A guide to understanding CDC

John Southall

Contents

Understanding CDC: a Q&A guide

Introduction

This paper presents a Question & Answer overview designed to explain our thoughts on how Collective Defined Contribution (CDC) schemes work. Its purpose is to clarify key concepts, highlight potential misconceptions, and promote a deeper appreciation of the mechanics, economics and trade-offs underlying CDC arrangements.
It is not an "anti-CDC" or a "pro-CDC" document. Rather, it is pro-coherence and pro-transparency, aiming to offer a balanced analysis. CDC schemes can deliver valuable outcomes for members, but like any financial arrangement, they raise important questions that deserve thoughtful examination of trade-offs.
Transparent understanding supports better decision-making, more informed debate, and fairer scheme design. By laying out the mechanics and our reasoning in an accessible Q&A format, this document aims to contribute to that goal.

A note of caution in advance: this paper is aimed at readers who want to engage seriously with the technical detail of CDC schemes. In our view, there are several misconceptions concerning CDC, and these can only be fully addressed by being explicit about the underlying mechanics and trade-offs.

Four key themes emerge:

(1) CDC achieves higher expected pensions for members by taking more investment risk and pooling longevity risk. While longevity risk is reduced, investment risk is not removed – rather it is redistributed from older to younger members.

(2) CDC involves smoothing changes in payment levels. This improves short-term stability, but increases long-term uncertainty

(3) Intergenerational fairness hinges on appropriate valuations for individuals - not just expected outcomes. We explain steps we believe could make CDC fairer.

(4) There is no free lunch but rather pros and cons of CDC schemes versus other pension arrangements that we summarise.

The paper is split into five parts plus an appendix. Each section builds on the last. The deeper you go, the more technical the paper becomes because the devil is in the detail of the design of CDC. Part I establishes mechanics. Part II develops a coherent definition of fairness and identifies key design challenges. Part III examines risk, sequence effects, and long-term uncertainty. Part IV contrasts whole-of-life CDC with retirement-only CDC (R-CDC), and Part V brings the strands together with concluding remarks.

Part I: CDC Mechanics

Q: What are CDC schemes?

A: CDC scheme blends elements of traditional DB and DC schemes.
There are two broad forms of CDC scheme: *whole-of-life* CDC and *retirement-only CDC* (R-CDC). Whole-of-life CDC covers both the accumulation and retirement phases whereas R-CDC applies only from the point of retirement using funds accumulated elsewhere (in DC, or a transfer-out from DB).

Like DB, contributions are pooled into a single collective fund for all members, rather than being held in individual ring-fenced accounts as DC does. In the case of R-CDC, members make a single contribution at the point of retirement.

With all types of CDC schemes, the scheme is designed to pay a target income for life but, unlike-DB, benefits can go up or down depending on experience. CDC schemes don't receive deficit contributions from sponsors, so any adverse or positive experience is ultimately absorbed by members.
UK CDC schemes also cannot operate with buffers; instead, target benefits are adjusted so that the present value of liabilities is recalibrated to equal the market value of assets at each valuation, which must occur at least annually.

CDC schemes tend to invest more aggressively than other pension scheme arrangements. A whole-of-life scheme could hold 90% of its assets in equities, for example. This leads to higher expected returns and higher expected outcomes – that's the main driver of the expected uplifts versus DC you may have heard about. At the same time the benefit adjustments are typically designed to result in only modest *short-term* changes in pension levels but, as we shall see, lead to potentially larger longer-term changes in pension levels.

Q: How do CDC schemes adjust benefits?

A: Regulations require that CDC schemes value *aggregate* liabilities using a best-estimate basis i.e. the benefits the scheme expects to pay out are valued using expected investment returns as discount rates[1].
With no employer to top up the fund if it underperforms, or potentially extract any surplus if it outperforms expectations, any bad or good experience must be reflected in benefits payable to members. The scheme's assets and liabilities must be valued at least annually, and benefits are adjusted on a regular basis so that the aggregate liabilities equal the assets of the scheme. These adjustments can be achieved:

- By changing the indexation rate of the scheme, i.e. the expected rate that pensions will be increased by each year; or
- By a one-off "scaling" adjustment – a certain percentage uplift or haircut to all future expected payments; or
- Some combination[2].

It is normally considered a key feature of CDC schemes that they attempt to use the indexation rate of the scheme as the primary lever – this is to reduce the *short-term* volatility of pension levels for members.

[1] There are some technical nuances around what we mean by "best-estimate" and "expected", and how discounting works. However, these details don't change the main message here, so we've kept things simple for the purposes of this Q&A.
[2] Some hybrid is likely in practice to prevent the indexation rate of the scheme becoming too high or low.

Q: Equity markets always seem to recover eventually. Does that mean any "losses" i.e. cuts in projected CDC benefits, are only temporary?

A: No. As explained above, CDC schemes value aggregate liabilities using expected returns as discount rates. By construction, this means that any change in assets relative to liabilities is *permanent in expectation* until benefits are adjusted. You can't expect the funding position to bounce back by itself, otherwise it wouldn't be a genuine best-estimate liability basis. After the adjustment happens the change to benefits is expected to be permanent[3].

Of course, it's important that expected return estimates used in the valuation basis are realistic. That could include potentially increased expected returns following a fall in equity markets (or decreased expected returns following an rise in equity markets). As such, the way expected equity returns are set implicitly assumes a degree of "mean reversion" in returns.

To understand how mean reversion works, it's helpful to think about two extremes. First imagine a hypothetical scheme invested entirely in matching bonds (assuming this were possible) i.e. where benefits line up perfectly with bond cashflows. If the value of assets fell by 20%, the value of liabilities would also fall by 20%. In that case, no benefit cuts would be needed because returns on the matching bonds fully mean-revert. This makes sense since no matter what happens in markets the scheme will be able to make its payments[4].

Equities are different because their returns do not fully mean-revert. For a scheme invested purely in equities, a 20% fall in assets might be offset by a 6% fall in liabilities due to higher discount rates (say), with the remaining 14% needing to be dealt with via cuts to expected benefits.

Exactly how *much* mean reversion exists in equity returns is an interesting and important question, but it would be unreasonable – both theoretically and empirically – to believe that equity returns fully mean revert. In Appendix A we outline one possible way of setting expected equity returns that allows for an appropriate (and limited) degree of mean reversion.

This is not to say that equity markets don't have a positive expected risk premium or are not likely to outperform cash and bonds in the long run; that is not a contradiction.

Q: Do CDC schemes share investment and longevity risk?

A: CDC schemes, including R-CDC schemes, share longevity risk[5]. However, describing investment risk as being "shared" or "pooled" in CDC schemes is, at best, loose language and, at worst, misleading[6]. Rather, investment risk is *transferred* or *shifted* from older to younger members to the extent that changes in the indexation rate are used to adjust benefits.

To explain the mechanics of how investment risk is transferred, suppose assets *relative to liabilities* fall by 18% and the scheme duration is 20 years. As explained earlier, this requires a cut to benefits that is permanent in expectation. This might be achieved by cutting the indexation rate of the scheme from 3% per annum to 2% per annum. This 1% per annum cut has a greater impact on longer-dated expected cashflows (payable to younger members) than nearer-term ones (payable to older ones).

[3] Also, pensioners would receive lower payments for a period even if it weren't permanent.
[4] Demographic and other risks aside
[5] But to the extent that death benefits are provided, they necessarily dilute the degree of longevity pooling.
[6] "Pooling" applies when risks have a degree of independence across members and adverse outcomes for some are offset by favourable outcomes for others. "Sharing" suggests symmetry, contemporaneous give-and-take and fairness by construction. Neither of these reflect how investment risk works in CDC.

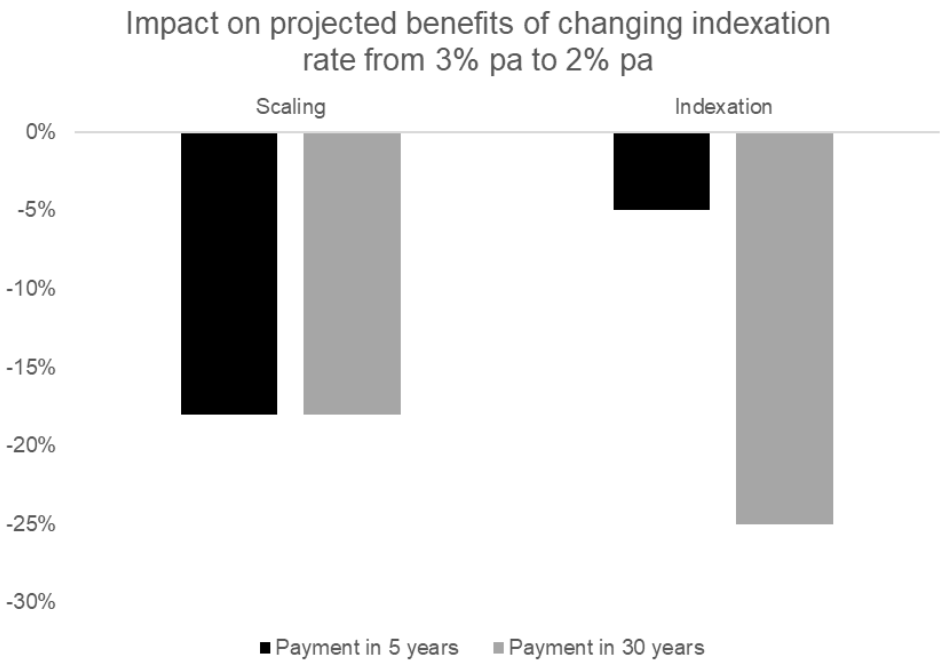

Source: L&G calculations. 20-year duration scheme.

Older members are effectively de-risked by shifting their risk to younger members. For a CDC scheme 100% invested in equities, younger members have investment risk exposure akin to levered equity[7]. This is not necessarily a bad thing – indeed it's like getting a mortgage on a home when you're young such that you have levered exposure to property. But it's important to understand what's really going on.

Q: Hang on, CDC are collective vehicles where everyone experiences the same adjustments to their pension each year. Doesn't that mean there is the same upside and downside for the young as the old?

A: It's easy to see why people might think there is the same upside and downside regardless of age. Suppose there is a reduction in the indexation rate from (say) 3% pa to 2% pa. Looking at a future year - say year 10 from now - all expected payments at that time are reduced by about 9%[8], regardless of member. But note:

- the change in indexation rate is permanent in expectation,
- payments in the more distant future are impacted more (e.g. in year 30 it would be a c25% reduction),
- younger members have more distant payments on average; and
- the value of a members' benefits is the sum of all their discounted future payments.

Younger members[9] are consequently exposed to greater upside and downside.

[7] The amount of effective leverage depends on the investment strategy and duration of the scheme, the age of the member and exactly what the benefit adjustment mechanism is. However later in this paper we show an implicit glidepath for an example scheme.
[8] Calculated as 1 – (1.02/1.03)^10 = 9.3%.
[9] We say "younger" rather than "young" given that e.g. a 65-year-old could be a relatively old member in a whole-of-life scheme but a relatively young member in a R-CDC.

Part II: Fairness: definitions, sources and design challenges

Q: How exactly do you define "fairness" and why do you use such a definition?

A: For the purposes of this paper, we define fairness to mean that members "should neither gain nor lose in risk-adjusted terms at the point of exchange" - i.e., that the present value of the expected benefits awarded to a member, discounted using a rate that properly reflects their risk, is equal to the contributions paid. It applies to accruals in whole-of-life CDC, transfers in or out of the scheme, and on entry to R-CDC. This is:

- an ex-ante concept: fairness is about whether the deal is fair at the point of entry, not whether outcomes happen to be good or bad ex post (i.e. with the benefit of hindsight). As an analogy, consider choosing between a fixed-rate mortgage and a variable-rate mortgage. With hindsight one choice will end up looking like the better option, depending on how rates move. But the point is that the fairness of the choice should be judged at the point the contract is entered into - based on the expected cost and the risks borne - not by whether rates subsequently rise or fall.
- accounts for risk: two arrangements can have the same expected benefits but very different risk. Higher expected benefits are only fair if they compensate for higher risk.
- avoids opaque subsidies: if the present value of expected benefits (risk-adjusted) exceeds a member's contribution then someone else must be subsidising them at that point in time.

Q: What are the potential sources of unfairness in CDC?

A: Unfairness, as we define it, can arise from a few different areas:

- If exchange terms, such as accrual rates in a whole-of-life CDC scheme, ignore members' ages this is unfair because it fails to reflect the time value of money. Contributions made earlier should earn more pension because they are invested for longer, even for a scheme that only uses scaling to adjust benefits (and so doesn't transfer investment risk to younger members). As an example, if expected real returns are 5% pa then a 25-year-old should be awarded approximately 7 times[10] as much expected pension[11] for a given contribution than a 65-year-old, purely due to the time value of money. Age-related accrual is therefore essential for fairness.
- Investment risk is transferred from older to younger members to the extent that changes in the indexation rate of the scheme are used to adjust benefits. If this isn't compensated for in the terms of exchange, that can lead to unfairness.
- A related concern is a risk that new entrants to the scheme reduce or stop, preventing current younger members from passing investment risk on once they become older.
- There is assumption uncertainty. If actuaries underestimate longevity risk, for example, then younger members disproportionately bear the cost.
- Finally, different members will face different mortality rates which may be predictable to some degree based on factors such as postcode, gender, medical history, salary etc. Ideally these would be allowed for to prevent e.g. poorer members subsidising richer members (given richer members are likely to live longer). If health inequalities widen, this will likely lead to greater cross subsidy from poorer members to richer ones. However, there is a trade-off with simplicity, ease of governance and member understanding.
-

These considerations apply to accruals but also apply to transfers in or out, or the terms of entry into a R-CDC scheme.

It's worth noting that DB schemes are not typically fair under our definition, especially if accrual rates are flat. But - as we also comment on later - fairness isn't the only relevant criterion in scheme design.

[10] 1.05^(65-25) = 7.04
[11] Adjusted for inflation

Q: Does sharing risks between members provide fairer outcomes over time?

A: As mentioned, CDC schemes "share" longevity risk but *transfer* investment risk. Even longevity pooling isn't about "fairness" though – it's just a trade-off: longevity hedging in exchange for reduced death benefits (i.e. for inheritance).

As a general principle, fairness under uncertainty can only be judged *before* outcomes are known. Ex-post, once uncertainty resolves, difference alone doesn't tell you if a process was fair.

In our view, "unfairness" in CDC arises if some cohorts *systematically* face worse risk-adjusted prospects due to scheme design, rather than chance. Feeling unfairness after the outcome is known is human but is not a coherent way to judge any process with uncertain results[12].

None of this is to say that schemes shouldn't aim to reduce unrewarded risks as much as possible i.e. minimise unnecessary uncertainty. Later we will explore if CDC schemes can achieve this.

Q: I've read that best-estimate assumptions help avoid cross subsidies between different generations of joiners. Is that wrong?

A: Best-estimate discounting is not fair for individuals in general. This is due to the way investment risk is transferred into the future in any scheme that uses changes in indexation rates to adjust benefits. Near-term expected payments should be discounted at rates close to risk-free, whereas long-term expected payments should be discounted using a rate that allows for a high investment risk premium[13]. Depending on a member's age, their mix of short and long-term benefits will differ. If this isn't allowed for in the discount rate of present value calculations (when money is exchanged for expected benefits) this can be a source of unfairness.

The attractiveness of benefits isn't just a function of their expected level – their uncertainty matters too.

Q: With best-estimate accruals, won't any unfairness average out for a young member who joins the scheme and remains in it for life?

A: It could offset for some members who join young and stay in a scheme with a pipeline of young new entrants. Their accrual terms may be "too low" when they are a young employee but then "too high" when they are older employee. It could happen for a particular individual that this perfectly cancels. That can happen in special cases - and may be more likely to occur in some industries - but in general it won't even out.

Q: Supposing accrual is calculated using best-estimate assumptions, is it a problem if future accrual stops? Is this why critics of CDC say it has Ponzi-like features?

A: Describing any CDC scheme as Ponzi-like is inaccurate given transparent rules for how a scheme will work, even if the implications of those rules for risk-adjusted outcomes are not well understood. However, the point that critics are getting at is that incumbents are *not* indifferent between best-estimate accrual continuing or stopping. Incumbents should prefer the former because young new entrants take on some of their investment risk without

[12] This is "outcome bias" - the tendency to judge a decision as good or bad based on how it turned out, rather than on whether it was reasonable given the information and risks known at the time. Pro poker players are great at avoiding outcome bias (very stoic in the face of individual gains and losses), as are good active fund managers.
[13] The easiest way to perform the calculations is with derivative valuation techniques.

reward via their accrual rates. As such the concern is a reliance on a pipeline of new entrants to maintain risk-adjusted outcomes.

If the pipeline dries up then although expected outcomes can be sustained (if the scheme maintains its expected investment strategy[14]), risk for incumbents must go up relative to if the pipeline had continued. Note that just because expectations can stay the same doesn't mean members should be equally happy – risk matters too. Note there are many plausible reasons a multi-employer CDC scheme might not exist indefinitely. These could include the industry winding down, CDC falling out of fashion (perhaps following sustained poor equity returns such as were seen in Japan from the early 1990s until recently), or bad press due to members not getting what they expected.

Q: How could you calculate the fair value of expected benefits for a member?

A: This is a good question because the ability to calculate fair accruals and transfer terms is central to addressing intergenerational unfairness in CDC schemes, widely seen as their most serious challenge. Crucially, it allows this to be done without sacrificing the features that make CDC attractive, which is why the calculation method is of real practical importance, even if it is necessarily technical.

A fuller answer to how to calculate fair accruals and transfer terms is given in Appendix B, but here we make three comments:

- The main thing to capture in the valuation of these expected cashflows is that, assuming benefits are adjusted primarily via adjustments in the indexation rate, near-dated cashflows should be discounted at a rate close to risk-free, whereas long-dated cashflows should include a high investment risk premium in the discount rate.
- The same result can be achieved computationally using the same techniques that have been developed to value financial derivatives. This can allow for complicated, hybrid benefit adjustments
- It's difficult to allow for all aspects of potential unfairness, but our methodology at least allows for the time value of money and investment risk transfer.

Q: Don't the "gateway tests" ensure fairness?

A: No. There are some initial and ongoing regulator tests that CDC schemes need to satisfy. However, they are weak tests because they don't use a fair value basis (it's all best estimate), and the comparison of the value of active member's benefits is with *member* contributions only, rather than also including employer contributions. Just because a scheme meets the gateway tests doesn't mean to say it is "fair", although the tests may help avoid the most severe cases.

Q: Are CDC schemes inherently and irreconcilably unfair?

A: No. At least in principle there is no additional amount of risk foisted on younger members that can't be compensated for by higher accrual rates (or other exchange terms). The challenge is calculating and applying those rates, whilst keeping explanations simple for members. Some sources of potential unfairness are easier to allow for than others.

[14] If the scheme instead de-risks relative to what was expected, then expected outcomes fall.

Q: Is fairness all that matters?

A: No. Simplicity also matters and there can be a trade-off. However, it's important to at least understand and attempt to quantify that trade-off in scheme design. Transparency also matters. For example, if attractiveness or fairness is judged purely based on expectations – neglecting risk – some members could get a poor deal without realising it.

Part III: Risk and uncertainty in CDC

Q: In DC, outcomes depend heavily on investment performance and economic conditions near retirement which are outside members' control. I've heard that CDC reduces/removes sequencing risk for members? If so, could CDC schemes potentially fairer in a broader sense?

A: It's commonly claimed that CDC reduces sequence risk (and timing risk, which we consider later as a related point). If CDC did reduce these risks, that would be positive in terms of risk-adjusted outcomes, even though it isn't "fairer". But it's unclear that CDC achieves this.

The benefit -adjustment mechanisms CDC schemes use (i.e. primarily indexation changes) mean that the *effective* investment exposure a CDC member is exposed to forms an implicit glidepath as they age. For a whole-of-life CDC scheme invested 100% in equity, that glidepath is levered equity initially (when young, since investment risk is transferred to them) followed by gradual derisking with age[15]. This de-risking glidepath is prone to sequence risk, particularly due to the de-risking that persists during retirement.

The chart below illustrates this for a specific example[16] of a 45-year-old deferred member who retires 20 years later (at age 65). Average experience is the same for both lines, but the assumed sequence of experience is changed. For the solid line, there is good experience in each of the first 20 years, followed by bad experience in each of the following 20 years. For the dashed line, the experience is reversed with bad performance following good.

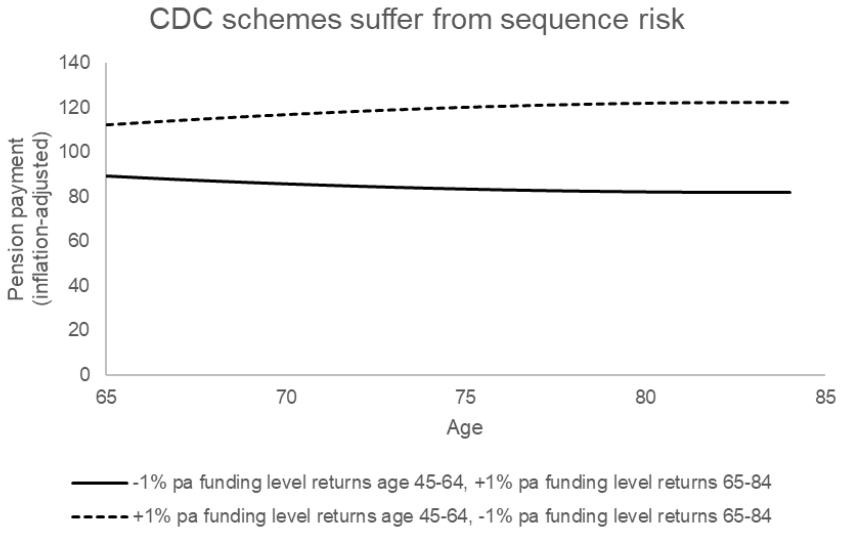

Here a "bad" year is assets *under*performing liabilities by 1% whereas a "good" year is assets *out*performing liabilities by 1%[17]. As you can see there is a very material impact on pensions paid to the member despite average experience being the same. The member is better off if the positive experience happens earlier because an increase in the indexation rate is more valuable the longer it applies for.

[15] In principle (if you could lever equity and hedge longevity whilst taking investment risk), you could replicate outcomes in DC.
[16] 20-year duration scheme. Assumes a pure indexation approach is used to adjust benefits. Source: L&G calculations.
[17] Asset returns relative to liability returns are 'funding level returns'.

Q: What does this implicit glidepath look like?

A: It depends on scheme specifics, but the chart below illustrates an implicit glidepath for a whole-of-life CDC scheme 100% invested in equity, duration 20 years, retirement age 65 and where the benefit-adjustment mechanism is purely changes in the indexation rate:

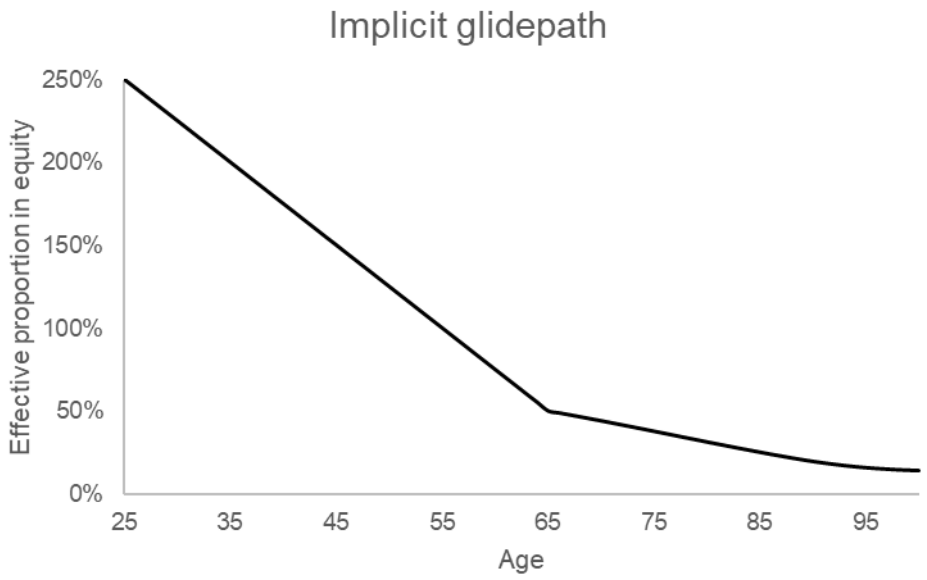

Q: I've heard about "hypothecated" or "notional" glidepaths for members – are they the same thing?

A: No. Hypothecated glidepaths notionally assign a derisking glidepath to each member of the scheme which may be used:

- to determine the overall scheme investment strategy, via a weighted sum of the hypothecated allocations for each member at a point in time. This can mechanically lead to lower duration schemes taking less risk.
- in best-estimate calculations as if the "funds backing a member" move in line with the hypothecated glidepath.

There is no issue with the first bullet, but the second bullet is problematic.
A hypothecated/notional glidepath does *not* reflect the effective investment risk that an individual faces. As an example, if the benefit adjustment mechanism is pure scaling (no changes in the indexation rate) then all members face the same investment risk at a given point in time, regardless of age or any hypothecated glidepath.
A second problem, that we've already identified, is that best-estimate calculations for individuals are flawed because they don't allow for investment risk transfer.

Q: Do CDC schemes reduce the timing risk that members of DC schemes face of de-risking at a bad time?

A: The context here is that a common claim is that CDC reduces the timing risk you see in DC schemes. In DC, a member may be derisked into bonds or annuities (which are essentially bonds plus longevity insurance), only for equities to subsequently perform well relative to bonds. The intuition is that CDC may reduce timing risk thanks to investment risk sharing.
There are grounds to be sceptical of this claim, given:

- As we explained earlier, investment risk is more accurately described as "transferred" rather than "shared".

- The concern over the timing of de-risking is a regret risk[18], identified ex-post (i.e. with the benefit of hindsight). Ex-ante it is far less clear that it is a good or bad time to derisk.
- In DC, members have the flexibility to derisk earlier or later if they have strong views on what isn't priced into markets. But in CDC they have no choice over their de-risking glidepath.
- As we saw earlier CDC schemes also derisk their members over time (there is an implicit glidepath) and expose them to sequence risk.

Q: Does the smoothing mechanism in CDC schemes mean that all pension payments are much more stable?

A: Adjusting benefits via changes in the indexation rate means that year-on-year changes in accrued pensions, and payments made from the scheme, are relatively modest. As such, short-term stability in payment levels is improved. However, the risk doesn't magically disappear and long-term uncertainty in payment levels increases. There's no free lunch. Members suffer worse long-term downside even if they get there smoothly. Thirty years of 2% per annum increases instead of 3% per annum increases amounts to a c25% reduction in expected payment levels, for example.

The chart below illustrates this with a fan chart (percentiles by age) of spending outcomes for a 65-year-old member of a R-CDC scheme[19] under two different ways of adjusting benefits: scaling or indexation:

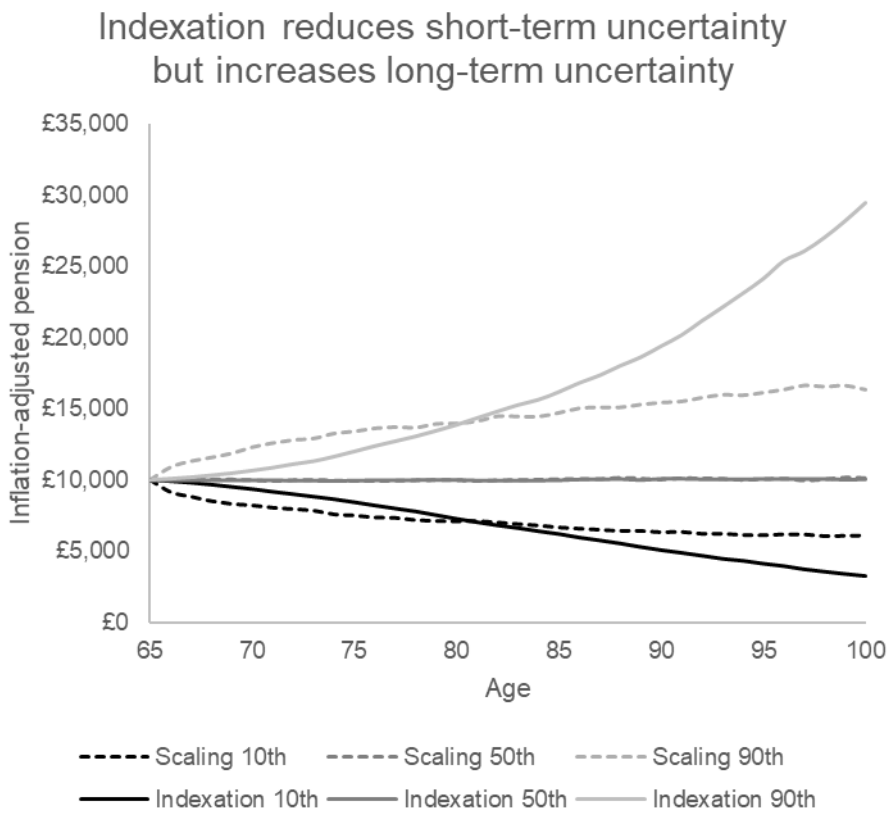

[18] Regret risk is the risk that a decision-maker later feels dissatisfied about a choice because an alternative would have performed better in hindsight, even if the original decision was reasonable at the time. It's closely related to outcome bias but not identical: outcome bias is a *cognitive error* in judging decisions based on outcomes, whereas regret risk relates to the *emotional risk* of feeling regret when outcomes turn out worse than an alternative.

[19] Source: L&G calculations. Duration 10 years and assumed perpetual. Asset vs liability volatility assumed to be 7% pa. Projection starts at retirement aged 65 with a pension of £10k with an initial indexation rate equal to expected inflation. Uses 1,000 simulations.

Q: I've seen analysis that suggests that CDC schemes can provide much higher income than other schemes with lower risk. Is that analysis wrong?

A: Such analysis is potentially misleading. The issue concerns how "risk" is defined. It's common to see risk depicted as the variability in *short-term* changes in pension levels. But CDC mechanically (via the benefit-adjustment mechanism of altering the indexation rate) *chooses* not to change short-term pension levels. This comes at the expense of longer-term variability (the can is kicked down the road) as you can see from the graph in the answer to the previous question. The analysis effectively uses the width of the funnel at the start to measure risk, but this isn't the whole story.

For those who prefer back tests to simulations, some analysis from LCP shows the evolution of real pension levels over various 20-year historic periods. This also clearly shows low short-term variability but high long-term dispersion in pension levels.

Q: But could there be any behavioural advantages to such smoothing of short-term changes?

A: Yes. Indexation in CDC schemes does not remove risk in terms of the ultimate outcomes members face – rather it redistributes it. However, it does change how that risk is experienced. By spreading adjustments gradually over time, CDC schemes tend to produce smoother paths for pension payments over time, rather than sharp and visible cuts. Many members are likely to prefer this smoother experience, even if long-term pension levels are more uncertain, because sudden large reductions are more salient and harder to adjust to or plan around.

Q: Couldn't DC members manage their spending in a similar way?

Yes. Following unexpected experience in decumulation, a DC member could choose to adjust how quickly they expect their spending to grow over time, rather than making an immediate cut to spending. The key difference with CDC relates to investment risk transfer: the effective duration of an individual DC member's spending needs shortens as they age, whereas in CDC it is the duration of the scheme as a whole that matters.

Q: What about "time-diversification" of equity returns? Can't CDC schemes be a better deal for both younger and older members?

A: The idea of time diversification is that longer horizons make risky assets "safer". This is a highly contested notion that is rejected by mainstream financial economics, at least in the way it's commonly used in pensions debates. Some material on this is included in Appendix C for interested readers.

But even if it holds, it is a property of financial markets and so available to all investors (including DC ones), not just CDC schemes. In other words, irrespective of if its right or wrong, it's arguably not a differentiator.
It's also important to remember the implicit glidepath that we saw earlier. From an investment perspective,[20] the advantage CDC has is that it can effectively leverage younger member's equity exposure without using derivatives[21]. This is not the same as taking equity risk for longer.

Views on time-diversification also do not impact the mechanics, modelling or arguments made so far in this paper.

[20] Not a longevity pooling perspective
[21] However, whether this is a good idea is debateable and depends on the nature of human capital. You can read more about this here: DC glidepaths: a random dog walk down Wall Street

Q: Do CDC schemes address adequacy problems?

A: CDC schemes can improve expected outcomes. For example, headline uplifts of roughly 20–60% are commonly quoted when comparing whole-of-life CDC to DC outcomes, and lower uplifts for R-CDC.

However, it is important not to assume that fact means that CDC schemes can address pension adequacy concerns. Like all financial products, CDC arrangements remain subject to the fundamental trade-off between risk and return, and they do not provide inherently superior risk adjusted outcomes.

Ultimately, the size of an individual's pension pot is far more important than the specific retirement vehicle used. If adequacy issues could be solved simply by taking on more investment risk to chase higher returns, then no one would need face financial hardship in retirement. Indeed, if equities truly carried no risk of long-term underperformance, young investors could leverage their exposure indefinitely to secure any desired pension level.

Part IV: Whole-of-life CDC vs Retirement-only CDC (R-CDC)

Q: How does R-CDC differ from whole-of-life CDC?

A: In R-CDC there is a single point of exchange of money for expected benefits – at retirement. This removes the age-related accrual problems. Also, it's expected that pensioners must remain in the scheme, so we don't need to worry about valuing transfers out[22].

Ideally, we should still use a risk-adjusted discount rate to calculate the exchange terms on entry. New entrants have a longer duration than the scheme in aggregate and so will be transferred some investment risk (how much depends on how much indexation is used in the benefit adjustment mechanism of the scheme). However, the risk-adjustment is smaller than that needed in whole-of-life because the duration of an R-CDC scheme is shorter and the scheme will typically invest less aggressively. As such, neglecting risk adjustments out of a desire for simplicity may be more acceptable for R-CDC than for whole-of-life schemes.

A key benefit of CDC is longevity pooling whilst taking investment risk, and retirement is when this really matters given much higher mortality rates then than before retirement.

Selection issues may be more acute in R-CDC than whole-of-life CDC because membership is voluntary and concentrated at a single point in time (retirement), when differences in health, longevity expectations and preferences are larger. That creates stronger scope for who chooses to join to systematically differ from who does not. If enhanced terms aren't offered for those in poorer health, then such members may avoid R-CDC leaving behind healthier retirees. This may require more conservative longevity assumptions and consequently lower expected pension levels.

Q: Is R-CDC potentially preferable to whole-of-life?

You could argue for this, although it isn't clear cut. The case for R-CDC over whole-of-life CDC is:

- Most of the less-controversial benefits of CDC are available in R-CDC, especially:
  - the removal of complex investment and spending decisions for members in retirement
  - the ability to pool longevity risk at the same time as seeking excess investment returns, which is particularly useful in retirement (when mortality rates are much higher)
- R-CDC suffers far less than whole-of-life CDC from potential intergenerational unfairness issues associated with investment risk transfer.
- Employees are not required to grasp a wholly new model relative to DC during the accumulation phase. R-CDC is also easier to frame as an alternative to annuities.
- R-CDC would be a member choice and thus will avoid some points concerning inappropriateness for some members.

On the other hand, the uplifts R-CDC to expected pensions are smaller (given their shorter duration and likely less aggressive investment strategy), and there are potentially greater selection issues. R-CDC also doesn't have the potential upside of being invested in levered equity when you're young.

[22] Whole-of-life CDC schemes must allow transfers out before retirement - members have a statutory right. After benefits have crystallised, they, and R-CDC schemes, are unlikely to permit transfers out because of the heightened selection risk involved (members would likely do it upon discovering they are seriously unwell).

Q: Are R-CDC schemes similar to with profits funds?

A: The underlying mechanics and governance differ but there are some similarities in terms of the economic effects. Both R-CDC and with-profits aim to smooth outcomes for members over time, but they do so in fundamentally different ways. In with-profits funds, smoothing is largely achieved through insurer discretion - via mechanisms such as annual and terminal bonuses, the retention of surpluses in good years, and tools like market value reductions to protect remaining policyholders. In R-CDC, by contrast, there is no separate buffer or insurer balance sheet: benefits are adjusted explicitly so that the present value of scheme liabilities is recalibrated to equal the market value of assets at each valuation, typically using changes to the indexation rate. As a result, smoothing in R-CDC arises from the benefit-adjustment rule itself rather than managerial discretion. That said, the economic effect - shifting risk through time rather than eliminating it - can look similar in practice.

Part V: Final thoughts

Q: Are CDC schemes worth it? Overall, what are the key pros and cons of CDC schemes?

A: Whether CDC schemes are worth it is ultimately subjective. They do not offer a free lunch. A summary of the trade-offs is given below:

Advantages	Disadvantages
Higher expected pensions than DC, typically, due to more aggressive investment	Overall investment uncertainty goes up
Stabilises short-term movements in pension levels due to benefit-adjustment mechanism, or recognition of mean reversion in returns. This leads to a smoother experience over time.	Using changes in the indexation rate to adjust benefits leads to greater uncertainty in pension levels in the long run (investment risk transferred from older to younger members)
Longevity pooling whilst taking investment risk	Higher longevity pooling leads to lower death benefits
No complicated decisions for members (re investment strategy, spending)	Care is needed re intergenerational unfairness (but in theory can be allowed for in accrual and other exchange terms)
Potentially easier to communicate as it provides a pension income, albeit not a guaranteed one	Risk that new entrants/accrual dries up
Could suit large workforces (scale needed) with low employee turnover and the employer/industry expects to exist for decades	Likely to be a poor fit for small workforces orwith high turnover or where the employer/industry could be short lived
	Lack of flexibility for members versus DC. This includes no capacity to tailor investments to personal ethical or religious preferences (for example, Sharia compliant or other ethical constraints).
	Reputation risk should CDC perform poorly or fall out of favour
	Higher overall costs vs DC

Q: Do any other researchers share similar views on the mechanics, economics and trade-offs of CDC?

A: Yes – we developed our ideas and models independently but there is some overlap in our thinking with "Making CDC work for the UK," a collaborative project between The Pensions Policy Institute and King's College London, particularly regarding how benefits should be valued (using derivative valuation techniques). There is also overlap with Professor Catherine Donnelly's CDC research in which she describes a transfer of investment risk, from older to younger members.

Appendix

Appendix A: Setting expected returns on equity

In setting expected returns, key questions include:

- Does the approach lead to asset vs liability volatility that's reflective of long-term risk (due to mean reversion)? Ideally any fall in assets relative to liabilities should be permanent in expectation (until benefits are adjusted), otherwise we are not using genuine expected returns. To do this we need to reflect any mean reversion.
- How do we ensure any 'fundamental' approach based on equity valuations doesn't lead to excessively high or low expected return relative to risk-free rates?

Historic "variance tests" on real equity returns can give us some steer on the strength of mean reversion. Shiller data 1871-2024 indicates there could be a c30% compression in annualised volatility when moving from 1-year to 20-year horizons:

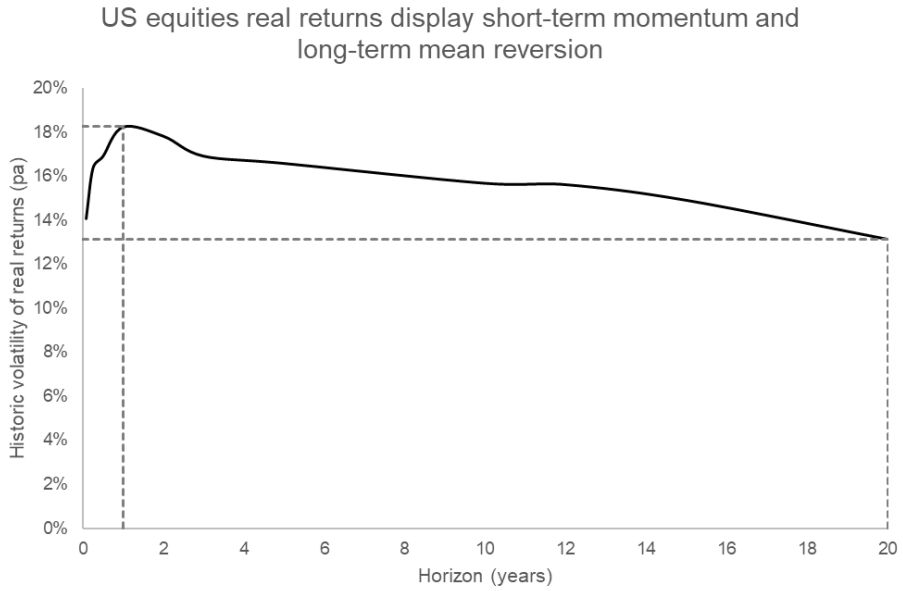

Two common ways of setting expected returns on equity are:

(1) Risk-free rates plus a static equity risk premium (ERP)

(2) A 'fundamental' approach where the expected real return on equity is taken to be the inverse cyclically adjusted price earnings ratio (CAPE).

Under approach (2), expected returns tend to increase after a market crash (because prices fall more than earnings) whereas under approach (1) they don't. Approach (2) tends to overstate mean reversion whereas (1) tends to understate it.

Our proposed method combines the two approaches by using approach (2) to start with but placing bounds on the ERP that we've called FLOOR and CAP. The resulting formula is:

Nominal expected return = risk-free + min(CAP, max(FLOOR, 1/CAPE + expected inflation – risk-free))

To determine appropriate values for CAP and FLOOR we back-tested historic asset versus liability volatility for different values of the FLOOR, CAP and scheme duration. We found that using a FLOOR of 0% and a cap of 8% appears a reasonable approach of getting a volatility term structure broadly in the right ballpark, whilst also preventing the ERP from being extremely high or low. This is just one reasonable approach.

Appendix B: Calculating fair accruals and other transfer terms

How exactly can fair accruals (and other transfer terms i.e. exchanges of money for expected benefits) be calculated? The first question we have to answer is actually "expected benefits under what conditions?". Given fair accruals are *not* based only on expectations (they also allow for risk), whether you assume future accrual occurs impacts expected pensions in respect of existing accrual. Once conditions are assumed, we can value the expected pensions that arise under that assumption.

For the purposes of our research to date, we've assumed that we *expect* the scheme to operate in perpetuity, with the scheme's duration static. Expected pensions correspond to the payments a member receives in the perpetual scheme if investment returns and longevity then deliver in line with expectations.

One way to calculate the value of these expected benefits would be to estimate what the appropriate discount curve looks like. Typically, this would involve discounting near-term payments using rates that are close to risk-free, whereas long-term payments would use a discount rate that includes a high risk premium.

However, a more elegant - and computationally straightforward - approach is to use stochastic risk-neutral valuation techniques. The steps involved are:

- Transform simulated asset returns to a risk-neutral framework.
- Calculate the simulated benefit payments.
- Discount the simulated payments using risk-free rates.
- Take the mean average across simulations.

One thing to note is that if you calculate fair accrual in this way, it does not mean the fair value of incumbents existing accrual should is indifferent to new entrants and accrual stopping. This is because fair value accruals are only neutral subject to the assumptions made. In this case they are conditional on perpetuity of the scheme, but that assumption fails as soon as accrual stops.

Appendix C: On time diversification

- Nobel prize winner Paul Samuelson vigorously disputed the dogma of long-termism, which says that the riskiness of stocks diminishes as time passes. He showed why time horizons matter less than commonly thought. The Economist magazine has a good article on this.
- Tim Gordon's "The Price of Actuarial Values," based on Exley Mehta and Smith's 1997 paper, addresses the "myth" that risk can be diversified over time in section 4.2.8. Section 4.2.6 on the "long term" is also relevant.
- One way to see the problem is to note that if equities were clearly so attractive in the long-term then investors would sell long-term bonds and buy equities until that was no longer the case.
- None of this means that the chance of equities beating bonds can't increase with the time horizon. Investors already know this and it's already in the price. It isn't a contradiction because chance alone is not enough to judge risk. As Cliff Asness writes in why not 100% equities, "It just ain't interesting to show the higher expected return asset has generally a higher realized return with the probability of winning (by some margin) increasing with your time horizon. It's finance 101. It's actually just math 101."