

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:

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

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 ÷ 90	0.11
Parking	20 ÷ 69	0.29
Kitchen	25 ÷ 48	0.52
Budget variance	20 ÷ 30	0.67
Landscape	18 ÷ 21	0.86
Pet policy	45 ÷ 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.

Revision #13

Created 2026-08-24 16:59:29 UTC

Updated 2026-09-02 22:44:17 UTC by Matt Elder