Making Rules Matter:

Toward a Theory of Structure-Oriented Collaboration

Abstract

Purpose: This paper introduces structure-oriented collaboration (SOC) as an emerging practice and research program for explaining how explicit, ongoing, prescriptive expectations function as participant-facing navigational aids in real collaborative systems. It targets a gap in rule theory: many existing accounts explain compliance, incentives, or norms, but under-specify how rules become usable as obligations from the participant's standpoint, especially in peer-to-peer and self-managing contexts.

Approach: The framework is developed through longitudinal, practice-based first-person inquiry within Holacracy as a revelatory context, using abductive concept formation, iterative writing, and repeated disconfirmation attempts through practitioner and client application. The method is explicitly positioned as hypothesis-generating rather than outcome-confirming.

Results/claims: SOC advances several testable claims. Collaboration exhibits a predictable drift toward explicit structure as coordination must span time, distance, and diversity. Legitimacy is modeled as two interacting types—structural and experiential—whose alignment predicts how obligations are encountered. Rule functioning is analyzed via a four-part "structure system" (exploring, encoding, enforcing, evolving), plus four participant standpoints (agent, author, enforcer, analyst). A minimal normative constraint—non-alienation—evaluates whether structures preserve participatory agency.

Implications: A structure-first diagnostic sequence is predicted to be both higher-leverage and more humane than people-first explanations. Compensatory patterns (strategic non-enforcement, official "shoulds," norm-reliance) are treated as diagnostic signals of under-differentiated structure, with predictable byproducts including shadow authority, politics, and onboarding cost.

Keywords: structure-oriented collaboration; rules; explicit expectations; legitimacy; accountability; enforcement; self-management; participant navigability

Making Rules Matter: Toward a Theory of Structure Oriented Collaboration

Rules are everywhere. They determine when work officially begins and ends, which side of the road cars drive on, and what counts as a legal contract. And these are only the explicit ones. Implicit rules shape when to speak, what "late" means, and how strictly specific contractual clauses are enforced. Taken together, explicit and implicit expectations enable collective action across space, time, and individual differences (North, 1990; Ostrom, 1990). Without rules—formal or informal—modern life would collapse. Rules do not merely reflect social order; in their enactment, they actively reconstitute it (Bourdieu, 1977; Giddens, 1984).

However, current work in political theory, deontic logic, game theory, sociology, and organizational studies do not answer some important questions about how rules work (when they do) as navigational aids for participants. This paper argues that this lens is understandably absent because other theories did not have to account for experiences revealed through new social practices in self-management and peer-to-peer collaboration, which place particular emphasis on explicit rules and agreements as coordinating[1] mechanisms. Critically, extensive experience with self-management practices has revealed deeper questions and answers about how rules work—when they do—that are relevant for any governance system and any collaboration of two or more people working toward a shared purpose.

About This Paper

This paper presents structure-oriented collaboration (SOC) as an emerging practice and research program concerned with how rules work as tools in human collaboration. Its animating claim is that by orienting to ongoing, explicit, and prescriptive expectations participants can both better understand and better shape their own collaboration systems. Throughout this paper, the terms "rules," "structures," and "explicit expectations" refer to the same underlying phenomenon:

[1] The term "collaboration" is used broadly to refer to situations in which two or more people work together over time towards a shared purpose (Bedwell et al., 2012). Occasionally, the term "coordination" is used but no significant distinction between collaboration and coordination is conveyed.

ongoing, explicit, prescriptive expectations. This definition is elaborated in the section on defining structure.

The primary impetus for the approach's emergence is the author's extensive experience with the self-management practice Holacracy, which served as a revelatory context foregrounding advantages and challenges in orienting collaborative efforts around explicitly defined expectations. The hypothesis is that this experience exposed psychological and social phenomena unknown to other theories and were therefore unaccounted for by them, but which are broadly applicable to all collaborative efforts.

This approach's central claim is that all collaborative endeavors are caught in a developmental drift towards the creation and maintenance of ongoing, explicit, and prescriptive expectations because of the inherent need to integrate individual differences (i.e. diversity) and the need to coordinate across time and space. Further, by better understanding the ever-present potential of the drift toward explicit structure, the web of norms that support structure, and how different types of legitimacy interact, people are better able to navigate and influence collaborative situations and systems they are in.

The Status of the Approach

What SOC means is not yet settled. Even now, the approach is difficult to pin down or summarize in a comprehensive and fully precise way. Different aspects of the approach are currently being developed in parallel, and different features have been granted greater or lesser emphasis depending on the context of the presentation. Key elements of the approach remain intentionally open, with the expectation that their content will be refined through continued empirical observation, conceptual clarification, and practical application.

For this reason, SOC should not be understood as a fully articulated model of rules or collaboration. Rather, it is best characterized as an emerging practice and research program concerned with how explicit, ongoing, prescriptive expectations function within real collaborative

systems. Its central aim is not to offer a unified theory of rules or human coordination, but to surface a minimal set of distinctions that make collaborative environments more navigable for participants in the system. These distinctions are evolving, and their boundaries are expected to shift as the framework encounters new domains and counterexamples.

It would therefore be premature to specify exhaustively what SOC is and is not committed to, as an umbrella orientation under which multiple models and applications may eventually fall. With this qualification in mind, the present work advances a plausible initial formulation of the framework, while acknowledging that its claims may be revised, further nuanced, or recast in future iterations. To be considered plausible, however, any such formulation must remain grounded in lived collaborative experience and consistent with the most general commitments motivating the approach—namely, that coordination breakdowns are often misdiagnosed as psychological or cultural failures, and that many such breakdowns become intelligible only when examined at the level of explicit structure.

Early development of SOC has been guided less by abstract theorizing than by sustained engagement with operational environments, where roles, policies, agreements, and accountability mechanisms are continuously enacted, contested, and revised. From this perspective, collaboration is treated not primarily as an interpersonal phenomenon, but as a system of navigable expectations embedded across institutional, interactional, and individual registers.

In this sense, SOC does not attempt to replace existing psychological, cultural, or organizational accounts of collaboration, but instead to complement them by making visible a layer of coordination that is typically implicit: the internal mechanics by which rules acquire meaning. The framework remains deliberately minimal at this stage, prioritizing descriptive coherence and early empirical testability over theoretical completeness.

An Example of Structure-Orientation

For the purposes of an introductory example, imagine a manager who is having difficulty getting a few team members to speak up during meetings and is unsure how some of their ideas are landing. Three broad types of intervention are available.

A psychology-oriented solution would target individual capacities and choices. The manager might periodically turn to the quieter people and ask, “Any thoughts on this?” or “Anything you’d add?” If they decline, the manager moves on without pressing, and follows up in one-on-one meetings, encouraging and reassuring them that disagreement is welcome. This approach can work when the main constraint is personal comfort or confidence, though it tends to make changes to an individual’s capacities and beliefs a requirement to achieve collective outcomes.

A culture-oriented solution would target shared norms. The manager might remind the team that “we value candor” or say the group goal is a “speak-up culture.” The manager praises people who offer pushback and reminds quieter members in private that being vocal is part of the team’s culture. This approach can gradually legitimize the desired behavior, though it tends to lock in one side of a polarity without telling participants when, for example, agreement or discretion should be prioritized over speaking-up.

A structure-oriented solution shifts the burden from courage and mind-reading to a repeatable interaction pattern. The manager might add a “Reaction Round” process to the meeting, defined as: “Going one at a time, anything you have to add is welcome, and no crosstalk, responses, or interruptions until that person is finished.” The manager facilitates the process rigidly but neutrally. In doing so, they create a protected and predictable space for contribution, clarifies the rules of engagement, and makes participation easier by focusing the alignment effort on procedural requests, not changes to participant’s beliefs or values. Structure, as used here, does not replace psychology or culture; it operationalizes the intent so the team can reliably produce the behavior even when energy or trust fluctuates. The essence of SOC is

taking this same structure-oriented approach and making it applicable within any collaborative situation or system.

Origins of the Approach

This approach did not emerge through abstract theorizing but through sustained engagement with practice—specifically, through long-term work within the self-management practice Holacracy, an organizational operating system that intends to make a group's shared structure and the governance mechanics to create it unusually explicit. For a time, it became popular with high profile companies like Medium and Zappos, and by some estimates is being practiced by more than 1,000 companies worldwide (Greenfield, 2015; Weirauch et al., 2023). However, Holacracy is not treated here as a normative ideal or a preferred governance model, but merely as a social practice acting as a revelatory context: one in which dynamics that are typically absorbed tacitly in other settings are rendered visible through explicit formalization (Bernstein et al., 2016; Lee, 2019). Rather than promoting Holacracy, SOC can be accurately understood as a post- or trans-Holacracy approach that provides a direct and substantive critique of Holacracy, while also suggesting concrete remedies. However, this stance is only an incidental consequence rather than an intentional aim.

Holacracy originated as an evolutionary branch of sociocracy and agile software development methods (Robertson, 2014). It provides an organization-wide and explicit framework specifying how decisions are made, how roles evolve, and how authority is exercised, all defined in its "constitution." This framework is presented as enabling any kind of organization of any size in any industry a roadmap for adopting continuously evolving, rule-bound self-management rather than through managerial discretion alone (Robertson, 2015). By operationalizing explicit expectations as roles, accountabilities, priorities, duties, and policies, Holacracy aims to remove many of the informal buffers that ordinarily compensate for weak or ambiguous structure. Expectations are written down. Authority is specified. Pathways for changing structure are articulated. As a result, participants cannot easily rely on personal

charisma, relational smoothing, or implicit hierarchy to resolve coordination problems. When the rules succeed, they do so on their own terms, and when they fail, the failure is difficult to work around through typical leadership interventions.

Notably, Holacracy is not unique in these respects. Many approaches have demonstrated how to operationalize an orientation toward explicit structure. Approaches like sociocracy and agile frameworks that make explicit working agreements and various forms of participatory governance all share something important: they make expectations referable, contestable, and revisable through explicit processes rather than relying on culture, charisma, or positional authority alone (Lee, 2019). These approaches collectively fit what Constantine labeled as the “open” organizational system paradigm (Constantine, 1993, 1995).

Research on Holacracy and related self-managing systems suggests several possible advantages when the practice is implemented effectively. Empirical and practitioner-based studies report improvements in role clarity, decision transparency, and organizational adaptability, particularly in environments requiring rapid learning and responsiveness (Bernstein et al., 2016; Lee & Edmondson, 2017). These outcomes align with broader research on peer-to-peer systems and self-managing teams linking formalized autonomy to increased engagement and coordination under certain conditions. Holacracy’s visibility increased further through its inclusion in Laloux’s (2014) analysis of “teal” organizations, where it was presented as an example of a social practice capable of supporting self-management, evolutionary purpose, and wholeness at scale.

However, limitations of Holacracy have also been documented. Common critiques include the high cognitive load imposed by the system, particularly during early adoption; the formalization of language and interaction to a degree experienced by some participants as alienating; and uneven outcomes that depend more on cultural readiness and coaching quality than on the formal rule set itself (Bernstein et al., 2016; Lee & Edmondson, 2017; Moeslein &

Laulitz, 2019). It was through directly experiencing advantages and limitations of Holacracy practice that the models presented here as SOC were initially considered.

Critically, the ambition of SOC extends well beyond the specific practice of Holacracy. As an emerging research program, it aims to make Holacracy practice cognizable and its successes and failures diagnosable within a single explanatory frame, but only as it attempts to do so with any collaborative system. The design constraint of the distinctions and models offered here is that they should apply to all collaborations within its scope including self-management, but also conventional management hierarchies and everything in-between.

The Role of Culture in Valuing Structure

As a particularly extreme form of structure-oriented practice, Holacracy depends not just on defining explicit expectations, but also living into them, and revising them as often as necessary to maintain legitimacy. Indeed, success and failure of Holacracy implementations seem to hinge upon how much the collective made explicit references in their internal communications to the structure they had created, a habit substantiated in Lee's (2019) empirical research. However, what has not yet been empirically substantiated, but which acts as an axiom for the structure-oriented approach, is that the mediating condition for success or failure is *how much a collective's culture values its structure*. More specifically, the key claim is that the health and stability of a group's Holacracy practice is reflected and reinforced by its habits and norms around evoking structure interactively in communications.

These norms express the group's collective beliefs about the importance, meaning, or utility of making their own expectations of each other explicit. A culture with a low value on explicit structure, for instance, would likely leave many formal expectations unwritten, treat all rules as equivalent in function and severity, assume expectations should be internalized rather than externally documented, and view explicit enforcement as a sign that something has gone wrong. Conversely, a culture with a high value on explicit structure would likely write down most or all formal expectations, distinguish among different types of rules and their associated

functions, and treat reminders and enforcement of explicit agreements as supportive rather than threatening.

The basis of this claim is that no matter how well designed the structures are, they do not self-define or self-enforce (Ostrom, 1980). Seen through this lens, one of Holacracy's innovations is that it provides many *structural solutions* to structural issues. For example, agents in the system are not going to trust a structure that is inconsistently defined, and Holacracy provides its own constitution of structure-language with stable definitions that everyone can understand (e.g., "An accountability is an ongoing activity") (Robertson, 2015). Similarly, agents are not going to trust a structure that is inaccurate or outdated, and Holacracy gives every circle a way to ensure its own inner structure is as accurate and current as it can reasonably be. Also, agents are never going to be able to objectively determine when something needs to be captured explicitly and when it can remain implicit, but Holacracy's tension-driven approach allows anyone in the system a pathway to propose a structural change while ensuring its final form has the consent of everyone impacted, or in lieu of that, a repair pathway.

But—and this is the key point—Holacracy never identified the *cultural* factors needed. Holacracy provides sufficient structural elements needed to make structure matter but neglects or willfully ignores the cultural elements. When Holacracy is understood not merely as a set of rules but as a practice that requires a sustained discipline toward referencing, updating, and acting from explicit structural agreements, its successes and failures become intelligible in ways that purely external analysis was unlikely to reveal. So, while SOC does not necessarily promote Holacracy, its central claims fully align with and support it.

The Missing Referents

An experienced Holacracy practitioner once said, "trying to explain Holacracy to someone who has never experienced it is like trying to explain how to change the oil in your cat." It is an expression that captures the multi-leveled complexities and constraints involved in communicating esoteric experience. This is relevant for understanding SOC's approach to rules

and collaboration because most accounts of rules and rule-following treat compliance and defiance as the only stances available to agents in relation to rules. Yet this framing obscures forms of lived rule-contact that become salient in full self-managing contexts.

One is the experience of using an explicit structure as an enabling constraint—one that reduces uncertainty, lowers coordination costs, and lets action proceed without continuous negotiation or oversight. This experience is familiar in traffic systems, games, and technical protocols, but it is far less commonly theorized in peer-to-peer collaboration. As a result, many theories lack the conceptual resources to recognize when rules are experienced not as instruments of control but as liberating structures that enable freedom. Through metaphors this experience can be suggested (e.g., how the lines on a highway constrain and thereby improve safety and traffic flow), but it is understandably difficult for many to intellectually transpose that felt sense into a work context without some form of direct experience.

Closely related, and equally under-modeled, is the complementary experience of holding the line—enacting accountability by leaning on an explicitly consented-to expectation rather than on personal authority or discretionary judgment. This stance is central to facilitation in self-managing systems, and to forms of structure-oriented leadership in which legitimacy flows from agreement and transparency rather than from status or position.

A third experience is also common in mature rule environments: encountering a binding expectation that one does not prefer, and may actively dislike, while still experiencing it as legitimately obligating. In this experience, the agent does not feel personally targeted and does not become resentful or alienated. They comply, or they pursue a legitimate pathway for revision, much as a rushed driver can dislike a red light while still treating it as having the right to bind them. Because these experiences are unevenly distributed in adult collaboration, they rarely appear as primary cases accounted for in rule theory, favoring constructs like “norms” within which explicit rules are merely a supporting mechanism. The claim here is that the absence of these experiential referents constrains what existing frameworks can see.

This epistemological constraint can be understood in semiotic terms. As Wilber (2001) argues, conceptual frameworks are limited not only by what they assert but by the experiential signifieds available to them; without an adequate experiential range, signifiers remain abstract and their referential power collapses. Where the enabling-constraint experience and non-dominating accountability are absent from the analyst's experiential repertoire, rules are inevitably interpreted through narrower lenses—typically as instruments of compliance, incentives, or control. A self-managing context however, makes it more obvious how all participants are necessarily engaged in the meaning making processes of legitimacy. The result is not merely incomplete explanation but systematic misrecognition of entire classes of rule-functioning phenomena.

Scope and Exclusions

Before getting into the specific distinctions and models SOC proposes, some framing is needed. This paper is a theory of deliberate collaboration, not a theory of relationships in general. The object of analysis is collaborations in which explicit, ongoing, prescriptive expectations are part of the medium of work. The focus is meso-level: it targets the internal mechanics of collaboration regimes—how obligations are authored, interpreted, enforced, contested, and revised within the collaboration itself. Meaning, this paper is not primarily concerned with interpersonal relationship dynamics as such. It does not attempt to explain attachment patterns or intimacy processes. It does not attempt to generalize across all human relating. Any relational examples that appear are used only as analogies for clarifying obligation mechanics, not as claims about relationship psychology.

Nor is this paper a field theory of collaboration. That is, it is not aimed at modeling predominantly emergent coordination carried by tacit alignment, high shared context, or moment-to-moment synchrony. Some collaborations do operate largely through implicit mutual attunement and they may be the preferred mode in certain domains. The present inquiry targets a different class of situations: deliberate coordination where participants must reliably navigate

explicit expectations, and where those expectations are treated as actionable obligations rather than mere descriptive norms. An analogy can be drawn to the difference between fast, automatic cognition and slow, reflective cognition that are called type 1 and type 2 decision making (Evans & Stanovich, 2013). The present focus sits on the “slow” side: it examines how people deliberately construct, reference, and repair explicit expectations, and how those expectations become usable as guides for action and accountability.

Finally, this paper is not a theory of society-level governance or political legitimacy. Terms such as legitimacy, alienation, and consent appear, but they are used in a constrained sense and at a constrained scale. The targets are rules and legitimacy as experienced and navigated inside collaboration systems, not local or state authority or macro-political obligation.

Core Models

To follow the claims advanced in this paper, several core distinctions must be provisionally understood. These distinctions are offered as the current working vocabulary of the approach; they are expected to be refined, and their boundaries may shift as the framework develops. Additionally, they are presented sequentially, but their interconnections and dependencies, while evident implicitly to the author, are still somewhat imprecise in how they are organized. Therefore, one of the research agendas is to further specify how all the models SOC generates can best be ordered and presented for different contexts.

The Drift Toward Explicitness

The first model presents the central conceit of the approach: explicit expectations are required to resolve diversity challenges. Several conjectures support the approach, and while they cannot be fully justified in this paper, they are necessary to establish if anything intelligible about SOC is to come through. They are presented here in their most economical form.

People need to coordinate across time and space. Coordination is the core problem of collaboration created by task interdependence, and it becomes harder as work spans distance, time, and individual differences since shared context and mutual knowledge tend to

degrade without deliberate mechanisms. Collaboration reliably generates artifacts and practices that stabilize expectations across contexts (Cramton, 2001; Hinds & Bailey, 2003; Malone & Crowston, 1994). The embodied nature of reality means on some level words and symbols are necessary to communicate across time and space.

Integrating and aligning individual diversity is a central collaboration challenge. Collaboration is not only a problem of aligning tasks but also of integrating heterogeneous perspectives, heuristics, and interpretations. Diversity can be a performance advantage in problem solving precisely because different cognitive repertoires explore the solution space differently, but that same heterogeneity increases the need for explicit coordination about meaning, roles, and expectations (Harrison & Klein, 2007; Hong & Page, 2004; Page, 2007).

There is a predictable drift toward making expectations explicit. As diversity increases, and if they do not want to rely on mindreading, groups need to make their shared expectations explicit. Those expectations are expressed in both descriptive (i.e., informational, observational, etc.) and prescriptive (i.e., injunctive, obligatory, etc.) ways, with the recurring prescriptive expectations being encoded into procedures, rules, and standards. This is done because informal memory and tacit, high-shared context alignment do not scale reliably. This drift is visible in classic accounts of bureaucracy as rule-governed administration and in empirical work treating formalization as a stable structural dimension that increases as coordination demands rise (Adler & Borys, 1996; Pugh et al., 1968; Weber, 1922/1978).

The drift is inherently awkward. Making expectations explicit is awkward because one party usually feels the need for more clarity before another. This means that even a sincere desire for clarification, arising as it typically does from one party before the other, is usually misinterpreted as pedanticism or some form of second-order power play (Brown & Levinson, 1987). At the very least, it tends to signal distrust, which is then met with distrust in turn. The point becomes clear by imagining the felt experience of making a low stakes promise to a friend and then being asked to put it in writing. This awkwardness also helps explain why explicitness

tends to increase under diversity. Hall (1976, 1990) argued that low-context communication is more common where people regularly transact across unfamiliar groups, because shared background knowledge cannot be assumed. In those liminal settings, clarity is not merely a cultural preference; it is a coordination requirement.

So, the drift toward explicitness can be read in two ways. It can look like a parochial bias in favor of low-context, Western styles (e.g., the United States), or it can be understood as a predictable adaptation that emerges when heterogeneous participants must coordinate without relying on tacit alignment. SOC adopts the second reading.

Explicitness-potential exerts a gravity on collaborative efforts. Coordination is not directionless: even when implicit alignment is functioning well, the underlying terrain is rarely neutral. As distance increases across space or time, and as heterogeneity increases across perspectives and lived experience, coordination tends to drift toward more explicit structures. Unlike physical gravity, however, this drift is not frictionless; rendering expectations explicit requires sustained effort. Collaborations nonetheless orient in that direction because the alternatives—default centralization of authority, substantive work stalling under the guise of harmony, or indefinite improvisation until participation erodes—are often more costly. The point is not that explicit structures are always desirable—they are not—but that recognizing this inherent directionality changes what is possible. It allows participants to anticipate thresholds at which small changes (a new member, a handoff, a time gap) are likely to flip smooth flowing implicitness into costly ambiguity; to normalize the growing need for clarity as a predictable shift in conditions rather than a failure of trust or competence; and to spend effort deliberately preparing for structural clarification so dysfunction can even become visible.

Therefore, if shared expectations drift toward explicitness, then the significant questions are how well participants will navigate that drift, how well will they use them, and how will they know when they are not needed. Critically, explicitness is not treated in SOC as inherently superior to implicitness. Coming away with that impression would be a

complete misreading of the approach. The label “structure-oriented” is meant only to convey that explicitness is a predictable feature of collaborative terrain. A conclusion which is intended to be somewhat intuitively obvious but one that is also empirically verifiable (DeChurch & Mesmer-Magnus, 2010). It does not ask whether a group will encounter the need for explicit structure, but when and under what conditions. And if that encounter arrives—whether prompted by a new member, a scaling threshold, a coordination failure, or simply the passage of time—the operative question shifts from whether to make expectations explicit to how well the group is equipped to create and navigate them. The mental models that follow are oriented by those questions.

Defining Structure

The second constitutive model of SOC is how this approach defines “structures” and “rules.” As a theoretical and methodological orientation concerned with understanding how rules function, SOC must clarify what it is treating as its central object of inquiry. Rather than resolving the complexity of the term “rule” through a stipulative definition, this approach clarifies how the term is already being used and relied upon in real collaborative situations. Following Rudolf Carnap’s (1950) account of *explication*, the approach adopts a mode of refining an imprecise concept while preserving continuity with ordinary and institutional use, prioritizing analytic adequacy over definitional fiat or mere cataloging.

Across disciplines, “rule” refers to a wide range of phenomena: legal mandates (Hart, 1961), strategic constraints (Schelling, 1960), constitutive and regulative structures (Searle, 1969), organizational policies (March & Olsen, 1989), common pool governance arrangements (Ostrom, 1990, 2005), cultural norms (Durkheim, 1893/2014; Sherif, 1936), computational algorithms (Knuth, 1968; Lessig, 1999), voluntary constraints (Suits, 1978), decision heuristics (Gigerenzer & Selten, 2001), and moral injunctions (Raz, 1979). These usages overlap in a family-resemblance fashion, but they do not converge on a single analytic object. As a result,

debates about rules often proceed past one another, with disagreement driven by divergent pragmatic interests embedded in competing definitions (Gallie, 1956; Schauer, 1991).

To address this, SOC introduces orienting definitions that progressively narrow analytic focus. At the broadest level, and in keeping with the natural language tradition in philosophy, rules may be understood simply but accurately as *requirements*: articulated standards specifying what must or must not be done under certain conditions. A second definition treats rules as *independent reasons for action* which means a rule influences an individual's behavior (when they do) by applying a felt sense of normative force. Said more simply, people will follow rules because they are rules (Schaurer, 1991). This distinction corresponds to what March and Olsen (1989) describe as a "logic of appropriateness" which is reasoning based on social expectations rather than a "logic of consequences" which is reasoning based on instrumental consequences. A third definition, which anchors SOC's analytic focus, defines a rule as any ongoing, explicit, and prescriptive expectation.

Ongoingness. This dimension distinguishes rules from one-off directives or isolated commitments. An ongoing expectation is designed to travel across time and space; meaning it must remain locatable, interpretable, and action-guiding even when the original authorities are absent and local contingencies vary (Hart, 1961; Schauer, 1991). Enduring rules disproportionately recruit deliberative, reflective processing because participants must reapply a generalized expectation to new particulars without relying on the situation to carry the norm (Evans & Stanovich, 2013; Kahneman, 2011).

Explicitness. This dimension distinguishes rules from tacit norms or background habits by making expectations referable—something people can point to, invoke, contest, or revise together. Norms are real and powerful, but they operate differently: they are inferred rather than articulated, enforced indirectly rather than explicitly invoked, and revised implicitly rather than deliberately. The combination of ongoingness and explicitness means that structures are

recorded somewhere, typically in written language, connecting them to theories of textual agency that show how texts can perform functions (Cooren, 2004).

Prescriptive expectation. This dimension distinguishes obligation from information, prediction, or preference. Prescriptive expectations establish standards that claim some authority over action—standards that participants can be held to, justify themselves against, or contest when circumstances change. The purpose of this distinction is not to separate "what is" from "what ought to be" as rigid philosophical categories, but to cultivate attentiveness to when descriptive beliefs and prescriptive propositions should be distinguished. Under this definition, rules include laws and policies, but also agreements, strategies, role definitions, priorities, projects, targets, heuristics, and values—provided these are operationalized prescriptively as obligatory considerations. The decisive question is not how an expectation is labeled but how participants use it to guide decision- and sense-making. The term prescriptive expectation (what one believes should happen) is used intentionally to preserve a clear parallelism with descriptive expectations (what one believes is likely to happen).

Two Types of Legitimacy: Structural and Experiential

Legitimacy, as observed through this approach, is an explanatory construct for why officially valid expectations can remain inert or resisted and why informally issued expectations can nonetheless bind powerfully. In its current formulation, SOC models legitimacy as operating across two analytically distinct types—structural and experiential—and treats the alignment or misalignment between them as a primary source of intelligibility for lived obligation in collaborative systems. This model also represents a formulation aligned with existing theories of rules and legitimacy theory work but presents a novel construction of its own.

The structural type of legitimacy concerns whether an expectation is valid according to the system's own second-order procedures: recognized pathways for authorization, standing, interpretation, exception, and revision. From this perspective, an expectation is valid when it is

created in the manner specified by higher-order norms that govern the production of lower-order norms (Hart, 1961; Kelsen, 1967).

The experiential type of legitimacy concerns whether the expectation is encountered by agents as having *the right to obligate them*. This is not the same as whether the expectation is liked or endorsed. Rather, it concerns whether the expectation shows up as properly issued from the agent's standpoint—typically consensual, noncoercive, but is actually equivalent to whatever the normative standard the agent may have. This parameter is described in Key Claims as the principle of non-alienation. An agent can therefore judge an expectation to be structurally legitimate while denying its experiential legitimacy: the expectation may be recognized as officially valid and yet interpreted by the agent as somehow lacking the right to bind them.

These two types of legitimacy provide two vectors that yield four logical configurations. When the registers align, obligation can be stable without requiring continuous persuasion or coercion. When structural legitimacy is high but experiential legitimacy is low, the same formal structure may become inert or resented even while remaining "valid" in the system-facing sense. When structural legitimacy is low but experiential legitimacy is high, expectations that lack structural legitimacy can nonetheless bind through social pressure, charisma, threat, or informal dependency—functioning as shadow authority that is difficult to contest precisely because it is not encoded through explicit channels. When neither is present, little normative force is generated.

This interdependent dynamic between legitimacy types explains why the same expectation can be encountered differently depending on history: low-trust and high-trust contexts alter how agents interpret the same structural move, and those interpretations in turn shape whether legitimacy or trust is reinforced or reduced over time.

In addition, while most theories of rules consider defiance and compliance as the two most significant indicators of how effective a rule is, the model of the types of legitimacy

highlight that the behavioral axis is not sufficient. Conflating behavioral alignment with normative standing is a recurrent error in many theories of rules and social order. Importantly, the behavioral quadrants (upper right and lower right in the AQAL model) are equally considered, but the point is that compliance or defiance is downstream and semantically underdetermined: the same outward behavior can be generated by materially different normative conditions. Illustratively, identical compliance behavior can express distinct configurations of legitimacy:

(a) Structural legitimacy and experiential legitimacy: compliance (healthy).

(b) Structural legitimacy without experiential legitimacy: compliance (resentful).

(c) Experiential legitimacy without structural legitimacy: compliance (exploitative).

(d) No legitimacy at all: compliance (fear-based, situational).

Same behavior, different meaning. That is precisely why this approach does not treat compliance as the hinge and instead locates the hinge in the participant's experiential obligation without excluding behavioral concerns.

The Structure System: Exploring, Encoding, Enforcing, and Evolving

SOC must supply not only claims but usable analytic lenses that make the target phenomena visible and one such lens, offered here in initial form, is viewing structure as a system of norms. The Structure System model provides a practical model for reasoning about how rules come into being, acquire binding force, and sustain or lose legitimacy over time. It begins with the acknowledgement that rules efficacy depends on how they are interpreted, enacted, and maintained within shared practices. To make this visible, the model distinguishes four analytically separable but dynamically interdependent dimensions: Exploring, Encoding, Enforcing, and Evolving (See Figure 1). Each aspect corresponds to a distinct dimension in a web of norms that support (or fail to support) the experiential legitimacy of a given rule and is oriented by a characteristic question that shapes its contribution. The four dimensions of this model represent SOC's primary diagnostic framework, though their boundaries may be redrawn as it is tested in new domains.

Exploring addresses a critical question: is a structural intervention warranted at all? This dimension resists the assumption that every issue requires a new rule, policy, or agreement. It attends to the nature of the issue, who experiences it, and which stakeholders should be invited

Figure 1

The Structure System

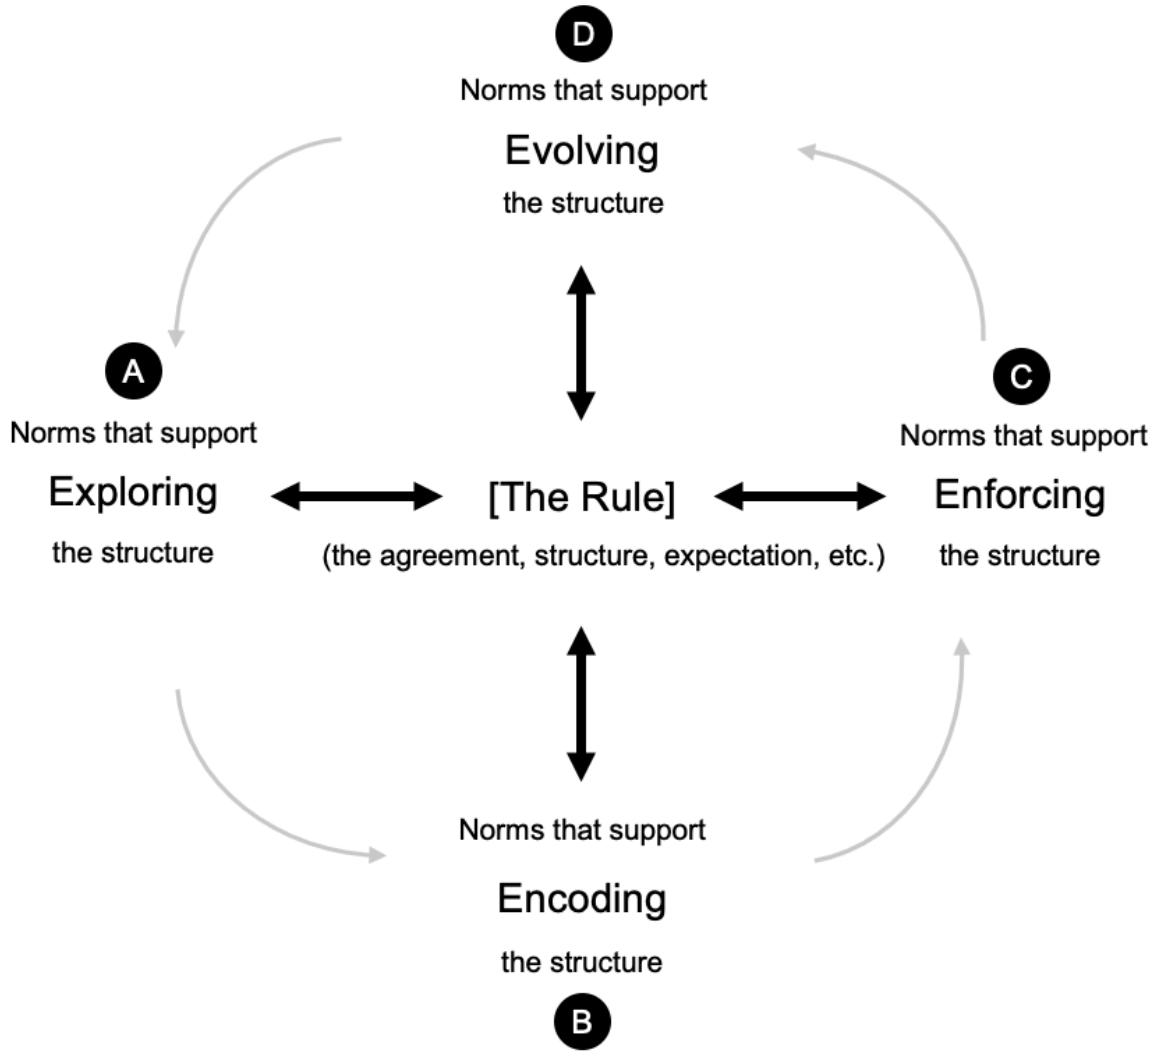

into the inquiry. It also examines existing structures, tools, or environmental conditions that already shape the situation, and explicitly includes the possibility that the appropriate response may be nonstructural. When deliberate exploring is absent or truncated, systems predictably accumulate misaligned expectations that draw down on perceived legitimacy.

Encoding addresses the question: when precisely does an expectation become binding? Encoding marks a transition from information, suggestion, or recommendation to official obligation. It raises questions of standing and legitimacy: who has authority to encode the expectation, through what process, and with what recordability and communicability? Encoding also includes formulation choices—how specific or general the expectation should be and how much interpretive discretion it affords—elements which are improved through effective exploring practices. Expectations that are difficult to revise require a higher legitimacy threshold at encoding, whereas expectations designed for easy modification (i.e. evolving) carry lower risk.

Enforcing addresses the question: how does a structure continue to matter in practice? In this approach, enforcement is not synonymous with punishment, nor is its primary function compliance. Enforcing is the ongoing work of stabilizing a structure's salience as an independent reason for action. This includes responding to violations and calibrating consequences, but also providing reminders, adjusting access or tooling, and supporting sensemaking. Enforcement failures often reflect upstream deficiencies—vague encoding or phrasing, an insufficient or unclear justification, or missing change pathways—elements which weaken the enforcer's confidence in enacting correction. Further, enforcing concerns clarify what should be encoded through questions like what can realistically be identified and enforced? Is the system prepared for the administrative burden required to manage this expectation well?

Evolving addresses the question: how does a structure remain legitimate over time if there is not a way to change or remove it? This dimension maintains relevance and systemic health by preventing expectations from persisting beyond their usefulness. Review cycles, expiration dates, trigger conditions, and opt-out provisions reduce the stakes and approval thresholds, enhance the confidence of agents and enforcers to lean on their interpretations and judgment. Evolving is especially critical where affected agents were not meaningfully involved in

initial exploration or encoding, because it provides a mechanism through which legitimacy can be restored or renegotiated through lived experience.

For example, systems like Holacracy and sociocracy strengthen experiential legitimacy through consent-based procedures, but consent is often unavailable or impractical in many systems (i.e. parent-child collaborations, traditional management hierarchy, etc.). One solution, therefore, is to replace the whole system of governance so that consent can be achieved. But another solution, made clear in SOC, is the recognition that consent is only one of many possible pathways for increasing experiential legitimacy or reducing alienation potential (alternatives include adding exception pathways, changing the rule from a requirement to a required consideration, lowering evolving costs, etc.) And therefore, collaborations have far more options than they realize for making their own rules matter to the participants who constitute them.

Across all four dimensions, legitimacy functions as a mediating condition between explicit expectation and practical obligation. When experiential and structural legitimacy are present, agents tend to experience structures as legitimate constraints making it more likely they will produce the rule's or the system's intended results.

The Four Participant Roles: Agent, Author, Enforcer, and Analyst

The preceding models describe how expectations are structured, differentiated, and sustained. But none of these models specifies who is encountering the expectation and from what functional standpoint. This matters because the same rule generates different saliences, different risks, and different characteristic tensions depending on the position a participant occupies in relation to it. A model of structure that does not account for these positional differences cannot diagnose where a breakdown is occurring or what kind of repair is even available. Current rule theories tend to use a binary rule-makers versus rule-followers distinction, which likely reflects the presuppositions of non-self-managing contexts in which people are equated with positional roles, but peer-to-peer collaborations embrace and

operationalize the simple fact that a person who defines a rule may also be subject to it. Therefore, a model is needed to appropriately account for the functionally distinct roles.

SOC addresses this by distinguishing four participant perspectives: agent, author, enforcer, and analyst. These are not formal job titles or fixed identities; they are functional standpoints that any participant may occupy—often simultaneously—in relation to a given expectation. A manager who authored a policy may also be an agent bound by it, and an enforcer responsible for responding to deviations, and an analyst reporting to a superior on whether it is working. The perspectives are analytically separable even when they converge in a single person, because each foregrounds a distinct set of concerns.

The agent is the participant to whom the expectation applies—the one for whom the rule is intended to generate a sense of obligation. The agent's characteristic tensions concern intelligibility, consistency, fairness, predictability, and feasibility. Can I determine what is expected of me? Is this expectation applied consistently? Is compliance realistic given my actual conditions? When agents cannot answer these questions, the rule's binding force degrades regardless of its formal validity. Navigability, from the agent's standpoint, means the ability to orient to an obligation and act on it without having to infer hidden requirements or rely on interpersonal relationships to determine what counts as compliance.

The author is the participant with standing to officially encode an expectation—typically the one who felt the tension that led to the expectation's creation in the first place. A conventional manager drafting a policy is acting as an author; but so are circle members in a Holacracy governance meeting who propose and accept a new role or accountability. The author's characteristic tensions concern compliance and the administrative overhead of managing the expectation because they want the rule to do work for them to reduce coordination costs (for themselves or the system). However, self-managing contexts also make it clear that since this is the case, authors also implicitly bear responsibility for the rule’s enactment as it is their tension it is intending to solve.

The enforcer is the participant positioned to respond to deviation in a way that preserves the rule's structural legitimacy—that is, *making the words matter*. Enforcement may include reminders, boundary-setting, consequence assignment, escalation, or repair invitations. Critically, the enforcer's standpoint is active well before any violation occurs, because the question of what would happen if the expectation is not followed is part of what makes the expectation intelligible and stable in the first place. The enforcer's characteristic tensions concern why agents are not following the rule and what constitutes an appropriate response to a violation—questions that implicate both the expectation's encoding quality and the available enforcement pathways.

The analyst is the participant positioned to describe, evaluate, or diagnose what is happening structurally and culturally with how the rule is enacted. This might be an HR practitioner assessing how a policy is being enforced, an outside consultant brought in to examine a governance breakdown, or a team member who steps back from their other roles to ask whether the structure system around a rule is working as intended. Critically, in SOC, the analyst remains a participant in the system—this perspective does not signify a detached external observer. The analyst's characteristic tensions concern identifying where in the system a problem is occurring, which phase of the structure system is implicated, and what, if anything, should be done about it.

Together, these four perspectives operationalize what SOC means by participant navigability. Existing analytic grammars—such as Ostrom's (2005) Institutional Analysis and Development framework—are designed primarily for cross-institutional comparison by external researchers, a purpose that privileges standardization and observational objectivity. That design is legitimate and productive for its intended use, but is optimized for understanding at a distance, not to aid participant navigation.

SOC's participant perspectives are designed to complement such frameworks by providing a diagnostic interface that participants themselves can use: a shared language

through which agents, authors, enforcers, and analysts can identify what is salient from their respective standpoints, locate the source of a coordination breakdown, and determine what kind of structural repair is available—without requiring a heroic amount of interpersonal negotiation or social power to make expectations predictable.

This orientation toward participant-facing theory also reflects a methodological commitment. Tensions—understood here, following Holacracy's usage, as a felt sense of the gap between current reality and a sensed possibility—serve as the primary operational input through which participants surface misalignments in the Structure System. Once tensions are treated as legitimate inputs rather than private emotions, the operative questions are who is feeling the tension, what it is, and from which functional role are they feeling it. The four roles supplies that precision. It clarifies not only what a participant is experiencing, but what kind of response the system needs: an encoding repair, an enforcement recalibration, an evolutionary review, or a diagnostic reframe. In this way, the participant perspectives are not an optional taxonomy appended to the framework; they are the navigational interface through which the Structure System's phases become actionable for the people who must enact them.

The Three Registers: Individual, Interactional, and Institutional

Last, but certainly not least, is the model of how SOC spans three independent registers of practice and analysis. Most frameworks for understanding or improving collaboration operate at a single scale. Organizational governance systems design roles and decision-making protocols but go silent on how individuals manage their own commitments or navigate a difficult conversation. Communication methods address real-time dialogue but do not design authority structures. Personal productivity systems organize individual attention but offer no account of institutional coordination. In practice, however, participants move constantly between these scales—managing their own commitments, engaging in dialogue with others, and operating within collective structures—often within the same hour.

SOC addresses this by modeling collaboration across three analytically distinct but structurally homologous registers: the individual, the interactional, and the institutional. These registers are not proposed as a novel ontological claim; they correspond roughly to familiar distinctions among personal, interpersonal, and organizational domains. What is proposed is that the same obligation mechanics—how expectations are explored, encoded, enforced, and evolved—recur across all three, and that recognizing this recurrence has both explanatory and practical consequences.

The individual register concerns how a person manages their own commitments, tracks tensions, and structures their environment to support outcomes they wish to achieve. The exemplar practice is likely David Allen's Getting Things Done (GTD) approach, whose training and coaching companies, not incidentally, practice Holacracy. The commonalities between Holacracy's paradigm and GTD's was expressed by Allen himself when he described Holacracy as "GTD for organizations" (David Allen Company, 2024). SOC then, is merely the attempt to more fully describe what those commonalities are, and provide evaluative criteria from which they can be judged.

The individual register also encompasses internal negotiation among competing priorities and commitments—what might be understood, by analogy to Schwartz's (1995) Internal Family Systems model, as something akin to *inner collaborative systems* whose parts must coordinate under shared purpose. Exemplar practices for this sub-type of the individual register are Kelley's true purpose work (2009) and Kegan & Lahey's immunity-to-change mapping (2009).[2]

The interactive register concerns real-time and asynchronous communication between participants—dialogue, meetings, and the back-and-forth through which expectations may be clarified, negotiated, contested, and encoded. This register foregrounds a paradigm of

[2] Notably, the immunity-to-change mapping process does not explicitly use inner parts but uses a functional equivalent construct of "competing commitments."

interactivity oriented around making one's desires and requests explicit. Exemplar practices for request-based structurally sensitive communication protocols are Nonviolent Communication (NVC) (Rosenberg, 2015) and Meta-Relating (Porcelli, 2019). In addition, facilitation methods, consent-seeking processes, and objection-integration protocols all operate at this register by structuring how contributions are solicited, sequenced, and resolved. One notable collection of these paradigmatic facilitation processes presented in Lipmanowicz & McCandless's book appropriately titled "Liberating Structures" (2014).

The institutional register concerns a collective's governance architecture—roles, policies, decision-making protocols, authority structures, and the second-order processes through which these are created and revised. As mentioned, exemplar practices include agile project management, Holacracy and sociocracy. This register addresses the broader context outside any specific interaction, though it remains relevant within them. In the current formulation, SOC does not draw a sharp analytic boundary between "team" and "organization" at this register, treating the distinction as a matter of scale rather than kind—an expression of the participant-navigability principle, which holds that distinctions should help participants orient rather than impose categorical divisions that do not track their experience.

The analytic value of this three register model lies in two properties. First, it enables learning transfer: because the same obligation mechanics operate across registers, skill or insight gained at one scale develops capacity at the others. A person who becomes skilled at managing individual commitments can locate structurally homologous dynamics in organizational governance, and vice versa. This also means that institutional dynamics are not exotic phenomena requiring positional authority to understand; they are available in a participant's own experience at any scale. Second, the registers are independently deployable. A participant can adopt structure-oriented practices like GTD at the individual register without requiring institutional adoption, and an organization can implement structural interventions without requiring that every individual adopt the full framework. The registers do not impose a

coordination requirement or a critical-mass threshold for basic functionality and improving performance at one register does not degrade performance at the others, yet it still allows for the possibility of combining effects.

The point is that practices such as GTD, NVC, Holacracy each operate at one register and that SOC does not compete with them but seeks to make visible the structural patterns that operate underneath and across them, offering a shared grammar through which their respective contributions can be understood. Many practices express different versions of identifying meaningful obligations, making those obligations explicit, and then using those explicit structures as tools for decision- and sense-making, and while some have sensed some similarity in practices at different registers, as David Allen did regarding the relationship between GTD and Holacracy, no one has articulated what those similarities are. This is one aim of SOC.

Taken together, all the previously presented core models, from the drift towards explicitness to the three registers, provide a small sampling of the theories generated from this approach. They are meant to serve as introductory explanations, as each one would require a dedicated paper to describe sufficiently, but are provided to convey a general sense of SOC as the broader meta-theory.

Methods and Approach

This paper is grounded in longitudinal first-person inquiry conducted within the formalized self-management practice Holacracy. And the purpose of the paper is to describe SOC as a meta-theory, and the aim of SOC is to render the internal mechanics of rule enactment and legitimacy-sensing intelligible in a way that can be scrutinized by others. But this research program was initially just an attempt to understand why Holacracy seemed to work in some groups and not in others.

At this point, a positionality note is warranted. My experience with Holacracy was not “I tried it and reflected on it.” I was deeply embedded in Holacracy’s central development-and-diffusion organization, where I worked as one of the few Certified Holacracy

Master Coaches. My individual contributions included refining Holacracy's coaching and training materials, creating asynchronous online learning programs, and conducting Holacracy Coach certification assessments with which I certified a majority of the world's Certified Holacracy Coaches personally. The depth and extent of that experience made certain collaborative mechanics unusually visible. It also created predictable risks, including overidentification and selective attention which could make those contexts seem more representative than they were. Therefore, throughout the formulation of this approach, I treated the emerging distinctions as defeasible and repeatedly tried to break them, both through direct disconfirmation attempts and through adversarial testing with practitioners and clients. Though admittedly, primarily still in self-managing contexts.

In terms of Integral Methodological Pluralism (IMP), the paper primarily enacts a Zone 1 ("looking from") method oriented toward Zone 7 ("looking at") phenomena (Wilber, 2006) (see Figure 2). Meaning I used first-person inquiry to examine the lived enactment of explicit rules and the participant-level sensing of legitimacy through which a collaborative system is reproduced and revised in practice. Importantly, this research method choice was not deliberately weighed against alternatives but was enacted as a reflection of my own perspectival bias toward phenomenological meaning-making.

The point of this disclosure isn't to stipulate domain expertise or authority, but to be transparent in the actual method used to develop the approach. Different methods disclose different domains of data, and that responsible metatheoretical work states its methodological commitments along with their limits (Bhaskar et al., 2015). In that spirit, this paper treats first-person analysis as a legitimate hypothesis-generating access condition. It also treats later multi-method research as future work rather than an implicit claim already satisfied.

In IMP terms, the injunction enacted was sustained participation in the practice paired with disciplined reflective analysis. Practically, the core analytic procedure was iterative concept formation in writing. I generated candidate distinctions and then repeatedly stress-tested them through explicit attempts at disconfirmation, through contestation within a community of practice and through description and application with clients to see whether the distinctions remained

Figure 2

Simplified Kosmic Address

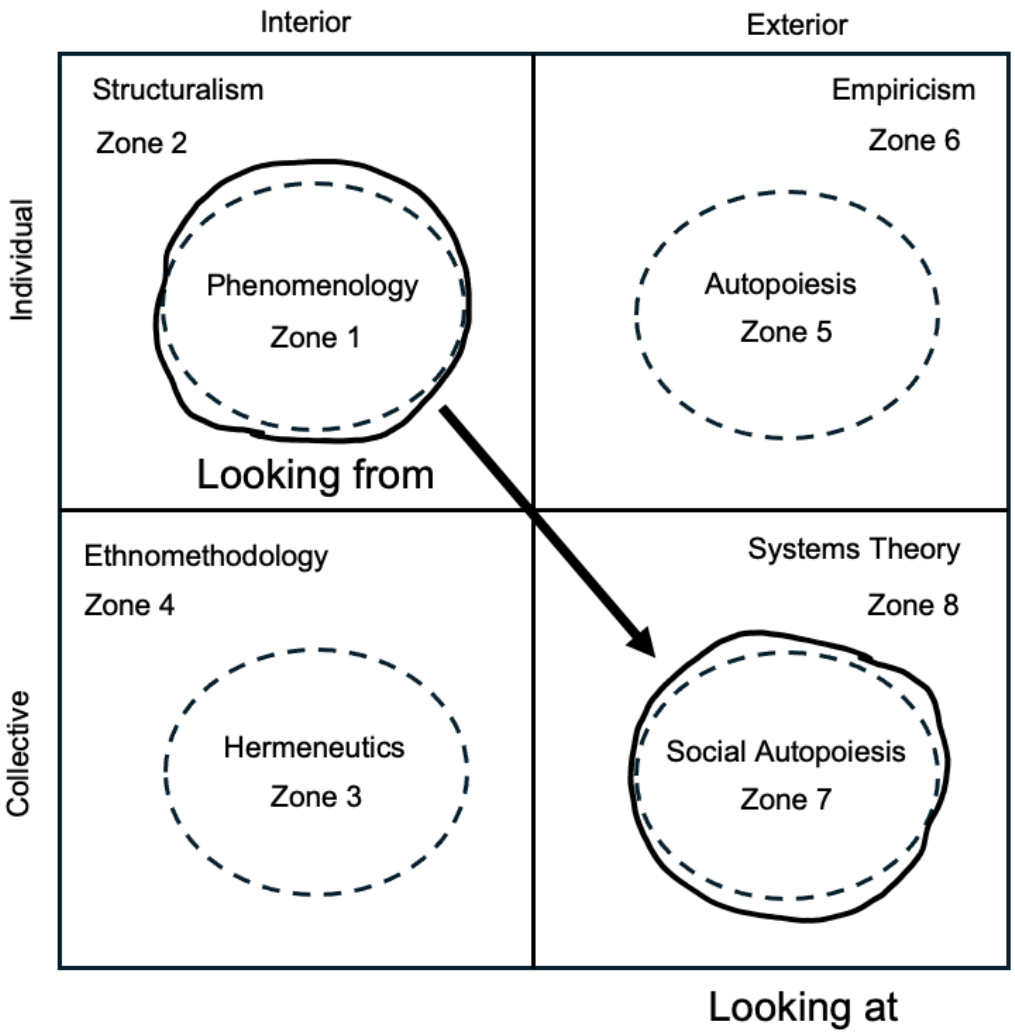

useful as evidenced if and when my distinctions and models were subsequently presented back to me.

The data for this inquiry is therefore practice-based and first-person. It consists of reconstructed episodes, iterative conceptual models—shared and tested publicly—in which distinctions either held up, required revision, or were dropped. The main analytic outputs are explicit conceptual distinctions and claims about interactional mechanics. The validity standard is likewise practice-near: in this paper, “validity” primarily means explanatory traction and recognizability for informed readers. It does not mean outcome confirmation.

A predictable limitation is that the paper does not provide independent outcome measures and does not offer external validation in the conventional empirical sense. Therefore, the warrant standard used here is abductive. Following Peirce (1903/1998), the threshold for consideration is plausibility and pursuit worthiness—whether the framing renders salient “surprising facts” intelligible enough to justify the cost of further inquiry. On that basis, the contribution is best read as a framing of an emerging approach that is now ready for critique through peer review and for follow-on testing through multiple methods, including empirical research (Hedlund, 2010).

Finally, the scope claim is intentionally structured to match both IMP and the abductive posture. The paper offers conceptual distinctions intended to travel across a broader class of collaborative systems and to be tested in diverse contexts. Follow-on testing can occur through peer review, practitioner replication, and targeted empirical studies that bring in additional zones and additional measures (Bhaskar et al., 2015; Wilber, 2006). It is for these reasons SOC is presented as both an integral meta-theory and a research program.

Evaluative Presuppositions

This analysis, much like 2nd tier constructive developmental meaning making, is oriented toward frameworks and organizing principles for integrating differences without losing normative capacity. However, neither is value neutral. In the spirit of IMP, this approach treats every disclosure of a phenomenon as occurring from within a standpoint, with associated

methodological choices and normative loadings. The claims presented herein should be read as constrained by a small set of explicit (of course they would be) evaluative presuppositions.

These presuppositions do not function as comprehensive moral axioms and are not presented as a deductive foundation for the conclusions, but function as orienting principles that shape what counts as a "good" description, what counts as a "problem," and which tradeoffs are treated as admissible or unacceptable. In practice, they operate as criteria for judgment when interpreting structures, agreements, and enforcement dynamics in collaborative systems. They foreground legitimacy and participant experience alongside effectiveness.

Principles Guiding This Approach

Congruency and coherence. This principle holds that any conceptual distinctions should be internally coherent and mutually congruent, and that the practical enactments implied by those distinctions should align with what participants can do in lived collaboration. One enactment is the demand that words and actions remain aligned enough to preserve trust and legitimacy. Another enactment is methodological: when evaluating structure, the approach prioritizes structural critiques—such as assessing whether a regime's expectations are calibrated to balance stability and flexibility—rather than merely reflecting aspirational ideals. As Austin (1975) put it, "Accuracy and morality alike are on the side of the plain saying that our word is our bond" (p. 10).

Structure-orientation. This principle holds that, when diagnosing coordination problems and proposing repairs, the first explanatory and practical investigations should be placed into the role of explicit expectations. This does not deny psychological or cultural factors; it establishes a priority. Where other approaches default to motivation, personality, or culture as primary levers, this approach treats structural conditions as the initial locus of analysis and then integrates other factors as needed.

Participant navigability. This principle holds that the primary test of a conceptual distinction is whether it improves participant navigability. A distinction should help participants

orient, interpret, and resolve their sensed tensions in the situations they face. This implies that participants are the primary consumers of theory, not merely objects of explanation. Accordingly, rather than centering the classic is–ought dispute at the level of abstract structure, the approach locates normativity at the interaction level by tracking how tensions are experienced and resolved from within the four participant perspectives. This principle consciously prioritizes meaning making over technical precision.

Non-alienation. This principle holds that collaboration regimes should avoid adding barriers to legitimacy. In SOC, experiential legitimacy refers to the participant's ability to recognize a structure's standing to obligate them without damage to self-concept and without undermining the regime's viability. This principle functions as a normative constraint: it is the standard against which structures and structure systems are evaluated, rather than relying only on effectiveness outcomes.

Embrace all structures. SOC treats collaboration as shared navigation within a constraint-laden environment. In this framing, "structure" is not limited to rules and agreements. It also includes any identifiable constraints embedded in the material and technical environment, such as system permissions, but also elements of the physical (e.g. room layouts) or temporal (e.g. time-boxes and deadlines) environment. This inclusion of a broader class of structures is necessary to include even if the analytic focus is still rules as a real-life example can illustrate.

During a Holacracy governance meeting, a participant wanted to restrict the ability of regular employees to access a sensitive software system. The participant asked their meeting facilitator if the proper governance construct was what Holacracy calls a "domain," which acts as an official policy granting a single role or circle centralized control over some asset or process. "Yes," the facilitator responded, a domain was the proper Holacracy structure to do that, however the facilitator had a question: "aren't you the systems administrator for that system?" "Yes," the participant responded. The facilitator then asked, "and do you have the ability to adjust permissions of who has access through an admin panel?" "Yes, I do," they answered.

The facilitator then clarified, “Well, if you have the technical ability to restrict their access via system permission, then you should just do that. We don’t want to have domains and policies if we can help it, and while that threshold should remain quite low, if you can shape the environment through available materials, then that’s a more elegant solution.”

So, even though the emphasis is on ongoing explicit and prescriptive expectations, it cannot be cognizable for participants in the real world unless it accounts for broader classes of structures. It is this broader definition of structures that makes SOC interoperable with adjacent bodies of work, including human factors, user-centered design, performance support, habit formation, and decision architecture. The emphasis on rules remains prominent in practice, largely because mature theory about rule-function in lived collaboration is comparatively underdeveloped, but the scope of SOC however, is intentionally general as an inquiry into structures as such, with rules as a salient but not exclusive subset.

Key Claims

The following claims represent the current state of the approach’s core propositions. They are advanced as testable hypotheses rather than settled conclusions. As SOC continues to develop, these claims may be revised, narrowed, or reframed considering new evidence.

Changing structures is usually more effective and humane than trying to change people

In collaborative systems, interventions that modify explicit structures will, in most cases, produce more reliable improvements in coordination than interventions aimed primarily at changing individual traits, attitudes, or motivations. This claim is grounded in a well-documented attribution bias: people tend to over-attribute others’ behavior to dispositions while under-weighting situational constraints, even when contextual factors plausibly account for the observed action (Ross, 1977; Gilbert & Malone, 1995). When that bias governs diagnosis, a predictable pattern follows: recurring coordination problems are interpreted as deficits in effort or character, and the dominant remedies become coaching, exhortation, or selection. By contrast, a structure-oriented intervention targets the conditions under which action is interpreted and

performed. It changes what is feasible, what is salient, and what is required, which can shift behavior without requiring continuous persuasion or exceptional personal discipline (Lewin, 1936).

The practical consequence is that structural modification is often a higher-leverage and more scalable move. A structural change can reduce variance across many participants at once, because it alters shared constraints rather than relying on individualized uptake. This is a central theme in systems-oriented accounts of error and performance: what appears as “human error” at the surface is frequently shaped by latent conditions, such as unclear task boundaries or misaligned incentives, which are more effectively addressed through redesign than through intensified blame or retraining alone (Reason, 2000). The claim is empirically evaluable: in comparable settings with recurrent breakdowns, a structure-first intervention should, on average, reduce rework and reduce recurring conflict frequency more than a people-first intervention, holding resourcing constant and measuring outcomes over a defined time window.

The same shift also tends to be more humane. Because it relocates explanatory weight from character to context, it reduces moralization and blame, which are common byproducts of dispositional diagnosis (Ross, 1977). It also preserves participant agency by treating individuals as capable actors operating within constraints, rather than as objects to be corrected. Finally, by making expectations more legible, a well-designed structure can reduce cognitive load and reduce the need for social guesswork, which lowers the interpersonal cost of coordination even when obligations remain demanding. This does not imply that individual development is irrelevant, or that personality has no place, but it predicts that, where recurring coordination failures are present, redesigning the structures that shape action will usually be the more effective and more humane starting point.

Non-alienation as normative criterion for evaluating structures systems

In its current formulation, SOC reframes legitimacy around a minimal design constraint: a structure should not degrade participatory agency. Participatory agency means an agent can

orient to binding expectations, act without having to infer hidden rules, and regain intelligible standing when friction occurs. Non-alienation is not endorsement or agreement. It is a “do no harm” standard for rule implementation.

Non-alienation is treated as a parameter rather than a single mechanism. Systems can realize it through different pathways, given context and constraints. What matters is whether the overall structure keeps agents out of guesswork and out of exposure to unconstrained discretion. This framing aligns with legal and social theory. Fuller (1969) argued that law fails as law when basic conditions such as clarity and congruence are violated. Empirical work on procedural justice shows that people comply more when processes are predictable, neutral, and respectful, even when they reject outcomes (Tyler, 2006). Political theorists likewise ground legitimacy in constraints on arbitrariness, emphasizing freedom from domination rather than consent (Pettit, 2012).

Within this approach, a structure violates non-alienation when it reliably produces standing-degrading conditions. These include illegibility, arbitrariness, punitive ambiguity, and unbounded enforcement discretion. These are disqualifiers, not merely missing features. Conversely, participatory agency can be supported through substitutable design moves. Some make obligations legible. Some stabilize change. Some align stated rules with operative practice. Some constrain enforcement discretion. Some provide restoration through repair or appeal. The evaluative question is whether the system, taken as a whole, preserves agents as participants who can orient, act, and recover without mind-reading or degradation.

This principle applies across the Structure System. Exploration keeps expectations grounded in actual conditions. Encoding makes obligations legible and authority explicit. Enforcement constrains discretion and maintains congruence between rule and response. Importantly, SOC allows for flexibility in recognizing that practical realities can rarely if ever allow all dimensions to meet this intention simultaneously. Instead, it views them as tradeoffs which can be locally managed. Evolution provides structured ways to surface tensions and adapt

commitments over time. The invariant question remains: does the structure, as enacted, preserve agents who can orient, act, and make meaning without being reduced to guesswork or exposure to arbitrary power?

When structures are under-differentiated, systems will compensate in other ways.

SOC treats selective enforcement and "culture problems" as partially explainable by a functional constraint internal to structure: prescriptive expectations must balance stability and flexibility. Stability supports coordination across time and distance. Flexibility supports workability under local variation. The claim is that when the structure system cannot represent that balance explicitly—by distinguishing kinds of expectations, pathways for exceptions, or different change rates—the system will reintroduce flexibility indirectly through predictable adaptations. These adaptations are not random, and they are not reducible to bad faith. They are incentive-consistent solutions to under-differentiated structure, and they can be evaluated using structure-based criteria such as legibility, congruence between rule and response, and boundedness of enforcement discretion.

Strategic non-enforcement is the pattern where the rule stays formally stable, but enforcement becomes flexible. Applicability and consequence become situational, and exceptions occur through discretionary judgment rather than explicit exception pathways. The predicted signature is a widening gap between published expectations and experienced consequences. Over time, accountability weakens because the rule no longer reliably predicts the response. Shadow power increases because authority migrates to informal channels—who can grant an exception, whose interpretation prevails, and what insiders know about "how it really works."

Official shoulds are the pattern where flexibility is purchased at the point of encoding. Expectations are written as "should" rather than "must," or otherwise left obligation-ambiguous, so the system can invoke the expectation without committing to consistent enforcement or clearly defined exceptions. The predicted signature is obligation contestability: accountability

conversations shift from “did you meet the requirement?” to “was it really a requirement?” This increases interpretive load, makes enforcement discretionary by default, and allows leadership direction without corresponding structural commitment.

Leaning on norms is the pattern where stability is maintained socially while flexibility is maintained interpretively, because expectations are not formalized but are still enforced through informal signals, status dynamics, or local correction. The predicted signature is onboarding cost and politics: newcomers must infer the real expectations, and disagreements become disputes about what is “normal” rather than about a shared, inspectable commitment. This can also punish difference, because deviation reads as a personal failure rather than as evidence that the expectation itself needs differentiation.

On this account, the presence of these three patterns is evidence about the structure system’s expressive capacity. Where differentiation is low, these compensations become more necessary. Where differentiation is higher—through explicit expectation types, explicit exception-by-request pathways, and differentiated change rates—the system can preserve stability without smuggling flexibility through uneven enforcement and can preserve flexibility without pushing authority into shadow channels. Critically, this places the normative anchor inside the concept of structure itself, rather than inside any particular culture or value system.

Implications

Structure-first leverage. Recurring coordination breakdowns should be treated, first, as hypotheses about the structure system, rather than as evidence of deficient motivation or character. Practically, this implies a structure-first sequencing discipline: modify explicit expectations and the pathways that make them usable before escalating to people-directed levers such as coaching.

Humaneness as legibility and bounded discretion. Humaneness, on this account, is not primarily an affective posture but a property of the collaboration’s legibility and constraint of arbitrariness. A structure is more humane insofar as it reduces participants’ need to infer hidden

rules, and limits exposure to unconstrained discretion, thereby lowering blame dynamics while preserving participant agency even under demanding obligations.

Legitimacy as a minimal design constraint. Evaluation should include a minimally normative criterion—non-alienation—as a constraint on what counts as an acceptable structure system. This implies that legitimacy can be assessed without appealing to endorsement or value alignment: the operative question is whether participants can orient to binding expectations, act without mind-reading, and restore intelligible standing after conflict.

Compensatory patterns as diagnostic signals. Patterns such as strategic non-enforcement, official "shoulds," and reliance on norms should be treated as diagnostic signals of under-differentiated structure, especially a limited capacity to represent the stability–flexibility balance explicitly. Their presence may imply predictable byproducts—shadow authority, politics, and onboarding cost—because flexibility is being "purchased" indirectly.

Limitations

As an initial formulation of an emerging research program, this approach is necessarily provisional and self-reflective. Several limitations warrant explicit acknowledgment, beyond the general qualification that the framework's claims are expected to be revised through continued inquiry.

First, this approach does not claim that explicitness is always desirable or that implicit coordination is inherently inferior. On the contrary, one of the central risks SOC seeks to illuminate is the misapplication of explicit prescriptive expectations in contexts where they displace tacit competence, erode relational trust, or introduce unnecessary rigidity. By foregrounding explicit structure as an object of analysis, the approach aims to clarify when such structures enable coordination and when they interfere with it. The framework is diagnostic: it helps identify mismatches between collaborative context and structural form, including cases where restraint in articulation is a condition of success rather than a limitation.

Second, this approach treats legitimacy as a central explanatory variable but does not provide a full normative theory of legitimacy. The two-register account is descriptive and interpretive: it explains when rules are experienced as binding, not when they ought to be. Developing normative criteria for legitimate rulemaking remains an open task that would require deeper engagement with moral philosophy, political theory, and domain-specific practice.

Third, this approach emphasizes participant interpretation while acknowledging the risks of partiality, blind spots, and asymmetries of power. Integrating participant-centered interpretation with external evaluation remains a methodological challenge, particularly in contexts where lived legitimacy diverges sharply from formal authorization. Future research will need to develop mixed-method approaches capable of triangulating lived experience, institutional design, and observed outcomes without assuming any single vantage point is epistemically complete.

Finally, this approach raises questions about scale. While the analysis spans interpersonal collaboration, organizations, and sociotechnical systems, further work is needed to examine how binding force operates differently across micro-, meso-, and macro-levels. Algorithmic governance and global coordination regimes, in particular, may require additional conceptual refinement. The central claim at this stage is that through the crucible of self-management, these distinctions became necessary. Through further reflection and empirical testing, the approach has reached a point where those distinctions must be made simpler and more requisite rather than extended outward indiscriminately—a constraint that reflects the framework's commitment to remaining deliberately minimal until its core propositions have been more thoroughly tested.

Conclusion

Contemporary collaboration relies increasingly on rules to coordinate action across time, distance, and difference. At the same time, distrust of institutions, resistance to imposed structure, and frustration with opaque governance systems continue to grow. Existing theories

explain parts of this tension but lack a shared object of inquiry capable of integrating issues of interpretation, power, legitimacy, and practice.

SOC is proposed as a response to this condition. By treating explicit, ongoing prescriptive expectations as its central object, this approach offers the beginnings of a field-level approach for explaining when rules function as binding structures, why they fail when they do, and under what conditions they enable collaboration rather than impede it. Its contribution, at this stage, is not to prescribe optimal systems or governance forms in the abstract but to specify the functional consequences of different structural choices in contexts where the drift toward explicitness is historically and developmentally predictable. This allows for continual recalibration between stability and flexibility, recognizing that balance is not achieved once but maintained through ongoing adjustment.

In doing so, SOC aims to support a more coherent intellectual commons—one capable of bridging disciplines without flattening them and of informing practice through a shared analytic grammar rather than a single prescriptive ideology. The framework remains deliberately minimal and open to revision, prioritizing descriptive coherence and empirical testability over theoretical completeness. Whether applied to organizational governance, algorithmic systems, or everyday collaboration, the study of explicit prescriptive expectations offers a promising site for understanding and navigating the common challenges of working together.

References

Adler, P. S., & Borys, B. (1996). Two types of bureaucracy: Enabling and coercive. Administrative Science Quarterly, 41(1), 61–89.

Austin, J. L. (1975). How to do things with words (2nd ed.). Harvard University Press. (Original work published 1962)

Bicchieri, C. (2006). The grammar of society: The nature and dynamics of social norms. Cambridge University Press

Bernstein, E., Bunch, J., Canner, N., & Lee, M. (2016). Beyond the Holacracy hype. Harvard Business Review, 94(7–8), 38–49.

Bettman, J. R., Luce, M. F., & Payne, J. W. (1998). Constructive consumer choice processes. Journal of Consumer Research, 25(3), 187–217.

Bhaskar, R., Esbjörn-Hargens, S., Hedlund, N., & Hartwig, M. (Eds.). (2015). Metatheory for the twenty-first century: Critical realism and integral theory in dialogue. Routledge.

Bourdieu, P. (1977). Outline of a theory of practice (R. Nice, Trans.). Cambridge University Press.

Brandom, R. (1994). Making it explicit: Reasoning, representing, and discursive commitment. Harvard University Press.

Brehm, J. W. (1966). A theory of psychological reactance. Academic Press.

Brown, P., & Levinson, S. C. (1987). Politeness: Some universals in language usage. Cambridge University Press.

Campbell, D. T. (1979). Assessing the impact of planned social change. Evaluation and Program Planning, 2(1), 67–90. (Original work published 1976)

Carnap, R. (1950). Logical foundations of probability. University of Chicago Press.

Chernev, A., Böckenholt, U., & Goodman, J. (2015). Choice overload: A conceptual review and meta-analysis. Journal of Consumer Psychology, 25(2), 333–358.

Constantine, L. L. (1993). Work organization: Paradigms for project management and organization. Communications of the ACM, 36(10), 35–43.

Constantine, L. L. (1995). Constantine on Peopleware. Yourdon Press.

Cooren, F. (2004). Textual agency: How texts do things in organizational settings. Organization, 11(3), 373–393.

Cramton, C. D. (2001). The mutual knowledge problem and its consequences for dispersed collaboration. Organization Science, 12(3), 346–371.

Dahlstrom, M. F., & Ho, S. S. (2012). Ethical considerations of using narrative to communicate science. Science Communication, 34(5), 592–617.

David Allen Company. (2024, September 10). Episode #275: GTD and Holacracy intersections. Getting Things Done. https://gettingthingsdone.com/2024/09/gtd-and-holacracy-intersections/

DeChurch, L. A., & Mesmer-Magnus, J. R. (2010). Measuring shared team mental models: A meta-analysis. Group Dynamics: Theory, Research, and Practice, 14(1), 1–14.

Durkheim, É. (2014). The division of labor in society (W. D. Halls, Trans.). Free Press. (Original work published 1893)

Edmondson, A. (1999). Psychological safety and learning behavior in work teams. Administrative Science Quarterly, 44(2), 350–383.

Espeland, W. N., & Sauder, M. (2007). Rankings and reactivity: How public measures recreate social worlds. American Journal of Sociology, 113(1), 1–40.

Evans, J. St. B. T., & Stanovich, K. E. (2013). Dual-process theories of higher cognition: Advancing the debate. Perspectives on Psychological Science, 8(3), 223–241.

Fox, C. R., & Rottenstreich, Y. (2003). Partition priming in judgment under uncertainty. Psychological Science, 14(3), 195–200.

Fuller, L. L. (1969). The morality of law (Rev. ed.). Yale University Press.

Gallie, W. B. (1956). Essentially contested concepts. Proceedings of the Aristotelian Society, 56, 167–198.

Giddens, A. (1984). The constitution of society: Outline of the theory of structuration. University of California Press.

Gigerenzer, G., & Selten, R. (Eds.). (2001). Bounded rationality: The adaptive toolbox. MIT Press.

Gilbert, D. T., & Malone, P. S. (1995). The correspondence bias. Psychological Bulletin, 117(1), 21–38.

Goodhart, C. A. E. (1975). Problems of monetary management: The UK experience. Papers in Monetary Economics, 1(1), 1–20. Reserve Bank of Australia.

Greenfield, R. (2015, March 30). Holawhat? Meet the alt-management system invented by a programmer and used by Zappos. Fast Company.

Hall, E. T. (1976). Beyond culture. Anchor.

Hall, E. T., & Hall, M. R. (1990). Understanding cultural differences. Intercultural Press.

Harrison, D. A., & Klein, K. J. (2007). What's the difference? Diversity constructs as separation, variety, or disparity in organizations. Academy of Management Review, 32(4), 1199–1228.

Hart, H. L. A. (1961). The concept of law. Oxford University Press.

Hedlund, N. (2010). Integrally informed qualitative research. [Unpublished manuscript].

Hick, W. E. (1952). On the rate of gain of information. Quarterly Journal of Experimental Psychology, 4(1), 11–26.

Hinds, P. J., & Bailey, D. E. (2003). Out of sight, out of sync: Understanding conflict in distributed teams. Organization Science, 14(6), 615–632.

Hong, L., & Page, S. E. (2004). Groups of diverse problem solvers can outperform groups of high-ability problem solvers. Proceedings of the National Academy of Sciences, 101(46), 16385–16389.

Hume, D. (2000). A treatise of human nature. Oxford University Press. (Original work published 1739–1740)

Hyman, R. (1953). Stimulus information as a determinant of reaction time. Journal of Experimental Psychology, 45(3), 188–196.

Iyengar, S. S., & Lepper, M. R. (2000). When choice is demotivating: Can one desire too much of a good thing? Journal of Personality and Social Psychology, 79(6), 995–1006.

Johnson, E. J., & Goldstein, D. G. (2003). Do defaults save lives? Science, 302(5649), 1338–1339.

Kegan, R., & Lahey, L. L. (2009). Immunity to change: How to overcome it and unlock the potential in yourself and your organization. Harvard Business Review Press.

Kahneman, D. (2011). Thinking, fast and slow. Farrar, Straus and Giroux.

Kelley, T. (2009). True purpose: 12 strategies for discovering the difference you are meant to make. Transcendent Solutions Press.

Kelsen, H. (1967). Pure theory of law (M. Knight, Trans.). University of California Press.

Knuth, D. E. (1968). The art of computer programming (Vol. 1). Addison-Wesley.

Kuhn, T. S. (1962). The structure of scientific revolutions. University of Chicago Press.

Lakatos, I. (1978). The methodology of scientific research programmes. Cambridge University Press.

Laloux, F. (2014). Reinventing organizations. Nelson Parker.

Lee, M. Y. (2019). Holacracy and the transformation of organizational governance [Doctoral dissertation, Harvard University].

Lee, M. Y., & Edmondson, A. C. (2017). Self-managing organizations: Exploring the limits of less-hierarchical organizing. Research in Organizational Behavior, 37, 35–58.

Lessig, L. (1999). Code and other laws of cyberspace. Basic Books.

Lewin, K. (1936). Principles of topological psychology. McGraw-Hill.

Lewin, K. (1951). Field theory in social science. Harper & Brothers.

Malone, T. W., & Crowston, K. (1994). The interdisciplinary study of coordination. ACM Computing Surveys, 26(1), 87–119.

March, J. G., & Olsen, J. P. (1989). Rediscovering institutions: The organizational basis of politics. Free Press.

Moeslein, K., & Laulitz, T. (2019). Self-management at scale: The Holacracy experience. In Handbook of research on management and organizational history. IGI Global.

North, D. C. (1990). Institutions, institutional change and economic performance. Cambridge University Press.

Ostrom, E. (1990). Governing the commons: The evolution of institutions for collective action. Cambridge University Press.

Ostrom, E. (2005). Understanding institutional diversity. Princeton University Press.

Ostrom, V. (1980). Artisanship and artifact. Public Administration Review, 40(4), 309–317.

Page, S. E. (2007). The difference: How the power of diversity creates better groups, firms, schools, and societies. Princeton University Press.

Payne, J. W., Bettman, J. R., & Johnson, E. J. (1988). Adaptive strategy selection in decision making. Journal of Experimental Psychology: Learning, Memory, and Cognition, 14(3), 534–552.

Peirce, C. S. (1998). Pragmatism as the logic of abduction. In The essential Peirce (Vol. 2, pp. 226–241). Indiana University Press. (Original work published 1903)

Pettit, P. (2012). On the people's terms: A republican theory and model of democracy. Cambridge University Press.

Porcelli, M. (2019, April 21). The three modes of conversation: A roadmap to navigating interactions. MetaRelating. https://www.metarelating.com/articles/the-three-modes-of-conversation

Power, M. (1999). The audit society: Rituals of verification. Oxford University Press.

Pugh, D. S., Hickson, D. J., Hinings, C. R., & Turner, C. (1968). Dimensions of organization structure. Administrative Science Quarterly, 13(1), 65–105.

Raz, J. (1979). The authority of law: Essays on law and morality. Oxford University Press.

Reason, J. (2000). Human error: Models and management. BMJ, 320(7237), 768–770.

Robertson, B. (2014, July 28). History of Holacracy: The discovery of an evolutionary algorithm. Holacracy. https://www.holacracy.org/blog/history-of-holacracy/

Robertson, B. J. (2015). Holacracy: The new management system for a rapidly changing world. Henry Holt.

Rosenberg, M. B. (2015). Nonviolent communication: A language of life: Life-changing tools for healthy relationships (3rd ed.). PuddleDancer Press.

Ross, L. (1977). The intuitive psychologist and his shortcomings: Distortions in the attribution process. In L. Berkowitz (Ed.), Advances in experimental social psychology (Vol. 10, pp. 173–220). Academic Press.

Schatzki, T. R. (2002). The site of the social: A philosophical account of the constitution of social life and change. Pennsylvania State University Press.

Schauer, F. (1991). Playing by the rules: A philosophical examination of rule-based decision-making in law and in life. Oxford University Press.

Scheibehenne, B., Greifeneder, R., & Todd, P. M. (2010). Can there ever be too many options? A meta-analytic review of choice overload. Journal of Consumer Research, 37(3), 409–425.

Schelling, T. C. (1960). The strategy of conflict. Harvard University Press.

Searle, J. R. (1969). Speech acts: An essay in the philosophy of language. Cambridge University Press.

Sherif, M. (1936). The psychology of social norms. Harper.

Simon, H. A. (1956). Rational choice and the structure of the environment. Psychological Review, 63(2), 129–138.

Slovic, P. (1995). The construction of preference. American Psychologist, 50(5), 364–371.

Strathern, M. (1997). “Improving ratings”: Audit in the British university system. European Review, 5(3), 305–321.

Tversky, A. (1972). Elimination by aspects: A theory of choice. Psychological Review, 79(4), 281–299.

Tversky, A., & Kahneman, D. (1981). The framing of decisions and the psychology of choice. Science, 211(4481), 453–458.

Tyler, T. R. (2006). Why people obey the law (2nd ed.). Princeton University Press.

Weber, M. (1978). Economy and society: An outline of interpretive sociology (G. Roth & C. Wittich, Eds.). University of California Press. (Original work published 1922)

Weirauch, L., Galliker, S., & Elfering, A. (2023). Holacracy, a modern form of organizational governance: Predictors for person–organization fit and job satisfaction. Frontiers in Psychology, 13, Article 1021545.

Wilber, K. (2001). Sex, ecology, spirituality: The spirit of evolution (2nd ed.). Shambhala.

Wilber, K. (2006). Integral spirituality: A startling new role for religion in the modern and postmodern world. Shambhala.