FTP between scheduled payouts

Russ Oxley | Research paper | 4 September 2026

A fair transfer can leave the surviving pool unable to make its next transfer fairly, even where a different fair allocation preserves that next step. This paper develops and tests calculations for selecting and checking allocations that meet specified continuation requirements in heterogeneous pools.

The results support a practical division of work. Use ordinary FTP when an inexpensive check proves that every immediate successor remains feasible. Use a more general constrained allocation where that check identifies a problem requiring further investigation. For a stated number of cumulative deaths, a separate conservative certificate can cover every possible intervening allocation and deceased identity.

The calculations cover chronological deaths under a declared mortality model. They do not establish exact fairness for every sequence of calendar death batches. The payout horizon and treatment of uncovered events must still be specified as part of the contract.

The three calculations

Three calculations have been implemented and checked.

Component	Question answered	Evidence
Ordinary-policy check	Does this fair common-weight allocation leave every immediate successor feasible?	Independent interval calculation, linear in membership; 4,148 sampled states passed
Constrained allocation	Which allocation is closest to the reference while preserving next-step feasibility?	Exact four-member optimum; numerical heterogeneous cases through 100 members
Death-budget certificate	Is a fair next allocation available after every allowed redistribution through K deaths?	Outward-rounded certificates on the existing 1,000- and 10,000-member inputs

The first check is a filter, not a competing allocation objective. If the ordinary entropy optimum already meets the continuation requirement, adding that requirement cannot improve on it: the optimum remains unchanged. A failed sufficient check requires a stronger calculation; it is not automatically a reason to change the allocation or a proof of impossibility.

The larger pools usually needed no adjustment in these tests

We used the existing synthetic account populations and mortality inputs from the ONS UK life tables. The retired populations span ages 55-100; the mixed population spans 30-100. Published annual death probabilities supply constant mortality intensities, which are frozen for each one-year experiment. No new members, investment returns, ageing updates or payouts occur within that year. These controls isolate the effect of repeated mortality transfers. They are not a fitted projection of an actual pension membership. [1]

At every realised state, we recalibrated ordinary FTP, independently bounded its expected member gains and losses, and checked all possible immediate successor states. We then processed the next simulated death. The check covers every potential next decedent at that state, rather than only the decedent selected in that simulation.

Population	Simulated one-year paths	States audited	Largest certified next-successor loss share
Retired, 1,000 members	20	966	7.720%
Mixed ages, 1,000 members	20	819	6.235%
Retired, 10,000 members	5	2,363	0.987%

The loss share is a member's expected forfeiture as a proportion of the pool's total expected forfeiture at the next single death. Fair redistribution is feasible when no member exceeds 50%. The displayed percentages are upper bounds, rounded upwards, and remain well below that boundary.

All 4,148 state audits passed the £0.001 expected-error tolerance per member. The largest bound was below £8e-11. The 45 paths contained 4,103 deaths; the reported run took approximately 25 seconds. These bounds measure numerical fairness under the stated model, not mortality-model accuracy or the dispersion of realised outcomes.

The implication is limited but substantive. In these diversified states, the extra feasibility constraint does not require a different allocation. A cheap check can establish that fact. The tests do not estimate the probability of ever encountering a difficult state: each visited state is certified, but unvisited histories are not covered by that statement.

The replay uses fractional pound balances stored in binary64. Its independently checked allocation formula is defined on each frozen state. Integer-penny settlement, ledger posting and accumulated rounding in actual payments remain separate from this experiment.

The constrained method handles a case ordinary FTP cannot

The four-member example has a unique entropy optimum subject to the chosen continuation limit. Fairness, conservation, successor feasibility and the optimality conditions are checked exactly with rational arithmetic. This strengthens a feasible-witness result to an optimality result for this specified example.

The balances are 7, 6, 4 and 3. Members are equally likely to die next. Requiring every successor balance to remain at or below 9.75 puts its largest expected-loss share at no more than 48.75%, leaving a margin below the 50% boundary. The resulting transfer table differs from the earlier feasible witness and has a lower KL objective. Fairness, estate conservation, successor feasibility and the optimality conditions are all checked with rational arithmetic. A smaller KL objective is a property of the chosen selection rule; it is not a saving in fees or an improvement in expected member wealth.

A deliberately concentrated 100-member example tests heterogeneity. One member starts with 45% of expected forfeiture, another with 10%, and the remaining 45% is spread across 98 members. Mortality intensities vary between 0.9 and 1.1 in arbitrary common time units; balances are chosen to produce those shares. This is a concentration stress, not a representative large pension pool.

Ordinary FTP creates a successor with approximately 53.42% of expected forfeiture concentrated in one member. That successor cannot make its next transfer fairly. The constrained solution targets a maximum of 49.9% across all successors, while maintaining the current fairness and estate equations. One continuation restriction binds. The numerical solve, including a separate feasibility calculation, took about 3.6 seconds.

The general solution is numerical, so it also needs an independent acceptance test. A separate interval calculation audits the policy obtained by normalising each candidate column to distribute its deceased estate exactly. It bounds every member's expected error and every successor's expected-loss share. This certifies the proposed allocation's fairness tolerance and next-step feasibility; numerical optimality is assessed separately from the solver diagnostics. The four-member rational example has the stronger exact optimality proof.

For the 100-member stress, that independent audit took about 0.06 seconds and bounded the largest successor loss share below 49.901%. The expected member error was below 5e-15 in the example's balance units. The ordinary counterpart passed current fairness but was rejected for its successor concentration.

The general optimiser stores a full donor-recipient table. Its entropy iterations use memory proportional to the square of membership; the separate feasibility programme tested through 100 members has a larger coefficient count. These tests cover smaller or concentrated cases and do not demonstrate a replacement for the compressed large-pool batch solver.

A guarantee covering more than the sampled histories

The strongest easy-to-state safeguard asks a deliberately demanding question: after any set of at most K deaths, and any nonnegative redistribution of their estates, will a fair next single-death allocation still exist?

For each surviving member, the worst redistribution sends all the intervening estates to that member. We can calculate the resulting concentration without listing the possible death histories. The identities creating the greatest pressure are the K largest values of a simple score combining balances and mortality intensities. The appendix gives the expression.

This gives two implementations. A direct calculation examines those scores for each member. A faster sufficient bound sorts two lists once. Both use outward arithmetic, so a pass covers rounding error as well as the algebraic restriction.

Population	Death budget certified by the fast bound	Additional direct check passed
Retired, 1,000 members	12	13
Mixed ages, 1,000 members	10	11
Retired, 10,000 members	142	143

For example, the direct 10,000-member result says that every permitted state after up to 143 deaths still admits a fair next single-death allocation. It makes no assumption that the intervening credits were distributed by the particular FTP rule. The operating policy must nevertheless choose fair allocations: availability alone does not make its choices fair.

The direct checks took approximately 0.2, 0.3 and 6.8 seconds. The listed direct budgets are selected successful checks, not claims to be the largest possible. A failed bound does not establish that a specified fair policy fails. It may reflect a highly concentrated redistribution that that policy would never make.

That explains why the 1,000-member year-long replays can remain comfortably feasible after more than 13 deaths. The universal certificate admits much more extreme allocations than the rule used in those paths.

What the next payout date changes

A common proportional payout does not reset relative feasibility. If every surviving account retains the same positive proportion, every feasibility slack is multiplied by that proportion. A common investment return has the same scaling property. Member-specific distributions require their own state-transition calculation.

Member-specific payouts can change the result because they change relative exposures. With equal mortality, accounts (6,2,2) are infeasible; reducing the first account to 3 produces the feasible state (3,2,2). Conversely, (4,3,3) is feasible, but reducing the last two accounts to 1 produces the infeasible state (4,1,1). These are comparisons of specified payment schedules, not authority to change members' entitlements to make the mathematics work.

A calendar deadline differs from a death-count limit. Under fixed positive independent exponential intensities, every death order has positive probability before any positive deadline. Exact event fairness on every such history therefore retains the closed-run-off restriction when there are no intervening payouts or other state changes. Shortening the window reduces extreme-event probabilities; it does not remove those histories from the model.

The calendar coverage calculation encloses the probability of exceeding a selected death budget under the declared hazards. It uses an alternating-series bound for each member’s death probability and a positive recursion for the count tail. It does not obtain a tiny tail probability by subtracting an approximate cumulative probability from one.

Population and certified death budget	Chance of exceeding it within 30 days	Within 90 days
Retired 1,000: 13 deaths	At most 0.006781%	At most 30.397%
Mixed ages 1,000: 11 deaths	At most 0.039445%	At most 40.597%
Retired 10,000: 143 deaths	Below 0.000001%	At most 3.272%

These are illustrative windows, not assumed ladder terms. They show where the broad certificate is strong and where it becomes too conservative. They are probabilities of leaving that certificate's coverage, not probabilities that FTP becomes infeasible. In particular, the year-long ordinary-policy replays continue to pass after the broad death budgets have been exceeded.

A tail probability tells us how much of the model lies beyond the certificate. It does not specify how those estates would be settled. The research implementation can hold an uncovered event for further calculation, but that operational hold is not a completed contractual treatment of the event.

What follows for the research programme

The results comprise an exact constrained optimum, a tested heterogeneous optimiser, a separate table audit, a test that accepts ordinary FTP when admissible, and a conservative certificate covering all histories within a declared death budget.

The material remaining work concerns the contract and the general batch implementation. The ladder's actual member-specific distributions need to be inserted into the state transitions. Admission and mortality updates need explicit review points. Calendar batches need their own conditional identity laws and continuation constraints; chronological exponential death probabilities cannot be substituted for finite-period mortality odds.

The next practical design question is therefore the coverage and terminal treatment attached to each payout interval. We should measure the consequence of alternative specified rules, including rare-event handling, before asserting that the pension design can honour fairness throughout that interval. The results here make that comparison computable in several relevant cases. They do not establish empirical materiality for every pool or academic priority over the existing literature.

Technical appendix

A. The inexpensive check on ordinary FTP

Let s_i be balances, λ_i fixed positive exponential mortality intensities, and w_i the candidate ordinary FTP weights. Write W for the sum of weights, A for the sum of λ_i s_i, and H for the sum of λ_i w_i. After member j dies, define $d_j=s_j/(W-w_j)$. Survivor i holds s_i+d_j w_i.

The total next loss intensity is

$$A'_j = A + d_j(H - \lambda_j W).$$

Let $d_{\max}$ be the largest d_j. A sufficient condition covering every immediate successor is

$$\min_j A'_j > \max_i [2\lambda_i\,(s_i + w_i d_{\max})].$$

This relaxes the relationship between the deceased identity and the recipient when taking the extrema. Consequently, failure is inconclusive, but a strict pass proves all successor Hall inequalities. The calculation takes linear work and is independently enclosed by interval arithmetic. Current fairness is checked separately using expected credits under $p_i=\lambda_i$ divided by the sum of intensities.

B. The death-budget certificate

After deceased set D, conservation and nonnegative transfers give survivors s_i+x_i, where the x_i sum to the original balances of D. For a tested survivor i, the future feasibility slack is total surviving loss intensity minus twice i's loss intensity. Its minimum puts all x into i. The resulting exact condition for arbitrary redistributions through K deaths is

$$A - 2\lambda_i s_i \geq \mathrm{TopK}_{j\neq i}\left[(\lambda_i + \lambda_j)s_j\right] \quad \text{for every } i.$$

Here K leaves at least two survivors. Strict inequalities certify strict feasibility after every covered history. A faster sufficient bound replaces the right side by the sum of the K largest λ_j s_j, excluding i, plus λ_i times the K largest s_j, excluding i. Its two lists can select different decedents, which is the source of conservatism.

The direct method requires quadratic work in membership for a selected K; the separated method requires sorting and then linear work. Exact small-state checks independently enumerate vertices of the redistribution simplex. An underflow regression ensures that a reported deceased-set witness cannot include its tested survivor.

C. The exact constrained optimum

With balances (7,6,4,3), equal next-death probabilities and a successor balance ceiling of 9.75, the optimal credit table is:

Recipient	Death 1	Death 2	Death 3	Death 4
1	0	11/4	255/104	187/104
2	15/4	0	135/104	99/104
3	15/8	15/8	0	1/4
4	11/8	11/8	1/4	0

Rows and columns sum exactly to the starting balances. At every unconstrained permitted cell the credit equals $a_i\ b_j$, with a=(2805/208,1485/208,15/8,11/8) and b=(1,1,2/11,2/15). The two binding cells have ratios $a_1\ b_2/C_{12}$=255/52 and $a_2\ b_1/C_{21}$=99/52. These exceed one, giving positive multipliers for the upper bounds.

Drop the common probability factor one quarter and the separable-reference terms, which are constant under the row and column constraints. Row and column multipliers minus log a_i and minus log b_j, and upper-bound multipliers log(255/52) and log(99/52), then satisfy stationarity and complementary slackness. All feasibility identities are rational. Convexity proves optimality; strict convexity proves uniqueness. Multipliers can equivalently be rescaled by one quarter if that factor is retained; reference factors can be absorbed into the row and column multipliers.

The general interval table audit treats candidate entries as weights within each donor column. Its mathematical credit is the deceased balance times that entry divided by its column sum. This enforces exact estate conservation by definition, even if the numerical solver's raw column sum has a small residual. The resulting normalised policy is the one whose fairness and successor states are audited. The generic numerical optimum and its audited normalisation are not claimed to have an exact KKT certificate.

D. Reproduction and assumptions

All mortality intensities supplied to certificates are declared binary64 model inputs. In the population experiments they are computed once from the existing daily probabilities and then held fixed. The arithmetic guarantee begins with those declared values; it does not certify the logarithm used to estimate them or their empirical accuracy.

The interval routines assume correctly rounded IEEE-754 basic arithmetic, gradual underflow and no unsafe reassociation. The calendar routine additionally requires its stated range for the alternating expansion. The mathematical checks use independent rational calculations on small cases. None of these conditions validates source death records or posts a payment ledger.

The numerical procedures comprise a constrained optimiser, a rational fixture, a table certificate, a policy replay with interval checks, a death-budget gate, and payout and calendar-coverage calculations. The scope and arithmetic assumptions of each procedure are stated above.

[1] Office for National Statistics, National life tables: UK, 2022–2024. Population period mortality rates supply the source inputs for the synthetic population experiments.

Russ Oxley directed the research and its economic interpretation. AI systems contributed to derivations, calculations, code and drafting. Each paper states the assumptions and checks supporting its conclusions. Responsibility for the published work rests with the author.