

Briefs on Agenda Voting

- [Choosing How We Choose](#)
- [The Method of Equal Shares](#)
- [Sequential Phragmén](#)

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.

The Method of Equal Shares

How it works, what it guarantees, and what it costs

A briefing for the Nyland Process Group

About this document. Written by Claude, an AI assistant made by Anthropic, working from the research literature at Matt's direction. It's a companion to [Choosing How We Choose](#) and [Sequential Phragmén](#).

The idea in one sentence

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 don't actually spend anything. The "spending" is bookkeeping done during the vote count. But that bookkeeping is what makes the method fair, and it's what distinguishes it from every method that simply ranks items by popularity.

Where it comes from

The Method of Equal Shares was introduced by Dominik Peters, Grzegorz Pierczyński and Piotr Skowron in 2021, building on earlier work by Peters and Skowron. It's used in practice for participatory budgeting — deciding which citizen-proposed projects a city funds with a limited budget — in some European municipalities, and there's public reference material and open-source implementations at equalshares.net.

Our situation is structurally identical to participatory budgeting. Their projects have costs; our agenda items have durations. Their budget is money; ours is 90 minutes. This is not a case of borrowing a metaphor from a distant field — it's the same mathematical problem with different units.

What makes it more than a clever heuristic is that it satisfies a formal 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. Skowron's own illustration of what this fixes: if 51% of voters support only "blue" projects and 49% support only "green," a standard highest-votes method gives everything to blue and throws away every green vote.

The ballot

Ordinary. This is worth emphasizing, because it's the most common misunderstanding: **all the complexity lives in the counting, none of it in the voting.**

NYLAND AGENDA BALLOT – Community Meeting, Oct 12

Agenda time available: 90 minutes

Check every item you want on this agenda.

Check as many as you like – checking more does NOT weaken your other checks.

- [] 10 min Set date for fall workday
- [] 20 min Guest parking policy
- [] 25 min Common house kitchen scope decision
- [] 20 min Budget variance review
- [] 18 min Landscape Team report & questions
- [] 45 min Pet policy – first discussion

Member ID: _____

Ballots can be filled out ahead of time, dropped in a box, or completed in the room. The method is a fixed calculation over a fixed set of ballots — it doesn't care when or where they were filled out, and voters never respond to each other.

A worked example

Assume 90 members vote and we have 90 minutes of agenda time. Each member's share is exactly **1 minute** — convenient numbers, close enough to our real ones.

Suppose the community has three loose centers of gravity:

- **Cluster A — buildings, budget, infrastructure: 48 members**
- **Cluster B — land, gardens, outdoor space: 21 members**
- **Cluster C — broadly engaged, no strong center: 21 members**

Nobody chose these clusters. They're just where people's attention naturally sits.

The submitted items and who approved them:

Item	Duration	Approvals	Who
Set fall workday date	10 min	90	everyone
Guest parking policy	20 min	69	A + C
Common house kitchen scope	25 min	48	A
Budget variance review	20 min	30	part of A
Landscape Team report	18 min	21	B
Pet policy — first discussion	45 min	26	scattered

What simple approval ranking produces

Sort by approval count, fill top-down:

- Workday (90 approvals, 10 min) → **10 minutes used**
- Guest parking (69, 20 min) → **30 used**
- Kitchen scope (48, 25 min) → **55 used**
- Budget variance (30, 20 min) → **75 used**
- Pet policy (26, 45 min) → doesn't fit in the remaining 15
- Landscape report (21, 18 min) → doesn't fit in the remaining 15

Final agenda: workday, parking, kitchen, budget variance. Four items, 75 minutes.

Cluster B — twenty-one members, a quarter of the community — gets nothing. Their approvals were counted and discarded. If the clusters are stable, this repeats next month and the month after.

What Equal Shares produces

Everyone starts with 1.00 minute of share. At each step we ask: **which item is cheapest per supporter, among those its supporters can still afford?**

Round 1. Prices per supporter:

Item	Calculation	Price each
Workday	$10 \div 90$	0.11
Parking	$20 \div 69$	0.29
Kitchen	$25 \div 48$	0.52
Budget variance	$20 \div 30$	0.67

Item	Calculation	Price each
Landscape	$18 \div 21$	0.86
Pet policy	$45 \div 26$	1.73 — <i>more than anyone has</i>

Cheapest is the workday item. **Scheduled.** Every member pays 0.11 and now holds 0.89. *(10 minutes used.)*

Round 2. Parking is cheapest at 0.29 per supporter, and its 69 supporters each hold 0.89. **Scheduled.** Clusters A and C drop to 0.60; cluster B still holds 0.89, because none of their money went to parking. *(30 minutes used.)*

Round 3. Kitchen scope costs 0.52 per supporter and cluster A holds 0.60 — affordable. Budget variance costs 0.67 and its supporters hold 0.60 — *not* affordable. Landscape costs 0.86 and cluster B holds 0.89 — affordable. Of the two affordable items, kitchen is cheaper. **Scheduled.** Cluster A drops to 0.08. *(55 minutes used.)*

Round 4. Budget variance still costs 0.67, and cluster A is now nearly broke at 0.08 — they've had three items and spent almost their whole share. Landscape costs 0.86 and cluster B holds 0.89. **Landscape is scheduled.** Cluster B drops to 0.03. *(73 minutes used.)*

Round 5. Nothing remaining is affordable. Stop.

Final agenda: workday, parking, kitchen scope, landscape report. Four items, 73 minutes.

What just happened

The Landscape Team report, with **21 approvals**, beat the budget variance review, with **30 approvals**.

Not because it was more popular. Because its supporters hadn't gotten anything yet, and the budget variance supporters had already received three of the four items on the agenda.

That is the entire point of the method, and it is the thing that no amount of clever scoring can reproduce. You cannot get this result from any rule of the form "give each item a number and take the best ones," because at the moment the fourth slot is decided, budget variance genuinely *is* the more popular item. What changed isn't the item's merit — it's what its supporters are still owed.

Why "cheapest per supporter"?

At every step, the method buys the item that asks the least of each of its supporters. That rule deserves to be stated as a principle rather than buried as an implementation detail, because it is doing nearly all of the work:

|

Among the items that can be afforded, schedule the one that costs its supporters the least each.

Two reasons it's the right rule.

First, it makes everyone's share go furthest. Spending a group's minutes where they buy the most per minute leaves the maximum purchasing power intact for everyone else — including for groups who haven't been served yet.

Second, and more important for us: **it is the only part of the method that isn't forced.** Equal shares, splitting an item's cost among its supporters, rolling over what's unspent — each of those follows almost inevitably once you decide meeting time is a commonly-held resource. The selection rule is a genuine choice, and the obvious alternative is to work through candidate items in some fixed order, buying whatever is affordable as you reach it.

That alternative fails, and it fails in an instructive way. Working in a fixed order can drain a group's entire share on an item they share with everyone, leaving them nothing for the item only they wanted — while another group gets both. Choosing the cheapest-per-supporter item first prevents exactly that.

And there's a governance argument on top of the mathematical one. A fixed scanning order has to be set by someone. Whoever sets it holds real influence over outcomes, in a place nobody would think to audit. The argmin rule needs no such order — which means it removes the last discretionary handle from the process, rather than relocating it somewhere less visible.

Fairness across meetings, not just within one

A recurring worry: an item that keeps not quite making the cut, month after month.

The instinct is to give waiting items a bonus that grows with age. **That aims at the wrong target.** Boosting an *item* for age rewards whoever is most persistent about resubmitting — resubmit long enough and you win regardless of support. Boosting a *member* who keeps not getting anything rewards the people we're actually worried about.

With Equal Shares this is nearly free: **let unspent share roll over.** A member who got nothing this month walks into next month with two minutes instead of one. A member who got three items walks in with one. Over a year, everyone gets roughly their fair share of the community's attention, automatically, without anyone tracking grievances.

This has been studied under the name *perpetual voting* (Martin Lackner, 2020), which formalizes fairness across a sequence of decisions rather than within one. One of its axioms is called

"bounded dry spells" — the guarantee that every voter gets a satisfying outcome on a regular basis, which is our anti-starvation requirement stated precisely.

There's a pleasing historical precedent. In 1918 Gottlob Frege — better known as the founder of modern logic — drafted an election law, rediscovered in a Jena library in 1998, in which votes cast for losing candidates carried over to the next election while winning candidates incurred a cost. It provably produces proportional representation over the long run. He was doing this a century before anyone had the computers to run it.

Honest costs and caveats

It needs a spreadsheet. Hand-runnable once, tedious monthly. Realistically we'd want it automated within a few cycles. This is a genuine adoption cost and shouldn't be minimized.

It can leave time unspent. Our example ended with 17 minutes free and no remaining item small enough to fit. The method is what mathematicians call *non-exhaustive*. We'd need an agreed rule: either accept a shorter meeting, use the slack for a reserved Process Group item, or fill it greedily with the highest-approval item that fits. (The more principled completion is to raise everyone's starting share slightly and re-run, but it's harder to explain and I'd skip it.)

Carryover requires identifiable ballots. Fully anonymous ballots and fairness-across-time are incompatible — you cannot carry forward a share you can't attribute. A stable pseudonymous member number, issued once, is the usual compromise. This is a real tradeoff and the group should decide it deliberately rather than discover it in month three.

It is not manipulation-proof. No reasonable voting method is; this is a theorem, not an oversight. The specific trick available is withholding approval from an item you're confident will pass anyway, so your share stays free for something niche. It's mild — it only moves your own share, and it backfires if you misjudge — and plain approval voting has the identical incentive.

A long item with narrow support may never make it. Pet policy in our example (45 minutes, 26 supporters) can never be afforded, because 26 members' full shares only add to 26 minutes. This is arguably correct behavior — a 45-minute discussion wanted by 29% of the community is a real claim on everyone's evening — but it means long items need either broad support or a route that bypasses the vote. Splitting such items into shorter first-pass discussions is the natural fix, and probably a good idea anyway.

The property that may matter most

Equal Shares produces a **receipt**.

For every scheduled item, you can state exactly which members' shares paid for it and how much each contributed. When someone asks "why did that get on the agenda and mine didn't," there is a real, checkable answer — not "it got more votes," but "here is the arithmetic, and here is where your share went instead."

For a community whose stated worry is that people feel unheard, that transparency may be worth more than the fairness theorem behind it.

Sequential Phragmén

A proportional method with fewer moving parts

A briefing for the Nyland Process Group

About this document. Written by Claude, an AI assistant made by Anthropic, working from the research literature at Matt's direction. Companion to [Choosing How We Choose](#) and [The Method of Equal Shares](#). It covers an alternative to Equal Shares that reaches similar outcomes with substantially less apparatus.

The idea in one sentence

“Each item's minutes are charged to the members who asked for it, split so the members' running totals come out equal. At each step, schedule the item that leaves the highest running total lowest.”

Or, less precisely but more memorably: **schedule the item that spreads the community's time most evenly.**

Each member carries one number — call it their *tally* — representing minutes of agenda time charged to them so far. Nobody is penalized for a high tally; it's an accounting device, and it resets to zero at the start of each meeting's calculation.

Where it comes from

Lars Edvard Phragmén was a Swedish mathematician who published this method between 1894 and 1899, as part of a public debate about how Sweden should elect its parliament. The problem under discussion was that simple block voting hands every seat to the largest party, and the country was looking for a method that would give minorities representation — which is recognizably our own concern, arrived at from the other direction.

It is not a theoretical curiosity. Versions of Phragmén's method were adopted into Swedish election law, used for distributing seats within parties, and taken up for municipal elections from 1922. As of the mid-2010s, it was still part of Swedish election law today, in a reduced role. That is

roughly a century of continuous use in a functioning democracy.

What happened recently is not the method's invention but its translation. The English-language literature on voting largely developed the Single Transferable Vote instead and overlooked Phragmén's work for most of the twentieth century. Svante Janson published a detailed English presentation in 2016, and in 2017 Brill, Freeman, Janson and Lackner connected it to the modern mathematical framework for proportional representation, proving what it guarantees.

For our purposes the track record matters more than the theorem. A century of real elections is evidence about how a method behaves in practice — where it confuses people, where it produces results that feel wrong — of a kind no proof supplies. By comparison, the Method of Equal Shares has roughly five years of municipal use behind it.

It has one property that matters here and that Equal Shares lacks: it needs **no budget parameter, no carried balances, and no rule for filling leftover time.**

The ballot

Identical to the Equal Shares ballot. Approval, uncapped, with durations shown.

NYLAND AGENDA BALLOT – Community Meeting, Oct 12

Agenda time available: 90 minutes

Check every item you want on this agenda.

Check as many as you like – checking more does NOT weaken your other checks.

[]

10 min

Set date for fall workday

[]

20 min

Guest parking policy

[]

25 min

Common house kitchen scope decision

[]

20 min

Budget variance review

[]

18 min

Landscape Team report & questions

[]

45 min

Pet policy – first discussion

No member ID needed. Because nothing carries between meetings, ballots can be fully anonymous. This is a real advantage over Equal Shares, where carrying unspent share forward requires knowing whose ballot is whose.

The calculation

For each candidate item, compute:

A horizontal bar chart with a light gray background. A blue segment is on the left, followed by a white segment, and then a light gray segment. The blue segment represents a small fraction of the total length.

$$\text{new tally} = (\text{item's duration} + \text{sum of its supporters' current tallies}) \div (\text{number of supporters})$$

Schedule whichever item gives the smallest result, and set all of that item's supporters to that new tally. Repeat until nothing else fits.

That's the entire method. One formula, applied repeatedly.

A worked example

Same instance as the Equal Shares document: 90 members, 90 minutes. Clusters A (48 members, buildings and budget), B (21, land and outdoors), C (21, broadly engaged). Everyone starts at tally 0.

Round 1. With all tallies at zero, the formula reduces to duration ÷ supporters:

Item	Calculation	New tally
Workday	$10 \div 90$	0.11
Parking	$20 \div 69$	0.29
Kitchen	$25 \div 48$	0.52
Budget variance	$20 \div 30$	0.67
Landscape	$18 \div 21$	0.86
Pet policy	$45 \div 26$	1.73

Lowest is the workday item. **Scheduled.** All 90 members go to tally 0.11. *(10 minutes used.)*

Round 2. Now the tallies are in play. Parking has 69 supporters currently at 0.11, so $(20 + 69 \times 0.11) \div 69 = \mathbf{0.40}$ — the lowest. **Scheduled.** Clusters A and C go to 0.40; cluster B stays at 0.11, since none of parking's minutes were charged to them. *(30 minutes used.)*

Round 3. Kitchen: $(25 + 48 \times 0.40) \div 48 = \mathbf{0.92}$. Landscape: $(18 + 21 \times 0.11) \div 21 = 0.97$. Budget variance: $(20 + 30 \times 0.40) \div 30 = 1.07$. Kitchen wins. **Scheduled.** Cluster A goes to 0.92. *(55 minutes used.)*

Round 4. Landscape at 0.97 now beats budget variance at 1.59 — cluster A has been charged a lot and cluster B almost nothing. Pet policy needs 45 minutes and only 35 remain. **Landscape is scheduled.** Cluster B goes to 0.97. *(73 minutes used.)*

Round 5. Budget variance needs 20 minutes; 17 remain. Nothing fits. Stop.

Final agenda: workday, parking, kitchen scope, landscape report.

This is the same agenda Equal Shares produces on the same ballots — reached without assigning anyone a budget, without tracking balances between meetings, and without a completion rule.

The self-check

Final tallies: cluster A at 0.92, cluster B at 0.97, cluster C at 0.40.

Multiply out: $48 \times 0.92 + 21 \times 0.97 + 21 \times 0.40 = \mathbf{73 \text{ minutes}}$, exactly the scheduled total. The tallies always sum to the minutes scheduled, which makes a useful arithmetic check when running this by hand.

The tallies also tell a readable story. A and B end up near each other: both clusters had a topic scheduled and both were charged roughly evenly for it. C ends lower because no item exclusive to them was on the ballot this month — they weren't shortchanged by the method, there was simply nothing of theirs to schedule. That's a backlog question, not a counting question.

Why this is the simpler option

The arithmetic is comparable to Equal Shares. The simplification is in everything that surrounds it:

Question	Equal Shares	Sequential Phragmén
How big is each member's budget?	must be decided	doesn't exist
What happens to unspent share?	carryover rule needed	nothing carries
Can someone bank influence by not voting?	yes — needs a cap	no
What if turnout varies?	needs adjustment	no adjustment
What if time is left over?	happens often — needs a completion rule	rarely happens
Can ballots be anonymous?	no, if carryover is used	yes

That last row is worth dwelling on. Every one of those questions is a decision the Process Group would have to make, defend, and later explain to someone unhappy about an outcome. Sequential Phragmén does not raise them.

On the leftover-time question specifically: across 400 simulated meetings, Equal Shares left an average of **29.7 of 90 minutes** unscheduled, while Sequential Phragmén left **5.7**. Equal Shares hands out budgets in advance and can strand purchasing power; Phragmén buys as soon as something becomes affordable and rarely leaves usable time behind.

Honest costs

The formal guarantee is weaker. Equal Shares satisfies a technical property called Extended Justified Representation; Sequential Phragmén satisfies a weaker relative called Proportional Justified Representation. This probably should not matter to us - these axioms were shaped as much by what could be mathematically proven as by what we might actually want. Satisfying the property most people would actually want ("no group could do better for themselves by taking its share and spending it alone") is an open research problem that no efficient method is known to satisfy.

Phragmén doesn't work with ballots more complicated than approval ballots. We're not planning to use ballots more complicated than approval ballots, so that's probably moot.

"Tally" is a less intuitive metaphor than "spending." Equal Shares can be explained as a purse of minutes that members spend, which is concrete. A tally of minutes charged to you is a step more abstract, and someone will reasonably ask why having a topic scheduled counts *against* you. The answer — that it's a record of community time used on your behalf, so we can keep it even — is satisfying once heard but does need saying.

No canonical priority order. With items of a single fixed length, a longer Phragmén agenda is always the shorter one plus additions, which would give a natural ranking. With varying durations that breaks: in simulation, a 90-minute agenda contained the 75-minute agenda only 76% of the time. So we can't publish a ranked list and cut at the line.

Running it on a whiteboard

Practical for something like 40 voters and 20 items.

The trick is that supporters' tallies are *equalized* every time an item is scheduled, so after a few rounds only three or four distinct tally values exist across the whole community. Don't track 40 individual numbers. Track the distinct values and how many people hold each — a table of three or four rows. Each round is then a handful of small multiplications, and the sum-check above catches arithmetic slips.

And because nothing carries forward, the whole thing can be run in public, on real ballots, with no record of who voted for what.