Choosing How We Choose

Ways of setting the agenda for community meetings

A briefing for the Nyland Process Group

> **About this document.** This was written by Claude, an AI assistant made by Anthropic, over the course of a long conversation with Matt Elder in which he set the questions, pushed back on weak answers, and directed what was worth pursuing. The research behind it is Claude's; the judgment about what mattered for Nyland is largely Matt's. Matt edited it further.
>
> Companion documents go deeper on two methods: The Method of Equal Shares and Sequential Phragmén.

What we're actually deciding

We've agreed on the general shape: members submit agenda item requests, those requests sit in a backlog, and the community somehow selects which ones make it onto a given meeting's agenda.

"Somehow" is carrying a lot of weight in that sentence. There are three decisions hiding inside it:

1. **What does a ballot look like?** How does a member express what they want?
2. **How do ballots become an agenda?** Given everyone's input and a fixed amount of meeting time, which items get scheduled?
3. **What counts as fair?** This sounds like the abstract question, but it turns out to determine the other two — and it's a values question, not a technical one.

Most groups settle (1) and (2) by instinct and never notice they've made a choice about (3). We'd do better to take (3) first.

A useful picture throughout: **meeting time is a shared resource, like the common house or the garden beds.** We have roughly 90 minutes. Items have different sizes. Not everything fits.

Part 1: The ballot

Three families of ballot get proposed for this kind of thing.

Ranked choice. Members rank items in order of preference. This works well when you're picking *one* winner from several candidates — which is what it was designed for. It works poorly here, because we're not picking a winner. We're packing a set of items into a time budget, and most of the ranking information can't actually be used by any sensible packing method. Ranked ballots are also the slowest and most error-prone to fill out on paper in a room. **Not recommended.**

"You get K votes." Each member gets, say, three votes to distribute among submitted items. Intuitive and familiar. Its hidden cost is that it forces people to *ration*: if you genuinely think five items deserve discussion, you must withhold support from two of them. Worse, it invites second-guessing — "I won't vote for the parking item, it'll pass anyway, I'll save my vote for something that needs it." We end up measuring people's predictions about other people rather than their actual preferences. **Workable but lossy.**

Approval voting. Members check every item they want on the agenda, with no limit. Simple to explain, fast to fill out, nothing to ration, no strategy required. Its one real limitation is that it doesn't capture *intensity* — an item you'd walk over broken glass for and one you're mildly curious about get the same checkmark. **Recommended**, and it's what every method below assumes.

Two details that matter more than they look:

- **Print each item's estimated duration on the ballot.** "Do you want this 45-minute discussion?" is a genuinely different question from "do you want this discussion?" Durations function as prices, and people should see prices before they choose.
- **State plainly on the ballot that checking more items does not weaken your other checks.** People trained on "you get three votes" will otherwise under-check out of a mistaken folk theory about vote dilution.

Part 2: The problem that shapes everything else

Once we have approval counts, the obvious move is: rank items by number of approvals, fill the agenda from the top down until time runs out.

That's what almost every organization does, and it has a flaw that's easy to miss and hard to reverse.

The 51/49 problem

Imagine our community splits, softly and without anyone intending it, into two rough centers of gravity. Fifty members care most about buildings, budget, and infrastructure. Forty care most about land, gardens, and outdoor common space. Nobody is hostile. People just find different things urgent.

Under "rank by approval count," the fifty-member cluster's items win every slot, every meeting, indefinitely. The forty-member cluster's approvals are collected, tallied, and then discarded — month after month, with perfect procedural regularity.

The predictable result is that those forty members conclude voting doesn't do anything for them and stop voting. Now they're twenty-five active voters instead of forty. The gap widens. Nobody did anything wrong, no rule was broken, and the process is impeccably democratic on its face.

And it has manufactured exactly the "I don't have a voice here" feeling we set out to prevent.

This isn't a hypothetical edge case. It is the default behavior of majority-rule selection applied repeatedly to a stable population, and it's the reason the field of computational social choice spent the last decade developing alternatives.

Why the obvious patches don't work

The natural instinct is to patch the scoring: give older items a bonus, divide approvals by duration, weight by urgency. All of these are worth discussing on their own merits, but **none of them touch the 51/49 problem.** They all still amount to "compute a number for each item, then take the best ones," and any method of that shape hands the larger cluster the top slots regardless of what the numbers are.

Fixing it requires a structurally different kind of method — one where an item's standing depends on *how much its supporters have already received.*

That single requirement is what separates the options below into two families.

Part 3: The options

A. Rank by approval count

Count the checkmarks, sort, fill the agenda top-down.

Simple, familiar, instantly explainable, and requires nothing but a tally sheet. If the 51/49 concern strikes the group as theoretical for a community like ours, this is the right answer and we should stop here rather than build machinery we don't need.

The honest case against: our stated worry is that people will feel they lack a voice. This method is the one most likely to produce that feeling, and it produces it slowly enough that we won't notice until the pattern is well established. It also wastes information — every approval beyond the number needed to win a slot does nothing at all.

B. The Method of Equal Shares

Every member gets an equal share of the meeting's minutes. An item gets scheduled when the members who approved it can jointly afford its duration out of their remaining shares. Members never actually spend anything — the "spending" is bookkeeping done during the count — but that bookkeeping is what produces fairness.

This is a real method from the research literature (Dominik Peters, Grzegorz Pierczyński and Piotr Skowron, 2021), currently used for participatory budgeting in several European municipalities. It satisfies a formal fairness property called Extended Justified Representation. Roughly: if a group all want the same set of items, and the group is large enough that its members' combined shares could pay for that set, then at least one of them ends up at least as well served as that set would have made them. It's a floor on what a cohesive minority can be denied — not a cap on what a majority can take, which is a distinction worth keeping in mind.

It also produces a *receipt* — for every scheduled item, you can say exactly whose shares paid for it, which is a transparency property we may value more than the math.

Costs: this method is quite complex! The counting genuinely needs a spreadsheet or a script, it can leave time unspent and so needs a rule for filling gaps, and carrying unspent share between meetings requires ballots to be identifiable rather than anonymous. **See the companion document for the full explanation and a worked example.**

C. Sequential Phragmén

During item selection, each member carries a running tally of community minutes charged to them. When an item is scheduled, its duration is divided among the members who approved it, split so their tallies come out equal. At each step, we schedule whichever item leaves the highest tally lowest — in short, **the item that spreads the community's time most evenly.**

Same ballot as Equal Shares, and in testing it reaches very similar agendas. The reason to consider it is everything it *doesn't* require: no budget to assign, nothing carried between meetings, no cap to prevent members banking influence by staying away, no adjustment when turnout varies, and — because it almost always fills the available time — no rule for what to do with leftover minutes. Each of those is a decision the Process Group would otherwise have to make and later defend.

It also permits **fully anonymous ballots**, which Equal Shares cannot if we want fairness to carry across meetings.

It also has a track record that none of the other options here can match. The Swedish mathematician Lars Edvard Phragmén published it in the 1890s, during a national debate about giving minorities representation in parliament; versions of it were written into Swedish election law, used for municipal elections from 1922, and it remains in the law today. English-language researchers only caught up with it in 2016–17. A century of real elections tells us more about how a method behaves in practice than any proof does.

Costs: its formal fairness guarantee is a notch weaker than Equal Shares', the researchers behind Equal Shares argue it doesn't extend cleanly to items of differing lengths, and the idea of a "tally

charged to you" is a step more abstract than "minutes you spend." **See the companion document for the full explanation and a worked example.**

D. Take turns — a lottery

Draw members at random. Each one, in turn, names an item from the backlog. Continue until the agenda is full.

No ballots. No tallying. No software, ever. Nothing to manipulate, because there is no way to game a lottery. Everyone has exactly equal claim, and the fairness is obvious to anyone watching rather than requiring an explanation.

The cost is real: it discards all information about *how many* people want something. An item 60 members want and an item 1 member wants have different likelihoods of being selected, but both are possible, and any single meeting could come out lopsided. Over a year the variance evens out, but individual meetings may not feel optimized.

This deserves serious discussion rather than a polite mention. It competes with Equal Shares on the fairness axis at roughly one-hundredth the complexity, and the group's objections to it will reveal a great deal about what we actually want from the process. (Selection by lot has a long democratic pedigree — it was the primary mechanism for filling office in classical Athens, precisely because it was immune to the accumulation of influence.)

E. Ask for objections instead of preferences

The traditional cohousing approach, and the incumbent we'd be departing from. Someone prepares an agenda; the meeting asks whether anyone has a principled objection. Legitimacy comes from the absence of serious objection, not from tallied approval.

It's dramatically cheaper than anything above, and it handles the case where an agenda is actively wrong much better than voting does.

Its weakness is precisely our concern: **nobody objects to an agenda that simply never contains their topics.** Consent is an excellent instrument for catching things that are wrong and a poor one for detecting quiet dissatisfaction. That's why it probably isn't our answer — but it's the status quo, and it deserves an explicit reason for being decided against rather than being assumed away.

F. Let someone curate it

The agenda as an edited artifact — like a magazine issue, a conference program, or a court's docket. Legitimacy comes from the curator being accountable and their criteria being public, not from a procedure being provably fair.

This frame raises a real objection to everything else in this document: **voting can only surface what people already know they want.** It cannot surface a thing nobody wants to discuss but that we need to discuss. It cannot surface something people would care intensely about if they understood the stakes. No voting rule fixes this, because the information simply isn't on the ballots.

The most interesting hybrid is the U.S. Supreme Court's certiorari process. By longstanding convention, four of nine justices suffice to put a case on the docket. That's not a vote — it's a deliberately tuned threshold that lets a minority force the full body's attention. A Nyland version might be: any six members who co-sign an item get it scheduled, no vote required.

Whatever we choose, this frame argues for keeping *something* outside the selection mechanism — a reserved slot the Process Group fills on judgment, or a standing rule that certain categories (safety, legal deadlines, financial obligations) bypass selection entirely.

G. Measure how much people care

Methods exist that let people express intensity by spending a budget on votes at escalating cost — quadratic voting is the developed version. Attractive in principle, and it fixes approval voting's real limitation.

I'd advise against it here. These methods are notoriously vulnerable to small groups coordinating, and that vulnerability is worst in exactly our situation: a small community where everyone knows everyone and three families talking at dinner can meaningfully move an outcome. The comprehension cost is also higher than Equal Shares, not lower.

H. The option that dissolves the problem

Every option above assumes plenary minutes are the scarce good. That scarcity exists partly because we've assumed items must reach a full community meeting at all.

If the real constraint is *attention* rather than *minutes*, the leverage isn't in a better allocation method — it's in reducing what needs allocating. Giving committees genuine decision authority within defined areas. Gathering objections in writing before meetings so that floor time handles only unresolved ones. Routing informational items to the newsletter and never to the agenda.

Elinor Ostrom's research on communities that successfully govern shared resources found that durable ones are almost always organized in *nested* layers, with decisions made at the smallest level that can competently make them. An elaborate selection mechanism at the top layer can be a symptom of too much reaching that layer, rather than a fix for it. Worth ruling out before we invest in machinery.

Part 4: What we're really choosing

Here's the thing worth sitting with before any of the details.

- **Equal Shares** treats each member as a **claimant with an equal entitlement to a shared resource.**
- **Sequential Phragmén** treats each member as a **share of the community's time that should be spent evenly.**
- **Lottery** treats each member as a **holder of equal probability.**
- **Consent** treats each member as a **holder of a veto.**
- **Curation** treats each member as a **beneficiary of someone's good judgment.**
- **Approval ranking** treats each member as **one unit of preference in an aggregate.**

Those are all different theories of what a Nyland member *is*. Whichever mechanism we adopt will quietly teach the community one of them over the next several years — not through anyone explaining it, but through a thousand small experiences of asking for something and getting it or not getting it.

That is the actual choice in front of us, and it is not a technical one. The arithmetic is downstream.

Questions for the group

1. **Is our primary worry efficiency or voice?** (Meetings cost too much for what they deliver, versus people don't feel they can affect what we discuss.) These pull toward different methods. If it's mostly efficiency, the whole selection question may be secondary to option G.
2. **Does the 51/49 concern feel real for Nyland, or theoretical?** If theoretical, simple approval ranking is much easier and we should just do that. This is the hinge question.
3. **How much complexity are we willing to tolerate in the counting**, given that the *voting* stays equally simple either way? Equal Shares has compelling properties but probably requires opaque automation.
4. **Are we willing to give up ballot anonymity** in exchange for fairness that works across meetings rather than only within one? Carrying unspent share forward requires knowing whose ballot is whose.
5. **What stays outside the vote regardless?** Safety items, legal deadlines, a reserved Process Group slot, a co-signature threshold?
6. **What's our review commitment?** Whatever we pick, we should agree now to look at it again after — say — six meetings, with explicit willingness to change it. Adopting a method provisionally is far easier than adopting one permanently, and it lowers the stakes of this decision considerably.

Revision #12
Created 2026-08-24 04:48:08 UTC
Updated 2026-09-03 19:05:54 UTC by Matt Elder