GOT-V1-3b-0

0. Part 3b: Prelude

Power, Responsibility, and Trust

Reader introduction • Non-canonical. The Part 3b papers govern technical definitions, arguments and evidence.

1. If This Were Possible

What if we could benefit from systems beyond our individual understanding while keeping a recoverable account of how they shape our lives?

We depend on knowledge, decisions and work performed elsewhere. That dependence enables cooperation, specialisation and capabilities no individual could reproduce. We use instruments we did not build, records we did not create and judgements whose supporting work we cannot continually repeat.

Part 3a asked how we can proceed without complete understanding. Part 3b takes that predicament into a shared world. The people who shape an outcome, those who rely on the process and those who bear its consequences may be different people. A decision that saves work in one place can create exposure in another.

We propose to examine this dependence through three grounded relations: Power, Responsibility and Trust. If their conditions can be located and tested, we could preserve the benefits of reliance while making its consequences more open to examination and correction.

2. Why This Is Difficult

A decision may reach us as a finished answer. Behind it lie observations, classifications, retained records, rules and actions performed at different times by different systems. The person who delivers the result may have contributed only one transformation in that history. The person affected may see almost none of it.

Our language can close the account too early. The operator was responsible. The institution has authority. The system is trusted. Such statements can allocate work and make coordination possible. Yet they can also let a familiar name stand in for an unexamined pathway. A role becomes mistaken for a causal contribution; an established reputation carries reliance into a use it has never supported.

Trust compounds the difficulty because it saves us work precisely by allowing some checking to remain in the background. We need that economy. We also need to recover its grounds when conditions change or the consequences become more demanding.

3. What We Propose

General Organisation Theory and The Epistemic Method propose to give these load-bearing words distinct, inspectable functions. Each directs attention to a different question about the organised relations through which information becomes consequential.

Power concerns the physically supported potential to make what happens next differ through organised discrimination. A system may possess that capacity before using it. Its effect can be small or probabilistic; the question is whether the distinction can change the downstream relation under meaningful comparison.

Responsibility concerns which consequential relations can be attributed to that organisation under the conditions that existed. It follows the operations that made distinctions matter. A contribution may be an applied rule or a maintained condition that changes another system's available options, where the relevant dependence can be demonstrated.

Trust concerns treating a relation as dependable enough for a particular use while uncertainty remains. It reduces the need to repeat earlier checking. Audit re-examines whether the relation still supports that reliance and must be capable of changing what follows.

These are empirical, testable and substrate-agnostic proposals. Their scope comes from the relation demonstrated in each case. A switching circuit, a cell and an institution require different investigations; a shared language makes those investigations comparable without making their capacities equivalent.

4. What This Part Inherits

Part 1 supplies Non-Trivial Computation: the system's organisation makes what happens next depend on a distinction. Part 2a grounds the physical pathways through which signals abstract causal history into available relation. Part 2b distinguishes registration, retention and application. Part 3a disciplines the conclusions and actions that finite access can support.

Part 3b inherits this architecture and its comparison requirements. A retained classification can shape a later decision long after its original circumstances have disappeared from view. Tracing its contribution means recovering how it remained operative, what later systems did with it and which limits or correction conditions survived.

5. Why Three Papers

On the Origin of Power establishes capacity before its exercise. It restricts the relevant alternatives to those physically supported by the system's organisation. A small discriminating event can redirect much larger flows of available energy. The paper separates this potential from its application, its resulting consequence and stronger claims about control or success.

On the Origin of Responsibility follows attributable contributions through the consequential pathway. Responsibility Provenance records where organised distinctions did that work. It can extend across systems and time without becoming one quantity divided among everyone involved. A later transformation does not erase an earlier contribution, and an earlier contribution does not absorb every subsequent use.

On the Origin of Trust explains how finite systems proceed without reopening every dependency. It develops reliance on work performed elsewhere, the people and records that carry it, and the conditions for renewed scrutiny. It also asks how reliance becomes overextended and what must change for damaged trust to become justified again.

6. What You Need to Grant – and What You Do Not

The invitation is to examine these relations before asking them to carry richer judgements. A circuit need not deserve blame to have an attributable switching function. Human duties can legitimately extend beyond direct causal contribution; assigning a duty does not establish the causal history. These primitives leave further political, moral and legal questions to their own arguments.

Nor does the framework demand suspicion of every expert. Demonstrated competence can justify preferential reliance for a particular question, provided its strength survives the path by which it reaches us. Permission to decide remains a separate matter. Attribution must also respect what was available at the time, including warnings and correction routes that actually existed.

7. Where to Attack It

Consider a proposed temperature sensor and heating controller. Vary temperature around the switching threshold, then bypass the sensing pathway while preserving the power supply, switching hardware and other relevant conditions. Test whether the output's dependence on temperature changes. Simply breaking the heater would not isolate the claimed discrimination.

Trace the sensor reading, retained calibration and switching rule separately. Which transformation changes the output under comparison? The final relay does not automatically contain the whole contribution, and the calibration's existence does not prove that it mattered. Attribution must follow the dependence actually recovered.

Now examine reliance on that calibration between checks, within a stated operating range and tolerance. Introduce a detectable offset and compare the reading with an independent reference. Can the discrepancy reopen checking and restrict output or trigger recalibration? A logged warning that changes nothing does not demonstrate the promised correction function.

This is an illustrative comparison, not a reported experiment. Attack the framework where its distinctions fail to constrain the account: potential that cannot change the outcome, attribution unsupported by comparison, or supposedly calibrated reliance that persists after its own stopping condition is met. Each claim must remain vulnerable to the evidence relevant to it.

8. If the Architecture Holds

If these relations survive scrutiny, they could make our dependence on complex systems more accountable. We could distinguish who shaped an outcome from who was assigned a duty, and preserve justified reliance while exposing the conditions under which it should narrow or stop.

For intelligent machines, grounding language must extend into the pathways through which outputs direct action. A classification can close an option; a retained score can shape a later decision. The affected person may never have chosen to trust the system. An auditable account must therefore preserve the grounds of reliance and the exposure of those who bear its consequences.

The three papers establish the proposed relations. Part 3c then examines how their conditions can remain recoverable as records, categories and outputs travel and become authoritative. The human question remains: when our systems act on information we have chosen to trust, who bears what remains unresolved, and how can they challenge it?

Canonical source note

Canonical basis: On the Origin of Power (GOT-V1-3b-3), On the Origin of Responsibility (GOT-V1-3b-4), and On the Origin of Trust (GOT-V1-3b-5), all Issue 1. Their definitions, qualifications and evidence govern this introduction. The controller example is illustrative.

Prepared under GOT-SD-PRELUDE-01, Issue 2, with the Part 3 Issue 1 concept map, glossary, guardrail register and bibliography/relation register as secondary checks. The book supplies wider orientation; the issued preludes supply the presentation precedent.