

Intergenerational cross-subsidies in UK collective defined contribution (CDC) funds

John Armstrong, James Dalby and Catherine Donnelly

November 22, 2024

Abstract

We evaluate the performance of and level of intergenerational cross subsidy in single-employer and multi-employer collective defined contribution (CDC) schemes which have been designed to be compatible with UK legislation. The single-employer scheme captures the essential features of the Royal Mail CDC scheme, which is currently the only UK CDC scheme. We find that the schemes can be successful in smoothing pension outcomes while outperforming a DC + annuity scheme, but that this outperformance is not guaranteed in a single-employer scheme. There are significant intergenerational cross-subsidies in the single-employer scheme. These qualitatively mirror the cross-subsidies seen in existing defined benefit schemes, but we find the magnitude of the cross-subsidies is much larger in single-employer CDC schemes. The multi-employer scheme is intended to minimize such cross-subsidies, but we find that such subsidies still arise due to the approximate pricing methodology implicit in the scheme design. These cross-subsidies tend to cancel out over time, but in any given year they can be large, implying that it is important to use a rigorous pricing methodology when valuing collective pension investments.

1 Introduction

Collective Defined Contribution (CDC) pension schemes are arriving in the UK, after a long, uncertain beginning [19]. These are pension schemes in which the contribution rate is fixed, as a percentage of salaries, and the benefits paid out vary. The idea is that the scheme members share both investment and longevity risk in order to have a better outcome in retirement. The adjustment lever in a CDC pension scheme, that benefits are adjusted to keep the contributions unchanged, is different to that in defined benefit (DB) pension schemes, in which contributions are adjusted to keep the benefits unchanged. The CDC scheme uses the same lever as a defined contribution (DC) pension scheme.

In this paper we study two different CDC scheme designs. The first design we consider is a single-employer scheme closely modelled on that of the Royal Mail CDC plan. This is the first CDC scheme to launch in the UK and has been the

driving force behind the introduction of the legislation enabling CDC schemes in the UK [19]. Combined with the scale of the scheme, which is anticipated to have well over 100,000 members, the importance and significance of the Royal Mail CDC is clear. It is the prototype presented by the actuaries involved in its design and testing to other actuaries. This scheme is designed as a replacement for a previous DB scheme and so is intended to produce a similar profile of contributions and benefits to that scheme.

The second design we will consider is a multi-employer CDC scheme. No such scheme exists in the UK at present, so we develop a scheme design which we believe reflects the common understanding of UK pension professionals on how multi-employer CDC might operate. To achieve this, we have worked with the Pension Policy Institute who have conducted a series of interviews on CDC pensions with a broad range of stakeholders and who have used these insights to publish two briefing notes on CDC. In addition, we have used the expertise of a project advisory group drawn from across the pension industry and academia. The consensus is that multi-employer scheme design is intended to minimize cross-subsidies between generations to prevent one employer potentially subsidizing others. It is also envisioned as a DC replacement rather than as a DB replacement, which too implies there should be low levels of intergenerational cross subsidies.

Owadally et al. [14] discuss comprehensively CDC plans in the UK with respect to policy, as well as their risks and the ensuing advantages and disadvantages. They model an earlier proposal for the Royal Mail CDC plan, but the scheme design has since changed. They also compare the outcomes for a single member of their model mature CDC plan against typical DC-based alternative pension schemes. They find their CDC plan can pay higher and more stable pensions compared to the alternative schemes considered.

This work is an attempt to understand better the broad implications of the Royal Mail CDC plan, and to compare them to multi-employer CDC. We capture the broad design of the Royal Mail scheme, but not all of it. In particular, death and spousal benefits are omitted and we assume that benefit cuts and bonuses that must occur when the fund is significantly over- or under-funded take affect immediately. Our mortality modelling is also simplified: we assume no systematic longevity risk and that the fund is large enough for individual longevity risk to be ignored. We do not consider any heterogeneity in the multi-employer fund in our simulations in this paper, but we do show how that can be incorporated into the scheme design if desired. We have discussed results from simulating heterogeneous funds in [18].

The collective risk-sharing in both CDC schemes we consider is done through annual pension increases. Every year, a pension increase on accrued benefits is declared, and the increased pension amount becomes the accrued benefit of each member. The same annual pension increase is declared on all members' accrued pension. It is this uniform declaration which makes the scheme a collective risk-sharing one. We will refer to our schemes as *shared-indexation schemes* to highlight this common feature of their design.

It is important to note that these schemes are not the only possible CDC

designs. For example, post-retirement-only CDC schemes which operate as pooled annuity funds have gained increasing attention from the UK industry. Even within whole-life schemes, there are many different possible designs. It is also true that the ‘best’ plan from a purely financial point of view, is probably not the ‘best’ plan if no-one wants it. It may also not be the best plan if it is so complicated that members do not understand what they will receive. However, this paper does not consider these aspects and instead focuses on the financial characteristics of shared-indexation schemes. In particular, the focus is on intergenerational risk-sharing.

Intergenerational risk-sharing is a fundamental part of CDC schemes. It is also the most controversial. The idea is that risks, such as volatile investment returns and uncertain future lifetimes, are shared across the generations of the pension scheme. If this can be done fairly, then all generations with the same risk characteristics will face the same risk-return trade-off.

We will see that a dominant source of cross-subsidy in the single-employer scheme results from the fixed contribution rate and nominal-benefit which is accrued each year. This favours investors who are close to retirement as they receive the same benefit as younger investors for the same contribution, despite their being less time for their investments to grow. This mirrors the behaviour of the DB scheme that the single-employer scheme is designed to replace. Thus, these cross-subsidies should be seen as an intentional design feature of the single-employer scheme. However, we find that the level of cross-subsidy is considerably higher in a single-employer scheme than a DB scheme. An age-related benefit accrual would reduce the cost of this anomaly on later generations. The lack of financial fairness on an accrued basis distinguishes a Royal Mail-type CDC scheme from those studied in the literature. For example, [9], which shows a greater financial fairness of their studied collective scheme, is based on CDC plans which are seen in the Netherlands.

The multi-employer scheme is designed to avoid such obvious subsidies and is instead intended to be actuarially fair in the sense that each year members’ contributions match their new benefit entitlements. However, the implicit pricing methodology used to achieve this is only approximate and so cross-subsidies still appear. It is intuitively clear that this must occur for the fund to succeed in its goal of risk-sharing. The design of the fund 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 have the same feature [2, 4, 7] and they can also out-perform alternative pension options by taking on more investment risk.

In shared-indexation CDC schemes, transfer of risk between generations must imply some transfer in the market value of assets, and so some transfer of wealth must occur to achieve smoothing. The existence of these cross-subsidies is not immediately apparent from the central-estimate pricing formulae that one is required to use in valuing UK CDC funds. Instead, one must use a rigorous pricing methodology such as risk-neutral pricing to compute them. This source of cross-subsidies will occur in both single- and multi-employer schemes,

but is more apparent in multi-employer schemes where it is the main source of intergenerational cross-subsidy. These cross-subsidies are stochastic and tend to cancel to a large extent over a lifetimes' investment. Nevertheless, we find that in any given year these subsidies may be surprisingly large, with some generations losing up to 40% of the value of their contributions in some years. This implies that a significant employer subsidy may be required to ensure that it is always in each generations interest to invest in a multi-employer shared-indexation scheme.

An additional source of intergenerational cross-subsidy in single-employer CDC schemes arises from what we will call *infinite-horizon effects*. A single-employer CDC scheme is designed so that if one uses a deterministic economic model, it will move into a steady-state which achieves a desired pension outcome. However, there is no reason to assume that in a steady-state the value of the pension received by each generation will match the value invested. In a steady state, it is possible for each generation to either consistently subsidize, or consistently be subsidized by subsequent generations. As a simple example, consider a situation where each generation consumes the next generations' pension investments, with the first generation also consuming their own pension investments. This would result in a steady-state over an infinite-time horizon, but if the scheme is terminated at any time, the final generation will lose out. We see that because of these infinite-horizon effects, it is possible for a CDC scheme to yield lower pension outcomes in a steady-state than an annuity, even though the CDC scheme is investing in assets with higher returns.

We also evaluate the cross-subsidies that occur when actual investment returns differ from their predictions. Since the level of pensions paid out depend on future predictions of returns, poor predictions of those returns lead to pension payments which 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.

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 realised contributions and investments are known, this model describes what the resulting benefits will be. In the case of the multi-employer CDC scheme, this therefore gives a precise operationalisation of the broad shared understanding among pension professionals of what a multi-employer CDC scheme should entail. In Section 2.3, we explain briefly how our understanding of CDC funds corresponds to the UK regulations.

Second, in Section 3, we simulate the behaviour of these funds using the stochastic economic model described in [1]. This allows us to compare the outcomes of these schemes with other possible investment choices such as DC funds and a pooled annuity fund.

Third, in Section 4, we study the extent of intergenerational cross subsidy in

CDC funds. To do this we use the much simpler Black–Scholes model to generate stock-price scenarios and assume other risk-factors grow deterministically. The advantage of this model is that one can compute unambiguous market values for any derivative product, and this includes the income stream one derives from a CDC pension. This allows us to quantify which generations do best from a CDC scheme and which do worst. We also derive closed-form formulae for the intergenerational cross-subsidies and infinite-horizon effects that occur in DB schemes. These provide a useful benchmark for the performance of the single-employer scheme and a simple model that explains why the cross-subsidies and infinite-horizon effects are larger in single-employer CDC schemes.

Our key findings are that: multi-employer shared-indexation schemes are not able to target a given level of annual pension increases as effectively as single-employer schemes; multi-employer shared-indexation schemes have substantially lower intergenerational cross-subsidies than single-employer schemes; multi-employer 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; single-employer schemes do not necessarily outperform annuities; multi-employer schemes can markedly outperform single-employer schemes because they do not experience the same infinite-horizon effects.

2 The operation of shared-indexation CDC funds

In this section, we describe in precise mathematical terms the operation of what we call shared-indexation CDC funds. Some aspects of the operation are common to both single- and multi-employer funds and we describe these first. We then explain the different approaches taken in the single- and multi-employer cases.

We assume there are M distinct types of individual in the CDC fund: possibly of different ages, employers, or other demographic factors. At any time, t , ($t \in \mathbb{Z}_{\geq 0}$) each type of individual ξ ($0 \leq \xi \leq M - 1$) has accrued a nominal benefit amount $B_t^{\xi, \text{cum}}$. Individuals of type ξ who have retired will receive this as their pension in year t , but for individuals who have not yet retired it represents their currently accrued entitlement. We write $\mathbf{1}_t^{R, \xi}$ to be the indicator function taking the value 1 if individuals of type ξ have retired at time t and 0 otherwise. So the pension received by individuals of type ξ is $\mathbf{1}_t^{R, \xi} B_t^{\xi, \text{cum}}$. The initial entitlements satisfy

$$B_0^{\xi, \text{cum}} = 0$$

for all ξ .

Each year, a *prevailing indexation level* is selected, the aim being to increase all benefit entitlements at this level of indexation. We will write this indexation level as $h_t^n = (1 + h_t)(1 + i_t) - 1 \approx 1 + i_t + h_t$ where i_t represents CPI growth. We will call h_t the *indexation above inflation*. The scheme is designed to target a specific level for the indexation above inflation, say $h_t \approx 1\%$ and hence $h_t^n \approx \text{CPI} + 1\%$.

However, it will not always be possible to apply reasonable levels of indexation, so each year one also defines a bonus/benefit-cut level of θ_t . The intention is that in normal years $\theta_t = 1$, indicating no bonus or benefit cut, but if assets and liabilities become too mismatched, benefit cuts or bonuses may be awarded. An individual's benefit entitlement before new contributions are added is given by

$$B_{t-}^{\xi, \text{cum}} := (1 + i_t)(1 + h_t)\theta_t B_{t-1}^{\xi, \text{cum}} \quad \text{if } t > 0. \quad (1)$$

Write A_{t-} for the value of assets at time t before new contributions are added or benefit payments are made. The growth rate h_t is selected to ensure that a central estimate of the future liabilities matches the value of A_{t-} , as we will now explain in detail.

We write $p^\xi(t_1, \ell)$ for the probability of an individual of type ξ surviving to time $t_1 + \ell$ given that they are alive at time t_1 . Each type of individual ξ , has an associated investment strategy that determines what proportion of their investments they will place in risky assets. This gives rise to an associated discount rate $P^\xi(t_1, \ell)$ which can be used in a central-estimate pricing methodology to calculate the value at time t_1 of payments that will be received at $t_1 + \ell$. The investment strategy of the fund as a whole will be a liability-weighted combination of the investment strategies for each individual. At each time point t , we have a projection $i(t, \ell)$ for CPI at the future time-point $t + \ell$. We write N_t^ξ for the number of surviving individuals in the fund of type ξ at time t , all of whom receive a pension that year.

We require for $t > 0$ that

$$A_{t-} = \theta_t \sum_{\xi=0}^{M-1} \sum_{\ell=0}^{\infty} I_t(h, \ell) P^\xi(t, \ell) N_t^\xi B_{t-1}^{\xi, \text{cum}} p^\xi(t, \ell) \mathbf{1}_{t+\ell}^{R, \xi}. \quad (2)$$

The discount rates satisfy

$$P^\xi(t, \ell) = \prod_{k=0}^{k < \ell} (1 + r_{P, \xi, t, k})$$

where $r_{P, \xi, t_1, k}$ is the rate of return predicted at time t_1 on assets over the period $(t_1 + k, t_1 + k + 1)$ when the assets are invested according to the strategy associated with individuals of type ξ . The increment I is given by

$$I_t(h, \ell) = \prod_{k=0}^{k < \ell} (1 + i(t, k))(1 + h).$$

We note that we have not made clear how the prediction is made for asset returns or how CPI is projected: these are modelling issues that we will discuss later in Section 3 and Section 4.

The equation (2) determines h_t and θ_t for $t > 0$ in a way that we will now describe. First we try to solve the equation for h with $\theta_t = 1$. If the resulting value is greater than some upper bound h^+ , we cap h and award a one-off benefit

payment. In this case $h_t=h^+$ and θ_t is chosen to ensure that (2) holds. Similarly, if we find that the nominal growth h_t^n is negative, we choose $h_t = -i(t, 0)$ and apply a benefit cut by choosing θ_t so that equation (2) holds.

An important difference between equation (1) and (2) is that the former is a contractual equation that specifies that all fund members must receive the same level of indexation each year. Equation (2), on the other hand, is model dependent and captures the requirement that the indexation level must be chosen, as far as possible, to ensure that assets and anticipated liabilities are equal.

Each type of individual makes a contribution C_t^ξ at time t and receives an additional benefit amount B_t^ξ . So

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

The calculation of benefits and contributions is the key difference between single- and multi-employer schemes and we will discuss how these calculations are performed later in a separate section for each fund type. However, we can continue to define some common notation for both fund types. The assets of the fund after new payments have been received and annual pension payments have been made is denoted A_{t+} and is computed as

$$A_{t+} := A_{t-} + \sum_{\xi=0}^{M-1} N_t^\xi (C_t^\xi - B_t^{\xi, \text{cum}} \mathbf{1}_t^{R, \xi}) \quad (3)$$

We may compute asset values at all times using the equation

$$A_{t-} = \begin{cases} 0 & t = 0 \\ (1 + r_t)A_{(t-1)+} & t > 0. \end{cases}$$

Here r_t denotes the return realised by the fund on their investments over the time period $(t-1, t]$.

2.1 Single-employer schemes

We now describe the design of a single-employer scheme which is intended to provide an alternative to a DB scheme.

Let S_t^ξ denote the salary of individuals of type ξ at time t . We will write $\mathbf{1}_t^{C, \xi}$ for the indicator taking the value 1 if individuals of type ξ make a contribution in year t and 0 otherwise. The contribution level and benefit amount in our single-employer model satisfy

$$B_t^\xi = \frac{1}{\beta_t} S_t^\xi \mathbf{1}_t^{C, \xi} \quad (4)$$

and

$$C_t^\xi = \alpha_t S_t^\xi \mathbf{1}_t^{C, \xi} \quad (5)$$

for some values α_t and β_t . These equations should be seen as contractual: by publishing their values each year a fund specifies exactly what nominal benefit will be received in exchange for each contribution. However, the intention of the fund is that, so long as modelling assumptions do not change, α and β will remain constant. Thus, the model used by the fund will provide fixed values of α and β , but the investor should expect that these values may have to be varied over time if the fund needs to change its modelling assumptions.

The contribution in equation (5) is what we will call the *full contribution*, as it consists of a *direct contribution* by the employee and an additional contribution by the employer. Formally, the employer's additional contribution is not allocated to any specific employee, but only to the fund as a whole. This enables money to be distributed unevenly across generations, which, as we discuss later, is an essential feature of any "DB-lite" type scheme.

When the fund is first created, a benefit level β is selected together with a target for the indexation above inflation h_t . The value of α is then chosen to ensure that, if all risk-factors follow their central estimates, the fund will converge to a steady-state which achieves this target level of indexation.

Computing the value of α from β and the target level of indexation will require a number of additional modelling assumptions which we now describe. We assume that these modelling assumptions also hold in our simulations (except where we state explicitly otherwise).

We assume that the type of an individual ξ is equal to their age. The earliest one can join the fund is at age x , an individual receives their first pension payment at age $x + T$ and is certain to be dead at age $x + \omega$. We will assume that there is an equal number of individuals contributing to the fund of all ages from x up to $X = T + x - 1$. Each year, a new generation of individuals age x arrive of the same size. We may therefore write down the various indicator functions we have discussed above:

$$\begin{aligned} \text{age}(\xi, t) &:= X - \xi + t, \\ \mathbf{1}_t^{R,\xi} &:= \mathbf{1}_{\text{age}(\xi,t) \geq T+x}, \\ \mathbf{1}_t^{C,\xi} &:= \mathbf{1}_{\text{age}(\xi,t) \in [x, T+x)}. \end{aligned} \tag{6}$$

We will assume that mortality is constant. This may be written as:

$$p^\xi(t, \ell) =: p(\text{age}(\xi, t), \ell). \tag{7}$$

We will assume that all investor receive the same salary S_t , regardless of age. We will assume that central estimates for the long-term values of risk-factors such as CPI are unchanging, and will refer to these central estimates as steady-state rates. We will write N^a for the expected number of members of age a in the fund each year, r^S for the steady-state return on growth assets, r^B for the steady-state return on gilts and g for the steady-state wage-growth.

Proposition 1. *Assuming deterministic mortality and other risk-factors growing at the steady-state rates, the following relationship must hold between α , β*

and the steady-state value of prevailing indexation level h^n .

$$\begin{aligned}
& \sum_{a=x}^{x+\omega-1} \left\{ \sum_{\ell=0}^{\omega-1} (1+h^n)^{a+\ell+1-x} (1+r^{\text{net}})^{-1} P_{a+1}(\ell) \hat{N}^{a+\ell+1} (1+g)^{x-a} \hat{B}_a^{\text{cum}} \frac{1}{\beta} \mathbf{1}_{a+\ell+1 \geq T+x} \right. \\
& \left. - \sum_{\ell=1}^{\omega-1} (1+h^n)^{a+\ell-x} P_a(\ell) \hat{N}^{a+\ell} (1+g)^{x-a} \hat{B}_{a-1}^{\text{cum}} \frac{1}{\beta} \mathbf{1}_{a+\ell \geq T+x} \right\} \\
& = \sum_{a=x}^{x+\omega-1} \hat{N}^a \alpha \mathbf{1}_{a \in [x, T+x)}
\end{aligned} \tag{8}$$

where we have made the following definitions:

$$\hat{B}_a^{\text{cum}} := \sum_{i=0}^{T-1} (1+g)^i (1+h^n)^{-i} \mathbf{1}_{a-x-i \geq 0}, \tag{9}$$

$$P_a(\ell) := \prod_{0 \leq k < \ell} (1 + \pi_{a+k} r^S + (1 - \pi_{a+k}) r^B)^{-1}, \tag{10}$$

$$L^a := \sum_{\ell=1}^{\omega-1} (1+g)^{x-a} (1+h^n)^{a+\ell-x} P_a(\ell) \hat{N}^{a+\ell} \hat{B}_a^{\text{cum}} \mathbf{1}_{a+\ell \geq T+x}, \tag{11}$$

$$\pi^{\text{net}} := \frac{\sum_{a=x}^{x+\omega-1} \pi_a L_a}{\sum_{a=x}^{x+\omega-1} L_a}, \tag{12}$$

$$r^{\text{net}} = \pi^{\text{net}} r^S + (1 - \pi^{\text{net}}) r^B. \tag{13}$$

A proof is given in Appendix A.

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 this target rate is any precise sense the “average” rate of increase that will occur in simulations. The proposition simply provides a reasonable heuristic one can use to determine the relation between costs and benefits for the fund. In a deterministic model, h will tend to the target level as $t \rightarrow \infty$. In a stochastic model, we expect that h will be a mean-reverting process, with longterm mean given approximately by the target level.

2.2 Multi-employer schemes

In contrast to a single-employer scheme, a multi-employer CDC scheme is intended as an alternative to a DC fund. 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.

In a multi-employer scheme the benefit amount is designed so that assets match liabilities and so is computed from the contribution using the equation:

$$C_t^\xi = \sum_{\ell=0}^{\infty} I_t(h, \ell) P^\xi(t, \ell) B_t^\xi p^\xi(t, \ell) \mathbf{1}_{t+\ell}^{R, \xi}. \tag{14}$$

If one uses the above equation to compute benefit amounts, one automatically has that assets and liabilities (computed using central estimates) also match after payments have been received and made:

$$A_{t+} = \sum_{\xi=0}^{M-1} \sum_{\ell=1}^{\infty} ((1+i(t,\ell))(1+h_t))^\ell P(t,\ell) N_t^\xi B_t^{\xi,\text{cum}} p^\xi(t,\ell) \mathbf{1}_{t+\ell}^{R,\xi}.$$

If the assets then grow according to the centrally-estimated rate then equation (2) will automatically hold at the next time-period.

However, there is a clear concern with equation (14) that different types of individuals will be charged different amounts for the same future cashflows unless the quantities $P^\xi(t,\ell)$ are the same for all types of individual. Some of the pension providers we spoke to were considering developing multi-employer CDC funds using equation (14) to determine benefits with different discount rates for different generations. Using the techniques of Section 4, one can compute that this results in large intergenerational cross-subsidies (resulting in expected instantaneous profits or losses of the order of 75% for the most affected generation). This runs against the design goal of minimizing intergenerational cross-subsidies.

To avoid this, one can use the overall investment strategy of the fund to determine the discount rate when valuing liabilities. However, if the investment strategy is supposed to be liability-weighted we now have an interdependence between our definition of the liabilities and our definition of the investment strategy. In principle, one could attempt to resolve this by solving a fixed-point problem to identify an overall investment strategy which yields a consistent solution. However, we felt that this approach was unnecessarily complex and computationally challenging. The purpose of pursuing a liability-weighted strategy is simply to take reasonable account of any changes in the age profile of a fund when deciding what proportion to invest in risky assets. There is no reason to believe that an exact liability-weighted strategy is in any sense optimal, so computing an exact fixed point might be overkill.

Instead, in calculating the benefits of a multi-employer fund, we chose a heuristic to find an investment strategy that varies with time, but is the same for all investors. We will describe these heuristics later when we discuss our modelling and the resulting investment strategy is plotted in Figure 1.

Notice that for the multi-employer fund there is no target level of indexation. Equation (14) is designed to ensure that (2) will hold at the next time-period so long as all risk-factors follow their central estimates. So one can, loosely, say that the design of a multi-employer fund results in values of h_t which “on average” will not change over the next time-period. This means that h will follow a random walk. This contrasts with the dynamics of h_t in a single-employer fund which are designed to revert towards the target level of indexation. In a multi-employer fund one can choose h_0 to match some target level, but one cannot expect h_t to have any particular tendency to revert back to this level. Instead, 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 3.4

2.3 Relationship to regulations

We briefly describe how our equations correspond to the UK CDC regulations [8]. Equation (1) captures clause 17.4 (c), that benefit adjustments “must be applied to all the members of the scheme without variation”. Equation (2) is essentially a mathematical formulation of the remainder of section 17 (1-6). These rules specify the general principles that: the assets should match the liabilities (17.4(e)); that central estimates should be used throughout the calculations (17.3); that the same increase should be applied each year (17.4 (a)) and should include the projected change in inflation (17.5 (b)).

The regulators’ interpretation of the regulations, [15], clarifies that “All qualifying benefits in a section must: be provided by reference to the same rate or amount for all members; have the same rate or amount of contributions paid by all members; have the same rate or amount of contributions paid by the employers; have the same normal pension age for all members”. These requirements are captured by equations (4), (5) and (6).

Sections (8-13) of the regulations allow the scheme to make benefit cuts as necessary. There are complex rules on how benefit cuts should be made if they are spread over a number of years. We have for simplicity assumed benefit cuts and bonuses are made immediately.

The required relationship between the contribution rate and the level of benefits is hard to identify from the regulations. However, [15] states that: “A qualifying benefit is subject to periodic adjustment, in accordance with the scheme rules, to achieve a balance between the value of the available assets and the required amount (i.e. the amount expected to be required to provide benefits collectively)”. While most of this idea is already expressed by equation (2), considering the contribution rate is also an important part of how “balance” might be achieved “collectively”. However, there is no mention in this statement of achieving any particular target rate. Nevertheless, in our discussions with industry representatives there appears to be a general understanding that single-employer schemes do attempt to target a particular indexation level.

3 Simulations of CDC funds

3.1 Modelling details

For the simulations detailed in this section, we simulate all risk factors using a minor variation of the economic scenario generator (ESG) described in [1] (see Appendix C 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 3.1, in order to match the Office for Budget Responsibility figures for September 2024 [6].

Stock growth	7.73%
Wage growth	3.83%
CPI growth	2.00%
Index-linked bond growth	4.36%

Table 1: Long-term medians in the model

Although this scenario generator allows one to simulate population mortality, we do not use this feature and instead use the IFA's S1PMA tables. This means we assume that life-expectancy is unchanging throughout the simulation and so we do not need to adjust the relationship between benefits and contributions to account for changing life-expectancy data. For simplicity, we assume that all individuals survive until their first year of retirement and receive the first pension payment. From then on, we follow the IFA tables. We assume throughout that the fund sizes are infinite and so we work with the proportion of survivors in place of the number of survivors, with this proportion changing deterministically.

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).

In our simulations, members join the scheme at age $x = 25$ and retire at $T = 65$. The fund is closed to new contributions after 100 years.

Unless stated otherwise, the single-employer fund we consider targets an annual benefit increment of $\text{CPI} + 0\%$ and an annual nominal benefit accrual of $\beta = 80$. Choosing the target indexation level of $\text{CPI} + 0\%$ aids comparison with other schemes that target a constant retirement income, but we also discuss the results of simulations with a different target level.

When the target indexation is $\text{CPI} + 0\%$, we can calculate a required contribution rate of $\alpha = 6.34\%$. This is then used as the contribution rate for all the other fund types we simulate.

Benefit increments are capped at $\text{CPI} + 5\%$ for both types of CDC fund, and the floor is set at 0% for both. The single-employer fund uses a life-styling strategy for each member of 100% in the risky-asset up to age 65, tapering linearly to 0% at 85. Our industry consultation suggested that a benefit cap

of $\text{CPI} + 2\%$ would be considered a more reasonable level for a multi-employer scheme, but as we explain in more detail later, choosing the lower cap results in a decreasing median pension in retirement.

To choose the strategy for the multi-employer fund, we simulate the single-employer fund using a constant economic model where all risk-factors grow at the median values of our ESG. 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 multi-employer fund. As shown in Figure 1, the result is that our multi-employer fund follows a similar strategy to the median strategy adopted by the single-employer fund.

Alongside all our computations we perform various tests to validate our models. See Appendix E for details of some of the key sanity checks.




Figure 1: (Left) fan diagram showing the deciles of the proportion invested in risky assets in the single-employer CDC fund each year. (Right) the deterministic proportion invested in risky asset each year in the multi-employer scheme.

3.2 Annual increments

In Figure 2, we plot a fan-diagram of the prevailing level of indexation above inflation 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 single-employer scheme is approximately zero until the closure of the scheme. At the end of the scheme we are invested entirely in risk-free assets, so h will remain constant unless a benefit cut occurs due to changes in inflation pushing the total $h + \text{CPI}$ below zero. A bonus is not possible because these are determined by the value of h rather than $h + \text{CPI}$. 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 an infinite fund. For finite funds, 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 multi-employer 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 3, 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 are significant.




Figure 2: Fan diagrams of the indexation level above inflation, h_t , in each year of the simulation.




Figure 3: Fan diagrams of the benefit increase/decrease over inflation including bonuses/cuts in each year of the simulation.

3.3 Pensions for a typical generation

We define the *replacement ratio* to be the ratio between one's income in retirement and one's final salary, taking into account CPI. Figure 4 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 have retired

by the time the fund closes, but the fund has had sufficient time to settle into an approximately “steady state”. In fact, as can be seen from Figure 5, 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 single-employer CDC scheme targeting $\text{CPI} + 0\%$, the median replacement ratio is constant in retirement, but the width of the fan-diagram continues to increase throughout retirement. Since the median income is constant in real-terms, the CDC scheme can be viewed as on average achieving the target.

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

For comparison, we show the income in retirement of a DC-plus-annuity scheme. In this scheme, individuals invest at 100% in the risky asset until 55, then taper linearly to 0% in the risky asset at 65 when they purchase an annuity. The price of the annuity is calculated by using 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 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 55, tapering down to 33% at 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 the cost, C , of an annuity which pays one unit ever year using the mortality distribution used in the simulation and the discount rate r . The pension paid out in the pooled annuity fund is equal to the value of the assets in the fund divided by C . As Figure 4 shows, this scheme results in an approximately constant real-terms income in retirement, but is able to offer a slightly better income than an annuity. 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 annuitisation, show lower median outcomes than the single- and multi-employer schemes, but do have less risk later in retirement.

A plot of the median lifetime-mean replacement ratio for each generation is given in Figure 5, together with a plot of the mean lifetime-mean replacement ratio. The replacement ratio is calculated in proportion to the number of years invested. 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 single-employer scheme, the earliest generations receive a proportionately better pension than those in the steady state and the final generations receive a worse pension. The multi-employer scheme works in a manner which is closer to being actuarially fair: the earliest

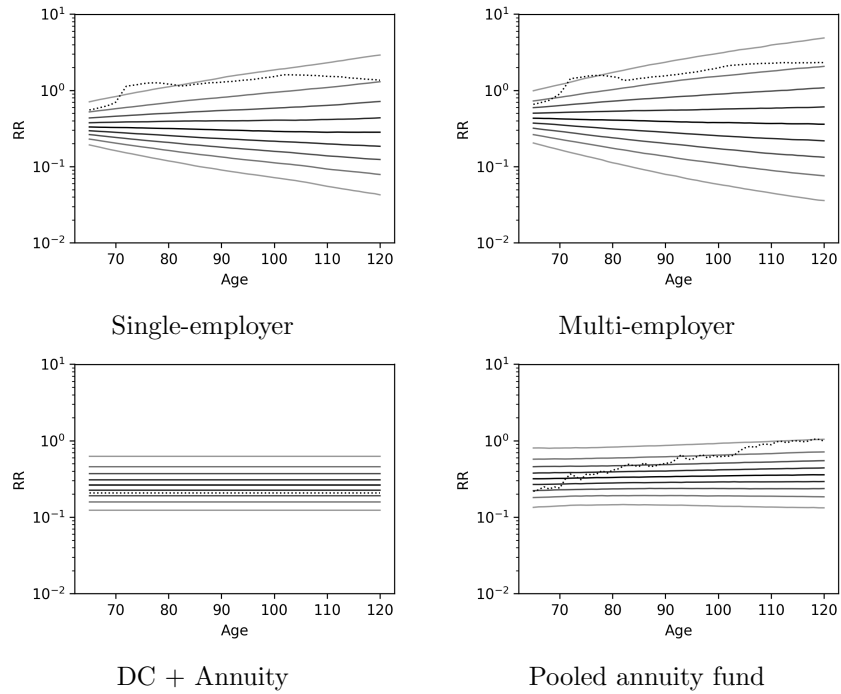


Figure 4: Fan diagrams of the log replacement ratio by age for the 60th generation.




Figure 5: Left: Median lifetime-mean replacement ratio. Right: Mean lifetime-mean replacement ratio. Both plots are generated using our economic scenario generator and a target of CPI + 0% for the CDC funds.

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 employment before the scheme was closed in year 100.

As a simple numerical comparison, we give the median lifetime-mean replacement ratio for the 60-th generation. This is 32.6% for the single-employer CDC scheme, 42.6% for the multi-employer CDC scheme, 26.3% for the DC scheme followed by annuitization, and 33.0% for the pooled annuity fund. By this metric, the single-employer scheme outperformed DC and annuity by a factor of 23%. The comparable result assuming a constant economic model is 12% and is shown as Scenario A in Table 2.

3.4 Changing the target level

So far we have considered the single-employer scheme with a target indexation level of $\text{CPI} + 0\%$. 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.




Figure 6: Plots of median income (log scale) against age for generation 60. Left: for a single-employer scheme with the target level of indexation being varied. Right: for a multi-employer scheme with the level at which bonuses are made being varied.

We plot the results of simulations targeting different levels of indexation above inflation on the left of Figure 6. As we increase the target level, we increase also increase the level at which bonuses are awarded by the same amount. When targeting $\text{CPI} + 0\%$ the median retirement income decreases throughout retirement. If the median income in retirement were to increase at a 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 multi-employer scheme does not incorporate a target level, only an initial level h_0 , we cannot adjust the pension increments in a multi-employer 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 CPI + 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. On the right hand side of Figure 6 we have shown a plot of the median income in retirement for a multi-employer scheme when indexation is capped at CPI + 2% and when it is capped at CPI + 5%. As this plot shows, even if we have an upper bound of CPI + 5% on indexation, the median income in retirement in the multi-employer scheme does not quite manage to keep up with inflation.

4 Intergenerational cross-subsidies

4.1 Intergenerational risk sharing

CDC funds are designed to 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.

To see how this is achieved, consider an analytically tractable model where mortality is exponentially distributed with force of mortality λ , and where inflation, asset growth and indexation of benefits both occur in continuous time with rates \bar{i} , $\bar{\mu}$ and \bar{h} respectively. The liability for an individual retiring at time T evaluated using the CDC methodology would then be

$$L = \int_T^\infty e^{(\bar{i} + \bar{h} - \bar{\mu} - \lambda)t} dt = \frac{e^{(\bar{i} + \bar{h} - \bar{\mu} - \lambda)T}}{\bar{i} + \bar{h} - \bar{\mu} - \lambda},$$

where we assume $\bar{i} + \bar{h} - \bar{\mu} - \lambda < 0$ to ensure convergence. We may then compute that

$$\frac{1}{L} \frac{dL}{d\bar{h}} = \frac{-1 + (\bar{i} + \bar{h} - \bar{\mu} - \lambda)T}{\bar{i} + \bar{h} - \bar{\mu} - \lambda}.$$

Since $\bar{i} + \bar{h} - \bar{\mu} - \lambda < 0$, we see that this is increasing in T . This shows that the relative sensitivity of an individual's liabilities to changes in h is increasing in time to retirement. For more general models, we can define the h -duration of the liabilities to be given by this partial derivative and, just as for the standard duration of a bond portfolio, this duration will increase for longer-dated cashflows.

In Figure 7, we compute the steady-state of a single-employer fund assuming a constant economic model using a recursive approach described in Appendix

B. We then compute the effect of a 10% increase and a 10% drop in the total fund value on the value of liabilities for each generation in a single-employer scheme in its steady state.

These changes are effectively hidden from the membership. From any member’s perspective, the value of h changes from 0 to 0.79% for the upward shock and the from 0 to -0.93% for the downward shock. The accrued benefit amount therefore increases by 2.79% and 1.07% respectively. Thus members are not necessarily aware of any fall in the asset value. While the value of the current accrued pension has changed by radically different proportions for different generations, this is not obvious to the members. Ignorance is not necessarily a bad thing either; [13] makes a lucid and compelling argument for the focus to be on a member’s income in retirement and not on their underlying value.




Figure 7: Percentage changes in liabilities for each generation in a single-employer CDC fund when the total fund-value changes by 10%. The risk-factors are assumed to grow as in Table 3.1 and the fund is assumed to be in its steady-state before the change in total fund-value.

Intuitively, if one transfers risk between generations without changing the return, then one is in effect transferring wealth between generations. The pricing formula (2) does not consider the fact that h varies with time and so does not capture this transfer of wealth. Thus, in order to get a complete picture of intergenerational transfers of wealth, we must use a pricing methodology which takes into account the fact that h is stochastic.

To do this, we simulate the different types of CDC scheme in a model with constant deterministic inflation, wage-inflation and returns on riskless assets, along with a single-risky asset following geometric Brownian motion. In addition, we assume that the fund is infinitely large and that there is no-systematic longevity risk, so that mortality can be treated deterministically. We will refer to this henceforth as a Black–Scholes model. When simulating the dynamics of CDC funds in this model, the “central-estimates” and “projections” required are computed using the mean values predicted by the model.

In a Black–Scholes model, the pension an investor 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 are assuming that we know what others will contribute to the fund, so, in effect, we are assuming that employer contributions are sufficiently high to ensure that there is never any incentive for employees 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 amount paid by the investor. For readers who are unfamiliar with the theory of risk-neutral pricing, we give a very brief account of the parts of the theory we need in Section D.

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 liabilities implied in equation (2), even when one uses the same discount rate for all investors. This is because equation (2) counter-factually assumes a deterministic trajectory for h_t .

In all simulations in this section, the deterministic growth-rates match those of the long-term views used in Section 3. The volatility of the stock and the mortality distributions also match those of Section 3. Similarly, we assume that the fund is closed to new contributions in year 100.

4.2 Single-employer schemes

For the simulation of the single-employer scheme in this section, we assume that the investment strategy associated with an individual depends on their age, with 100% investment in the risky asset up to age 65, tapering linearly to 0% invested in the risky asset at age 85. Since our single-employer scheme is designed primarily as a replacement for a defined-benefit scheme, we expect it to contain significant intergenerational cross subsidy.

The decision to provide a fixed proportion of salary as 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 that occur within DB schemes are well understood. The Equality Act 2010 contains specific provisions to ensure that DB schemes are lawful despite the age-discrimination inherent in such cross-subsidies [12].

At each time, the amount paid by each investor can be compared to the risk-neutral price of the cashflows they will receive. We will call the difference the *instantaneous profit or loss* of investing in the scheme, as, theoretically, an investor could create a replicating portfolio for the cashflows and so guarantee this profit and loss. Notice that this is a profit and loss relative to the full contribution consisting of both the employee and employer’s contribution. Since the employer’s contribution is not specifically allocated to any individual,




Figure 8: Expected instantaneous profit and loss for investors in a single-employer CDC fund, by age and year of operation, evaluated using 20,000 Monte Carlo scenarios.

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 multi-employer schemes and we wish to use a consistent approach to all schemes.

Let $X_s^{t,g}$ denote a random cashflow experienced by an investor in generation g at time s , as a result of their pension payments at time t . $X_t^{t,g}$ will be negative, representing the pension payment itself and $X_s^{t,g}$ will be positive for s in retirement, representing pension payments. The instantaneous profit or loss at time t for generation g is given by the random variable

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

Here $E^{\mathbb{Q}}$ denotes the expectation taken in the risk-neutral measure \mathbb{Q} . We will write \mathbb{P} for the physical measure. We define the *expected instantaneous profit or loss* to be the \mathbb{P} -measure expectation of this, namely:

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

To compute this value, we simulate economic scenarios using the stock drift for 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 8 shows the expected instantaneous profit and loss made by each cohort in each year when the fund is open to contributions. As one would expect, older investors benefit considerably from the single-employer CDC structure.

In Figure 9, we show the aggregate effect for each generation. The older investors at the start of the scheme have a large total expected instantaneous profit, while investors who only contribute early in their career (before the scheme closes) have a considerable expected loss. Let us emphasize that it is a deliberate design feature to give older investors at the start of the scheme higher pension payments than other generations. This is because in the DB scheme this is designed to replace, they will have been receiving lower payments earlier in their career in expectation of higher payments towards retirement. Thus, the pattern qualitatively reflects the existing asymmetries of DB funds.

In Figure 10, we compute the total risk-neutral price of the cashflows received by each generation, evaluated at time 0, that is

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

We see that during the mid-part of the scheme, the risk-neutral value at time 0 of each generations pension cashflows is net negative, although we recall from




Figure 9: Total lifetime expected instantaneous profit and loss for investors in a single-employer CDC fund by generation, evaluated using the data points in figure 8, with linear-interpolation for missing years.




Figure 10: Total lifetime discounted value of net cashflows at time 0 for each generation as a proportion of one year's salary. Calculated with 10^5 Monte Carlo simulations. 95%-confidence interval shown (but barely perceptible).

Figure 9 that the expected instantaneous profit and loss is net positive. One can interpret these results as the earlier generations deciding to speculate 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 10 is equal to zero. However, as we have remarked, generations in the approximate steady-state of the scheme receive an amount with a negative risk-neutral value at time 0. This is possible because the relationship between contributions and benefits in the single-employer scheme is not chosen to ensure any relationship between the value of lifetime contributions and the value of lifetime benefits. One should therefore not expect these to match, even in a model where all assets grow deterministically. The steady-state design of a single-employer scheme considers an infinite horizon and it is possible, in an infinite steady state, either for each generation to always subsidize (or be subsidized by) the subsequent generation. Hence, the equations determining a steady-state do not guarantee financial fairness. We will refer to such phenomena as *infinite-horizon effects*.

As a result of these infinite-horizon effects, it is possible to find economic parameters such that the steady-state of a single-employer CDC fund gives worse average outcomes than a DC scheme followed by annuitisation. Indeed, this can be done in a constant economic model where equity investments deterministically outperform bonds. Thus, the heuristic argument that single-employer CDC funds outperform annuities because the smoothing effects allows them to invest in higher-growth assets is not always valid. Table 2 gives an example of the age-profile and lifestyling strategy required to achieve this in a constant economic model matching Table 3.1. The details of how this computation is performed are given in Appendix B.

In Section 4.4, we study the cross-subsidies and infinite-horizon effects that exist in DB schemes and we see that they are considerably smaller. Thus, while the single-employer CDC behaves in a qualitatively similar fashion to a DB scheme, it is very different quantitatively.

One practical consequence of the observations in this section are that employer contributions must be high in a single-employer CDC scheme in order to guarantee that it is always in an employees interests to contribute to the scheme. In addition, a younger investor in a single-employer CDC scheme should ask what, if any, recompense they will receive for subsidizing older generations if they change employer. Similarly, a younger investor 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 more pressing for investors in a single-employer CDC scheme.

4.3 Multi-employer schemes

Just as for our simulations using the ESG, in this section we chose the investment strategy for the multi-employer CDC fund to match the observed time-

Scenario	A	B	C
Age join workforce	25	18	18
Retirement age	65	67	67
CDC lifestyling start	65	67	Never
CDC lifestyling end	85	87	
CDC replacement ratio	0.361	0.412	0.412
DC lifestyling start	55	57	57
DC lifestyling end	65	67	67
DC replacement ratio	0.337	0.418	0.408
CDC improvement over DC	12%	-1%	6%

Table 2: The replacement ratios in three different scenarios of age profile and lifestyling strategy computed in a constant economic model. The comparison is between a single-employer CDC fund and a DC scheme followed by an annuity purchase. Recall that the annuity purchase includes a 5% premium that is added to represent the insurers need to cover systematic longevity risk

dependent strategy of a single-employer scheme in a deterministic version of the economic model. This ensures that we are using the same discount rate for all investors in equation (2) and so our results highlight the discrepancy between the mathematically rigorous methodology of risk-neutral pricing and the heuristic central-estimate approach implicit in equations (2) and (14).

Figures 11 and 12 show the expected instantaneous profit and loss for the multi-employer fund. Note that the right-hand chart in Figure 12 shows the slice of the 3D plot Figure 11 taken at year 50.

The intergenerational disparities are, as one would expect, much smaller than those seen in a single-employer fund, but are still non-trivial. Of the first members of the fund, the oldest-generation does particularly well and the youngest-generation does particularly badly.

The instantaneous profit and loss is a random variable, so we should consider more than simply its expected value. In Figure 13, we show 50 different scenarios for the instantaneous profit and loss of investors by age at time $t = 50$. This plot is generated using nested Monte Carlo simulations. We use the physical measure to generate 50 different stock price paths up to time $t = 50$ and then perform a Monte Carlo pricing calculation using the pricing measure to calculate the profit or loss made at time $t = 50$.

What these charts show is that intergenerational cross-subsidies still occur in a multi-employer CDC scheme. This may mean that it is not feasible to create a multi-employer scheme if there are significant differences in the ages of employees between employers. This could be mitigated by using separate sections for each employer and using longevity insurance between sections to ensure scalability.

We would not have predicted the specific shape of Figure 11 and cannot provide a complete explanation for it. The highly non-linear determination of h




Figure 11: Expected instantaneous profit and loss for investors in a multi-employer CDC fund, by age and year of operation, evaluated using 20,000 Monte Carlo scenarios.




Figure 12: Expected instantaneous profit and loss for investors in a multi-employer CDC fund. Left shows total profit and loss by generation. Right shows profit and loss in year 50 by age. The left plot was evaluated using the points in Figure 11, with linear-interpolation between years. The right plot was evaluated with 10^5 Monte Carlo scenarios and the 95% confidence interval is shown.




Figure 13: 50 random scenarios simulated in the physical measure showing the instantaneous profit or loss made by each age group which invests in year $t = 50$. Each curve represents a different scenario. The profit and loss in each scenario is computed using a nested Monte Carlo simulation in the pricing measure with 5×10^4 samples.

in equation (2) makes it extremely challenging, if not impossible, to predict such features except through simulation. Thus, although the design of CDC funds is often considered to be “simple” we would question this characterisation.

4.4 Cross subsidies in defined-benefit schemes

To understand the infinite-horizon effects in a single-employer CDC scheme, it is helpful to consider the behaviour of DB and DC schemes when there is no investment in risky assets and when wage inflation and inflation are constant. In this case, we can compute closed-form formulae for the instantaneous profit and loss of each generation and for the ratio of DB performance to DC performance.

Theorem 1. *Suppose all individuals contribute an amount $C(1 + g)^t$ to a DB fund each year from the time they join a scheme to retirement n years later. Each year, their current benefit entitlement grows to match inflation and they receive an additional benefit entitlement $B(1 + g)^t$. If the scheme is operating in a steady-state in a constant economic model with interest rate r , inflation i and wage inflation g , then the instantaneous percentage profit and loss made by an individual who has been employed for k years is:*

$$\left(\frac{n\alpha^k(1 - \alpha)}{\alpha^n - 1} - 1 \right) \times 100\%, \quad \text{where } \alpha := \frac{1 + r}{1 + i}. \quad (15)$$

The ratio of the resulting benefit entitlement achieved in this DB scheme to the annuity payments of an individual who contributes into a DC scheme which

invests only in riskless assets before purchasing an index-linked annuity at retirement is given by

$$\frac{(\alpha - 1)(i + 1)n\alpha^n(r - g)((i + 1)^n - (g + 1)^n)}{(r + 1)(i - g)(\alpha^n - 1)((r + 1)^n - (g + 1)^n)}. \quad (16)$$

In this formula, we assume that there is no additional fee charged when purchasing the annuity to cover systematic longevity risk. In addition, limits should be taken when the denominator vanishes.

As $r \rightarrow g$, the ratio in equation (16) tends to 1.

See Appendix A for a proof.

Theorem 1 shows that infinite-horizon effects also exist in DB schemes, but will be small if the risk-free rate is close to wage inflation. This is the case for the OBR assumptions shown in Table 3.1. In Figure 4.4, we plot the ratio of the DB and DC payoffs when all rates match the OBR assumptions, except for the risk-free rate which we allow to vary. As can be seen, if the risk-free rate differs significantly from wage inflation, the infinite-horizon effects may be large. In a single-employer CDC scheme, one would intuitively expect similar infinite-horizon effects to a DB scheme but where r is chosen to match the returns in risky assets. This explains why infinite-horizon effects are an important issue to consider for single-employer CDC funds, but have not attracted much attention in the study of DB funds.




Figure 14: The ratio of the benefits from a DB fund and a DC fund investing in the risk-free asset as a function of the risk-free rate. All other parameters are as in Table 3.1. The dotted line indicates that the ratio is equal to 1 when the risk-free rate is equal to wage inflation.

Similarly, we see from equation (15) that the level of intergenerational cross-subsidy will typically increase with r . This explains why the intergenerational cross subsidies seen in Figure 8 are so large. In a comparable defined benefit scheme, we find from equation (15) that the profit and loss for the youngest generation would be -39.4% each year, and for the oldest generation it is 53.0% .




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, the chart on the right shows the case when stock returns are 1% greater than expected.

We may read off from this that 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 of the youngest generation. In our single employer scheme, Figure 8 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.

4.5 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 single-employer 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 are enough to significantly change which generation benefits the most from the scheme.

In the multi-employer 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.

5 Conclusions

The single-employer CDC fund seeks to mirror the high levels of intergenerational cross subsidy seen in DB schemes, but we find that the intergenerational cross subsidy is greater in the single-employer CDC scheme. The single-employer

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. While such effects also exist in DB schemes, they will be small. In order to reproduce similar levels of cross subsidy to the DB schemes they replace, single-employer CDC schemes would need to offer varying benefit entitlements by age. This should reduce the infinite-horizon issues experienced in these funds and result in higher average retirement incomes.

The multi-employer fund does not target a specific benefit level in the same way as a single-employer 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 multi-employer CDC design results in much lower intergenerational cross subsidies than a single-employer scheme, though they are still present. The level of cross subsidies in multi-employer schemes are difficult to predict as they arise from differences between the risk-neutral and best-estimate pricing methodologies. As a result, one cannot apply the intuitions of best-estimate pricing to determine what these cross-subsidies are likely to be.

Since the multi-employer scheme has much lower intergenerational cross subsidies than a single-employer scheme, the infinite-horizon effects are small. This explains why the average retirement income is higher in multi-employer schemes.

This paper only considers shared-indexation CDC schemes. Some industry figures we spoke to suggested that giving all generations the same indexation was a form of “fairness”, but we believe that, due to duration effects, shared-indexation represents a transfer of risk between generations and so inherently also represents a transfer of wealth between generations. Instead, the primary sense in which CDC schemes are fair, is simply a consequence of the fact that the same rules are applied over time. This means, that once a steady-state is achieved, different generations will receive outcomes with a similar probability distribution. This is a very weak form of fairness, and even this is only guaranteed at the beginning of the scheme, and not at the time an employee joins the scheme. Multi-employer schemes do additionally ensure (at least approximately), that in the steady-state each generation receives, on average, the total amount they put in. Single-employer schemes do not provide this approximate guarantee. By contrast, the pooled annuity fund gives an example of a collective scheme which provides a much stronger fairness guarantee: 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.

Acknowledgments

We would like to thank the PPI and the members of our project advisory board for their feedback on earlier drafts of this work. This research is funded by Nuffield grant FR-000024058. High Performance Computing was performed using [3].

References

- [1] Sergio Alvares Maffra, John Armstrong, and Teemu Pennanen. Stochastic modeling of assets and liabilities with mortality risk. *Scandinavian Actuarial Journal*, 2021(8):695–725, 2021.
- [2] Jan Bonenkamp and Ed Westerhout. Intergenerational risk sharing and endogenous labour supply within funded pension schemes. *Economica*, 81(323):566–592, 2014.
- [3] CREATE. King’s Computational Research, Engineering and Technology Environment (CREATE), 2024.
- [4] Jiajia Cui, Frank De Jong, and Eduard Ponds. Intergenerational risk sharing within funded pension schemes. *Journal of Pension Economics & Finance*, 10(1):1–29, 2011.
- [5] C Donnelly. Inter-generational crosssubsidies in the uk’s first cdc pension scheme.
- [6] Office for Budget Responsibility. Fiscal risks and sustainability report – september 2024 – charts and tables: supplementary tables, 2024.
- [7] Christian Gollier. Intergenerational risk-sharing and risk-taking of a pension fund. *Journal of Public Economics*, 92(5-6):1463–1485, 2008.
- [8] UK Government. The occupational pension schemes (collective money purchase schemes) regulations 2022, uk statutory instruments 2022 no. 255, 2022.
- [9] Jurre Haan, Zina Lekniute, and Eduard HM Ponds. Pension contracts and risk sharing—a level playing field comparison. *Available at SSRN 2741542*, 2015.
- [10] J Michael Harrison and David M Kreps. Martingales and arbitrage in multiperiod securities markets. *Journal of Economic theory*, 20(3):381–408, 1979.
- [11] J Michael Harrison and Stanley R Pliska. Martingales and stochastic integrals in the theory of continuous trading. *Stochastic processes and their applications*, 11(3):215–260, 1981.
- [12] UK Statutory Instrument. The equality act (age exceptions for pension schemes) order 2010 (si 2010/2133), 2010.
- [13] Robert C Merton. The crisis in retirement planning. *Harvard Business Review*, 92(7/8):43–50, 2014.
- [14] Iqbal Owadally, Rahil Ram, and Luca Regis. An analysis of the dutch-style pension plans proposed by uk policy-makers. *Journal of Social Policy*, 51(2):325–345, 2022.

- [15] The Pension Regulator. Collective defined contribution (CDC).
- [16] Myron Scholes and Fischer Black. The pricing of options and corporate liabilities. *Journal of political Economy*, 81(3):637–654, 1973.
- [17] Amanda Turner and Dirk Zeindler. *Stochastic finance: an introduction with examples*. Cambridge University Press, 2023.
- [18] John Upton. Quantifying multi-employer and single employer cdc outcomes, 2024.
- [19] Lauren Wilkinson. How might cdc develop in the uk?, 2008.

A Proofs

Proof of Proposition 1. If the fund tends to a steady state $\lim_{t \rightarrow \infty} h_t^n = h^n$ when assets grow deterministically, we must have

$$\lim_{t \rightarrow \infty} \frac{A_{(t+1)-}}{A_{t+}} = (1 + r_P^{\text{net}})$$

where r^{net} is the steady-state return for the entire fund. We will show at the end of end of the proof that in a steady state the proportion invested in the risky asset is given by equation (12), explaining equation (13).

Using equations (2) and (3) we must require

$$\begin{aligned} & \lim_{t \rightarrow \infty} \frac{1}{A_{t+}} \sum_{\xi=0}^{M-1} \sum_{\ell=0}^{\infty} (1 + h^n)^{\ell+1} P^\xi(t+1, \ell) N_{t+1}^\xi B_t^{\xi, \text{cum}} p^\xi(t+1, \ell) \mathbf{1}_{t+\ell+1}^{R, \xi} \\ &= \lim_{t \rightarrow \infty} \frac{1}{A_{t+}} (1 + r_P^{\text{net}}) \sum_{\xi=0}^{M-1} \left(N_t^\xi C_t^\xi \right. \\ & \quad \left. + \sum_{\ell=1}^{\infty} (1 + h^n)^{\ell+1} P^\xi(t, \ell) N_t^\xi B_{t-1}^{\xi, \text{cum}} p^\xi(t, \ell) \mathbf{1}_{t+\ell}^{R, \xi} \right) \end{aligned}$$

If wage growth is a deterministic constant g , then pension contributions will be $\alpha(1 + g)^t \mathbf{1}_t^{P, \xi} S_0$. Hence we require:

$$\begin{aligned} & \lim_{t \rightarrow \infty} \frac{1}{A_{t+}} \sum_{\xi=0}^{M-1} \left\{ \sum_{\ell=0}^{\infty} \frac{(1 + h^n)^{\ell+1}}{1 + r_P^{\text{net}}} P^\xi(t+1, \ell) N_{t+1}^\xi B_t^{\xi, \text{cum}} p^\xi(t+1, \ell) \mathbf{1}_{t+\ell+1}^{R, \xi} \right. \\ & \quad \left. - \sum_{\ell=1}^{\infty} (1 + h^n)^{\ell+1} P^\xi(t, \ell) N_t^\xi B_{t-1}^{\xi, \text{cum}} p^\xi(t, \ell) \mathbf{1}_{t+\ell}^{R, \xi} \right\} \\ &= \lim_{t \rightarrow \infty} \frac{1}{A_{t+}} \sum_{\xi=0}^{M-1} (1 + g)^t N_t^\xi \alpha \mathbf{1}_t^{P, \xi} S_0. \end{aligned}$$

We now assume that mortality is deterministic too and replace N_t^ξ with \underline{N}_t^ξ , the expected number of survivors at time t .

$$\begin{aligned} & \lim_{t \rightarrow \infty} \frac{1}{A_{t+}} \sum_{\xi=0}^{M-1} \left\{ \sum_{\ell=0}^{\infty} \frac{(1+h^n)^{\ell+1}}{1+r_P^{\text{net}}} P^\xi(t+1, \ell) \underline{N}_{t+\ell+1}^\xi B_t^{\xi, \text{cum}} \mathbf{1}_{t+\ell+1}^{R, \xi} \right. \\ & \quad \left. - \sum_{\ell=1}^{\infty} (1+h^n)^{\ell+1} P^\xi(t, \ell) \underline{N}_{t+\ell}^\xi B_{t-1}^{\xi, \text{cum}} \mathbf{1}_{t+\ell}^{R, \xi} \right\} \\ & = \lim_{t \rightarrow \infty} \frac{1}{A_{t+}} \sum_{\xi=0}^{M-1} (1+g)^t \underline{N}_t^\xi \alpha \mathbf{1}_t^{P, \xi} S_0. \end{aligned}$$

We may rewrite this in terms of an age-at-time- t index, a , in place of ξ . To do this we write $P_a(\ell)$ for the discount rate experienced in a steady-start from age a to age $a + \ell$, explaining equation (10). This yields:

$$\begin{aligned} & \lim_{t \rightarrow \infty} \frac{1}{A_{t+}} \sum_{a=x}^{\infty} \left(\sum_{\ell=0}^{\infty} \frac{(1+h^n)^{\ell+1}}{1+r_P^{\text{net}}} P_{a+1}(\ell) \underline{N}_{t+\ell+1}^{X-a+t} B_t^{X-a+t, \text{cum}} \mathbf{1}_{t+\ell+1}^{R, X-a+t} \right. \\ & \quad \left. - \sum_{\ell=1}^{\infty} (1+h^n)^{\ell+1} P_a(\ell) \underline{N}_{t+\ell}^{X-a+t} B_{t-1}^{X-a+t, \text{cum}} \mathbf{1}_{t+\ell}^{R, X-a+t} \right) \\ & = \lim_{t \rightarrow \infty} \frac{1}{A_{t+}} \sum_{a=x}^{\infty} (1+g)^t \underline{N}_t^{X-a+t} \alpha \mathbf{1}_t^{P, X-a+t} S_0. \end{aligned}$$

Rewriting the indicator functions in terms of age and noting that the expected number of individuals alive of age a at time t is a known constant, \hat{N}^a , and that this is zero for $a \geq x + \omega$, we have

$$\begin{aligned} & \lim_{t \rightarrow \infty} \frac{1}{A_{t+}} \sum_{a=x}^{x+\omega-1} \left\{ \sum_{\ell=0}^{\omega-1} \frac{(1+h^n)^{\ell+1}}{1+r^{\text{net}}} P_{a+1}(\ell) \hat{N}^{a+\ell+1} B_t^{X-a+t, \text{cum}} \mathbf{1}_{a+\ell+1 \geq T+x} \right. \\ & \quad \left. - \sum_{\ell=1}^{\omega-1} (1+h^n)^{\ell+1} P_a(\ell) \hat{N}^{a+\ell} B_{t-1}^{X-a+t, \text{cum}} \mathbf{1}_{a+\ell \geq T+x} \right\} \\ & = \lim_{t \rightarrow \infty} \frac{1}{A_{t+}} \sum_{a=x}^{x+\omega-1} (1+g)^t \hat{N}^a \alpha \mathbf{1}_{a \in [x, T+x)} S_0. \end{aligned} \tag{17}$$

An individual in generation $\xi = X - a + t$ is of age a at time t , and was of age x at time $x - a + t$. Let the index i denote the number of years an individual had been employed when they make a given contribution. Their age was $x + i$, so the time that has passed since they made the contribution is $a - x - i$. Using

this information, we may write down their accumulated benefit

$$\begin{aligned} B_t^{X-a+t, \text{cum}} &= (1+g)^{x-a+t} \sum_{i=0}^{T-1} (1+g)^i (1+h^n)^{a-x-i} \frac{1}{\beta} S_0 \mathbf{1}_{a-x-i \geq 0} \\ &= (1+g)^{x-a+t} (1+h^n)^{a-x} \frac{1}{\beta} S_0 \hat{B}_a^{\text{cum}}. \end{aligned}$$

This explains our defining equation (9) for \hat{B}_a^{cum} . The steady-state equation then becomes

$$\begin{aligned} &\lim_{t \rightarrow \infty} \frac{1}{A_{t+}} \times \\ &\sum_{a=x}^{x+\omega-1} \left\{ \sum_{\ell=0}^{\omega-1} \frac{(1+h^n)^{\ell+1}}{1+r^{\text{net}}} P_{a+1}(\ell) \hat{N}^{a+\ell+1} \times \right. \\ &\quad \left. (1+g)^{x-a+t} (1+h^n)^{a-x} \hat{B}_a^{\text{cum}} \frac{1}{\beta} S_0 \mathbf{1}_{a+\ell+1 \geq T+x} \right. \\ &\quad \left. - \sum_{\ell=1}^{\omega-1} (1+h^n)^{\ell+1} P_a(\ell) \hat{N}^{a+\ell} (1+g)^{x-a+t} (1+h^n)^{a-1-x} \hat{B}_{a-1}^{\text{cum}} \frac{1}{\beta} S_0 \mathbf{1}_{a+\ell \geq T+x} \right\} \\ &= \lim_{t \rightarrow \infty} \frac{1}{A_{t+}} \sum_{a=x}^{x+\omega-1} (1+g)^t \hat{N}^a \alpha \mathbf{1}_{a \in [x, T+x]} S_0. \end{aligned}$$

Eliminating the factor of $\frac{S_0(1+g)^t}{A_{t+}}$ to obtain a time-independent equation yields equation (8).

All that remains is to calculate α^{net} , the steady-state proportion invested in risky-assets. We know from the definition of the fund's investment strategy that α^{net} should satisfy

$$\alpha^{\text{net}} = \lim_{t \rightarrow \infty} \frac{\sum_{a=x}^{\infty} \pi_a L_t^a}{\sum_{a=x}^{\infty} L_t^a}.$$

with L_t^a equal to the fund's liability to individuals of age a at time t after benefits have been paid.

The liability L_t^a satisfies

$$\begin{aligned} L_t^a &= \sum_{\ell=1}^{\omega-1} (1+h^n)^\ell P_a(\ell) \hat{N}^{a+\ell} B_t^{X-a+t, \text{cum}} \mathbf{1}_{a+\ell \geq T+x} \\ &= \sum_{\ell=1}^{\omega-1} (1+g)^{x-a+t} (1+h^n)^{a+\ell-x} P_a(\ell) \hat{N}^{a+\ell} \hat{B}_a^{\text{cum}} \mathbf{1}_{a+\ell \geq T+x}. \end{aligned}$$

When calculating α^{net} we can cancel the time-dependent term, motivating our definition of L^a in equation (11), and justifying the definition for π^{net} in equation (12).

□

Proof of Theorem 1. Write A for the cost of a unit index-linked annuity paying from retirement to death for an individual who purchases this annuity when they join the scheme. The additional liability of the fund in year t is

$$\sum_{k=0}^{n-1} \alpha^k AB(1+g)^t$$

which must match the additional contribution $nC(1+b)^t$. We compute

$$C = \frac{1}{n} \sum_{k=0}^{n-1} \alpha^k AB = \frac{1}{n} \left(\frac{\alpha^n - 1}{\alpha - 1} \right) AB. \quad (18)$$

The percentage profit and loss of an individual who has been employed for k years is

$$\left(\frac{AB\alpha^k}{C} - 1 \right) \times 100\%.$$

Substituting in our expression for C gives equation (15).

Write $b(t)$ for the total annual benefit accrued by year t for an individual who joins the DB scheme in year t_0 . It satisfies the difference equation

$$b(t) = B(1+g)^t + (1+i)b(t-1), \quad b(t_0-1) = 0.$$

Solving this (and assuming $i \neq g$) we compute that the total annual benefit accrued by the final year of employment is

$$\frac{B(g+1)^{t_0} ((i+1)^n - (g+1)^n)}{i-g}.$$

The formula remains valid if one takes the limit $i \rightarrow g$. Multiplying this by $(1+i)$ and solving for B using the steady state equation (18) we find that their first pension payment will be

$$B^{\text{cum,DB}} := \frac{C(i+1)nr(1+g)^{t_0} ((i+1)^n - (1+g)^n)}{A(i-g)((r+1)^n - 1)}.$$

The total annual benefit accrued by an individual who joins the DB scheme in year t_0 is

$$\begin{aligned} \sum_{k=t_0}^{t_0+n-1} B(1+g)^k &= \frac{B(1+g)^{t_0} ((1+g)^n - 1)}{g} \\ &= \frac{Cnr(r+1) ((g+1)^n - 1) (g+1)^{t_0}}{Ag((r+1)^n - 1)} \quad \text{by (18)}. \end{aligned}$$

If an individual contributes an amount $C(1+g)^t$ into a fund which grows at the rate r every year from $t = t_0$ up to retirement then the fund value satisfies the difference equation

$$f(t) = C(1+g)^t + f(t-1)(1+r), \quad f(t_0-1) = 0.$$

Solving this, the value of their fund at the beginning of their final year in employment is

$$f(t+n-1) = \frac{C(1+g)^{t_0} ((1+r)^n - (1+g)^n)}{r-g}$$

The formula remains valid if one takes limits as $r \rightarrow g$. The cost of a unit index-linked annuity at retirement is $\alpha^n A$, so the benefit they will achieve is

$$B^{\text{cum,DC}} := \frac{C(r+1)\alpha^{-n}(g+1)^{t_0} ((r+1)^n - (g+1)^n)}{A(r-g)}.$$

By comparing this with the steady-state equation (A) we obtain equation (16).

The final claim can be checked using the identity

$$\lim_{a \rightarrow b} \frac{a^n - b^n}{a - b} = b^{n-1}$$

□

B Recursive form of steady-state equations

As an alternative method of calculating the steady-state behaviour of the fund, one can observe that the investments in the fund at each time in the steady state should be proportionate to current. This allows one to recursively compute the accumulated benefit of one age-group at an instant in time by multiplying the accumulated benefit of those who are a year younger, applying wage inflation and then adding the benefit accrued in that year. From this it is possible to recursively compute all the assets and liabilities of a fund in its steady state up to a scale factor proportional to current salaries. The details of the computation are given in Table 3. We have used essentially the same notation as in the rest of the paper except that in Table 3 we use subscript a to denote indexes by age at a fixed instant in time.

The computation in Table 3 is easily implemented in a spreadsheet. One can then use the spreadsheet optimization tools to find the value of α which ensures that assets remain equal to liabilities after payments, investment growth and wage inflation. This is done by solving the equation:

$$(1+g) \sum_{a=x}^{\omega} L_a = \sum_{a=x}^{\omega} A_a^+.$$

C Economic scenario generator

To generate economic scenarios we initially used the scenario generator described in [1]. However, we then found that there was a surprisingly high probability of benefit cuts at start-up of a CDC scheme. This could be traced to the handling

Term	Interpretation	Value
S	Salary	
g	Annual wage inflation	
r_a	Asset returns from age a to $a + 1$	
p_a	Probability of dying between ages a and $a + 1$	
C_a	Is an individual contributing at age a ?	$\begin{cases} 1 & a \geq X \text{ and } a < T \\ 0 & \text{otherwise} \end{cases}$
R_a	Is an individual retired at age a ?	$\begin{cases} 1 & a \geq T \\ 0 & \text{otherwise} \end{cases}$
I_a	Contribution of those age a	$\alpha S C_a$
B_a	Additional benefit at age a	$\frac{1}{\beta} S C_a$
$B_a^{\text{cum},-}$	Benefit accrued by age a before new contribution	$B_a^{\text{cum},-} = \frac{B_{a-1}^{\text{cum}}(1+i+h)}{1+g}$
B_a^{cum}	Benefit accrued by age a	$\begin{cases} 0 & a = x - 1 \\ B_a^{\text{cum},-} + B_a & \text{otherwise} \end{cases}$
N	Proportion of original cohort surviving to age a	$\begin{cases} 1 & a = x - 1, \\ N_{a-1} p_{a-1} & \text{otherwise} \end{cases}$
O_a	Payment out to those of age a	$R_a B_a^{\text{cum}} N_a$
ℓ_a	Liability per survivor age a , per unit benefit accrued	$\begin{cases} 0 & a = \omega, \\ \ell_{a+1} p_a r_a (1+i+h) + R_a & \text{otherwise} \end{cases}$
L_a	Liability before new contributions to individuals of age a	$\ell_a B_a^{\text{cum},-} N_a$
A_a	Assets held after contributions and benefit payments	$L_a - O_a + I_a$
A_a^+	Assets held at start of next period	$A_a(1+r_a)$

Table 3: Recursive formulae for assets and liabilities in a steady state for a constant economic model. Where no formula is given, the term is determined directly by the input data.

of CPI inflation in the model of [1], which allows deflation to occur with a relatively high probability. See Figure 8 of [1].

To address this issue we performed an additional transformation of the variable used to model CPI, similar to the transformations used for the yields to maturity in the original model. The original paper defined a variable

$$I_t = \log \left(\frac{\text{CPI}_t}{\text{CPI}_{t-1}} \right)$$

and used this as one of the vector components in a vector auto-regressive model. We defined a transformed variable

$$\tilde{I}_t = \log(I_t + \mu)$$

where $\mu = 0.01$ and used this as a vector component of the model in place of I_t .

We then calibrated the model using the same data set as [1] and applied the views described in Table 3.1 to determine the long-term medians of the risk factors.

D Risk-neutral pricing

We give a very brief account of the risk-neutral pricing methodology for derivatives traded in the Black–Scholes model. The theory of pricing in these models was developed in [16] using partial differential equation methods. Their approach was subsequently given a probabilistic interpretation, with the papers [10] and [11] providing a thorough account of the resulting theory of risk-neutral pricing in markets driven by diffusion processes. The theory is mathematically sophisticated, but is a standard topic in financial mathematics postgraduate courses and is described in many textbooks such as [17]. In this section we will describe *how* to price derivatives in the Black–Scholes model, but we will not attempt to justify why the methodology is correct, beyond the briefest remarks.

In the Black–Scholes model, one assumes that the logarithm of the stock price, z_t , can be simulated using the following recursive formula:

$$z_t = z_{t-\delta t} + \left(\mu - \frac{1}{2}\sigma^2 \right) \delta t + \sigma\sqrt{\delta t} \epsilon_t \quad (19)$$

where δt is a chosen time step, μ and σ are constants and the ϵ_t are independent, identically distributed standard normal random variables. The constant σ represents the volatility of the stock and μ determines the expected returns. The $\frac{1}{2}\sigma^2$ is chosen to ensure that the expected stock price at time t is $\exp(z_0 + \mu t)$. This model is called the *physical probability model* and it represents our assumptions about the probabilities of different scenarios occurring.

The only other asset in the Black–Scholes model is a risk-free bond which grows at a continuously compounded rate of r . It is assumed in the Black–Scholes model that you can trade in continuous time. In other words, one can

trade on any discrete-time grid and one can also consider limiting strategies as $\delta t \rightarrow 0$.

A derivative contract with maturity T on the stock is simply some chosen function of the path $\mathbf{z} := (z_t)_{t \in [0, T]}$. In practice, a derivative can only depend upon the values of z at discrete time points. We will assume these points all lie on an evenly spaced grid $\{0, \delta t, 2\delta t, \dots, T\}$.

The theory of pricing in the Black–Scholes model shows that, subject to extremely mild technical conditions, the price of any derivative can be computed by simulating the payoff using a slightly different probability model called the *pricing probability model* or the *risk-neutral model*. In the pricing probability model one simulates the stock price using the following recursion formula:

$$z_t = z_{t-\delta t} + \left(r - \frac{1}{2}\sigma^2\right)\delta t + \sigma\sqrt{\delta t}\epsilon_t \quad (20)$$

Note that the difference between equation (19) and equation (20) is that we have replaced μ with the risk-free rate r . According to the theory of risk-neutral pricing, the price of a derivative with payoff function $f(z_0, z_{\delta t}, z_{2\delta t}, \dots, z_T)$ is equal to:

$$E^{\mathbb{Q}}(e^{-rT}f(z_0, z_{\delta t}, z_{2\delta t}, \dots, z_T)) \quad (21)$$

where $E^{\mathbb{Q}}$ denotes the expectation with respect to the risk-neutral model. This should be contrasted with $E^{\mathbb{P}}$ which represents the actual expected value of the payoff, which will normally be different.

This makes it very easy to price derivatives by Monte Carlo methods. Simply simulate a large number of stock price paths in the risk-neutral measure and use this to approximate the expected value of f in the risk-neutral measure. This is how we have computed the value of the payoff from CDC contracts in this paper. We have used the Central Limit Theorem to estimate confidence intervals for the values we compute.

These Monte-Carlo pricing calculations are simple to code, but naturally one asks why the pricing formula (21) is correct. This the challenging point. The essential idea is to show that it is possible to replicate any derivative payoff by trading only in the stock and the risk-neutral bond, so long as one is allowed to trade in continuous time. This is the approach Black and Scholes took, and this allowed them to show that the price of derivatives can be computed using partial differential equation methods. A celebrated theorem due to Feynman and Kac shows that one can compute certain partial differential equations by Monte Carlo methods. If one applies their theorem to Black and Scholes' partial differential equation one obtains equation (21). What is important is that as one can replicate any derivative at a cost given by the risk-neutral price, this means that it is the only possible price for the derivative in any market that is free from arbitrage.

E Testing

The results in this paper depend heavily upon simulations whose results are hard to predict. In particular, Figure 13 is difficult to validate as we have little intuition as to what to expect. However, there are some meaningful tests we can perform to validate our simulations.

- (i) If we perform a simulation using an economic model where all risk-factors are constant, a single-employer CDC fund exactly hits the target level of indexation.
- (ii) The total area under the graph in Figure 10 is close to zero.
- (iii) The plots for the pooled annuity fund scheme and the DC + Annuity scheme are as one would expect.
- (iv) Plots generated using the ESG and using a Black–Scholes model with the same parameterisation are close.
- (v) Our plots of the trajectories for a single-employer CDC scheme demonstrate the properties such as smoothing one expects of a CDC fund. Our plots of the intergenerational cross subsidies occurring in a single-employer CDC scheme have the qualitative properties one expects.
- (vi) The funds remaining at the end of a simulation are approximately zero.
- (vii) Our Python codebase was developed independently from an R codebase for modelling single-employer CDC funds created for [5]. It reproduces the results of that paper.
- (viii) The results for the steady-state of the single-employer CDC simulation and the results for the annuity simulation have been checked using the spreadsheet as described in Appendix B. These also match the closed-form formulae computed for DB simulation.
- (ix) When a pricing chart similar to Figure 8 is created, but with all investment in the risky asset, the results match the analytic formulae computed for a DB fund.

We have, as far as possible, written our code in a modular fashion so that each of these tests increases confidence in the codebase overall.