



Intergenerational cross-subsidies in UK collective defined contribution funds

John Armstrong¹, James Dalby¹ and Catherine Donnelly²

¹King's College London, UK; and ²Heriot Watt University, UK

Corresponding author: John Armstrong; Email: john.armstrong@kcl.ac.uk

(Received 23 April 2025; revised 01 January 2026; accepted 15 February 2026; first published online 25 March 2026)

Abstract

We evaluate the performance and level of intergenerational cross-subsidy in flat-accrual and dynamic-accrual collective defined contribution (CDC) schemes, which have been designed to be compatible with UK legislation. In the flat-accrual scheme, all members accrue the benefits at the same rate, irrespective of age. This captures the most significant feature of the Royal Mail Collective Pension Plan, which is currently the only UK CDC scheme. The dynamic-accrual schemes seek to reduce intergenerational cross-subsidies by varying the rate of benefit-accrual schemes in accordance with the age of members and the current funding level. We find that these CDC schemes can often be successful in smoothing pension outcomes postretirement while outperforming a defined contribution scheme followed by annuity purchase at the point of retirement. However, this out-performance is not guaranteed in a flat-accrual scheme, and there is little smoothing of projected pension outcomes before retirement. There are significant intergenerational cross-subsidies in the flat-accrual scheme, which qualitatively mirror the cross-subsidies seen in defined benefit schemes, but the magnitude of cross-subsidies is much larger in flat-accrual CDC schemes. The dynamic-accrual scheme design seeks to reduce such cross-subsidies, but we find significant cross-subsidies still arise due to the approximate pricing methodology used to determine the benefits.

Keywords: CDC; Collective defined contribution

1. Introduction

Collective defined contribution (CDC) pension schemes have arrived in the UK after a long and fragile genesis (Wilkinson, 2024). These are pension schemes that share risks, such as longevity, investment, and inflation risk, collectively among the scheme members. This is different from individual defined contribution (DC) pension schemes, where members bear only their own risk. Unlike defined benefit (DB) pension schemes, the other main form of pension provision in the UK, benefit payments in CDC schemes are not guaranteed; indeed, they are expected to fluctuate. It is important to understand what risks are being borne by which members in CDC schemes and to quantify them, to reduce and mitigate against unfairness between different members.

We study two different whole-life CDC scheme designs. The first is a stylization of the Royal Mail Collective Pension Plan (RMCPP), a single-employer scheme that is the first CDC scheme in the UK. The second models the likely design of the multiemployer schemes that will launch in the next couple of years. Our aim is to compare these two scheme designs and evaluate the extent to which they succeed in their goals: providing an income for life, minimizing intergenerational unfairness, and smoothing pension outcomes. This is important to analyze because if they do not succeed in their goals, then the CDC industry will fail in the UK. The UK's success or failure is likely to influence the development of CDC schemes in other countries too.

© The Author(s), 2026. Published by Cambridge University Press on behalf of The Institute and Faculty of Actuaries. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (<https://creativecommons.org/licenses/by/4.0/>), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

In each of these schemes, the collective risk-sharing is done through declaring the same increases on all members' benefits, a UK regulatory requirement for CDC schemes. We refer to such schemes as *shared-indexation schemes* to highlight this common feature of their design.

The first shared-indexation scheme is a flat-accrual scheme, in which the contributions paid into the scheme are a fixed percentage of salaries. In exchange, they receive a deferred retirement income linked to their current salary. This scheme is closely modeled on that of the RMCPP, the first, and currently only, UK CDC scheme in existence. The design of RMCPP has strong similarities with the DB scheme it replaces. The desire to launch the RMCPP has been the driving force behind the introduction of the legislation enabling CDC schemes in the UK (Wilkinson, 2024). Combined with the scale of the scheme, which has over 100,000 members, the importance and significance of the RMCPP is clear. It is the prototype for CDC plan design in the UK, for UK regulations and regulators, for actuaries, employers, and workers' unions.

The second shared-indexation scheme we consider is a dynamic-accrual CDC scheme, which is the anticipated design of the multiemployer whole-life CDC schemes that are due to be authorized in 2026 in the UK (Wilkinson, 2024; Upton, 2024). Multiemployer decumulation-only CDC schemes, which have recently been consulted on by the UK's Department of Work and Pensions, are also expected to follow a dynamic-accrual approach. In such a scheme, members can choose how much to contribute and, in return, receive a deferred retirement income that reflects financial market conditions at the time the contribution was paid.

Due to the novelty of UK CDC schemes, with the RMCPP launched in October 2024, there is little literature on their unique design. A 2023 report prepared for the Association of British Insurers (ABI) (Taylor & Ward, 2023) summarizes numerous industry studies of similar designs, The Government Actuary's Department (2009); AON (2013); Popat *et al.* (2015); Taylor & Ward (2023); Blake (2016); Willis Towers Watson (2020), as well as providing its own additional analysis. However, there are important differences in details between the scheme designs in these studies and the RMCPP. Notable differences include the use of buffers and funding corridors. Of these studies, only two perform stochastic modeling (The Government Actuary's Department, 2009; Popat *et al.*, 2015). In the academic literature, Owadally *et al.* (2022) discuss CDC plans in the UK with respect to policy, as well as their risks and the ensuing advantages and disadvantages, and model an earlier proposal stochastically. They also compare the outcomes for a single member of their CDC plan model against typical DC-based alternative pension schemes when the scheme is mature. They find that their CDC plan can pay higher and more stable pensions compared to the alternative schemes considered.

This paper contributes to the literature in several ways. First, in Section 2, we describe an explicit mathematical model for shared-indexation CDC schemes. This model describes the contract of such a scheme. Once the realized contributions and investments are known, this model describes what the resulting benefits will be. In the case of the dynamic-accrual CDC scheme, this gives a precise operationalization of the broad shared understanding among pension professionals of what a dynamic-accrual CDC scheme should entail.

In describing a flat-accrual scheme, we follow the key design features of the RMCPP, but there are differences that we now summarize. Only a single-life pension in retirement is valued in our model – spousal benefits, death benefits, and the UK's tax-free cash lump sum payment made at the point of retirement are omitted. We include an upper bound on indexation, which is not present in the RMCPP, but is required for dynamic-accrual and so aids comparison between the two scheme types. We assume that benefit cuts and bonuses that must occur when the fund is significantly over- or under-funded take effect immediately, rather than over time, as is done for benefit cuts in the RMCPP. In our model, the relationship between contribution rate and level of benefits is determined by choosing a long-term target rate of indexation. In the RMCPP, these levels are fixed in the Trust Deed and Rules (Slaughter and May, 2024), but it is not specified how they are determined. Nevertheless, the member handbook does suggest that “The aim is to provide increases to help keep up with the cost of living” (Royal Mail Collective Pension Plan, 2024, p. 67).

We endeavor to make a clear separation between the contract of a CDC fund and our modeling assumptions. Section 2 describes the contract. Section 3 describes modeling assumptions, including demographic and economic factors.

Having established the contract of CDC funds, the second way our paper contributes to the literature is by studying the extent of intergenerational cross-subsidy in CDC funds. This is the topic of Section 4.

We find that flat-accrual schemes and dynamic-accrual schemes contain very different levels of cross-subsidies. A dominant source of cross-subsidy in the flat-accrual scheme results from the combination of a fixed contribution rate and a fixed-accrual rate. This favors participants who are close to retirement, as they receive the same benefit amount as younger members for the same contribution, despite there being less time for their investments to grow. This qualitatively mirrors the behavior of the DB scheme that the RMCPP flat-accrual scheme is designed to replace. However, when examined quantitatively, we find that the level of cross-subsidy is considerably higher in a flat-accrual scheme than in a DB scheme.

Since the earliest generations in a flat-accrual scheme receive an income that is worth more than the value of their lifetime contributions, subsequent generations must receive an income that is worth less than their contributions. We find that the flat-accrual scheme broadly tends to a long-term steady-state situation where typical generations receive a retirement income worth less than their contributions. One can think of this as later generations paying interest on the debt built up by overpaying the earliest generations. We refer to this effect as “drag.” It may appear counterintuitive that a scheme can enter a steady state that is not actuarially fair. This is possible because over an infinite time horizon, one cannot argue that the total money paid into the scheme must equal the total money paid out, so we will also refer to this as an “infinite-horizon effect.”

The dynamic-accrual scheme is designed to avoid the obvious cross-subsidies of a flat-accrual scheme. It is instead intended to be *actuarially fair*, in the sense that members’ contributions match their new benefit entitlements at the time of each contribution. However, the implicit pricing methodology that has been proposed in industry to do this is only approximate, and so cross-subsidies still appear. These cross-subsidies are stochastic in nature, meaning one cannot predict in advance that generations will lose out.

We briefly consider two ways to mitigate the cross-subsidies we have identified. If a flat-accrual scheme is designed to replace a DB scheme, one might wish to design a scheme that has similar levels of cross-subsidy to the DB scheme it seeks to replace. This could be done using an age-based accrual rather than a flat-accrual. In a dynamic-accrual scheme, we show how to statistically calibrate a pricing formula to obtain fairer prices for benefits than the formula proposed by the industry. These mitigations are able to manage the level of cross-subsidies without departing from a shared-indexation design, but do not eliminate them entirely.

We use two distinct methodologies to examine the cross-subsidies. The cross-subsidies in a flat-accrual scheme and the notion of drag can be studied using a constant economic model (by which we mean a deterministic model where inflation, wage inflation, interest rates, and stock growth over each period are all constant). This allows us to find analytic results. However, the cross-subsidies in a dynamic-accrual scheme require stochastic modeling. In this case, we use a Black–Scholes economic model in order to ensure that any pension income can be priced unambiguously using the theory of risk-neutral pricing. The cross-subsidies then emerge from the difference between the true risk-neutral price and the discounting formulae used to determine how much members are charged for their pension benefits. In the case of flat-accrual schemes, stochastic modeling confirms the results obtained using a constant economic model.

Cross-subsidies have been studied in other types of CDC schemes, and the high cross-subsidies and consequent drag are a specific consequence of the flat-accrual design. For example, Haan, Lekniute, and Ponds (2015) study a Dutch-style CDC scheme and show a greater financial fairness of those pension schemes.

The third way we contribute to the literature is by simulating the behavior of these funds using a stochastic economic model and comparing the performance to other scheme designs. This is done in Section 5, using the model described by Alvares Maffra, *et al.* (2021). We compare the performance with the alternatives of DC schemes and a pooled annuity fund. Similar stochastic models have been used to study CDC designs that have been proposed for use in the UK, e.g., Popat *et al.* (2015); Owadally *et al.* (2022), but not for the current shared-indexation design of the RMCPP.

In all our modeling in this paper, we assume there is no systematic longevity risk, and there are sufficient numbers of members so that individual longevity risk is eliminated through pooling. We do not allow for any heterogeneity in the schemes in our simulations. However, we do show how heterogeneity can be incorporated into the scheme design if desired; results from simulating heterogeneous funds are discussed in Upton (2024).

The fourth contribution to the literature is studying the extent to which shared-indexation schemes are able to smooth pension incomes. This is considered in Section 6.

The design of shared-indexation schemes is such that younger generations experience greater volatility in the value of their future benefits, while older generations experience a smoothing in this value. This allows the fund to invest in riskier assets while still providing a smooth income in retirement. Other CDC plans studied in the literature have the same feature, and they can also outperform alternative pension options by taking on more investment risk (Bonenkamp & Westerhout, 2014; Cui *et al.*, 2011; Gollier, 2008).

We find that the shared-indexation designs do provide some smoothing of benefits after retirement, but they do not smooth projected benefits before retirement by as much as has been suggested (AON, 2013). The discrepancy between our results and the literature can be attributed to the fact that none of the studies we examined attempted to compute projected benefits using a stochastic model, and the deterministic approximations used to project benefits are not sufficiently accurate for this purpose.

Our fifth contribution is in Section 7, where we consider how a scheme designer might control the outcomes in a CDC fund. We find that it is quite easy in a flat-accrual scheme to target a particular growth rate for median incomes in retirement by adjusting the relationship between benefits and contributions. This flexibility is not available in a dynamic-accrual scheme, making it difficult to design a scheme so that it will produce a desired level of indexation.

Our final contribution, in Section 8, is to briefly consider model risk by evaluating the cross-subsidies that occur when actual investment returns differ from their predictions. Since the level of pensions paid out depends on future predictions of returns, poor predictions of those returns lead to pension payments that are, in hindsight, either too low or too high. Consequently, future generations have higher or lower pensions. This type of cross-subsidy can go either way: it may benefit earlier or later generations to join. However, it is also true that once the plan begins its life, the cost due to previous predictions deviating from their observed expected values could be calculated.

Our key findings are that:

- dynamic-accrual schemes are not able to target a given level of annual pension increases as effectively as flat-accrual schemes;
- dynamic-accrual schemes have substantially lower intergenerational cross-subsidies than flat-accrual schemes;
- dynamic-accrual schemes still feature some level of intergenerational cross-subsidy as a result of the use of approximate pricing formulae in the design of CDC funds;
- flat-accrual schemes do not necessarily have higher pension outcomes than DC followed by life annuity purchase at retirement;
- dynamic-accrual schemes can markedly outperform flat-accrual schemes because they do not experience drag effects;

- shared-indexation CDC schemes do not result in as much smoothing of projected pension benefits as has been suggested in the literature;
- nominal benefits are not an accurate means for estimating projected benefits.

2. The contract of shared-indexation CDC funds

In this section, we detail the contractual aspects of shared-indexation CDC funds. The aspects of their operation that are common to both flat- and dynamic-accrual funds, namely, the annual indexation of accrued pensions, are described first. We then explain the different approaches taken in the flat- and dynamic-accrual CDC funds to the accrual of benefits. In online Appendix D, we explain briefly how our description of CDC funds corresponds to the UK regulations.

2.1 Elements common to all shared-indexation CDC funds

2.1.1 Scheme membership and contributions

To describe the operation, we assume there are a total of M different types of individual investing in the fund over time, indexed by $\xi \in \{0, 1, 2, \dots, M-1\}$. All individuals of type ξ share the same demographic and financial characteristics. For the numerical examples in this paper, ξ is simply a generation index (Upton (2024) studies the situation when individuals vary by both generation and sex). Each individual of type ξ who is below the retirement age makes a contribution C_t^ξ at time t and receives, in exchange, an additional nominal benefit amount B_t^ξ in retirement.

2.1.2 Cumulative benefit and new benefit accrual amounts

At any time t , each individual of type ξ has accrued a cumulative nominal benefit of annual amount $B_{t-}^{\xi, \text{cum}}$ before new contributions. At time 0, this is equal to 0 and we will describe inductively below how the value is computed at subsequent times. Assuming the value $B_{t-}^{\xi, \text{cum}}$ is known at time t , after paying a contribution at time $t \in \mathbb{N}_0$, the individual accrues an additional annual nominal benefit of amount B_t^ξ . This increases their cumulative nominal benefit to

$$B_t^{\xi, \text{cum}} := B_{t-}^{\xi, \text{cum}} + B_t^\xi.$$

Members do not accrue benefits until they begin contributing to the scheme (which may be at times $t > 1$; the notation considers all members who are ever in the scheme, even those not currently in the scheme). Retired individuals are paid $B_t^{\xi, \text{cum}}$ at time t . For individuals who have not yet retired, $B_t^{\xi, \text{cum}}$ represents their currently accrued benefit entitlement before allowance for future pension increases and benefit adjustments.

2.1.3 Pension increase mechanism

Each year, all accrued pensions are increased by the same pension increase factor $\theta_t(1 + h_t^{\text{nom}})$, in which θ_t is a *bonus level* and h_t^{nom} is the *nominal indexation rate*. It is currently a UK regulatory requirement that scheme members must receive identical adjustments to their accrued benefits. This is the defining characteristic of a shared-indexation scheme.

The procedure for determining the pair $(h_t^{\text{nom}}, \theta_t)$ is described next. Denote by $q_t = \frac{\text{CPI}_t}{\text{CPI}_{t-1}} - 1$ the annual effective rate of change in the Consumer Price Index, CPI_t , from time $t-1$ to time t . Define the *real indexation rate* h_t , at time t , as the solution to

$$1 + h_t^{\text{nom}} = (1 + h_t)(1 + q_t).$$

After the pension increase factor is determined at time t , the total benefit accrued by each individual of type ξ becomes

$$B_{t-}^{\xi, \text{cum}} := \theta_t(1 + q_t)(1 + h_t)B_{t-1}^{\xi, \text{cum}}, \quad \text{for } t \geq 1. \quad (1)$$

2.1.4 Calculation of pension increases

Broadly, to calculate the real indexation rate h_t at time t , first, the accrued benefits are projected from time t onwards assuming that the future real indexation rate remains h_t at all future times $s \geq t$. Discounting the projected pension payments back to time t , their total value is equated with the market value of the scheme's assets. The resulting expression can be solved for h_t .

The design spreads the impact of market shocks over time, through changes to long-term indexation. We explain in more detail below how this smooths pensions in retirement. The projection of accrued pensions using an annual pension increase rate reflects the approach taken in UK DB schemes, although DB schemes typically apply a predetermined rate, such as CPI. The design of the RMCPP was strongly influenced by DB scheme principles, embedding those features into its structure. As UK regulations were developed to authorize the RMCPP, this approach became embedded in the regulatory framework.

However, there are some critical elements necessary to complete the calculation. The real indexation rate h_t is constrained to lie in the range $[-q_t, h_t^{\text{upper}}]$, in which $h_t^{\text{upper}} \geq -q_t$ is chosen by the scheme manager. Although in the RMCPP, a flat-accrual scheme, $h_t^{\text{upper}} = \infty$, we will assume a finite upper limit for our studied flat-accrual scheme. In the considered dynamic-accrual scheme, $h_t^{\text{upper}} < \infty$ and the goal is to ensure that the nominal indexation rate awarded at time t on accrued benefits is neither negative nor too high. The use of such constraints on the value of h_t is embedded in regulation. It is viewed by practitioners as necessary to ensure that the scheme provides planned benefit increases that will be acceptable to members.

If the solution h_t to the expression involving the scheme's asset value and discounted benefits would imply breaching either of these limits, then the real indexation level is set equal to the relevant limit and a bonus level θ_t is applied to the accrued benefit entitlements.

The intention is that, when the limits on h_t are not breached, then $\theta_t = 1$ and changes to benefits arise only from the indexation rates. If, for example, the solution implied $h_t < -q_t$, then $h_t := -q_t$ and the bonus level falls below one, i.e., $\theta_t < 1$, corresponding to a nominal benefit cut. The calculation of the value of θ_t is explained more fully below.

Let:

- N_t^ξ be the number of surviving individuals of type ξ in the fund at time t .
- $p_{t,k}^\xi$ be the probability that an individual of type ξ survives from time t to time $t+k$, given that they are alive and age x at time t . It is assumed that future lifetimes are independent random variables.
- $v_{t,k}^\xi$ be the cumulative discount factor for an individual of type ξ , that discounts 1 unit paid at time $t+k$ back to time t . Let $\hat{R}_{t,k}^\xi$ be the expected annual effective rate of return at time t of the assets invested, in the time period $[\max\{t+k-1, t\}, t+k)$, in line with the investment strategy associated with individuals of type ξ . Then, for $k \in \mathbb{N}_0$,

$$v_{t,k}^\xi = \prod_{n=0}^k (1 + \hat{R}_{t,n}^\xi)^{-1}.$$

Note that $\hat{R}_{t,0}^\xi = 0$. We will describe how the asset mix associated with an individual is determined in our models in Section 3.4. When the scheme operates in practice, the appropriate rates of return for different asset classes are determined by regulatory reporting standards, but in our simulations, they are determined by our models. The only asset classes used in our simulations are an equity market index and index-linked bonds.

- $\hat{q}_{t,k}$ be the projection made at time t , of the annual change in CPI from time $t+k-1$ to time $t+k$, for $k = 0, 1, 2, \dots$. Note that $\hat{q}_{t,0} = q_t$, the actual change in CPI over the year to time t .

- $I_{t,k}(h)$ be the projection made at time t of the cumulative nominal indexation rate from time t to time $t + k$, using a constant real indexation rate h , for $k = 0, 1, 2, \dots$, which is calculated as

$$I_{t,k}(h) := \prod_{n=0}^k (1 + \widehat{q}_{t,n})(1 + h).$$

- $\mathbf{1}_t^{\xi, \text{Ret}}$ be the deterministic indicator function taking the value 1 if individuals of type ξ have retired at time t , and 0 otherwise.
- A_{t-} be the value of assets at time t , just before new contributions are added or benefit payments are made at time t .

We now define $\mathcal{L}_{t-}(h, \vartheta)$ to be the discounted value of the accrued benefits at time t , just before new contributions are added or benefit payments are made, where the real indexation rate used to project the accrued pensions from time $t - 1$ is h , and the bonus level applied at time t is ϑ . Thus,

$$\mathcal{L}_{t-}(h, \vartheta) := \vartheta \sum_{\xi=0}^{M-1} N_t^{\xi} B_{t-1}^{\xi, \text{cum}} \sum_{k=0}^{\infty} p_{t,k}^{\xi} \nu_{t,k}^{\xi} I_{t,k}(h) \mathbf{1}_{t+k}^{\xi, \text{Ret}}. \quad (2)$$

Equating the asset value with the discounted benefits yields a function of the unknown real indexation rate h and the unknown bonus level ϑ , i.e.,

$$A_{t-} = \mathcal{L}_{t-}(h, \vartheta) \quad (3)$$

A solution for the pair (h, ϑ) is first sought by setting $\vartheta := 1$ in (2) and solving for the constant h . There are three possible outcomes at each time t .

- If $h \in [-q_t, h_t^{\text{upper}}]$, then $h_t := h$ is the real indexation rate declared at time t , and no bonus is awarded, i.e., $\theta_t := 1$.
- If $h < -q_t$, then the real indexation rate declared at time t is $h_t := -q_t$. The expression (2) is solved for ϑ with the free variable $h := -q_t$. The solution is the declared bonus, i.e., $\theta_t := \vartheta$, and it corresponds to a nominal benefit cut, i.e., $\theta_t < 1$.
- Similarly, if $h > h^{\text{upper}}$, then $h_t := h^{\text{upper}}$ and expression (2) is solved for ϑ with the free variable $h := h^{\text{upper}}$. The solution is the declared bonus, that is, $\theta_t := \vartheta$, corresponding to a uniform increase in the accrued benefits, i.e., $\theta_t > 1$.

The net effect is to ensure that assets match discounted liabilities at all times, that h is appropriately constrained, and that benefit cuts and bonuses only occur if the constraints would be violated.

In Figure 1, we show a stylized picture of how the projected benefits change under two different pension increases. The dotted line depicts the projected benefits over time just before a shock occurs. In the example shown, markets have performed better than expected. Suppose that the concomitant increase to benefits can be distributed either as a uniform one-off increase by changing the one-off bonus level ϑ , shown as a dashed line, or by changing the nominal indexation rate, shown as a solid line.

The shaded regions illustrate the projected retirement benefits that are expected by two different generations. The left-hand colored block in Figure 1 shows the income received during retirement for a generation that has already retired. Its area represents the total income. The relative increase in their benefits when the nominal indexation rate h increases, depicted with lighter shading, is small.

Similarly, the right-hand colored block indicates the retirement income for a generation that is yet to retire. The paler shaded area is much larger, showing that their retirement benefits are much more sensitive to changes in h . The increase in their annual pension in retirement is higher than that of the retired generation, compared to before the shock. If the shock results in a reduction




Figure 1. A stylized diagram of how increased benefits might be distributed among different age groups when the market outperforms expectations.

in benefits, then the young would have a larger relative decrease in their pension at retirement, compared to the old.

Thus, using annual indexation rates results in lower volatility in the annual benefit amount paid to older generations. In exchange, there is a greater fluctuation in projected benefits for the younger generations. In other words, there is a risk transfer from the older to the younger generations. This argument is heuristic and should be seen as motivational; a more precise argument is made in Section 6.

In contrast, using only a one-off bonus level means that annual benefit amounts change by the same percentage. No pension smoothing would take place.

However, in our shared indexation schemes, the actual adjustment to benefits is a combination of an indexation rate (which is compounded annually to project future benefits) and a one-off bonus (which is applied only at its calculation time and is not reapplied in each projected year). The combination of using both indexation rates and one-off bonuses means that there are limits on the transfer of risk from the old to the young.

In RMCPP, nominal benefit cuts (corresponding to $\theta_t < 0.95$) are spread over two or three years with the goal of trying to keep reductions to at most 5% each year. Positive bonuses (i.e. $\theta_t > 1$) can occur if they are used to mitigate against these multiyear reductions. In our flat-accrual scheme design, benefit cuts and bonuses are applied in full immediately, for simplicity.

The design does not allow individuals to customize their investment strategy in line with their risk preferences. The benefit increases received by different individuals are uniform across members and are not dependent on their chosen investment strategy. The purpose of associating an investment strategy with different individuals is to ensure that the asset mix of the fund will change if the demographics of the fund change significantly (as happens during start-up and shut-down of the scheme). When the scheme closes, the fund will invest more heavily in bonds. This is broadly intended to compensate for the reduced smoothing that will occur, but is only heuristically justified.

The inability to accommodate different risk preferences is a weakness of the shared-indexation design. However, proponents of the approach argue that most members want all decisions to be made for them, making a one-size-fits-all approach acceptable.

2.1.5 Asset and discounted benefit value dynamics

The scheme's asset value evolves as

$$A_{t-} = \begin{cases} 0 & \text{if } t = 0, \\ (1 + R_t)A_{t-1} & \text{if } t = 1, 2, \dots \end{cases}$$

Here, R_t denotes the return realized by the fund on its investments over the time period $(t-1, t]$.

The assets of the fund at time t , after new payments have been received and annual pension payments have been made, are

$$A_t := A_{t-} + \sum_{\xi=0}^{M-1} N_t^{\xi} (C_t^{\xi} - B_t^{\xi, \text{cum}} \mathbf{1}_t^{R, \xi}). \quad (4)$$

After new payments have been received and annual pension payments have been made at time t , the discounted benefit value is

$$\mathcal{L}_t := \sum_{\xi=0}^{M-1} N_t^{\xi} B_t^{\xi, \text{cum}} \sum_{k=1}^{\infty} p_{t,k}^{\xi} v_{t,k}^{\xi} I_{t,k}(h_t) \mathbf{1}_{t+k}^{\xi, \text{Ret}}. \quad (5)$$

Note that the benefits paid out at time t are excluded from the calculation of $\mathcal{L}_t$. Additionally, as θ_t is already incorporated into $B_t^{\xi, \text{cum}}$, the bonus level is neutral (i.e., $\theta_t = 1$) in this calculation. For brevity, we will sometimes call the quantity $\mathcal{L}_t$ the liability at time t , although, since the fund does not promise any particular pension payment, it is not strictly speaking a liability.

It is straightforward to show that

$$\begin{aligned} \mathcal{L}_t = \mathcal{L}_{t-}(h_t, \theta_t) &= \underbrace{\sum_{\xi=0}^{M-1} N_t^{\xi} B_t^{\xi, \text{cum}} \mathbf{1}_t^{\xi, \text{Ret}}}_{\text{Benefits paid out at time } t} + \underbrace{\sum_{\xi=0}^{M-1} N_t^{\xi} B_t^{\xi} \sum_{k=1}^{\infty} p_{t,k}^{\xi} v_{t,k}^{\xi} I_{t,k}(h_t) \mathbf{1}_{t+k}^{\xi, \text{Ret}}}_{\text{Discounted new benefits accrued at time } t}. \end{aligned}$$

Since $A_{t-} = \mathcal{L}_{t-}$, substitution from (4) leads to

$$\begin{aligned} \underbrace{A_t - \mathcal{L}_t}_{\substack{\text{Asset value less} \\ \text{discounted benefits} \\ \text{at time } t}} &= \underbrace{\sum_{\xi=0}^{M-1} N_t^{\xi} C_t^{\xi}}_{\text{Contributions paid at time } t} \\ &\quad - \underbrace{\sum_{\xi=0}^{M-1} N_t^{\xi} B_t^{\xi} \sum_{k=1}^{\infty} p_{t,k}^{\xi} v_{t,k}^{\xi} I_{t,k}(h_t) \mathbf{1}_{t+k}^{\xi, \text{Ret}}}_{\text{Discounted new benefits accrued at time } t}. \end{aligned} \quad (6)$$

The last expression shows that if new contributions do not equal the value of the newly accrued benefits, then there will be an immediate scheme gain or loss at time t . This is in spite of the asset value equaling the discounted benefits immediately before the contributions are paid at time t , i.e., $A_{t-} = \mathcal{L}_{t-}(h_t, \theta_t)$.

For this reason, the contribution rate is chosen so that, when a flat-accrual CDC scheme has a stable population, the right-hand side of equation (6) is zero. As the same contribution rate is used in the dynamic-accrual scheme for our numerical illustrations, there are different changes in the scheme's funding position immediately after new contributions are made, depending on whether the scheme is a flat- or dynamic-accrual scheme, as discussed in Sections 2.2 and 2.3.

2.2 Flat-accrual scheme design

A flat-accrual CDC scheme has a design similar to a typical UK DB scheme. Like DB schemes, income benefits in flat-accrual CDC schemes are accrued at a constant proportion of salaries in exchange for a fixed rate of salaries being paid as contributions. This means that a member who is approaching retirement accrues the same amount of annual benefit as a member at the start of their career, assuming they pay the same contribution amount.

This leads to a difference in the value of the benefits accrued to the value of the contributions made in respect of those accrued benefits. Younger active members contribute more than the value of the benefits they are accruing, whereas older active members contribute less.

Let S_t^ξ denote the salary of individuals of type ξ , at time t . We write $\mathbf{1}_t^{C,\xi}$ for the indicator function taking the value 1 if individuals of type ξ make a contribution in year t , and 0 otherwise.

Let the constant $1/\beta$ be the accrual rate, that is, the proportion of salaries accrued as an annual benefit at time t by all currently contributing members. Then, the benefit amount accrued at time t , by each individual of type ξ , due to the contributions made at time t in the flat-accrual model is

$$B_t^\xi = \frac{1}{\beta} S_t^\xi \mathbf{1}_t^{C,\xi}, \quad (7)$$

which has a discounted value at time t of

$$\frac{1}{\beta} S_t^\xi \mathbf{1}_t^{C,\xi} \sum_{k=1}^{\infty} p_{t,k}^\xi v_{t,k}^\xi I_{t,k}(h_t) \mathbf{1}_{t+k}^{\xi, \text{Ret}}. \quad (8)$$

Let the constant $\alpha > 0$ be the proportion of salaries, which are paid at time t in respect of all members who make a contribution at time t . Then, the amount of contribution made at time t , in respect of each individual of type ξ , is

$$C_t^\xi = \alpha S_t^\xi \mathbf{1}_t^{C,\xi}. \quad (9)$$

It is assumed here that neither the contribution rate α , nor the accrual rate $1/\beta$, changes over time. In reality, both may change as a result of management decisions, but this is anticipated to happen infrequently.

The contribution rate α is chosen so that the right-hand side of equation (6) is zero when the scheme has a stable population. This approach implies that there is collective financial fairness at the time of each contribution payment. However, there is no individual financial fairness at each time. The contribution C_t^ξ is, in general, not equal to the discounted value of the newly accrued benefit, which is seen by comparing (8) to (9). Moreover, the total contributions may not match the total additional discounted benefits. Consequently, new benefit accrual at time t causes a sudden change in the scheme's overall funding position, which can be calculated using (6). This movement away from a 100% funding position is only rectified at the next year's valuation. The gain or loss due to the new contributions is then shared among all the scheme members through the pension increase factor. As a result, the level of indexation will typically vary from year to year, even if one assumes a constant economic model.

If we assume a deterministic model, then this gives rise to dynamical equations for h_t . Suppose we allow the number of types of individuals, M , and the time, t , to tend to infinity. Assume that the demographics of the fund and the economic factors all converge to long-term limits. In that event, we expect that h_t will also converge to a long-term limit h^∞ given by solving the equation

$$\lim_{t \rightarrow \infty} \frac{\sum_{\xi=0}^{\infty} \frac{1}{\beta} S_t^\xi \mathbf{1}_t^{C,\xi} \sum_{k=1}^{\infty} p_{t,k}^\xi v_{t,k}^\xi I_{t,k}(h^\infty) \mathbf{1}_{t+k}^{\xi, \text{Ret}}}{\sum_{\xi=0}^{\infty} \alpha S_t^\xi \mathbf{1}_t^{C,\xi}} = 1. \quad (10)$$

The numerator in this equation is determined by equation (8) and represents the total discounted benefits at time t . The denominator is determined by equation (9) and represents the total contributions made at time t . If the demographics of the contributors to the scheme stabilize from some time τ onwards, the fraction on the left-hand side will become constant for $t \geq \tau$, so instead of the limit, one can simply work with values of t that are chosen to be sufficiently large. In the demographic model, we use in our simulations, one may take any $t \geq 0$.

In a constant economic model, equation (10) can be used to determine the ratio between the contribution rate α and the benefit accrual rate $\frac{1}{\beta}$, which gives a desired target level of indexation $h^\infty = h^{\text{target}}$. More generally, in a stochastic, but stationary, economic model, one can compute median estimates for all terms and use equation (10) to roughly estimate the ratio of contributions to benefits that will achieve a target long-term median indexation. Conversely, however, one chooses the ratio of contributions and benefits, this will determine the long-term median indexation of the fund.

In a flat-accrual scheme, the mismatch in a given year between a member's discounted benefit and their contributions may be substantial. It is important to consider whether this is fair to members. In practice, the contribution calculated in equation (9) is really the sum of the member contribution and the sponsoring employer's contribution made on behalf of a particular member (as is the case in (11), too). However, formally, the employer's contribution is not allocated to any specific member, only to the fund as a whole. This enables money to be distributed unevenly across generations, which, as we discuss in detail later, is an essential feature of this type of "DB-lite" scheme.

If a member is in the scheme for their entire working life, one might imagine that lifetime benefits will match lifetime contributions. However, in general, this will not be the case. This is because the earliest generations receive disproportionately high benefits and later generations must, in effect, pay interest on the debt this creates. As a result, later generations receive less than they pay in, an effect we call drag. It is only in the long-term steady state that all scheme members are treated equally, and in this steady state, no member will receive as much benefit as they pay in. We quantify this effect in Theorem 1 below.

The use of flat-accrual means that if a member leaves the scheme early, they can expect to have overpaid for the value of what they can transfer out of the scheme. This is considering the contributions paid by both member and employer. Similarly, it is not fair for existing members if someone close to retirement joins and their contributions are not sufficient to cover the value of their accrued benefits. This perception of unfairness may exist even if the employer's contribution is not attributed to each individual member, but rather to the scheme as a whole. Existing members would suffer a decrease, albeit likely small, to make up for this individual deficit. Lastly, the calculation of the annual pension increases on accrued benefits can lead to significant changes in accrued pensions, depending on the scheme's maturity and asset value. We discuss this further in Section 4.

2.3 Standard dynamic-accrual scheme design

The term dynamic-accrual is intended to suggest that accrual rates vary with market conditions. We define a *standard dynamic-accrual CDC scheme* to be one where the value of the new benefit accrued by a contribution matches the amount of that contribution.

A standard dynamic-accrual CDC scheme is intended as an alternative to a DC fund and could be well-suited to multiemployer schemes. In a multiemployer scheme, it is important to minimize intergenerational cross-subsidies as different employers may have different demographics of their workforce and would presumably not wish to provide an overall subsidy to another employer (see McNally (2023)).

In our studied dynamic-accrual scheme, the new benefit accrued at time t by the corresponding contribution C_t^ξ made at time t , is given by defining the price of a unit of nominal benefit by

$$\hat{V}_t^\xi := \frac{1}{I_{t,0}(h_t)} \sum_{k=1}^{\infty} I_{t,k}(h_t) v_{t,k}^\xi p_{t,k}^\xi \mathbf{1}_{t+k}^{R,\xi}, \quad (11)$$

and then choosing B_t^ξ to solve

$$C_t^\xi = B_t^\xi \hat{V}_t^\xi. \quad (12)$$

By equation (6), this ensures that

$$A_t - \mathcal{L}_t = 0. \quad (13)$$

The basis used in (11) is the same as that used for the calculation of the pension increase awarded at time t . Assuming all interest rates are positive, the ratio of benefits to contributions will decrease with age, reflecting the time value of money.

Using the valuation basis to calculate the new accrued benefits means that there is no change in the scheme's overall funding position due to new benefit accrual. Since the new contributions made at time t are equal to the value of the new accrued benefits, the right-hand side of (6) equals zero. Consequently, if the valuation basis is borne out in practice, then the asset value will automatically equal the discounted benefit value at the next and subsequent time periods, for the selection $(h, \vartheta) := (h_t, 1)$, i.e., equation (3) holds at the next time.

For the contributions paid into the scheme at time 0, an assumption is needed about the level of future pension increases in order to calculate the first benefits accrued. We assume that the real indexation rate is a constant, h^{target} , at time 0 and hence set $h_0 := h^{\text{target}}$ in (11). For contributions made at time $t = 1, 2, \dots$, the real indexation rate h_t calculated at time t_- is used in (11).

We have seen that when designing a flat-accrual fund and in a mean-reverting economic model, it is possible to choose a target level of indexation h^{target} and set the contribution and benefit accrual rates to ensure h_t tends to revert to this level in the long term. There is no equivalent mechanism to set long-term indexation in a standard dynamic-accrual fund. Setting $h_0 = h^{\text{target}}$ will have a negligible influence on the values of h_t for large t . Instead, if one wishes to target a long-term median level for h_t , one must choose the bounds on h_t to limit its fluctuations and use benefit cuts and bonuses to constrain h_t within reasonable limits. We will see a numerical example of this in Section 7.

As nominal indexation rates, discount rates, and expectations of future lifetimes will vary over time, the amount of benefit calculated using (11) will also vary continuously for a fixed contribution amount. For occupational pension schemes in the UK, it is usual to have infrequently changing conversion factors for members. This is unlike life insurance companies, which may update their life annuity conversion rates daily. Although the trustees of a CDC pension scheme may prefer to have a rarely changing table of values for members, we do not explore the implications here, as we focus on dynamically changing benefit accrual only.

In a flat-accrual scheme, $v_{t,k}^\xi$ varies with ξ and will ultimately be a function of a member's age to give a lifestyle strategy (see Section 3.4). Some of the pension providers we spoke to were considering developing dynamic-accrual CDC funds where v depends upon ξ in the same way. However, as all funds are pooled and subject to the same risks, if one allows v to vary with ξ , this will result in different generations being charged different amounts for the same future cashflows. This runs against the design goal of minimizing intergenerational cross-subsidies, as can be confirmed numerically using the techniques of Section 4, where we explore cross-subsidies in detail. To avoid this issue, we selected an overall strategy for the dynamic-accrual fund that was designed to reflect the cumulative effect of each generation's lifestyle strategy and chose the discount rate $v_{t,k}^\xi$ to match this overall strategy, and hence to be independent of ξ . See Section 3.4.1.

3. Modeling details

3.1 Demographic assumptions

We will assume that the following modeling assumptions hold in all our simulations except where we explicitly state otherwise.

Table 1. Long-term medians in our economic models

Economic variable	Long-term median
Stock growth	7.73% p.a.
Wage growth	3.83% p.a.
CPI growth	2.00% p.a.
Index-linked bond growth	4.36% p.a.

The earliest entry age to the scheme is age $x_0 := 25$. An individual receives their first pension payment at age $\text{NRA} := 65$ and pension payments are made annually in advance. Every individual is assumed to be dead at age $\omega := 120$, the limiting age. The fund is closed to new contributions after 100 years.

In our simulations, the type $\xi \in \{0, 1, 2, \dots, M-1\}$ represents the generation number of the individual. At time 0, there are $\text{NRA} - x_0 + 1 \leq M$ generations in the scheme, with generation $\xi = 0$ of age $\text{NRA} - 1$, generation $\xi = 1$ of age $\text{NRA} - 2$, and so on, up to generation $\xi = \text{NRA} - x_0 + 1$ who are age x_0 at time 0. At time 1, generation $\xi = \text{NRA} - x_0 + 2$ joins the scheme at age x_0 , and the new entrants continue to flow steadily into the scheme at each integer time. Thus, we have the following equations:

$$\begin{aligned} \text{age}(\xi, t) &:= \text{NRA} - \xi + t - 1, \\ \mathbf{1}_t^{R, \xi} &:= \mathbf{1}_{\text{age}(\xi, t) \geq \text{NRA}}, \\ \mathbf{1}_t^{C, \xi} &:= \mathbf{1}_{\text{age}(\xi, t) \in [x_0, \text{NRA})}. \end{aligned} \quad (14)$$

for $\xi = 0, 1, 2, \dots, M-1$.

There is the same number of individuals in each generation when they join the scheme, so that the number of survivors satisfies $N_0^0 = N_0^1 = \dots = N_0^{\text{NRA}-x_0+1} = N_1^{\text{NRA}-x_0+2} = N_2^{\text{NRA}-x_0+3} = \dots$. Additionally, before reaching age NRA , each individual in generation ξ who makes a contribution at time t is paid the annual rate of salary $S_t^\xi := S_t$ at that time. Therefore, there are no promotional salary increases. Salaries increase annually at time t at the rate g per annum, such that $S_t = S_0(1+g)^t$.

All scheme members survive to age NRA . Thereafter, an individual who is alive at age x has the same probability of dying over the next year, regardless of their generation. In our simulations, the survival probabilities from age NRA are computed using the S1PMA tables produced by the Institute and Faculty of Actuaries' Continuous Mortality Investigation.

Let us write $p(x, n)$ for the probability of an individual surviving $n \geq 0$ years from age x , given that the individual is alive at age x . Then,

$$p_{t,n}^\xi =: p(\text{age}(\xi, t), n). \quad (15)$$

We assume throughout that there are sufficient members so that individual longevity risk can be perfectly hedged at all ages. As a consequence, we use the proportion of survivors, rather than the number of survivors, in our calculations.

3.2 Economic assumptions

We use a number of economic models with varying levels of sophistication in this paper, but all the models share the same values for the long-term median of various risk factors, as shown in Table 1. Where applicable, values are chosen to match the figures of the UK Office for Budget Responsibility (2024), for September. The stock volatility in all stochastic models is 15.3% per annum to match the model of Alvares Maffra *et al.* (2021).

3.3 Scheme design assumptions

3.3.1 Bounds on the nominal indexation rate

For both CDC schemes, the real indexation rate at time $t \in \mathbb{N}_0$ is constrained to lie in the range $h_t \in [-q_t, 0.05]$, so that $h_t^{\text{upper}} := 0.05$. Any further adjustments to benefits required to equate the discounted benefits with the asset value are dealt with by changing the bonus level θ_t to a value other than one.

Consequently, real indexation rates granted on benefits are capped at 5% for both types of CDC fund, so nominal rates are capped at approximately $q_t + 5\%$. The floor for nominal rates is set at 0%. Our industry consultation suggested that a benefit cap of $q_t + 2\%$ would be considered a more reasonable level for a standard dynamic-accrual scheme, but as we explain in more detail later, choosing the lower cap results in a decreasing median pension in retirement. The RMCPP does not have an upper bound on h_t , but we have included one to aid comparison between the fund types.

3.3.2 Benefit accrual rate

For the flat-accrual CDC scheme, the benefit accrual rate is set to $1/\beta = 1/80$. This matches the value for the RMCPP. For the dynamic-accrual scheme, the benefit accrued depends on the amount of contribution and is determined using equation (11).

3.3.3 Setting the contribution rate

To aid comparability between scheme types, the same contribution rate α is used for both the flat-accrual and dynamic-accrual CDC schemes.

Within the flat-accrual CDC scheme, α is computed as the rate required to deliver the desired benefits. To do this, it is assumed that the membership is a stable population and that, within the chosen economic model, the economic variables have attained their long-term medians, which are displayed in Table 1. We refer to these central estimates as *steady-state rates*.

We now assume the fund is targeting a long-term indexation rate h^∞ that matches the long-term CPI growth. Regulations require schemes to target at least CPI growth, and according to our industry consultation, schemes are expected to target this, as otherwise the cost of nominal benefits would be higher, potentially making the scheme appears less attractive. Using equation (10) for the resulting constant economic model, we may compute α . The resulting contribution rate is $\alpha = 4.84\%$. This is then used as the contribution rate for all the other fund types we simulate. Sensitivity testing shows that the contribution rate is most sensitive to the CPI and stock growth rate assumptions. For example, jointly increasing CPI by 1% per annum and lowering the stock growth rate by 1% per annum gives a contribution rate of just over 8%. Further lowering the stock growth rate by another 1% results in a contribution rate of around 11%.

Since the argument used to derive the contribution rate assumes a deterministic model and does not take into account benefit cuts or bonuses, one cannot expect that a future nominal indexation rate of CPI is in any precise sense the “average” rate of increase that will occur in simulations.

The contribution level in our simulations is much lower than in the RMCPP. The Trust Deed and Rules (Slaughter and May, 2024) specify a total contribution rate of 14.9% of pensionable pay (10.9% from the employer, 4% from the member) but do not state how these contribution rates were derived. Thus, no explicit target rate of indexation is set in the Trust Deed and Rules. Since we are not considering issues such as spousal benefits, it would not be appropriate to directly use the RMCPP contribution levels in our simulations. Nevertheless, since this contribution rate is much larger than the value of α we use, one would expect that the indexation achieved in the RMCPP will, on average, be above q . The previous Royal Mail DB scheme targeted the Retail Price Index, which can be approximated as $q + 1\%$.

3.4 Investment strategy: Flat-accrual scheme

The investment objective of RMCPP is to give scheme members exposure to the returns on return-seeking assets (which are defined in RMCPP's Trust Deed and Rules (Slaughter and May, 2024, Schedule 1, C.3)). Specifically, the objective is to invest in line with a discounted benefit-weighted investment strategy, after a cash investment that is sufficient to meet the scheme's liquidity requirements is met. The discounted benefit-weighted investment strategy is an interpolation between 100% low-risk assets, defined as bonds and other assets that have an investment-grade credit rating, and 100% return-seeking assets. Return-seeking assets are assets that have an expected return in line with that on the FTSE All World Index, 50% hedged to sterling and weighted according to market capitalization, with an expected volatility of absolute return in the range [65%, 100%] of the same index.

The consequent RMCPP investment strategy is similar to a typical default investment strategy of a DC pension scheme, but with a higher risk strategy that extends well into retirement. Members aged 67 years or younger are assumed to invest entirely in return-seeking assets with respect to their discounted benefit value. Those aged 90 years or older invest entirely in low-risk assets. Linear interpolation between these two extremes is done for members between the ages of 67 and 90 years. The aggregate scheme investment into return-seeking assets is the average of the percentage across the membership, weighted by each member's discounted benefit. The discounted benefit values at the last scheme valuation are used in the calculation.

We assume that our shared indexation CDC schemes have a similar objective. However, we ignore cash liquidity requirements and assume that the schemes invest in a combination of risky stocks and index-linked bonds.

The flat-accrual CDC scheme uses a life-styling strategy for each member of 100% in the risky asset up to age NRA years, tapering linearly to 0% at age NRA + 20. These percentages are based on the discounted benefits of each member at the last valuation date. Aggregating the percentages over the entire scheme membership gives the scheme's investment in the risky asset. The remaining assets are invested in index-linked bonds.

3.4.1 Investment strategy: dynamic-accrual scheme

The dynamic-accrual CDC fund applies a deterministic investment strategy. It is calculated as the median investment strategy of the flat-accrual CDC scheme. This is both for reasons of simplicity and to have broadly the same level of investment risk in both scheme types.

To choose the strategy for the standard dynamic-accrual fund, we simulate the flat-accrual fund using a constant economic model where all risk factors grow at the median values. From this, we can read off the overall proportion invested in risky assets each year, and then use this to determine the investment strategy of the standard dynamic-accrual fund. As shown in Figure 2, the result is that our standard dynamic-accrual fund follows a similar strategy to the median strategy adopted by the flat-accrual fund.

Alongside all our computations, we performed various tests to validate our models. See online Appendix C for details of some of the key sanity checks.

4. Cross-subsidies

4.1 Deterministic modeling for flat-accrual CDC

On an intuitive level, the decision to provide a fixed proportion of salary as a benefit for a given level of contribution can be expected to benefit investors who are approaching retirement. This is because younger investors have longer to invest their assets than older investors, so, to be financially fair, one would expect younger investors to receive a higher nominal benefit entitlement for a given contribution. Similar existing intergenerational cross-subsidies occur within DB schemes and are well understood. The Equality Act (Age Exceptions for Pension Schemes) (2010) contains




Figure 2. (Left) Fan diagram showing the deciles of the proportion invested in risky assets in the flat-accrual CDC fund each year. (Right) The deterministic proportion invested in risky assets each year in the standard dynamic-accrual scheme. The economic model used for this simulation is as described in Section 5.

specific provisions to ensure that DB schemes are lawful despite the age-discrimination inherent in such cross-subsidies.

At each time, the amount paid by each investor can be compared to the present value of the additional cash flows they will receive. We will call the difference the *instantaneous profit or loss*. Notice that this is a profit and loss relative to the full contribution consisting of both the employee and the employer's contributions. Since the employer's contribution is not specifically allocated to any individual, one could phrase this effect in terms of different age groups receiving different effective employer contributions, rather than in terms of a profit or loss. We use the terminology of profit and loss because this seems more natural when discussing dynamic-accrual schemes, and we wish to use a consistent approach to all schemes.

If we assume a deterministic economic model, we can quantify the instantaneous profit or loss analytically. For interest rates i and j and positive integer $n > 0$, define

$$a_{\overline{n}|}^{i,j} := \sum_{k=1}^n \left(\frac{1+j}{1+i} \right)^k \quad \text{and} \quad \ddot{a}_{\overline{n}|}^{i,j} := \sum_{k=0}^{n-1} \left(\frac{1+j}{1+i} \right)^k.$$

Theorem 1. Consider a flat-accrual CDC scheme that has a stable population. Before retirement, each individual contributes the amount αS_t at time t , where α is the contribution rate as a proportion of salary. Contributions are made in advance for an integer n years by each member, before retirement. The new annual benefits accrued at time t is $1/\beta \times S_t$ per member, where $1/\beta$ is the accrual rate.

Assume a constant economic and demographic model, in which $S_t = (1+g)^t$ and accrued benefits receive annual pension increases at the constant effective rate q p.a., both before and after retirement. Asset returns are a constant effective rate R per annum. It is assumed that the scheme is funded 100% at all times under this model. For the latter assumption to hold, the contribution rate α is calculated by equating the contributions paid at time t with the new benefits accrued at time t , so that the right-hand side of equation (6) is zero.

Under this contribution rate, the instantaneous percentage gain made by a member who is exactly k years from retirement is the additional value of the new benefits accrued by the member over the contribution made by that member, i.e.,

$$\left(\frac{n}{a_{\overline{n}|}^{R,q}} \left(\frac{1+q}{1+R} \right)^k - 1 \right) \times 100\%. \tag{16}$$




Figure 3. The ratio, BR, of the benefits from a CDC fund and a DC fund investing in the same asset, (a) as a function of the interest rate R used for discounting, and (b) as a function of the number of years n from joining the scheme until retirement. All non-varying parameters are as in Table 1. In (a), $n = 40$. In (b), R is equal to the stock growth rate given in Table 1. The dotted lines in (a) indicate that the ratio is equal to 1 when the risk-free rate is equal to wage inflation or price inflation.

A negative value indicates a loss. As $R \rightarrow q$, the instantaneous percentage gain tends to zero.

Suppose the member could alternatively have invested their contributions in an individual DC fund and then used the accumulated DC funds at retirement to purchase a single-life annuity that increases at the rate q p.a. Assume that there is no additional fee charged when purchasing the annuity to cover systematic longevity risk.

The benefit ratio (BR) at retirement is the ratio of the annual benefits from the CDC scheme to the equivalent from the individual DC fund at the first time of retirement. It is given by

$$BR = \left(\frac{1+g}{1+R} \right)^n \frac{n a_{\bar{n}|}^{g,q}}{a_{\bar{n}|}^{R,q} \cdot \ddot{a}_{\bar{n}|}^{R,g}}. \quad (17)$$

Finally, as $R \rightarrow g$ and as $R \rightarrow q$, $BR \rightarrow 1$.

See online Appendix A for a proof. Note that since this result only refers to a single asset, we can use a constant economic model to compute the value of benefits unambiguously without worrying about the possibility of arbitrage. As we explain in Section 4.3, one must be cautious when using a constant economic model to evaluate benefits in a CDC fund. Since this theorem is predicated on the assumption that indexation is constant, it can only be expected to give reasonable estimates for a CDC fund when indexation is maintained at its long-term average level. Thus, the result estimates the average behavior of a CDC fund, but the instantaneous profit may be somewhat different from these estimates in any particular year.

We see from equation (16) that the level of intergenerational cross-subsidy will typically increase with the annual interest rate R . If we assume R is given by the nominal growth of index-linked bonds, using the values in Table 1, we find from equation (16) that the instantaneous profit for the youngest generation would be -39.4% each year, and for the oldest generation it is 53.0% . We may read from this that, in a DB scheme, the benefit entitlement received each year by the oldest generation is worth approximately 2.5 times as much as the benefit entitlement received by the youngest generation. If we assume R is given by equity growth and repeat the calculation, we can estimate that, in a flat-accrual CDC scheme, the benefit entitlement of the oldest generation is worth almost 9 times as much as that of the youngest generation.

Equation (17) quantifies what we call *infinite-horizon effects* that can occur in both flat-accrual CDC schemes and DB schemes. It shows that when a flat-accrual CDC scheme is in its steady state, it may not be actuarially fair.

To see this, in Figure 3 (a), we plot the ratio of the CDC and DC annual pension payments when all rates match the values in Table 1, except for the rate of returns on investment, which

we allow to vary. The rate of indexation in this figure is equal to price inflation. This ratio is also equal to the ratio of the value of fixed-accrual CDC benefits to payments in a constant economic model. When the return rate R is less than wage inflation but higher than the rate of indexation in the steady state of the scheme, each generation receives a pension that is worth slightly more than the value of their accumulation contributions. This is possible because they are able to consume some of the savings of subsequent generations, who then take some of the savings of the next generation, and so on *ad infinitum*. Because wage inflation dominates indexation, this results in each generation being better off in the steady state. This is, of course, a form of Ponzi scheme. If subsequent generations decided they no longer wanted to participate in the scheme, resulting in its closure, the fund would leave its steady state, and the remaining members would receive an income worth less than they had paid in. We call this an infinite-horizon effect because it is only sustainable if one assumes the fund operates over an infinite time period.

As R increases above wage inflation, the level of intergenerational unfairness increases. As early generations have been overpaid, the fund begins to accumulate a debt and the payments that can be sustained in the long term are reduced by the interest that must be paid on this debt. If the fund invests primarily in risky assets, resulting in a large value for R , the value of the lifetime benefit received in a flat-accrual scheme will be less than the accumulated value of the contributions made over their working lifetime.

If one wants to ensure that the scheme is actuarially fair over each investor's lifetime and that the funding position is unchanged after new contributions are paid, one must have either that R is equal to the level of indexation or that R is equal to wage inflation. Neither of these options is desirable in a flat-accrual CDC scheme.

Theorem 1 shows that infinite-horizon effects would also exist in hypothetical fully-funded and fully-hedged DB schemes, but would be small if the risk-free rate is close to wage inflation. This is the case for the OBR assumptions shown in Table 1.

In Figure 3 (b), we plot the ratio of the CDC and DC payoffs as a function of n , the number of years from joining the scheme to retirement. This is equal to the ratio of the CDC and DC payoff in the steady state of the schemes. In this plot, we have taken all rates to match the OBR assumptions, with R equal to stock growth. The plot shows that the longer the time over which contributions are made, n , the worse the value of the flat-accrual CDC scheme. This explains why Upton (2024) found that flat-accrual CDC sometimes barely outperforms a DC scheme followed by annuitization. In that report, it was assumed that individuals join the scheme at age 20 and retire at age 67, giving a value of $n = 47$, substantially higher than the value $n = 40$ used in this paper.

Our results also explain the wide variation in the estimates of the potential benefits of CDC that have been found in different studies. Willis Towers Watson (2020) gives a particularly high estimate for the benefit of CDC schemes, but its methodology tacitly assumes actuarial fairness over each member's lifetime, and so it cannot be applied to flat-accrual schemes.

4.2 Stochastic modeling of cross-subsidies in flat-accrual schemes

In order to get a complete picture of intergenerational transfers of wealth, we must use a pricing methodology that takes into account the fact that h will vary over time. The liability given by formula (5) is an actuarial estimate of the value of the benefits that are expected to be received by the members of a CDC scheme. However, this formula is only an approximation as it does not consider the fact that the nominal indexation rate h varies with time. Instead, it assumes that h continues at its current level.

To get a complete picture, we simulate the different types of CDC schemes in a Black–Scholes model with constant rates of price inflation and wage inflation. We assume that the scheme membership is an infinitely large, stable population and that there is no systematic longevity risk, so that mortality can be treated deterministically. The continuously-compounded annual return on

the riskless asset is a constant r . The price process of the risky asset follows geometric Brownian motion with constant drift μ and constant volatility σ . When simulating the dynamics of CDC funds in this model, the “central-estimates” and “projections” required are computed using the mean values predicted by this Black–Scholes model.

In the Black–Scholes model, the pension that a generation receives can be viewed as a path-dependent derivative contract in the stock price. This is because once one knows the full history of the stock price, one can unambiguously compute the payments received in retirement. Here, we assume that we know what others will contribute to the fund. In effect, we are assuming that employer contributions are sufficiently high to ensure that there is never any incentive for members not to contribute and, similarly, that there is no incentive to leave the scheme. Subject to this assumption, the theory of derivative pricing in the Black–Scholes model then allows one to compute, at any time, an unambiguous market value for these payments. We can then directly compare the value of the benefits received to the contributions paid in respect of each scheme member.

The central point is simply that there is a single, undisputed pricing methodology in the Black–Scholes model, and it does not correspond exactly to the “central-estimate” approach to pricing benefits implied in equation (5), even when one uses the same discount rate. In effect, applying risk-neutral pricing allows us to calculate the monetary value of cross-subsidies experienced by each generation of members. It also allows us to estimate the accuracy of equation (5) as an estimate of future benefits.

The Black–Scholes model gives a toy model in which we can rigorously derive prices and demonstrate the pricing inaccuracies that can occur from the use of discounting formulae. In online Appendix E, we repeat some of our calculations using a stochastic interest rate model. This provides evidence that the average levels of cross-subsidies we calculate do not depend heavily on the precise risk-neutral model.

Fix a time t and consider the cashflows that arise from time t for all members of generation ξ , due to the contributions they pay at time t . Let $B_s^{t,\xi}$ denote the total annual amount of nominal benefit at time $s \geq t$ that was accrued by generation ξ at time t , by their contributions paid at time t , i.e.,

$$B_s^{t,\xi} := \begin{cases} B_t^\xi & \text{if } s = t; \\ \theta_s(1 + h_s^{\text{nom}})B_{s-1}^\xi & \text{if } s > t. \end{cases}$$

Let $X_s^{t,\xi}$ denote the total random cashflow paid or received by all members of generation ξ at time $s \geq t$, as a result of their contributions paid at time t . We have

$$X_s^{t,\xi} := \begin{cases} -C_t^\xi N_t^\xi & \text{if } s = t; \\ B_s^{t,\xi} N_s^\xi \mathbf{1}_s^{R,\xi} & \text{if } s > t. \end{cases}$$

Let $E_t^{\mathbb{Q}}$ denote the expectation taken in the risk-neutral measure $\mathbb{Q}$, conditional on the information available at time t . Write $\mathbb{P}$ for the physical measure. The instantaneous profit at time t for generation ξ , is given by the random variable

$$E_t^{\mathbb{Q}} \left(\sum_{s=t}^{\infty} e^{-rs} X_s^{t,\xi} \right),$$

with a negative value indicating a loss.

We define the *expected instantaneous profit at time t* to be the $\mathbb{P}$ -measure expectation of this, namely,

$$E^{\mathbb{P}} \left(E_t^{\mathbb{Q}} \left(\sum_{s=t}^{\infty} e^{-rs} X_s^{t,\xi} \right) \right). \quad (18)$$




Figure 4. Expected instantaneous profit for investors in a flat-accrual CDC fund, by age and year of operation, evaluated using 20,000 Monte Carlo scenarios.

To compute this value, we simulate economic scenarios using the stock drift under the physical measure, $\mathbb{P}$, up to the time of retirement. We then perform the remaining simulation using the stock drift for the pricing measure, $\mathbb{Q}$. This means that the expected instantaneous profit at a particular time t can be computed using a single Monte Carlo simulation, but a separate simulation is needed for each time.

Figure 4 shows the expected instantaneous profit made by each generation in each year when the fund is open to contributions. As one would expect, older investors benefit considerably from the flat-accrual CDC structure. This figure indicates that the average benefit entitlement of the oldest generation is worth approximately 10 times more than the average benefit entitlement of the youngest generation, which is in line with the estimate of 9 times using Theorem 1. We conclude that, despite using a constant economic model, Theorem 1 is able to account for most of the intergenerational cross-subsidy in the flat-accrual scheme.

In Figure 5, we show the aggregate effect for each generation,

$$\sum_{t=0}^{\infty} E^{\mathbb{P}} \left(E_t^{\mathbb{Q}} \left(\sum_{s=t}^{\infty} e^{-rs} X_s^{t,\xi} \right) \right).$$

Under the latter measure, there are two groups that have the most extreme outcomes. The first group consists of the older members who joined when the scheme started; they have a large total expected instantaneous profit. These members joined when they were close to the retirement age of 65 years. For example, generation $\xi = 0$ is age 64 years and generation $\xi = 1$ is age 63 years at time 0. Their contributions would not, in isolation, have funded all their retirement benefits.

However, as the generation number increases, the contribution history also increases and the total expected instantaneous profit declines. The nadir is reached by generation $\xi = 39$, who are the first generation to make contributions over the entirety of their working life.

We emphasize that it is a deliberate design feature to give older members at the start of the scheme's life a pension payment that is not entirely funded by their *de facto* individually attributable contributions. The employer's contribution is to the scheme as a whole, and is not owned by individual members. This mirrors what happens in UK DB pension schemes, which the flat-accrual CDC scheme is designed to replace.




Figure 5. Total lifetime expected instantaneous profit for investors in a flat-accrual CDC fund by generation, evaluated using the data points in Figure 4, with linear interpolation used for missing years.




Figure 6. Total lifetime discounted value of net cashflows at time 0 for each generation as a proportion of one year's salary, calculated with 100, 000 Monte Carlo simulations. A 95%-confidence interval is shown.

The second group with extreme outcomes comprises the members who joined the scheme shortly before it closed, when they were quite young. Their contributions were higher than the value of their accrued benefits, and thus, they have a considerable expected loss.

We compute the total risk-neutral value of the cashflows received by each generation, evaluated at time 0, that is

$$E^{\mathbb{Q}} \left(\sum_{t=0}^{\infty} \sum_{s=t}^{\infty} e^{-rs} X_s^{t,\xi} \right),$$

with the results shown in Figure 6. We see that when the scheme membership is fully mature (which we identify as the middle phase when all members have a full contribution history), the risk-neutral value at time 0 of each generation's pension cashflows is net negative. In contrast, Figure 5 shows that the expected instantaneous profit is net positive. One can interpret these results as the earlier generations speculating on stock growth on behalf of subsequent generations, and also using some of the subsequent generations' capital to increase their own pensions.

The total area under the graph in Figure 6 is equal to zero. However, as we have remarked, when the scheme is fully mature, members receive an amount with a negative risk-neutral value

at time 0. In other words, their contributions are higher than the value of the benefits they receive. This is a result of the infinite-horizon effects discussed above.

We conclude that, for flat-accrual schemes, the use of a stochastic pricing methodology reproduces the results found using a constant economic model. One practical consequence of our observations is that employer contributions must be high in a flat-accrual CDC scheme in order to guarantee that it is always in an employee's interests to contribute to the scheme. In addition, a younger, potential member of a flat-accrual CDC scheme should ask what, if any, recompense they will receive for subsidizing older generations if they change employers. Similarly, a younger member should ask what will be done if a scheme is closed, to ensure that they receive the intergenerational cross-subsidy they expected. These are equally valid questions for members of DB schemes to ask, but they are far more pressing for potential members of a flat-accrual CDC scheme.

4.2.1 Age-related accrual

To reduce the cross-subsidies of a flat-accrual scheme, one might consider an age-related accrual scheme where the amount of annual benefit accrued decreases with age. For an individual of age x with annual salary rate S_t , suppose at time t their additional annual benefit is of amount

$$B_t = \frac{1}{\beta} (1 + d)^{\text{NRA} - 1 - x} S_t,$$

for some constant discounting level d . In Figure F.1 in the online appendix, we show the expected instantaneous profit cross-subsidies that occur across different generations when $d = 3.5\%$. This level was chosen as it is approximately equal to the difference between stock growth and bond growth rates, so it can be expected to reduce the cross-subsidies to a similar level to that seen in a DB scheme invested in gilts (as can be seen by applying Theorem 1). We will also see later that this scheme reduces the drag effect seen in the flat-accrual scheme.

4.3 Dynamic-accrual schemes

We now compute the instantaneous profit for standard dynamic-accrual schemes using a stochastic economic model. If our formula for liabilities priced the value of benefits in a CDC fund perfectly accurately, we would expect the instantaneous profit to be zero. Thus, our results evaluate the discrepancy between the mathematically rigorous methodology of risk-neutral pricing and the heuristic central-estimate approach implicit in equation (5).

As described in Section 3.4.1, we chose the investment strategy for the standard dynamic-accrual CDC fund to match the observed time-dependent strategy of a flat-accrual scheme in a deterministic version of the economic model. This ensures that we are using the same discount rate for all members in equation (5).

Figure 7 shows the expected instantaneous profit for the standard dynamic-accrual fund.

The intergenerational disparities are, as one would expect, much smaller than those seen in a flat-accrual fund (note the difference in the vertical scale between Figures 4 and 7), but are still nontrivial at fund startup.

The instantaneous profit is a random variable, so we should consider more than simply its expected value. In Figure 8, we show 50 different scenarios for the instantaneous profit of members by their age at time $t = 50$. This plot is generated using nested Monte Carlo simulations. We use the physical measure $\mathbb{P}$ to generate 50 different stock price paths up to time $t = 50$, and then perform a Monte Carlo pricing calculation using the pricing measure $\mathbb{Q}$ to calculate the profit made at time $t = 50$.

To understand why equation (11) is an inaccurate pricing formula, we note that it assumes h remains constant. If the stock volatility is sufficiently small that the upper and lower barriers for




Figure 7. Expected instantaneous profit and loss for members in a standard dynamic-accrual CDC fund, by age and time, evaluated using 20,000 Monte Carlo scenarios.




Figure 8. Left plot: 50 random scenarios simulated in the physical measure showing the instantaneous profit made by each member at time $t = 50$, according to their age. Each curve represents a different scenario. Interpolation is used between integer ages to obtain the curves. The profit in each scenario is computed using a nested Monte Carlo simulation in the pricing measure with 50,000 samples. Right plot: the residual of the model defined by (19), expressed as a percentage of V_t^ξ , the Monte Carlo estimate of (11).

h are never hit, this approximation will be true on average. However, in practice, the barriers are very likely to be hit, and so the long-term average value of h is not determined by the current value of h , but by the choice of barriers. Since (11) depends upon values of h projected a considerable time into the future, we cannot expect the discounting formula to be accurate. As a result, one expects that when h_t is high, equation (5) will overprice benefits. When h_t is low, it will underprice benefits.

To test this explanation, and to also obtain a more accurate pricing model, first recall the value of a unit of nominal benefit at time t , $\hat{V}_t^\xi$, defined in (11). Let us write V_t^ξ for the value computed using our Monte Carlo calculation. We then perform a linear regression to find a model of the form

$$\log(V_t^\xi) \sim c_0 + c_x \text{age}(\xi, t) + c_h h_t + c_{x,h} \text{age}(\xi, t) h_t, \quad (19)$$

Table 2. Estimated coefficients of the linear model (19) computed for both the empirical value of 1 unit of nominal benefit, $\log(V_t^\xi)$, and the discounting-formula estimate, $\log(\hat{V}_t^\xi)$

	$\log(V_t^\xi)$	$\log(\hat{V}_t^\xi)$
c_0	−1.952	−1.412
c_x	0.06808	0.05855
c_h	26.06	72.76
$c_{x,h}$	−0.1597	−1.002

using the same 50 scenarios used to obtain Figure 8. This gives a total of $2000 = 50 \times 40$ data points for the regression. All coefficients are significant with $p < 0.1\%$, and their values are summarized in Table 2. For comparison, the coefficients for a similar model for $\log(\hat{V}_t^\xi)$ (defined in equation (11)) are also shown. Comparing the coefficient c_h confirms our prediction that when h_t is large, benefits will be overpriced by $\hat{V}_t^\xi$.

The R^2 -statistic for the model for $\log(V_t^\xi)$ is 0.9984. A plot of the residuals as a percentage of the values of V_t^ξ is shown in Figure 8. We conclude that estimating V_t^ξ using equation (19) will be considerably more accurate than using the estimate $\hat{V}_t^\xi$. The R^2 -statistic for the model for $\log(\hat{V}_t^\xi)$ is 1.0000.

Treating age as a categorical variable did not significantly reduce the maximum value of the residuals, so we conclude that the benefits of using a more complex model will be limited. A precise calculation of the value of the benefits would be path dependent, and so, there is a limit on what can be achieved using only the age and h_t as input.

4.3.1 Statistically-calibrated accrual

It is natural to consider pricing benefits using the statistical model given by the coefficients in Table 2, instead of using the liability value in (5). One cannot simply charge the amount given by V_t^ξ in equation (19), as then equation (13) will be violated. Instead, one can charge individuals an amount $c V_t^\xi$ where the constant c is chosen to ensure that equation (13) holds. We refer to this method of pricing benefits as *statistically-calibrated* dynamic-accrual.

As one might expect, simulations show this succeeds in reducing the average intergenerational cross-subsidies. This is illustrated in Figure F.2 in the online appendix. It also reduces the random pricing errors to a size similar to the residuals of our statistical model.

One potential downfall with this approach is that in any incomplete market model, there may not be universal agreement on the price of benefits. However, it seems reasonable to expect that the prices obtained by using a Black–Scholes approximation to an incomplete market model will yield fairer outcomes than using the liability formula (5). We see in the sequel that using the model with coefficients given by Table 2 yields slightly better outcomes in a steady state when we perform the simulation in a more realistic market model, see Figure 12. One could, if desired, attempt to obtain a yet more accurate pricing model by calculating the value of benefits in this new scheme and then recalibrating the linear model.

5. Simulations in a richer economic model

5.1 The economic model

So far, we have used a constant economic model and a Black–Scholes model for parsimony. In particular, when computing cross-subsidies, the Black–Scholes model allowed us to value benefits

unambiguously. In this section, we will evaluate the overall performance of CDC schemes without considering the value of benefits, and this allows us to use a richer economic model.

For the simulations detailed in this section, we simulate all risk factors using a minor variation of the economic scenario generator (ESG) described in Alvares Maffra *et al.* (2021) (see online Appendix B for further details). This scenario generator allows one to input views about long-term median values for different risk factors, and these are chosen as shown in Table 1. Although this scenario generator allows one to simulate population mortality, we do not use this feature and continue to use the S1PMA tables produced by the continuous mortality investigation.

To compute projected values of CPI to apply in equation (2), we use median values as these are simple to project in our ESG. Our ESG is designed to output index-linked long-term bond yields, and we use these combined with the long-term median estimate for CPI to determine the discount rate for riskless assets in equation (2). We also assumed that the bond portfolio is index-linked and chosen to match the liabilities exactly. The returns on the bond portfolio can then be computed by pricing this portfolio using the index-linked long-term bond yields supplied by the ESG. The net result of these assumptions is that if the fund is invested 100% in bonds, then h will remain constant unless the lower bound on indexation is hit due to changes in CPI.

In our ESG, the stock follows geometric Brownian motion, so we can compute the predicted mean returns and use this to compute the appropriate discount rate for risky assets in equation (2).

All plots in this section were performed using 20,000 Monte Carlo simulations.

There is always a significant subjective element in developing such a rich economic model, but we have confirmed that the qualitative results we identify also hold in a Black–Scholes model and in the SIR model. We have shown the resulting plots for the SIR model in online Appendix E.

5.2 Annual increments

In Figure 9, we plot a fan diagram of the real indexation rate h in different years of our simulation. In all fan diagrams, in this paper, we plot all nine deciles together with an example scenario to indicate the volatility.

The median value of h for the flat-accrual scheme is approximately zero until the closure of the scheme. At the end of the scheme's lifetime, it is invested entirely in risk-free assets, so h will remain constant unless a benefit cut occurs due to changes in inflation pushing the total $h + q$ below zero. A bonus is not possible because these are determined by the value of h , rather than $h + q$. This is why, at the closure of the scheme, the lower percentiles in the distribution of h begin to increase. The constancy of h , excluding benefit cuts, is a consequence of our decision to model a fund with an infinitely large membership. For finite populations, h becomes increasingly volatile at scheme closure as liabilities cannot be perfectly hedged and due to the short duration of these liabilities; this requires comparatively large adjustments in h .

The dynamic-accrual scheme behaves similarly. As the 95th percentiles are very close to the cap and floor levels, it seems that bonuses and benefit cuts each occur on the order of once every twenty years, excluding the period after scheme closure.

The indexation level h does not include information about cuts or bonuses. In Figure 10, we plot the actual increase or decrease to the nominal benefit amount each year, including cuts and bonuses. One sees that the impact of cuts and bonuses is significant.

5.3 Pensions for a typical generation

We define the *replacement ratio* to be the ratio between the annual benefits at a given time during retirement and the annual salary rate exactly one year before retirement, indexing the salary with CPI to the given time. Figure 11 shows fan diagrams of the logarithm of the replacement ratio for the 60th generation in a number of different schemes. The 60th generation is chosen as they




Figure 9. Fan diagrams of the real indexation rate, h_t , in each year of the simulation.




Figure 10. Fan diagrams of the benefit increase/decrease over inflation, including bonuses/cuts in each year of the simulation.

have retired by the time the fund closes, but the fund has had sufficient time to settle into an approximate “steady state.” In fact, as can be seen from Figure 12, the CDC funds never truly enter a steady state, but nevertheless, broadly similar outcomes are experienced by generations 60 to 80.

We see that in a flat-accrual CDC scheme targeting $q + 0\%$, the median replacement ratio is close to constant in retirement, but the width of the fan diagram continues to increase throughout retirement. Since the median income is close to constant in real terms, the CDC scheme can be viewed as, on average, achieving the target.

In the dynamic-accrual CDC scheme, median real-terms pensions remain constant over time and are somewhat higher than for the flat-accrual scheme. The dynamic-accrual scheme appears to straightforwardly outperform the flat-accrual scheme in the sense that for all the plotted percentiles, the dynamic-accrual scheme always gives a better outcome.

The age-based accrual scheme introduced in Section 4.2.1 performs somewhere between flat- and dynamic-accrual schemes. Although it is not clear from this picture, we will see later that the statistically-calibrated scheme slightly outperforms the standard dynamic-accrual scheme.

For comparison, we show the income in retirement of a DC-plus-annuity scheme. In this scheme, individuals in a DC scheme invest 100% of their assets in the risky asset until age 55, and nothing in the riskless asset. They then taper linearly to 0% in the risky asset at age 65, when they purchase an index-linked single-life annuity. The price of the annuity is calculated by using




Figure 11. Fan diagrams of the log replacement ratio by age for the 60th generation.

the discount rate for long-term bonds together with the SP1MA mortality distribution used in the simulation. An additional charge of 5% is then applied to represent the charges that would need to be levied by the annuity provider to account for systematic longevity risk.

As another comparison, we show the income in retirement of a pooled annuity fund. In this scheme, the fund invests at 100% in the risky asset until age 55, tapering down to 33% at age 65. Each year in retirement, the assets of anyone who dies are shared with the survivors. To compute the consumption rate, suppose that the expected return on assets in a given year is R . Compute




Figure 12. Left plot: Median lifetime-mean replacement ratio. Right plot: Mean lifetime-mean replacement ratio. Both plots are generated using our economic scenario generator and a target of $q + 0\%$ for the CDC funds.

the cost of a level single-life annuity which pays one unit at the start of each year using the mortality distribution used in the simulation and the discount rate R . The pension paid out from the pooled annuity fund is equal to the value of the assets in the fund divided by the life annuity cost. As Figure 11 shows, this scheme results in an approximately constant real-terms income in retirement, but is able to offer a slightly better income than a DC-plus-annuity scheme. The pooled annuity fund benefits from lower costs (as there is no need to charge for systematic longevity risk) and higher returns (as there is some investment in the risky asset), but this does result in some annual fluctuations in pension income. The pooled annuity fund and the DC fund with annuitization both show lower median outcomes than the flat- and dynamic-accrual schemes, but do have less risk later in retirement.

Let ζ_t^ξ denote the replacement ratio of the pension received by individuals of type ξ at time t . We define the lifetime-mean replacement ratio for individuals of type ξ as

$$\left(\frac{\sum_t \mathbf{1}_t^{C,\xi}}{\text{NRA} - x} \right) \frac{\sum_t N_t^\xi \zeta_t^\xi \mathbf{1}_t^{R,\xi}}{\sum_t N_t^\xi \mathbf{1}_t^{R,\xi}}.$$

The first factor scales the value according to the number of years for which a generation has contributed in order to facilitate comparison between generations.

A plot of the median lifetime-mean replacement ratio for each generation is given in Figure 12, together with a plot of the mean lifetime-mean replacement ratio. Assuming that investors are risk-averse, the median lifetime-mean replacement ratio seems a better measure of the quality of a pension.

Using the median lifetime-mean as a metric in the flat-accrual scheme, the earliest generations receive a proportionately better pension than those when the scheme is fully mature, and the final generations receive a worse pension. The dynamic-accrual scheme works in a manner that is closer to being actuarially fair: the earliest generations receive a proportionately lower pension as their investments have not had time to grow; the latest generations receive a proportionately larger pension as all their payments were made early in their employment before the scheme was closed in year 100. One can see the behavior of the dynamic-accrual scheme is similar to that of the pooled annuity fund and the DC + annuity scheme, both of which are actuarially fair.

When the scheme is fully mature, the dynamic-accrual CDC scheme gives a higher median lifetime-mean replacement ratio than the flat-accrual scheme. This can be attributed to the drag effect discussed in Section 4.1.

The age-based accrual scheme is intended to approximately replicate the level of cross-subsidies seen in a DB scheme. The first generation receives a lower median lifetime-mean replacement ratio than those in the steady state of the scheme, but this will be compensated for by the fact that their pension income is less risky. As a result of reducing the cross-subsidies, the drag effect is reduced. This confirms the intuition that it is overpaying the earliest generations that causes the comparative underperformance of flat-accrual CDC.

The statistically-calibrated scheme contains slightly less intergenerational cross-subsidy than the standard dynamic-accrual scheme, and this results in slightly better performance in the steady state of the scheme.

As a simple numerical comparison, we give the median lifetime-mean replacement ratio for the 60th generation. This is 32.6% for the flat-accrual CDC scheme, 42.6% for the dynamic-accrual CDC scheme, 26.3% for the DC-plus-annuity scheme, and 33.0% for the pooled annuity fund. By this metric, the flat-accrual scheme outperformed DC and annuity by a factor of 23%.

6. Smoothing

CDC funds transfer the fluctuations in predicted pension incomes from older generations to younger generations. The idea is that the fund can then invest a greater proportion in risky assets without this resulting in large fluctuations in pension income for those who have already retired.

Let us call the partial derivative of the log of the current market value of an individual's benefit with respect to h , the h -duration of the benefit.

$$h\text{-duration} := \frac{\partial \log(V_t^{\xi})}{\partial h}$$

This definition mirrors that of the modified duration of a bond. We can calculate the h -duration of a member's benefit using the coefficients in Table 2, and we see that h -duration increases with time to retirement. Intuitively, this occurs for the same reason that a bond's duration increases with time to maturity.

In a CDC fund in years with no benefit cuts, fluctuations in asset values are accounted for by changing h . As a result, the higher the h -duration of a member's benefits, the more sensitive it will be to changes in the asset price. This means that individuals who are close to retirement will experience fewer fluctuations in their pension income than individuals who are a long way from retirement. This is a key design goal of shared-indexation CDC, as it results in smoother pensions in retirement by shifting the risk of asset volatility from older to younger members.

Notice that the coefficient $c_{x,h}$ in our linear models, is much lower in the model for $\log(V_t^{\xi})$ than it is in the model for $\log(\hat{V}_t^{\xi})$. This means that if one attempts to estimate how much smoothing will occur by looking at the sensitivity of liabilities to h , one will significantly over-estimate the amount of smoothing that occurs.

Figure 13 (a) illustrates an example scenario that shows how the projected mean lifetime-mean replacement ratio varies for one cohort over its lifetime. The roughness of the two shared-indexation designs is very similar to that of a DC + annuity scheme. The latter is somewhat smoother in the final years before retirement due to the life-styling strategy. Figure 13 (b) considers the same scenarios in the dynamic-accrual scheme and shows how the projected accrued benefit entitlement at retirement fluctuates with the stock price. We see that for younger cohorts, their projected benefit entitlement fluctuates slightly more than excess stock returns. For the oldest cohort, the projected benefit entitlement fluctuates approximately half as much as excess stock returns. We conclude that shared-indexation schemes do provide some smoothing of pension incomes before retirement, but it is a relatively modest effect compared to some estimates that have been made in the literature. For example, in Donnelly (2022), smoothing is explicitly




Figure 13. (a) Projected mean lifetime-mean replacement ratio for one typical cohort against age. (b) Annual fluctuation in projected mean benefit entitlement at retirement for the same cohort due to investment returns, plotted against the corresponding excess stock returns. Age is indicated by the shading of the points, with darker points used when the cohort was young. The lines through the origin with gradient 1 and $\frac{1}{2}$ are shown for reference.

estimated using such sensitivities with prices estimated using (5). As a result, this estimates much higher smoothing than is found through simulation.

AON (2013) contains a chart (Chart 5, p. 31) which suggests that in a CDC design similar to our flat-accrual scheme, projected benefits are almost entirely smooth in the run-up to retirement. This can be explained by their study computing projected benefits using a deterministic model rather than a simulation. Our results suggest that this deterministic calculation is likely to be inaccurate.

Smoothing of benefits may also be over-estimated due to a confusion between nominal benefit entitlements and projected benefit entitlements. Nominal benefit entitlements vary much more smoothly than projected benefit entitlements. This is because the change in projected benefits is a function of both the current nominal benefit and the level of indexation h_t . If one only looks at changes in nominal benefits each year, one will miss the changes in h_t . As an example of the potential for confusion, the Government Actuary's Department (2009) uses the phrase “projected benefits” to describe what we call nominal benefits.

Examining the example trajectories plotted in the fan diagrams in Figure 11, we see that similar levels of smoothing of pension outcomes are achieved across the CDC designs. While the CDC trajectories are smoother than those for the pooled annuity fund, this is a short-term smoothing rather than a long-term decrease in risk. As indicated by the funnel shapes of the CDC diagrams, risk increases significantly throughout retirement in our CDC funds. Popat *et al.* (2015) and Owadally *et al.* (2022) both seek to quantify the risk of CDC investments, but only measure the standard deviation of incomes at the time of retirement. As a result, these studies will have tended to understate the risks of CDC schemes.

7. Changing the target level

So far, we have considered the flat-accrual scheme with a target indexation level of $q + 0\%$. This aids comparisons with other schemes that target a constant income in retirement. However, to avoid frequent benefit cuts, one can target a higher indexation level, though this requires increasing contributions.

We plot the results of simulations targeting different levels of real indexation rate in a flat-accrual scheme (left plot of Figure 14). As we increase the target level, we also increase the level at which bonuses are awarded by the same amount. When targeting $q + 0\%$, the median retirement income decreases throughout retirement. If the median income in retirement were to increase at a




Figure 14. Plots of median income (log scale) against age for generation 60. Left plot: for a flat-accrual scheme with the target level of indexation being varied. Right plot: for a dynamic-accrual scheme, with the level at which bonuses are made being varied.

rate of 1% per year, it should increase by a total of about 73% from age 65 to 120. Thus, we see that the target levels do not correspond to the achieved rate of increase of median incomes in retirement. We note, however, that the method we are using to target a particular level of indexation is just a heuristic.

These results show that it would be better to use a line search over simulations to identify the required contribution rate. Since one can achieve different levels of indexation by varying the ratio of benefits to contributions, one should not be overly concerned that the heuristic is imperfect.

As the dynamic-accrual scheme does not incorporate a target level, only an initial level h_0 , we cannot adjust the pension increments in a dynamic-accrual scheme in the same way. If one wishes to choose a long-term target for indexation, one can attempt to do this by adjusting the boundaries for cuts and bonuses or by adjusting the proportion invested in risky assets. The former option may be unattractive as a fund currently targeting $q + 5\%$ will appear very expensive to investors. The latter may be unattractive as changing this will simultaneously change the risk profile of the fund. We have taken the approach of varying the boundary for bonuses in our simulations. The right plot of Figure 14 shows the median income in retirement for a dynamic-accrual scheme when indexation is capped at $q + 2\%$ and when it is capped at $q + 5\%$, plotted against age. As this plot shows, even if we have an upper bound of $q + 5\%$ on indexation, the median income in retirement in the dynamic-accrual scheme does not quite manage to keep up with inflation.

8. Scheme behavior when investment returns are not as expected

In Figure 15, we show the effect of running our simulations when the scenario generator is adjusted to generate data with stock returns on average 1% greater or less than expected, but where the fund is managed as before. In this case, one sees that if stock returns are greater than expected, later generations in the flat-accrual scheme benefit more than earlier generations. Similarly, if stock returns are lower than expected, later generations experience a greater penalty. The size of these effects is enough to significantly change which generation benefits the most from the scheme.

In the dynamic-accrual scheme, the fund responds in a way that appears broadly consistent with actuarial fairness: the longer an individual invests their funds, the more they benefit from unexpectedly high stock prices. Generations in the middle of the scheme all receive very similar pensions, even when the investment returns are not as expected.




Figure 15. Median lifetime-mean replacement ratio by generation when investment returns are not as expected. The chart on the left shows the case when stock returns are 1% lower than expected, and the chart on the right shows the case when stock returns are 1% greater than expected.

9. Conclusion

The flat-accrual CDC fund seeks to mirror the high levels of intergenerational cross-subsidy seen in DB schemes, but we find this subsidy is approximately four times greater in the flat-accrual CDC scheme. The flat-accrual scheme smooths income in retirement. It can lead to a higher average income in retirement than DC + annuity, but this will not necessarily occur due to infinite-horizon effects.

In order to reproduce similar levels of cross-subsidy to the DB schemes they replace, a better alternative is to offer varying benefit entitlements by age. This reduces the infinite-horizon issues experienced in these funds and results in higher average retirement incomes.

The standard dynamic-accrual fund does not target a specific benefit level in the same way as a flat-accrual scheme, and this makes it harder to control factors such as the frequency of benefit cuts or the median annual rate of pension increase. The dynamic-accrual CDC design results in much lower intergenerational cross-subsidies than a flat-accrual scheme, though they are still present. Much of the level of cross-subsidies in dynamic-accrual schemes can be predicted using a linear model in h and age, with a path-dependent residual.

Using this linear model to match benefits and contributions reduces the pricing errors. Because the dynamic-accrual scheme has much lower intergenerational cross-subsidies than a flat-accrual scheme, the infinite-horizon effects are small. This explains why the average retirement income is higher in dynamic-accrual schemes.

For both types of scheme, it is important to distinguish between nominal benefits and projected benefits. Nominal benefits should not be presented to members as though they were projected benefits. There is considerable scope for regulators to consider how accrued benefits should be presented to members to balance encouraging saving with communicating potential uncertainty. It is important that members are not misled, and there would be value in consistent presentation across scheme designs, both CDC and non-CDC.

This paper only considers shared-indexation CDC schemes. The pooled annuity fund gives an example of a collective scheme which provides a much stronger fairness guarantees than a shared-indexation fund: for all generations, the value of the cashflows received is equal to the value of the cashflows invested. In future work, we will investigate the extent to which the benefits of CDC funds can be reproduced while maintaining similarly strong guarantees of fairness.

Supplementary material. The supplementary material for this article can be found at <https://doi.org/10.1017/S1748499526100256>.

Acknowledgements. We would like to thank the PPI and the members of our project advisory board for their feedback on earlier drafts of this work. High performance computing was completed using CREATE (2024).

Data availability statement. Not applicable.

Funding statement. JA and JD were supported by a Nuffield grant FR-000024058. CD was funded by the IFoA's Actuarial Research Centre on the "Optimizing Future Pension Plans: Phase II" research programme.

Competing interests. None.

Ethical standards. The research meets all ethical guidelines, including adherence to the legal requirements of the study country.

Author contributions. All authors contributed to both the drafting and analysis and approved the final submitted draft.

References

- Alvares Maffra, S., Armstrong, J., & Pennanen, T. (2021). Stochastic modeling of assets and liabilities with mortality risk. *Scandinavian Actuarial Journal*, 2021(8), 695–725.
- AON. (2013). The Case for Collective DC. <https://tinyurl.com/2ax4s83k>.
- Blake, D. (2016). We Need a National Narrative: Building a Consensus around Retirement Income. Independent Review of Retirement Income. <http://www.pensions-institute.org/IRRIRReport.pdf>.
- Bonenkamp, J. & Westerhout, E. (2014). Intergenerational risk sharing and endogenous labour supply within funded pension schemes. *Economica*, 81(323), 566–592.
- CREATE. (2024). King's Computational Research, Engineering and Technology Environment (CREATE). <https://doi.org/10.18742/rnvf-m076>.
- Cui, J., De Jong, F. & Ponds, E. (2011). Intergenerational risk sharing within funded pension schemes. *Journal of Pension Economics & Finance*, 10(1), 1–29.
- Donnelly, C. (2022). Inter-generational crosssubsidies in the UK's first CDC pension scheme. Preprint available at https://pure.hw.ac.uk/ws/portalfiles/portal/100128075/Deliverable_3.1.pdf.
- Gollier, C. (2008). Intergenerational risk-sharing and risk-taking of a pension fund. *Journal of Public Economics*, 92(5–6), 1463–1485.
- Haan, J., Lekniute, Z. & Ponds, E. H. M. (2015). Pension Contracts and Risk Sharing—a Level Playing Field Comparison. Available at SSRN 2741542. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2741542.
- McInally, K. (2023). Multi-employer CDC schemes: can they be fair? Abdrn. <https://tinyurl.com/2ephu3y3>.
- Office for Budget Responsibility. (2024). Fiscal risks and sustainability report – September 2024 – charts and tables: supplementary tables. UK Office for Budget Responsibility. <https://obr.uk/frs/fiscal-risks-and-sustainability-september-2024/>.
- Owadally, I., Ram, R. & Regis, L. (2022). An analysis of the Dutch-style pension plans proposed by UK policy-makers. *Journal of Social Policy*, 51(2), 325–345.
- Popat, S. Curry, C., Pike, T. & Ellis C. (2015). Modelling Collective Defined Contribution Schemes. Pensions Policy Institute. <https://www.pensionspolicyinstitute.org.uk/media/200n2xuj/20151105-modelling-cdc-schemes.pdf>.
- Royal Mail Collective Pension Plan. (2024). The Collective Plan's Handbook. <https://tinyurl.com/mv43dukt>.
- Slaughter and May. (2024). Royal Mail Collective Pension Plan, Definitive Trust Deed and Rules. <https://tinyurl.com/23crh47b>.
- Taylor, S. & Ward, R. (2023). Collective DC modelling. Association of British Insurers. <https://tinyurl.com/2suvurkp>.
- The Equality Act (Age Exceptions for Pension Schemes) (2010). UK Statutory Instrument. Order 2010 (SI 2010/2133). <https://www.legislation.gov.uk/ukxi/2010/2133/schedule/1/made>.
- The Government Actuary's Department. (2009). Modelling Collective Defined Contribution Schemes. <https://tinyurl.com/4f9wdhj5>.
- Upton, J. (2024). Quantifying multi-employer and single employer CDC outcomes. Pensions Policy Institute, Briefing Note 139. <http://pensionspolicyinstitute.org.uk/research-library/research-reports/2024/briefing-note-139-quantifying-multi-employer-and-single-employer-cdc-outcomes/>
- Wilkinson, L. (2024). How might CDC develop in the UK? Pensions Policy Institute, Briefing Note 137. <https://www.pensionspolicyinstitute.org.uk/media/ceyauk4c/20240515-bn-137-how-might-cdc-develop-in-the-uk-final.pdf>
- Willis Towers Watson. (2020). How CDC pension levels compare with other types of schemes. <https://tinyurl.com/3nn3y5w7>.

Cite this article: Armstrong J, Dalby J and Donnelly C (2026). Intergenerational cross-subsidies in UK collective defined contribution funds, *Annals of Actuarial Science*, 20, 381–413. <https://doi.org/10.1017/S1748499526100256>