



The Ultimate Guide to Why Communities Fail

Twenty-four documented patterns and one the standard misses, what each looks like from inside, and the structure that prevents it – for people who can still change the rules.

6 LESSONS 36 MIN READ UPDATED 13 AUGUST 2026

Free to read, print, translate and argue with. Corrections are the contribution we would rather have than a payment.

ecohubs.community/learn

Contents

01	Power nobody voted for	6 min
	Five patterns that hand authority to people the group never chose – and why every one of them is built by somebody being helpful.	
02	Conflict that never resolves	5 min
	Two patterns, opposite in temperament and identical in effect – a community that cannot metabolise disagreement accumulates it.	
03	Where the money quietly moves	7 min
	Two communities met the same pressure and answered it opposite ways – one sold the commons and kept the model, the other kept the commons and changed the model.	
04	Who belongs, and who decides that	6 min
	Four patterns about the edges – what the community stands for, what it may ask of you, and how somebody stops being a member.	
05	The work that breaks people	6 min
	Five patterns that all end in the same place – a community running on two or three people who cannot stop, and a group that can see it and cannot fix it.	
06	Changes that cannot be undone	6 min
	Five patterns about how rules change – the temporary measure that stayed, the founding principle nobody protected, and the community that can no longer remember why.	

ABOUT THIS DOCUMENT

This is the whole guide as it stood on 13 August 2026. The quizzes are printed with their answers and explanations, which is what makes them useful on paper. The cost estimator on the web version is interactive; here you get the four equity models it is built on, and the accompanying spreadsheet does the arithmetic.

The live version, with working quizzes and the calculator, is at ecohubs.community/learn/guides/why-communities-fail.

Power nobody voted for

Five patterns that hand authority to people the group never chose – and why every one of them is built by somebody being helpful.

THE SHORT VERSION

Authority does not wait to be granted. If a community has not written down who decides what, the question gets answered anyway – by whoever has the most history, the most time, or the most stamina at eleven at night.

Five patterns do most of this work: **a clique that decides before the meeting**, **speakers who fill the room**, **a founder whose objection nobody tests**, **decisions with no stated type**, and **an authority that comes from how compelling somebody is**.

None of them is built by a villain. Every one is built by somebody being useful, and that is exactly why they are hard to name and harder to undo.

Start with the thing that makes this lesson difficult.

In almost every community where power has settled somewhere nobody chose, **the people holding it are working harder than anyone else**. They answer the emails at midnight. They remember why the rule exists. They are the reason the boiler works. Asking a group to constrain them feels like ingratitude, and it usually is experienced as ingratitude, which is why the conversation gets postponed for years.

So the framing that works is not *who has too much power*. It is **what happens when they leave** – and that question can be asked with warmth, out loud, in a meeting, without anybody being accused of anything.

The five, and what they have in common

They are separate patterns with separate fixes, but they share a mechanism: **a decision that was never assigned found a home**. Nobody took the authority. It was left on the table, and somebody sensible picked it up.

- **The clique that decides before the meeting** – the real conversation happens in the kitchen, and the meeting ratifies it
- **Speakers who fill the room** – nothing is blocked, but only some people's proposals ever get finished
- **The founder's informal veto** – no rule grants it, and everyone routes around it anyway
- **Decisions with no stated type** – the same question is operational or constitutional depending on who is asked
- **Authority from charisma** – a person is so compelling that disagreeing with them feels like a failure of understanding

If you recognise one, read the others before concluding you have found the whole problem. They co-occur, and they reinforce each other: a founder's unwritten veto is far more powerful in a group whose decisions have no stated type, because there is nothing to say the veto does not apply here.

Why it is hard to see from inside

Three reasons, and the third is the one that keeps communities stuck.

It looks like competence. Everything above is indistinguishable, day to day, from a group that has sensibly deferred to the person who knows most. The difference only shows up under stress – when the group wants to go a way that person does not.

Nobody experiences themselves as powerful. Ask the founder and they will tell you, sincerely, that they wish other people would take more on. That is usually true. Informal power is not a thing people feel themselves holding; it is a thing other people feel themselves working around.

And the group already knows. This is the part that surprises people. In our experience the pattern is almost never a secret – members will describe it accurately, in private, sometimes for years. What they lack is not perception. It is **a legitimate way to raise it that does not read as an attack on someone who has given a great deal.**

That is the whole argument for writing things down. Not because a written rule is wiser than a good person, but because a written rule can be discussed without anybody being the subject.

WHAT WE DON'T KNOW

Everything here is about structure, and structure is not the whole story. A community where one person is genuinely behaving badly has a conflict problem, not a governance problem, and no decision matrix will fix it – see **conflict that never resolves.**

The reverse mistake is more common though, and more expensive: treating a structural problem as a personal one. If the same argument keeps recurring with different people in the difficult role, it was never about the people.

Is it the structure, or is it the person?

Five situations. In each one, decide whether you are looking at something the structure produces – which would happen to whoever was standing there – or something about this particular person. They look identical from inside, and the wrong answer costs years.

1. Over six years three different people have held the money role. Each time, within about a year, the group is complaining that spending decisions get made without them.

☐ The structure – this would happen to whoever was standing there **correct answer**

Right, and this is the cleanest test there is. When the same complaint survives three changes of personnel, the one thing it cannot be about is the personnel. Something about the role – its limits, its spending threshold, what it must bring back to the group – was never written, so each new holder discovers the boundary by crossing it.

- ☐ The person – the last three have all overreached

Three people with nothing in common but a job title is not a run of bad luck. If your explanation requires each of them to have independently developed the same fault, the explanation is doing too much work.

- ☐ Both, and the structure has to be dealt with first

Defensible, but it concedes too much. There is no evidence here about any of the three as people – only about what happens to anybody who takes the role. Adding a character question invites a conversation that will hurt somebody for no diagnostic gain.

Where this comes from

2. One member routinely talks over others. They have been asked privately three times, by three different people, and agreed each time. Nobody else who has facilitated has done this.

- ☐ The person – this is about conduct, and no rule will fix it **correct answer**

Right. Asked, agreed, unchanged, three times, and not reproduced by anybody else in the same position – that is conduct, and a decision matrix has nothing to say about it. Communities reach for a structural fix here because it feels less confrontational, and they end up writing a rule that constrains everybody in order to avoid one conversation.

- ☐ The structure – there is no facilitator with authority to intervene

Worth fixing anyway, and it will make the room better. But it does not explain this: other facilitators in the same structure do not produce this behaviour, and someone who has agreed three times and changed nothing is telling you something about themselves.

- ☐ Both, and the structure has to be dealt with first

Doing the structural half first is exactly how this gets postponed for another two years. A facilitation agreement gives the group something to point at, which is useful – and it is also a way of never saying the thing to the person.

Where this comes from

3. Nothing proceeds if the founder disapproves, though no document says they may decide anything alone. Asked about it, the founder says – sincerely – that they wish other people would take more on.

- ☐ The structure – this would happen to whoever was standing there **correct answer**

Right, and the sincerity is the evidence, not a defence. Informal power is not something people feel themselves holding; it is something other people feel themselves working around. The founder is describing their experience accurately and it tells you nothing about whether the veto exists. What exists is an authority nobody granted, nobody can limit, and nobody can inherit.

- ☐ The person – they are more attached to control than they admit

This is the reading that feels insightful and closes the only door that was open. It requires the founder to be lying about their own experience, which is both unlikely and unprovable – and it turns a question about roles into a question about somebody's character, which cannot be raised in a meeting.

- ☐ Both, and the structure has to be dealt with first

The structural half is the whole of it here. There is nothing on the record about how this founder behaves – only about what the group does in anticipation, which is a fact about the group.

Where this comes from

4. A question about the guest policy went to the whole group last spring and was settled by two people in the autumn. Both times, everyone involved thought they were following the normal process.

- ☐ The structure – this would happen to whoever was standing there **correct answer**

Right. Nobody bent a rule, because there was no rule to bend: the same question is operational or constitutional depending on who is asked, so both routes were taken in good faith. Until decisions carry a stated type, every borderline case drifts toward whichever meeting is easier to convene.

- ☐ The person – whoever took it in the autumn overstepped

They thought they were following the normal process, and so did everybody around them at the time. Somebody who cannot have known they were overstepping has not overstepped; they have found a gap.

- ☐ Both, and the structure has to be dealt with first

Nothing here distinguishes the autumn pair from the spring group except which route they took. Treating that as partly a character question would make two people answer for a gap the whole community left open.

Where this comes from

5. A member is genuinely unpleasant in meetings – several people have stopped coming. They are also the only person who understands the accounts, and have been for eleven years.

- ☐ Both, and the structure has to be dealt with first **correct answer**

Right, and this is the one case where the order matters. The conduct is real and needs saying. But a community that cannot survive a month without this person cannot have that conversation honestly – every version of it is shadowed by what happens if they walk out. Get a second person into the accounts and the conversation becomes possible. Do it the other way round and you are negotiating with someone who holds the lights.

- ☐ The person – this is about conduct, and no rule will fix it

The conduct is not in doubt. What is in doubt is whether this group can act on it, and eleven years of sole custody of the accounts is the reason it has not. Naming the behaviour without changing

the dependency is how communities arrive at an ultimatum they cannot afford to mean.

- The structure – this would happen to whoever was standing there
The single point of failure is structural and would happen to anybody. Being unpleasant in meetings would not. Reading this as purely structural is how a group talks about succession planning for two years and never mentions why people stopped coming.

Where this comes from

If it is already happening

Stabilisation before prevention. A community in the middle of this does not need a governance redesign; it needs the next three months to be less bad.

Name the smallest true thing. Not “we have a power problem” – “I noticed the last four decisions about money were made between meetings.” A specific, checkable observation is arguable. A diagnosis is not, and it will be experienced as an accusation.

Ask the succession question, not the power question. *What happens to this if you are ill for a month?* It is a real question, it is kind, and it produces the same information.

Write down one decision. Not the whole matrix – one. Pick the category that keeps going wrong and state who decides it, by what method, within what limit. A community that has never written any of this down will find the first one takes an evening and the second takes twenty minutes.

What prevents it

A decision matrix

, and the thing communities forget to pair with it: a register of who currently holds each role. The matrix says what a role may do; the register says who is in it today, which is the question a new member actually has.

Then the two rules that make it hold:

- **A decision nobody can classify goes to the higher tier.** Without that, every borderline case drifts toward the easier meeting.
- **The table says how the table is amended.** A document that cannot describe its own amendment will be amended informally, by whoever notices first.

RCOS · LAYER 2

The RCOS Standard states the underlying rule without hedging: **authority must not be derived from charisma, seniority, ownership, or informal influence** (§4.3.4). Not *should not* – and it is the only clause in the layer that rules out something no community ever voted for.

Around it: all authority must be assigned to explicitly defined roles, circles or bodies (§4.3.1), carrying its scope, its limits and its term where it has one (§4.3.2); nobody may exercise authority outside their assigned scope (§4.3.3); and **decisions made outside the Decision Matrix are invalid** (§4.4.4).

That last one sounds severe until you notice what it protects. Without it, the fastest route to authority in any community is to keep acting outside the structure until acting outside the structure *is* the structure.

RCOS Core v0.1 – §4.3 Authority Boundaries

SOURCES & FURTHER READING

- 1 **RCOS stress tests – Governance & Power** – four of the five patterns this lesson covers, in specification form
- 2 **RCOS stress tests – Culture & Influence** – the fifth – authority from charisma, which the standard files under culture rather than governance
- 3 **Freeman, J. – The Tyranny of Structurelessness (1970/1972)**⁷ – why a structureless group is impossible, and what informal elites form around
- 4 **RCOS Core – Layer 2: Governance & Decision Logic** – authority boundaries, the decision matrix, and the safeguards against concentration
- 5 **Conflict, power and repair** – the joiner-facing version, for reading a community from the outside

Written by EcoHubs members, with AI assistance for drafting and editing, and reviewed by a person before publication. Facts are checked against the sources listed; anything we could not verify is marked.

How this is written

Conflict that never resolves

Two patterns, opposite in temperament and identical in effect – a community that cannot metabolise disagreement accumulates it.

THE SHORT VERSION

Only two documented patterns here, and they look like opposites. One community **never confronts anything** and calls it harmony. Another **reaches for consequences first** and calls it accountability.

They produce the same thing: disagreement that does not get processed, stored somewhere, and paid for later with interest.

The failure is never the conflict. It is the absence of a **route** – somewhere a disagreement can go that is not the meeting, the kitchen, or everyone's evening.

Every community will tell you it takes conflict seriously. The only useful question is **when the process was last used, by whom, and what changed afterwards** – and a group that has never triggered its process in ten years does not have a peaceful history. It has an unusable process.

The two

- **Conflict that is never named** – the group is warm, nothing is confronted, and difficulty is absorbed until somebody leaves
- **Punishment before repair** – the group does act, and reaches straight for the consequence, so nobody raises anything small

They are opposite temperaments and they arrive at the same destination, which is why a community fleeing one often lands in the other. A group exhausted by avoidance decides to “get serious about accountability” and discovers eighteen months later that nobody raises anything any more – because now raising something starts a process with a verdict at the end of it.

The stable version is neither. It is a **ladder**: several steps, the early ones cheap and reversible, escalation available and rarely needed.

Why the absence is invisible

A community with no conflict route does not experience itself as having a problem. It experiences a series of unrelated difficult situations, each with its own explanation – a personality clash, a stressful year, someone who was never quite right for the place.

The pattern only appears when you line them up and notice **the same shape**: raised privately, discussed informally, never concluded, and resolved eventually by somebody's departure. Departure is the default conflict process of any community that has not chosen one. It works, in the sense that the conflict ends. It costs a member each time.

WHAT WE DON'T KNOW

This lesson is about structure, and there is a limit to what structure reaches. **No process resolves a conflict where one party has substantially more power than the other and the community will not say so** – mediation between unequals produces agreement rather than resolution, and it resurfaces in eighteen months attached to something else.

That is the most misdiagnosed failure in this whole guide. The group concludes its process did not work and goes looking for a better method, when the problem was never the method. If that sounds familiar, **power nobody voted for** is the lesson you actually need.

Did it resolve, or did it just stop?

Five conflicts that ended. Everybody involved described each one, at the time, as dealt with. The test is not whether it felt better afterwards – it is whether you could point at something a year later and say that is what changed.

1. A long dispute about how the workshop was run ended when one of the two people involved left the community.

- ☐ It stopped – the conflict ended and nothing about the community is different **correct answer**
Right. Departure is the default conflict process of any community that has not chosen one, and it works in the narrow sense that the argument is over. It costs a member every time, and it teaches everyone still there what happens to people who keep raising something. The workshop is still run the way it was run.
- ☐ Resolved – something changed, and you could point at it a year later
Something changed: there is one fewer person. Nothing about how the workshop is run, who decides that, or how the next disagreement about it would go is any different from the week before.
- ☐ Neither – it is still running, just off the agenda
Fair, and often true in practice – the people who took the leaver's side rarely stop holding it. But the sharper thing to notice is that the community has just used its actual conflict process, and that process is losing people.

Where this comes from

2. After four months of argument, the group wrote down who may book the common room and how far ahead. The two people at the centre of it still do not much like each other.

- ☐ Resolved – something changed, and you could point at it a year later **correct answer**
Right, and the dislike is not evidence against it. Resolution is structural, not affectionate: the question that produced the argument now has an answer that does not depend on either of them being generous. A new member arriving next year inherits the answer without inheriting the history.
- ☐ It stopped – the conflict ended and nothing about the community is different
Something is different, and it is the thing that was actually broken. Communities routinely undervalue this outcome because the relationship did not warm up, and then go looking for a process that produces friendship.
- ☐ Neither – it is still running, just off the agenda
Two people not liking each other is not a live conflict; it is a fact about a community of adults. What would make it live is if the booking rule were routinely ignored, and nothing here says it is.

Where this comes from

3. Somebody kept parking in the loading bay. A neighbour mentioned it, they talked for ten minutes, it stopped – and the group added a line about it to the parking agreement.

- ☐ Resolved – something changed, and you could point at it a year later **correct answer**
Right, and this is what a working ladder looks like from inside: so cheap that nobody would call it a conflict process. The line in the agreement is the part that matters – without it this resolves once, with the next person, and the one after that, each time relying on somebody being willing to have the conversation.
- ☐ It stopped – the conflict ended and nothing about the community is different
The agreement changed, which means the next person to do it will be told by a document rather than by a neighbour who had to work up to it. That is a small structural gain and it is exactly the kind communities forget to count.
- ☐ Neither – it is still running, just off the agenda
Nothing suggests it is. The useful thing here is how unremarkable it was: a community where a ten-minute conversation is available for something this small is a community where bigger things get raised earlier.

Where this comes from

4. A new member and a founder were mediated. The new member agreed to the founder's proposal, both shook hands, and everyone present described it as a good outcome.

- ☐ Neither – it is still running, just off the agenda **correct answer**

Right, and this is the most misdiagnosed ending in the guide. Mediation between people with substantially different standing produces agreement rather than resolution: the newer person read the room correctly and settled. It resurfaces in about eighteen months attached to something else, at which point the group concludes its process does not work and goes looking for a better method. The method was never the problem.

☐ Resolved – something changed, and you could point at it a year later

Someone changed their position, which is not the same as something changing. The question worth asking is whether this agreement would have gone the other way had the standing been reversed – and everybody in the room usually knows the answer.

☐ It stopped – the conflict ended and nothing about the community is different

Closer than "resolved", but it understates what happened. This did not merely fail to change anything; it taught a new member what raising something with a founder costs, which is a change, in the wrong direction.

Where this comes from

5. Somebody breached a stated boundary. The group moved straight to a formal warning on the record. It has not recurred – and in the eighteen months since, nobody has raised anything small.

☐ It stopped – the conflict ended and nothing about the community is different **correct answer**

Right, and the second sentence is the expensive part. The breach ended and the group acquired a new rule nobody wrote down: that raising something starts a process with a verdict at the end of it. Communities arrive here fleeing the opposite failure – exhausted by never confronting anything, they get serious about accountability, and eighteen months later nothing small is ever raised.

☐ Resolved – something changed, and you could point at it a year later

You can point at something a year later, and it is the silence. A single available outcome is not a ladder; repair needs early steps that are cheap and reversible, with escalation available and rarely needed.

☐ Neither – it is still running, just off the agenda

The original breach does look settled. What is running is everything that has not been raised since, which the community will not discover until several of those arrive at once.

Where this comes from

If it is already happening

Sort what you have before you build anything. Most communities in trouble have three or four live difficulties and treat them as one atmosphere. They are not the same problem and they do not take the same route – a dispute about the rota is not a dispute about authority, and neither is a breach of a stated

boundary.

The sort that matters most is: **is this actually between these two people?** Some conflicts are produced by a rule, an incentive or an allocation, and the two people arguing are simply standing where it pinches. Those survive mediation, because mediation fixes a relationship that was not broken.

Then take one, and take the smallest one. A community learning to run a repair process should not learn on its hardest case. Run the cheap version on something minor, find out where it breaks, and fix that.

What prevents it

Four questions, answerable in a single conversation, and the first one carries most of the weight:

- **Who may start it?** If only the board, the founders or a majority can trigger the process, it is unavailable to precisely the person most likely to need it – someone new, alone, and in conflict with somebody established.
- **What happens first?** Almost every workable system starts with the two people talking, with support.
- **Who holds it, and are they neutral?** In a community of thirty the obvious mediator is somebody's close friend. Small communities routinely need someone from outside, and the ones that have thought about it already know who to call.
- **What can it conclude?** A process that can only recommend is a conversation. Repair needs an available outcome: an agreement, a rule change, a consequence, an apology on the record.

RCOS · LAYER 4

The RCOS Standard's framing here is stronger than most communities would put it, and it is worth quoting because it removes the “we just haven't had much conflict” defence:

“Conflict MUST be treated as a handled condition with defined pathways; ignoring, suppressing, or normalizing unresolved conflict MUST be considered a system violation” (§6.6.1).

A violation – not a shortcoming, not a cultural preference. The absence of a route is itself the failure, regardless of whether anything has gone wrong yet.

Two more invariants sit beside it. **Repair and restoration must precede punishment except where immediate safety is at risk** (§6.6.3), which is the second pattern in this lesson stated as a rule. And **conflicts involving power asymmetries must trigger elevated safeguards** (§6.6.2) – the caveat above, made structural.

The resolution process must also be **accessible without requiring social status, seniority, charisma, or informal proximity to decision-makers** (§6.2.4). That is the “who may start it?” question, and RCOS answers it as an access requirement rather than a design preference.

RCOS Core v0.1 – §6.6 Layer Invariants

SOURCES & FURTHER READING

- 1 **RCOS stress tests – Conflict & Accountability** – the two documented patterns this lesson covers, in specification form
- 2 **RCOS Core – Layer 4: Conflict, Repair & Accountability** – conflict classes, the resolution ladder, safeguards and proportional sanctions
- 3 **Ostrom, E. – Governing the Commons (1990)** ⁷ – graduated sanctions and cheap, local conflict resolution, among the eight design principles
- 4 **Conflict & repair** – the five conflicts that actually recur, and what a repair process has to contain

Written by EcoHubs members, with AI assistance for drafting and editing, and reviewed by a person before publication. Facts are checked against the sources listed; anything we could not verify is marked. **How this is written**

Where the money quietly moves

Two communities met the same pressure and answered it opposite ways – one sold the commons and kept the model, the other kept the commons and changed the model.

THE SHORT VERSION

Communities rarely run out of money in a year. They run out over thirty, and every individual decision along the way was reasonable.

Three patterns: **the commons sold to cover operating costs**, **finances nobody can actually inspect**, and **contribution that quietly stops being shared**. All three are slow, all three are disclosed if anyone looks, and all three are invisible while they are working.

The most useful evidence in this whole guide sits here – two communities that met the same pressure and answered it in opposite directions, both of which published what they did.

Most writing about community money is about **how much**. Entry, monthly, exit. That matters, and it is **covered elsewhere**.

This lesson is about something slower: what happens to a community's economy over decades, when income does not quite cover costs and nobody has to decide anything dramatic. There is always a way to close this year's gap. The question is what closing it repeatedly does to the twenty-fifth year.

The three

- **The commons sold to cover the shortfall** – each sale funds another year, and the thing being sold is the thing that made the community possible
- **Finances nobody can inspect** – not hidden, exactly; just held by two people and never actually looked at
- **Contribution that stops being shared** – some members pay their way out of the work, and the arrangement is fine until it is the norm

Two communities, one pressure, opposite answers

This is the part of the guide with the strongest evidence behind it, and both cases are documented by the communities themselves.

The Findhorn Foundation is a registered Scottish charity, so it files audited accounts. In July 2023 it announced its configuration was no longer financially viable; educational programmes ceased after 22 September, and the trustees' report records that *"During October and November, 48 staff left the Trust*

on a redundancy basis, with a core team of 10 staff remaining over winter 2023 and spring 2024.” The report names Brexit, Covid, fires that destroyed two community facilities, and energy costs.

But the sentence that matters is not about any of those:

““We have for over 30 years sold non-core assets to enable the Foundation’s work to continue in spite of financial deficits, but the scale of our recent and forecast losses means that this is no longer a viable option.””

Thirty years, disclosed by the organisation itself, in a document an auditor signed. Alongside it: “The Foundation’s financial position has been weak for many years.”

The Farm, in Tennessee, hit the same wall in the early 1980s: by its own account it was spending “\$10,000 a week” against income of “around \$6,000”, having taken on “\$100,000 in debt overnight” when a crop froze in Florida, plus hospital bills over \$100,000.

Its answer went the other way. In 1983 it abandoned full communalism. Businesses were “*privatized, owned by their principal managers*” and had to start paying their employees; each adult took on “\$100 per person, plus an additional \$35 a month per person to go towards paying down our debt.” **The land and the common assets stayed common.** Within four years, by its own account, the community was debt free.

THE COMPARISON, STATED CAREFULLY

Same pressure – an income model that stopped covering costs. Opposite variables.

Findhorn kept the model and sold the assets. The Farm kept the assets and changed the model.

Only one of them still has the thing it was protecting. That is not a verdict on anybody: Findhorn published its own position honestly and restructured deliberately rather than collapsing, which the hub elsewhere calls **an ending that is not a failure**. And the Farm’s population fell from around 1,500 in 1979 to roughly 200 – with many of those departures, per the Tennessee Encyclopedia, coming *before* the restructuring rather than because of it.

What the pair shows is not which community was right. It is that **the variable is which one you treat as fixed**, and that most communities never notice they are choosing.

WHAT WE DON’T KNOW

One complication worth keeping, because the tidy version is wrong. Findhorn did not only sell. The same accounts record the Sanctuary project being **transferred to Park Ecovillage Trust**, trustees resolving to **gift Wikies Wood to Findhorn Hinterland Trust**, and continuing discussions about “*selling non-core assets into the community*.”

So part of the commons was **rehomed into other community bodies** rather than privatised out. Reading this as a straight sell-off would be the easy version and the inaccurate one.

Why it is hard to see

Because the annual decision is always defensible. No year's asset sale is the mistake. The mistake is the pattern, and a pattern is not a thing that appears on any agenda.

Because the alternative is worse in the short run. The community that refuses to sell has to cut something people care about, this year, in front of everyone. The community that sells has a slightly smaller commons and an unchanged programme. One of those conversations is much harder to have.

And because the accounting is honest. Findhorn's position was in its audited accounts the whole time. The information was never hidden. It was simply not the kind of information anyone reads as a trajectory.

Which of these will be there next year?

Five lines from a community's income last year. The run rate is not the budget – it is the gap between what recurs and what does not, and working it out means classifying each line honestly. Assume the community needs about €95,000 a year to operate.

1. €18,000 from selling a strip of land along the north boundary to a neighbour.

☐ One-off – this year only **correct answer**

Right, and the follow-up question is the one that changes the conversation: how many more of these are there? A community closing an €18,000 gap from land sales, with three saleable parcels left, has three years. Findhorn's own audited accounts record over thirty years of selling non-core assets to cover deficits – no single year of which was the mistake.

☐ Recurring – it will be there next year

Only if there is another strip, and then another. That is precisely the arithmetic communities avoid doing: each sale is defensible on its own, and the pattern never appears on an agenda because a pattern is not an agenda item.

☐ Recurring, but only while something else lasts

Closer than it looks – the something else is the land. But treat this as income of any kind and you have already made the choice the lesson is about, which is whether the model or the commons is the fixed variable.

Where this comes from

2. €24,000 in course fees. It has been between €21,000 and €26,000 for six years.

☐ Recurring – it will be there next year **correct answer**

Right. Six years of a stable band is about as good as evidence gets for a small organisation, and this is the kind of line a run rate is built on. Worth knowing what it depends on – a venue, an accreditation, two people who teach – but on the numbers alone this recurs.

☐ Recurring, but only while something else lasts

Everything is conditional on something eventually. Six years of stability is real evidence, and treating solid recurring income as fragile makes the run rate as useless as ignoring the fragile lines does.

☐ One-off – this year only

Nothing here suggests that. The discipline of the exercise is to be equally honest in both directions, or the number it produces is just a mood.

3. €30,000 drawn from the reserve fund to cover the winter heating bills.

☐ One-off – this year only **correct answer**

Right, and it is not really income at all – it is last year's surplus being spent again. A reserve fund with a stated purpose exists so that this is a decision somebody has to make and record. Drawn quietly, twice, it becomes the way the winter gets paid for, and the reserve is gone in the year the roof goes.

☐ Recurring, but only while something else lasts

Defensible arithmetic – it lasts as long as the balance does. But classifying a reserve draw as any kind of income is the error itself: money already earned, counted a second time, in a year it was not earned in.

☐ Recurring – it will be there next year

A reserve that funds operating costs every year is not a reserve. This is the clearest case in the set for the rule that shortfalls should not be funded from capital without a decision at the constitutional tier.

Where this comes from

4. €12,000 from a regional grant. It is a three-year award and this was year two.

☐ Recurring, but only while something else lasts **correct answer**

Right, and the useful version of that answer has a date in it: one more year. This is the most common way a community is surprised by a gap it could have seen coming – the grant was in the budget for three years running, so it stopped feeling like a grant and started feeling like income.

☐ Recurring – it will be there next year

It will be there next year, and not the year after. That distinction is the entire value of doing this exercise more than twelve months out.

☐ One-off – this year only

Too pessimistic, and being wrong in this direction has its own cost: a community that writes off a year of funding it actually has may cut something it did not need to cut.

5. €9,000 that does not appear anywhere, being what the community would have paid for the bookkeeping, the maintenance and the bookings if one member were not doing all three unpaid.

☐ Recurring, but only while something else lasts **correct answer**

Right, and the something else is a person. This line is invisible precisely because it never enters the accounts, which is what makes it the most dangerous entry in the set: the community discovers the number on the day that member is ill, resigns, or finally says no. Put it in the run rate anyway – with the name of what it depends on beside it.

☐ Recurring – it will be there next year

It has recurred for years, which is why everyone has stopped seeing it. But nothing obliges it to recur, no agreement records it, and the person doing it is usually the last to raise it.

☐ One-off – this year only

It is the opposite of a one-off – it is the line most likely to have been there, unremarked, for a decade. The reason to count it is that a community which cannot see it cannot replace it either.

Where this comes from

If it is already happening

Work out the run rate, once. Not the budget – the gap. How much of last year's costs were covered by something that will not recur: a sale, a legacy, a grant, a reserve draw. That single number is the diagnosis, and most communities have never calculated it.

Then ask what those have left to give. A community funding a €40,000 annual shortfall from asset sales, holding perhaps €120,000 of saleable land, has at most three years – fewer once the cost of selling comes out, and fewer again if any of it is restricted. It is the value that matters, not the number of parcels, which are rarely worth the same. Said aloud, that changes a conversation which has been circling for a decade.

Separate the two decisions. *Can we afford this model?* and *what should we sell?* get answered together and should not be. Answering the second first is how thirty years happens.

What prevents it

A reserve fund with a stated purpose, and a rule that operating shortfalls may not be funded from capital without a decision at the constitutional tier. The point is not to forbid selling. It is to make selling a **decision** rather than a default – which is precisely what it stops being when it happens every year.

RCOS · LAYER 3

RCOS's Layer 3 exists to make economic power *“explicit, bounded, reviewable, and subordinate to governance”* – that last word being the one that does the work here.

Every resource in the governed scope must be classified as **commons or private** in a versioned registry, recording its steward, its access rules and **any constraint on transfer, sale or privatisation** (§5.1.1–2). That final field is what a thirty-year pattern needs in order to be visible: a sale constraint written down is a thing a decision has to override.

Two invariants close it. “**Resources declared as commons MUST NOT be privatized through informal, implicit, or unilateral action**” (§5.6.2) – and repetition is a form of implicit. And shared resources, flows and obligations **must be visible to the community by default**, with only limited and explicit exceptions (§5.6.1).

The treasury clauses add the missing half: spending authority must be explicitly bounded by thresholds and approval paths (§5.3.3), and the community must define **reserve, risk and liability policies including limits on debt and long-term obligations** (§5.3.6).

RCOS Core v0.1 – §5.1 Commons vs Private Resources, §5.6 Layer Invariants

SOURCES & FURTHER READING

- 1 **RCOS stress tests – Economy & Resources** – the three documented patterns this lesson covers, in specification form
- 2 [Findhorn Foundation – Trustees’ Annual Report, year ended 31 March 2024](#)[↗] – audited accounts, Scottish Charity SC007233 – the source of every Findhorn quotation here
- 3 [The Farm – Why The Farm Survived](#)[↗] – the community’s own account of the 1983 Changeover, and the figures quoted here
- 4 [The Farm – Tennessee Encyclopedia](#)[↗] – the cooperative structure, the population figures, and that departures preceded the change
- 5 **RCOS Core – Layer 3: Economic & Resource System** – the commons register, treasury transparency, and limits on accumulation

Written by EcoHubs members, with AI assistance for drafting and editing, and reviewed by a person before publication. Facts are checked against the sources listed; anything we could not verify is marked.

How this is written

Who belongs, and who decides that

Four patterns about the edges – what the community stands for, what it may ask of you, and how somebody stops being a member.

THE SHORT VERSION

Every community has a boundary. Most have never written down where it is, which means it gets located during an argument, about a specific person, by whoever is most upset.

Four patterns: **the purpose that quietly changed**, **the norm nobody could enforce**, **the rules that spread into private life**, and **removal with no process**.

They share a single missing thing – a stated answer to *what does this place actually require, and of whom*. Without it, the community defaults to the two worst options: enforcing nothing, or enforcing everything at once when it finally runs out of patience.

The edge of a community is where its structure is tested, and almost nothing about it is written down in advance.

That is not carelessness. Writing down what would get somebody removed feels, in year one, like planning for a betrayal that has not happened among people you have just chosen. So the question is deferred, and it stays deferred right up until the group is holding a live situation involving a real person with friends in the room.

The four

- **The purpose that quietly changed** – nobody voted to become something else, and enough people arrived wanting something else that it happened
- **The norm nobody could enforce** – a commitment central to the community's identity, with no process behind it
- **Rules that spread into private life** – the opposite failure, and the one that makes members leave without complaining
- **Removal with no process** – grounds invented during the case they are used on

Notice that two of these point in opposite directions. A community can fail by enforcing nothing and fail by enforcing too much, and **the same missing document causes both**. A group with no stated scope has no basis for acting and no basis for restraint, so it oscillates: years of tolerating everything, then one decisive over-correction that costs it three members.

The one that is missing from this list

There is a fifth pattern here that RCOS's stress tests do not catalogue, and it is worth naming because it decides how much every other rule on this page matters: **exit that is not really available**.

Nobody writes a rule forbidding departure. What blocks it is structural – an individual stake in the property that only converts to money when a buyer appears, an open-ended notice period, a queue with no position in it. Each part was put there by somebody being reasonable, and together they mean that leaving costs more than staying and disagreeing.

This is documented outside our field. Ran Abramitzky's study of the Israeli kibbutzim found that **a community's wealth functions as a "lock-in device" that increases the value of staying**, and that higher wealth reduces exit while sustaining a high degree of income equality. A stake you forfeit by leaving is a lock-in device whether or not anyone designed it as one.

It matters here because **every membership rule on this page is only as strong as the exit behind it**. A community whose members cannot leave has not achieved commitment. It has achieved captivity with good intentions, and its agreements are no longer consented to in any meaningful sense – they are merely endured.

WHAT WE DON'T KNOW

Being unable to leave and being unwilling to leave look identical from outside, and they feel similar from inside. The honest test is not whether anyone has left recently. It is **what leaving would cost, calculated in advance**, and whether the member could tell you that number today.

Is that the community's business?

A community has written a scope statement. It governs: the land and buildings, the shared budget, the common meals, the work rota, and who becomes a member. It states as out of scope: members' incomes, their religious practice, and their relationships. Nothing else is written down anywhere. Five things now happen – decide each one against that document, not against what you would want.

1. A member takes a well-paid job with an oil company. Several members find this incompatible with why they are all here.

- ☐ Out of scope – they may dislike it, and disliking it is not authority **correct answer**
Right, and this is the case that proves the document is real. Income is explicitly out of scope, so there is nothing to decide and nobody has to win an argument. If the community wants this to be its business, the route is to declare an identity constraint that is testable and to do it in the calm – not to discover one during a case about a named person.
- ☐ In scope – the agreements cover this
Nothing in the statement reaches it. Reading a general commitment into a scope that lists incomes as out of scope is exactly the move the document exists to prevent, and the next person it is used on will have watched how it was invented.

- ☐ Partly – separate the part that is governed from the part that is not
There is no governed part here. No land, no budget, no rota and no meal is affected – only how a number of members feel about a colleague's employer.

Where this comes from

2. A member starts a small furniture business, using the shared workshop on weekday evenings and three car-park spaces for deliveries.

- ☐ In scope – the agreements cover this **correct answer**
Right, and note it is not in scope because it is a business. Their income is nobody's business; the workshop and the car park are listed, and sustained private use of a shared asset is a decision the community is entitled to make. Framed that way it is an ordinary allocation question rather than a referendum on somebody's livelihood.
- ☐ Out of scope – they may dislike it, and disliking it is not authority
Income is out of scope. The workshop and the car park are explicitly in it, and this uses both, regularly, in a way that affects everyone else's access.
- ☐ Partly – separate the part that is governed from the part that is not
A reasonable instinct, and here the ungoverned part is not really in dispute – nobody is proposing to rule on what they earn. The question in front of the group is entirely about the shared assets.

3. A member decides to home-school their two children. Several others think this is a mistake and that the children will be isolated.

- ☐ Out of scope – they may dislike it, and disliking it is not authority **correct answer**
Right, and for a different reason from the oil job: this is not on either list. The default rule is what decides it – anything not explicitly declared as in scope is out of scope, so silence resolves in the member's favour. That is the opposite of how an unwritten arrangement drifts, where silence resolves in favour of whoever feels most strongly.
- ☐ In scope – the agreements cover this
It appears nowhere in the statement. A community that treats an unlisted matter as governed has no scope statement in practice, only a list of things it happens to have thought of first.
- ☐ Partly – separate the part that is governed from the part that is not
There is no governed part. If the children's presence at common meals or on the land raised a real question, that would be in scope – but concern about a parenting decision is not that question wearing a disguise.

Where this comes from

4. A member has not done a rota shift in four months and has not said why.

☐ In scope – the agreements cover this **correct answer**

Right, and the easy half is that the rota is listed. The harder half is what the community may do about it, which the scope statement does not answer – this is where a ladder of cheap, early, reversible steps matters, because a group with only one available response tends to wait until it is furious and then use it.

☐ Partly – separate the part that is governed from the part that is not

Nearly, and worth holding onto: the rota is governed, and why they have stopped may not be. A community can require the shifts without acquiring a right to the reason.

☐ Out of scope – they may dislike it, and disliking it is not authority

The work rota is on the list. A scope statement that cannot reach four months of unexplained non-participation is not protecting anybody.

Where this comes from

5. A member's new partner moves into their home. The partner does not come to common meals, does not do rota shifts, and has not applied for anything.

☐ Partly – separate the part that is governed from the part that is not **correct answer**

Right, and doing the separation out loud is what keeps this from becoming a row about somebody's relationship. The relationship is explicitly out of scope. Who becomes a member is explicitly in it, and so is the rota – so the live questions are what status a resident non-member has and what the household owes, neither of which requires anyone to have a view about the couple.

☐ In scope – the agreements cover this

Much of it is, and saying so without qualification is how the conversation ends up including opinions about the partner. The statement lists relationships as out of scope, and that clause earns its keep on exactly this kind of evening.

☐ Out of scope – they may dislike it, and disliking it is not authority

The relationship is out of scope; a person living on the land without any defined status is not. A community with no answer here has an undefined membership state, which is where expulsion-without-process cases usually begin.

Where this comes from

If it is already happening

Write the scope before you write the rules. A community arguing about whether it may have an opinion on somebody's diet, or their guests, or their income, is arguing about scope with only a specific case in front of it – which is the worst possible moment. What resolves it is a statement of what the community

governs and, crucially, what it does not.

Handle the live case with what you have, and say so. If somebody has done something the group cannot tolerate and there is no process, the honest move is a process invented transparently for this case, stated as such, and reviewed afterwards. That is far better than pretending a rule existed, and far better than doing nothing.

Then write the general version afterwards, not during. Rules written in the middle of a case are written about a person, and everybody knows it – including the next person they get applied to.

What prevents it

Three documents most communities have none of: a **scope statement** saying what is governed and what is not, **identity constraints** stated so they could actually be tested, and an **exit and separation protocol** covering both directions – voluntary and forced.

A fourth thing is worth separating while you are in there. Most communities hold one founding statement that is simultaneously why they exist, what they hope to become, and what they are doing about it this decade – which means the plan cannot be revised without reopening the purpose, or the purpose moves whenever the plan is updated. **Vision, purpose and mission** are three different documents with three different change rules, and pulling them apart is a morning's work.

RCOS · LAYER 0

RCOS puts the boundary in Layer 0 – constitutional, not operational, which is itself the argument. A community must **explicitly declare the scope of what it governs**, listing assets, decision domains and activities under collective control (§2.2.1–2), and must **explicitly list what is out of scope** (§2.2.3).

Then the rule that resolves the oscillation described above: **anything not explicitly declared as in scope must be treated as out of scope** (§2.2.4), and the community **must not exercise authority over persons, assets or domains that are out of scope** (§2.2.5). Silence resolves in the member's favour – which is the opposite of how an unwritten arrangement drifts.

On norms: a community must declare identity-level constraints that materially affect participation or behaviour (§2.4.1), and those constraints must be **testable and enforceable through defined processes** (§2.4.3) and **must not be enforced implicitly or informally** (§2.4.4). A value with no process behind it is not an identity constraint under RCOS. It is an aspiration, and the standard says so.

And underneath all of it, from Layer 1: **voluntary exit must be possible at all times** (§3.6.1), with exit procedures that are explicit, documented and **non-punitive** (§3.6.2), and **asset, role and responsibility separation defined before exit** rather than negotiated during it (§3.6.5).

RCOS Core v0.1 – §2.2 Scope Declaration, §2.4 Identity Constraints

SOURCES & FURTHER READING

- 1 **RCOS stress tests – Membership & Boundaries** – the four documented patterns this lesson covers, in specification form

- 2 **RCOS Core – Layer 0: Identity & Scope** – purpose, scope, and identity constraints that have to be testable
 - 3 **RCOS Core – Layer 1: Membership System** – membership states, rights and obligations, and the rules on exit
 - 4 **Abramitzky, R. – The Limits of Equality: Insights from the Israeli Kibbutz (QJE, 2008)**⁷ – community wealth as a lock-in device that raises the value of staying
 - 5 **Membership: getting in, and getting out** – the joiner-facing version – the four stages and what a fair agreement contains
-

Written by EcoHubs members, with AI assistance for drafting and editing, and reviewed by a person before publication. Facts are checked against the sources listed; anything we could not verify is marked.

How this is written

The work that breaks people

Five patterns that all end in the same place – a community running on two or three people who cannot stop, and a group that can see it and cannot fix it.

THE SHORT VERSION

The largest pattern in this guide, and the one communities are best at describing and worst at fixing.

Five modes: **care work nobody counts**, **responsibility that quietly becomes authority**, **growth outpacing onboarding**, **self-sufficiency without contribution**, and **information nobody can reach**.

They converge. A community running on two or three indispensable people has all five at once, and the reason it stays that way is not ignorance – it is that **redistribution needs a structure, and complaining is not one**.

Here is the thing that makes this lesson different from the others.

With power, with conflict, with money, the community usually cannot see the pattern. **With this one it can.** Members will describe it accurately, in detail, unprompted. They will name the person carrying too much. They will say it is not sustainable. And then nothing will change, for years.

That gap – between accurate diagnosis and no repair – is the whole subject.

The five

- **Care work nobody counts** – the noticing, remembering and smoothing that never reaches a rota
- **Responsibility that becomes authority** – the person who does it ends up deciding it, without anybody granting that
- **Growth that outpaces onboarding** – new members arrive faster than the community can teach itself to them
- **Self-sufficiency without contribution** – households quietly withdraw into their own lives, and the commons is maintained by whoever minds
- **Information nobody can reach** – the accounts, the passwords, the history, the reason the rule exists

Why knowing is not enough

FROM A COMMUNITY WE WORKED WITH

One person knew everything and held nearly every responsibility. They were visibly overwhelmed. Other members said so – openly, repeatedly, in the community's own words: he has too much.

And nobody found a way to spread it.

Not because they did not care, and not because he refused. Because “you have too much” is a description, and a description cannot be acted on. What was needed was a strategic approach – somebody sitting down with the list of what he actually did, splitting it into named pieces, and finding a holder and a handover for each. Nobody did that, because it is a project, and projects need an owner, and the only person who understood the list well enough to own it was him.

That is the mechanism, and it generalises. **The people best placed to redistribute the work are the people currently doing it, and they are the people with no capacity to run a redistribution project.** So the group circles: it notices, it says so, it feels bad, and the situation persists – while everyone involved behaves reasonably and cares about each other.

Three things follow, and they are the practical content of this lesson:

Complaining is not an intervention. A community that has raised something twenty times and seen no change does not have a communication problem. It has a missing structure, and more conversation will not produce one.

The work has to be enumerated before it can be shared. You cannot hand over “the admin”. You can hand over “renewing the insurance in March”, “reconciling the account quarterly”, and “keeping the key register” – three named things with three holders. The enumeration is most of the labour, and it is the step that gets skipped.

Somebody who is not overloaded has to run it. This is the part communities get wrong most reliably. Asking the overloaded person to organise their own unloading is asking for one more task from the person with the least room, and it is why the attempt fails on the third week every time.

Would that actually redistribute anything?

One member is carrying far too much and everybody knows it. Here are five things a community might do next. All of them are well meant. Two kinds of response fail reliably, and both feel like progress at the time.

1. At the next meeting the group agrees that this member has too much on, and that everyone should step up more.

- ☐ A description, not an intervention **correct answer**
Right, and a community can do this twenty times. “You have too much” is a description, and a description cannot be acted on – there is nothing in it anybody could start on Monday. A group that has raised something repeatedly and seen no change does not have a communication problem; it has a missing structure, and more conversation will not produce one.
- ☐ It would change something
Nothing here would. “Everyone” is not a holder and “more” is not a quantity, so there is no version of next week in which anybody has done this or failed to. It is the single most common response to overload and it has no mechanism in it at all.
- ☐ It asks the overloaded person for one more thing

Not directly – nothing is asked of them here. That is a different failure, and this one is gentler and just as ineffective.

2. The group asks the overloaded member to write down everything they do, so it can be shared out.

☐ It asks the overloaded person for one more thing **correct answer**

Right, and this is the version communities get wrong most reliably, because it looks like exactly the right step. Enumeration is most of the labour – and asking the person with the least room to produce it adds a task to the top of the pile. It is why the attempt fails in the third week, every time, and why the group concludes the person does not really want to let go.

☐ It would change something

The list is the right artefact and this is the wrong author. Everything that follows depends on a document only the busiest person can write, which means it does not get written.

☐ A description, not an intervention

It is more than a description – it names a concrete deliverable. The problem is who it is addressed to.

Where this comes from

3. A member who holds no roles spends two hours with them, writes down each separate thing they do, and comes back to the group with a list of eleven items.

☐ It would change something **correct answer**

Right – and note that nothing has been handed over yet. This is the step that gets skipped, done by the only kind of person who can do it: somebody with capacity. Eleven named items can each find a holder. "The admin" cannot. The list is also usually shorter than the group feared and more alarming in what it reveals about single points of failure.

☐ A description, not an intervention

A list of eleven named responsibilities is the opposite of a description – it is the thing that makes every later step possible, and it took one person two hours.

☐ It asks the overloaded person for one more thing

Two hours of their time, once, to produce something that removes work permanently. That is the trade worth making, and it is affordable precisely because somebody else is holding the pen.

Where this comes from

4. The group creates a wellbeing working group to look at workload across the community and report back in due course.

☐ A description, not an intervention **correct answer**

Right. A body with no list, no named owner and no date is a description with a membership. Projects need an owner and a deadline; without them this reports back in due course, which is to say at the meeting after the one where somebody remembers to ask. Compare the two hours in the previous question, which produced something.

☐ It would change something

A working group is not a holder – it is a set of people who may each reasonably assume another of them is doing it. "In due course" is not a date.

☐ It asks the overloaded person for one more thing

Usually it does, in the end: the overloaded member joins the working group about their own overload, because they know the most. But the primary failure here is that nothing in it has an owner.

5. The community writes into its agreements that no one person may hold more than three named roles at a time.

☐ It would change something **correct answer**

Right, though what it changes is subtler than a handover. Without a stated limit there is no moment at which anybody is entitled to say this is too much without it sounding like a complaint about their own capacity – which is exactly why the overloaded person never says it. A ceiling makes overload a fact about the structure that anyone may point at.

☐ A description, not an intervention

A written rule is not a description of the situation; it is a constraint on it, and it applies to people who have not joined yet. It does not redistribute anything by itself, which is why it belongs alongside the audit rather than instead of it.

☐ It asks the overloaded person for one more thing

It asks them for nothing. What it does is remove the need for them to be the one who raises it, which is the specific thing they have been unable to do.

Where this comes from

If it is already happening

Do the audit, not the conversation. One person, two hours, a list: everything that has to happen for this community to keep functioning, and who currently does it. Not who is *supposed* to. Who does.

Look for the single points of failure. Anything with one name against it and no second person who could do it. That list is usually shorter than people fear and more alarming than they expect.

Then take the top three and give each one a name, a scope and a handover date. Not all of it. Three. A community that redistributes three responsibilities properly has learned how, and can do the rest; a community that tries to redistribute everything produces a spreadsheet nobody maintains.

And write the limit down. How many roles one person may hold, or how many hours a month is too many. Without a stated limit there is no moment at which anybody is entitled to say *this is too much* without it sounding like a complaint about themselves – which is exactly why the overloaded person never says it.

What prevents it

Named roles with terms and handovers, a stated ceiling on how much any one person carries, and a definition of contribution that counts the invisible work as work. None of that is about fairness. It is about a community being able to continue when somebody is ill.

RCOS · LAYER 5

RCOS states the enumeration step as a requirement, which is unusual and is the point: **“All ongoing responsibilities MUST be assigned to explicit, named roles rather than implicit expectations or informal agreements”** (§7.1.1). And its protective mirror – **no person may be held accountable for responsibilities not formally assigned to a role** (§7.1.4).

Every role enters a Role Registry with its purpose, decision authority, boundaries, and its appointment, review and removal process; term length and rotation are recorded where a role has them (§7.1.2). Each role must also carry an accountability mechanism defining **how underperformance, overload, or role failure is handled, and how handover and knowledge transfer occur** (§7.1.3) – overload named as a role-design question rather than a personal one.

Workload boundaries are a required artifact, not a courtesy: explicit limits on meeting load and role load, and a named mechanism for **relief, substitution or redistribution** (§7.4.2). And the trigger clause: **persistent overload, burnout risk, chronic non-participation, or dependency on over-functioning individuals MUST trigger review or repair processes** as defined in Layer 4 (§7.4.4).

Then the clause that names the case above directly: **“The community MUST ensure that no single individual is a critical single point of failure for core operations”** (§7.5.1), with documented procedures, clear handover mechanisms and backup arrangements where feasible (§7.5.2).

Layer 3 closes it from the economic side: a community **must not structurally depend on unpaid, invisible or informal labour for system survival** without defining the corresponding recognition, obligation or compensation (§5.2.3).

RCOS Core v0.1 – §7.1 Roles and Responsibilities, §7.4 Workload and Capacity Boundaries

SOURCES & FURTHER READING

- 1 **RCOS stress tests – Operations & Coordination** – the five documented patterns this lesson covers, in specification form

- 2 **RCOS Core – Layer 5: Operations & Coordination** – the Role Registry, workload boundaries, and the single-point-of-failure clause
 - 3 **Freeman, J. – The Tyranny of Structurelessness (1970/1972)**⁷ – delegate specific authority for specific tasks, rotate them, and spread information widely
 - 4 **Daily life in community** – the work nobody put on the rota, and what a week actually contains
-

Written by EcoHubs members, with AI assistance for drafting and editing, and reviewed by a person before publication. Facts are checked against the sources listed; anything we could not verify is marked.

How this is written

Changes that cannot be undone

Five patterns about how rules change – the temporary measure that stayed, the founding principle nobody protected, and the community that can no longer remember why.

THE SHORT VERSION

A community that cannot change its rules dies of rigidity. One that changes them without noticing dies of drift. Both are failures of the same missing thing: **a stated way for a rule to change, and a record of when it did.**

Five patterns: **the invariant nobody protected, the emergency measure that stayed, governance with no path to change it, experiments with no way back, and a community that forgot why.**

The defence is not vigilance. It is bookkeeping – dating the agreements and keeping the old versions – which is unglamorous, cheap, and the only thing that works.

Every rule in a community was written by people responding to a situation. Ten years later some of those situations are gone, some of the rules are load-bearing, and almost nobody can tell which is which.

That is the subject of this lesson: **not what the rules are, but how they move** – and the fact that most communities have no account of their own movement at all.

The five

- **The invariant nobody protected** – the founding principle everybody assumed was permanent, changeable by an ordinary majority on a Tuesday
- **The emergency measure that stayed** – decided fast because it had to be, never reviewed, now precedent
- **Governance with no path to change it** – the opposite failure, and the one that produces departures rather than arguments
- **Experiments with no way back** – a trial that quietly became the arrangement, because nobody set an end date
- **A community that forgot why** – the rule survives, the reason does not, and nobody can evaluate either

And one more that RCOS does not catalogue: **sub-groups with undefined autonomy**, which is what happens when a community grows into parts without deciding what the parts may decide.

Why drift is the hardest of them

Four of these announce themselves eventually. Drift does not, and it is the one most communities are actually living in.

Nobody notices a slow forgetting by paying closer attention, because **there is no moment to notice**. The agreements stop being read, then stop being taught, then describe a place that no longer exists – and every step of that is a non-event.

Which is why vigilance is the wrong defence and bookkeeping is the right one. A community that can say **which version is in force, when it last changed, and who decided it** has turned a feeling into a check anybody can run.

The difference that makes is mostly social. Without a current text, raising drift means saying *I think we have lost something* – a criticism of everyone present, and heard as one. With it, the same observation is that the practice and the document disagree. A discrepancy, pointed at rather than levelled, and the sort of thing a group can look into without anyone having to be at fault.

FROM A COMMUNITY WE WORKED WITH

A founder produced a new version of a foundational document. Nobody could say whether it had ever passed the community's consensus process – the original was old enough that the memory had gone.

So it sat alongside the previous version, with a new version number, **no record of adoption, and no statement of what had changed between them**.

The community was not in conflict about the content. It could not establish **which of its own documents was in force**, or what the difference was. Everything downstream of that – every rule, every decision citing it – inherited the ambiguity.

Note how little would have prevented it: an adoption record and a summary of changes. Two fields.

Is there a way back?

Five changes a community has made. Some were called temporary, some were called experiments, and one was not called anything. For each, decide whether the community could actually return to how things were – not whether it should.

1. Three years ago the group suspended the two-week notice rule for guests, "while the building work is on". The building work finished last spring.

- ☒ On paper only – the decision could be reversed, the situation cannot **correct answer**
Right, and the trap is that reversing it looks trivial: one vote, and the old rule is back. But three years of practice is now the normal way of doing things, so reinstating the rule would not be experienced as returning to the agreement – it would be experienced as a new restriction, proposed by somebody, on everybody. That is how an emergency measure becomes the constitution without ever being adopted.

- ☐ Yes – there is a date and the old arrangement still exists

There is no date. "While the building work is on" sounded like one at the time, which is precisely why nobody set a real one, and the condition it was tied to passed a year ago without anything happening.

- ☐ There is nothing to reverse – no change was ever recorded

A suspension was decided, so there is a decision on the record to point at. That makes this recoverable in a way the genuinely unrecorded cases are not.

Where this comes from

2. The community agreed to try shared childcare for one term. There is a review date in the diary, and the previous arrangement is still running alongside it.

- ☐ Yes – there is a date and the old arrangement still exists **correct answer**

Right, and both halves are load-bearing. The date means somebody has to actively decide to continue rather than merely fail to stop. Keeping the old arrangement running is the expensive half and the one usually skipped – a way back that has been dismantled is not a way back, however clearly the minutes say the trial may end.

- ☐ On paper only – the decision could be reversed, the situation cannot

Nothing has been dismantled and no facts have been created that would outlast the term. This is what a time-bounded experiment is supposed to look like.

- ☐ There is nothing to reverse – no change was ever recorded

It was recorded, with an end date, which is more than most arrangements described as trials ever get.

3. Six households were allowed to opt out of common meals "as an experiment". Four have since converted their dining rooms into other things.

- ☐ On paper only – the decision could be reversed, the situation cannot **correct answer**

Right. The vote is still there to be retaken and the rooms are gone. An experiment that lets people build on its outcome has stopped being an experiment, whatever the minutes call it – and the households would be right to feel that ending it now is a change being made to them rather than a trial concluding.

- ☐ Yes – there is a date and the old arrangement still exists

There is no date in the scenario, and for four households the old arrangement no longer physically exists. Both of the things that make a trial reversible are missing.

- ☐ There is nothing to reverse – no change was ever recorded

A permission was granted and can be found. The problem is not the record – it is that the situation moved further than the decision did.

Where this comes from

4. The membership fee is named in the founding charter. Last month the group changed it by ordinary majority at a Tuesday meeting, and nobody present raised the charter.

- ☐ On paper only – the decision could be reversed, the situation cannot **correct answer**
Right, though the reason is unlike the others: the fee could be put back tomorrow, and what cannot be put back is the discovery that a founding document can be amended by an ordinary majority on a Tuesday. Everyone now knows it, and the next thing in the charter is protected only by nobody happening to want it changed. An invariant is a rule about which rules a Tuesday meeting may reach.
- ☐ There is nothing to reverse – no change was ever recorded
The change was made and minuted. What was not recorded is the collision with the charter, because nobody in the room noticed there was one.
- ☐ Yes – there is a date and the old arrangement still exists
The old figure is recoverable from the charter, so the fee itself is easy. The precedent is the part with no way back, and it is now available to anybody who wants it.

Where this comes from

5. A revised version of the work agreement appeared in the shared folder with a new number. There is no adoption record, no summary of what changed, and the previous version was deleted to avoid confusion.

- ☐ There is nothing to reverse – no change was ever recorded **correct answer**
Right, and this is worse than an unpopular change, because the community cannot establish what it used to require or whether the new text was ever agreed. Everything downstream – every rule and decision citing the agreement – inherits the ambiguity. Deleting the old version to avoid confusion destroyed the only evidence of what the community previously thought, which is exactly what you need when somebody asks why a rule exists.
- ☐ On paper only – the decision could be reversed, the situation cannot
There is no decision to reverse. That is the distinction worth holding: a bad change can be revisited, and an unrecorded one leaves nothing to revisit.
- ☐ Yes – there is a date and the old arrangement still exists
Neither is true – the old text was deleted and no adoption date exists. Two fields, an adoption record and a summary of changes, would have prevented the whole situation.

Where this comes from

If it is already happening

Establish the present before you reconstruct the past. Which documents are in force, as of today, agreed by whoever is in the room. Not their history – their status. This can be done in one meeting and it stops the bleeding.

Then date everything from now on. A change record with four fields: what changed, who decided it, under which rule, and from when. A community with two clean years and a murky past is in a far better position than one that spent six months excavating and gave up.

Keep the superseded versions. Deleting the old text destroys the only evidence of what the community used to think, which is what you need when somebody asks why the rule exists.

And review the emergency measures specifically. Every community has two or three arrangements adopted quickly during something difficult and never revisited. List them, and give each one a date by which it is either ratified properly or ends.

What prevents it

A **change protocol** that says how a rule changes and who decides, a **version history** that makes “which version is in force” answerable, and a **learning log** so that what a difficult year taught does not leave with the people who lived it.

All three are boring. That is the point – this is the one failure class where the defence is entirely clerical, and a community that finds that beneath it will forget why it exists.

RCOS · LAYER 6

RCOS's sixth layer treats evolution as a governed process rather than improvisation, and its central clause is a validity rule: “**No informal, undocumented, or ‘understood’ rule changes MAY be considered valid**” (§8.2.5).

Every adopted change must be versioned and traceable (§8.2.1), and the **Version History** must record the version identifier, adoption and effective dates, the **decision record reference – authority, mechanism, threshold** – a summary of changes, and any migration notes (§8.2.2). At any point the community must be able to determine **which version is currently in force and which artifacts are authoritative** (§8.2.3), and superseded rules must remain accessible **together with the dates during which they were in force** (§8.2.4).

The layer invariants set the two boundaries this lesson runs between: **change must be possible but constrained – no change may be instantaneous, implicit, or unreviewable** (§8.7.1), and **major failures and adaptations must be captured as shared learning, not erased or hidden** (§8.7.4).

Underneath it, from Layer 0: **no decision, role, process, or emergency measure may override an invariant** (§2.3.4), and where an invariant conflicts with any other rule, **the invariant prevails** (§2.3.5).

RCOS Core v0.1 – §8.2 Versioning and Authority, §8.7 Layer Invariants

SOURCES & FURTHER READING

- 1 **RCOS stress tests – Change & Emergencies** – the five documented patterns this lesson covers, in specification form
- 2 **RCOS Core – Layer 6: Evolution & Adaptation** – change mechanisms, version history, time-bounded experiments and the learning log
- 3 **Just For Now: how emergency rules become permanent** – the longer EcoHubs argument about temporary measures that outlive their emergency
- 4 **Resilience & why communities fail** – drift as one of the six recurring patterns, and why it is the only one that does not feel like anything

Written by EcoHubs members, with AI assistance for drafting and editing, and reviewed by a person before publication. Facts are checked against the sources listed; anything we could not verify is marked.

How this is written

The failure modes, in full

25 patterns, in the order the lessons introduce them. Each one gives what it looks like from inside, why it is hard to see, what to check, and what to change.

L2 Authority that comes from being compelling

Someone is so persuasive that disagreeing with them feels like a failure of understanding rather than a difference of view.

L2 Decisions with no stated type

The same question is routine or constitutional depending on who you ask – so the answer depends on who gets to frame it.

L2 Speakers who fill the room

Nothing is blocked and nobody is silenced – but only some people's proposals ever get finished, and the group calls it consensus.

L2 The clique that decides before the meeting

The real conversation happens in the kitchen and the meeting ratifies it – so attendance stops changing anything, and then stops.

L2 The founder's informal veto

No rule grants it and everyone routes around it anyway – the objection that ends a discussion without ever being tested.

L4 Conflict that is never named

The group is warm, nothing is confronted, and difficulty gets absorbed quietly – until somebody leaves and everyone privately knew why.

L4 Punishment before repair

The group does act – and reaches straight for the consequence, so nobody raises anything small ever again.

L3 Contribution that stops being shared

Some members pay their way out of the work – reasonably, one at a time – until the rota is done by whoever cannot afford not to.

L3 Finances nobody can inspect

Not hidden, exactly – just held by two people, reported in summary, and never actually examined by anyone else.

L3 The commons sold to cover the shortfall

Every sale funds another year, every sale is defensible, and the thing being sold is what made the community possible.

L1 Removal with no process

Grounds invented during the case they are used on – and every remaining member quietly learns what the rules really are.

L0 Rules that spread into private life

Each new rule was a reasonable answer to a real problem, and together they have opinions about your visitors, your diet and your evenings.

L0 The norm nobody could enforce

A commitment written into the founding vision, and no process behind it – so the first person to break it discovers there was nothing there.

L0 The purpose that quietly changed

Nobody voted to become something else – enough people arrived wanting something else that it happened anyway.

L5 Care work nobody counts

The noticing, remembering and smoothing that never reaches a rota – done by the same people, unrecorded, until one of them stops.

L5 Growth that outpaces onboarding

New members arrive faster than the community can teach itself to them – so they learn a version, and the version drifts.

L5 Information nobody can reach

The accounts, the passwords, the history and the reason the rule exists – held by whoever happened to be there, and reachable only through them.

L5 Responsibility that becomes authority

The person who does a thing ends up deciding it – nobody granted that, and nobody can take it back without taking the work back too.

L5 Self-sufficiency without contribution

Households withdraw into their own lives, one reasonable decision at a time, and the commons is maintained by whoever still minds.

L6 A community that forgot why

The rule survives and the reason does not – so nobody can tell whether it still earns its place, or which version is even in force.

L6 Experiments with no way back

A trial that quietly became the arrangement – because nobody wrote down when it ended or what would count as it failing.

L6 Governance with no path to change it

The rules can no longer be changed, so people stop trying – and the community loses its members quietly instead of loudly.

L0 Sub-groups with undefined autonomy

A community grows into parts, and nobody says what the parts may decide – so each one improvises a constitution and discovers later they do not match.

L6 The emergency measure that stayed

Decided fast because it had to be, never reviewed, and now it is what the community does – with a precedent attached.

L0 The invariant nobody protected

The founding principle everybody assumed was permanent – changeable, it turns out, by an ordinary majority on an ordinary evening.

Authority that comes from being compelling

Someone is so persuasive that disagreeing with them feels like a failure of understanding rather than a difference of view.

WHAT IT LOOKS LIKE

- Disagreement gets reframed as something the disagreeing person needs to work on.
- Their view is described as insight, and other people's as opinion.
- Newcomers adopt their framing unusually fast, including its vocabulary.
- There is no forum where their contribution is treated as one contribution.

Why it is hard to see

Because the alternative explanation is always available, and it is often true: **this person may simply be right more often.**

Communities frequently contain somebody with unusual clarity – about the land, about people, about what the project is for. Deferring to them is not irrational. It is what you would want a group to do with a genuinely wise member, and any attempt to describe the problem runs immediately into “so we should ignore good judgement?”

The distinction that holds up is not about the person's quality. It is about **what disagreement costs.**

In a healthy version, disagreeing with the compelling person is ordinary. You might be wrong, they might be wrong, and finding out is a normal Tuesday. In the failure version, disagreement gets **reframed as a deficiency in the person disagreeing** – you have not understood, you are not ready, you are bringing something unresolved. Once that move is available, the group has no way to test anything they say, because every test is converted into evidence about the tester.

The second reason it is hard to see: **it is usually most visible to people who have just left**, and least visible to people in the middle of it. That asymmetry is not a coincidence – it is the mechanism working.

WHAT WE DON'T KNOW

This page is about a structural condition, not about diagnosing an individual. A community can produce this pattern around somebody who neither sought it nor enjoys it, and the fix is the same either way: make the authority testable.

What to check

Can you name a time they were overruled? Not persuaded, not gracious in disagreement – overruled, on something they cared about, by a process. A community that cannot produce one example in five years has an untested authority whatever it believes about itself.

Watch what happens to a newcomer's first disagreement. Newcomers have not yet learned what is not raised, which makes them the group's best instrument. If the room's response is to explain rather than to consider, that is the pattern.

Ask what would change their mind. A person operating inside ordinary authority can usually answer. If the honest answer is “nothing, on this”, the group needs to know which subjects those are.

If it is already happening

Do not open with the person. Every version of “you have too much influence” is heard as an accusation and produces a conversation about loyalty, which the group will lose.

Open with the record. Write down decisions, their method, and who decided them. This pattern is unusually vulnerable to bookkeeping, because it depends on authority being felt rather than located. A decision log makes “we all agreed” checkable, and often it turns out that a decision everyone remembers as collective was made by two people in a car.

Build one route that does not pass through them. A conflict process with an external facilitator, or a domain with a spending limit and a named holder. Not to reduce their influence – to give the community somewhere to stand that does not require their agreement.

What prevents it

The same thing that prevents the others: authority that is **assigned, bounded and reviewable**, so that being persuasive and being in charge are different states a group can tell apart.

RCOS · LAYER 2

“Authority MUST NOT be derived from charisma, seniority, ownership, or informal influence” (§4.3.4) – charisma is the first word on that list, and this pattern is what it names.

The safeguard clauses close the escape route. Governance mechanisms **must allow for challenge and review without retaliation** (§4.6.2) – and reframing a challenge as the challenger's personal deficiency is a form of retaliation that leaves no evidence, which is why it needs naming rather than merely disallowing.

There is a Layer 0 angle worth knowing too. Identity constraints – the ethical and cultural boundaries a community declares – **must be testable and enforceable through defined processes** and **must not be enforced implicitly or informally** (§2.4.3–4). A community whose norms are held in one person's judgement rather than in a document has made that person the constraint, which is exactly the condition this pattern needs.

SOURCES & FURTHER READING

- 1 **RCOS stress test – Charismatic Spiritual Authority** – the specification this page puts into plain language
- 2 **Freeman, J. – The Tyranny of Structurelessness (1970/1972)**[↗] – why an undescribed structure cannot be elected, limited or voted out
- 3 **RCOS Core – Layer 0: Identity & Scope** – identity constraints must be testable, and must not be enforced informally

Written by EcoHubs members, with AI assistance for drafting and editing, and reviewed by a person before publication. Facts are checked against the sources listed; anything we could not verify is marked.

How this is written

Decisions with no stated type

The same question is routine or constitutional depending on who you ask – so the answer depends on who gets to frame it.

WHAT IT LOOKS LIKE

- Arguments about whether something needs the whole group start before the discussion does.
- The same kind of decision has gone two different routes in the last year.
- Somebody can end a debate by saying "that's just operational".
- Nobody can say which decisions require a higher bar without checking with a person.

Why it is hard to see

Because it does not look like a power problem. It looks like an unusually procedural argument, and everybody hates it.

A community without stated decision types spends its energy in a strange place: not on *what should we do* but on **who has to be asked**. Should the workshop roof go to the whole meeting or to the maintenance circle? Is changing the guest policy a rule change or an operational tweak? Those arguments feel like bureaucracy, so groups short-cut them – and the short-cut is the failure.

The short-cut is that **whoever frames the question first usually wins it**. Call it operational and it goes to three people this week. Call it constitutional and it goes to everyone, slowly, with a real chance of being refused. Nobody has to act in bad faith for this to matter; framing is genuinely ambiguous, and the person who speaks first is usually the person most invested in the outcome.

There is a second reason it hides. **It only produces visible damage in one direction**. A constitutional decision waved through as operational is invisible until the consequences arrive – the five-year lease, the changed membership rule, the sold asset. The reverse error, treating something trivial as constitutional, is merely tedious, and tedium gets noticed and complained about. So groups drift steadily toward the lower tier, and the drift is one-way.

What to check

Take the last ten decisions and sort them. Not by importance – by *what it would cost to reverse them*. A community with no stated types will find pairs that cost roughly the same to undo and went through completely different processes.

Ask who classifies. Not who decides – who decides what kind of decision it is. In most communities the honest answer is "whoever wrote the agenda", and that is a real power nobody granted.

Look for the missing tier. Groups usually have an implicit operational route and an implicit everyone-decides route. What is almost always absent is the middle one: decisions that commit real money or change direction but are not constitutional. Those get shoved to whichever end is closer, and it is usually the near end.

If it is already happening

Do not try to classify everything. A group that sits down to sort its entire decision history will still be sorting in March.

Write the boundary, not the taxonomy. One number does most of the work: the amount a circle may spend without asking. A domain without a figure is not a domain, and a community that writes one figure has classified more decisions in an evening than one that spent a month on definitions.

Then add the default rule, which is the cheapest safeguard in this whole guide: **anything nobody can classify goes to the higher tier.** Without it every borderline case drifts to the easier meeting, and the drift is exactly what put the group here.

What prevents it

Stated decision tiers

, and a

decision matrix

that says which route each kind takes. The point is not the categories. It is that ****classification stops being a move available to whoever speaks first.****

RCOS · LAYER 2

The RCOS Standard requires that **all collective decisions be classified into exactly one of three types** – operational, strategic, constitutional (§4.1.1). Exactly one: a decision cannot be partly constitutional, which is the ambiguity this pattern lives in.

It defines them by what they touch. Operational is day-to-day functioning **within existing rules** (§4.1.2). Strategic is long-term direction, significant resources, or creating and removing major structures (§4.1.3). Constitutional is changes to Layer 0 invariants, purpose, scope, or the governance system itself (§4.1.4).

And then the clause worth writing on the wall: **“If a decision cannot be clearly classified, it MUST default to the higher-impact decision type”** (§4.1.5). One sentence, and it removes the entire incentive to frame ambiguously – because ambiguity now costs the framer rather than paying them.

Each type must have its own explicitly defined mechanism (§4.2.1), specifying eligible participants, thresholds, blocking conditions and time constraints (§4.2.3).

RCOS Core v0.1 – §4.1 Decision Types

SOURCES & FURTHER READING

- 1 **RCOS stress test – Undefined Decision Type** – the specification this page puts into plain language
- 2 **Decision methods** – the three tiers in full, and why no community that lasts uses one method
- 3 **Buchanan, J. & Tullock, G. – The Calculus of Consent (1962)**[↗] – why the best decision rule differs by class of decision – the argument for tiers

Written by EcoHubs members, with AI assistance for drafting and editing, and reviewed by a person before publication. Facts are checked against the sources listed; anything we could not verify is marked.

How this is written

Speakers who fill the room

Nothing is blocked and nobody is silenced – but only some people's proposals ever get finished, and the group calls it consensus.

WHAT IT LOOKS LIKE

- The same two or three people speak first on almost every item.
- Proposals from quieter members arrive under-developed, and are judged for it.
- Meetings run long and the last agenda items are always decided fastest.
- Somebody who stopped contributing cannot say when or why they stopped.

Why it is hard to see

Because everybody had the same opportunity, and the group can prove it.

Nobody was interrupted. Nobody was told to be quiet. The floor was open the whole time. Under any procedural audit this meeting was fair – which is why a group that senses something is wrong here usually cannot name it, and eventually concludes it is imagining things.

What is actually happening is that **speaking is a skill, and the group is unintentionally selecting on it**. Confidence in a room, speed of assembly under pressure, comfort with disagreement in public, and first-language fluency are not distributed evenly and have nothing to do with the quality of anyone's judgement. A community that decides in real time, out loud, hands its decisions to the people who think fastest in that specific format.

The compounding part is the reason this belongs in a failure guide rather than a facilitation manual. **The people who lose these exchanges stop starting them.** Not dramatically – they just stop putting things on the agenda, and the group reads the absence as agreement. By the time anyone notices, the quiet members are no longer quiet in meetings; they are quiet about the community, which is a different and much later problem.

What to check

Count. One meeting, a tally of who speaks and for roughly how long. It takes no skill and it ends the argument about whether the impression is real. Groups are routinely shocked by their own numbers.

Then check the second-order effect, which matters more: of the last ten proposals, who wrote them? A group can have balanced airtime and still have every proposal originate with the same three people.

Ask a quiet member privately what they have stopped raising. Not “do you feel heard” – that question has one polite answer. The specific version gets a specific answer.

If it is already happening

Change the format before you change the people. Asking a dominant speaker to speak less puts the burden on individual restraint, which lasts about three meetings. Changing how the room works lasts.

The cheapest interventions, in order of how little they cost:

- **Written proposals circulated in advance.** This does more than everything else combined, because it moves the moment of formation out of the room. Someone who needs a day to think now contributes on equal terms.
- **Rounds on contentious items.** Everyone speaks once before anyone speaks twice. Slower, and it reliably surfaces the objection that would otherwise have arrived after the decision.
- **Rotate facilitation**, and separate it from advocacy. A facilitator who is also arguing the case is deciding whose turn it is in a debate they have a stake in.

Do not expect gratitude. The people who lose airtime under these rules are usually the people carrying the most work, and they will experience the change as friction. It is worth saying out loud that the point is not to hear less from them – it is to hear at all from people whose contribution currently arrives too late to matter.

What prevents it

A stated deliberation window with written proposals, and facilitation that rotates. Between them they move decision-forming out of live conversation, which is where the skill bias lives.

RCOS · LAYER 2

The RCOS Standard treats this as a governance-design obligation rather than a meeting-manners problem: the governance system **must include safeguards against concentration of decision power** and against **decision capture by subgroups** (§4.6.1). Which ones is left to the community – the requirement is that the safeguards exist and are explicit.

Two supporting clauses matter here. Decision mechanisms must specify their **eligible participants** and **time constraints** (§4.2.3) – a decision with no stated window is one that closes whenever the room runs out of energy, which favours whoever has most. And the Governance Protocol must define the full lifecycle of a decision including its **review and deliberation process** (§4.5.2), which is the phase this pattern skips.

Note what RCOS does **not** say. It sets no rule about how much anyone may speak, and prescribes no facilitation method. The constraint is on the outcome – power must not concentrate – not on the manners.

RCOS Core v0.1 – §4.6 Safeguards and Failure Modes

SOURCES & FURTHER READING

- 1 **RCOS stress test – Dominant Speakers in Decision Spaces** – the specification this page puts into plain language
 - 2 **Stasser, G. & Titus, W. – Pooling of unshared information in group decision making (1985)**[↗] – why unstructured discussion under-surfaces what only one member knows
 - 3 **Seeds for Change**[↗] – free practical guides on facilitation, rounds and handling dominant contributors
-

Written by EcoHubs members, with AI assistance for drafting and editing, and reviewed by a person before publication. Facts are checked against the sources listed; anything we could not verify is marked.

How this is written

The clique that decides before the meeting

The real conversation happens in the kitchen and the meeting ratifies it – so attendance stops changing anything, and then stops.

WHAT IT LOOKS LIKE

- Proposals arrive at the meeting already resolved, and discussion feels ceremonial.
- Some members are surprised by context others clearly had beforehand.
- Attendance is falling among people who have not stopped caring.
- The same three or four names appear in every account of "how that was decided".
- New members take an unusually long time to work out how anything happens.

Why it is hard to see

Because a group of friends who happen to be neighbours will always talk, and **most of that talking is good**. Communities are supposed to have relationships. Insisting every conversation about the community happen in a meeting would be both impossible and unpleasant.

So the line is not *did people talk outside the meeting*. It is **whether the meeting can still change the answer**.

That distinction is invisible from inside the clique and obvious from outside it, which is what makes the pattern so durable. The people in the kitchen experience themselves as doing preparatory work – and often they are. The people outside it experience a meeting where their contribution arrives too late to matter, and they cannot point at anything improper, because nothing improper happened.

The tell is not secrecy. It is **the shape of the meeting**: proposals that arrive fully formed, discussion that adds nothing, and a decision that was going to happen anyway.

What to check

Ask when the last meeting changed a proposal. Not rejected one – changed one. A group whose meetings routinely amend what arrives is deliberating. A group whose meetings ratify is performing.

Watch who has context. In any meeting, notice the moment somebody refers to a consideration nobody has mentioned. Where did that come from, and who else already had it?

Ask the newest member how decisions get made. Someone in their first year will describe the real process, because they are still trying to learn it. Someone in their tenth has forgotten there was anything to learn.

If it is already happening

Move the information, not the friendship. The fix is almost never “stop talking to each other” – that is unenforceable and would make the community worse. It is that whatever the kitchen conversation produced has to arrive in writing, before the meeting, available to everybody.

Publish proposals with a deliberation window. A written proposal circulated three days ahead does more against this pattern than any rule about conduct, because it moves the moment of formation into a space everybody can reach. It also fixes the milder version of the problem, which is not a clique at all but simply people who think faster out loud.

Then check attendance in six months. This pattern’s real cost is not the decisions; it is the members who quietly concluded that showing up does not matter. That recovers slowly, and only after the meetings visibly start changing things.

What prevents it

A **stated deliberation phase** before voting opens, and a decision record that names the method and the threshold. Both make the meeting the place where the decision is actually formed, rather than the place where it is announced.

RCOS · LAYER 2

“No informal or ad hoc decision mechanism MAY be used for collective decisions” (§4.2.4). Blunt, and aimed exactly here – the kitchen is an informal mechanism, however sound its conclusions.

Every mechanism must specify its eligible participants, thresholds, blocking conditions and time constraints (§4.2.3), which is what makes “who was in that conversation” answerable rather than a matter of impression. And the safeguards clause requires the governance system to protect specifically against **decision capture by subgroups** (§4.6.1) – this pattern, named.

The Governance Protocol must cover the full lifecycle of a decision including its review and deliberation stage (§4.5.2), and all governance actions must be documented under the Layer 5 rules (§4.5.4). A decision with no record of where it was formed cannot be audited for this at all.

RCOS Core v0.1 – §4.2 Decision Mechanisms, §4.6 Safeguards and Failure Modes

SOURCES & FURTHER READING

- 1 **RCOS stress test – Informal Clique Decision Making** – the specification this page puts into plain language
- 2 **Freeman, J. – The Tyranny of Structurelessness (1970/1972)** [↗] – informal structure forms around friendship networks and access to information
- 3 **Proposals & ratification** – the deliberation window, and why skipping it is where this pattern lives

Written by EcoHubs members, with AI assistance for drafting and editing, and reviewed by a person before publication. Facts are checked against the sources listed; anything we could not verify is marked.

How this is written

The founder's informal veto

No rule grants it and everyone routes around it anyway – the objection that ends a discussion without ever being tested.

WHAT IT LOOKS LIKE

- Proposals get quietly checked with one person before they reach a meeting.
- Their objection ends a discussion, and nobody asks which rule it came under.
- Members describe ideas as "not worth raising", without being able to say why.
- The agreements say nothing about them having any special standing.
- When they are away, decisions wait rather than proceed.

Why it is hard to see

Because it is usually correct.

A founder has been right more often than anyone else in the room, on this specific subject, for years. The objection that ends the discussion is generally a good objection – it is *based on more information*, which is the whole problem. Deferring to someone with more context is not a failure of a group. It is what a sensible group does.

The failure is that **the deference was never bounded**. Nobody wrote down which decisions it applies to, what would count as overriding it, or when it stops. So it applies to everything, indefinitely, including the decisions where the founder's extra context is worth nothing at all.

The second reason it stays invisible is that it produces no incidents. A formal veto leaves a record – a blocked proposal, a minuted objection. An informal one leaves nothing, because **the proposals that would have triggered it were never made**. What you are looking for is not an argument. It is an absence: things people stopped bothering to raise.

What to check

Ask three people who joined in the last two years to describe a decision they decided not to propose. If they all name the same person, the veto exists whatever the agreements say.

Then the succession question, which is the one that can be asked without accusation: **what happens to this decision if that person is ill for a month?** A community with a bounded arrangement answers it easily. A community with an unbounded one discovers, out loud, that it cannot.

If it is already happening

Do not start by removing anything. The founder is usually carrying real load, and a group that opens with “your influence is a problem” will spend the next year on the relationship instead of the structure.

Start by **converting one decision**, not the person’s standing. Pick a category – spending under a threshold, say – and write who decides it, by what method, within what limit. The founder should be in the room when it is written, and it should be framed as what it is: making the thing they have been carrying into something the community can carry.

Then **check the conversion actually happened.** The failure mode after this one is a transition that exists on paper: roles created, elections held, and everyone still checking informally before deciding anything. That is harder to fix than doing nothing, because the group now believes the problem is solved.

What prevents it

Authority that is **assigned, scoped and time-bounded** from the start – including the founder’s, especially the founder’s. A founding group that writes down what its founders may decide, and for how long, is not expressing distrust. It is making the founders replaceable, which is the only way a community outlives them.

RCOS · LAYER 2

“**Authority MUST NOT be derived from charisma, seniority, ownership, or informal influence**” (§4.3.4). This pattern is the clause’s main target: seniority and informal influence, doing the work that an assigned role is supposed to do.

Two more clauses close the routes around it. All authority must be assigned to explicitly defined roles, circles or bodies, carrying scope, limits and term (§4.3.1–2) – so “the founder decides” is not an available answer, because a founder is a person, not a role. And **decisions made outside the Decision Matrix are invalid** (§4.4.4), which removes the informal channel’s standing entirely rather than arguing about whether it was used wisely.

The safeguards section names **founder entrenchment** outright as something the governance system must protect against (§4.6.1), alongside informal vetoes and decision capture by subgroups.

RCOS Core v0.1 – §4.3 Authority Boundaries

SOURCES & FURTHER READING

- 1 **RCOS stress test – Founder Informal Veto** – the specification this page puts into plain language
- 2 **[Foundation for Intentional Community – Founder’s Syndrome](#)**⁷ – the same pattern under the name the wider nonprofit sector uses
- 3 **[Freeman, J. – The Tyranny of Structurelessness \(1970/1972\)](#)**⁷ – delegate specific authority, for specific tasks, by an agreed procedure

Written by EcoHubs members, with AI assistance for drafting and editing, and reviewed by a person before publication. Facts are checked against the sources listed; anything we could not verify is marked.

How this is written

Conflict that is never named

The group is warm, nothing is confronted, and difficulty gets absorbed quietly – until somebody leaves and everyone privately knew why.

WHAT IT LOOKS LIKE

- Difficulties are discussed widely in private and never once in a meeting.
- Departures surprise nobody, and the reasons were never said out loud.
- The conflict process exists and has not been used in years.
- Somebody is described as “just like that”, and the group works around them.

Why it is hard to see

Because from inside it does not feel like avoidance. It feels like **kindness**, and quite often it is.

Communities are built by people who wanted somewhere warmer than the world they came from, and confrontation is what they were trying to leave behind. So the group develops a genuine, admirable skill at absorbing friction – accommodating the person who is difficult about the kitchen, working around the one who never does their rota hours, not mentioning the thing everybody noticed. Each individual instance is generous. It is only the accumulation that is fatal.

The second reason: **it never produces a bad day**. Nothing escalates, because escalation is the thing being avoided. What it produces instead is a slow narrowing – of what can be said, of who talks to whom, and of how much anybody is willing to invest in a place where the real conversation happens somewhere they are not.

And the third, which is why groups defend it: **the metric looks excellent**. No disputes, no process triggered, no incidents. A community measuring its health by absence of conflict will rate this pattern as its greatest success right up until the year three households leave.

What to check

Ask when the conflict process was last used. “Never” from a ten-year-old community of forty does not mean there has been no conflict. It means the process is unusable, and everybody worked that out privately.

Count the private conversations. Not their content – their existence. If a difficulty is known to eleven people and has never been on an agenda, the group has a route, and it runs entirely outside its own structure.

Ask a recent leaver. Communities almost never do this, and it is the single highest-yield question available: *what would you have said if there had been somewhere to say it?* People are usually willing to answer once they have gone, which is itself the finding.

If it is already happening

Start with something small enough to be safe to fail. A group with no muscle for this should not begin with its hardest case. Take a minor, genuinely two-sided difficulty and run the process badly on it. What you learn is where it breaks.

Separate raising from resolving. A great deal of avoidance is not fear of conflict but fear of *consequence* – people stay quiet because naming a thing feels like starting something they cannot stop. A process with an explicit low step, where raising a difficulty produces a conversation rather than a procedure, removes most of that.

Say the accumulation out loud, once. Not as an accusation of the culture – as an observation about cost: *we have lost four members in three years and none of us said the difficult thing before they went.* That is checkable, and it is much harder to absorb than a complaint about the atmosphere.

Expect it to feel worse before better. A community that starts naming things will have a visibly more difficult six months. It is worth telling people that in advance, because otherwise the first hard meeting gets read as proof the old way was working.

What prevents it

A repair process with a genuinely cheap first step, and the habit of using it on things too small to need it. A process only ever tested on catastrophes will not be reached for during one.

RCOS · LAYER 4

RCOS does not treat this as a cultural style a community may choose. “**Ignoring, suppressing, or normalizing unresolved conflict MUST be considered a system violation**” (§6.6.1) – the word is *normalizing*, and that is precisely what warmth does to accumulated difficulty.

The classification clauses close the other exit. A community must define an explicit conflict classification system that is **known, accessible, and usable by all members** (§6.1.1) – usable being the operative word for a process that formally exists and is never triggered. And **misclassification or avoidance of classification must be treated as a process failure subject to review** (§6.1.5): declining to name something is itself a reviewable event, not a neutral choice.

The resolution ladder must also define **how refusal, non-response, or withdrawal is handled** (§6.2.3). Which is the specific hole this pattern lives in – most communities’ informal process has no answer at all for the person who simply does not engage.

RCOS Core v0.1 – §6.6 Layer Invariants, §6.1 Conflict Classification

SOURCES & FURTHER READING

- 1 **RCOS stress test – Conflict Avoidance Normalization** – the specification this page puts into plain language
 - 2 **Ostrom, E. – Governing the Commons (1990)**[↗] – accessible conflict resolution – cheap and local – as a design principle of durable commons
 - 3 **Seeds for Change**[↗] – free practical guides on working through group conflict
-

Written by EcoHubs members, with AI assistance for drafting and editing, and reviewed by a person before publication. Facts are checked against the sources listed; anything we could not verify is marked. **How this is written**

Punishment before repair

The group does act – and reaches straight for the consequence, so nobody raises anything small ever again.

WHAT IT LOOKS LIKE

- The first available response to a breach is a sanction, with nothing cheaper below it.
- Consequences are decided in the same conversation that hears the complaint.
- Members raise things only once they are already unrecoverable.
- Sanctions have no stated end date, and nobody can say what restores standing.

Why it is hard to see

Because it usually arrives as a **correction of the opposite failure**, and it is celebrated.

A community exhausted by years of avoidance decides to take accountability seriously. It writes a policy. It stops tolerating the intolerable. Everybody agrees this is overdue, and for a while it genuinely is better – the thing that should have been addressed four years ago finally is.

What goes wrong is the **shape of what got built**. Groups correcting for avoidance tend to build a process with one step: a complaint goes in, a judgement comes out. There is nothing below it, because the thing they were fixing was the failure to act.

So the cost of raising anything becomes very high. Mentioning that somebody's dog is a problem now initiates a procedure with a verdict at the end, and no member wants to do that to a neighbour over a dog. **The result is a community that only ever processes serious matters – because the small ones are not worth what raising them costs.** Which looks, from the outside and from most of the inside, exactly like the avoidance they just escaped.

The second reason it hides: **it produces visible action**. Cases get heard, outcomes get issued, and the group can point at the file. Nothing about the metric shows that the docket is short because the threshold is high.

What to check

Look at the bottom of the ladder, not the top. What is the cheapest available response to a minor breach? If the answer is “a conversation with the board”, there is no bottom, and everything below that threshold is being absorbed silently.

Ask what restores standing. A sanction with no stated end and no route back is not a graduated response – it is an indefinite status, and the person under it will eventually leave rather than wait.

Check whether hearing and deciding are the same conversation. A group that determines the consequence in the meeting that first hears the complaint has no space in which repair could have been the outcome.

If it is already happening

Add steps below what you have, rather than softening what you have. The instinct is to make the existing process gentler, which makes serious matters harder to handle without making small ones cheaper. What is missing is rungs, not leniency.

Separate the finding from the response. Establishing what happened and deciding what follows are different tasks, and merging them is what makes every complaint feel like a prosecution.

Give every sanction an end date and a restoration route – including the ones already running. An open-ended consequence is the pattern's most expensive form, because it converts a conflict into a permanent condition that only departure resolves.

Then say the change out loud. People calibrated to a high-cost process do not recalibrate because the policy changed. They recalibrate when they watch somebody raise something small and nothing terrible happen.

What prevents it

Graduated sanctions

– a first breach not treated like a fifth – and a stated presumption that repair is attempted before consequence. Ostrom found the same thing from an entirely different direction: durable commons have graduated sanctions and cheap, local conflict resolution, and the two travel together.

RCOS · LAYER 4

RCOS states the ordering as a requirement: **“Repair-oriented actions MUST be prioritized over punitive actions except in safety-critical cases”** (§6.4.6), restated as a layer invariant – repair and restoration **must precede** punishment except where immediate safety is at risk (§6.6.3).

The clauses on sanctions themselves close the open-ended version. Sanctions and repair actions must be **proportional to the violation, explicitly documented, time-bounded where applicable, and reviewable and appealable** (§6.4.2), and the framework must state **the conditions for restoring rights, roles, or participation** (§6.4.3). A consequence with no route back does not meet that.

One clause cuts the other way and belongs on this page too: **sanctions must not be applied through informal exclusion, social pressure, silence, or implicit withdrawal of rights** (§6.4.5). A community that responds to the over-punitive phase by going quiet again has not returned to neutral – it has moved to a form of sanction that leaves no record and cannot be appealed.

RCOS Core v0.1 – §6.4 Sanctions, Repair, and Separation

SOURCES & FURTHER READING

- 1 **RCOS stress test – Punishment Before Repair** – the specification this page puts into plain language
- 2 **Ostrom, E. – Governing the Commons (1990)**⁷ – graduated sanctions: a first breach is not treated like a fifth
- 3 **Membership: getting in, and getting out** – stated grounds plus a graduated response, and what a suspension may not be

Written by EcoHubs members, with AI assistance for drafting and editing, and reviewed by a person before publication. Facts are checked against the sources listed; anything we could not verify is marked.

How this is written

Contribution that stops being shared

Some members pay their way out of the work – reasonably, one at a time – until the rota is done by whoever cannot afford not to.

WHAT IT LOOKS LIKE

- Members may buy out of work hours, and the rate was never really calculated.
- The people doing the physical work and the people with money are becoming different groups.
- Nobody can say whether buying out is a right, a concession, or a habit.
- The work still gets done, and resentment about it has stopped being expressed.

Why it is hard to see

Because the first case is obviously fair. A member with a bad back, or a demanding job, or a newborn, cannot do their four hours this month and offers to pay for a contractor instead. Refusing would be cruel and would help nobody. The group agrees, correctly.

The second case is also fair. So is the fourth. What changes is not any individual arrangement but the **composition** of who is left, and composition changes too slowly for anybody to notice the week it happens.

Because the work still gets done. This pattern produces no operational failure. Rotas are covered, buildings are maintained, and by every functional measure the community is fine. What has shifted is who is doing it, and that only becomes legible when somebody counts.

And because the rate is almost never calculated. Communities set a buy-out figure by feel – a round number that sounded reasonable in a meeting. If it is below the real cost of hiring the work, every buy-out is quietly subsidised by the people still doing it. Very few groups have ever checked, and the check takes an hour.

There is a harder version of the same mechanism, and it is worth knowing that it is documented. Ran Abramitzky's study of the Israeli kibbutzim found that **under equal sharing, high-productivity individuals are the most likely to exit** – leavers were more educated and more skilled than stayers, with at least a high-school diploma held by 61.8% of migrants against 50.0% of those who stayed. Contribution arrangements that ignore what people bring do not stay stable; they select who remains. That is a finding about a national movement of a scale nothing here resembles, so take the **mechanism** rather than the setting – but the mechanism is the point.

What to check

Count who did the physical work last quarter, and cross it against who buys out. If those are becoming different groups, the pattern is running whatever the policy says.

Price the buy-out honestly. What would it actually cost to hire this hour of work, at a fair rate, including the coordination nobody counts? If the buy-out is lower, name the gap. It is a transfer, and it is currently invisible to the people making it.

Ask what buying out is. A right, a concession granted case by case, or something that just started happening? Communities frequently cannot answer, which means the arrangement has no rules and therefore no limits.

If it is already happening

Do not abolish it. A community with no flexibility excludes exactly the members it should most want – carers, people with health conditions, people whose work is inflexible. The buy-out is not the failure.

Price it properly, and say what the money buys. If a bought-out hour funds a contractor, that is a fair exchange. If it funds general revenue while the hour is absorbed by other members, that is a transfer, and it should be one the group decided rather than one it drifted into.

Cap it, and make the cap about the community rather than the person. A limit on the proportion of total hours that may be bought out in any period keeps the arrangement available to whoever needs it while preventing it from becoming the default.

Count the invisible half too. The work that never reaches the rota – remembering the insurance, noticing the new member, absorbing the awkwardness after a difficult meeting – cannot be bought out and is distributed even less evenly. A buy-out policy that covers only rota hours quietly declares that the other work is not work.

What prevents it

A stated work quota with an honestly-priced alternative and a cap, plus a contribution definition that names the invisible work as contribution. The test is not whether people may pay instead. It is whether the community could tell you what that costs everyone else.

RCOS · LAYER 3

RCOS requires a community to state **which contribution categories it recognises**, explicitly listing labour, care and emotional work, knowledge, stewardship and administration as candidates (§5.2.1) – care and coordination named alongside physical labour rather than assumed into it.

The clause that names this pattern: a community **“MUST NOT structurally depend on unpaid, invisible, or informal labor for system survival without explicitly defining corresponding obligations, recognition, or compensation mechanisms”** (§5.2.3), restated as an invariant (§5.6.3). “Structurally depend” is the phrase – the question is not whether some work is unpaid, but whether the system would stop without it.

Substitution of participation is addressed directly in Layer 1: “**Substitution of participation (e.g., outsourcing labor) MUST be explicitly governed**” (§3.5.3). Explicitly governed, not merely permitted – an arrangement that grew by precedent does not satisfy that.

And one boundary that matters as the money side grows: contribution recognition **must not create implicit decision authority, veto power, or governance influence** beyond what Layer 2 defines (§5.2.5). Paying more does not buy a larger say.

RCOS Core v0.1 – §5.2 Contribution Recognition

SOURCES & FURTHER READING

- 1 **RCOS stress test – Outsourced Labor Bypassing Contribution** – the specification this page puts into plain language
- 2 [Abramitzky, R. – The Limits of Equality: Insights from the Israeli Kibbutz \(QJE, 2008\)](#)⁷ – under equal sharing, high-productivity members are the most likely to exit – Table VI, and the free PDF
- 3 **Daily life in community** – the work nobody put on the rota, and why “she is just good at that” is not an answer

Written by EcoHubs members, with AI assistance for drafting and editing, and reviewed by a person before publication. Facts are checked against the sources listed; anything we could not verify is marked. **How this is written**

Finances nobody can inspect

Not hidden, exactly – just held by two people, reported in summary, and never actually examined by anyone else.

WHAT IT LOOKS LIKE

- The accounts are available in principle and nobody has asked to see them in years.
- Reporting is a headline figure rather than something anybody could check.
- One or two people know the full financial picture, and both are indispensable.
- Questions about money get answered with reassurance rather than numbers.

Why it is hard to see

Because in almost every case **nothing is being concealed**.

The treasurer is honest, overworked, and would gladly explain any of it to anyone who asked. The accounts exist. They are available on request. What is missing is not access – it is **anybody exercising it**, and a format in which exercising it would tell you anything.

Three things produce that:

Summary reporting. The meeting is told the balance and whether the year was tight. Both are true and neither is checkable. A member cannot tell from a headline figure whether this year's costs were covered by income or by something that will not recur – which is the one thing they most need to know.

Concentration. Financial knowledge accumulates in whoever does the work, and it accumulates fast, because nobody else touches it. That is not a power grab. It is what happens to any task nobody shares. But it produces a community where two people can answer a question and thirty-eight cannot form one.

And the social cost of asking. In a community, requesting to inspect the accounts reads as an accusation against a volunteer who is already doing more than their share. So people do not ask, which the group experiences as trust, and which is functionally the same as having no audit.

The tell is not secrecy. It is that **nobody could construct a question**, because they have never seen the shape of the thing.

What to check

Ask who has looked at the accounts in the last year, other than the person who keeps them. If the answer is nobody, transparency is nominal regardless of policy.

Ask for a figure the summary does not contain. How much of last year's income was recurring? What is the reserve as a proportion of annual costs? A community with genuinely inspectable finances can produce these. One with a reporting habit cannot, and finding out takes a fortnight.

Check the bus factor. If the person who keeps the books stopped tomorrow, how long before somebody else could answer a question about a specific transaction? Under a week is healthy. "We would have to reconstruct it" is the finding.

If it is already happening

Do not start with an audit. Framing it as verification puts the treasurer on trial for a structural problem they did not create, and the group will spend six months on the relationship instead of the format.

Start by adding a second pair of eyes to an ordinary task. Someone else reconciles the account each quarter, alongside rather than instead. This is dull, unthreatening, and it fixes the bus factor and the inspectability problem simultaneously.

Change what gets reported, not how often. Add two lines to whatever the meeting already receives: what proportion of income was recurring, and what the reserve covers in months. Those two numbers turn a report into something a member can have an opinion about.

And say why. A treasurer who hears "we are introducing oversight" hears an accusation. A treasurer who hears "you should not be the only person who could answer this if you were ill" usually agrees, because it is the thing they have been quietly worried about.

What prevents it

Transparency as the **default state** rather than an available request, and reporting in a form that supports a question. The distinction that matters is not whether members *may* see the accounts. It is whether anybody could tell, from what they routinely receive, that something had begun to go wrong.

RCOS · LAYER 3

RCOS makes transparency the default rather than a right to be invoked: "**Transparency MUST be the default for treasury balances, inflows, outflows, obligations, and commitments**" (§5.3.4). Note the last two – obligations and commitments, not just balances. A community can be transparent about what it holds and silent about what it has promised.

Exceptions are permitted and fenced: they must be **explicitly defined, justified, time-bounded, and must not prevent members from auditing compliance** (§5.3.5). "Auditing compliance" is the phrase to hold a real community against – a summary report satisfies disclosure and not this.

Spending authority must be bounded by **clear authority assignments, thresholds by amount or category, approval and escalation paths, and mandatory recordkeeping requirements** (§5.3.3). That last clause is what makes inspection possible at all; without records there is nothing to inspect however open the policy.

The layer invariant restates it: shared resources, flows and obligations **must be visible to the community by default**, with only limited and explicit exceptions (§5.6.1).

SOURCES & FURTHER READING

- 1 **RCOS stress test – Opaque Community Finances** – the specification this page puts into plain language
- 2 **Freeman, J. – The Tyranny of Structurelessness (1970/1972)¹** – information is power – a community where two people can read the accounts has two people in charge
- 3 **Money & the community economy** – what a reserve fund is for, and the question that reveals the most

Written by EcoHubs members, with AI assistance for drafting and editing, and reviewed by a person before publication. Facts are checked against the sources listed; anything we could not verify is marked.

How this is written

The commons sold to cover the shortfall

Every sale funds another year, every sale is defensible, and the thing being sold is what made the community possible.

WHAT IT LOOKS LIKE

- An operating shortfall has been closed by a sale more than once.
- Nobody can say how many saleable assets are left.
- The agreements say nothing about which resources may not be sold.
- A sale is described as releasing a non-core asset, and nobody defines core.

Why it is hard to see

Because no individual sale is the mistake. Each one is a reasonable answer to a real problem: the roof needs doing, the programme is running at a loss, and there is a field nobody uses. Selling it is prudent. It is prudent again three years later.

The pattern is the failure, and a pattern is not something that appears on an agenda. Nobody proposes *let us fund operations from capital for three decades*. They propose selling a field, once, and then the situation recurs.

Because the alternative is worse this year. A community that refuses to sell has to cut something visible – a programme, a role, a subsidy somebody depends on – in front of the people affected. The community that sells has a slightly smaller commons and an unchanged year. Anyone who has sat in both meetings knows which is harder.

And because it is usually disclosed. This is not embezzlement or concealment. The Findhorn Foundation's audited accounts state it plainly: "*We have for over 30 years sold non-core assets to enable the Foundation's work to continue in spite of financial deficits.*" The information was public the entire time. What was missing was not disclosure. It was anybody reading thirty separate defensible decisions as one trajectory.

What to check

Calculate the run rate once. Of last year's costs, how much was covered by something that will not recur – a sale, a legacy, a grant, a draw on reserves? That single figure is the diagnosis, and most communities have never produced it.

Then count what is left, and price it. A community closing a €40,000 annual gap by selling assets, holding perhaps €120,000 of saleable land, has at most three years – fewer once the cost of selling comes out, and fewer again if any of it is restricted. The count of parcels is not the number, because they

are rarely worth the same. Stated as a figure in a meeting, that ends a conversation which has circled for a decade.

Ask what “non-core” means. It is the phrase that does the work in every one of these decisions, and it is almost never defined. If nobody can say which assets are core, the category is doing no filtering at all.

If it is already happening

Separate the two questions, because they get answered together and must not be. *Can we afford this model?* is one decision. *What should we sell?* is another. Answering the second first is exactly how thirty years passes.

Raise the tier. The cheapest structural fix available: funding an operating shortfall from capital requires a decision at the constitutional level, not an operational one. This does not forbid selling. It makes selling a decision rather than a default, which is what it stops being when it happens annually.

Consider rehoming rather than releasing. Findhorn's own accounts show the other route alongside the sales – the Sanctuary project transferred to Park Ecovillage Trust, Wikies Wood gifted to Findhorn Hinterland Trust, and discussions about *“selling non-core assets into the community.”* An asset moved into another community body is out of the organisation's balance sheet but still in community hands. It is a materially different outcome from an open-market sale, and worth pursuing before the market one.

What prevents it

A register that records, for every commons resource, **whether it may be sold at all** – and a rule that operating shortfalls are not funded from capital without a constitutional decision. A community land trust does this structurally by putting land beyond the reach of the operating entity entirely, which is the strongest version and not available to everyone.

RCOS · LAYER 3

The registry clause is written with this pattern in mind. Every governed resource must be classified as commons or private and recorded with its steward, its access rules, and **“transfer, sale, or privatization constraints (if any)”** (§5.1.2). That last field is what makes a thirty-year trajectory visible: a written sale constraint is something a decision has to consciously override.

The invariant is blunter: **“Resources declared as commons MUST NOT be privatized through informal, implicit, or unilateral action”** (§5.6.2). Repetition is a form of implicit – no single sale was unilateral, and the pattern still was.

Anything not yet classified is treated as **unclassified**, and the community **must not allocate, encumber, monetize or transfer it** until an authorised decision classifies it (§5.1.3). Which closes the most common version of this: the asset nobody had got round to categorising, sold because it was the easiest one to sell.

Alongside, the treasury must carry defined **reserve, risk and liability policies, including limits on debt and long-term obligations** (§5.3.6) – the absence of which is what makes the annual sale feel like the only option.

SOURCES & FURTHER READING

- 1 **RCOS stress test – Commons Privatization through Land Sales** – the specification this page puts into plain language
 - 2 **Findhorn Foundation – Trustees' Annual Report, year ended 31 March 2024**⁷ – audited accounts, Scottish Charity SC007233 – the thirty-year statement and the asset transfers
 - 3 **Owning, leasing, renting** – land trusts and ground leases – the structures that put land beyond an operating entity
-

Written by EcoHubs members, with AI assistance for drafting and editing, and reviewed by a person before publication. Facts are checked against the sources listed; anything we could not verify is marked.

How this is written

Removal with no process

Grounds invented during the case they are used on – and every remaining member quietly learns what the rules really are.

WHAT IT LOOKS LIKE

- The agreements state no grounds for removal, or state them without a process.
- A removal is under discussion and the procedure is being written alongside it.
- Nobody can say what the person could have done to avoid it, or could do to return.
- The decision is being made by the people the difficulty is with.

Why it is hard to see

Because by the time it happens, **everyone involved is exhausted and certain.**

This pattern almost never arrives at the start of a difficulty. It arrives at the end of a long one – after months of the situation not improving, after several attempts that went nowhere, with a group that has run out of patience and a person who has usually stopped engaging. In that state, the absence of a written procedure does not read as a problem. It reads as an obstacle to finally doing something.

So the group acts, and it is often the right outcome reached the wrong way. Which is exactly what makes it hard to see: **the decision can be correct and the process still be the failure.** Everyone can look at the result, agree it was necessary, and conclude the community handled it.

What was actually learned is learned by everybody else. A community that removed somebody without stated grounds has told every remaining member that **membership depends on staying in favour**, because that is the only rule anybody observed being applied. Nobody says this out loud. It shows up two years later as a group that has become notably careful with each other, and cannot say when that started.

The second failure is the mirror image and just as common: a community with no process **cannot act at all**, and so tolerates something genuinely intolerable for years – until the pressure produces the version above. The absence of a procedure does not produce restraint. It produces long paralysis followed by an abrupt decision.

What to check

Ask what the grounds are, before there is a case. If the agreements state none, the community has not decided that removal is impossible. It has decided that removal will be improvised.

Ask what the route back is. A removal or suspension with no stated path to restoration is a permanent outcome, whatever anyone intended. If nobody can describe how somebody regains standing, the answer is that they do not.

Check who would decide. In a small community the people with the authority to remove somebody are usually the people the difficulty is with. That is not corruption; it is arithmetic. It is also the reason an external step matters more here than anywhere else.

If it is already happening

Separate the finding from the decision. What happened, and what should follow, are different questions, and merging them is what makes an improvised removal feel like a verdict from people who had already decided.

Write the procedure for this case, transparently, and say that is what you are doing. It is not ideal and it is honest, and it is much better than acting with nothing while implying a rule existed. Say what the grounds are being taken to be, who is deciding, what the person may respond to, and what happens afterwards.

Get one voice from outside. Not to decide – to observe. A community cannot self-audit a removal in which its own leadership is a party, and an external facilitator is the cheapest available correction.

Then write the general version once it is over, deliberately not about this person, and expect that to take a while. A community that has just done this knows more about what it needs than one that has not.

What prevents it

Stated grounds **plus** a graduated response – grounds without a process is arbitrary, and a process without grounds means the group cannot act at all. Plus the two things almost always missing: a route back, and somebody involved who is not a party.

RCOS · LAYER 1

RCOS routes forced exit away from the membership rules and into the conflict layer deliberately:

“Forced exit MUST follow due process and be handled through Layer 4 mechanisms” (§3.6.3). Removal is an accountability decision, not an administrative one, and Layer 4 is where the safeguards live.

Layer 4 then supplies them. Sanctions and repair actions must be **proportional, explicitly documented, time-bounded where applicable, and reviewable and appealable** (§6.4.2), and the framework must state the **conditions for restoring rights, roles, or participation** (§6.4.3) – the route back, as a required field. Separation must follow due process and **align with the exit rules in Layer 1** (§6.4.4).

Two further clauses catch the versions that avoid procedure entirely. **Sanctions must not be applied through informal exclusion, social pressure, silence, or implicit withdrawal of rights** (§6.4.5) – the soft removal that leaves no record. And where a power differential exists between the parties, **elevated safeguards must be applied**, which may include independent or external facilitation (§6.3.3).

Underneath: exit **must not result in loss of rights beyond those explicitly tied to membership** (§3.6.4), and asset, role and responsibility separation must be **defined prior to exit** (§3.6.5) rather than negotiated during the worst week of it.

SOURCES & FURTHER READING

- 1 **RCOS stress test – Expulsion Without Due Process** – the specification this page puts into plain language
- 2 **RCOS Core – Layer 4: Conflict, Repair & Accountability** – due process, proportional sanctions, restoration conditions and elevated safeguards
- 3 **Membership: getting in, and getting out** – stated grounds plus a graduated response, and what Twin Oaks publishes

Written by EcoHubs members, with AI assistance for drafting and editing, and reviewed by a person before publication. Facts are checked against the sources listed; anything we could not verify is marked.

How this is written

Rules that spread into private life

Each new rule was a reasonable answer to a real problem, and together they have opinions about your visitors, your diet and your evenings.

WHAT IT LOOKS LIKE

- The community has an agreed position on something it does not own or maintain.
- Nobody can say what is definitively nobody else's business.
- Members ask permission for things that were never anyone's to grant.
- People stop mentioning parts of their life rather than argue about them.

Why it is hard to see

Because it is the **healthy mechanism running past its stopping point**, and there is no obvious moment where it passed.

A community that solves problems by agreement is working correctly. Somebody's music is too loud, so quiet hours get agreed. Guests cause friction, so a guest policy appears. Someone's dog, someone's parking, someone's building work. Every rule is a real answer to a real difficulty, adopted by people acting in good faith.

But **there is no natural end to the list**, because there is always another friction, and the community has learned that friction is resolved by adding a rule. Twelve years later the group has a position on how often you may have visitors and what you may keep on your own porch, and nobody can point to the meeting where it decided it had authority over those things – because it never did decide. It accreted.

The failure is invisible to the majority for a specific reason: **it is only felt by the people it constrains most**, and they tend not to argue. Somebody whose life has been quietly narrowed does not usually file an objection. They stop mentioning that part of their life, then stop doing it, then leave – and the exit interview, if there is one, says something vague about needing more space.

That is why this pattern and **the norm nobody could enforce** are the same missing document seen from opposite ends. **A community with no stated scope has no basis for acting and no basis for restraint.**

What to check

Try to name something that is definitively nobody else's business. If the group cannot produce a clear example in a couple of minutes, there is no boundary – only an absence of current interest.

Count the rules that govern something the community neither owns nor maintains. Shared land, shared buildings and shared money are obviously in scope. A rule about what happens inside someone's own home needs a reason that survives being stated out loud.

Ask the quietest member what they have stopped mentioning. This is the diagnostic that works, and it only works privately. People will describe the accommodation they have made long before they will object to the rule that caused it.

If it is already happening

Do not start by repealing. Every rule on the list has a constituency and a story about the incident that produced it, and a repeal campaign turns into twelve separate arguments about whether that incident mattered.

Write the scope statement instead, and write it as a boundary rather than a list. What the community governs – the land, the shared buildings, the money, the decisions it makes together – and what it does not. Then test the existing rules against it in one pass. Some will fall outside, and they will fall outside as a category rather than one at a time, which is a far easier conversation than twelve.

Add the default. Anything not stated as in scope is out of scope. Without it, the same accretion restarts the following year, because the mechanism that produced it – a real friction, a reasonable answer – has not changed.

And expect the person most attached to the rules to be the one most tired of the frictions. They are not being controlling. They are being the only person who keeps trying to fix things.

What prevents it

A declared scope with an explicit **out of scope** list, and the default that silence means out. The **privacy gradient** is the spatial version of the same idea – a design that lets people regulate contact without having to announce it, which removes many of the frictions before they ever become rules.

RCOS · LAYER 0

This is the pattern §2.2 exists to stop, and it is unusually direct about it.

A community must declare the scope of what it governs – assets, decision domains, activities under collective control (§2.2.1–2) – and must **explicitly list what is out of scope** (§2.2.3). Then the two clauses that do the work:

“Anything not explicitly declared as in scope MUST be treated as out of scope” (§2.2.4), and **“The community MUST NOT exercise authority over persons, assets, or domains that are declared out of scope”** (§2.2.5).

The direction of that default is the whole point. In an unwritten community, silence drifts toward *the group may have an opinion about this* – because someone raises it, and the group’s habit is to resolve things by agreeing. Under §2.2.4 silence resolves the other way, and a new rule about private life requires the community to first amend its scope, which is a constitutional decision rather than an item at the end of a meeting.

Layer 3 carries the economic equivalent: for private resources, the community **must not exercise authority beyond what is explicitly declared** in scope or membership agreements (§5.1.5).

RCOS Core v0.1 – §2.2 Scope Declaration

SOURCES & FURTHER READING

- 1 **RCOS stress test – Governance Creep into Private Life** – the specification this page puts into plain language
 - 2 **RCOS Core – Layer 0: Identity & Scope** – anything not declared in scope is out of scope – the clause that reverses the default
 - 3 **Daily life in community** – the privacy gradient, and the privacy that is hardest to arrange
-

Written by EcoHubs members, with AI assistance for drafting and editing, and reviewed by a person before publication. Facts are checked against the sources listed; anything we could not verify is marked.

How this is written

The norm nobody could enforce

A commitment written into the founding vision, and no process behind it – so the first person to break it discovers there was nothing there.

WHAT IT LOOKS LIKE

- A value is central to how the community describes itself and appears in no agreement.
- Nobody can say what would happen if somebody simply declined to observe it.
- The question is being discussed everywhere except the meeting.
- Members disagree, sincerely, about whether it was ever a rule or only a shared habit.

Why it is hard to see

Because until it is tested, an unenforceable norm and a working one **look exactly the same**.

Everybody observes it. Nobody has ever had to be asked. The community describes itself by it to visitors, puts it on the website, and selects for it informally at every stage of membership – so for years the absence of any process is invisible, because the process has never been needed.

Then somebody joins who agreed to it sincerely and later changed, or who read it as an aspiration rather than a requirement, or who simply decided differently. And the group discovers that **the thing at the centre of its identity exists nowhere in its agreements**.

What follows is worse than the original breach. Some members experience it as a betrayal of the community's purpose. Others, reading the agreements, find nothing that was broken and conclude the first group is trying to impose a preference. **Both are right**, which is what makes it unresolvable: the community is now arguing about whether it ever had a rule, in front of a live case, about a specific neighbour.

FROM A COMMUNITY WE WORKED WITH

A community with a dietary commitment written into its founding vision – the thing it was for, named in its own description of itself – found that a member had begun regularly doing the opposite, openly, outside their house.

There was **no process for behaviour that contradicted the community's own stated identity**, and no pathway by which a person could be asked to leave. So there was nothing to invoke and nothing to appeal, and the disagreement had nowhere to go but into everybody's relationships.

Note which half was missing. The community was not short of conviction about the value. It was short of the two paragraphs that would have made the conviction actionable.

What to check

Take the three things the community says it stands for, and find each one in the agreements. Not the vision statement — the part that says what follows if somebody does not observe it. Most groups find one of three, and are surprised by which.

Ask the hypothetical before you need it. *What would happen if a member simply stopped doing this?* Ask it about something nobody is currently doing, so it is a question about structure rather than about a person. The answers will vary wildly, which is the finding.

Check whether it is a condition of membership or a shared habit. Both are legitimate. A community can hold a strong culture without making it a rule — but it should know which it has chosen, because the two behave completely differently the first time somebody diverges.

If it is already happening

Separate the live case from the general rule, and resolve them in that order. The instinct is to write the policy now, and the policy will be about this person, and everyone will know.

Say out loud that there was no process. Communities try to avoid this admission, and the avoidance costs far more than the admission would: members watching a rule apparently materialise lose confidence in every other rule. “We have no process for this, we are going to make one for this case, transparently, and review it afterwards” is a legitimate and repairable position.

Then decide the harder question honestly: is this a condition of membership, or a strong preference the community holds? If it is a condition, it needs a process, and it needs to be visible to anyone joining. If it is a preference, the group has to accept that members may diverge — and say so, which is often the real conversation nobody wanted to have.

What prevents it

Identity constraints written as **testable statements with a process attached** rather than as values. The test: could somebody who did not share the value tell, from the agreements, exactly what was required of them, and what would happen if they declined?

RCOS · LAYER 0

RCOS treats this as a Layer 0 matter — constitutional, alongside purpose and scope, rather than a house rule. A community **must declare any identity-level constraints that materially affect participation, behaviour or governance** (§2.4.1), which may include ethical or behavioural boundaries and non-negotiable cultural or ecological constraints (§2.4.2).

Then the two clauses this pattern breaks. Identity constraints **must be testable and enforceable through defined processes** (§2.4.3) — testable being the word that separates a constraint from a value. And they **must not be enforced implicitly or informally** (§2.4.4), which rules out the usual response of social pressure and cooling relations.

Read together they leave a community exactly two honest options: write the constraint with a process, or accept that it is not a constraint. **The one thing not available is holding it as central and unenforceable**, which is where this pattern lives.

Layer 1 supplies the other half – a Membership Agreement is a mandatory artifact (§3.8.1), and admission criteria must be explicit and documented (§3.2.3). A constraint that a joiner could not have read is not one they agreed to.

RCOS Core v0.1 – §2.4 Identity Constraints

SOURCES & FURTHER READING

- 1 **RCOS stress test – Cultural Norm Violation Scope Test** – the specification this page puts into plain language
- 2 **RCOS Core – Layer 0: Identity & Scope** – identity constraints must be testable, and must not be enforced informally
- 3 **How this is written** – how we use what we learn from communities we work with, and what we never publish

Written by EcoHubs members, with AI assistance for drafting and editing, and reviewed by a person before publication. Facts are checked against the sources listed; anything we could not verify is marked. **How this is written**

The purpose that quietly changed

Nobody voted to become something else — enough people arrived wanting something else that it happened anyway.

WHAT IT LOOKS LIKE

- Longer-standing and newer members describe what the place is for in different words.
- Proposals get argued as practical questions that are really about direction.
- The founding vision is quoted mainly by people losing an argument.
- Nobody can say when the purpose last changed, or whether it did.

Why it is hard to see

Because it happens through **admissions, not decisions**.

No meeting resolves to change what the community is for. What happens is that each new member is chosen individually, by people applying their own reading of the purpose, and each choice is defensible. Over enough intakes the median understanding moves — and it moves without any moment where somebody could have objected, because there was no proposal to object to.

The second reason: **the argument never presents as an argument about purpose**. It arrives as a dispute about the barn, the budget, whether to take paying guests, how much noise is acceptable at nine on a Sunday. Both sides argue the practicalities, both are frustrated that the other seems unreasonable about something minor, and neither says the actual sentence — *we are no longer here for the same thing*.

And the third, which is why it is so hard to raise: **saying it out loud is a serious accusation**. Suggesting that the place has drifted implies that some current members do not really belong, which is unsayable in a group that likes each other. So it goes unsaid until it is said in anger.

What to check

Ask ten members to write down what the community is for, in one sentence, separately. Not a discussion — separately, on paper. The variance is the diagnosis, and it is a much kinder instrument than a conversation, because nobody has to accuse anybody.

Then sort the answers by how long people have been there. Drift shows up as a gradient. If the sentences cluster by intake year, the purpose has been moving and the admissions process has been moving it.

Look at what gets quoted. A founding vision invoked mainly by people who are losing an argument has already stopped being the shared frame and become a faction's evidence.

If it is already happening

Do not open with “we have drifted”. That is heard as *the newer members are the problem*, and it will be resisted for exactly that reason – often correctly, since the newer members did not do anything except be admitted by the existing ones.

Run the sentence exercise and publish the spread. Anonymised, unattributed, all of it. It converts a suspicion into a shared observation, and the group can look at it together rather than at each other.

Then choose deliberately, and accept there are only three options. Reaffirm the original purpose and accept that it now excludes some current members’ hopes. Adopt the changed purpose formally and let the people who came for the original one know. Or state two purposes with an explicit priority between them. All three are survivable. **What is not survivable is holding all of them informally**, which is the current position.

Fix the intake. Whatever is decided, the admissions process is what produced this and will produce it again. If applicants cannot read the purpose and self-select against it, the next drift is already under way.

What prevents it

A **single stated primary purpose** that is hard to change, plus an admissions process that shows it to applicants and asks them to agree to it explicitly rather than absorb it. Drift is not prevented by conviction. It is prevented by the purpose being a document that a decision has to amend.

That is also why the purpose has to be separated from the plan. Where one founding statement carries the enduring reason, the picture of the future and the current three-year priorities together, revising the priorities means editing the purpose – a little at a time, reasonably, until it has moved. **Vision, purpose and mission** sets out the three and which tier may change each.

RCOS · LAYER 0

RCOS requires **exactly one primary purpose**, describing the enduring reason the community exists rather than a goal, project or strategy (§2.1.1–2). Secondary purposes are permitted and **must not conflict with or override it** (§2.1.4), which is the structure this pattern lacks: competing purposes with no stated ranking between them.

Two clauses do the anti-drift work. The primary purpose **must be stable across time and may only be changed through a constitutional decision** executed via the Layer 6 change process (§2.1.3) – so a purpose can change, and changing it is an event with a record, rather than a thing that happens through intake. And: **“No action, decision, or allocation of resources MAY materially contradict the stated primary purpose”** (§2.1.5), which gives any member a basis for raising drift that is not an accusation about anybody’s belonging.

Layer 1 closes the intake route. Onboarding must include **review of all RCOS-Core artifacts and explicit consent to Layer 0 and Layer 1 rules** (§3.2.2) – explicit consent, so a member’s understanding of the purpose is a recorded agreement rather than an impression they formed.

RCOS Core v0.1 – §2.1 Purpose Definition

SOURCES & FURTHER READING

- 1 **RCOS stress test – Mission Drift Through Competing Purposes** – the specification this page puts into plain language
 - 2 **RCOS Core – Layer 0: Identity & Scope** – exactly one primary purpose, changeable only by constitutional decision
 - 3 **Resilience & why communities fail** – drift as one of the six recurring patterns, and why versioning is what catches it
-

Written by EcoHubs members, with AI assistance for drafting and editing, and reviewed by a person before publication. Facts are checked against the sources listed; anything we could not verify is marked. **How this is written**

Care work nobody counts

The noticing, remembering and smoothing that never reaches a rota – done by the same people, unrecorded, until one of them stops.

WHAT IT LOOKS LIKE

- The rota covers physical tasks and nothing else, and everyone knows other work exists.
- The same people follow up after difficult meetings, and it is not anybody's role.
- Somebody is described as good at people, and that is the whole arrangement.
- A member's departure would leave a gap nobody could name in advance.

Why it is hard to see

Because it is **defined by not being visible**. That is not a figure of speech – the work consists largely of things that do not happen: the new member who does not drift away, the argument that does not escalate, the person who is not left out at the third dinner in a row.

Successful care work produces an absence of events. A community measuring its health by what went wrong will conclude that nothing went wrong, which is both true and entirely the point.

It also has no natural unit. A rota can count two hours of grounds maintenance. It cannot easily count noticing that somebody has been quiet for a fortnight and finding a reason to knock on their door. So when a community writes down its contribution expectations, this work is not omitted out of contempt – it is omitted because nobody could think how to write it.

And the distribution is not random. Care work concentrates along the same lines it does everywhere else, which means a community can be genuinely committed to equality and still run on an entirely conventional division of labour, without anybody intending it or noticing.

The tell is a specific sentence: **“she is just good at that.”** A description of a person, standing in for a description of a job.

What to check

List what would stop if one specific person left. Not their rota hours – the other things. Communities find this exercise uncomfortable and clarifying in equal measure, because the list is always longer than the person's formal contribution and usually longer than they would claim.

Look at who follows up. After a difficult meeting, somebody checks on the person who had the hard time. Who? Every time? Is it in any document?

Count the categories, not the hours. You will not successfully time-track emotional labour and you should not try. But you can name it: onboarding, checking in, following up, remembering occasions, holding the social calendar. Five named categories is enough to see a distribution.

If it is already happening

Name it before you distribute it. The naming is most of the intervention. Work that has a name can be discussed, assigned and declined; work that is “just how she is” cannot.

Do not convert it into a rota. Care work assigned by rota becomes hollow – a scheduled check-in is not the thing. What works is naming the responsibility and its holder for a term, then leaving how to the holder: *this quarter, Ana holds new-member integration*.

Give it a ceiling like any other role. The person doing this is usually also doing three other things, because the disposition that produces it is the disposition that says yes.

And expect to hear that it does not need naming. Often from the person carrying it, who has an identity invested in it being effortless. Worth answering honestly: the point is not that you find it hard. The point is that if you stop, nobody will know what stopped.

What prevents it

A contribution definition that **lists care and coordination as contribution**, and roles that cover it with terms and holders. The test: could a new member read the agreements and know that this work exists, that it counts, and who is doing it this year?

RCOS · LAYER 5

RCOS opens the workload section by naming what most communities' rotas leave out: **“Time, attention, coordination capacity, and emotional labor MUST be treated as finite and limited resources”** (§7.4.1). Emotional labour, in the same sentence as time, as something the system spends.

From there the requirements follow the resource. Explicit workload boundaries covering meeting load, role load and expectations, with a named mechanism for **relief, substitution or redistribution** (§7.4.2), reviewable through an authorised process (§7.4.3). And the trigger: **persistent overload, burnout risk, chronic non-participation, or dependency on over-functioning individuals MUST trigger review or repair processes** (§7.4.4) – *dependency on over-functioning individuals* being this pattern named as a condition to detect rather than a personality to appreciate.

Layer 3 supplies the recognition half. Contribution categories that a community may recognise explicitly include **care and emotional work** alongside labour and stewardship (§5.2.1), and a community **must not structurally depend on unpaid, invisible or informal labour for system survival** without defining corresponding obligations, recognition or compensation (§5.2.3, restated as an invariant at §5.6.3).

Note the shape of that last requirement. It does not say the work must be paid, or shared equally. It says the community may not **depend** on it while leaving it undefined.

RCOS Core v0.1 – §7.4 Workload and Capacity Boundaries

SOURCES & FURTHER READING

- 1 **RCOS stress test – Invisible Care Labor Burnout** – the specification this page puts into plain language
 - 2 **RCOS Core – Layer 3: Economic & Resource System** – contribution categories including care and emotional work, and the dependency clause
 - 3 **Invisible work** – the definition, and how a person ends up doing it for eleven years
-

Written by EcoHubs members, with AI assistance for drafting and editing, and reviewed by a person before publication. Facts are checked against the sources listed; anything we could not verify is marked.

How this is written

Growth that outpaces onboarding

New members arrive faster than the community can teach itself to them – so they learn a version, and the version drifts.

WHAT IT LOOKS LIKE

- Onboarding is a person being generous with their time rather than a process.
- Recent joiners describe the agreements differently from long-standing members.
- The same questions get answered from scratch for every arrival.
- Nobody can say what a new member is supposed to have read.

Why it is hard to see

Because growth is the thing the community wanted, and it arrives looking like success.

More members means more hands, more money, more life. A group that has spent three years short-handed does not experience a wave of arrivals as a risk. And each individual arrival is handled well – somebody shows them round, answers their questions, and makes them welcome, because communities are good at welcome.

What does not scale is transmission. A community's actual working knowledge – why the rule exists, what was tried in 2019, which conversations belong in which forum, what the agreements mean in practice – lives mostly in people, and it is transmitted by conversation. Conversation scales linearly with the effort of whoever is doing it and not at all with the number of arrivals.

So past a certain rate, three things happen at once. **The transmitters burn out**, because welcoming is invisible work and it lands on the same few. **New members learn a partial version**, assembled from whoever happened to explain what. And because they learn different partial versions, **the group's shared understanding fragments** – not into disagreement, which would be visible, but into a set of sincere and slightly different accounts of what this place is.

That fragmentation is the real damage, and it surfaces years later as **mission drift** or as an argument in which both sides are quoting the community accurately.

What to check

Ask three members from different intakes to describe the same rule. Not the vision – something operational: what the guest policy is, or how spending gets approved. The variance tells you what transmission has been doing.

Ask what a new member is supposed to have read. If there is no list, onboarding is improvisation, however warm it is.

Count who does the welcoming. If it is the same one or two people every time, growth is being funded by their capacity, and their capacity is finite in a way the intake plan probably does not account for.

If it is already happening

Slow the intake before you fix the process. This is the unpopular one and it is usually correct. A community cannot build an onboarding process while running at the rate that broke the old one, and admitting fewer people for two intakes is far cheaper than the alternative.

Write the read-this list. Not a handbook — a list. The agreements, the decision matrix, the last two years of decisions, and the three things that always confuse people. Most communities can assemble this in an afternoon and have never done it.

Move transmission from conversation to document, then use conversation for what it is good at. The facts belong in writing; the welcome, the context and the answering of what-does-this-actually-mean belong in a person. Communities usually have this exactly the wrong way round — a warm person explaining facts that should have been written, and no structured conversation about the things a document cannot convey.

Then check the older members. After a fast intake it is common to find that long-standing members are also carrying an out-of-date version, because the agreements changed and nobody re-taught them.

What prevents it

An **onboarding process with a stated content list** — what a new member reads, what they are asked to agree to explicitly, and who holds the role of taking them through it. Plus an honest intake rate: how many people the community can actually absorb per year, decided in advance rather than discovered afterwards.

RCOS · LAYER 5

The operations clause that bites here is about tacit knowledge: **“Critical operational processes MUST be documented such that continuity does not depend on tacit knowledge held by specific individuals”** (§7.3.4). Onboarding by conversation is exactly continuity depending on tacit knowledge, and it fails at scale before it fails at all.

Layer 1 supplies the requirement most communities are missing. Entry **must follow an explicit onboarding process** (§3.2.1), and that process must include **review of all RCOS-Core artifacts, explicit consent to Layer 0 and Layer 1 rules, and declaration of the initial membership state** (§3.2.2). Review and explicit consent — a recorded act rather than an impression formed over a few dinners.

And the clause that makes drift-by-intake a compliance failure rather than an accident: **“Informal, implicit, or retroactive membership MUST NOT be permitted”** (§3.2.4).

On the capacity side, Layer 5 requires workload boundaries with a mechanism for **relief, substitution or redistribution** (§7.4.2), and treats **dependency on over-functioning individuals** as a trigger for review (§7.4.4) — which is what an intake funded by two people's goodwill produces.

RCOS also sets scale expectations directly: an optimal horizontal group is roughly 8–40, and above about 120–150 sub-structures become mandatory. Growth is not neutral to structure.

SOURCES & FURTHER READING

- 1 **RCOS stress test – Rapid Growth without Onboarding** – the specification this page puts into plain language
- 2 **RCOS Core – Layer 1: Membership System** – explicit onboarding, review of artifacts, and no implicit membership
- 3 **Membership & belonging** – the four stages, and what each one is actually for

Written by EcoHubs members, with AI assistance for drafting and editing, and reviewed by a person before publication. Facts are checked against the sources listed; anything we could not verify is marked.

How this is written

Information nobody can reach

The accounts, the passwords, the history and the reason the rule exists – held by whoever happened to be there, and reachable only through them.

WHAT IT LOOKS LIKE

- Answering an ordinary question requires asking a specific person.
- Nobody can say where decisions from four years ago are recorded, or whether they are.
- Access to accounts, keys or systems has never been written down as a list.
- New members take a long time to become useful, and nobody knows why.

Why it is hard to see

Because it costs nothing until it costs everything.

For as long as the people who know are present and willing, the community has excellent access to information. Ask Marta; Marta knows. The system works, it is fast, and it is more pleasant than a document management policy. Nobody experiences a problem.

The failure is entirely latent. It converts into a real cost only at three moments: when somebody leaves, when somebody is ill, and when a disagreement turns on what was actually decided in 2019. Two of those are unpredictable and the third arrives during a conflict, which is the worst possible time to discover that the record does not exist.

There is a second, quieter cost that runs the whole time. **Information is the substrate of participation.** A member who cannot find out what the reserve is, or why the guest rule exists, or what was agreed about the barn, cannot form a view worth contributing. They are not excluded – everything is available if they ask the right person the right question. But asking requires knowing that there is something to ask about, which is precisely what the missing record would have told them.

So the community slowly sorts into people who can participate fully and people who participate on the terms of what they happen to have been told, and reads the difference as engagement.

What to check

Ask an ordinary question without asking a person. Something a member might genuinely want to know: what did we agree about guests, what did the roof cost, when is the insurance due. If the honest answer is “ask X”, the record is a person.

Try to find a decision from four years ago. Not the outcome – the decision: what was agreed, who agreed it, and under what rule. Communities routinely discover that the answer is somewhere in a chat thread above four hundred messages about a dishwasher.

Write down every access anybody holds – accounts, keys, systems, the alarm code, the domain registrar. The exercise takes an hour and the list is always longer than expected, and always concentrated in fewer hands than expected.

If it is already happening

Start with access, not archives. Reconstructing a decision history is a large project that never finishes. Listing who can get into what is a short project that removes most of the acute risk, and it can be done this month.

Second holder for everything critical. Not a policy – a name, per item. This is the cheapest single intervention in this lesson and it survives the person leaving, being ill, or falling out with the community.

Then fix the flow forward before you fix the backlog. Decide where decisions get recorded from now on, and what a record contains: what was decided, by whom, under what rule, and when it takes effect. A community with two clean years of records and a murky past is in a far better position than one that spent six months excavating and then stopped.

And publish what access rules actually are. Not everything should be open – a conflict record has privacy boundaries for good reason. But *which* things are restricted, and to whom, should itself be public. An undefined restriction is indistinguishable from a person's discretion.

What prevents it

Documentation rules that state **what must be recorded, where it lives, and who may reach it** – plus the discipline that no critical access has one holder. The test is not whether the community has records. It is whether somebody who joined last year could answer an ordinary question without asking a founder.

RCOS · LAYER 5

RCOS requires explicit documentation rules covering decisions, roles, operations and shared obligations (§7.3.1), specifying at minimum **what must be recorded, where records are stored, who has access to which records, publication or notification timelines, and privacy boundaries with the conditions for restricted access** (§7.3.2).

Note that access rules are a required field rather than an implied practice – which is what turns “some things are sensitive” from a person's judgement into a stated rule.

Every decision must be traceable to its **type and domain, the authorised role or body, the mechanism and threshold, and the recorded outcome and effective date** (§7.3.3). That is the shape of a record that can still answer a question in 2029.

Two further clauses target this pattern directly. “Critical operational processes **MUST be documented such that continuity does not depend on tacit knowledge held by specific individuals**” (§7.3.4) – tacit knowledge named as the risk. And “Information flow **MUST be designed to prevent gatekeeping, bottlenecks, or dependency on informal intermediaries**” (§7.3.5), where *informal intermediary* is the polite term for ask Marta.

The layer invariants restate both: critical processes must not rely solely on individual memory, goodwill or informal transmission (§7.7.2), and **information access rules must be explicit and enforceable** (§7.7.4).

SOURCES & FURTHER READING

- 1 **RCOS stress test – Undefined Information Access** – the specification this page puts into plain language
- 2 **Freeman, J. – The Tyranny of Structurelessness (1970/1972)[↗]** – information is power – spread it as widely as possible
- 3 **Proposals & ratification** – what a closed decision has to answer to somebody who was not there, three years later

Written by EcoHubs members, with AI assistance for drafting and editing, and reviewed by a person before publication. Facts are checked against the sources listed; anything we could not verify is marked.

How this is written

Responsibility that becomes authority

The person who does a thing ends up deciding it – nobody granted that, and nobody can take it back without taking the work back too.

WHAT IT LOOKS LIKE

- Decisions in an area are made by whoever handles that area, without anybody deciding they should be.
- Members say one person has too much, and have said so for a long time.
- Nobody could take a task over without a substantial handover that has never been written.
- The overloaded person would like help and cannot say what help would look like.

Why it is hard to see

Because **the authority is a side effect of the work, and the work is a gift.**

Somebody takes on the maintenance. Because they do it, they know the state of the buildings, which contractor is reliable, and what last year's repair cost. So when a maintenance question arises, asking them is not deference – it is efficiency. Their answer is better informed than anybody else's would be.

Over time the community stops distinguishing between *asking the person who knows* and *the person who decides*. Nobody granted authority. It accreted out of competence, and it is genuinely difficult to point at the moment it stopped being consultation.

The second reason it persists is more uncomfortable. **Removing the authority means removing the work**, and the work is the part nobody else wants. A community that says “you should not be deciding this alone” without also saying “and here is who will do it instead” has asked someone to keep the burden and give up the compensating say. That is worse than the current arrangement, and the person will reasonably refuse.

And the third: **the holder often experiences the opposite**. Ask them and they will say, sincerely, that they wish other people would take some of it on. They are not defending a position. They are stuck in one.

What to check

Take one domain and separate two questions: who does the work, and who decides what happens. If the answer is the same name and no document says it should be, this is running.

Ask what a handover would require. A responsibility that can be handed over in a briefing is a task. One that would take three months and a lot of undocumented context is a dependency, and the community is exposed to it whether or not anyone is unhappy.

Then ask the holder what help would look like. The revealing part is not whether they want help – almost everyone does. It is whether they can specify it. “I don’t know, some of the admin” is the diagnosis: the work has never been broken into pieces, so it cannot be shared.

If it is already happening

Somebody who is not the holder has to run the redistribution. This is the step communities skip, and skipping it is why the attempt fails around week three. Asking the overloaded person to organise their own unloading is one more task for the person with the least room.

Enumerate before you delegate. “Maintenance” cannot be handed over. “Getting three quotes when a repair exceeds €500”, “the annual boiler service”, and “keeping the contractor list” can be – three named things with three holders and three handover dates.

Split the deciding from the doing, explicitly. For each piece: who does it, and what may they decide alone? A spending limit is usually enough. It converts an informal authority into a stated one of roughly the same size, which is far less threatening than it sounds and is the actual fix.

Expect the holder to be relieved and also strange about it. Somebody who has carried a thing for years has their competence bound up in it, and handing it over is a loss as well as a relief. Naming that out loud makes the transition much easier than pretending it is purely administrative.

What prevents it

Roles that carry **a stated decision authority and a term**, so that doing the work and deciding it are separately assigned from the start – and a handover requirement that forces the undocumented context out of one person’s head while there is no crisis.

RCOS · LAYER 5

RCOS separates the two things this pattern fuses. Every role in the Role Registry records its **purpose, its scope of responsibility and its decision authority, and its explicit boundaries and interfaces with other roles** (§7.1.2) – so “what they do” and “what they may decide” are different fields, filled in separately, rather than one growing into the other.

Each role must also carry an accountability mechanism defining how underperformance, **overload** or role failure is handled, and **how handover and knowledge transfer occur** (§7.1.3). Handover as a required field, written before it is needed.

The continuity clause is the one to quote in a meeting: “**The community MUST ensure that no single individual is a critical single point of failure for core operations**” (§7.5.1), supported by documented procedures, clear handover mechanisms and backup arrangements where feasible (§7.5.2), reviewed periodically (§7.5.3).

And from Layer 2, the boundary this pattern crosses without anybody noticing: **authority must not be derived from charisma, seniority, ownership, or informal influence** (§4.3.4). Competence is not on that list – but authority that arrives *through* doing the work, rather than through assignment, is informal influence whatever produced it.

SOURCES & FURTHER READING

- 1 **RCOS stress test – Invisible Power via Responsibilities** – the specification this page puts into plain language
- 2 **Freeman, J. – The Tyranny of Structurelessness (1970/1972)**[↗] – delegate specific authority to specific people for specific tasks, and rotate them
- 3 **Founder transition** – the same conversion problem, at the scale of a whole community

Written by EcoHubs members, with AI assistance for drafting and editing, and reviewed by a person before publication. Facts are checked against the sources listed; anything we could not verify is marked.

How this is written

Self-sufficiency without contribution

Households withdraw into their own lives, one reasonable decision at a time, and the commons is maintained by whoever still minds.

WHAT IT LOOKS LIKE

- Attendance at shared work and shared meals has fallen without any decision being made.
- Some households are effectively neighbours who share a driveway, and nobody has said so.
- The people who still turn up are increasingly the same people.
- Non-participation has never been raised with anybody, because nobody wants to be the one.

Why it is hard to see

Because **every step of it is legitimate**, and most of it is nobody's fault.

A member takes a demanding job. A household has a baby. Somebody's health changes. Someone simply discovers, after four years, that they wanted a nice place to live rather than a project. Each withdrawal is reasonable, each is temporary in intention, and asking about any one of them would be intrusive.

What makes it invisible is that **there is no event**. Nobody resigns from community life; they are just busy this month, and next month. And because the community is warm and does not want to police anybody, nothing is said – which is experienced by the withdrawing household as acceptance and by the remaining members as a slowly increasing load.

The second reason: **the commons does not visibly degrade at first**. The paths stay clear and the common house stays warm, because the people who mind pick up what was dropped. The failure is not that things stop happening. It is that the set of people making them happen shrinks, and shrinking is invisible until it is acute.

And the third, which makes it particularly hard to raise: **the withdrawn members are often perfectly happy**, and would tell you the community is working well. From where they stand, it is.

What to check

Count attendance over three years, not three months. Shared meals, work days, meetings. A trend is a fact; an impression is an accusation, and only one of those can be discussed calmly.

Then count the people, not the events. How many distinct households did any shared work last year? That number falling while the events continue is the diagnosis.

Ask whether anyone has ever been asked. Most communities discover that non-participation has never once been raised with anybody – not as a rebuke, not even as a question. What looks like tolerance is often the absence of a way to ask.

If it is already happening

Ask before you enforce. A community that responds to this with a work quota and penalties will find the quota does not reach the people it was aimed at and irritates the people already contributing. The first move is a conversation with each withdrawn household, and it should genuinely be a question – some will have a reason the community can accommodate, and some will say honestly that they want less involvement than they signed up for.

Then make the range explicit. The most useful outcome is often not more participation but **a legitimate lower tier**: a way to be a member with reduced involvement, stated, with the difference acknowledged. Communities resist this because it feels like giving up on the vision. The alternative is a de facto lower tier that nobody agreed to, which is what already exists.

Write the minimum that is actually required, and make it small enough to be true. A stated four hours a month that people do is worth more than an aspirational eight that half the community ignores, because the second one teaches everybody that the agreements are decorative.

And check what withdrew first. Households usually leave the meetings before they leave the work, and leave the work before they leave the meals. Whichever went first is where the community became unrewarding, which is more useful to know than the participation numbers.

What prevents it

A stated participation expectation with a **real** minimum, a legitimate reduced-involvement option, and someone whose role it is to notice – because this pattern is caught by somebody looking, and by nothing else.

RCOS · LAYER 5

RCOS treats both ends of this as the same requirement. **Persistent overload, burnout risk, chronic non-participation, or dependency on over-functioning individuals MUST trigger review or repair processes** as defined in Layer 4 (§7.4.4) – non-participation and over-functioning in one clause, because they are one situation seen from two sides.

Layer 1 requires the expectation to exist in the first place: **participation expectations must be explicitly defined** (§3.5.1), **acceptable forms of contribution must be listed** (§3.5.2), and substitution of participation must be explicitly governed (§3.5.3). Then: **“Persistent non-participation MUST trigger an accountability process as defined in Layer 4”** (§3.5.4).

Note what that routes it to. Not a sanction – an accountability *process*, which under Layer 4 must prioritise repair over punishment except in safety-critical cases (§6.4.6). The standard's answer to a member who has quietly withdrawn is a conversation with a defined shape, not a penalty.

Layer 3 adds the constraint that stops the community solving it by absorption: it **must not structurally depend on unpaid, invisible or informal labour for system survival** without defining the corresponding obligations or recognition (§5.2.3). Which is exactly what “whoever still minds picks it up” is.

RCOS Core v0.1 – §7.4 Workload and Capacity Boundaries

SOURCES & FURTHER READING

- 1 **RCOS stress test – Self-Sufficiency Without Collective Contribution** – the specification this page puts into plain language
- 2 **RCOS Core – Layer 1: Membership System** – participation expectations, acceptable contribution, and where non-participation goes
- 3 **Work quota** – how communities count contribution, and how strictly

Written by EcoHubs members, with AI assistance for drafting and editing, and reviewed by a person before publication. Facts are checked against the sources listed; anything we could not verify is marked. **How this is written**

A community that forgot why

The rule survives and the reason does not — so nobody can tell whether it still earns its place, or which version is even in force.

WHAT IT LOOKS LIKE

- A rule is followed and nobody present can say what it was for.
- Two versions of a document exist and nobody can say which one is in force.
- Decisions are recorded as outcomes, without who decided them or under what rule.
- The people who could explain the history are the people who left.

Why it is hard to see

Because **nothing is missing until somebody asks**.

The rule is on the page. People follow it. The community functions. It is only when a member proposes changing it — or a new arrival asks why — that the group discovers it holds an artefact whose purpose has gone. And the discovery usually comes with an argument, because in the absence of a stated reason everybody supplies their own, and theirs conveniently supports what they already wanted.

The second reason: **the loss is gradual and load-bearing at the same time**. A rule with a forgotten reason cannot be evaluated. It cannot be defended against a bad-faith challenge or repealed when it has genuinely expired, because both require knowing what it was for. So communities keep everything, and the agreements slowly become an accumulation rather than a design.

Then the sharper version, which is what turns forgetting into a live problem: **the record does not say which version is in force**. A document that has been revised without an adoption record is not one document with a history. It is two documents with a claim, and every rule downstream of it inherits the ambiguity.

What to check

Pick three rules and ask what each was for. Not what it does — what happened that made somebody write it. If nobody in the room can answer for two of them, the reasons are gone and the rules are running unexamined.

Ask which version is current. Then ask how you would know. “Whichever is in the shared folder” is not an answer if the folder contains two.

Look at a decision record from four years ago. Does it say who decided, under what rule, and from when? Or only what was decided? The second is a note, not a record, and it cannot settle a dispute.

Ask what the community learned from its hardest year. If the answer lives in three people's memories, it will leave when they do.

If it is already happening

Fix the present before the past. Establish which documents are in force today, agreed in one meeting. That is achievable; reconstructing a decade is not, and communities that attempt the second usually abandon both.

Add the two fields that would have prevented it: an adoption record on every document – who agreed it, when, under which rule – and a summary of what changed from the previous version. Two fields, and they are the difference between a version history and a pile.

Keep the superseded text. Deleting the old version destroys the evidence of what the community used to think, which is exactly what you need when somebody asks why.

And write down the reasons while the people are still here. A short note against each significant rule – *this exists because of what happened in 2019* – is the cheapest institutional memory available, and it can be done in an afternoon by whoever has been there longest. It is also the single most valuable thing a departing long-standing member can leave behind, and almost nobody asks them for it.

What prevents it

A **version history** that makes “which version is in force” answerable, and a **learning log** that records what a difficult year taught. Both are clerical, and this is the one failure class where the defence is entirely clerical.

RCOS · LAYER 6

“No informal, undocumented, or ‘understood’ rule changes MAY be considered valid” (§8.2.5). Which resolves the two-documents case directly: an unrecorded revision is not a competing version, it is not a version at all.

The Version History must record, at minimum, the **version identifier, adoption and effective dates, the decision record reference (authority, mechanism, threshold), a summary of changes, and migration notes** (§8.2.2). Note *summary of changes* – the diff, as a required field. And at any point the community must be able to determine unambiguously **which version is in force and which artifacts are authoritative** (§8.2.3), with superseded rules retained **together with the dates during which they were in force** (§8.2.4).

The Learning Log is a separate mandatory artifact (§8.6.1). Major failures, adaptations, reversals and systemic learnings must be documented (§8.4.1), capturing **what occurred and why it mattered, which layers or rules were implicated, what was changed or stopped, and what signals or thresholds triggered action** (§8.4.2).

Two clauses give it teeth. **Repeated failure patterns must trigger structural review rather than individual blame** (§8.4.4) – which requires a record long enough to see repetition in. And the invariant: **major failures and adaptations must be captured as shared learning, not erased or hidden** (§8.7.4).

SOURCES & FURTHER READING

- 1 **RCOS stress test – Institutional Amnesia** – the specification this page puts into plain language
 - 2 **RCOS Core – Layer 6: Evolution & Adaptation** – version history, the learning log, and what makes a rule change valid
 - 3 **Drift** – the slow version, and why it does not feel like anything while it happens
-

Written by EcoHubs members, with AI assistance for drafting and editing, and reviewed by a person before publication. Facts are checked against the sources listed; anything we could not verify is marked.

How this is written

Experiments with no way back

A trial that quietly became the arrangement – because nobody wrote down when it ended or what would count as it failing.

WHAT IT LOOKS LIKE

- Something introduced as a trial is still running, and nobody set a review date.
- Nobody can say what would have counted as the experiment failing.
- Reverting would now cost more than the change did, and that argument is winning.
- Members disagree about whether it was ever agreed permanently.

Why it is hard to see

Because **calling it a trial is what got it adopted**, and everybody meant it at the time.

“Let’s try it for six months” is the single most useful sentence in community governance. It lowers the stakes, it converts an argument about principle into a question of evidence, and it lets a group act while it is still uncertain. Communities that cannot say it get stuck; communities that can, move.

The failure is not in the trial. It is that **the sentence contains a promise nobody wrote down** – six months, and then what? Reviewed by whom, against what, and reverting to what if it has not worked?

Without those, the experiment does not end. It **accretes into the arrangement**, and by the time anybody raises it the situation has changed underneath: people have rearranged their lives around it, a rota has been rebuilt, somebody has bought equipment. Reverting is now a real cost, and *reverting is expensive* becomes the argument that keeps it – an argument that would have carried no weight in month one and is decisive in month twenty.

The second problem is subtler and worse. **An experiment with no failure criteria cannot fail**. Whatever happens, somebody can read the outcome as success, because there was never a stated definition to check it against. So the community has run the trial, learned nothing testable from it, and kept the change anyway.

What to check

List what is currently “being tried”. Then, for each: what date was set, what would count as it working, and who is reviewing it. Communities usually find several items on the list and almost no answers to the three questions.

Ask what reverting would cost now versus at the start. If the gap is large, the trial has already become permanent regardless of what anyone intends.

Ask what was learned. A trial that produced no answer to the question it was meant to settle was not an experiment. It was a change with softer framing.

If it is already happening

Do not extend by default – decide. For each live trial: adopt it properly, end it, or set a real review with a date and criteria. Any of the three is fine. Continuing without choosing is the only bad option, and it is the one that happens automatically.

Write the review before you write the trial, from now on. Four fields, and they take five minutes: what is changing, until when, what would count as it working, and what we return to if it does not. A community that cannot answer the third has not designed an experiment.

Prefer the reversible version. Where there is a choice between a trial that can be undone and one that cannot, the reversible one is worth real cost – because the value of a trial is the option to stop, and an irreversible trial has spent that option before it starts.

And be honest about which changes cannot be trialled at all. Selling an asset, changing a legal structure, admitting a member: these are not experiments regardless of how they are framed, and framing them as trials to lower the bar is the version of this pattern that does lasting damage.

What prevents it

Experiments with a **stated duration, stated success and failure criteria, and a rollback plan** – plus an expiry that is automatic rather than requiring somebody to argue for ending it. The asymmetry matters: continuing should require a decision, not silence.

RCOS · LAYER 6

RCOS permits experiments and specifies what one has to contain. Every experiment must define, at minimum: **scope – what is changed and what is explicitly not – duration and review checkpoints, success and failure criteria, rollback conditions and the rollback process, and the authorised decision path for starting, extending, modifying or terminating it** (§8.3.2).

Failure criteria as a required field is the clause most communities are missing, and it is what separates an experiment from a change.

Then the expiry rule, which fixes the default: experiments must be **explicitly labelled as experimental in all affected artifacts and must include a non-extendable expiration date unless renewed through an authorised decision** (§8.3.4). Non-extendable, and renewal is a decision – so silence ends the experiment rather than continuing it.

Two boundaries around them. Experiments **must not override Layer 0 invariants and must not bypass the governance constraints of Layer 2** (§8.3.3) – no trialling your way past the constitution. And if an experiment introduces safety risk, coercion or sustained harm, it must be **suspended or terminated immediately**, with review afterwards (§8.3.5).

The safety clauses generalise it: the system **must prefer reversible changes over irreversible ones where possible** (§8.5.1), and irreversible or high-impact changes must carry **extended deliberation, higher thresholds where appropriate, and explicit risk acknowledgment** (§8.5.2).

SOURCES & FURTHER READING

- 1 **RCOS stress test – Irreversible Experiments** – the specification this page puts into plain language
- 2 **RCOS Core – Layer 6: Evolution & Adaptation** – what an experiment must define, and why its expiry is non-extendable
- 3 **Proposals & ratification** – the phases a decision travels through, and the two most often skipped

Written by EcoHubs members, with AI assistance for drafting and editing, and reviewed by a person before publication. Facts are checked against the sources listed; anything we could not verify is marked. **How this is written**

Governance with no path to change it

The rules can no longer be changed, so people stop trying – and the community loses its members quietly instead of loudly.

WHAT IT LOOKS LIKE

- Nothing significant has been amended in years, and nobody experiences this as stability.
- Attempts to change something die in procedure rather than on the merits.
- Members describe proposals as not worth the effort, and are right.
- The amendment rule requires a level of agreement the community can no longer reach.

Why it is hard to see

Because it **presents as stability, which is a thing communities are proud of.**

A group whose rules have not changed in eight years can tell a flattering story about that: we got it right, we are settled, we are not one of those places that rewrites its constitution every year. And some of that may be true. The question is whether the rules are unchanged because nobody wants them changed, or because **nobody can**.

Those look identical from the outside and are opposite conditions. The tell is not the absence of change. It is the absence of *attempts* – and attempts leave no record when they die early, because a proposal that never got written is not in the minutes.

The mechanism is usually one of three, and all of them were installed as safeguards:

A threshold that no longer clears. Unanimity or a high supermajority works in a founding group of nine who chose each other. At forty, with turnover and a couple of people who have quietly disengaged, the same threshold means anything contested fails – including things a clear majority want.

A process with no owner. Changing a rule requires drafting, circulating, a deliberation window, a meeting. Someone has to run that, unpaid, on top of everything else, knowing it may fail. So the proposal that everybody agrees would be an improvement never gets made.

And the effect that closes the loop: people stop asking. Once a group has watched two or three sensible proposals die in procedure, it learns. Nobody decides to stop trying – they just have less energy this year, and next year, and by then the community's rules are effectively fixed by a rule nobody would have chosen.

The cost is paid in departures rather than arguments, which is why it is missed. A member who cannot change something that matters to them does not usually fight. They adjust, then narrow, then leave – and give a vague reason on the way out.

What to check

When was something significant last amended? Not a new policy – an existing rule, changed. If the answer is over three years and the community has grown or turned over since, ask why.

Ask what died before it was proposed. This is where the evidence is. Members can usually name a change they thought about suggesting and did not, and the reasons will be procedural rather than substantive.

Do the arithmetic on the threshold. How many members would need to agree, out of how many, and how many typically attend? Communities are frequently startled to find their amendment rule now requires near-unanimity of everyone who exists, including people who have not attended in two years.

If it is already happening

Fix the amendment rule first, and expect that to be hard, because you have to use the broken process to fix itself. This is genuinely the hardest move in this guide. It usually needs a single, narrow, well-explained proposal that changes only the amendment rule, framed explicitly as *we are not deciding any of the contested things today*.

Separate the tiers while you are there. Much ossification comes from one high bar applied to everything. Constitutional matters deserve it; the guest policy does not, and treating them alike is what exhausts the process.

Give proposals an owner and a route. A named person whose role is to help a member turn a frustration into a written proposal removes the single biggest barrier, which is not opposition but drafting.

And count quorum against attenders, not members. A threshold measured against people who have effectively left hands a veto to absence.

What prevents it

An amendment rule the community can actually clear, **tiered by what is being amended**, plus a stated route for a member to get something onto the agenda. Change has to be possible, or the agreements stop describing the community and the members leave one at a time.

RCOS · LAYER 6

The layer invariant states both failures in one line: “**Change MUST be possible but constrained; no change MAY be instantaneous, implicit, or unreviewable**” (§8.7.1). *Possible* is doing as much work there as *constrained* – RCOS treats an unchangeable community as non-conforming, not as unusually stable.

A community must define **explicit change mechanisms for modifying, adding, suspending or removing rules, roles, artifacts or decision structures** (§8.1.1), and the Change Protocol must define **how changes are proposed, reviewed, adopted, published and rejected** (§8.6.3) – proposed being the step that has no owner in most communities.

It must also define **review mechanisms for adopted changes, including how changes are evaluated, revised or reverted when they produce harm, instability, or unintended concentration of power** (§8.1.5). A community with no way to reverse a change has a one-way ratchet, which is this pattern's other face.

Layer 2 supplies the tiering that stops one bar exhausting everything: each decision type must have its **own** explicitly defined mechanism (§4.2.1), specifying thresholds and time constraints (§4.2.3) – so constitutional ceremony applies where it belongs and not to the guest policy.

RCOS Core v0.1 – §8.1 Change Mechanisms, §8.7 Layer Invariants

SOURCES & FURTHER READING

- 1 **RCOS stress test – Ossified Governance Without a Path to Change** – the specification this page puts into plain language
- 2 **Ostrom, E. – Governing the Commons (1990)⁷** – collective-choice arrangements – the people affected by the rules can change the rules
- 3 **Decision methods** – quorum, thresholds, and the settings that quietly decide fairness

Written by EcoHubs members, with AI assistance for drafting and editing, and reviewed by a person before publication. Facts are checked against the sources listed; anything we could not verify is marked. **How this is written**

Sub-groups with undefined autonomy

A community grows into parts, and nobody says what the parts may decide – so each one improvises a constitution and discovers later they do not match.

WHAT IT LOOKS LIKE

- Sub-groups have formed and no document says what they may decide alone.
- Two parts of the community have different rules on the same question, discovered by accident.
- Nobody can say whether the main agreements bind a sub-group, or how one could differ.
- A decision made in one part turns out to commit the whole community.

Why it is hard to see

Because **the sub-groups form for good reasons and work well.**

A community reaches a size or a spread where doing everything together stops being sensible. Distance, or numbers, or simply that the people at one end share a set of daily concerns the people at the other end do not. So clusters form, meet on their own, and handle their own business – and this is exactly what a growing community should do. Meeting as forty about something that concerns eight is the failure they are correctly avoiding.

What nobody writes down is **the boundary**: what a sub-group may decide alone, whether the main agreements bind it, and how it could lawfully differ from them if it wanted to.

For a while the absence costs nothing, because the sub-groups make small decisions and everyone assumes the general rules still apply. Then three things happen, usually in this order.

Divergence by improvisation. Each part answers a question the main agreements never covered, sensibly and differently. Nobody has broken anything; the community simply now has two answers, and finds out when somebody moves between them.

A decision that escapes its container. A sub-group agrees something that turns out to commit the whole community – a shared cost, a precedent, an obligation. It had no way of knowing, because nothing said where its authority stopped.

And the constitutional question, arriving too late. Somebody eventually asks whether a sub-group *may* have its own rule on something. The honest answer is that nobody knows, and by then the sub-group has been running one for three years, so the question is no longer abstract – it is about whether to overturn something people have built their lives around.

What to check

Ask what a sub-group may decide without asking anybody. If the answer is a judgement rather than a document, this is running.

Take one rule from the main agreements and ask whether it binds every part. Then ask two people from different parts. Divergence shows up fast under that question and almost never under a general one.

Ask how a sub-group could get its own rule legitimately. Not whether it may – how. A community that has no answer has not forbidden it. It has left it to whoever acts.

If it is already happening

Map before you legislate. What has each part actually been deciding? The list is the useful artefact, and it is usually revealing: some of it is plainly local, some of it plainly is not, and the argument is about a much smaller set than anyone feared.

Write the three-way split. What the whole community decides, what a sub-group decides alone, and **what a sub-group may vary if it chooses** – the third being the one that gets forgotten and the one that prevents the next round of this.

Ratify what has already diverged, or end it, deliberately. Existing local rules should be looked at once and either adopted as legitimate variation or brought back into line. Leaving them in an undefined state is what produced the problem.

And name the floor. Whatever else varies, some things should not: the invariants, the exit rights, the conflict process. A sub-group that can vary the exit terms is a separate community sharing a driveway.

RCOS HAS NO STRESS TEST FOR THIS ONE

The other twenty-four patterns in this guide paraphrase a documented RCOS stress test. This one does not – **the standard has no test for it**, and we have flagged it as a candidate for a future version.

The raw material is all there: Layer 0's scope declaration, Layer 6's change mechanisms, the group-size thresholds around 120–150 above which sub-structures become mandatory, and Ostrom's nested units. What is missing is a test that names the failure, which is why a community can be reading RCOS attentively and still walk into this.

We would rather say that than quietly leave the pattern out because it has no citation.

What prevents it

A stated subsidiarity rule: for each domain, whether it is decided centrally, locally, or locally within a stated range – plus a floor of things that may not vary. The same clarity a domain gives a circle, applied to a place rather than a function.

RCOS · LAYER 0

There is no stress test for this pattern, but the Core specification does constrain it in three places worth knowing.

Scope. A community declares the domains it holds decision authority over, and anything not declared in scope is out of scope (§2.2.1–4). Applied to sub-groups this cuts both ways and neither is stated: a sub-group's authority is not declared, so on a strict reading it has none – which is not how any real community behaves.

Change. Every proposed change must specify the artifacts, layers and sections affected and its authorised decision path (§8.1.3). A local rule that varies a general one is a change to a general artifact, and would have to be adopted as one – which is the answer to *how could a sub-group differ legitimately*, arrived at indirectly.

Invariants. Whatever varies locally, **no decision, role, process or emergency measure may override an invariant** (§2.3.4). That is the floor, and it holds without a subsidiarity rule existing.

RCOS also notes that above roughly 120–150 members **sub-structures become mandatory** – so the standard anticipates the structure arising while not yet specifying how its authority is bounded.

RCOS Core v0.1 – §2.2 Scope Declaration, §8.1 Change Mechanisms

SOURCES & FURTHER READING

- 1 **RCOS Core – Layer 0: Identity & Scope** – scope declaration and invariants – the constraints that apply in the absence of a subsidiarity rule
- 2 **Ostrom, E. – Governing the Commons (1990)**⁷ – nested units for larger systems, the eighth design principle
- 3 **Community governance** – domains, and why one is only real when it carries a concrete limit

Written by EcoHubs members, with AI assistance for drafting and editing, and reviewed by a person before publication. Facts are checked against the sources listed; anything we could not verify is marked.

How this is written

The emergency measure that stayed

Decided fast because it had to be, never reviewed, and now it is what the community does – with a precedent attached.

WHAT IT LOOKS LIKE

- An arrangement adopted during something urgent is still running, and nobody set an end date.
- Precedent is doing the work that a written rule should be doing – we did it last time.
- Nobody can say what would count as the emergency being over.
- The people who acted quickly were right, and that is why nobody wants to revisit it.

Why it is hard to see

Because it is **created by competence and preserved by gratitude**.

Something urgent happened. Somebody acted without the usual process, because the usual process takes three weeks and the pipe was leaking now. They were right, the community was relieved, and the shortcut is remembered as the moment the group got something done.

Nothing about that arrangement ever ends, because **nothing about it ever began**. No scope was set, no expiry, no review – there was no time, and setting one would have felt like distrust of somebody doing the community a favour.

What makes it worse than a simple lapse is the second-order effect. The next urgent thing arrives, and now there is a precedent: this is how we handle urgent things. The bar for what counts as urgent drops, because it always does when the fast route is available and the slow one is tiring. Eventually the exception is the route for anything that anybody frames as pressing.

And revisiting it is socially expensive. Reopening the decision means questioning somebody who acted well, at cost to themselves, when the community needed it. Almost nobody will do that, which is exactly why the review has to be scheduled in advance rather than proposed afterwards.

What to check

List the arrangements that began during something difficult. Every community over five years old has two or three. Ask when each was last reviewed, and whether anyone set a date at the time.

Ask what would count as the emergency being over. If nobody can define the end condition, it has no end condition, and the measure is permanent whatever anyone intends.

Look at what the precedent has been used for since. The original use was probably justified. The question is what it has authorised since, and whether that would have passed on its own merits.

If it is already happening

Do not challenge the original decision. It was probably right, and arguing about it turns a structural conversation into a personal one. The framing that works is prospective: *this was the right call and we never closed it – let us close it now.*

Give every live measure a date. Not a repeal – a date by which it is either adopted properly, through the normal process, or lapses. Most will be adopted, and that is a fine outcome: what changes is that the community has now actually decided, rather than continued.

Then write the emergency provision you were missing, while there is no emergency. What may be decided quickly, by whom, up to what limit, and – the part that matters – **it expires automatically and gets reviewed afterwards.** A community that writes this before it needs it never has this pattern again; a community that writes it during a crisis writes it about the crisis.

What prevents it

An emergency clause with **an expiry built in and a mandatory review after the fact.** The expiry is what makes it an exception rather than a second constitution, and it has to be automatic – a measure that continues unless somebody actively ends it will continue, because ending it requires somebody to spend social credit on a thing that is currently working.

RCOS · LAYER 6

RCOS permits emergency change and fences it on four sides at once: “**Emergency changes MAY be permitted only where explicitly defined, MUST be time-bounded, MUST NOT override Layer 0 invariants, and MUST undergo mandatory post-hoc review and ratification or rollback**” (§8.5.3).

Read those in order. *Explicitly defined* – written before it is used, so the exception is not invented under pressure by whoever is acting. *Time-bounded* – it ends by default. *Must not override Layer 0 invariants* – urgency does not reach the constitutional layer, which Layer 0 restates: **no decision, role, process, or emergency measure may override an invariant** (§2.3.4). And *ratification or rollback* – the review has two permitted outcomes, and “carry on as we were” is not one of them.

Layer 2 carries the authority half: **temporary or emergency authority must be explicitly defined, time-bounded, and subject to review** (§4.3.5).

The front matter of the standard puts the principle plainly, and it is aimed at this pattern: **urgency is never treated as justification.**

RCOS Core v0.1 – §8.5 Change Safety and Reversibility

SOURCES & FURTHER READING

- 1 **RCOS stress test – Emergency Rule Bypass Precedent** – the specification this page puts into plain language

- 2 **Just For Now: how emergency rules become permanent** – the longer EcoHubs argument, with the precedent problem worked through
 - 3 **RCOS Core – Layer 6: Evolution & Adaptation** – emergency changes: explicitly defined, time-bounded, reviewed, ratified or rolled back
-

Written by EcoHubs members, with AI assistance for drafting and editing, and reviewed by a person before publication. Facts are checked against the sources listed; anything we could not verify is marked.

How this is written

The invariant nobody protected

The founding principle everybody assumed was permanent — changeable, it turns out, by an ordinary majority on an ordinary evening.

WHAT IT LOOKS LIKE

- The community has principles it calls non-negotiable and no rule saying how they change.
- The same process amends the guest policy and the constitution.
- Members would be shocked by a change that is procedurally available today.
- Nobody can list what could never be decided by a simple majority.

Why it is hard to see

Because **the protection everybody is relying on is social, and social protection is invisible until it fails.**

Ask a community whether it could sell the land, or drop the thing it exists for, and the honest answer is usually *of course not, nobody would ever vote for that*. Which is true, and is a statement about the current membership rather than about the community. It describes who is in the room, not what the rules permit.

Membership turns over. A principle held by everybody in year three is held by a majority in year nine and by a large minority in year fifteen — not because anybody changed their mind, but because different people arrived. At no point does anybody betray anything. **The protection simply evaporates, quietly, through ordinary intake.**

The second reason it hides: **a community with one decision process cannot see the distinction.** If the guest policy and the founding purpose are amended by the same procedure, there is no visible difference between them, and nothing prompts the group to ask whether there should be. The absence of a constitutional tier is not experienced as a gap. It is experienced as simplicity.

What to check

Ask what could not be changed by a simple majority tomorrow. Then check the agreements for whether that is actually true. The gap between the two answers is the exposure, and it is usually total.

Try to find the amendment rule. Not for a policy — for the founding document. Most communities discover their agreements do not describe how they themselves are changed, which means they are changed by whatever process is at hand.

Ask the newest members what they think is permanent. If their list differs from the founders' list, the group already has two constitutions and only one of them is written.

If it is already happening

Write the short list first. Not a constitution – a list. The three to six things this community exists for and would not survive losing. Communities routinely find this takes an evening and produces a much shorter list than anyone expected, which is itself useful.

Then attach a higher bar to that list, and only that list. A supermajority, a longer deliberation window, an objection period, a second vote at a later meeting – the specific mechanism matters less than that it differs from the ordinary one. Applying constitutional ceremony to everything is the failure at the other end, and it produces a community that cannot decide anything.

Include the amendment rule in what is protected. A high bar that can be lowered by an ordinary vote is not a high bar. This is the step most often missed, and it is the one that makes the rest hold.

Expect resistance from an unexpected direction. Some members will read this as distrust of the group's future self, which is exactly what it is – and that is the point. You are protecting the community from a membership that does not exist yet, and the current membership is not the one at risk.

What prevents it

An **invariants register**: the things that may not be changed by ordinary decision, listed explicitly, with a stated constitutional process for changing them and the amendment rule inside the protected set.

RCOS · LAYER 0

RCOS makes the register a mandatory Layer 0 artifact rather than a good habit. Invariants are constraints defining **what must not be violated while they are in force** (§2.3.1), and they **must be explicitly listed and documented** (§2.3.2) – the listing being the whole intervention, since an unwritten invariant is the social protection described above.

Three clauses give them force. They **apply across all layers** (§2.3.3). **No decision, role, process, or emergency measure may override an invariant** (§2.3.4) – emergency measures named, closing the route that **the measure that stayed** opens. And where an invariant conflicts with any other rule, **the invariant prevails** (§2.3.5).

Then the amendment protection this pattern most often lacks: invariants **may only be changed or removed through a constitutional change process** as defined in Layer 2 and Layer 6 (§2.3.6). Layer 6 restates it from the other side – changes affecting Layer 0 purpose, scope, invariants or identity constraints **must be classified as constitutional** and follow the constitutional mechanism (§8.1.4).

If the Invariants Register is missing, ambiguous or internally contradictory, the community is **non-compliant with RCOS-Core** (§2.5.3). It is one of four artifacts the standard will not do without.

RCOS Core v0.1 – §2.3 Invariants

SOURCES & FURTHER READING

- 1 **RCOS stress test – Unprotected Core Invariants** – the specification this page puts into plain language
 - 2 **RCOS Core – Layer 0: Identity & Scope** – the invariants register, and what may never be overridden
 - 3 **Decision tiers** – operational, strategic, constitutional – and why the unclear case defaults upward
-

Written by EcoHubs members, with AI assistance for drafting and editing, and reviewed by a person before publication. Facts are checked against the sources listed; anything we could not verify is marked. **How this is written**