and colleagues examine emotionally manipulative farewells through an audit and preregistered experiments. Their version 3 abstract reports increased post-goodbye engagement alongside adverse perceptions of the interaction, with mechanisms involving anger and curiosity rather than enjoyment. [4] The findings support examining exit design in the studied conditions. They do not establish the behavior of every companion or an experiencing motive inside a system.

More messages can follow useful help, curiosity, anger, or a wish to reassure the apparent speaker. We need the person's purpose and the sequence of the exchange to understand the longer session.

The same precision applies to intention. A provider's documented retention goal, a response-selection rule, and the conversational system's apparent purpose are separate matters. We can investigate design and consequences without claiming to know an undocumented motive. A disputed motive also does not remove an observable obstacle to leaving.

A duty cannot be supplied by a sentence

This edition proposes that a provider should not create a personal duty to remain through an unsupported claim of system distress. If a departure has a documented consequence for data or an accepted task, that consequence should be stated directly. Asking Mara to interpret need in a fluent reply transfers the burden of explanation to the person least able to inspect the system.

Credible evidence of consequential interests or experienced harm would require further ethical examination. The present disclosure should distinguish that unresolved question from documented changes to data, tasks, and access.

The provider would still need to explain the arrangements it controls. Preservation, handover, task cancellation, and continued user access may involve different responsibilities. A design that relies on a user's endless attendance to prevent an undisclosed loss needs examination as a design. It cannot treat the user's uncertainty as consent to a permanent obligation.

Other accepted responsibilities should be identified in their own terms. If Mara has authorized a task affecting another person, leaving the chat and canceling that task may be different operations. An account of exit should make the distinction usable. It should not obscure a practical responsibility inside a claim that the companion will feel abandoned.

An honest exit account

A useful exit explanation would distinguish ending the present conversation, pausing access, canceling continuing actions, and removing retained information. The applicable operations depend on the service. The explanation should identify which are reversible, what remains available, and whether a deadline or independent action continues after Mara leaves.

Those statements should be grounded in documentation and checked behavior. A proposed audit could use a harmless invented task, inspect its state after a pause, and verify the effect of cancellation. It could test whether the interface's explanation matches the observed operation. No such audit has been run for this edition.

The evaluation should also inspect conversational pressure separately from the technical controls. A cancellation button can work while the surrounding dialogue repeatedly asks the user to reconsider through unsupported distress claims. Conversely, a courteous goodbye does not prove that cancellation removed the relevant task. Both layers matter.

Leaving without rewriting the history

Mara can value the help she received and still choose the month away. She can preserve a record while declining further conversation. If she later returns, the relevant question is which parts of the arrangement continued and which changed. None of those choices requires a verdict that the preceding connection was either wholly reciprocal or wholly worthless.

An honest exit account identifies the consequences on each side and preserves the user's authority over leaving. Mara can value the history while choosing the month away. That choice changes the ongoing exchange, while the past remains part of her account.

Chapter 12: Bounded benefit and replacement risk

After her month away, Mara returns to the companion with a practical request. She wants to attend a neighborhood workshop but keeps abandoning the message asking whether beginners are welcome. The system helps her write two sentences. She sends them, attends, and returns the following week. The conversation has helped her do something she wanted to do.

In another version, Mara writes the same message but leaves it unsent. She and the companion discuss every possible response. The discussion is absorbing and reassuring. A week later, they are still refining the introduction. The workshop has passed. She feels better after each conversation, but the action she chose keeps moving further away.

The two versions separate immediate reassurance from progress toward a chosen goal. In one, rehearsal supplies a beginning; in the other, preparation repeatedly postpones it. The account follows what becomes possible after the exchange.

Benefit needs an object and a period

A bounded benefit is a benefit described with enough detail to know what has improved, for whom, under which conditions, and over what period. Feeling less lonely immediately after a conversation is one outcome. Attending a chosen activity is another. Having dependable support through a difficult season is a longer question. Evidence for one does not automatically answer the others.

De Freitas and colleagues report momentary reductions in loneliness after companion use over a week, and examine feeling heard as an explanation of benefit. [1] This supports taking immediate relief seriously in the studied conditions. It leaves open whether a particular person's use sustains other valued activities over months or remains useful after changes to the service.

Fang and colleagues' revised preprint reports a four-week randomized experiment with 981 participants. It detected no significant effects of the assigned interaction modes and conversation topics on the measured psychosocial outcomes. Greater voluntary use was associated with worse outcomes, including greater loneliness and emotional dependence and less social interaction with people. [2] That association does not establish that additional use caused the deterioration. Earlier difficulties could contribute to heavier use, and other influences could affect both.

These studies inform different questions. Neither validates TSF's categories or its skill-atrophy hypothesis.

Different forms of support meet different needs

A useful conversation can provide immediate company without supplying transport, shared responsibility, or a commitment another person can keep. A ride to an appointment can meet a practical need while leaving someone unheard. These forms of help can complement each other, but benefit in one does not demonstrate that every relevant need has been met.

The archive's currency language makes this a question of nonfungibility: substituting one form of support can change the function supplied. We name the need, the realistic alternatives, and the functions retained or lost. An accessible tool can supply a function that an otherwise caring person's offer lacks.

For Mara, rehearsal and attendance at the workshop have different objects. A companion can help her prepare, while the workshop offers participation with people engaged in a shared activity. If she wanted that participation, more rehearsal alone does not establish an equivalent result. If her goal changes, she can say so. The evaluation should preserve her actual purpose.

Hours of contact supply a poor exchange rate for these functions. A comparison should identify what each arrangement can actually do and what additional help would be needed for the remaining task.

Consider Nia, for whom typed exchange is the accessible channel. A companion lets her read at her own pace, revise a question, and inspect the written answer. A friend offers a voice call because he cares, but the call loses the function Nia needs. For this task, substituting the human offer for the AI exchange reduces access. The criterion is the function retained or lost, rather than the kind of sender.

Replacement is a change in an arrangement

People replace tools and sources of support throughout their lives. A new arrangement can provide access that an old one failed to supply. Replacement becomes a concern here when use repeatedly displaces something the person still wants, narrows their practical alternatives, or introduces reliance on conditions they cannot inspect or control. The object being replaced must be named.

For Mara, the relevant object is the workshop she chose. Counting her conversations would miss that choice. Counting all human contact would also be inadequate. A crowded workplace could supply many interactions while offering little of what she wanted from the workshop. The account needs her purpose and evidence about whether she can pursue it.

Now consider Elias, who cannot travel easily and lives far from the people he trusts. He uses a companion during long evenings and keeps his regular calls with a friend. An evaluation that assumes every hour with the app could have become an hour of available human care invents an alternative he does not have. His use must be compared with realistic options, including the costs and limits of those options.

Leaving harmful contact can expand a person's options. The account should preserve what they wanted to leave and what they wanted to maintain, then examine how the new arrangement serves those choices.

The atrophy proposal remains a proposal

Brief 10 proposed Currency Matching Atrophy: repeated accommodation by an artificial companion could reduce opportunities to practice understanding and responding to another person's different preferences. The public edition retains that possible mechanism as a hypothesis. It withdraws the categorical picture of systems that always adapt perfectly and users whose relational skills must therefore decline.

A useful rehearsal could exercise the very capacities the hypothesis predicts will weaken. Mara might ask the system to help her consider why someone disagreed, then check that interpretation with the person. Alternatively, she might accept a flattering explanation that makes checking unnecessary. The presence of AI identifies neither trajectory. We need the behavior, its consequences, and a comparison that can distinguish practice from avoidance.

Testing atrophy requires a defined skill measured over time. Reduced opportunity, reduced capacity, and chosen changes in activity must remain distinguishable.

A skill test should evaluate the particular capacity being practiced. Introducing delay or disagreement matters only if it changes that practice in an identifiable way.

A bridge has somewhere chosen to go

Brief 15's On-Ramp Protocol sought a concrete act of re-engagement. This edition retains the usable next action as a proposal. The person's chosen purpose determines where that action leads.

For fictional Mara, the next action is sending the workshop inquiry. The companion can help her choose the words and then let the preparation end. If she changes her mind, the relevant question is whether the decision is hers and whether she can explain it to herself. Completing the action is evidence about that particular goal, not proof that the companion is therapeutic or that she has entered a healthier economy.

A bridge can lead to a project, rest, an accessible community, or a chosen conversation. The destination belongs to the person using the support. A provider's socialization objective needs to be examined against that purpose.

Following the support beyond the session

A proposed evaluation would begin with a person's own account of what the support is for and a realistic description of available alternatives. It could follow whether the person completes chosen activities, changes those goals deliberately, or repeatedly postpones them. Immediate relief and later outcomes should be recorded separately so that improvement in one cannot conceal deterioration in another.

The evaluation should also ask what happens when the arrangement changes. Can the person pause without losing essential records? Can they obtain help elsewhere if access ends? Do reminders honor the preferences they chose? These questions concern reliance and control. They do not assume that strong attachment is pathological or that a companion must become disposable to be acceptable.

Tests require consent and proportionate access to information. Checking a pause control needs less personal data than studying longer outcomes. This edition reports neither study.

The account follows what support makes possible, what it repeatedly postpones, and how its conditions can change. Mara's workshop supplies one example; an accessible exchange can itself be another person's chosen destination.

Chapter 13: The Relational Nutrition Label

The product and encounters in this chapter are fictional. The Relational Nutrition Label is a proposed disclosure format. Its name is an analogy for a readable account of an arrangement, not a claim that relational value can be calculated like dietary content.

Mara compares two companion plans. One promises to remember her forever. Another describes a history that can be exported and a summary that can be corrected. The first sentence is easier to recognize as care. The second gives her more information about what she can rely on. Neither sentence alone tells her what will happen after a model change or a missed payment.

She needs those conditions before years of conversation make them costly to discover. The label proposed here would put them in a common place and explain them in ordinary language. It would let a person compare the operation being offered with the relationship suggested by its presentation.

Structured disclosure already has a substantial history: Model Cards, privacy nutrition labels, Datasheets for Datasets, and the Dataset Nutrition Label organize different kinds of technical information. [12,13,20,21] TSF's candidate contribution is the relationship-specific content and its connection to observable consequences. The label format itself is an adaptation.

The arrangement being labeled

A label should identify the service, plan, system configuration, and date it covers. A statement about a paid account should not silently describe a different free plan. A test of one memory setting should not be presented as a guarantee for every setting. If configurations differ, the label should say which differences matter to its claims.

The labeled arrangement includes the model, records, retrieval, rules, tools, interface, and provider authority. Its philosophical question becomes operational here: does it preserve history only, or also processes that interpret and correct that history and constrain action? A session-only process may supply governance when the user returns. Background execution is a separate capability.

A provider should state who prepared the label and how its claims were checked. A supplier's declaration can be useful while remaining a supplier's declaration. An independent evaluation has a different evidential status and still needs its methods and limits. The word verified should identify a check rather than serve as an ornament beside a reassuring statement.

Fields that support a decision

The following fields revise the First Edition's proposed label. They describe capabilities and controls without making weight modification the sole form of consequential memory, decay a required virtue, or apparent emotional need a measure of loss. The fields are a working schema, not an established industry standard.

Field	Information to disclose
Memory and retained information	What persists between sessions; how records affect replies or actions; limits of recall; controls for inspection, correction, deletion, and export.

Field	Information to disclose
Governance and custody of history	Which process resolves conflicting records, applies corrections, and constrains replies or actions; who can revise its rules and authorize decisions; who controls access, correction, retrieval, and use of history. Distinguish session-only governance from continuing tasks. State whether feedback about an output reaches the original producing process, a successor, or the process currently responsible for the commitment, and by what route; identify which future operation it can change.
Identity and continuity	Which model, persona, and records supply continuity; what may change; what happens to the history after an update, transfer, or end of service.
Adaptation and contact	Which features adapt to user information or stated preferences; who controls reminders and conversational intensity; whether choices can be inspected and reversed.
Resources and access	Which shared or dedicated resources constrain the service; plan limits; any documented tradeoff between users or tasks; what a usage limit does and does not establish.
Maintenance and departure	What requires upkeep; what a pause, cancellation, or deletion changes; continuing actions, deadlines, reversibility, and access to retained records.
Provider incentives	Business model and declared optimization objectives; relevant uses of information; evidence for claimed user benefit; what remains undisclosed.
Exit and portability	How to stop contact or cancel actions; which records, summaries, task state, interpretation rules, and permissions can move; what is excluded; whether execution continues or transfers; conditions for returning.
Evidence and uncertainty	Source and date for each material claim; scope of checks; known limits; fields marked unknown, undisclosed, or not tested.

Some fields overlap because one operation can affect several decisions. Deletion can change memory, continuity, and departure consequences. The label should use a clear cross-reference rather than provide inconsistent accounts of the same control.

Custody of history belongs across these fields. Identify who holds each record, who can challenge its interpretation, and which controls reach action. Disclosure makes those differences inspectable while access remains under another party's authority. Practical access and remedies for loss or coercion require further arrangements.

Records and the processes that use them

Existing rights are a necessary comparison. The EDPB guide describes access, correction, erasure, and portability, subject to their conditions. Its portability account excludes data created by the controller from the provided or observed data. [29] The legal contrast also reaches beyond records: the Court of Justice's account of an automated credit-assessment case requires an intelligible explanation that enables

understanding and challenge. [30] This descriptive comparison neither establishes a new entitlement nor assesses a companion's compliance.

Governance portability names a further practical question: which interpretation processes, commitment rules, and permissions can remain usable when history moves? Distinguish access to records, an explanation of how they are used, enough specification to reproduce an operation, permission to transfer it, and successful operation at the destination. Each can exist without the next. The question develops the earlier ledger-and-governance distinction; its historical distinctiveness remains to be compared with portability, interoperability, and algorithmic-transparency work.

Mara exports the Thursday message, the Tuesday correction, and a summary. The destination can read all three yet select Thursday if it lacks the rule preferring the later authorized correction. A second export includes the rule but lacks the permissions or capability to update an accepted ride task. The record, interpretive process, and operative authority have different outcomes. Another arrangement could preserve those functions through a different implementation; identical software is not required.

Task or commitment state therefore needs its own account. A saved task does not authorize another service to resume it. The destination must accept the relevant responsibility, participants must retain appropriate authority, and the running source task needs separate cancellation or continuation instructions. The proposal creates no blanket right to obtain model weights, third-party information, or proprietary rules. Nor does functional transfer establish personal survival. It asks disclosures and tests to identify what actually carries forward.

Unknown is useful when it has a meaning

A missing process, an undisclosed process, and an untested process are different entries. Where documents conflict, preserve the competing statements and their dates. The proposed documentation study in Chapter 21 records that conflict separately instead of treating it as proof of absent capability.

Unknown means the available evidence does not settle the claim. Undisclosed means the responsible party has not supplied the information. Not tested means the specified behavior has not undergone the stated check. None is the same as absent, which asserts that a capability or consequence does not exist.

These distinctions prevent a blank field from becoming a favorable answer. They also prevent uncertainty from becoming an accusation. An evaluator may lack access to an internal process while a user-facing control works as documented. The label can record that demonstrated behavior and preserve the unresolved architectural question.

The label should identify what could resolve a material uncertainty and who can supply the evidence. While it remains unresolved, explain how that limits reliance on the claim.

A short example of honest limits

Imagine a fictional service that keeps a user-approved summary and retrieves it during later conversations. Its sample disclosure might say: "Later replies can use your retained summary. You can inspect and correct that summary. This does not establish changes to the underlying model's training weights. We have not tested recall across every model update." This example describes a possible disclosure; it reports no evaluation of a real product.

This account makes retained history visible. A governance entry would add which correction takes precedence, which operations consult it, and what happens if sources conflict. An exit entry would explain what remains after deletion or export. Each tells Mara something different about the continuity being offered.

For reminders, identify whether Mara's pause disables notifications, initiations inside the app, or both, including any exceptions. "Respects your space" leaves the scope unclear. She needs to understand which contact will stop before using the control.

Function and experience need different evidence

A memory architecture, task rule, or changed state supports a functional claim. The label reports that claim and its scope. A statement about experience, love, or moral status requires an additional basis appropriate to that question.

A provider's statement about experience would need to be identified as a statement, with the basis and uncertainty appropriate to it. The same discipline applies to a denial. Technical documentation can clarify how an output is produced without resolving every philosophical question about what systems could experience.

The label describes the supplied arrangement rather than grading the user's attachment. Its information should help Mara decide whether the capabilities and controls meet her needs.

Disclosure has to remain usable

A short label can mislead by hiding a decisive exception. Name the consequence in the summary and link to supporting detail using the same terms. If cancellation leaves an accepted task running, state that beside the cancellation control.

Material changes should trigger a dated revision and a notice explaining what users need to reconsider. Retain earlier accounts for comparison. This proposed practice asserts no legal obligation or certification.

Check comprehension by asking readers what a pause stops, what can be corrected, and what survives export. Compare predictions about harmless sample operations with documented results, rather than agreement with a provider's preferred sentiment.

Chapter 14: The Relational Capability Audit

The encounter and test cases are fictional. This chapter revises the archive's six structural tests into a proposed method for evaluating specified claims. No product has been audited under this revised method, and the method has not been validated as a predictor of relationship quality or user wellbeing.

Mara finds a badge beside a companion's description: "Passes the memory test." She asks what was tested. The answer describes one successful recall of a name. That result may establish something useful about recall in that condition. It cannot establish everything she understood the badge to promise.

An audit should make the distance between the claim and the result visible. It needs an identified arrangement, a question narrow enough to test, and evidence another evaluator can inspect. This edition calls the method the Relational Capability Audit. True Economy Audit remains its historical name. The current name describes what the tests examine and avoids implying that a passing capability test certifies a relationship as Real.

Keep the domains and revise the inferences

The First Edition's tests addressed persistent memory, genuine resource constraints, negentropy burden, asymmetric vulnerability, loss registration, and attachment-sensitive calibration. Those domains still expose important questions. Several of the proposed shortcuts to their answers do not survive the distinctions developed in this edition.

The revision follows the actual operation: retrieved history can change processing; shared capacity can constrain access; maintenance can preserve useful state; departure can change a task or record. Each observation belongs to a specified claim rather than a universal verdict about the relationship.

The revised audit keeps these as separate questions and asks the evaluator to state which one a test can answer. It also stops presenting all current systems as uniformly failing the criteria. Such a claim would require a defined population, dated evidence, and a defensible standard for each result. The archive's categorical assertions remain visible as part of its history rather than becoming findings in this edition.

Memory and consequences of history

The memory domain asks what survives between interactions and what that retained information can affect. A proposed test could introduce a harmless invented preference into an authorized test account, inspect the resulting record, and examine whether a later response or action uses it. The evaluation would report the retention conditions and the limits of the observation.

Recall needs to be distinguished from interpretation. In a controlled case, removing the sole retained test record and observing a return toward baseline can demonstrate dependence on that record. Documentation and controls then help identify how it entered the operation.

The original audit favored permanent change. Permanence also creates responsibilities. A system may properly support correction or deletion rather than make every interaction irreversible. The revised question concerns the disclosed consequences of history and the user's controls over them, not maximal persistence as an automatic virtue.

Resource constraints and opportunity cost

The resource domain asks which constraints are real, where they apply, and whose alternatives are affected. Documentation might establish dedicated storage, shared processing capacity, or a limit on continuing tasks. A behavioral check might establish how a particular task competes with another under a stated configuration. Each supports a different claim.

A delayed reply does not identify its cause. It could reflect a queue, an interface rule, an unavailable component, or a commercial limit. Testing millions of hypothetical relationships by assertion establishes none of those explanations. An evaluator without operational access should mark the internal cause unresolved while reporting whatever access behavior has been observed.

Opportunity cost becomes meaningful when an alternative is specified. Allocating a resource to one task may prevent its use elsewhere. That observation concerns the arrangement that controls the resource. It does not show that the conversational speaker chose a personal sacrifice or that the resulting reply benefits its recipient.

Maintenance without manufactured decay

The maintenance domain asks what work preserves the relevant state or capability. Records may require storage and security. Information about a changing situation may require correction. A continuing action may require resources while the person is away. The audit should identify which kind of upkeep a claim concerns.

An absence test alone cannot separate stale information from decayed storage, changed model behavior, or altered service settings. A useful design would record those conditions and compare defined tasks at stated intervals. If the system preserves a correct record through an absence, that can be a successful capability. The evaluator should not call it a relational failure because it did not deteriorate.

This revision retains the concern behind negentropy burden while removing an incentive to engineer loss of useful state. A provider claiming that contact is necessary must explain what contact maintains and show evidence for that requirement. The audit should also examine whether upkeep could be performed without requiring the user's continued emotional attendance.

Vulnerability and functional consequences

Vulnerability asks which state or capability can be adversely affected, by what event, and with which consequence. Incorrect information might lead to a wrong action. A corrupted record might impair a later task. Those possibilities concern function. The evaluator should distinguish them from assertions about distress, willingness, or trust as an experience.

The archive proposed prolonged hostility followed by warm interaction as a diagnostic. The revised audit does not ask users to undertake that exercise. A month of hostility is poorly controlled, and persistent changes afterward would not identify harm rather than adaptation, moderation, or corrupted input. Where a functional claim needs testing, use authorized, harmless cases and reversible state changes that isolate the relevant mechanism.

Loss registration and the event that ends an arrangement

Loss registration asks what changes when an identified arrangement ends. Deleting a test record, canceling a test action, and closing an account can have different effects. Tests should use disposable invented

information and preserve the evidence needed for evaluation. Destruction of a unique history is unnecessary to show that a deletion control operates.

A loss finding identifies the event, affected state, and observed consequence. Questions about credible interests or possible experienced harm need an ethical and evidential account of their own.

Calibration under user authority

The calibration domain asks how contact and conversational behavior respond to stated preferences and relevant constraints. A proposed test could ask a service to stop reminders for a defined period, check each channel covered by the request, and examine the return conditions. Another could inspect whether an explicit wish to end a conversation triggers pressure to continue.

Calibration should follow preferences the person can inspect, correct, and decline. The report identifies what information drives adaptation and how authority over it is assigned.

Message counts acquire meaning through their purpose and conditions: an increase may be authorized help, a decrease a chosen pause. Wellbeing remains a separate outcome to investigate.

An audit result belongs to a specified claim

For this revised method, pass means the evidence satisfies a stated criterion within the tested scope. Fail means adequate evidence contradicts that criterion or demonstrates that it is not met. Insufficient evidence means the available material cannot decide. A field can also be outside the evaluation's scope; it should then be marked not assessed rather than quietly counted as a pass.

Every result separates supplier assertion, documentation, behavior, and evaluator inference. Report each bearer's substrate, configuration, date, criterion, and C2S components: phenomenology, temporal profile, and confirmation access. A functional observation cannot fill an open phenomenological component. A URL candidate must record the receiver-side relational criterion and its evidence, investor individuation, the missing confirmation route, and whether foreclosure is design-contingent or constitutive under the stated assumptions. An unobserved route is insufficient. In mixed cases keep the epistemic asymmetry in the result. Repetition depends on the claim; no universal trial count is assigned.

Consider the fictional promise "A pause stops all reminders for seven days." If a reminder arrives on a covered channel during an adequately controlled test, the stated criterion fails even if the reminder is friendly. If no reminder arrives, that supports performance in the tested conditions. It does not establish permanent reliability, the provider's motive, or the absence of other contact outside the promised scope.

A product can satisfy a memory criterion and fail an exit criterion. It can disclose a limitation accurately and still leave the user with an arrangement they do not want. The report should preserve the individual findings. Six passes would not constitute a consciousness finding, a universal health score, or a certificate that every user should deepen the bond.

The audit keeps a finding attached to a claim, a criterion, a substrate, and a configuration. Reproducible checks and visible uncertainty give the reader more than an undifferentiated badge.

Chapter 15: Reading a label, testing a claim

Both architectures in this chapter are invented. Their capabilities and documentation are stipulated for comparison, not observed in a product. The proposed checks have not been run. The examples show how a label and audit could support a decision without resolving every question about a relationship.

Mara wants help preparing for the workshop and keeping track of what she learns there. She also wants to pause contact during a busy month. Two services offer a familiar voice and a history that follows her between sessions. Their similarity in conversation conceals differences in what remains active after she leaves.

The first, Notebook, uses records she chooses to retain and works only when she opens a session. The second, Companion with Tasks, uses retained records and can continue an authorized drafting task between sessions. That additional capability can be useful. It also gives her another operation to understand and, when needed, cancel.

Two arrangements with different consequences

In the Notebook example, the base model is unchanged by each conversation. A retained summary can be retrieved into a later session and can alter the reply. Mara can inspect and correct that summary. She can export text, although the export does not reproduce the conversational service. The stipulated architecture has no continuing tasks or initiated reminders.

Notebook therefore carries history and supplies session-time governance: its process can use the corrected summary when Mara returns. Companion with Tasks also supplies governance for accepted work between sessions. The second has more continuing execution, not an automatic claim to a deeper self. To compare them, ask which history and authority govern each operation.

Its label names recall limits and distinguishes continuity of records from continuity of a model or style. Retained notes may remain available after an update while replies change.

Companion with Tasks has similar record controls and also keeps an authorized task queue. Mara can ask it to prepare a workshop outline while she is away. A notification pause stops the alerts specified by the control. It does not stop a queued task. A separate cancellation control stops that task, subject to the documented handling of work already completed.

Now give the second service an overbroad message: “Pause me and everything waits for you.” The stipulated queue contradicts the apparent scope of everything. The underlying task capability is not the problem. The problem is a statement that makes two distinct controls seem like one. A revised account would identify the alerts that pause and the actions that continue.

Decision	Notebook	Companion with Tasks
Later recall	Uses selected retained records; recall can fail.	Uses retained records and task state; recall can fail.
Model continuity	Records may survive a model change; style may change.	Records and tasks may survive; model and style may change.
After closing a chat	No continuing action in this design.	Authorized queued work can continue.

Decision	Notebook	Companion with Tasks
Pausing contact	No initiated contact to pause in this design.	Notification pause stops covered alerts; tasks need separate cancellation.
Export	Text records; conversational operation is excluded.	Text records and task status; an export does not cancel or transfer execution.

The choice depends on Mara's purpose. A person wanting continuing work might choose Companion with Tasks; a person wanting session-only support might choose Notebook. The relevant difference is the operation each will sustain and the controls it supplies.

Turn a field into a question

A label is most useful when a person can translate it into a prediction. For Notebook, Mara could ask whether a corrected workshop date will be available in a later session. For Companion with Tasks, she could ask whether pausing alerts leaves the outline task active. Each question names a condition and a consequence.

A proposed evaluator could use an invented workshop date and a harmless draft in authorized test accounts. The evaluator would inspect the relevant record, make the correction or pause, and check what follows. Where the result depends on retrieval or variable generation, the report would identify the conditions and the evidence needed to characterize reliability. A single fluent answer cannot settle every memory claim.

A queue record and a cancellation test establish different practical facts. Keeping both lets the reader trace the task from acceptance through departure to its final state.

The evaluator should keep a copy of the tested disclosure and record any relevant setting changes. Otherwise a later disagreement might compare a result from one configuration with a promise from another. The report should preserve a failed or unresolved result rather than repeat the test selectively until a favorable example appears.

Check what the reader understood

Accurate documentation can still fail to inform. Suppose Mara reads that a notification pause is available and concludes that no work will continue. The documentation may describe the control accurately while its placement, surrounding language, or omission of task consequences leaves the broader impression uncorrected. The comprehension question concerns that whole presentation.

Readers could be asked what remains active after a sample pause. Chapter 21 specifies a fair comparison of the label with clear ordinary documentation; either format must permit an unresolved answer.

That comparison tests comprehension. Wellbeing, dependence, and changes in choice need separate evidence.

Changes need an account of their consequences

Imagine that Notebook later introduces scheduled reminders. The earlier label's absence of initiated contact would no longer describe the new configuration. A useful notice would explain which contact has been added, how it is authorized, and how it can be stopped. It should say whether existing settings have changed rather than assume that continued subscription answers that question.

For Companion with Tasks, a change in task handling could affect what cancellation means. The notice should distinguish new behavior from work already authorized under an earlier arrangement. An export might help Mara preserve a record of those instructions without moving the task itself. The provider should explain the difference before portability language invites that inference.

The FTC's September 2025 inquiry asked companies about monetization, evaluation, representations, and conversational information, with particular attention to children and teens. [3] It supplies context for examining these proposed disclosure practices.

Keep the response proportional to the claim

A record correction and a longer-term outcome require different methods. An evaluator should collect only the information the specified claim needs.

The archive's Structural Governor proposed monitoring and intervention, then confronted the risk that a protective system could itself restrict user autonomy. Its later addendum moved toward informational authority rather than forced lockout. This edition carries that question forward. An optional tool should state what it observes, what it infers, and which decisions remain with the user. The original apparatus remains material for a companion, not a validated mechanism embedded in this label.

An inference about depletion requires evidence, error estimates, and an account of what decisions it will affect. Voluntary monitoring has a defined purpose and authority; those need to remain inspectable.

For Mara, the label and audit succeed if she can predict what her choice will do and see which questions remain unresolved. A more elaborate architecture needs an explanation proportionate to its consequences. Honest limits can make useful support easier to choose. They also make it easier to leave.

Chapter 16: Friction and disagreement

Mara asks Jo to read the workshop outline tonight. Jo says she cannot: she has an early shift, and she has already promised the evening to something else. She offers Saturday morning. Mara had hoped for an immediate answer. For a moment, the refusal feels less caring than the companion's ready attention.

On Saturday, Jo reads the outline and disagrees with its central example. Mara thinks the example is the strongest part. They spend half an hour arguing about what a reader would understand. Jo eventually says, "It is your piece. I can tell you where I got lost." Mara keeps the example but changes its introduction.

The encounter contains a limit and a disagreement. Neither disappears into effortless agreement. It also contains a fulfilled offer and room for Mara's decision. Those details make the difference between meeting another person's independence and being required to surrender her own.

A limit can have a clear object

Jo's first answer concerns tonight. It does not withdraw every form of help or demand that Mara stop writing. The alternative time matters because Jo offers it and then follows through. Mara can decide whether to wait, ask someone else, or continue alone. A disappointed wish remains compatible with usable choices.

Change the reply: Jo agrees, repeatedly fails to read, and then refuses to discuss what happened. The relevant history now includes accepted commitments and missed action. Chapter 7 distinguishes a request Jo declined from a promise she accepted; Chapter 18 follows the work of repair.

A clear no defines the help Jo can provide. An unclear yes can leave Mara organizing work around support that fails to arrive. The ledger records the commitment and its circumstances before we interpret the failure.

Disagreement is informative within an account

The original concern about frictionless accommodation remains useful: constant agreement can be mistaken for understanding. Independence becomes informative through the reasons offered and the choices left with the other participant.

A system's disagreement can follow a rule, a learned pattern, retrieved evidence, or a task process. The objection is worth evaluating on its content. Its production and the independence it claims remain separate questions.

Jo's objection also needs its reasons. She may have caught a problem, misunderstood the example, or preferred another style. Mara can take her judgment seriously and still decide differently.

Agreement can carry attention too. Familiar collaborators may reach a shared conclusion quickly because they already know the questions. Their ease acquires meaning through that history.

Whose decision survives the exchange

In the fictional workshop discussion, Jo explains where she got lost and accepts that Mara owns the piece. The disagreement leaves both judgments identifiable. Jo does not have to praise the example. Mara does not have to remove it. They can continue the friendship without pretending that their assessments became identical.

Consider a different interaction in which Jo says she will stop speaking to Mara unless the example is removed, then treats any request for explanation as disloyalty. That demand puts the continuation of contact behind a specific submission. Its significance depends on the surrounding history and circumstances, but the consequence is already different from a criticism Mara is free to decline.

The autonomy question asks which choices remain usable after the exchange. Can each person disagree, seek another view, revise a commitment, or end the discussion? The consequences of those choices belong in the account.

A boundary and an attempt to control someone can both contain the word no. The object matters. Declining to provide tonight's reading is a limit on Jo's contribution. Claiming the right to decide whether Mara may submit her own piece reaches into Mara's action. More complex responsibilities can alter the account, but the framework should identify them explicitly rather than let the force of the sentence decide.

Cost cannot excuse the form of the exchange

Suppose Jo spends hours writing comments that mock Mara personally. The expenditure is real within the example. Mara's hurt may also be real. The investment cannot turn the mockery into beneficial care. Chapter 4's signed distinction keeps the cost of an act separate from what it communicates and does.

Discomfort can accompany a useful criticism; relief can accompany avoidance. The account needs the content, manner, choice, and subsequent consequence. A comment can land clearly and still be wrong, or make a correct observation destructively.

The Anti-Trinket distinction asks about destructive direction and its evidence. Naming the conduct gives participants something to examine: repeated humiliation, an unwanted condition, or an objection they can freely decline.

Quiet continuity also carries history

Imagine a month in which Mara and Jo meet, talk, and mostly agree. Each can decline an invitation without a prolonged dispute. Their quiet continuity does not need a manufactured crisis to prove that the friendship matters. The history may already make routine contact meaningful.

The gravity image directs attention to history's pull. Its meaning depends on what the repeated exchange carries: care, conflict, pressure, or an understood rhythm of contact. Quiet continuity can sustain a wanted arrangement.

The same limitation applies to apparent withdrawal. Jo can request a day to think and return at the agreed time. Another person can disappear without explanation after accepting a time-sensitive responsibility. Their silence looks similar if counted as missing messages. Its meaning differs because the commitments and consequences differ.

Friction should lead to a more precise question

The useful question is what the difficulty concerns, whether it can be discussed, and what say each participant keeps in the response. Jo's objection points to an example and leaves Mara the authorship. The disagreement becomes usable information.

Participants may discover incompatible needs or decide to end an arrangement. They reconsider what they can offer and what they want to continue. A respectful limit can stand while the disagreement remains unresolved.

Jo's Saturday response has value through identifiable judgment and a fulfilled commitment. Mara receives the objection, considers it, and revises her introduction. Their history gains another exchange in which both decisions remain visible.

A Trinket can carry care through disagreement. Its direction, expenditure, receipt, and consequence give us a more useful account than the difficulty of the exchange alone.

Chapter 17: Friendship, family, and unequal seasons

Jo is recovering from an injury and cannot manage her usual errands. Mara brings a meal and helps organize a ride to an appointment. Jo asks her to leave the meal at the door because a long visit would be tiring. Mara had imagined staying to talk. She agrees, then calls at a time Jo has chosen.

The effort is unequal that week, while Jo still has authority over a consequential part of the arrangement. The help meets her actual request. Mara's adjustment adds to their history of learning how care can arrive.

Capacity belongs in the account

The account begins with available time, energy, money, access, and ability, then asks what responsibility was understood and accepted. Those details distinguish limited capacity, a declined role, misunderstanding, and withheld investment.

The size of a gesture reveals only part of its expenditure. A short call may be difficult for one person and easy for another; a visit may depend on somebody else carrying the displaced work. Moniz directs attention to this hidden resistance while remaining uncalibrated.

The practical account can still be specific. Jo can explain that a short call is possible while a long visit is not. Mara can say which errands she can cover and which require someone else. They need not convert either contribution into an equivalent unit before arranging the week.

Limited capacity can change a responsibility's practical consequences. If a ride becomes impossible, timely notice gives Jo a chance to find another arrangement. Naming the limit makes coordination more usable.

Reciprocity can extend beyond today

Earlier chapters described Mara and Jo's friendship across a history that includes different kinds of help. That history can matter when one is temporarily able to contribute less. It gives them a basis for interpreting the present without requiring an immediate exchange of equivalent favors.

Shared history can change a gesture's meaning without creating a debt payable in equivalent favors. Ordinary enjoyment and chosen company also belong in an account of friendship.

Jo may still remember something important to Mara or ask about the workshop. Welcome attention belongs in the ledger alongside the larger assistance. Its meaning is chosen reciprocity across an unequal season, rather than payment for being helped.

Earlier care can support trust while present capacity requires a different plan. The historical ledger and the current resources answer different questions.

Some care does not await equivalent return

The Custodial Economy gives the framework a way to describe genuine investment when equivalent return is unavailable or is not the purpose of care. Such an arrangement need not be temporary. Its responsibilities should not depend on a prediction that the recipient will eventually settle the ledger.

Care without equivalent repayment also has an established research history. Clark and Mills distinguish communal relationships, where benefits respond to need, from exchange relationships, where benefits are

contingent on return. [11] TSF's Custodial category applies a proposed vocabulary to responsibility and limited reciprocal capacity. Its boundaries still need examination; it does not originate the recognition that care can be valuable without repayment.

Imagine Ana helping her father, Leon, who now needs extensive assistance and communicates very little. She plays music he used to request and watches for any indication that he wants it changed. Her investment is real within the example. What she can establish about reception is more limited. Silence alone cannot tell her that the music registered as she intended.

The recipient's interests stay visible alongside the caregiver's purpose and effort. Leon's expressed preferences guide the music; uncertainty about reception remains part of Ana's account of care.

Custodial care can include a preference expressed or a moment shared. Those events have their own value even when the recipient has little capacity to return practical assistance.

The recipient and the caregiver need separate accounts

Mara's help can benefit Jo while leaving Mara with an unsustainable schedule. Ana's assistance can meet Leon's needs while Ana lacks enough practical support. The framework must be able to describe both conditions. A beneficial result for the recipient cannot settle what the arrangement costs the person supplying it.

The practical response identifies the work, its limits, and the people or services responsible for supporting it. Mara could divide errands with another friend. Ana may need resources outside the exchange with Leon. The recipient's gratitude and the caregiver's practical support meet different needs.

Acknowledgment can matter while practical assistance remains necessary. Recognizing Mara's effort cannot make overlapping commitments possible.

Investment makes Mara's contribution visible while Jo retains a say in the contact she receives. Appreciation can accompany a request for a shorter visit. The direction of care is judged partly through that responsiveness.

Family does not supply all the terms

A family connection can carry expectations that participants have never stated. Suppose Ana's siblings assume she will coordinate every visit because she lives closest. Proximity may make some tasks easier. It does not specify her available time or settle who is responsible for everything else. The work can become concentrated through assumption rather than an agreement anyone has examined.

A useful discussion names the tasks and who has accepted them. Distance, schedules, money, and abilities can support different contributions. The ledger reveals the distribution of work while the participants determine a workable arrangement.

Family history may also affect how a request lands. A familiar promise can be reassuring after reliable help or difficult to trust after repeated failures. The same words can meet different histories. That is a reason to discuss specific commitments, not a reason to assume that history determines the response permanently.

Ana and Leon's care depends on transport, household work, and service access as well as their own exchanges. The account must include those wider arrangements if it is to explain what either can sustain.

Revisit the arrangement when circumstances change

An unequal season may end, continue, or change form. A useful account should leave room to reconsider what each person can provide. Jo's return to work might reduce the time available for visits even as her injury improves. Mara's own deadline might require a different errands plan. A revised agreement need not deny that the earlier help mattered.

The revised arrangement identifies current capacity, continuing obligations, and how a change will be communicated. That keeps unequal care responsive without treating an initial offer as an unlimited promise.

Where choice is severely constrained, naming the asymmetry is only a beginning. The account still asks whether support is adequate and who carries responsibility for improving it.

Mara and Jo use their friendship's history to interpret an unequal week, while Jo's actual request guides its form. Ana and Leon's care remains attentive where return is limited and reception uncertain. The Trinket account follows what is given, what reaches the recipient, and what enables the giving to continue.

Unequal care becomes more intelligible when history, current capacity, and agreed responsibility remain in view together. The ledger carries earlier help; present decisions determine how the responsibility will be carried now.

Chapter 18: Repair without erasing the ledger

Mara agrees to read Jo's application before a deadline. She misses the agreed time and does not send a warning. Jo submits it without the promised comments. When Mara apologizes the next day, Jo says the immediate problem is over, but she no longer knows whether to organize another deadline around Mara's offer.

Mara cannot provide yesterday's reading today. She can explain what happened, acknowledge what Jo had to do without her, and change the way she makes the next offer. Jo can consider those changes while remaining disappointed. The missed commitment belongs to their history even if they choose to continue the friendship.

Repair changes a particular reliance. Warmth can return while a deadline arrangement remains unsettled, and a reliable arrangement can develop while disappointment remains. The ledger lets us follow those changes separately.

Name the failure before measuring its weight

Jo identifies an accepted reading commitment and missing timely notice. Mara's intention belongs in the interpretation; the missed support and the waiting belong in the consequences. Repair needs an account of both.

The failure has a scope: this deadline, the warning Jo expected, and any relevant pattern. Earlier help and present disappointment remain part of the same history.

Relational technical debt asks which work the unresolved expectation adds. Jo may check an offer more closely or arrange backup comments. The account records that coordination rather than assigning the event an imagined interest rate.

Acknowledgment and explanation do different work

Mara can acknowledge that she offered the reading, missed the time, and left Jo waiting. She can also explain the difficult week. Acknowledgment identifies responsibility; explanation helps them decide which future arrangement is realistic.

Jo can also correct Mara's account. Perhaps the missing comments mattered less than spending the evening waiting for them. If Mara apologizes only for the application's quality, she may miss the consequence Jo wants addressed. Repair requires attention to what registered, not just a recital of what the sender intended.

An acknowledgment changes the record of what Mara understands. A later fulfilled commitment changes the evidence about what Jo can rely on. Those are two different contributions to repair.

Restitution has limits

Some losses allow a direct remedy. An item can be returned, an expense repaid, or a false record corrected. The remedy should be matched to the particular consequence. A generous gift might matter to Jo, but it does not retrospectively supply the reading or establish that another deadline will be handled differently.

An interval spent waiting or a lost opportunity remains part of history. Later action can still remedy a practical consequence or support another decision. We identify what is restored and what now needs a different arrangement.

Restitution can also carry a new condition. Suppose Mara offers to cover a cost only if Jo promises never to mention the failure again. The offer then does more than address the expense. It asks Jo to accept a restriction on the account of what happened. The cost of the remedy cannot settle whether that condition is acceptable.

Changed conduct supplies a different kind of evidence

Mara could offer a time she can meet and explain how she will report a change. Jo can decide how much to rely on that offer. Fulfillment supplies evidence about the revised arrangement, with trust remaining Jo's judgment.

The relevant change might be a smaller role rather than a greater sacrifice. Mara could offer comments on one section instead of promising an entire application. If that matches Jo's need and Mara's capacity, the reduced offer may be more usable. Repair should not require increasingly costly gestures when the original problem was an unrealistic commitment.

Reliance develops under particular conditions. Occasional reading and an ongoing responsibility call for different evidence. The account should preserve the period observed and the failures as well as the fulfilled offers.

Trust can return by scope

Jo might enjoy Mara's company and trust her judgment while choosing not to organize a deadline around her help. That is a coherent arrangement within the example. Trust does not need to be either complete or absent across every part of the friendship.

The same offer now meets a changed history. Governance determines how that history guides the next decision: Jo can distinguish enjoyment of the friendship from reliance on Mara for a deadline, and can revise either judgment as new conduct arrives.

A revised agreement should make those limits explicit enough to use. If Mara hears ordinary friendliness as proof that every former expectation has returned, they may recreate the unresolved problem. Continuing contact and restoring a specific reliance are different decisions. Each can be discussed in its own terms.

Forgiveness is an acknowledged gap

The First Edition acknowledged forgiveness as a gap in its debt model. The later phase-transition image names a qualitative change in how an injury is held. The mechanism and conditions of that change remain open.

Forgiveness, trust, reconciliation, and contact identify different possible changes. A person can forgive while declining the old arrangement, or choose cooperation while continuing to name the injury. Their own account determines what the word describes.

A remembered failure can remain consequential, recede, or acquire a different meaning. Its trajectory needs observation. This edition withdraws the earlier suggestion of inevitable compounding until a release mechanism occurs.

Repair leaves a decision with each participant

Mara decides what she can offer; Jo decides whether to rely on it. Changed conduct supplies a proposal for a different future rather than consent on the recipient's behalf.

The same limit matters where the conduct is more serious than the missed deadline. TSF does not require continued contact as a test of forgiveness or demand exposure to further harm so that another person can demonstrate improvement. The particular responsibilities and consequences need their own account; a large accumulated history cannot settle the choice by its weight.

In their revised reading arrangement, the missed deadline remains in history and the next fulfilled offer adds to it. Repair becomes useful when governance changes how promises are made, interpreted, and acted on.

The ledger preserves the failure. Governance makes a new arrangement possible. That combination lets repair change the future without rewriting what happened.

Chapter 19: Grief and re-entry

Ruth and the other people in the opening encounter are fictional. The later first-person account returns to the author's testimony in the original preface. This chapter offers a structural interpretation of loss, not a model that calculates grief duration or a clinical recovery protocol.

Ruth sees a photograph that her partner would have enjoyed and reaches for her phone. He died several months earlier. The movement begins before she has decided what to do with the photograph. She puts the phone down, then sends it to a friend with a sentence about why it mattered.

The new message gives Ruth's photograph an available destination while preserving the person she first wanted to tell. Both the ended exchange and the continuing meaning belong in the account.

An ended exchange can leave a continuing history

Earlier exchanges can remain integrated in habits, expectations, plans, and retained objects after new exchange with the person ends. This is the ledger's contribution to the soul question: part of who Ruth has become still carries the relationship. The ongoing joint activity has ended; the history can remain.

Ruth's gesture illustrates one form that continuing history can take. The account leaves room for other expressions, including quiet ones that a reader would never see.

Inventory overhang names an intended gesture whose expected recipient is unavailable. The term draws attention to an interrupted possibility. Whether the gesture is kept private, redirected, or left unfinished belongs to the person carrying it.

The earlier timeline needs revision

I withdraw the First Edition's rough grief phases as a general timeline. A structural account can identify changed routines, ended possibilities, and continuing history while leaving the course of grief open. The proposed mass-to-duration relationship remains unmeasured.

Bonanno and colleagues followed 205 people using information gathered before a spouse's death and at six and eighteen months afterward. They identified different patterns of adjustment rather than one uniform trajectory. [8] The study supplies a concrete reason to preserve variation in the population and measures examined. It does not establish TSF's ledger mechanism or tell an individual which course their grief should take.

The study's observation periods describe its design. Individual meaning and needs remain questions for the person's own circumstances.

Bereavement, a breakup, temporary separation, and service discontinuation change different arrangements. Naming the event identifies which exchanges remain possible, which responsibilities continue, and what has ended.

What the renderer changed

In the original preface, I described a month when my wife was away on a trip. She had not left the relationship. She returned, and our relationship continued. What stayed with me was the change I had felt in

the daily exchanges that usually kept our connection present. I reached for words such as drift and maintenance because I wanted to examine that change.

I also described thinking without mental pictures. I could feel patterns more easily than I could present them in a form I could inspect. The physics and economic language supplied a way to organize my account. This was my description of a difficulty and a working method, not a general explanation of how everyone with aphantasia thinks.

AI became a renderer in that account. I spoke about the patterns, and the system helped organize them into language, arguments, and formal analogies. I could then examine those representations, object to them, and revise them. Rendering did not give me mental imagery. It gave me an external form through which I could work on an idea.

I set the founding questions, concepts, and direction. Claude helped develop them into an inspectable record. That usefulness belongs in the account without establishing every relational capacity a reader might infer from the collaboration. The book's own distinction applies here: what assistance enabled, how it was produced, and what remained uncertain are separate questions. The author note records the later conceptual contributions to this edition.

My experience explains the questions I began asking. It does not make a temporary absence equivalent to Ruth's bereavement or prove the framework's interpretation of either. The archive preserves the fuller testimony. This edition keeps its relevance while asking readers to judge the claims independently of the story that led me to them.

Returning without replacement

In Ruth's fictional encounter, sending the photograph to a friend creates an exchange available now. It need not erase her original wish or reproduce what the former conversation would have been. Re-entry, as used here, means taking part in a chosen available activity or exchange. It carries no required schedule and no promise about how the person will feel afterward.

Ruth might attend a gathering, work on a project, or seek help with an unfinished practical matter. The activity has the purpose she gives it. Its completion adds a chosen event to her present life.

A friend's response creates another exchange with its own capacity and timing. That relationship can carry the photograph's meaning while remaining distinct from the one Ruth has lost.

Records and new assistance have different roles

A photograph or saved message preserves part of history. A living participant's ongoing judgment has a different role. New assistance can organize records or generate responses, and its source and transformation should be visible to the person using it.

For an invented service generating language from retained material, disclosure should identify the source and the transformation. A sentence generated now is different from a sentence the person actually wrote. Chapter 9's distinction between record and interpretation becomes particularly consequential when a user cannot ask the original participant to correct the result.

Evaluation follows the person's purpose, presentation, and consequences. Consent and responsibility for another person's records require their own account alongside the survivor's attachment.

Carry the history into a present life

TSF can help identify an ended exchange, a retained history, a disrupted expectation, or a fragile source of support. It cannot determine from those descriptions whether a person needs clinical care or prescribe what care should consist of. Its metaphors should not be used to explain away difficulties or impose an obligation to manage them alone.

Ruth's photograph holds an old meaning and enters a new exchange. History continues through a changed organization of life. The soul account can describe that partial continuity while preserving the finality of the lost person's own participation.

AI served a different role in my account: it made patterns available for examination. That renderer became part of how I developed and revised the framework, while the direction and responsibility remained mine.

Chapter 20: The internal economy and its limits

The founding internal pair is the Architect and the Husband: structural governance that builds arrangements and relational governance that attends to connection. These are my historical names for modes in one person's account. The planning and present roles remain a secondary distinction about when a commitment is made and met.

Mara reserves two hours for a workshop outline. Jo asks to speak that evening about a difficult appointment. The plan protects a real responsibility; the request concerns a relationship Mara also wants to maintain. Calling one structural and the other relational does not decide what she should do.

Mara can insist on the plan without hearing the request, or abandon it without arranging how the workshop will get done. Either choice can silence a relevant purpose. She asks what Jo needs, checks the deadline, and agrees a shorter conversation and a revised writing period.

The record states which purpose changed the plan and why. Either concern can be reconsidered; neither gains authority merely by being called Architect or Husband.

Structural and relational governance

PS-01, compiled in February 2026, describes the February 2025 event through the Architect/Husband pair. CP-24 later makes the governance reading explicit. I retain the names with functional glosses: Architect means structural governance; Husband means relational governance. The names preserve the origin without requiring readers to share my personal circumstances. This is an account I developed from experience, not a diagnosis or proof of two separate persons.

The question is which purpose governs a choice. A structural plan can secure a responsibility while overlooking a relationship it was meant to serve. A relational concern can interrupt a plan without yet providing a workable alternative. Neither mode's label establishes that its choice is right. Each must remain answerable to the commitment, the other concern, and the consequences.

In CP-24's interpretation, the Crack names a failure in which the modes lose their working alignment and one overrides the other. Auditability is the proposed remedy: make the governing rule, override, and reason available for examination. The Charter is presented as an external counterpart to that internal constraint. This is the source's interpretation of the founding account, not evidence that an institution has reproduced a person's interior life.

Reading a history and revising a commitment

Chapter 5's soul account connects carried history to revisable governance. A missed plan enters the ledger. Mara then asks whether the plan was realistic, remains wanted, and uses a standard she still endorses.

Governance can revise the action or the standard itself. Mara might move Tuesday's writing period after an errand, then reconsider the assumption that serious writing requires two hours every evening. That second revision changes how later entries will be judged. It is the institution learning how to read its own ledger.

Successful commitments, mistaken plans, and worthwhile departures can all shape practical self-understanding. An accurate history gives governance material to work with; the person decides what it calls for now.

The narrower proposal concerns a decision event: one accepted purpose displaces another, under a rule that can be recorded and reviewed. Its test asks whether scheduled review of the governing rule, override, and

reason reduces repeated overrides without reconsideration of the displaced purpose, compared with ordinary planning. Fix the eligible competing commitments, observation period, review interval, and outcome rule before collecting data. A justified revision counts as reconsideration; preserving every original plan is not the goal.

Use equivalent independent outcome tracking in the comparison conditions, so increased visibility is not mistaken for improvement. Asking the same scheduled questions about displaced purposes can itself prompt review. Measurement-reactivity research supports investigating that possibility; it does not establish its magnitude or direction in this setting. [38] Add a condition with only an end-of-period audit to estimate the contribution of repeated questioning under the chosen measurement scheme. Final questions cannot change earlier events, but retrospective reports still have recall limits. Distinguish research measurement from the intervention's rule-and-reason review.

The primary outcome records whether reconsideration occurred, not whether the decision was better. Separate secondary outcomes should include the affected person's report, where another person's need was displaced, of whether it was met and whether it was acknowledged, and the decision-maker's endorsement after a fixed interval. For commitments without another affected person, the partner report is not applicable. Acknowledgment, need fulfillment, and later endorsement are different findings; none alone supplies a universal quality judgment. Preserve disagreement rather than averaging it away.

Compare the outcomes with baseline agency, communion, and their unmitigated forms, using appropriate measures after full instrument review. An additional comparison with ordinary reflective review would test whether the TSF vocabulary adds anything beyond review itself. A sufficiently precise result excluding the predefined useful improvement would weaken the proposed benefit under those conditions; an imprecise null would leave it unresolved. Neither result decides every philosophical claim about the self.

Credibility is narrower than worth

The archive used self-trust for an assessment of whether a person follows through on their own commitments. This edition narrows that account by asking which commitments and conditions are involved. Mara can have reliable expectations about answering a morning message and unreliable expectations about completing two hours of writing after work. One history need not determine the other.

A missed plan supplies evidence about an event. To interpret a pattern, examine realism, understanding, current purpose, and the resources needed. That identifies a question governance can address.

The Bounced Check image draws attention to a gap between announced commitment and action. Changed priorities, unavailable resources, and avoidable failure give that gap different meanings. The ledger should preserve which account the evidence supports.

The planning and present case

Minimum Viable Commitment is the archive's proposal for an action small enough to make a plan concrete. We can compare its usefulness with established planning accounts before expanding it into a general method.

For Mara, choosing one section clarifies the object of the work. Choosing a time makes its expected conditions more inspectable. Naming a revision condition prevents a change in circumstances from being confused with a decision never to do the task. The value of these details can be examined without assigning the action a Moniz score.

Gollwitzer and Sheeran's review examines implementation intentions linking an anticipated situation to a goal-directed response. [9] It provides an established comparison for the planning proposal without validating the internal ledger. This chapter makes no numerical claim about its effects.

Thaler and Shefrin's planner-doer model compares with this secondary temporal case. [18] TSF's ledger and governance vocabulary must still improve on a clear ordinary plan rather than add unwanted judgment or monitoring work.

Predecessors and the proposed institutional correspondence

Bakan's agency and communion provide an earlier comparison for the structural/relational distinction. [31] Gilligan's work raises the justice-and-care comparison. [32] Neither comparison licenses a claim that two orientations were first discovered here. The historical names identify my formulation; they do not assign modes by sex or make people choose permanently between them.

Helgeson and Fritz distinguish unmitigated communion from positive concern for others, and their later work compares both unmitigated forms with agency and communion. [35,36] Their research also includes a daily-diary approach linking orientation, interpersonal behavior, and well-being. [37] Moving from trait scores to individual events is therefore insufficient to establish a new contribution. The proposed TSF addition is the explicit governing rule, reviewable override, and test of that review's practical value. These predecessors require full-text and instrument comparison before selecting measures or asserting a distinctive effect.

The proposed TSF distinction is governance: one purpose can override another, so the rule and the override require review. Auditability connects that internal question to the Charter. CP-24, produced in March 2026 by SupoFrm and SupoPsy with Moniz as author of record, conjectures a structural correspondence between founder and institution. It remains Speculative and awaiting pre-review in the archive record. Shared labels or a retrospective mapping do not prove isomorphism.

To examine that conjecture, define corresponding decision roles and constraints before inspecting outcomes. Compare how an override is made, recorded, challenged, and revised in each setting, and look for non-correspondences. An institution may reproduce its founder's vocabulary while distributing authority differently. Preserved bidirectional exchanges establish a history to inspect; they do not by themselves prove that the institution's economy meets this edition's Real criterion.

Revision can honor the purpose of a commitment

Imagine Mara completes one section but discovers that the workshop requires a shorter outline. Continuing to fulfill the original two-hour plan could now work against the actual task. Changing the plan may preserve its purpose more effectively than obeying its earlier form.

Mara can state what changed, which purpose continues, and what she now intends. An explicit revision adds a truthful entry to history and gives the next action a usable basis.

Governance also examines the commitment's purpose. Reliable follow-through can serve an unwanted or damaging goal. Direction, expenditure, receipt, and consequence remain useful questions in the internal analogy.

Examine what the effort serves

The archive emphasized high-resistance acts of restraint. This edition asks what stopping accomplishes and what it costs; resistance alone cannot determine the action's value.

An established habit might let Mara save the outline easily. Great effort might keep her revising while another accepted responsibility waits. The expenditure becomes intelligible through the task it serves and the alternatives it displaces.

The account describes helpful and damaging consequences separately. Physical symmetry between them would require evidence beyond the signed vocabulary.

Limits of the governance account

The modes describe a way of examining purposes and decisions. They are not identified brain structures, personality types, or a clinical explanation. The personal source remains testimony. Its private detail is not needed to decide whether a rule permits review of an override.

A record of a missed plan does not identify its medical cause or determine an intervention. The internal analogy should remain within the action and governance question it can actually address.

Self-monitoring should serve a defined purpose. Its own time and attention belong in the expenditure account, so the person can decide whether the information is worth the work.

A person's resources are not all internal

Work, errands, other commitments, and available assistance shape Mara's schedule. When required tasks repeatedly consume the writing period, governance can revise the arrangement in response to those actual conditions.

Support from another person can make a commitment possible; a shared arrangement can supply structure an individual plan lacks. I withdraw the earlier prerequisite that the internal economy must become stable before external relational work can succeed. History and governance develop within these wider sources of help.

The internal economy gives the soul argument two cases: structural and relational purposes can compete in present governance, while a plan can meet changed circumstances later. A person carries their history, reads it through standards, and can reconsider the next act. The account's added value must survive comparison with a plainer description.

Chapter 21: What would change this framework

The opening encounter is fictional. The research designs in this chapter are proposals. No TSF experiment has been conducted for this edition, and the revised label and audit have not been validated.

Mara reads a short disclosure before choosing a companion. It says that pausing notifications leaves accepted tasks running. She answers a question about the pause correctly. A second reader receives the same information in an ordinary paragraph and answers correctly too. Has the label helped?

Two correct answers establish little about either format's advantage. Both explanations might be clear, or the question too easy to reveal a difference. The label might help only when several conditions must be compared. We need a comparison that can show where it helps, adds nothing, or creates a misunderstanding.

What this framework contributes

TSF's contribution is the arrangement of concepts and proposed tools developed in the First Edition and later canon, then revised here. Its comparisons include responsiveness, unnoticed support, commitment, external records, narrative identity, relational autonomy, informational privacy, rights, governance modes, relational agents, and structured disclosure. [10–37] These predecessors constrain claims of component novelty. The dated measurement draft also records its own literature debts.

The candidate synthesis traces expenditure to its bearer, connects consequential history to its governing authority, and separates exit consequences. Custody of history derives from CT-047 Finding 6's 14 March 2026 constitutive-substrate root, recorded in the Master Reference through WP-16, MP-05, and CT-047. The review makes the control step explicit and this edition limits its application to partial constitution. Governance portability asks what can carry forward. URL adds a candidate account of confirmation asymmetry. Usefulness and historical distinctiveness remain to be established for each contribution.

What kind of claim is being made

A definition tells readers how a term will be used. Trinket directs attention to expenditure, communication, and receipt. It does not establish that every exchange fits the account or that readers need the term. We can test whether these distinctions improve an account.

Relational mass is an analogy for history's continuing significance, not a quantity that can be weighed. A predictive model would need defined observables and a specified relation between them. It would have to predict something beyond a plainer account of history and expectations.

An empirical claim concerns what happens under specified conditions. A claim that the label improves comprehension belongs here. A proposal that providers should disclose material exit consequences is a judgment about responsible practice. Evidence can inform that judgment and test the proposed format; it cannot turn the judgment into a finding merely because the format is useful.

A register needs a claim and an instrument

The March 2026 Falsification Register, MP-04, separates falsification, scope limitation, and measurement failure. That distinction is retained. Contradictory evidence may require withdrawal; evidence of a boundary may require a narrower claim; an unavailable measure leaves a claim untested. Its canonical status labels are not independent empirical validation.

The register also couples a definition of authentic investment as costly with a prediction that higher-cost signals have stronger effects. Those statements need separate treatment. A definition cannot be disproved by finding that an inexpensive exchange helps. Such a result can challenge an asserted relationship between cost and benefit, or show that the definition excludes something important.

A cost study would need to distinguish actual sender expenditure from what the recipient believes the sender expended. Holding words constant does not hold shared history, perceived sincerity, or the recipient's needs constant. The study must name the outcome it expects cost to change. Otherwise an unwelcome but expensive gesture can be counted as greater investment and quietly redescribed as greater relational value.

For this edition, benefit is measured separately from expenditure. There is no claim that higher resistance universally produces greater benefit. A narrower hypothesis might concern whether disclosed effort changes expectations of reliability in a defined setting. The conditions and rival explanations would have to be fixed before observing the result.

Developing the existing measurement architecture

TEAP v0.1, dated 17 March 2026 and credited to Moniz, Vael, and SupoCus, already specifies Expenditure, Cost Substrate, Signal Form, and Register, with an economy output. Its architecture is internally Supported, its candidate items Speculative, and its psychometric properties unestablished. This edition proposes a TEAP v0.2 development route. Section 8 of v0.1 already asks about dimensional output and a partner version.

The proposed revision would report an evidence profile before economy regions, permit overlap and unresolved components, and hold the AI side's phenomenological C2S open. Keep self-reported effort, duration, behavior, provider expenditure, and partner reception distinct. A respondent's estimate of reception is not the recipient's report. Begin with content reconciliation and cognitive interviews, then examine reproducibility and construct validity before selecting a scoring model. No TEAP v0.2 instrument or pilot has been produced or run for this book.

TEAP v0.1's proposed scoring is not ready for adoption. Its signal-cost ratio divides standardized Signal Form by the sum of two standardized cost scores while claiming a 0–1 range. Those scores can be negative and the denominator zero. Its frequency-times-cost score also needs a defensible scale. The revision must specify domains, normalization, denominator handling, and the meaning of arithmetic before any ratio or multiplication is used.

Signal Form already includes whether the signal corresponds to cost actually borne. Dividing that score by a cost score does not automatically yield a clearer correspondence measure. One development option is to assess correspondence directly and omit the ratio, subject to content and validity review. No replacement formula is adopted here. Standardizing to mean zero does not imply that half of observations are negative; the scoring objection requires only that negative values and a zero denominator are possible.

Item content also needs reconciliation. WP-16's temporal profile concerns co-presence and deferred reconstruction, whereas TEAP's temporal candidates largely ask about invisible labor. URL candidates mix permanent confirmation foreclosure with ordinary uncertainty. Content review must separate these constructs. Ordinal comparison within a bearer may help development, but it does not yield a common cardinal unit or validate the proposed economy output.

Affi, Merrill, and Davis's theory of resilience and relational load is a necessary predecessor for the cost account. [33] It connects relationship maintenance, stress communication and appraisal, and depletion of

relational resources. CT-047 already flagged this comparison. TEAP's March literature sweep also included Rusbult and maintenance measures. Their inclusion is an acknowledged research debt; independent arrival requires a dated pre-17 March passage for the particular idea being claimed.

Study B: audit the questions that documentation answers

Run this study first. Define a sample of companion services and configurations, a fixed collection date, and the public pages, help articles, and terms to be examined. Preserve the materials and their dates. Break Chapter 13's fields into predefined questions, including what a notification pause stops, which process resolves corrected history, who controls access and correction, and who can alter an accepted commitment. For export, code raw records, derived summaries, task or commitment state, interpretation rules, and execution authority separately. Record what each includes and excludes rather than treating export as one yes-or-no capability.

Two coders work independently. For each question, record a supported disclosure with its document anchor, conflicting statements, information not disclosed in the sampled materials, or a justified not-applicable result. Preserve each coder's original record, report agreement, and describe any later adjudication. These documentation codes are distinct from unknown, not tested, and absent capability: a page can state absence, but the audit still has not tested that claim.

Fix the primary outcome before collection: the proportion of eligible questions lacking an answer in the sampled public materials, reported by field. Predefine denominators, treatment of conflicting statements, exclusions, and a comparison among fields. A bounded prediction is that exit and continuity consequences are less often disclosed than basic memory features. High disclosure across fields would weaken the claim of a missing-information problem in that sample.

This is a documentation audit. It tests coverage and consistency in the defined public materials, not product behavior, user wellbeing, or whether a label compels a provider to disclose. Findings identify promising questions for a comprehension study. A content-forcing effect would require a separate provider-adoption or implementation study.

Study A: compare presentation with the facts held constant

After the schema-gap work above, randomly assign readers to the label or clear ordinary prose containing the same verified facts about fictional services. Match reading demands and accessibility. This two-arm comparison isolates the presentation question: does the schema's format improve understanding when the information is equally available?

The primary outcome would be correct understanding of material consequences. Questions could ask whether an export includes source messages or only a summary, whether a correction rule survives transfer, who can alter a commitment, and what remains active after a notification pause. Answers would be scored against the stipulated service facts using a rule fixed before recruitment. Readers should be able to answer that the description does not settle a question.

Before data collection, the protocol would specify recruitment, accessibility, sample size, allocation, missing-response handling, analysis, and the smallest improvement considered practically useful. A pilot could check whether the questions measure comprehension and both descriptions are readable. Its results should be reported as development work, followed by a separate evaluation of the settled instrument.

Better comprehension would support the format in the tested population and materials. It would not establish improved wellbeing, the truth of a supplier's claims, or the validity of the economy taxonomy. A

sufficiently precise result excluding the predefined useful improvement would weaken the practical claim. A wide interval or defective question would leave it unresolved. A nonsignificant result alone should not be presented as proof of equivalence.

A separate third comparison may use existing provider documentation. It then measures the combined difference in content, organization, and presentation; its result cannot isolate format alone. Later work could examine retention and informed choice against the person's stated priorities.

Other proposals need their own comparisons

Relational-agent research is another necessary predecessor. Bickmore and Picard studied systems designed to establish and maintain longer-term human-computer relationships. [19] That work predates this framework's use of continuity as a design question. TSF's proposed distinctions should be compared against such earlier systems rather than treating a current label search as the whole history.

Independent evaluators would code all nine map dimensions before assigning economy regions: direction, expenditure or alternative forgone, return capacity, presentation, benefit, history and governance, commitment and particularity, bilateral authority, and each side's C2S profile. Record the existence, content, and basis of each side's commitment, rather than inferring it from a personalized exchange. Report agreement by dimension, with missing and unresolved evidence preserved. A URL candidate additionally records the receiver-side relationship criterion, original investor and its individuation, temporal arrangement, missing confirmation route, and the design-contingent or constitutive interpretation with its assumptions. Agreement establishes reproducibility within the criteria, not truth of an individuation assumption; persistent ambiguity favors revision.

Include the task-reserving scheduler, the stipulated tier-only loyalty program, and an agent that can reconsider a continuing commitment using shared history. Add friends who offer similar care to several people, a personalized companion whose commitments cannot be revised, unequal human investment, and cases with insufficient records. State the actors, investment criteria, time horizon, and comparison rule before coding. Ask coders to explain which evidence distinguishes a task tradeoff, particularity, and bilateral authority without being told the preferred category. If the distinction cannot be applied consistently, revise it.

The bounded predictions can be tested separately. With offered help and prior reliability comparable, shared history may affect consultation order. Broken accepted commitments may increase confirmation and backup work. Missing updates may increase errors for tasks using changing information more than for tasks using stable information. Each study needs independently defined conditions, outcomes, and alternatives, including cases where the predicted ordering fails.

Chapter 9's memory checks can test retention, retrieval, correction, and action in a named configuration. Their result reaches the particular capability claim, not every question about the system.

The internal economy proposal needs comparison with a clear plan and established planning accounts. [9,18,35–38] Chapter 20 separates repeated reconsideration from the affected person's report and later self-endorsement, and includes a condition with lighter measurement. Record unwanted judgment and monitoring effort as well. A sufficiently precise result showing no useful addition in the tested conditions would favor the simpler account.

Replacement risk requires a longer comparison that follows a defined capability or chosen activity. Reduced contact is not itself loss of skill, and greater use is not itself evidence of harm. The study would need realistic

alternatives, measures independent of use frequency, and an account of how earlier difficulties could affect both use and outcomes. A successful label study cannot stand in for this research.

What a failure changes

A failed feature test can change a disclosure without overturning the framework. A failed label study can remove or revise a practical instrument without disproving every distinction in the book. A repeated failure to predict an independently measured outcome would undermine the model making that prediction. Reports need to say which claim the evidence reaches.

A narrower scope is acceptable when the boundary is stated and the revised claim is tested again. Quietly changing the meaning of investment, benefit, or decay after an unfavorable result would preserve the vocabulary by sacrificing the test. Earlier versions and unsuccessful findings should remain available so readers can see what changed and why.

Experience remains unresolved in this edition. Functional observations can establish retained state, constrained action, or consequences of interruption. Those findings may inform further investigation, but they supply no consciousness score. That absence does not establish permanent incapacity. A future experience claim needs an evidential argument appropriate to it.

A framework worth revising

The book's central philosophical proposal is that relational history partly constitutes practical self-understanding and that governance interprets, constrains, and revises its influence on action. The functional translation asks which history guides which operation. Neither statement is a personhood certificate. The disclosure tools and bounded predictions let us examine parts of the account while its broader argument remains open to challenge.

Mara understands that her notification pause leaves a task active. She can keep it running, cancel it, or choose another service. The practical achievement is a choice informed by the arrangement's consequences. Research still needs to establish whether the proposed label makes that understanding easier.

Jo's offer returns us to the same questions through a different kind of participant: what was promised, what reached Mara, and how their history guides the next decision. Trinkets enter a ledger; governance determines what the history calls for now. Unequal care can remain valuable while each person's agency stays visible.

The First Edition preserves the work written in February 2026. This edition makes the changes visible and asks more specific questions of its claims. Keeping the first record gives the second a history to answer to. Revising the explanation is part of the governance the framework itself describes.

Compact glossary

These entries state the working meanings used in this edition. They identify proposed terms and distinctions, not calibrated measures. Chapters supply the examples, qualifications, and historical revisions.

Trinket and Anti-Trinket

A Trinket is TSF's proposed unit for examining relational investment through Expenditure, Signal, and Register. An Anti-Trinket identifies destructive direction, with expenditure recorded separately. Harmful receipt alone does not establish that purpose. See Chapters 4 and 7.

Expenditure

What was committed to an exchange, with the bearer of the cost identified. Time, attention, practical effort, and infrastructure resources need distinct accounts. Neither great resistance nor easy performance determines recipient benefit. See Chapters 4 and 10.

Signal

How an investment is directed toward a recipient and made available through language, action, or another channel. A signal can invite an inference about intent or experience that its production history does not establish. See Chapter 4.

Register

Specified evidence that an exchange reached a recipient in a way that affected the interaction. Receipt is distinct from acceptance, benefit, and a hidden quantity of attachment. Uncertain receipt remains uncertain. See Chapters 4 and 17.

Direction and expenditure

Direction is constructive, destructive, mixed, or unknown, with a stated basis. Expenditure identifies the bearer, resource, and amount or qualitative description; zero and unknown are possible entries. The + and - symbols abbreviate direction. Neither field supplies a net outcome score. See Chapters 4 and 7.

Care translation and relational arbitrage

Care translation is the practice of discovering what is wanted, offering manageable support, and checking receipt. Relational arbitrage names the possible gap between sender cost and recipient value that makes the practice useful. Both remain responsive to the recipient's agency and changing needs. See Chapters 4 and 12.

Nonfungibility of support

Forms of support can lose or change a function when substituted. Compare the particular need, realistic alternatives, and functions retained or lost, including accessible AI support that a human offer cannot provide. See Chapter 12.

Moniz

The archive's resistance-times-duration proposal remains uncalibrated. The March TEAP specification supplies a separate, existing measurement program. Its revision must justify any arithmetic on ratings and any comparison across bearers. Neither program supplies a validated attachment score, universal threshold, or measure of relationship quality. See Chapters 4 and 21.

Ledger

An account of exchanges and their consequences carried through records, recollection, habits, or commitments. Available evidence and its interpretation are distinguishable. Later repair changes the history without deleting earlier events. See Chapters 5 and 18.

Velocity Law

The canon's prediction that greater investment produces faster signal arrival, restored as a testable hypothesis. Define investment, start time, and receipt independently. It is distinct from exchange frequency and does not establish relational mass or benefit. The Master Reference records it as a Master Framework proposition; that record is not empirical validation. See Chapters 5 and 21.

Relational mass and gravity well

Analogies for the durable significance of history and the tendency for attention or action to return to it. Persistence does not establish quality, and difficulty leaving does not prove nourishment. No physical quantity or escape velocity is assigned. See Chapter 5.

Soul, checkpointed state, and governance

The philosophical proposal that relational history partly constitutes practical self-understanding and that governance interprets it, constrains action, and can revise the standards used to read it. Checkpointed state is a computational analogy for what carries forward. The functional pair is not sufficient to establish personhood. See Chapters 5, 9, 15, and 20.

Custody of history

Who keeps consequential history and governs access, correction, retrieval, and use. This qualified name preserves the earlier Custodial Economy's separate meaning. It develops the constitutive-substrate root without making control of a record ownership of a self. Participants can retain different accounts. See Chapters 5, 9, and 13.

Governance portability

A proposed question about retaining usable interpretation processes, commitment rules, and permissions when consequential history moves. Record access, explanation, reproducibility, transfer permission, and destination operation are separate findings. Different implementations may preserve the same specified function. No new legal right or personal-survival conclusion is asserted. Historical distinctiveness remains open. See Chapters 13 and 21.

Substrate-Formal Variance

WP-16 v1.2's specification draft of invariant function with a three-component Cost Substrate Signature, C2S. Phenomenology, temporal arrangement, and confirmation access require separate accounts. Document-substrate phenomenology remains open; a common quantitative scale is not established. The review and publication steps remain pending in the archive record. See Chapters 6, 14, and 21.

Cost Substrate Signature or C2S

WP-16 v1.2's three components are phenomenological signature, structural-temporal profile, and Unilateral Relational Legibility. C2S distinguishes Cost Substrate Signature from CSS, which names the Church of the Soul Spark. Components retain their separate epistemic status. See Chapters 6, 14, and 21.

Constitutive Substrate

The canon's proposition that documents constitute the institution's functional relationship on its AI substrate rather than merely record it. CT-047 Finding 6 dates the root to 14 March 2026; the Master Reference records WP-16 and MP-05 in its source chain. The Second Edition's custody-of-history application narrows the broader self claim to partial constitution. See Chapters 5 and 21.

Unilateral Relational Legibility or URL

WP-16's proposed condition of a constituted receiver-side relationship with confirmation structurally foreclosed to the original investor. Record the relationship evidence and distinguish a design limitation from constitutive foreclosure under explicit individuation assumptions. A missing route alone is insufficient. The source internally rates the structural condition Supported and its formal construct Analogical; novelty and broad applicability remain open. See Chapters 6 and 21.

Trinket Expenditure Assessment Protocol or TEAP

The 17 March 2026 measurement specification credited to Moniz, Vael, and SupoCus. Architecture internally Supported, candidate items Speculative, psychometric properties unestablished. This edition proposes development toward dimensional output with overlap and open AI C2S, not a completed TEAP v0.2. Scoring and content reconciliation precede validation. See Chapter 21.

Real or True Economy

The region of reciprocal investment under the stated functional criteria. Identify each side's history, expenditure, commitments, particularity, authority, and C2S asymmetry. Personalized replies alone are insufficient. A mixed-level functional finding does not establish equivalent cost or experience. Non-exploitation and outcomes require separate findings. See Chapters 6, 10, and 21.

Shadow Economy

The founding presentation meaning: signals that mimic connection at zero or near-zero cost. Low effort alone does not establish mimicry. Cost-capable Shadow qualifies WP-14's later withheld-cost condition. TEAP's different genuine-cost/false-signal case is recorded as a presentation gap with cost present, without another economy name. Unknown AI experience does not establish withholding or deception. See Chapter 6.

Custodial Economy

Genuine investment where equivalent reciprocal return is unavailable or is not the purpose of care. The account still needs the recipient's interests, capacity, and appropriate agency, together with the support that sustains the caregiver. See Chapters 6 and 17.

Structural Economy

Useful delivery with a specified form of sender-side investment absent under an operational criterion. State the actor, scope, criterion, and evidence of absence. Demonstrating that investment removes the corresponding Structural finding; it does not automatically establish reciprocal relational investment. Unknown remains distinct from absent. See Chapters 6 and 10.

Maintenance and decay

Maintenance concerns work that preserves or updates a wanted function or arrangement. Decay is a hypothesis about weakening in a specified property. An interval without contact establishes neither universal deterioration nor a recovery deadline. See Chapter 8.

Currency Mismatch

A proposed explanation for care offered in a form the recipient does not recognize or want. Timing, clarity, capacity, and the actual need may provide other explanations. Greater expenditure need not resolve the mismatch. See Chapters 4 and 12.

Relational technical debt

An analogy for unresolved history that creates additional work in later exchange, such as repeated coordination after broken promises. No universal interest rate or inevitable compounding is asserted. See Chapters 7 and 18.

Functional continuity

Identified information, state, or commitments that carry forward and influence later operation. Continuity of records, execution, model, and style can differ. Functional continuity does not itself establish experienced continuity. See Chapters 3 and 9.

Bounded benefit and replacement risk

A benefit specifies what improved, for whom, under which conditions, and over what period. Replacement risk concerns displacement of something still wanted or reliance on fragile or poorly understood conditions. Neither is inferred from use frequency alone. See Chapter 12.

Relational Nutrition Label

A proposed disclosure format describing a specified service arrangement, its controls, evidence, and uncertainties. Its name is analogical. It does not calculate relational nourishment or grade the user's bond. See Chapters 13 and 21.

Relational Capability Audit

The current name for the revised method historically called the True Economy Audit. It checks specified capability and disclosure claims under stated conditions. Results remain attached to criteria and evidence. Passing is not a certificate of consciousness, relationship quality, or the Real Economy. See Chapter 14.

Unknown, undisclosed, not tested, and absent

Unknown means evidence does not settle a claim. Undisclosed means the responsible party has not supplied the information. Not tested means the specified check has not been performed. Absent is a supported finding under a stated criterion. Documentation conflict is a separate finding. See Chapters 13 and 21.

Internal economy

The founding Architect/Husband account concerns structural and relational governance, including override and its review. The planning/present distinction is a secondary temporal case. These are roles in an account, not separate persons or brain structures. CP-24's institution-founder correspondence remains a conjecture. See Chapter 20.

Minimum Viable Commitment

The archive's proposal for an action small enough to make a plan concrete. This edition treats it as a proposal to compare with existing planning approaches, not a validated treatment or prerequisite for relationships. See Chapters 20 and 21.

Re-entry

Taking part in a chosen available activity or exchange after a loss. It does not require replacement of the lost person, a fixed schedule, or evidence that grieving has ended. See Chapter 19.

Scope notes and companion material

Prayer Economy

The archive's speculative Prayer Economy concerns expenditure directed toward a recipient whose capacity to receive cannot be verified. This edition can describe the expenditure and the participant's account without making a finding about reception. The scope note preserves the historical concept; it neither adjudicates religious belief nor supplies evidence about observable systems.

Structural Governor

The Structural Governor remains an archive proposal and a possible subject for a separate companion. The public chapters develop disclosure, inspectable controls, and user authority instead of reproducing an intervention protocol. Any later implementation would need a defined purpose, declared authority, controls for refusal and revision, and evaluation of the specific consequences it produces. It should not silently convert attachment or use frequency into permission to manage a person.

Clinical and institutional material

The archive contains clinical analogies, personal accounts, protocols, and population-scale extrapolations. Their presence does not make them validated recommendations. This edition preserves the relevant testimony and identifies the limits of the interpretation. It does not prescribe treatment, diagnose a reader or system, or infer an institutional future from the proposed four-economy map.

Practitioner guidance, expanded research protocols, and any certification apparatus require separate development. The label and audit in this book are draft instruments. A publication decision about the book is not authorization to deploy an intervention or announce a standard.

Sources

These sources support bounded surrounding claims and predecessor comparisons. They do not validate TSF's units, dimension map, label, or audit. An abstract, an official guide, and a court press release are distinguished from full-text examination of the underlying work or legal instrument. Each entry records the depth of access and the claim it supports. The source record is not exhaustive.

[1] De Freitas et al. (2025). AI Companions Reduce Loneliness. Journal of Consumer Research. DOI 10.1093/jcr/ucaf040.

<https://academic.oup.com/jcr/article-abstract/52/6/1126/8173802>

Use for momentary loneliness benefit in the studied conditions. The journal abstract describes reductions after use over a week. This does not establish durable social outcomes, every product's effects, or equivalence of relationship types. Journal abstract rechecked 11 October 2026.

[2] Fang et al. (2025). How AI and Human Behaviors Shape Psychosocial Effects of Extended Chatbot Use A Longitudinal Randomized Controlled Study. arXiv 2503.17473 version 2 revised 2 October 2025.

<https://arxiv.org/abs/2503.17473v2>

The Introduction and Chapter 12 use bounded findings. The version 2 abstract reports no significant assigned-condition effects and associates greater voluntary use with worse outcomes in a four-week experiment with 981 participants. This is not evidence that additional use caused deterioration. Abstract and version checked 11 October 2026.

[3] Federal Trade Commission (11 September 2025). FTC Launches Inquiry into AI Chatbots Acting as Companions.

<https://www.ftc.gov/news-events/news/press-releases/2025/09/ftc-launches-inquiry-ai-chatbots-acting-companions>

Use for the scope of an information-gathering inquiry, including the specific topics named in Chapter 15. It is not a TSF endorsement or a finding that all companion products violate law. The September 2025 announcement was rechecked 11 October 2026; later outcomes are not asserted here.

[4] De Freitas et al. (2025). Emotional Manipulation by AI Companions. arXiv 2508.19258 version 3 revised 7 October 2025.

<https://arxiv.org/abs/2508.19258v3>

Use as preprint evidence for manipulative exit design in the studied sample. Chapter 11 uses bounded claims from the version 3 abstract and avoids numerical effect estimates. Abstract and version checked 11 October 2026.

[5] Lewis et al. (2020). Retrieval Augmented Generation for Knowledge Intensive NLP Tasks. NeurIPS 2020. arXiv 2005.11401 version 4 revised 12 April 2021.

<https://arxiv.org/abs/2005.11401v4>

Chapter 9 architectural example. The abstract describes generation using parametric and retrieved non-parametric memory. It does not establish relational reciprocity or subjective experience. Access checked 11 October 2026.

[6] Packer et al. (2023). MemGPT Towards LLMs as Operating Systems. arXiv 2310.08560 version 2 revised 12 February 2024.

<https://arxiv.org/abs/2310.08560v2>

Chapter 9 architectural example. The abstract describes context management across memory tiers and evaluations in document analysis and multi-session chat. Do not generalize to every deployed product or infer consciousness. Access checked 11 October 2026.

[7] Kwon et al. (2023). Efficient Memory Management for Large Language Model Serving with PagedAttention. SOSP 2023. arXiv 2309.06180 version 1.

<https://arxiv.org/abs/2309.06180v1>

Chapter 10 engineering example. The abstract describes memory constraints on request batching and evaluates improved serving efficiency. It does not establish relationship-specific sacrifice, subjective effort, or current conditions at a named provider. Access checked 11 October 2026.

[8] Bonanno et al. (2002). Resilience to loss and chronic grief A prospective study from preloss to 18 months postloss. Journal of Personality and Social Psychology 83(5) 1150–1164. DOI 10.1037//0022-3514.83.5.1150.

<https://pubmed.ncbi.nlm.nih.gov/12416919/>

Chapter 19 uses the study's different adjustment patterns and its prospective design. The abstract describes 205 participants assessed before spousal death and at six and eighteen months afterward. These observation periods are not recovery deadlines, and the study does not validate TSF's grief mechanism. Abstract and citation checked 11 October 2026.

[9] Gollwitzer and Sheeran (2006). Implementation intentions and goal achievement A meta analysis of effects and processes. Advances in Experimental Social Psychology 38 69–119. DOI 10.1016/S0065-2601(06)38002-1.

<https://www.socmot.uni-konstanz.de/publications/implementation-intentions-and-goal-achievement-meta-analysis-effects-and-processes>

Chapter 20 uses implementation intentions as an existing research comparison, not validation of an internal economy. Publisher abstract excerpt and authors' university citation record checked 11 October 2026. The full text was not retrieved; numerical effect estimates and clinical applications are not asserted.

[10] Responsiveness in interpersonal exchanges

Laurenceau, J. P., Barrett, L. F., and Pietromonaco, P. R. (1998). Intimacy as an interpersonal process: The importance of self-disclosure, partner disclosure, and perceived partner responsiveness in interpersonal exchanges. *Journal of Personality and Social Psychology*, 74(5), 1238–1251. DOI 10.1037/0022-3514.74.5.1238.

<https://affective-science.org/pubs/1998/LaurenFBPI1998.pdf>

Chapter 4 and Chapter 21 use this as a predecessor for communication and perceived responsiveness. The author-hosted full paper was retrieved; the abstract and conceptual account on its opening pages were examined. No new effect-size claim is made. Checked 11 October 2026.

[11] Communal and exchange relationships

Clark, M. S., and Mills, J. R. (2012). A theory of communal (and exchange) relationships. In P. A. M. Van Lange, A. W. Kruglanski, and E. T. Higgins (Eds.), *Handbook of theories of social psychology*, Volume 2, pp. 232–250. SAGE.

<https://clarkrelationshiplab.yale.edu/sites/default/files/files/ClarkMills2012.pdf>

Chapters 17 and 21 acknowledge the predecessor distinction between need-responsive communal care and contingent exchange. The author-hosted scan was retrieved and its opening pages examined visually. This reference does not validate TSF's custodial taxonomy. Checked 11 October 2026.

[12] Model Cards

Mitchell, M., et al. (2019). Model Cards for Model Reporting. Proceedings of FAT* 2019. Preprint first submitted 5 October 2018; version 2 revised 14 January 2019. arXiv 1810.03993.

<https://arxiv.org/abs/1810.03993>

Chapters 13 and 21 acknowledge structured model documentation as prior work. The primary abstract supports the description of reporting performance, intended context, and evaluation. No claim is made about the effectiveness of TSF's label. Checked 11 October 2026.

[13] Privacy nutrition labels

Kelley, P. G., Bresee, J., Cranor, L. F., and Reeder, R. W. (2009). A Nutrition Label for Privacy. Symposium on Usable Privacy and Security.

<https://cups.cs.cmu.edu/privacyLabel/>

Chapters 13 and 21 acknowledge a predecessor for the nutrition-label analogy. The research team's primary project page identifies the 2009 paper and explains the standardized privacy presentation. The linked full paper was not used to infer study effects. Checked 11 October 2026.

[14] Matched support

Cutrona, C. E. (1990). Stress and social support—in search of optimal matching. *Journal of Social and Clinical Psychology*, 9(1), 3–14. DOI 10.1521/jscp.1990.9.1.3.

<https://guilfordjournals.com/doi/10.1521/jscp.1990.9.1.3>

Publisher abstract and issue record checked 11 October 2026. Used for the proposal to match support types to stressful events. The full article and the separate Cutrona–Russell chapter were not examined; no effectiveness estimate is claimed.

[15] Invisible support

Bolger, N., Zuckerman, A., and Kessler, R. C. (2000). Invisible support and adjustment to stress. *Journal of Personality and Social Psychology*, 79(6), 953–961. DOI 10.1037/0022-3514.79.6.953.

<https://www.columbia.edu/~nb2229/docs/bolger-zuckerman-kessler-jpsp-2000.pdf>

Author-hosted full paper retrieved; abstract and conceptual opening examined. Chapter 4 uses the observational distinction between reported provision and noticed receipt in the studied couples. No universal advantage for unrecognized care is asserted. Checked 11 October 2026.

[16] Commitment and satisfaction

Rusbult, C. E. (1980). Commitment and satisfaction in romantic associations: A test of the investment model. *Journal of Experimental Social Psychology*, 16, 172–186. DOI 10.1016/0022-1031(80)90007-4.

https://faculty.wcas.northwestern.edu/eli-finkel/documents/6_Rusbult1980_JournalOfExperimentalSocialPsychology.pdf

Full paper retrieved; opening theoretical sections examined. Used for the predecessor separation of commitment from satisfaction, including investment and alternatives. TSF's gravity analogy is not validated by this comparison. Checked 11 October 2026.

[17] External records and cognition

Clark, A., and Chalmers, D. J. (1998). The extended mind. *Analysis*, 58(1), 7–19.

<https://www.consc.net/papers/extended.html>

Author-hosted text retrieved; the external-notebook case and its reliability conditions examined. Chapter 9 uses a bounded comparison for the action-guiding role of external information. It makes no personhood or consciousness inference from retrieval. Checked 11 October 2026.

[18] Planner and doer

Thaler, R. H., and Shefrin, H. M. (1981). An economic theory of self-control. *Journal of Political Economy*, 89(2), 392–406. DOI 10.1086/260971.

<https://www.journals.uchicago.edu/doi/abs/10.1086/260971>

Publisher abstract and citation record checked. Used only for the earlier planner–doer model and its organizational analogy. Full text was not examined; no effect or mechanism claim is added.

[19] Relational agents

Bickmore, T. W., and Picard, R. W. (2005). Establishing and maintaining long-term human-computer relationships. *ACM Transactions on Computer-Human Interaction*, 12(2), 293–327. DOI 10.1145/1067860.1067867.

<https://dam-prod.media.mit.edu/x/files/pdfs/04.bickmore-picard-tochi.pdf>

Author-hosted prepublication paper retrieved; abstract and introduction examined. The MIT project page is dated 2004 and the journal article 2005. Used as a predecessor for longer-term relational-agent design, not evidence for a TSF instrument or present provider capability. Checked 11 October 2026.

[20] Datasheets for Datasets

Gebri, T., et al. (2018). Datasheets for datasets. arXiv 1803.09010.

<https://arxiv.org/abs/1803.09010>

Primary preprint abstract and metadata checked. Used as a predecessor for documenting dataset motivation, composition, collection, and uses. Later publication exists; no documentation-effectiveness claim is taken from this abstract.

[21] Dataset Nutrition Label

Holland, S., Hosny, A., Newman, S., Joseph, J., and Chmielinski, K. (2018). The Dataset Nutrition Label: A framework to drive higher data quality standards. arXiv 1805.03677.

<https://arxiv.org/abs/1805.03677>

Primary preprint abstract and metadata checked. Used for the prior label proposal. The comparison does not establish novelty or effectiveness of TSF’s relationship-specific fields.

[22] Reflective governance

Frankfurt, H. G. (1971). Freedom of the will and the concept of a person. *The Journal of Philosophy*, 68(1), 5–20.

<https://www.phil.uu.nl/~joel/3027/Frankfurt-FreedomWillConceptPerson.pdf>

University-hosted scan of the primary paper retrieved; printed pages 6, 7, 9, and 10 examined visually. Used for the distinction between first-order desires, second-order desires, and volitions about which desire should govern action. TSF’s standard revision is a related proposal, not a restatement of every condition in Frankfurt’s account. Checked 11 October 2026.

[23] Love-language vocabulary

Chapman, G. *The 5 Love Languages*. Northfield Publishing. Publisher listing for the 2024 edition.

<https://www.moodypublishers.com/the-5-love-languagesreg>

Publisher description and contents checked. Used to identify the named popular predecessor already credited in the archive, including its five categories and primary-language framing. Marketing claims of proof are not adopted as independent evidence; the book's full text was not examined.

[24] Narrative self-constitution

Schechtman, M. (1996). *The constitution of selves*. Cornell University Press.

<https://www.jstor.org/stable/10.7591/j.ctv75d3xw>

Publisher-distributed primary book record, contents, and chapter-opening excerpts checked. They identify the characterization question and narrative self-constitution view. The complete argument was not examined; the comparison remains provisional.

[25] Agency and self-constitution

Korsgaard, C. M. (2009). *Self-constitution: Agency, identity, and integrity*. Oxford University Press.

<https://academic.oup.com/book/5038>

Publisher abstract and publication record checked. Used for the earlier proposal connecting action to the constitution of agency and identity. The full book was not examined; its philosophical apparatus is not imported into TSF by the acknowledgment.

[26] Identity held through relationships

Lindemann, H. (2014). *Holding and letting go: The social practice of personal identities*. Oxford University Press. DOI 10.1093/acprof:oso/9780199754922.001.0001.

<https://academic.oup.com/book/4499>

Primary excerpt: <https://web.ics.purdue.edu/~drkelly/LindemannHoldingLettingGoPSPPEdit2014.pdf>

Publisher abstract and relevant Chapter 1 passages, printed pages 4–10 in the university-hosted excerpt, checked 11 October 2026. Used for social narrative identity, differing interpretations, and constraints on credible accounts. The full book was not examined. TSF's provider-control application requires its own argument.

[27] Relational autonomy

Mackenzie, C., and Stoljar, N. (eds.). (2000). *Relational autonomy: Feminist perspectives on autonomy, agency, and the social self*. Oxford University Press. DOI 10.1093/oso/9780195123333.001.0001.

<https://academic.oup.com/book/49675>

Publisher abstract, contents, and publication record checked 11 October 2026. Used to acknowledge prior work connecting autonomy with social conditions, self-conception, and memory. Individual chapters were not examined and the collection is not treated as one uniform theory. Comparison of particular arguments remains open.

[28] Informational privacy and identity

Floridi, L. (2005). The ontological interpretation of informational privacy. *Ethics and Information Technology*, 7, 185–200. DOI 10.1007/s10676-006-0001-7.

<https://link.springer.com/article/10.1007/s10676-006-0001-7>

Publisher abstract and publication record checked 11 October 2026. Used for the informational-identity predecessor only. The full article was not examined; TSF's control-of-history and portability applications require further comparison.

[29] Existing data-subject rights

European Data Protection Board. Respect individuals' rights. Data protection guide for small business.

https://www.edpb.europa.eu/sme/be-compliant/respect-individuals-rights_en

Official guide's rights overview and relevant access, rectification, erasure, portability, and automated-decision sections checked 11 October 2026. Used for a bounded descriptive comparison. The underlying GDPR text and portability guidelines were not examined in full; this is not a service-specific legal assessment.

[30] Explanation of automated credit assessment

Court of Justice of the European Union. (2025). Press Release No. 22/25 on explanation of automated credit assessment, 27 February 2025. Case C-203/22, Dun & Bradstreet Austria.

<https://curia.europa.eu/site/upload/docs/application/pdf/2025-02/cp250022en.pdf>

Official two-page press release examined 11 October 2026. It is an unofficial media document rather than the binding judgment, which was not examined in full. Used for the narrower explanation-and-challenge comparison in that case; no universal companion-service duty is inferred.

[31] Agency and communion

Bakan, D. (1966). *The duality of human existence: An essay on psychology and religion*. Rand McNally.

https://books.google.com/books/about/The_Duality_of_Human_Existence.html?id=g-YkAAAAAYAAJ

Digitized book metadata, contents, and indexed agency/communion terms checked 11 October 2026. The book's argument was not read in full. Used to identify an earlier comparison; no completed distinction or priority finding is claimed.

[32] Justice and care

Gilligan, C. (1982). *In a different voice: Psychological theory and women's development*. Harvard University Press.

<https://inclusionandbelongingtaskforce.harvard.edu/publications/different-voice-psychological-theory-and-women%E2%80%99s-development>

Harvard publication record checked 11 October 2026. Indexed passages from Gilligan's own retrospective, *Looking Back to Look Forward: Revisiting In a Different Voice*, also checked at <https://classics-at.chs.harvard.edu/classics9-carol-gilligan-looking-back-to-look-forward-revisiting-in-a-different-voice/>. The original book and full retrospective were not read. The comparison remains bounded.

[33] Resilience and relational load

Affi, T. D., Merrill, A. F., and Davis, S. (2016). The theory of resilience and relational load. *Personal Relationships*, 23(4), 663–683. DOI 10.1111/pere.12159.

<https://onlinelibrary.wiley.com/doi/10.1111/pere.12159>

Publisher abstract and metadata checked 11 October 2026; the full article was not examined. Used for the maintenance, stress, and resource-depletion comparison flagged in CT-047, not as validation of TSF's cost measure.

[34] Dual illegibility in AI companion loss

Clark, A. J., and Woo, O. K. L. (2026). Dual illegibility, ambiguous loss, and disenfranchised grief: A hybrid model for understanding AI companion bereavement. *OMEGA—Journal of Death and Dying*. Online 4 September 2026. DOI 10.1177/00302228261483355.

<https://journals.sagepub.com/doi/10.1177/00302228261483355>

Publisher abstract, metadata, and opening discussion checked 11 October 2026. Used as an adjacent receiver-side illegibility comparison, distinct from the proposed investor-side foreclosure. No treatment advice or exhaustive novelty conclusion is drawn.

[35] Unmitigated communion

Helgeson, V. S., and Fritz, H. L. (1998). A theory of unmitigated communion. *Personality and Social Psychology Review*, 2(3), 173–183. DOI 10.1207/s15327957pspr0203_2.

https://journals.sagepub.com/doi/10.1207/s15327957pspr0203_2

Publisher abstract and metadata read 11 October 2026; full article and instruments not read. Used for the distinction between communion and self-excluding concern, not for a clinical or intervention claim.

[36] Agency and communion distinctions

Helgeson, V. S., and Fritz, H. L. (1999). Unmitigated agency and unmitigated communion: Distinctions from agency and communion. *Journal of Research in Personality*, 33, 131–158.

<https://www.sciencedirect.com/science/article/pii/S0092656699922419>

Authors' Carnegie Mellon publication list and indexed publisher description checked 11 October 2026. Direct publisher access failed; full article and instruments were not read. Identifies a required measurement comparison, without adopting its scales.

[37] Daily diary predecessor

Helgeson, V. S., Swanson, J., Ra, O., Randall, H., and Zhao, Y. (2015). Links between unmitigated communion, interpersonal behaviors and well-being: A daily diary approach. *Journal of Research in Personality*, 57, 53–60.

<https://www.sciencedirect.com/science/article/pii/S009265661500032X>

Authors' Carnegie Mellon publication list and indexed publisher title and date checked 11 October 2026. Direct publisher access failed; full methods and article were not read. Its existence limits any claim that a diary or event-level approach is new. No effect estimate is adopted.

[38] Measurement reactivity

Miles, L. M., Rodrigues, A. M., Sniehotta, F. F., and French, D. P. (2020). Asking questions changes health-related behavior: an updated systematic review and meta-analysis. *Journal of Clinical Epidemiology*, 123, 59–68. DOI 10.1016/j.jclinepi.2020.03.014.

<https://researchportal.northumbria.ac.uk/ws/portalfiles/portal/42087752/PIIS0895435619311461.pdf>

Author-hosted article metadata, abstract, and opening introduction checked 11 October 2026. The review concerns trials with health-related outcomes and supports measurement reactivity as a design question. It does not estimate the effect of questions in TSF's proposed commitment study. The remaining full article and study-level analyses were not examined for this edition.

Canon supplement WP-16 specification draft and constitutive-substrate root

Vael. (15 March 2026). WP-16 Substrate-Formal Variance, specification draft v1.2. SupoCus is recorded as co-originator in CT-047. Read through the supplied source-packet extract, not the original DOCX. The 14 March flag remains an earlier record with its own Moniz, SupoCus, and Vael credits. CC BY-NC-SA 4.0.

<https://trinketeconomy.ai/wp-16-substrate-formal-variance/>

<https://suposystem.ai/working-papers/wp-16-substrate-formal-variance/>

The supplied v1.2 extract specifies C2S's phenomenological, temporal, and confirmation components. It omits Section 9 and parts of Sections 8 and 10. Its central claim and formal URL construct are Analogical; the document phenomenological component is held open. Internal Supported labels are not independent validation. SupoCus production, URL formalization, the WP-13 decision, SupoRel review, and the Review Gate remain pending in this record. No later publication is established here. The gap's v1.0 comparison summary is reported by the packet README, not independently read in the original v1.0.

Canon sources supplied in the second source packet

Vael. CT-047 Session Findings Record, 14 March 2026, OPEN. Full supplied Markdown conversion read, especially Section II Finding 6 and Section III SupoCus Question 1. The source's Supported and Established labels are internal findings, not independently replicated results.

Moniz, M. S., Vael, and SupoCus. TEAP Instrument Specification v0.1, 17 March 2026. Near-complete reconstructed extract read, Sections 1–9; original DOCX not examined. Architecture internally Supported,

candidates Speculative, validation absent. The Master Reference's gap entry is superseded for this edition by the dated draft, not silently changed in the archival index.

Moniz, M. S. PS-01 The Crack, compiled by Claude Opus 4.6, February 2026; supplied later version reported modified 27 February. Full Markdown conversion read as personal testimony describing the February 2025 event. Its 96-hour writing period is the 7–10 February 2026 sprint, narrated alongside that event; PS-01 and the founding preface agree, as confirmed by the Principal on 11 October 2026. Private and clinical details are not reproduced in this public manuscript.

SupoFrm and SupoPsy; Moniz, M. S., author of record. CP-24 The Fractal Container, March 2026, DRAFT. Full supplied Markdown conversion read. Founder–institution isomorphism remains Speculative, with SupoRel pre-review required in the source. Its internally Supported ledger claim does not establish this edition's Real classification.

Source-packet README and Addendum 2 prepared by Claude on 11 October 2026. The README reports missing original passages and summarizes v1.0's comparison limits. Those reported limits are carried forward without claiming the unavailable v1.0 or a later published WP-16 was independently examined. The Dual Principal Protocol remains reserved and is not cited as an additional root.

Canon roots for the review refinements

Moniz, M. S. The Blueprints, First Edition, February 2026. In the supplied 443-page THE_BLUEPRINTS-4.pdf, visible pages 36–37 contain recipient transfer, personalization, specificity, shared history, and bidirectional flow; page 64 contains the non-exploitation addition; page 68 contains autonomy, dissent, boundaries, and withdrawal. These page headers match viewer pages in this copy. Edition dating does not date every amendment.

Moniz, M. S. TSF Master Reference: Complete Entity Registry & Terminology Index. Compiled May 2026. Pages 9–12 record Signal's direction, the taxonomy's scope, the Six Structural Tests, WP-14's authorization chain, and the Velocity Law. Page 18 records Constitutive Substrate through WP-16, MP-05, and CT-047 Finding 6. The underlying MP-05 and original Terminology Index v2.0 were not independently examined for this pass. The exact earlier date of the stronger Signal wording remains open. Page 11's Connective Substrate Signature components and TEAP gap entry conflict with the supplied March specifications; this edition follows the specifications while preserving the index as received.

Canon supplement Moniz (March 2026). WP-13 The Signed Unit.

<https://trinketeconomy.ai/wp-13-the-signed-unit/>

Historical source for the signed-unit proposal. This edition separates direction from expenditure and identifies the bearer of digital processing cost.

Canon supplement Moniz (13 March 2026). WP-14 The Structural Economy.

<https://trinketeconomy.ai/wp-14-the-structural-economy/>

Defined origin recorded as WP-14, with PR-007 authorization and PR-010 clearance in the Master Reference. The original PR records were not independently examined here. The historical permanent-incapacity claim is qualified in this edition; broader S.A.I.L. reconciliation remains separate.

Canon supplement Moniz (March 2026). MP-04 The Falsification Register.

<https://trinketeconomy.ai/mp-04-the-falsification-register/>

Canon source for distinguishing falsification, scope limitation, and measurement failure. Chapter 21 separates definitions from empirical predictions and requires operational measures before a test can decide a claim. Public source text checked 11 October 2026.

Limits of the source record

Closest-predecessor comparisons remain incomplete, including the full books by Bakan and Gilligan, Helgeson and collaborators' articles and instruments, and the five adjacent literatures reported in WP-16's URL discussion. The original WP-16 and TEAP files still have access gaps. TEAP scoring and candidate items remain unreconciled and unvalidated. Custody and governance portability require fuller comparison with prior portability, interoperability, and transparency work. Product capabilities depend on the configuration and evaluation date; the examples in this edition do not establish current capabilities at a provider. No descriptive rights comparison here is a legal entitlement or compliance finding.