

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.

Revision #8

Created 2026-08-24 04:46:10 UTC

Updated 2026-08-24 17:55:48 UTC