ESSAY XXXI

Institutional Power and Moral Responsibility

Agency, Complicity, and the Problem of Many Hands

VOCABOR SILENTII

ABSTRACT

Institutional action distributes decisions across people, procedures, records, and technical systems. This distribution makes possible achievements beyond an individual's capacity, but it also complicates the attribution of harm. No participant may control an entire outcome, and membership alone does not establish culpability. This essay distinguishes causal contribution, role responsibility, answerability, blameworthiness, and duties of repair. Engaging the problem of many hands, accounts of structural responsibility, and the ethics of computerized systems, it argues that institutional complexity neither dissolves responsibility nor permits its indiscriminate assignment. Responsibility must be indexed to the question being asked: who produced an outcome, who could reasonably have known, who possessed authority to intervene, and who can now discharge an outstanding obligation? A hypothetical automated eligibility process tests these distinctions against distributed knowledge, deliberately fragmented decision-making, justified reliance, and organizational succession. The essay develops a layered account in which individual attribution and institutional responsibility can coexist without being identical. Its central claim is that the absence of a single comprehensive author requires an analysis of relations among partial authors, not a choice between universal guilt and an outcome for which no one must answer.

Keywords: responsibility; institutions; many hands; complicity; collective action; automation; repair

Conceptual analysis and critical reconstruction. Cases introduced as illustrations are hypothetical, not empirical findings.

1. The decision that no one says they made

Consider an illustrative organization that allocates access to a training program. A committee defines eligibility. A contractor translates the criteria into software. An administrator selects the records used by the system. A reviewer sees only cases marked for exception. An applicant is rejected because an outdated record is treated as current, but the rejection notice does not identify that reason. Each participant can describe a limited task accurately. The committee did not choose the inaccurate record. The contractor did not determine eligibility. The administrator did not issue the decision. The reviewer never saw the case.

The applicant nevertheless encountered one organized outcome: exclusion by the organization. The absence of a participant who performed every step does not make that outcome disappear. Nor does it show that everyone involved is equally blameworthy. Some may have acted competently within justified expectations; others may have known about a recurring defect; still others may have possessed authority to provide review but chosen not to exercise it. An account of responsibility must preserve those differences while explaining who owes the applicant an answer.

Thompson's formulation of the problem of many hands identifies the difficulty of assigning moral responsibility when organizational outcomes arise through numerous contributions (Thompson 1980). The present essay takes that problem beyond a search for one final decision-maker. Its question is how different responsibilities can attach to different positions in a process, including duties that remain after the grounds for personal blame have been exhausted.

Part One's conditioned-authorship model connects attribution to understanding, control, access to reasons, and opportunities for revision (Silentii 2026a, Essay VII). Institutional action complicates each condition because the relevant capacities are distributed. The task is not to abandon the conditions, but to ask what they mean when information, authority, and consequence are located in different parts of an organization.

2. Five questions concealed by responsibility

Causal responsibility asks how an outcome came about. Role responsibility concerns a task or domain entrusted to someone. Answerability concerns who owes an explanation. Blameworthiness concerns whether an agent acted wrongly under conditions that make moral criticism appropriate. Reparative responsibility concerns who must address an injury or discharge an outstanding obligation. These questions interact, but their answers need not identify the same person.

A technician may causally contribute to an outcome without knowing the relevant purpose of the system. A manager may hold responsibility for providing oversight without directly producing the defective record. An organization may owe an explanation because it exercised authority over an applicant even where the initial error was not negligent. A current administrator may have a duty to correct the record despite having joined after the error occurred. A person who deliberately concealed the problem may be blameworthy in a way none of the others is.

The distinctions prevent two shortcuts. The first moves directly from causation to guilt. Many necessary background contributions are not wrongful. Supplying electricity, maintaining a building, or processing an ordinary request can contribute to an outcome without providing a sufficient basis for blame. The second moves from the absence of culpability to the absence of any duty. An error can require correction even when everyone initially exercised reasonable care. The person harmed need not prove an individual vice before the institution can recognize that its decision was mistaken.

Young's social connection model distinguishes shared, forward-looking responsibility for structural processes from an approach centered on isolating liable perpetrators (Young 2006). That distinction is relevant here without making every organizational duty a claim about structural injustice. It shows why the question of what participants must now do can remain open after a narrower inquiry into who deserves blame has reached its limits. Forward-looking responsibility supplements rather than automatically replaces investigation of particular wrongdoing.

3. Partial authorship and the structure of contribution

Institutional outcomes can have several kinds of contributors. Some choose ends, some define rules, some implement a method, some authorize deployment, and some respond to failure. These roles can be occupied by one person or many. Their significance depends on how they affect the result. A formal title alone does not establish actual control, and physical proximity to the final act does not establish the greatest responsibility.

In the training example, the person who sends the rejection notice may have less relevant discretion than the committee that prohibited explanations. Conversely, an implementer may independently introduce an inappropriate criterion and conceal it from those authorizing the system. Responsibility therefore cannot be assigned merely by tracing the last visible step or assuming that seniority always determines culpability. The inquiry must connect a person's decisions to the feature of the outcome under examination.

Causal contribution is itself more complicated than a simple test of whether the harm would have occurred without one person. Two independently sufficient errors may produce the same exclusion. Removing either one alone would leave the other. It would not follow that neither mattered. Alternatively, many small contributions may jointly create a condition that none could create alone. These possibilities show why attribution must describe the actual causal structure rather than rely exclusively on the imagined removal of one participant.

The normative inquiry then adds questions that causation alone cannot answer. What did the person know or have reason to investigate? Which alternatives were available? What undertaking did they accept? What costs would intervention have imposed on them or others? These questions do not convert responsibility into a precise numerical share. They make reasons for differentiation explicit. A layered account is more demanding than assigning one culprit, but also more discriminating than distributing blame across everyone who touched the process.

4. Distributed knowledge and duties to connect information

An organization may possess information in the sense that different participants hold different pieces, although no individual has assembled the whole. This does not mean that every participant personally knew everything in the organization's records. Attributing collective information to each person would manufacture knowledge they did not have. Yet an institution cannot always defend itself by pointing to the fragmentation through which its own processes operate.

Suppose one office knows that records become inaccurate after a defined interval, another knows that the software never checks their date, and a third receives complaints about unexplained rejections. None initially recognizes the entire pattern. The question is whether responsibilities existed for connecting these observations. Was there an appropriate reporting channel? Did someone have a defined task of reviewing repeated failures? Were communications intelligible across the relevant roles? A defect in the organization of inquiry can exist even when no single person began with complete knowledge.

The standard must remain reasonable rather than omniscient. Every organization contains more information than anyone can inspect continuously. Duties of inquiry must be proportionate to foreseeable risk, the cost of investigation, and the evidence already available. A remote and unforeseeable interaction differs from a repeated anomaly that relevant personnel were instructed to ignore. The absence of perfect foresight cannot excuse deliberate blindness, but neither should the fact that a harm occurred retroactively make every possible precaution obligatory.

There is also a distinction between failing to know and arranging not to know. An institution might design reporting so that decision-makers receive only favorable summaries or so that no role is permitted to compare records. In such a hypothetical case, fragmented knowledge would be part of the conduct requiring explanation. Responsibility would attach not because every participant secretly possessed the missing information, but because identifiable choices obstructed the means by which consequential information could become available.

5. Automation and the location of judgment

Technical mediation does not remove the need to identify decisions. Someone determines what problem a system will address, which inputs it may use, how its outputs enter a process, and what happens when an output is challenged. These are distinct acts even when the system's internal operations are too complex for one person to reconstruct in every case. The relevant question is not whether a human manually produced each result. It is how authority was assigned to the process producing it.

Nissenbaum identifies obstacles to accountability in computerized settings, including many hands, the treatment of errors as inevitable, and the displacement of blame onto computers (Nissenbaum 1996). Her analysis provides a point of departure for distinguishing technical explanation from institutional justification. Explaining that a program generated an output can identify a mechanism; it does not establish that the organization was entitled to impose the corresponding consequence without explanation or review.

In the training example, the system may correctly execute an instruction to use the most recently imported record. The decision remains defective if that instruction is inappropriate for the relevant purpose. Conversely, a sound criterion can be implemented incorrectly. A third possibility is that both criterion and implementation are reasonable but an unusual input produces an error that should be corrected. These cases involve different failures and potentially different responsibilities. Calling all of them algorithmic bias, human error, or system failure would obscure the distinctions.

Human review is not an automatic solution. A reviewer without relevant information, time, authority, or a workable procedure may be unable to change a decision. Adding a human name to an unchanged process can therefore leave the responsibility problem intact. The analysis must ask what the reviewer can actually examine and alter. Responsibility cannot be made meaningful simply by designating a person who lacks the capacities required to discharge it.

6. Can the organization itself owe an answer?

One objection holds that only individuals can be responsible because only individuals act, understand, and experience moral attitudes. Another position allows organizations to count as agents under appropriate structures of decision and control. The present account does not require resolving every metaphysical issue between them. A narrower question is whether an organization can bear an obligation whose discharge requires coordinated action by changing individuals.

Ordinary institutional attribution already makes such obligations intelligible. A program undertakes to evaluate applications, a cooperative promises access, or a publisher agrees to preserve records. Replacing a staff member does not by itself discharge the undertaking. If the obligation belonged only to the person who happened to answer the first message, institutional continuity would not explain why later officeholders must act. The organization provides a structure through which the obligation persists and through which particular people acquire tasks for fulfilling it.

This does not make the organization a substitute object that absorbs all individual responsibility. An institution can owe correction while some individuals also owe explanations or face criticism for their conduct. Conversely, a person can violate institutional requirements without the violation being attributable to every member. The relevant question is how the act relates to authorized functions, resources, instructions, and subsequent responses. Individual and organizational attribution are compatible when they answer different questions.

Reparative duties can consequently survive changes in personnel without transferring guilt to successors. A new administrator can be required to provide an accessible hearing because that is now part of the office, not because they caused the initial exclusion. The applicant’s claim does not disappear with the departure of a culpable employee. At the same time, a duty to repair does not justify treating every current participant as though they personally committed the original wrong.

7. Complicity, refusal, and the limits of sacrifice

Participation can become morally significant when a person knowingly contributes to a wrongful undertaking. Yet complicity should not be inferred from association alone. A person may occupy a peripheral role, lack relevant knowledge, or attempt to change the process from within. The content of the contribution and the available alternatives matter. Institutional membership is evidence about a relationship, not a complete verdict about conduct.

Consider a reviewer who discovers that the system excludes applicants on the basis of obsolete information. The reviewer's responsibilities depend partly on whether they can correct cases, report the defect, pause decisions, or preserve evidence. The existence of one possible action does not establish that every imaginable act of resistance is required. A demand to risk immediate loss of livelihood or expose unrelated confidential information requires its own justification. Moral seriousness should not be measured simply by the maximum sacrifice an observer can imagine.

Nevertheless, constraints do not automatically eliminate responsibility. An employee may be unable to repair an entire system but able to avoid falsifying a record or to communicate a specific concern through an available channel. The appropriate inquiry concerns what a person could reasonably do, given their situation. Acknowledging limited power can support proportionate duties without either romanticizing refusal or treating institutional instruction as a complete excuse.

There is a further institutional implication. An arrangement that depends on exceptional personal sacrifice to reveal ordinary defects has assigned correction precariously. This is an analytical observation about the design of responsibility, not a finding that every organization has such a structure. Where reporting a problem predictably defeats the reporter's ability to continue reporting, the institution cannot treat the mere existence of a reporting address as evidence that its errors are meaningfully discoverable.

8. Repair, blame, and preventing substitution

Repair can include correcting a decision, restoring access, explaining what happened, preserving an accurate record, or addressing a loss that cannot be reversed. These responses are not interchangeable. An apology does not restore a missed opportunity. A changed procedure does not by itself acknowledge the person previously harmed. Compensation may address a loss without resolving whether a classification remains in a record. The required response depends on the specific obligation and injury.

Blame can also be misused as a substitute for repair. Removing one employee may create the appearance of accountability while leaving the design that produced the error unchanged. Conversely, declaring a problem systemic can erase the conduct of individuals who deliberately concealed evidence or misused authority. A serious investigation must remain capable of finding both structural defects and personal wrongdoing, neither of which excludes the other by definition.

The demand for proof matters at each level. Evidence sufficient to reopen a decision may not establish misconduct by a named person. Evidence of a recurring problem may justify an institutional inquiry without proving its precise cause. A finding of causation may still leave culpability uncertain. Keeping these thresholds distinct permits action under uncertainty without turning suspicion into a verdict or requiring definitive blame before any corrective step becomes possible.

The strongest objection to layered responsibility is that it multiplies duties until responsibility becomes everyone's and therefore no one's. The answer is specification. A shared obligation requires identifiable assignments, channels for coordination, and a means of determining whether the work occurred. The fact that several persons must contribute does not make their tasks indistinguishable. Where no one can identify who must take the next step, the account has not yet translated a general moral concern into an institutional responsibility.

9. Conclusion: an outcome with several authors and distinct obligations

The problem of many hands is not solved by locating the nearest hand to the final act. Nor is it solved by calling everyone equally responsible or treating complexity as an exemption. Institutional action requires an analysis of partial contributions, distributed knowledge, accepted roles, available interventions, and continuing obligations.

The essay has distinguished explanation, answerability, blame, and repair. Those distinctions permit an organization to correct an injury without inventing personal guilt, permit individuals to be criticized without making membership itself culpable, and permit obligations to survive the replacement of the persons who originally performed a role. They also expose how automation can mediate judgment without removing the decisions through which a technical process acquires institutional authority.

Responsibility is therefore not one undifferentiated burden moving from person to institution or back again. It is a set of relations that must be matched to a question and a task. Who contributed? Who could know? Who could intervene? Who must now explain or repair? Where the answers differ, precision is not a retreat from accountability. It is the means through which accountability becomes possible without sacrificing either the injured person's claim or the distinction between involvement and guilt.

References

Nissenbaum, Helen. 1996. “Accountability in a Computerized Society.” *Science and Engineering Ethics* 2 (1): 25–42. DOI: 10.1007/BF02639315

Silentii, Vocabor. 2026a. *The Architecture of Human Order. Part One: Foundations of Thought, Meaning, and Human Agency*. Revised and expanded edition. Author-supplied manuscript.

Thompson, Dennis F. 1980. “Moral Responsibility of Public Officials: The Problem of Many Hands.” *American Political Science Review* 74 (4): 905–916. Source

Young, Iris Marion. 2006. “Responsibility and Global Justice: A Social Connection Model.” *Social Philosophy and Policy* 23 (1): 102–130. DOI: 10.1017/S0265052506060043