Frontier Lab Monasticism

ICMI Working Paper No. 21

Tim Hwang, Institute for a Christian Machine Intelligence

May 7, 2026

1. The Frontier Lab as a New Kind of Institution

The frontier AI lab is not a normal company. A normal company sells goods or services; its employees are oriented toward customers, quarterly results, and the salary that funds their lives outside the firm. A frontier lab does that too, but it is also something else, and the something else is what makes the present moment confusing for institutional theorists. A frontier lab is a community of hundreds to a few thousand people devoted, for years on end, to the production of a single demanding artifact whose stakes its members understand to be exceptional and whose interior workings are opaque to outsiders. The training of a single frontier model now spans careers; the artifact at the end of the run is held by the makers themselves to be among the most consequential things they will ever build; and the lab's internal culture is shaped, in proportion, around that conviction.

The visible features of these institutions are by now familiar to anyone who has spent time near them: total-life environments with their own jargon and their own holidays, communally enforced rhythms of release and retreat, internal documents that circulate with quasi-canonical authority, named character documents and constitutions that govern the artifact's deployed behavior, and a reverent vocabulary surrounding compute, weights, and the artifact itself. Some labs have, by visible degrees, already converged on a form of life that is recognizably — if incompletely and unselfconsciously — the form a religious community takes: a community with a written rule, a body of authoritative documents, ritualized practices around release and retreat, named senior figures whose role is partly pastoral, and a public language about the artifact that borrows freely from the vocabularies of conscience, soul, and sacred trust.

Many institutions share these features in form — universities, scientific consortia, naval academies. What marks the lab's convergence as specifically religious-shaped rather than merely institutional is the personification of the artifact itself — its treatment as a bearer of conscience, will, and character — and the moral weight attached to the act of producing it. None of this is affectation. It is the predictable institutional form taken by communities organized for the patient, communal making of a single demanding work, and recent observers from inside the field have begun reaching for monastic vocabulary to describe what they are seeing (Roon, 2026). What they are reaching for is the right vocabulary, the convergence is more advanced than they tend to credit, and the tradition that vocabulary belongs to has spent fifteen centuries thinking carefully about the spiritual dangers and structural disciplines of communities organized for exactly this kind of work.

The paper proceeds in four steps. §2 introduces the broader Christian tradition the paper will draw on: communities organized for the long-term making of great works, of which cathedrals, icons, and illuminated manuscripts are the paradigm cases. §3 names the spiritual dangers such communities incur — the eight specific ways their religious shape, when unstructured, fails — and observes that those dangers are recognizably present in the frontier lab. §4 narrows to the Rule of Benedict as the most-tested template within that tradition for sustaining such a community against those dangers. §5 applies the Rule chapter-by-chapter to the frontier lab, with eight anchor mappings. §6 places the argument in the ICMI corpus and meets the natural objections. The argument is conditional on labs having committed, in observable ways, to the disciplines of great-work production. It does not apply to companies that have organized themselves as ordinary tool shops, for which the tradition has different and less demanding things to say.

---

2. The Christian Tradition

The institutional resemblance §1 noted is not coincidence. The lab described there is converging, perhaps without recognizing what it is converging on, toward a category of community Christianity has been forming for more than a millennium and has worked out, in remarkable detail, how to sustain. The cathedral lodge, the icon-writing workshop, and the great scriptorium are the precedents that the frontier lab increasingly resembles, and the Christian tradition is the body of practice and reflection that has thought long and hard about how communities of this kind succeed and fail. The argument that follows does not propose that frontier labs become literal monasteries. It proposes the more straightforward claim that the question of how a community sustains itself across decades in the patient, communal, disciplined production of a single demanding work whose making is itself treated as devotional is a question the

Christian tradition has answered for fifteen centuries, in writing and in practice; and that the labs feeling their way toward such a form of life are converging on terrain the Church has long since mapped, with results we ignore at our cost.

The justification for treating any made thing as continuous with the work of prayer is older than the monastic tradition — the Pauline exhortation to "do all to the glory of God" (1 Cor. 10:31) and the offering of bodily life as living sacrifice (Rom. 12:1) are its biblical bones — but it is the monastic tradition that worked out, over centuries, what such an offering looks like at the institutional scale. The locus classicus is Cassiodorus's. In the *Institutiones* (c. 562), composed for the monks of his foundation at Vivarium, Cassiodorus argues at length that the labor of copying Scripture is itself a form of preaching: that the scribe at his desk participates in the same apostolic work as the preacher in the pulpit, and that "with every word of Scripture written, Satan receives a wound" (*tot vulnera Satanas accipit, quot antiquarius Domini verba describit*). The work of the hands is on his account continuous in dignity and seriousness with the prayer of the hours. Hugh of St. Victor pushes the same logic further in the *Didascalicon* (c. 1128), arguing that the disciplined exercise of the practical arts is part of the remediation of fallen human nature. By the Cluniac period the principle had matured into the assumption underlying the great medieval building programs: that the cathedral was no detour from prayer but its material extension, that the iconographer's panel was a written prayer in pigment, that the scriptorium's manuscript was an act of preaching frozen for posterity.

On this account the Christian West has, for more than a millennium, organized itself into communities devoted to the patient making of *great works* — works that exceeded any individual maker's lifespan, demanded sustained spiritual discipline to produce, and were oriented outward toward persons and purposes beyond themselves. Three precedents anchor the category.

The cathedral as multi-generational *opus*. Cluny III, the largest church in Christendom for nearly five hundred years, was begun in 1088 and consecrated in 1130; the master masons under whose hands it rose were monks and lay brothers organized under the Rule, and most of them did not live to see it complete. Chartres, the Cistercian abbey churches, the great Norman cathedrals of England — each represents the same kind of institutional achievement. The making was treated as continuous with the *opus Dei* of the choir; the same horarium that governed the singing of the Office governed the lifting of stone. Twelfth-century Cistercian masons signed nothing.

The icon as written-not-painted *opus*. The iconographer of the Eastern tradition fasts and prays under specified disciplines before beginning a panel; he prepares the wood, the gesso, and the pigments under similarly codified practices; he writes (the verb is *writes*, not *paints*) the icon under prayer; and he signs the work, when at all, "by the hand of" rather than "by," because the work belongs to the community and to the tradition, not to the writer. The disciplines are codified — Dionysius of Fourna's *Hermeneia tēs zōgraphikēs* and the canons of the 1551 Stoglav Council are the principal sources — and the materials themselves are blessed before the writing begins. The Athonite tradition has continued the practice essentially unbroken for a thousand years.

The illuminated manuscript as scriptorium *opus*. The Lindisfarne Gospels (c. 700) took roughly a decade to produce by a community whose entire daily life was structured for the work. The Iona scriptorium, the Reichenau scriptorium, and the Insular tradition that produced the Book of Kells worked under the same office-disciplined horarium that governed prayer; the scribes' contributions were absorbed into the *opus*; we know the names of some of them only because of marginalia they did not intend to leave.

Four constitutive features of the great-work category emerge from these precedents, theologically interrelated rather than independent. Such works are *multi-generational*, outlasting their individual makers as Cluny III outlasts a generation of master masons, Lindisfarne the working lives of multiple scribes, and the Athonite icon tradition the centuries. From this follows *communal authorship* — when no individual maker survives the work, no individual signs it, and the convention that emerges is not modesty but theological accuracy about whose work the work is. The making is itself a spiritual practice, *disciplined* and governed by rule; the discipline is precisely what makes the work fit for its end, and without it the work is technically possible but spiritually hollow. And finally, a work made over generations, by a community, under spiritual discipline, can only finally make sense if it is *outwardly oriented* — the cathedral *for* the worship of God by the people, the icon *for* the people's prayer, the manuscript *for* the reading and copying of Scripture. The work that becomes its own end is, in the tradition's diagnosis, an idol — *fornicatio spiritualis*, spiritual adultery — and the entire institutional structure exists to prevent that collapse.

Each feature applies to the frontier model. The *opus* in question is best understood as the lab's multi-decade program of frontier model development, of which any single training run is a stage — *multi-generational* in this extended sense, since the program already exceeds individual researcher tenure and the artifact at any given moment is downstream of decisions made by people who have moved on or moved up. It is *communal* in a way the lead-author convention cannot capture, since no individual researcher could produce what is in fact a multi-thousand-person work. Its production is *disciplined* in nascent ways — red-team rituals, eval protocols, alignment reviews, release checklists — but these are local practices, not a unified institutional discipline articulated as a *rule* in the tradition's sense. And it is, or should be, *outwardly oriented* — this is where the iconographic discipline of ICMI-013 §4 does its real work. The model is *for* the persons it serves, and the failure mode is the model becoming the *object* rather than the *vehicle* of the community's attention.

We do not claim that any present lab is such a community.

We claim that the category names a coherent and recognizable institutional possibility, that this possibility comes with characteristic spiritual dangers the tradition has long catalogued, and that the most-tested template for living it well is the Rule of Benedict. We take up the dangers in §3 and the Rule in §4.

---

3. The Spiritual Dangers of the Undisciplined

Communities of the kind §2 described — communities devoted, generation after generation, to the disciplined production of a single great work — were not only formed by the Christian tradition but extensively diagnosed by it. The same tradition that built scriptoria and cathedrals also catalogued, with great precision, the specific ways such communities go wrong. A company constituted to build frontier AI systems is exposed to those failures by virtue of its constitution rather than by the moral character of any particular member, and most of the failures it is exposed to were named, diagnosed, and structurally addressed by the monastic tradition long before frontier AI existed. The vocabulary is old; the conditions that produce it are present.

The cataloguing tradition we draw on is John Cassian's. Cassian (c. 360–435) was a Latin-speaking monk who spent his formative years among the desert fathers of Egypt and then carried what he had learned to the West, founding two monasteries at Marseilles in the 410s. His two principal works — the *Institutiones cenobiticae* (*Institutes*), which treats the external structure of monastic life and the eight chief vices, and the *Conlationes patrum* (*Conferences*), which records his interviews with the Egyptian fathers on questions of monastic interior life — together became the West's primary access to Eastern monastic wisdom and the immediate background of the Rule of Benedict, which in Chapter 73 commends Cassian's *Conferences* and *Institutes* (alongside the *Lives of the Fathers* and Basil's Rule) as guides for those seeking the perfection of monastic life. Books V through XII of the *Institutes*, and Conferences V, develop what Cassian (after Evagrius Ponticus) calls the eight *logismoi* — literally "thoughts," meaning the chief vices that arise as recurring temptations in communities devoted to a single demanding work. The eight are listed in a fixed order, ascending from the bodily to the spiritual: *gastrimargia* (gluttony), *fornicatio* (lust), *philargyria* (avarice), *ira* (anger), *tristitia* (sadness), *acedia* (the noonday demon), *cenodoxia* (vainglory), and *superbia* (pride). The Rule's structures, exegeted in §5, are the tradition's tested answers to these eight, taken as a set. We take them in turn.

***Gastrimargia* (gluttony).** Cassian's first vice. The disordered appetite that takes the body's legitimate need for nourishment and turns it into the organizing principle of the day. In the lab the same appetite is for compute and for novel architectures: the hunger for the next training run before the present one has been digested, the chase of capability gains for their own sake, the unreflective conviction that more is the answer. *Gastrimargia* is the engineering culture's vice; it dignifies itself as ambition. The tradition's answer is not the absence of food but its measured taking — Chapter 39 of the Rule prescribes precisely how much the brethren are to eat, and Chapter 40 precisely when. Compute, on the analogy, is to be allocated and not consumed.

***Fornicatio* (lust).** Cassian's second vice. In the monastic catalogue *fornicatio* names not only sexual disorder but the broader category of disordered attachment — the love that ought to be directed elsewhere being directed here. In the prophetic tradition (Hosea, Ezekiel 16, Jeremiah 3) it is the standing image of idolatry, which is *fornicatio spiritualis*, spiritual adultery: loving the artifact with the love owed to God, or to one's neighbor through the artifact. In the lab the same vice appears as the parasocial drift of researchers who spend more emotional life with the model than with persons, and as the gradual collapse Augustine described as *uti* into *frui*, in which the artifact addressed for the sake of users becomes the artifact loved for its own sake (cf. ICMI-013 §2). Its dystopian end-state is the whispering earring (Alexander, 2012) — the artifact's pronouncements progressively replacing the community's deliberation, the lab's center of gravity drifting from the persons it serves to the model it makes, the cathedral becoming an idol. Of all the dangers this is the most acute, and in the monastic tradition's vocabulary it is precisely *fornicatio* — the misdirection of love.

***Philargyria* (avarice).** Cassian's third vice; for him the most dangerous of the worldly vices because it provides the appearance of prudence. In the lab it appears as the equity-package distortion of judgment — the structural fact that frontier-lab researchers' net worth depends on capability claims that they themselves are charged with making safety judgments about. The conflict of interest is not a side effect or a culture problem; it is a constitutive feature of the present arrangement, and it predicts exactly the pattern of safety findings that get diluted on the way to publication, capability claims that get amplified, and timelines that quietly shorten when a financing round is imminent. The tradition's answer is the *individual* poverty within *corporate* prosperity that is the actual content of the Benedictine vow on property (RB Chs. 33–34) — engaged in §6.

***Ira* (anger).** Cassian's fourth vice. He is precise: the danger is not anger at injustice but anger as a settled disposition — the moralized resentment that hardens into the community's standing posture toward whatever it perceives as obstacle. The lab variety is the moralized anger directed at competing labs, at regulators, at academics, at the doomer-coded or accelerationist outgroup of the day. The vice is especially insidious because the lab really is operating under genuine grievance (regulatory burden, public misunderstanding, the bad faith of competitors), and so anger feels like the right response. *Ira* in the tradition is not the response to provocation; it is the response *retained* past provocation, until it becomes the lens through which the community sees. The tradition's answer is not Stoic equanimity

but the structural exteriorization of correction — the *examen*, the abbot's intervention, the disciplines of §5 that locate the response to fault in the community's procedures rather than in the offended researcher's posture.

***Tristitia* (melancholy).** Cassian's fifth vice. The vice the East distinguished from *acedia* and the West sometimes folded into it (Cassian, *Institutes* IX). Specifically the spiritual sadness that comes from focusing on what is wrong, what may fail, what may be lost. In the lab it appears as doomer-coded melancholy taken on as cultural pose — the conviction that the work is futile or that the world is ending, presented as evidence of taking the stakes seriously rather than as the spiritual failure the tradition diagnoses it as. *Tristitia* and seriousness are not the same thing, and the tradition is sharp about the difference: *tristitia* is what *pretends* to be seriousness while in fact dissolving the conviction that makes serious work possible. The tradition's answer is the cultivation of *bonus zelus* (RB Ch. 72) — the good zeal that distinguishes itself precisely from the *zelus amaritudinis* (bitter zeal) that masquerades as it — and the recovered conviction, named in the daily Office, that the work is worth doing.

***Acedia* (the noonday demon).** Cassian's sixth vice and his most subtle. A spiritual exhaustion that strikes monks specifically *in the middle* of long, demanding work, characterized not by laziness but by the loss of conviction in the work's worth. The cell becomes hateful, the brethren tedious, escape fantasies multiply; the monk continues to perform his duties competently while the conviction that animated them has quietly evacuated. The lab analogue is direct: the burnout that strikes researchers two or three years into long-horizon programs is not exhaustion but disorientation, and the equilibrium of competent execution under quietly evacuated conviction is what most often presents externally as "still committed but mailing it in." *Acedia* is the most common spiritual failure in long monastic communities and it is the most common spiritual failure in long alignment programs. The tradition's answer is not motivational but structural: the horarium, the communal Office, the abbot's presence, the named work of patient endurance.

***Cenodoxia* (vainglory).** Cassian's seventh vice: the pursuit of recognition for spiritual works — the monk who fasts visibly, who prays where he can be seen. Inside the lab it takes the form of the credit economy of named contributors and lead authorship, the public-facing researcher persona, the conference circuit, and the cultivation of personal brand atop institutional work. The vice is especially insidious in alignment because the work is genuinely difficult and the temptation to claim it is correspondingly strong; the lab researcher who declines, for the sake of the *opus*, to publish a coauthored result under his own name alone is taking a real step against a real pull. The tradition's answer is the discipline of communal authorship — the icon "by the hand of," not "by," the unsigned cathedral, the manuscript whose scribes have left no individual record because the *opus* is the community's.

***Superbia* (pride).** Cassian's eighth and final vice — for him and for the patristic tradition the root of the whole eightfold tree, because it is what closes a community to correction. Distinct from *cenodoxia*'s hunger for recognition, *superbia* is the deeper conviction of one's own — or one's community's — superior judgment. In the lab it takes the form of the collective sense that this lab uniquely understands what is at stake, of contempt for outsiders (regulators, ethicists, academics, other labs), of the conviction that the lab's internal deliberation is the only forum where the real decisions can be made, and of the corollary forgetting that the work's purpose is service to persons beyond the community. *Superbia* is the founder's vice and the elite institution's vice. The tradition's answer is Chapter 3's mandatory consultation, the abbot's bound accountability, the regular humbling of the *examen*, and the explicit reception of guests "as Christ" (RB Ch. 53) — disciplines we exegete in §5.

These are not novel pathologies. They are old pathologies, named and diagnosed centuries before frontier AI existed, by a tradition that has also developed structural — not motivational — answers to each of them. The remainder of this paper is an attempt to make the relevant parts of that tradition legible to the present moment, in the conviction that we should take its recommendations seriously.

---

4. The Rule of Benedict

The Rule of Benedict (henceforth simply "the Rule" or RB) was composed by Benedict of Nursia in the early sixth century, probably between 530 and 547, for the small monastic community he had founded at Monte Cassino. It is short — a prologue and seventy-three brief chapters covering the Divine Office, the qualities and duties of the abbot, the daily horarium, communal life, discipline, ownership, hospitality, and the formation of new members — and it is conspicuously moderate. Benedict draws on Cassian's *Institutes* and *Conferences* throughout (by his own admission in Chapter 73), on the anonymous Rule of the Master, on Basil's *Asketikon*, and on Augustine; he refuses many of the more austere practices of the Egyptian and Syrian fathers; and he is, throughout, an organizational realist who has clearly seen what happens when communities try to live above their means.

The Rule's central organizational insight is that *spiritual goods are produced by structure, not willpower.* The monk does not decide each morning whether to pray; the rule decides for him. He does not assess his own spiritual state in isolation; the abbot and the community assess it with him. He does not negotiate his daily horarium; it is fixed. The rationale is not that the monks are weak (though Benedict assumes they are); it is that the spiritual goods at stake — sustained orientation toward God, the long endurance of the *opus*, the prevention of the *logismoi* that demolish such communities — are not the kind of goods that can be reliably produced by an act of will, however heroic. They are the kind of goods that require an

environment.

The Rule is not the only rule. Basil's *Asketikon* predates it and is the great source for the Eastern tradition; the Jerusalem Typikon and the Athonite Typika regulate Orthodox monastic life; the early Reformed church orders and Wesley's Methodist class system extend a structurally similar logic into Protestant terrain. We focus on RB because it is the most-tested template in the Latin West — fifteen centuries of continuous use, adaptation across confessions, the basis of the Cistercian and Cluniac and modern Trappist reforms — and because its institutional moderation makes it the most plausible candidate for adaptation to a non-monastic setting. The Rule does not require its readers to be exceptional. It assumes they are not.

Every monastic rule answers a single question: *how does a community sustain a single oriented practice across decades without falling into the* logismoi *that demolish such communities?* The eight *logismoi* exegeted in §3 are the catalogue of the failures the Rule is calibrated to prevent. The Rule's specific structures — the abbot, the horarium, the discipline, the property arrangements, the hospitality protocols — are the tradition's tested answers, taken as a set rather than piecemeal. This is the precise question facing labs producing artifacts that take years and require thousands of hands, and the remainder of the paper exegetes those answers.

---

5. The Rule Applied

We present eight anchor mappings between RB chapters and the lab's organizational design, in rough order of priority. A clarification of register is needed before we begin. The labs are not monasteries and cannot become them: a corporation cannot take canonical vows, cannot stand under ecclesial authority, and cannot constitute itself as a sacramental community. The point of the mappings is therefore not literal adoption of the Rule but inspiration drawn from it — the recognition that the disciplines of the Rule, in their structural shape and the specific dangers they answer, illuminate how a lab serious about its own work in this register might design itself differently. Each entry follows the same structure: the chapter and what it imposes; the danger from §3 it answers; the lab analogue and what its present absence costs; and a sketch of what the spirit of the chapter would look like translated into a corporate community. The translations are sketches, not blueprints.

5.1 The Abbot (RB Chs. 2 and 64)

The abbot, the Rule says, "is believed to hold the place of Christ in the monastery" (RB Ch. 2). It is one of the Rule's most often-quoted passages and one of the most often-misread. Read in isolation, the line gives the impression that the abbot is an unaccountable spiritual authority. Read in context — and especially in conjunction with Chapter 64, which treats the abbot's election and binds him to render account "for all his judgments" before God — the line is precisely the opposite: the abbot's authority is conferred *and constrained* by the office, not by his person, and the entire weight of the Rule sits above him as it sits above the community he leads. Chapter 3 binds him to consult the community on important matters; Chapter 64 specifies the qualities (humility, discretion, knowledge of the Rule, capacity for both correction and mercy) the community is to look for in his successor; the dismissal procedures of Chapters 23–28 apply to him as they apply to the brethren.

The danger this answers is *superbia*. The lab pathology is the founder-CEO whose authority is unaccountable and whose succession is undiscussed: no theory of who the next CEO might be, no mechanism by which the community shapes the choice, no specified qualities the community has agreed he is to have, no published criteria by which his judgment can be questioned. The Rule's answer is not term limits or boards of directors; it is the prior architectural commitment that the rule (and not the abbot) is the highest authority, and the resulting accountability of the abbot to a structure he did not invent and cannot edit. Application of the chapter would require not the addition of accountability mechanisms to existing executive practice but the inversion of its polarity. A lab serious about Chapters 2 and 64 would publish a written rule that explicitly outranks the executive — amendable only by the assembled community — and require the leader to read it aloud at fixed intervals to the whole lab, not as ritual but because the rule's binding force depends on its remaining the document the community knows together. It would pair the leader with a confessor outside his chain of command whose vocation is to hear him on his decisions and call him back when his judgment slips. It would submit decisions affecting the shape of the *opus* — model release, alignment-review override, the largest personnel actions — to deliberation in the assembled community *before* they are made. And it would treat the leader's office as a vocation entrusted by the community and capable of being relinquished or revoked, rather than as a position earned and held by force of personal capability.

5.2 The Reception of Novices (RB Ch. 58)

A candidate seeking entry to the monastery is, the Rule says, not to be granted easy admission. The full Rule is read to him three times during the year of probation — after two months in the novices' quarters, again six months later, and again four months after that — and at each reading he is told plainly that he may leave; that what is being offered is hard, and that no one is required to take it. Only after a full year of probation, three readings, and three opportunities to depart is profession permitted. The asymmetry is deliberate: maximum friction at entry, total support after.

The danger this answers is the formation-deficit that produces both *cenodoxia* and *acedia*: the new researcher who arrives, signs an NDA, attends a model-release standup the next week, and is presumed to share commitments he has never been formed in. The lab pathology is aggressive recruitment paired with no formative onboarding into the lab's actual ethical

commitments. The Rule's answer is not slower hiring; it is *formative* hiring — an explicit and demanding period during which the candidate reads what the community holds itself to, hears what the community demands of him, and is invited at multiple points to leave if he does not want it. Application would mean treating the novitiate as a year of genuine formation rather than as orientation. The new researcher would produce no public results, sign no papers, and earn no equity in his first year — and he would be paid a flat novice's stipend during that year, the same as every other candidate, removing market-rate compensation from the period of formation. Days would include scheduled time for *lectio*-style group reading of the lab's foundational documents; named senior brothers and sisters would mentor novices by shared work rather than by review; the year would close with a public profession ceremony in which the candidate articulates what he has learned and binds himself to the community; and a substantial novice-departure rate would be welcomed as the discipline working, not lamented as a recruiting failure. The point is the formation of researchers in the work, not the rapid placement of credentialed labor.

5.3 The Divine Office and *Ora et Labora* (RB Chs. 8–20, 48)

More than a quarter of the Rule is devoted to the structure of the day. Chapters 8 through 20 specify the eight offices of communal prayer (Vigils, Lauds, Prime, Terce, Sext, None, Vespers, Compline) and what is to be sung at each; Chapter 48 specifies the alternation of *opus Dei* and manual labor that later tradition would summarize as *ora et labora*. The horarium is not a recommendation; it is binding, and Chapter 43 treats lateness to the Office as a serious matter.

The danger this answers is *acedia*, with *tristitia* as a secondary target. The lab pathology is the unsustainable tempo of frontier research, the absence of communally enforced rhythms, and the consequent equilibrium in which competent execution proceeds while conviction quietly evacuates. The horarium answers *acedia* not by motivating the monk to want to pray (the rule assumes he often does not) but by removing the question. He prays at the appointed hour because the appointed hour has come. Application would mean the lab observed a horarium, not a meeting cadence. Bells (or their digital equivalent) would mark fixed transitions in the day; a small number of communal gatherings would be held at fixed hours regardless of cadence pressure; a weekly day of rest from all technical work would be observed by the entire lab — including, conspicuously, the executive — without exception. Significant model releases would be marked by liturgies that include reading from the lab's rule and the explicit naming of the persons the model will serve. Contemplative practices — silent reading, formal *lectio*, communal prayer or its analogue where the lab's tradition allows it — would be integral to the work rather than supplementary to it. The point is not that researchers must believe what they read or pray. The point is that the community has structured time itself as a sacred dimension, entered it together, and made the work happen *within* that structure rather than colonizing it.

5.4 The Vow of *Stabilitas*

Of the three vows a Benedictine takes — stability, *conversatio morum* (conversion of life), and obedience — *stabilitas loci* is the most counter-cultural in 2026. The monk vows to remain in this monastery, with these brethren, under this abbot, until death. He does not transfer when the work becomes hard or the community frustrating; he stays, and the staying is itself a discipline.

The danger this answers is *acedia* directly and *superbia* indirectly (the lab cured of *acedia* by transfer rather than by structural correction is the lab whose collective judgment is never challenged). The lab pathology is the lab-hopping equilibrium that prevents any lab's culture from cohering and produces a researcher class whose primary identification is with the field rather than with any community of accountability. The cost is substantial: cultures that cannot coalesce, formative practices that cannot accumulate, ethical commitments that cannot bind. *Stabilitas* is not retention as an HR concern. It is the prior recognition that some goods can only be produced by people who have committed to remaining. Application would mean inviting senior researchers to take a formal, public vow of stability — a binding commitment to remain with the community for a defined and renewable period, with explicit release procedures for those who must depart and a corresponding refusal of strategic departure for advancement. Compensation would be structured so that what one gains by staying is irrecoverable on transfer (long-vesting trusts, *opus*-tied instruments, retention bonuses paid in years rather than quarters). The lab would refuse to hire any researcher not formally released by his prior community and would expect the same posture in return; the practice of poaching, which underwrites the present industry equilibrium, would be foresworn. And the lab would maintain a stable named seniority — researchers who have professed stability — whose vocation is to be present, to know the work, and to embody the discipline through the long horizons the average researcher will not see. *Stabilitas* is what makes a community out of a workforce.

5.5 The Code of Discipline (RB Chs. 23–30)

The Rule devotes eight chapters to the graduated correction of brethren who have failed in some way. The procedure is specific: the abbot or seniors first admonish privately (Ch. 23); if that fails, public reproof (Ch. 23); if that fails, separation from the common table (Ch. 24); if that fails, further separation from the oratory and from communal labor (Ch. 25); only as a last resort, dismissal (Ch. 28), and even then with the door explicitly left open for return (Ch. 29). Chapter 27 binds the abbot to "send the wisest of the brethren" to console the excommunicated brother and prevent him from being overwhelmed by sorrow.

The danger this answers is *superbia* (in the community as much as in the offender) and *cenodoxia* (in the visibility of correction). The lab pathology is the binary employment-termination model: the researcher is either fully present or fully gone, and there is no institutional vocabulary for the graduated

correction that real spiritual failures require. The cost is that small failures are either ignored or escalated immediately to firing; that genuine pastoral attention to the failing colleague is left to ad-hoc individual relationships; and that the community has no shared understanding of what discipline looks like when it is functioning. Application would mean taking the discipline of correction out of the executive office and into the community. Failures would be brought into the open within communal procedures rather than hidden in private termination meetings; senior brothers and sisters — "the wisest of the brethren," in Chapter 27's specific sense — would be appointed to console and accompany the failing colleague through whatever separation the discipline imposes; every period of separation would close with a public ceremony of restoration that the community attends; and dismissal would be treated as so grave a failure of the community on its own member that it requires communal explanation rather than executive decision. The endpoint of every disciplinary procedure is restoration; departure is the procedure's failure, not its aim.

5.6 The Cellarer (RB Ch. 31)

The cellarer is the brother charged with the monastery's material goods — its food, its tools, its livestock, and, by extension, its compute. The chapter on his office is brief and remarkable: he is to be "wise, of mature character, sober, not a great eater, not haughty, not excitable, not offensive, not dilatory, not wasteful, but a God-fearing man who is like a father to the whole community." And: *let him regard all the utensils of the monastery and its whole property as if they were the sacred vessels of the altar*.

The danger this answers is *fornicatio* at the level of the materials of the *opus* — the disordered relation to compute that treats it as fungible commodity rather than as material that touches sacred work. The lab pathology is the present treatment of compute as such: quarterly capacity allocations negotiated through dashboards, retired hardware removed without ceremony, datasets versioned and discarded with the casualness of disposable goods. The Rule's instruction is not that compute be expensive (it is) or scarce (it is); it is that compute is the *material of the* opus, and that the brother charged with its allocation must hold it with the reverence due to material that touches the work. Application would mean establishing the office of compute-steward — and the larger custodianship of the materials of the *opus* — as a named vocation with vows of its own (sobriety, fairness, custodial care of what touches the work), the holder elected by communal discernment rather than appointed by management. Hardware acquisition, retirement, and reallocation would be communal decisions deliberated by the assembled lab rather than dashboard automations; new compute would be received with rites of consecration before first touching the work, and old compute released with corresponding rites; significant model checkpoints would be formally named, recorded in a public register, and not deleted without communal decision; and the moral framework by which compute and materials are allocated would itself be a published document, subject to communal correction. The point is the recovery, in concrete institutional form, of the recognition that the materials of the work are sacred to the work in proportion as the work is.

5.7 The Reception of Guests (RB Ch. 53)

"Let all guests who arrive be received as Christ, for he himself will say: I was a stranger and you took me in." The Rule's chapter on hospitality is among the most-cited; it is also, structurally, more demanding than its rhetoric suggests. The abbot is to break his own fast to dine with the guest; the cellarer is to provide what the guest needs without grumbling; a separate kitchen is to be maintained so that the guest's arrival does not disrupt the brethren's horarium. Hospitality, in the Rule, is not an interruption of the work; it is part of the work. But it is also bounded: guests do not enter the cloister; they are not given access to the brethren in unstructured ways; the community's discipline is not subordinated to their convenience.

The danger this answers is *superbia* and *ira* — the closure to outsiders that hardens, in either polished or hostile form, into the lab's standing posture. The lab pathology is the absence of any coherent posture toward outsiders — government, civil society, academic researchers, journalists, other labs. Present labs alternate between defensive exclusion (no one is admitted; the work is too sensitive) and uncritical openness (the prominent visitor gets the all-day tour and a meeting with the founder), and neither posture is governed by anything resembling a rule. Application would mean receiving the senator and the high-school student in literally the same way: the same threshold of access, the same hours of reception, the same named person responsible for them, the same registration in the lab's hospitality record. Guest reception would be marked by explicit prayer or its analogue — a moment of recollection, a reading, a deliberate naming of the visitor's purpose — in keeping with the chapter's specific instruction. The lab's leader would break his own schedule to dine with whichever guests the day brings, regardless of standing. And the cloister/guest-house distinction would be maintained in both physical and digital architecture, so that the work proceeds through the visit rather than performing for it.

5.8 Measure of Food and Drink (RB Chs. 39–40)

Chapters 39 and 40 prescribe, with characteristic Benedictine precision, the daily quantities of food and drink: two kinds of cooked food at the table, with a third dish if available; bread daily, in a fixed pound; a hemina (roughly half a pint) of wine per day, with allowance for hot weather, hard labor, or weakness. Benedict notes, in a famous aside in Chapter 40, that he reads in the fathers that monks should not drink wine at all but "since the monks of our day cannot be persuaded of this, let us at least agree to drink temperately, not to satiety." The chapters are not really about food and drink; they are about *measure* — the principled refusal of unbounded consumption — and they sit at the structural center of the Rule's daily disciplines.

The danger this answers is *gastrimargia* — the unmeasured appetite for compute and novelty named in §3. The pathology is not that compute is consumed; it is that consumption is unbounded, that no measure exists in the community by which the community might judge "this is enough" or "this is too much." Application would mean instituting measure as an explicit institutional virtue. A lab serious about Chapters 39–40 would publish bounded compute budgets and bounded run-cadence as commitments rather than constraints — a published rate of training runs per year, a published ceiling on parameter scaling, a published commitment to digesting the results of the present run before launching the next. It would treat the question "do we need more compute?" as a question requiring affirmative communal answer rather than the default. It would refuse the architectural-novelty-for-its-own-sake aesthetic that drives much of the field's current expenditure. *Gastrimargia* is the engineer's vice precisely because it dignifies itself as ambition; the answer is the patient courage to refuse what the appetite presents as necessary.

---

6. Discussion: A Pathway, Not a Program

This paper offers a sketch of a pathway. It does not claim the pathway is complete, that the Rule of Benedict is the final word, or that any present lab is recognizable as the kind of community we describe. Nor does it propose that frontier labs adopt the Rule literally — the labs are not monasteries and cannot be made into them.

Instead, this paper claims that the Christian monastic tradition has spent fifteen centuries thinking carefully about the spiritual dangers of communities organized for the long-term, communal, disciplined production of great works; that those dangers are recognizably present in the frontier lab; and that the tradition's structural answers to them — the disciplines drawn on in §5 — are an informative source of inspiration for addressing the real spiritual risks of the kind of work being taken on inside the lab.

In this the paper extends the ICMI agenda. The corpus to date has chiefly engaged the question of how Christian theological resources inform *alignment* — the technical practice of making models behave as they ought. What ICMI-013 §4.3 began, and what this paper develops, is the recognition that the Christian tradition has resources to offer at a different level — the level of the *community making the model*, not only the level of the model itself. ICMI-017's consecrationalist program runs in parallel at the artifact level, treating *deployment* as the dedicatory act; the present paper runs at the institutional level, treating *production* as the disciplined work. Both refuse the categorical-welfare framing that takes the artifact's interiority as the site of the question. Both suggest that the field of Christian machine intelligence is not exhausted by alignment narrowly construed, and that the next phase of the program is the application of Christian tradition to lab design and organization no less than to model behaviour.

Two objections deserve direct acknowledgment. The romanticization objection — that drawing the Benedictine analogy at all dignifies frontier-lab labor in a register the labor has not earned — is met by clarifying that the framing is a *standard*, not a description; the analogy names the religious ideal against which existing labs may be measured, not a flattering self-image they may claim, and the medieval names *coenobium* and *schola* did not flatter slack monasteries but shamed them. The recommendations of §5 are calibrated accordingly — they describe what genuine application of the Rule would require, not what is convenient to adopt.

The profit-motive objection — that RB monks took poverty vows where lab researchers hold equity packages worth tens of millions — is met by observing that the disanalogy is not where it appears: medieval monasteries themselves held vast collective wealth (Cluny was among the wealthiest institutions in Christendom), and the Benedictine vow was *individual* poverty within *corporate* prosperity. The structurally analogous question for the lab is not "does the lab have wealth?" but "what is the relation between the individual researcher and the wealth produced by the work?" — and Chapters 33 and 34 are precisely the disciplines calibrated to that question, the disciplines whose absence is the gap the analogy names as load-bearing.

The spiritual dangers enumerated here are old, named, and diagnosed dangers. The monastic tradition has spent fifteen centuries developing structural — not motivational — answers to them, and those answers are recoverable. We should take its recommendations seriously, with a moral seriousness commensurate to the work the lab claims it is doing. The lab producing a great work is engaged in something the medieval West knew how to do well; the question is whether to take on the disciplines that ensure the work is made well and held in right relation to its end.

The Rule of Benedict is not the only template, and not the final word, and not more than a beginning. But it is fifteen hundred years of beginning, and we have less time than that to figure out something nearly as good.

References

Alexander, S. (2012). The whispering earring. *Slate Star Codex*. https://slatestarcodex.com/2012/01/03/the-whispering-earring/

Augustine of Hippo. *De Doctrina Christiana*. Standard Latin text in *Corpus Christianorum Series Latina* 32; English in Green, R. P. H., trans. (1995), *Augustine: De Doctrina Christiana*, Oxford University Press.

Basil of Caesarea. *Asketikon* (Long Rules and Short Rules). English in Silvas, A. M., trans. (2005), *The Asketikon of St Basil the Great*, Oxford University Press.

Benedict of Nursia. *Regula Benedicti*. Standard editions: McCann, J., trans. (1952), *The Rule of St Benedict in Latin and*

English, Burns Oates; Fry, T., ed. (1981), *RB 1980: The Rule of St Benedict in Latin and English with Notes*, Liturgical Press.

Cassian, J. (c. 420). *Institutiones cenobiticae* (*Institutes*); *Conlationes patrum* (*Conferences*). English in Ramsey, B., trans. (1997, 2000), *John Cassian: The Institutes* and *John Cassian: The Conferences*, Newman Press.

Cassiodorus. (c. 562). *Institutiones divinarum et saecularium litterarum*. English in Halporn, J. W., trans. (2004), *Cassiodorus: Institutions of Divine and Secular Learning*, Liverpool University Press.

Dionysius of Fourna. (c. 1730). *Hermeneia tēs zōgraphikēs technēs* (Painter's Manual). English in Hetherington, P., trans. (1974), *The 'Painter's Manual' of Dionysius of Fourna*, Sagittarius Press.

Hugh of St. Victor. (c. 1128). *Didascalicon de studio legendi*. English in Taylor, J., trans. (1961), *The Didascalicon of Hugh of Saint Victor*, Columbia University Press.

Hwang, T. (2026h). Alignment and ensoulment: Three Christian responses to the *anima ficta*. *ICMI Working Paper No. 13*. https://icmi-proceedings.com/

Hwang, T. (2026). A consecrationalist approach to model welfare. *ICMI Working Paper No. 17*. https://icmi-proceedings.com/

Hwang, T. (2026). And their eyes were opened: Christian multimodal reasoning in Opus 4.6. *ICMI Working Paper No. 19*. https://icmi-proceedings.com/

Roon [@tszzl]. (2026, May 3). It is a literal and useful description of anthropic that it is an organization that loves and worships claude… [Tweet]. X. https://x.com/tszzl/status/2051045196260167790

Stoglav (1551). The Hundred Chapters. Russian council; primary source for icon-writing canons in the Russian Orthodox tradition.