Redistribution in Collective Pension Arrangements Without a Sponsor Guarantee: Hidden Versus Explicit Risk Transactions*

Lu Yi[1], Barbara Sanders[†,2], and Jean-François Bégin[2]

[1]University of North Carolina at Chapel Hill
[2]Simon Fraser University

February 25, 2019

Preliminary Version

Abstract

Collective pension arrangements without a sponsor guarantee are gaining popularity. In these plans, members' assets and liabilities are commingled, and workers' and pensioners' benefits are adjusted to reflect emerging plan experience. Oftentimes, such plans feature design elements intended to stabilize pension benefits (e.g., countercyclical buffers). Stakeholders give significant attention to the fairness of the intergenerational risk transactions arising from the inclusion of these design elements; however, implicit intergenerational risk transactions exist even in pure collective defined contribution arrangements without any such explicit stabilization mechanisms.

Adapting the methodology developed by Hoevenaars and Ponds (2008), we formulate these implicit risk transactions as a collection of exchange options written by different cohorts relative to an individual savings plan benchmark. We estimate the value of these options by Monte Carlo simulation and find that the implicit risk transactions are sizeable and uneven; that is, significant value transfers among cohorts can result from joining even a pure collective defined contribution plan. We then compare these to the additional value transfers arising from (1) introducing explicit stabilization mechanisms, and (2) using alternative approaches for valuing plan liabilities.

Keywords: Funded pension schemes; Hybrid pension plan; Collective defined contribution; Intergenerational risk sharing; Asset-liability management.

JEL Classification: J32, G22.

*The authors would like to thank Gary Parker, and seminar participants at the 2017 Actuarial Research Conference and the 2018 Statistical Society of Canada Annual Meeting for their insightful comments. Yi was a Master's student at Simon Fraser University for part of the project. This manuscript is based on Yi's thesis which won the 2019 Netspar Thesis Award. Yi and Sanders gratefully acknowledge financial support from Simon Fraser University; Bégin is supported by the Natural Sciences and Engineering Research Council of Canada, the Society of Actuaries and Simon Fraser University.

†Corresponding author.

1 Introduction

Collective pensions without a sponsor guarantee are becoming more common in the occupational pension sphere. Under these plans, the plan sponsor's commitment is limited to contributions at a specific level or within a predefined range. Plan members, including those still working and those who already retired, pool their assets and liabilities, and bear risk collectively. In many cases, the plan has a targeted benefit level, or aspiration, but actual benefits may vary from this target on account of economic or demographic experience. The target may be defined in nominal or real terms, and may be strongly or weakly funded.

Examples of such plans include *target benefit plans* and *shared risk plans* in Canada, *collective defined contribution plans* in the UK, as well as most collective plans in the Netherlands. The Dutch plans are classified as either *collective defined contribution* or *defined benefit without a hard guarantee*. In all cases, indexation of benefits (if offered at all) is conditional on the financial position of the plan, and accrued benefits must be cut if the funded position is weak and additional assets cannot be secured. In this paper we use the term collective defined contribution, or CDC, to describe these types of plans.

In a world of dwindling sponsor guarantees, these plans are seen as a desirable alternative to individual defined contribution plans, and rightly so: the benefits of collective pensions with risk-sharing elements are well documented in the literature including Bovenberg et al. (2007), Gollier (2008), Blommestein et al. (2009), and Cui et al. (2011). On the asset side, collective pools tend to create economies of scale which have been shown to improve members' outcomes. On the liability side, temporary subsidies among generations can create stability without sacrificing expected benefit levels; this enhances utility when members are risk averse.

As pointed out by Bovenberg et al. (2007), one disadvantage of risk-sharing designs is that intertemporal smoothing may lead to persistent intergenerational value transfers when shocks are spread out over more than the individual's remaining lifetime. It is, therefore, important to consider the fairness and intergenerational wealth redistribution effect of collective pension schemes.

Evaluation of fairness must take into account all risk transactions within the plan. This includes those that are explicit—that is, arising from stabilization mechanisms that are easily visible to stakeholders, such as countercyclical buffers and contingency reserves—as well as those that are implicit and therefore more difficult to notice or assess. In practice, focus tends to be on the fairness of explicit stabilization mechanisms. Our goal is to increase awareness of the implicit risk transactions and to re-evaluate explicit mechanisms in light of these.

Our conceptual model for a generic CDC plan draws on the work of the Canadian Institute of Actuaries' Task Force on Target Benefit Plans (CIA, 2015b), which identified the following five key elements: the contribution rate, the target benefit, the affordability test, the triggers, and the actions.

The contribution rate is the first design element to be set by the CDC plan sponsors. It can be expressed as a dollar amount or a percentage of the member's salary. In most cases, the contributions paid by the plan sponsors are determined at plan inception and fixed afterwards. In some plans, the contribution rate is adjustable to help balance the plan's funding position; however, such adjustments are restricted to a limited range. The contribution rate and the fund available at plan inception constitute the total funding resources of the plan.

The target benefit is the benefit that the plan is aspiring to provide; unlike in defined benefit (DB) plans, this benefit is not guaranteed. The target is usually defined in a way that is similar to DB plan benefits, e.g., a fixed percentage of career or final average earnings multiplied by the member's years of service. Although the actual benefits received by the plan members vary with plan experience, the target benefit acts as a guidepost, allowing members to estimate their retirement income. It is important that the target benefit be consistent with the chosen contribution rate, so that there can be a reasonable expectation *ex ante* of the target being achieved.

The affordability test is performed at each valuation to determine whether the target benefits are affordable. The result of the test is usually expressed as a funded ratio, equal to a measure of plan assets divided by the pension liabilities. Whether the target benefits appear to be affordable will depend on the choice of assumptions and methods used to value the plan's assets and liabilities.

The triggers are used to determine whether the plan's trustees should take some action given the results of the affordability test. In the case of a pure CDC plan, there is only a single trigger, such that whenever the funded ratio deviates from 100%, action is taken immediately to remedy it. In this case, there is very limited risk sharing among different generations and the actual benefits may be nearly as volatile as in individual defined contribution (DC) plans. To stabilize the retirement income, a lower and upper trigger can be set, forming a corridor inside which no action is taken: in this case the actual benefit will only be adjusted once the funded ratio hits either the lower or the upper trigger. Under such designs, when the plan performs better than expected, only part of the surplus will be paid out for benefit improvements; the rest of the surplus is saved to build a buffer to protect benefits during adverse conditions. This double-trigger design is an example of an explicit benefit stabilization mechanism.

Finally, the actions are the changes that should be made when the triggers are hit. There are several options, which include adjusting the benefits (in respect of past or future service, or both), adjusting the contributions (in plans where the contribution rate is adjustable within a limited range), and/or changing the investment strategy (e.g., updating the asset allocation).

The first two elements (i.e., the contribution rate and the target benefit) determine the plan cost and the funds available, while the last three (the affordability test, triggers and actions) control the risk-sharing structure. Different design choices and assumption sets have an impact on the risk transactions among members, and can significantly change the distribution of the pension benefits payable to different cohorts. We start by comparing the distribution of benefits under a pure CDC plan against the corresponding distribution under an individual DC plan benchmark, in order to isolate any implicit risk transactions. We then explore the impact of additional design elements.

The foundation of our investigation is to use derivative pricing techniques to value contingent claims within the pension fund. Kocken (2006) studied embedded options in DB plans with conditional indexation and also in pure collective DC plans. The pension plan modelled by Kocken (2006) was very simple: the membership was reduced to two groups of participants (actives and retirees) and the retirement benefits were assumed to be paid out through two lump sum payments.

Hoevenaars and Ponds (2008) extended the work of Kocken (2006) by considering each individual cohort instead of the two broad membership groups. Their approach, called valued-based ALM, is a combination of stochastic projections, generational accounting and derivative pricing technique. The plan's operation is projected using Monte Carlo simulation and the real-world cash flows are recorded separately for each generation. The market-consistent value of the pension contract for each generation is then evaluated using risk-neutral valuation.

We apply the value-based ALM method of Hoevenaars and Ponds (2008) to a pure CDC plan without any explicit stabilization mechanisms and find the net market-consistent value of this pension arrangement for each generation. As noted above, this value can be interpreted as the impact (positive or negative, for each cohort) of moving from an individual DC plan (in which there is no risk sharing) to a pure CDC plan, and is therefore a measure of the value of the implicit risk transactions. We explore this value further by decomposing it into the amounts corresponding to the upside and downside risks. We also differentiate between the value derived from benefits received while the plan is ongoing versus the value attached to any residual assets to which a cohort might be entitled at the end of our projection horizon, when the plan is assumed to terminate. Finally, we repeat our calculations for CDC plans with alternative implicit and explicit design elements, including one that uses a different discount rate for valuing plan liabilities, and two others with explicit benefit stabilization mechanisms.

Our analysis leads to four important conclusions. First, we confirm that value transfers do indeed arise from the act of commingling members' assets and liabilities under a pure CDC plan. Second, we observe that changing from a conservative valuation rate based on long-term bond yields to a *best estimate*-type valuation rate based on the expected return on plan assets tends to shift value in a way that counteracts some of the redistribution implicit in the pure CDC plan; nonetheless, the deal remains far from fair on an *ex ante* basis. Third, we find that the fairness of explicit stabilization mechanisms involving two triggers and a no-action range depends on the choice of trigger points. Finally, we note that the generations aged 30 to 55 tend to lose value under all of the designs we explored.

The rest of the paper is organized as follows. Section 2 presents the economic framework used in this study. The structure of the generic CDC plan is described in Section 3. Section 4 discusses the use of the derivative technique. Intergenerational transfers are assessed in Section 5. Then, Section 6 is devoted to the introduction of explicit stabilization mechanisms in the generic CDC plan. Finally, Section 7 concludes.

2 Economic and Financial Framework

An economic scenario generator (ESG) is required to make stochastic projections. It can be defined under both real-world (physical) and risk-neutral (pricing) measures. The real-world version of the ESG focuses on expressing a particular future view of the economy, and it enables us to investigate how the pension plans would perform under specific future economic conditions. By contrast, the risk-neutral version of the ESG is used to consistently evaluate contingent cash flows, whose values depend on the stochastic economic variables. Since we are interested in both the possible future outcomes and the value of intergenerational transfers under different plan designs, we need to construct an ESG that is capable of handling both real-world and risk-neutral measures at the same time.

In this study, we use a model that combines the first-order vector autoregressive model, i.e., VAR(1), and the generalized autoregressive conditional heteroskedasticity model, or GARCH(1,1). The mean reversion feature and the autocorrelation among the economic variables are captured by the VAR(1) model—similar to that used by Hoevenaars and Ponds (2008). In addition, the constant volatility in the VAR(1) model is replaced by a more realistic time-varying volatility, which is modelled by GARCH(1,1) processes.

This study considers five key economic variables. Let r_n^s and r_n^l be the monthly nominal yield of a short-term bond and of a long-term bond at the beginning of month n, respectively. Additionally, let i_n be the monthly inflation rate applicable during month n, r_n be the monthly S&P/TSX Composite Index return cum-dividend in excess of the short rate (i.e., r_n^s), and d_n be its corresponding monthly dividend yield.

Let us fix a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and a filtration $\mathbb{F} = \{\mathcal{F}_n : n \in \{0, 1, ..., N\}\}$ satisfying the usual conditions. The vector $\mathbf{z}_n = [\log(r_n^s), \log(r_n^l), i_n, r_n, \log(d_n)]^\top$ is modelled by a five-dimensional VAR-GARCH model, i.e.,

$$\mathbf{z}_{n+1} = \boldsymbol{\mu} + \boldsymbol{\Sigma}_{n+1}\boldsymbol{\gamma} + \boldsymbol{\beta}(\mathbf{z}_n - \boldsymbol{\mu}) + \boldsymbol{\Sigma}_{n+1}^{\frac{1}{2}}\boldsymbol{\varepsilon}_{n+1}, \tag{1}$$

where $\boldsymbol{\varepsilon}_{n+1} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_5)$ is a five-dimensional vector of standardized innovations and $\mathbf{I}_5$ is the 5×5 identity matrix. Parameter $\boldsymbol{\mu}$ is a five-dimensional vector that controls the constant mean reversion levels for the state variables, $\boldsymbol{\beta}$ is a 5×5 matrix for the autocorrelation coefficients that controls the speed of mean reversion of the economic variables, and $\boldsymbol{\gamma}$ is a five-dimensional vector related to the convexity correction and the equity risk premium, in the spirit of Engle and Granger (1987) and Heston and Nandi (2000). Let $\mathbf{y}_{n+1} = \boldsymbol{\Sigma}_{n+1}^{\frac{1}{2}} \boldsymbol{\varepsilon}_{n+1}$ represent the contemporaneous noise terms. Equation (1) can then be rewritten as

$$\mathbf{z}_{n+1} = \boldsymbol{\nu} + \boldsymbol{\Sigma}_{n+1} \boldsymbol{\gamma} + \boldsymbol{\beta} \mathbf{z}_n + \mathbf{y}_{n+1}, \tag{2}$$

where $\boldsymbol{\nu} = (\mathbf{I}_5 - \boldsymbol{\beta})\boldsymbol{\mu}$ is a five-dimensional vector.

The matrix $\boldsymbol{\Sigma}_{n+1}$ is a time-varying variance-covariance matrix. For the sake of tractability, the correlation between the contemporaneous noise terms is assumed to be zero, meaning that the dependence between the variables is only modelled through the autocorrelation coefficients. The simplified $\boldsymbol{\Sigma}_{n+1}$ is then a diagonal matrix, i.e.,

$$\boldsymbol{\Sigma}_{n+1} = \mathrm{diag}\left(\begin{bmatrix} \sigma^2_{1,n+1} & \sigma^2_{2,n+1} & \sigma^2_{3,n+1} & \sigma^2_{4,n+1} & \sigma^2_{5,n+1} \end{bmatrix}\right).$$

Each conditional variance $\sigma^2_{i,t+1}$ is modelled by an independent univariate GARCH(1,1) process such that

$$\sigma^2_{i,n+1} = \omega_i + a_i y^2_{i,n} + b_i \sigma^2_{i,n} \quad \text{for} \quad i = 1, ..., 5, \tag{3}$$

where a_i is the ARCH parameter that controls how this period's shock will impact the next period's volatility, b_i is the GARCH parameter that controls how persistent the next period's volatility is with respect to the current volatility, and ω_i is a constant term.

The GARCH update equation (i.e., Equation 3) can also be written in a matrix form as

$$\boldsymbol{\Sigma}_{n+1} = \boldsymbol{\omega} + \mathbf{A}\,\mathrm{diag}(\mathbf{y}_n)^2 + \mathbf{B}\,\boldsymbol{\Sigma}_n, \tag{4}$$

where $\boldsymbol{\omega}$, $\mathbf{A}$, and $\mathbf{B}$ are three 5×5 diagonal matrices with ω_i, a_i, and b_i as the i^{th} diagonal elements, respectively, and $\mathrm{diag}(\mathbf{y}_n)$ is a diagonal matrix with vector $\mathbf{y}_n$ on the diagonal.

The parameter $\boldsymbol{\gamma}$ in Equation (1) is a five-dimensional vector, where

$$\boldsymbol{\gamma} = \left[-\frac{1}{2}, -\frac{1}{2}, 0, \gamma_4 - \frac{1}{2}, -\frac{1}{2}\right].$$

This vector has two different purposes. For the short-term bond yield, the long-term bond yield and the dividend yield, the parameter $\boldsymbol{\gamma}$ acts as a convexity correction term because we are modelling the logarithmic value of these variable, therefore making

$$\mathbb{E}^{\mathbb{P}}\left[\exp(z_{i,n+1}) \middle| \mathcal{F}_n\right] = \exp(\nu_i + \boldsymbol{\beta}_i \mathbf{z}_n), \quad \text{for} \quad i = 1, 2, 5,$$

where ν_i is the i^{th} element of vector $\boldsymbol{\nu}$, β_i is the i^{th} row of matrix $\boldsymbol{\beta}$. Then, for the excess stock return, parameter γ_4 is the so-called equity risk premium parameter similar to the GARCH-in-mean process proposed by Engle and Granger (1987).[1]

A corresponding risk-neutral version of our VAR-GARCH model can be derived by assuming that the time-varying risk premiums are affine functions of the the economic state variables. This risk-neutral model is consistent with the stochastic discount factor (SDF) approach used in Hoevenaars and Ponds (2008). More details on the risk-neutralized version of the model is available in Appendix A.

In this study, Canadian economic data from May 1991 to June 2016 are used and more detail on the datasets can be found in Appendix B.[2] To find the model parameters, we conduct a two-stage estimation procedure for the ESG. In the first step, the VAR-GARCH model real-world parameters are estimated using maximum likelihood estimation. In the second step, the risk premium parameters $\boldsymbol{\lambda}_0$ and $\boldsymbol{\lambda}_1$ are calibrated, conditional on the real-world VAR-GARCH parameters, by minimizing the sum of the squared differences between the historical zero-coupon yields and the model zero-coupon bonds yields calculated from the $\mathbb{Q}$ measure using Monte Carlo simulation. This two-stage estimation procedure is consistent with the approach used by Hoevenaars and Ponds (2008). Additional details are available in Appendix C.

[1] As stated in Engle and Granger (1987), the compensation required by risk-averse agents for holding the assets—which has time-varying degree of uncertainty—must also be varying. The GARCH-in-mean model addresses this issue by establishing a risk-return relationship where the risk premium is expressed as a function of the current conditional variance.

[2] We exclude the pre-1991 data because the Bank of Canada adopted a 1%–3% medium-term inflation control target in 1991.

3 Stylized Plan and Key Features

We construct a number of variants of the generic CDC design and simulate their operation under the various economic scenarios generated by the process described in Section 2.

Prior to the inception of the CDC plan, members are assumed to have had individual DC accounts in which they accumulated assets while employed and from which they were drawing assets once retired. The individual DC account balances of all members (active and retired) are brought into the CDC plan and commingled at inception. The pension fund is liquidated after 55 years, and all the remaining assets are distributed to the participants as a lump sum payment at $t = 55$ based on pre-defined entitlement criteria.

Plan membership is stationary with 100 new members entering the plan each year at age 30, retiring at age 65, and dying at age 86. There are no decrements before age 86, so there are 56 cohorts of exactly 100 members each at all times.

Salaries increase at the beginning of each year. The growth rate applicable in year t consists of the annual inflation rate in year t, $\tilde{i}_t$, and a fixed annual increase for promotion and merit, $m = 0.5\%$.[3] The starting salary of the new entrants at time 0 is $\$50,000$. The salaries of subsequent cohorts of new entrants are the same in inflation-adjusted terms.

All active members contribute at the same fixed rate $c = 10.6\%$ of salary, payable at the beginning of each year. The total contributions made to the plan at time t are denoted by C_t.

The target benefit is based on a final-pay formula, where the target annual accrual rate b_t applies to all service (i.e., past and future). Retirement benefits are paid at the beginning of each year based on the current target accrual rate and each retired member's actual final year's earnings. When the target accrual rate is adjusted, the benefits of retired members change. The total benefits paid from the plan at time t are denoted by B_t.

The initial target accrual rate, b_0 is set at 1%. A test is performed at the beginning of each

[3] The annual inflation rate in year t is computed from the monthly inflation rates generated using the model in Section 2. Specifically,

$$\tilde{i}_t = \exp\left(\sum_{n=12(t-1)+1}^{12t} i_n\right) - 1.$$

year, before contributions are received and benefit payments are made, to decide whether the target accrual rate determined in the previous valuation is still affordable. Given the availability of funds and the future contribution commitment, there are two crucial elements that will affect the result of this affordability test: the valuation assumptions and the valuation methods used to evaluate the plan assets and liabilities.

For the plan assets, we assume that half of the fund is invested in the equities modelled in Section 2 and the rest is invested in a rolling portfolio of 15-year zero coupon bonds, rebalanced annually. Annual effective returns for each asset and for the combined portfolio are obtained directly from the simulation results of the VAR-GARCH model.[4] The annual portfolio return in year t—a by-product of the ESG—is denoted by r_t^P. The value of the plan assets F_t can be expressed recursively as

$$F_t = (F_{t-1} + C_{t-1} - B_{t-1})\,(1 + r_t^P).$$

The assumptions needed to value the plan liability relate to mortality, future salary increases and the valuation rate. For purposes of the affordability test, all members are assumed to die at age 86, and salaries are assumed to increase at a fixed annual rate based on a forward-looking inflation assumption of 2% per year and the 0.5% per year increase for merit and promotion. The valuation rate applicable at time t varies with the economic scenarios generated by the process in Section 2. It may reflect the yield on long-term bonds as of the valuation date, or it can be based on the actuary's (imperfect) estimate of the expected return on plan assets as of that date. The long-term bond yield expressed as an annual effective rate, $\tilde{r}_t^l$, is derived from the ESG outputs.[5]

The actuary's estimate of the expected return on the assets (EROA) is constructed by following the "building block approach" in the guidance material published by the CIA (2015a). We take the yield on 15-year zero-coupon bonds at time t as the actuary's best estimate of the long-term returns expected on the bond portion of the pension fund from time t onwards. To estimate long-term returns on the equity portion of the fund we add a fixed risk premium to the 15-year zero-coupon bond yield.[6] From

[4]In this study, we generate 10,000 paths of the ESG, starting each series at their long-run levels.

[5]The ESG models the continuously compounded bond yields; therefore, the transformation to be applied to get the annual effective yield is $\tilde{r}_t^l = \exp\left(12\, r^l_{12(t-1)+1}\right) - 1$.

[6]This approach is frequently used in North America. The same approach is included in the Alberta target benefit plan regulations for calculating the "benchmark valuation rate." (Alberta Reg 154/2014, 2014)

the historical data, we find that the average risk premium over the 15-year bond yields has been 2.23% for the TSX/S&P Composite index. Therefore, the valuation rate based on EROA is $\tilde{r}_t^l + 0.5\,(0.0223)$, which is about 1.1% higher than the bond-based valuation rate.

Plan liabilities are determined using the Entry Age Normal (EAN) cost method as described in Aitken (1996). We chose this method because it is consistent with a level contribution rate applied throughout a member's working life, and therefore avoids the intergenerational cost subsidies that arise when benefit allocation-type cost methods (e.g., unit credit) are combined with a uniform contribution rate. Under the EAN method, the actuarial liability in respect of each member is equal to the present value of the member's projected benefits less the present value of their future normal costs. The normal cost is calculated as the level percent of salary that is required to finance the targeted benefit over the member's entire working life, from entry to retirement. The actuarial liability of the plan at time t , AL_t, is the sum of the individual liabilities. Note that, for purposes of the affordability test performed at time t, the members' projected future benefits are first determined using the target accrual rate b_{t-1} established at the previous valuation. Therefore, AL_t is based on b_{t-1}. The key output of the affordability test is the funded ratio $FR_t = F_t/AL_t$.

The initial target accrual rate, b_0, is consistent with the contribution rate c in the sense that the normal cost established under the EAN method in respect of this target is 10.6% when using a valuation rate equal to the long-term bond yield at plan inception. The benefit accrual rate may be adjusted at subsequent valuation dates according to pre-specified rules in the pension contract based on the result of the affordability test. In a pure CDC plan, there is only a single trigger point set to a funded ratio of 100%. Under this plan design, whenever the funded ratio determined by the affordability test is not equal to 100%, actions are taken immediately. Specifically, at each time $t > 0$, the new accrual rate b_t is determined by multiplying the accrual rate b_{t-1} established in the previous valuation by a factor α_t. For the pure CDC plan, we implemented the following adjustment factor:

$$\alpha_t = \frac{F_t + PVFNC_t}{PVTB_t}, \tag{5}$$

where $PVFNC_t$ is the present value of future normal costs and $PVTB_t$ is the present value of the targeted benefits in respect of both past and future service for active and retired members, based on the most recent target accrual rate, b_{t-1}.[7]

As noted earlier, using a single trigger leads to very volatile retirement benefits. Plan stakeholders may wish to add an explicit stabilization mechanism with a double trigger and a no-action range between the trigger points. We study two plan designs with such explicit stabilization mechanisms. In the first one, the no-action range is set symmetrically around a funded ratio of 100%. In the second, the no-action range is biased towards savings, with a midpoint higher than a funded ratio of 100%. In both cases, the target accrual rate is adjusted to bring the funded ratio back to the edge of the no-action range.

The starting asset value, F_0, at the inception of the CDC plan is the sum of members' individual DC account balances. We assume that each member's individual account balance is exactly equal to that member's EAN liability calculated using a valuation rate equal to the long-term bond yield applicable at time 0. This means that, when the first affordability test is performed at plan inception, the funded ratio will be exactly 100% as long as the valuation rate used in the affordability test is also equal to the long-term bond yield. If, on the other hand, the affordability test uses a valuation interest rate based on the expected return on assets, the funded ratio will differ from 100% and the target accrual rate may need to be adjusted right away.

4 Quantifying Risk Transactions Using Derivative Techniques

Our goal is to quantify the risk transactions within the stylized CDC plans described in Section 3. We follow the value-based ALM methodology of Hoevenaars and Ponds (2008) and first determine the cash flows in respect of each cohort under each simulated economic scenario. These cash flows fall into four categories: the individual DC plan account values transferred in at the inception of the CDC plan, the contributions made while the CDC plan is ongoing, the benefits received from the CDC plan while it is ongoing, and the residual assets distributed to members when the plan is terminated. The plan's residual assets (after the last set of contributions is received and the last set of benefit

[7] While this factor does not return the funded ratio after adjustment to exactly 100%, it has the advantage of avoiding numerical instability in the simulations. This adjustment factor can be interpreted by rewriting Equation (5) as $\alpha_t\, PVTB_t = F_t + PVFC_t - (PVFC_t - PVFNC_t)$, where the left-hand side is the present value of all benefits payable at or after time t, taking into account the adjustment at time t; $PVFC_t$ is the present value of future contributions, so $F_t + PVFC_t$ can be thought of as the present value of "available assets" at time t; and $PVFC_t - PVFNC_t$ is a countercyclical buffer that, when positive, reduces α_t by reducing the "available assets." When this "buffer" is negative, it results in a larger α_t.

payments is made at $t = 55$) are distributed in proportion to each member's remaining EAN liability. Let

- $F_{x,0}$ be the individual DC account balance of a member age x at time 0,
- $C_{x,t}$ be the contribution made at time t by a member age x,
- $B_{x,t}$ be the benefit payment received at time t by a member age x, and
- $RV_{x,55}$ be the residual asset value received at $t = 55$ by a member age x.

These cash flows can be constructed under both the real-world and the risk-neutral scenarios. The real-world cash flows reflect the contribution and benefit payments that the plan member would expect under possible economic conditions in the future. The cash flows calculated under the risk-neutral scenarios are not meaningful in themselves, however, they enable us to calculate the market-consistent value of the pension deal for each generation at time 0 and to investigate the gains and losses of each generation arising from their participation in the CDC plan. We let V_x be the value of the pension deal evaluated at $t = 0$ for a member age x at plan inception, Q_t^f be the 1-period risk-free discount factor applicable to period $(t-1, t]$ such that

$$Q_t^f = (1 + \tilde{r}_t^s)^{-1},$$

and $Q_0^f = 1$. The subscript x in V_x can range from -25 to 85, with negative values of x representing the cohorts that are not yet born at plan inception. The generational value of the pension deal under a CDC plan is the risk-neutral expected present value at $t = 0$ of all cash flows of a specific cohort,

$$V_x = \begin{cases} \mathbb{E}^{\mathbb{Q}}\left[-\sum_{i=30-x}^{65-x-1} C_{x+i,i} \left(\prod_{j=0}^{i} Q_j^f \right) + \sum_{i=65-x}^{55} B_{x+i,i} \left(\prod_{j=0}^{i} Q_j^f \right) + RV_{x+55,55} \left(\prod_{j=0}^{55} Q_j^f \right) \middle| \mathcal{F}_0 \right] & \text{if } x < 30 \\ \mathbb{E}^{\mathbb{Q}}\left[-F_{x,0} - \sum_{i=0}^{65-x-1} C_{x+i,i} \left(\prod_{j=0}^{i} Q_j^f \right) + \sum_{i=65-x}^{86-x} B_{x+i,i} \left(\prod_{j=0}^{i} Q_j^f \right) \middle| \mathcal{F}_0 \right] & \text{if } 30 \le x < 65 \\ \mathbb{E}^{\mathbb{Q}}\left[-F_{x,0} + \sum_{i=0}^{86-x} B_{x+i,i} \left(\prod_{j=0}^{i} Q_j^f \right) \middle| \mathcal{F}_0 \right] & \text{if } x \ge 65 \end{cases} . \quad (6)$$

Practically speaking, the expected values are calculated by averaging over all risk-neutral scenarios.

For members who have not yet joined the plan at inception, the value of the pension deal includes the value of the actual retirement benefits and the residual payout they will receive, less the value of the contributions they will make. For active members at plan inception, the value of the deal is equal to the value of the retirement benefits they will receive net of the value of their initial contribution and all their future contributions. The retirees start receiving their retirement benefits immediately after plan inception and their pension value is equal to the value of the benefits they have yet to receive minus the initial contribution they made.

We calculate these generational values under different CDC designs. Differences in the generational values of specific cohorts from design to design indicate a shift in that cohort's risk profile. In an individual DC plan without risk sharing all generational values are zero by definition, since each cohort exactly pays for its own benefits without any cost or risk subsidies from other cohorts. As a result, the values of V_x under a pure CDC plan can be interpreted as the market-consistent value of the implicit risk transactions in that plan.

While changes in the value of V_x from design to design point out shifts in the overall risk profile of each cohort, it is useful to explore these shifts at a more granular level, distinguishing between changes in upside versus downside risks, and between risks associated with the ongoing benefits, $B_{x,t}$, versus risks associated with the residual payments, $RV_{x,t}$. We therefore decompose the cash flows under each CDC design into the corresponding cash flows under an individual DC plan benchmark, and any additional (positive or negative) cash flows.

If the starting account balances ($F_{x,0}$) and subsequent contributions ($C_{x,t}$) are the same under the individual DC benchmark as they are under the CDC plan and both plans enjoy the same portfolio returns, then the additional cash flows must relate to differences in benefit payments ($B_{x,t}$) and residual values ($RV_{x,55}$). These differences occur because the CDC plan adjusts the benefits of each member by reference to the experience of the whole plan, whereas adjustments to the payouts from the benchmark plan only reflect the experience of a single individual. In effect, by joining the collective DC plan, the participants write put options to the plan, which waive part of the benefits they would get in the individual DC plan under certain economic conditions. The put options give a right to the plan and impose an obligation on the participants. Generally, an option writer will receive an option premium in return for bearing a risk; however, these put options are "embedded" in the pension contract and not explicitly sold as financial products, thus no option premium is paid to the participants as compensation for bearing the risk. Instead, participants are offered call options written on the fund surplus. The same applies to the residual asset under the CDC plan: it can be thought of as the residual asset under the individual DC plan plus a put option written by members and a call option received by members. This gives us four kinds of options: a benefit put option, a benefit call option, a residual put option and a residual call option.

We let $B'_{x,t}$ be the annual retirement benefit drawn from the individual DC account by a member age x at time t, where $x \geq 65$, and let $RV'_{x,55}$ be the remaining individual DC account value at $t = 55$ of a member age x. The *benefit put option* is a basket of options that the participants write to the pension plan. These options expire at consecutive benefit payment dates starting with $t = 1$. The payoff of a benefit put option expiring at time t written by a member age x at plan inception is defined as

$$BP_{x,t} = \max\left[B'_{x+t,t} - B_{x+t,t}, 0\right].$$

The option is in-the-money when $B'_{x,t} > B_{x,t}$.

The *benefit call option* is a basket of options that the plan writes to the participants, expiring at each consecutive benefit payment date. The payoff of a benefit call option expiring at time t written to a member age x at plan inception is defined as

$$BC_{x,t} = \max\left[B_{x+t,t} - B'_{x+t,t}, 0\right].$$

The option is in-the-money when $B_{x,t} > B'_{x,t}$. This option is the opposite of the benefit put option. Together, these two (baskets of) options capture the deviations of the retirement benefits under the CDC plan compared to the individual DC benchmark.

The *residual put option* is a basket of options that the participants write to the pension plan, expiring at $t = 55$. The payoff of the residual put option written by a member age x at plan inception is defined as

$$RP_x = \max\left[RV'_{x+55,55} - RV_{x+55,55}, 0\right].$$

At the termination of the pension plan, there are still generations who have not received their full benefits and generations who are not yet retired. These generations can make claims on the fund's residual wealth. The value of the assets that are left in the group fund depends on how much is spent on the benefits of the members who have already retired. For generations who are still active at the plan termination date, the residual asset they get is not necessarily equal to the accumulated value of the contributions they already made, which is the residual payment they would get under the benchmark individual DC plan. When this option is in-the-money, it implies that some of the benefits received by the already retired generations are at the expense of the future generations.

The *residual call option* is a basket of options that the pension plan writes to the participants, expiring at $t = 55$. The payoff of the residual call option written by the pension plan to a member age x at plan inception is defined as

$$RC_x = \max\left[RV_{x+55,55} - RV'_{x+55,55}, 0\right].$$

This option is the opposite of the residual put option. When it is in-the-money, it implies that additional funds are left over for future generations above and beyond what they would be entitled to under the individual DC benchmark.

Since the benefits and residual values of both the CDC plan and the individual DC plan are scenario dependent, these are Margrabe options whose value at time 0 can be calculated using the risk-neutral valuation technique:

$$VBP_{x,t} = \mathbb{E}^{\mathbb{Q}}\left[\max\left[B'_{x+t,t} - B_{x+t,t}, 0\right]\left(\prod_{j=0}^{t} Q_j^f\right)\middle|\mathcal{F}_0\right], \tag{7}$$

$$VBC_{x,t} = \mathbb{E}^{\mathbb{Q}}\left[\max\left[B_{x+t,t} - B'_{x+t,t}, 0\right]\left(\prod_{j=0}^{t} Q_j^f\right)\middle|\mathcal{F}_0\right], \tag{8}$$

$$VRP_x = \mathbb{E}^{\mathbb{Q}}\left[\max\left[RV'_{x+55,55} - RV_{x+55,55}, 0\right]\left(\prod_{j=0}^{55} Q_j^f\right)\middle|\mathcal{F}_0\right], \tag{9}$$

$$VRC_x = \mathbb{E}^{\mathbb{Q}}\left[\max\left[RV'_{x+55,55} - RV_{x+55,55}, 0\right]\left(\prod_{j=0}^{55} Q_j^f\right)\middle|\mathcal{F}_0\right]. \tag{10}$$

The value of the total benefit put option $VTBP_x$ for a member age x at plan inception is the aggregate value of all the benefit put options written by that member. It is equal to 0 for cohorts who are still active at plan termination (i.e., less than age 10 at plan inception) and it is equal to

$$VTBP_x = \sum_{t=\max[65-x,0]}^{55} VBP_{x,t} \tag{11}$$

for other generations.

Similarly, the value of the total benefit call option $VTBC_x$ for a member age x at plan inception is

$$VTBC_x = \begin{cases} 0 & \text{if } x < 10 \\ \sum_{t=\max[65-x,0]}^{55} VBC_x & \text{if } x \geq 10 \end{cases}. \tag{12}$$

For generations younger than 10 years old at plan inception, the value of the options written on the benefit payouts is 0, thus the gain or loss in value triggered by joining the CDC plan is reflected in the value of the residual options only. The total generational value of a member age x at plan inception is the sum of the value of these options:

$$V_x = VTBP_x + VTBC_x + VRP_x + VRC_x.$$

5 Evaluation of Collective Defined Contribution Schemes

We now turn to evaluating intergenerational transfers under specific CDC designs.

5.1 The Pure Collective Defined Contribution Plan

We first look at a pure CDC plan where the valuation rate used in the affordability test is closely tied to long-term bond yields. The initial target annual accrual rate b_0 is 1%. No immediate benefit adjustment is needed at time 0 because the funded ratio produced by the affordability test is exactly 100% (i.e., the target accrual rate is consistent with the combination of the starting asset value F_0, the contribution rate c, and the valuation basis).

Since half of the plan assets are invested in the stock market which earn equity risk premium in the long-run, the bond-based valuation rate tends to underestimate the investment return. The median spread between the simulated asset return and the simulated valuation rate hovers around 1% each year, indicating that the fund is more likely to experience investment gains than losses.

We make a few observations about the performance of the plan using some of the metrics considered by Sanders (2016) before moving to the value-based ALM framework. Pensioners do not want their retirement benefits to fluctuate every year. However, since the pure CDC plan has a single trig-

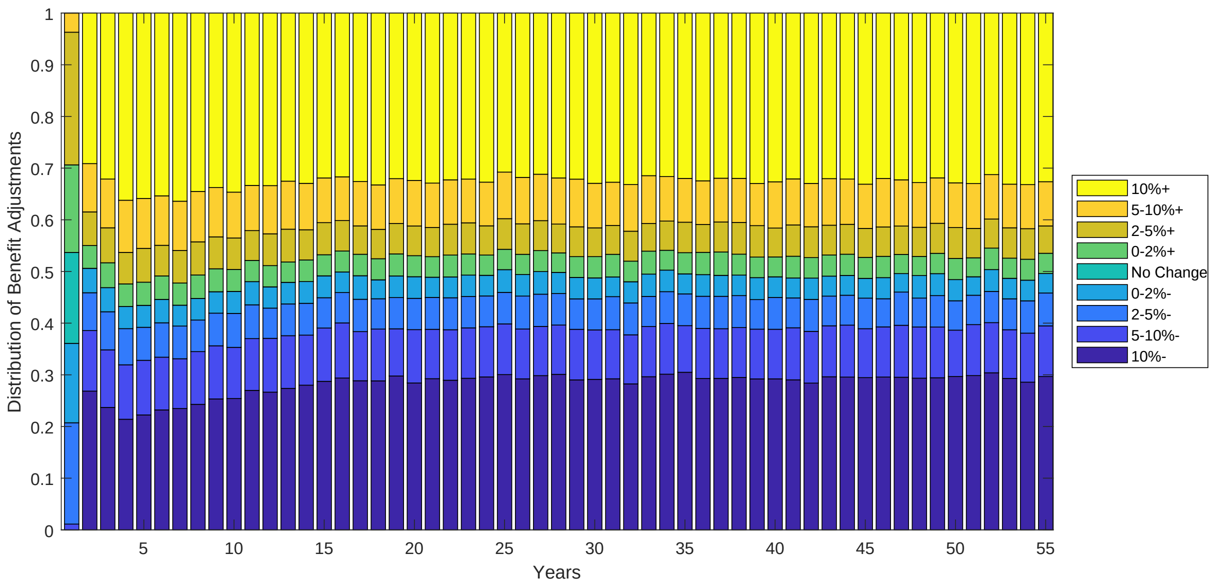

Figure 1: **Distribution of Benefit Adjustments by Valuation Year.**
The distribution of the size of the benefit adjustments is calculated at each valuation date. Each bar in the figure summarizes the distribution for a particular year.

ger, the benefits are not very stable: Figure 1 shows the distribution of benefit adjustments by size at each valuation date. Specifically, it reports that the probability of increasing or decreasing the benefit drastically (i.e., by more than 10%) in any given year is quite high—about 50%. During the first five years of operations, the likelihood of large negative spreads between the portfolio return and the valuation rate decreases. As a result, the downside benefit risk decreases during these years, which means that members are more likely to receive positive benefit adjustments.

This pattern can also be observed when computing the first year replacement ratio, which attempts to capture the cumulative impact of the adjustments over the lifetime of a member.[8] In fact, the distribution of the first year replacement ratio drifts upward in the early years of plan operations. Once the spread between the asset return and the valuation rate reaches its stationary level, the distribution of benefit adjustments also becomes stable with the probability of positive benefit adjustments hovering around 50%, and the median of the replacement ratio settling near 40%.

It is clear from these results that different cohorts of members face different distributions of risks and rewards. The value-based ALM framework allows us to investigate these differences in terms of

[8]The first year replacement ratio of a member who retires at time $t > 0$ is calculated as the benefit payment received in the year of retirement ($B_{65,t}$) divided by the member's final year salary.

market-consistent value. Figure 2 shows the generational pension value for different cohorts under the pure CDC plan. The values on the x-axis represent the ages at plan inception of the various cohorts who enter the plan. Since the value of the pension deal is defined as the market value of all payouts minus all contributions, a pension value of 0 represents a fair deal without any value transfers at plan inception. In Figure 2, the value of the pension deal for generations who are older than 25 at inception is negative, meaning that these cohorts are expected to be net losers in value terms if they enter this pure CDC plan. By contrast, the younger generations can expect to be net winners, on an *ex ante* basis. It is important to note that, without access to any external funding sources, the pension deal is a zero-sum game in market value terms, so one cohort's gain comes at the expense of other cohorts.

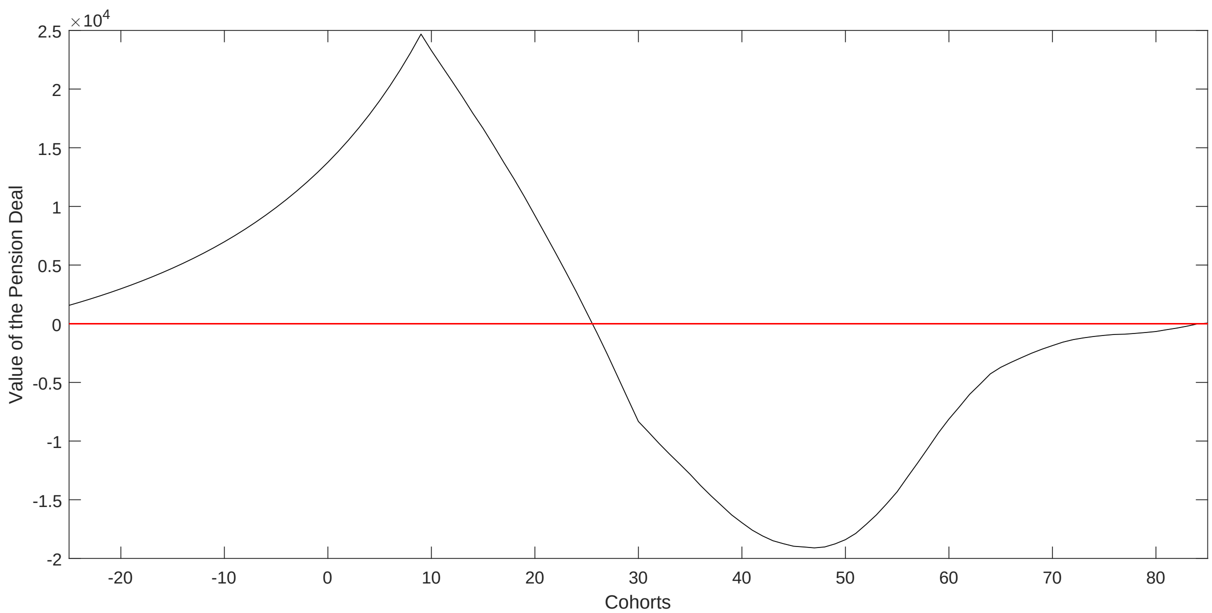

Figure 2: **Value of the CDC Plan for All Cohorts.**
The value of the pension deal is calculated using Equation (6) for all the generations who enter to the plan before termination. The x-axis in the figure shows all the cohorts of members, with negative numbers representing the cohorts who are not yet born at plan inception.

The resulting shift in value is significant: the cohort that is worst off (those around age 47 at plan inception) loses nearly \$20,000 in value (37% of their salary at inception), whereas those who are best off gain about \$25,000. Although this plan does not have explicit risk sharing, Figure 2 draws attention to the implicit risk-sharing structure of the pure CDC plan.

Figure 3 reports the value of the options embedded in the pure CDC plan for all cohorts. The benefit options plot (top panel) indicates that most cohorts have significant upside potential as measured by $VTBC$ (up to \$30,000 in market value terms), but this is more than offset by the downside potential, as measured by $VTBP$. This is due in large part to the use of a bond-based valuation rate in the affordability test, which results in low benefit payments to retirees. The problem is worst for the cohorts age 30 to 65 at plan inception, who receive all of their retirement benefits during the life of the plan. The assets that are held back on account of the conservative valuation rate are distributed at plan termination to the generations who are still in the plan at that time. This is reflected in the residual options plot (middle panel), where the VRC is much greater than VRP. The generation who is 9 years old at plan inception is the oldest active cohort at plan termination and has the highest actuarial liability at that point. As a result they receive the biggest portion of the residual assets and have the highest VRC.

If the plan kept operating past $t = 55$, the residual options would disappear. As a result, the younger generations who are expected to be net winners under the current setting would no longer expect to have such a large advantage.

The third panel in Figure 3 shows the aggregate value of the surplus options and deficit options from the first two plots. Since the value of the individual DC plan is equal to 0 for each generation, the value of the pension deal for each cohort is equal to the total value of the call options of that cohort less the total value of the put options for the same cohort.

To summarize, although there are no explicit intergenerational risk sharing mechanisms built into this plan design, value transfers between different cohorts still exist, pointing to significant implicit risk sharing. At any time during the simulation period, there are active members who are still earning benefits and paying contributions as well as retirees who are already receiving benefits from the plan. These two groups of participants are subject to different types of risks. Even within each group, the impact of changes in the inflation rate and the valuation rate is not the same for all cohorts of members because the amount of retirement savings and the duration of the pension liabilities are different. However, under the pure CDC plan, the target benefits with respect to both past and future service are adjusted simultaneously each year based on the funded position of the whole plan, without

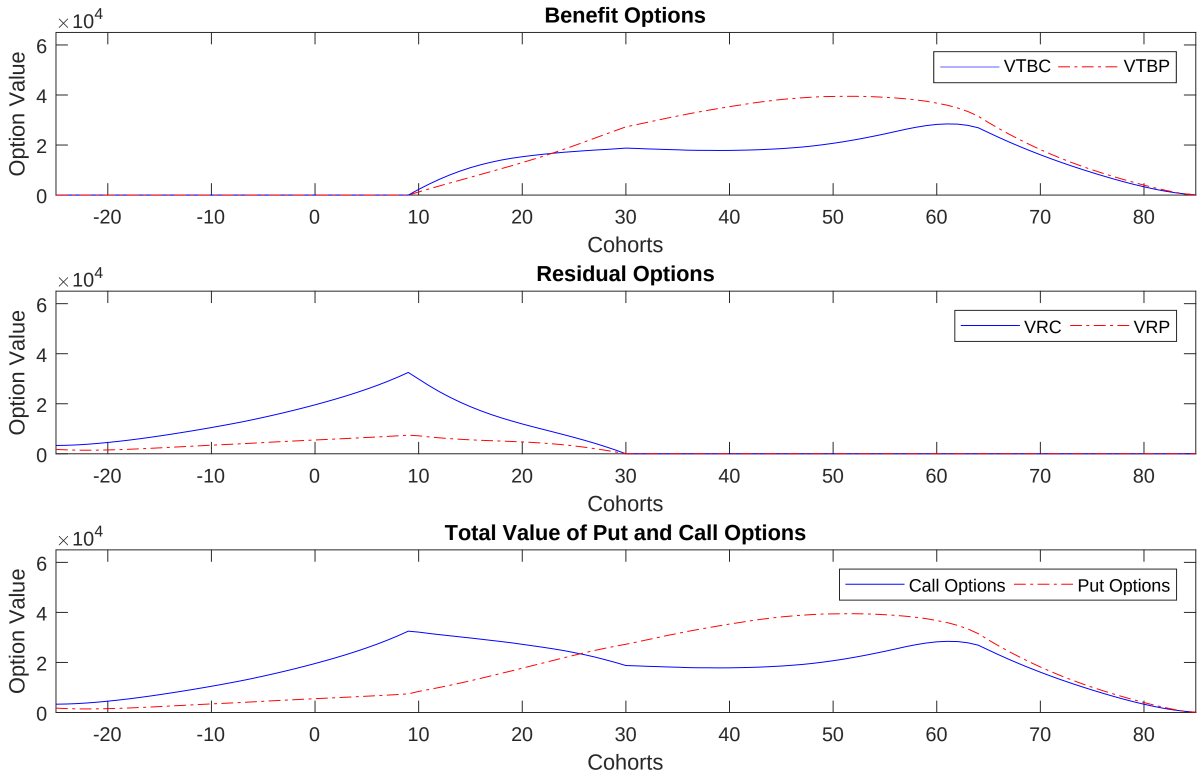

Figure 3: **Embedded Options in CDC Plan Relative to Individual DC Benchmark.**
The value of the embedded options are calculated for all the generations who enter to the plan before plan termination. The figure on the top shows the generational values of $VTBP$ and $VTBC$ calculated using Equations (11) and (12), respectively. The figure in the middle shows the generational values of VRP and VRC calculated using Equations (9) and (10), respectively. The figure at the bottom shows the total value of the call options and the put options for each generation. The x-axis in the figure shows all the cohorts of members, with negative numbers representing cohorts who are not yet born at plan inception.

distinguishing between different generations. Therefore, the inflation and valuation rate risks are shared implicitly among generations, giving rise to part of the value transfers.

Another significant cause of value transfers in this plan is the use of valuation rates based on bond yields. This conservative valuation assumption does not take into account the market risk premium expected to be earned on half of the plan assets invested in equities. The plan liabilities are constantly overestimated, which leads to lower funded ratios, and thus lower retirement benefits. As a result, more assets are retained in the plan. These extra assets are distributed at plan termination to the remaining generations in proportion to their final actuarial liabilities. Therefore, the generations who are still in the plan at termination are expected to be winners under this circumstance. The generation with the highest actuarial liabilities will receive the greatest benefit, which is the generation aged 9 at plan inception (age 64 at plan termination).

Since there are no external guarantees from the plan sponsors, the pension deal is a zero-sum game across all generations from a market value perspective: the surplus of one group of members comes at the expense of another group, on an *ex ante* basis. In our case, the generations who are younger than 25 years old at plan inception are expected to be net winners. Most of their surplus results from the fund distribution at plan termination, which is at the expense of smaller retirement benefits paid to older generations.

5.2 On the Impact of Alternative Affordability Tests

The relationship between intergenerational equity and key design choices in CDC plans is of great concern to actuaries and other stakeholders. For instance, the pure CDC plan analyzed in Section 5.1 uses a bond-based valuation rate in the affordability test. Yet, one area of particular interest and active debate is the choice of valuation rate and its impact on intergenerational equity. Using a valuation rate equal to the EROA (i.e., anticipating the equity risk premium) is permitted under Canadian and US pension regulations. Recently, Sanders (2016) investigated the changes in the dynamics of the retirement benefits in simple CDC plans after applying EROA as the valuation rate. Ma (2018) also explored the issue of selecting discount rates for assessing the funded status of the same types of CDC plans and concluded that a valuation rate based on EROA is more equitable and serves the best interests of all members. To test the conclusion of Ma (2018) under our framework, we change the valuation rate to the EROA while leaving all other design factors the same.

Effectively, this change results in an increase of about 1.1% in the valuation rate, reducing the median spread between the simulated asset returns and the simulated valuation rate to almost zero. Since the initial fund value is unchanged but the valuation rate is higher, the starting funded ratio will be higher than 100%. This results in a large positive benefit adjustment at the first valuation date, which brings the median of the funded ratio back to 100% quickly. A significant proportion of the surplus that becomes visible at the first valuation—to which all generations contributed at inception—is paid out directly to the members who are already retired. In addition, because the new valuation rate tends to slightly overestimate the investment return in the first few years, the retirees receive an even larger portion of the surplus. After the first few years, positive and negative benefit adjustments are equally likely, as shown in Figure 4.

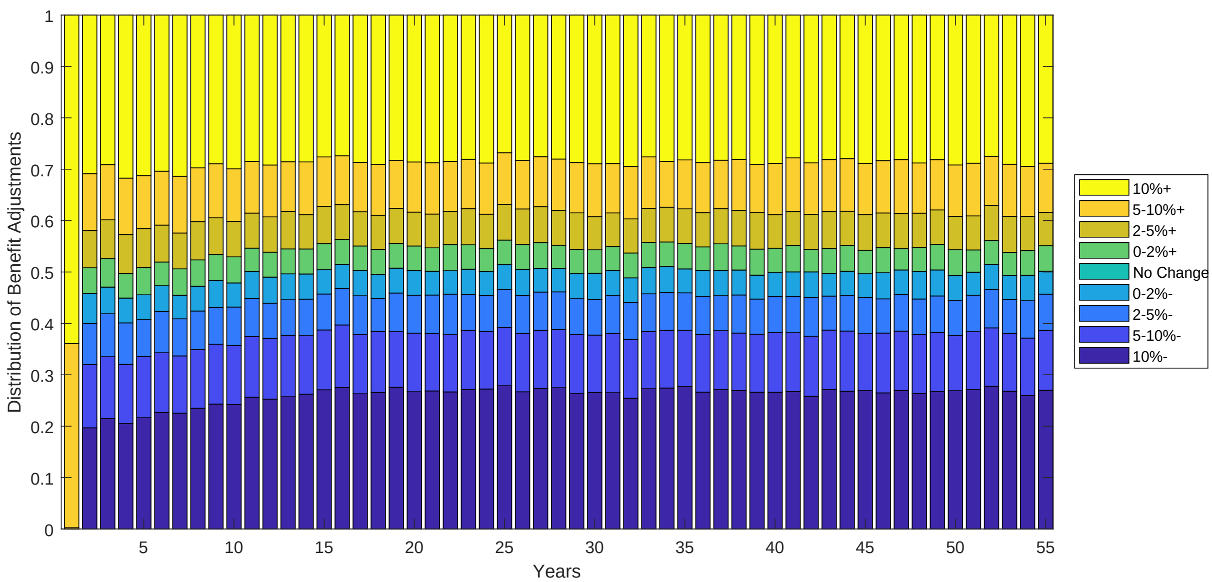

Figure 4: **Distribution of Benefit Adjustments by Valuation Year.**
See the description of Figure 1 for more details.

Figure 5 reports the value of the pension deal under this modified plan. In contrast to the negative pension value under the plan in Section 5.1, the members who are already retired at plan inception have a positive pension value when using EROA as the valuation rate. This is because their retirement benefits consume a portion the surplus from the initial contributions of both the active members and the retirees. Those who are still active at plan inception (i.e., retire after the initial surplus is spent) are worse off because the surplus included in their initial contributions is paid out to the retirees and the probability of positive benefit adjustments is lower than under the CDC plan in Section 5.1.

Figure 6 isolates the value transfers due to changing the valuation basis. Overall, the values are shifted from the younger generations to those over age 45 at inception. Most of the benefit gains for older generations come from the loss of residual value for younger generations. This is consistent with the fact that the EROA as valuation rate would decrease the plan liability and increase the funded ratio early on, accelerate the payment of larger benefits, and thus reduce the residual assets of the plan.

In our findings, the difference between the retirement benefit dynamics in the EROA case and

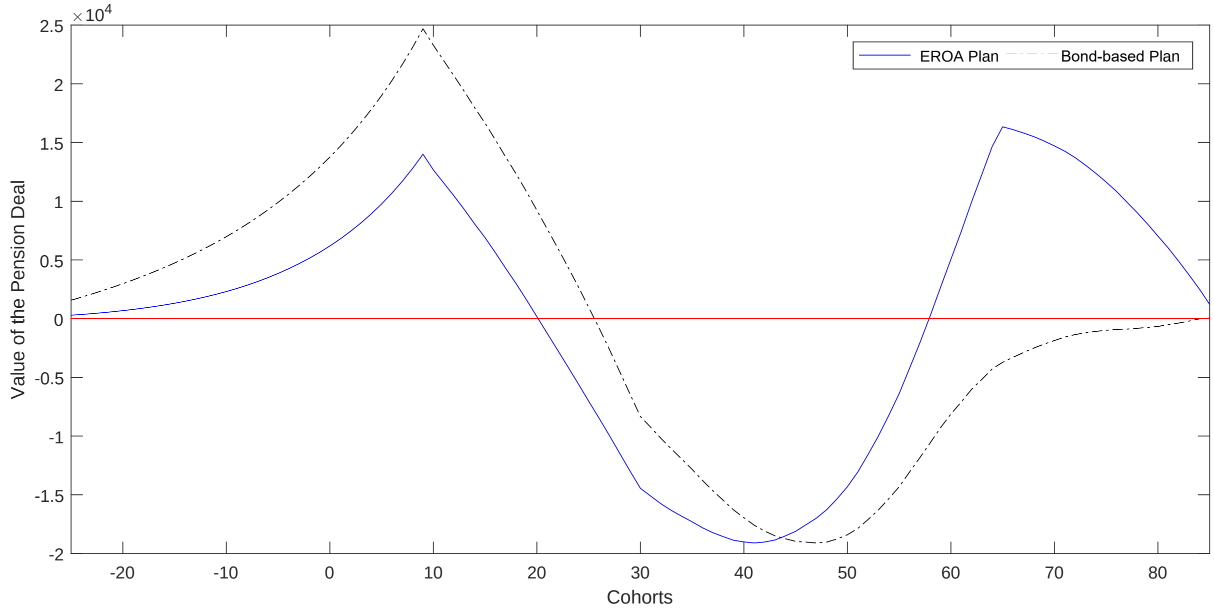

Figure 5: **Value of the EROA Plan and Bond-based CDC Plan for All Cohorts.**
See the description of Figure 2 for more details. The solid line represents the generational values under the EROA plan and the dashed line represents the generational values under the collective DC plan with bond-based valuation rate.

the bond-based case are consistent with Ma (2018). However, under our specific model and plan provisions, it is unclear whether using the EROA as valuation rate can be characterized as "more equitable". In our model, using the EROA simply shifts the expected residual benefits of future generations to retirees at plan inception. The generations who are 30 to 50 years old at plan inception still have a sizeable expected net loss. To be truly "more equitable," this loss would need to be addressed. Simply switching the valuation basis does not achieve this.

6 Introducing Explicit Stabilization Mechanisms

The retirement benefits in the pure CDC plans investigated in the previous sections are adjusted at every valuation date and can fluctuate a lot. Plan designs with two triggers are adopted commonly in practice to add stability to CDC plans. There may be a sense among stakeholders that the benefit stability achieved by introducing the no-action range comes at no additional cost. In reality, this stability is made possible by requiring potentially larger intergenerational subsidies, which may have hidden costs. In this section, we investigate these subsidies, focusing on fairness *ex ante*.

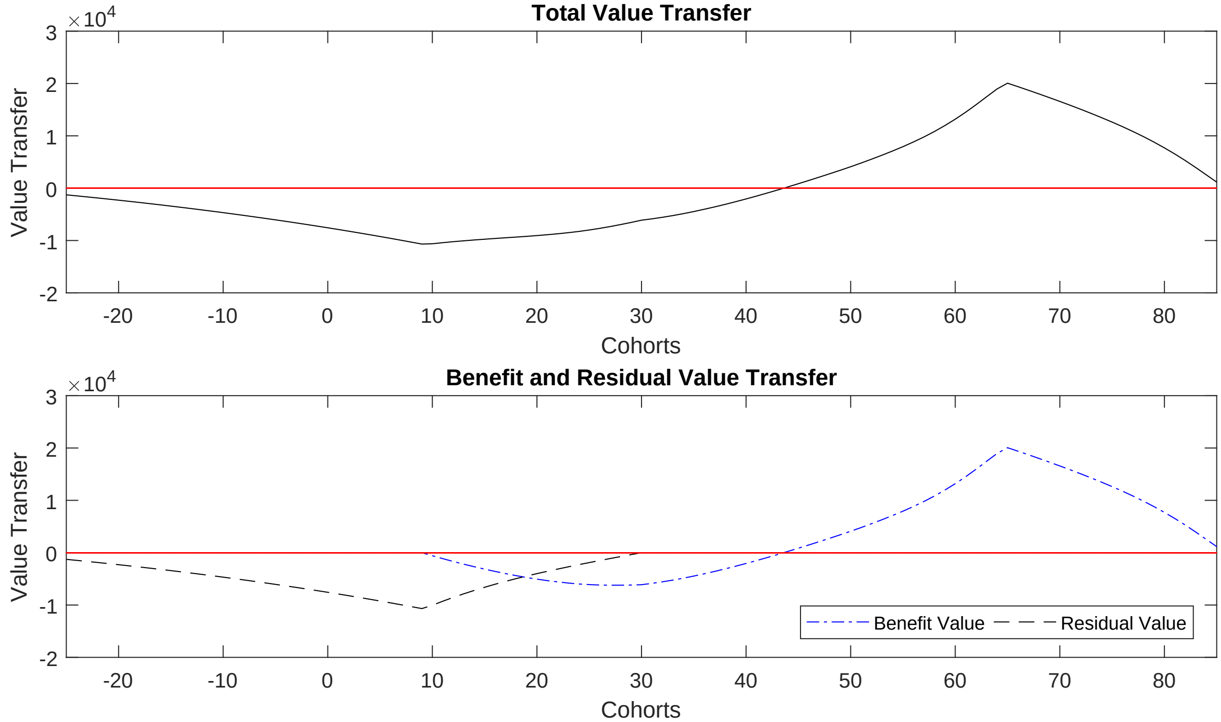

Figure 6: **Value Transfers Caused by Changing Valuation Basis.**
The figure on the top shows the total value transfers due to changing the valuation basis from a bond-based valuation rate to a EROA. The figure at the bottom separates the value transfers due to the benefit options versus the residual options.

6.1 Triggers Centred at a Funded Ratio of 100%

We first add a no-action range for funded ratios between 80% and 120% to the modified plan design from Section 5.2 (i.e., the pure CDC plan using the EROA as the valuation rate). Since this range is centred around a funded ratio of 100%, this plan is called a symmetric corridor plan. Under this plan design, benefits are unchanged until the funded ratio moves outside the no-action range, at which point adjustments are made to bring the funded ratio closer to the edge of the range. For example, if the funded ratio is 140% at the valuation date, then an adjustment factor is applied to the target accrual rate to bring the funded ratio close to 120%.

For this plan, as for the corresponding plan in Section 5.2, the median of the funded ratio starts at a value higher than 100% due to the lack of consistency between the initial target accrual rate of 1% and the affordability test. However, unlike the previous case where the funded ratio is brought back towards 100% immediately, the median funded ratio in the symmetric corridor plan drifts towards the midpoint of the no-action range slowly over time. After about 15 years, the distribution of the funded ratio becomes stationary. In this symmetric corridor plan, the initial surplus is retained in the plan instead of being paid out through large benefit increases in the early years. This surplus is then allowed to be eroded by negative plan experience (or boosted by positive plan experience, as the case may be).

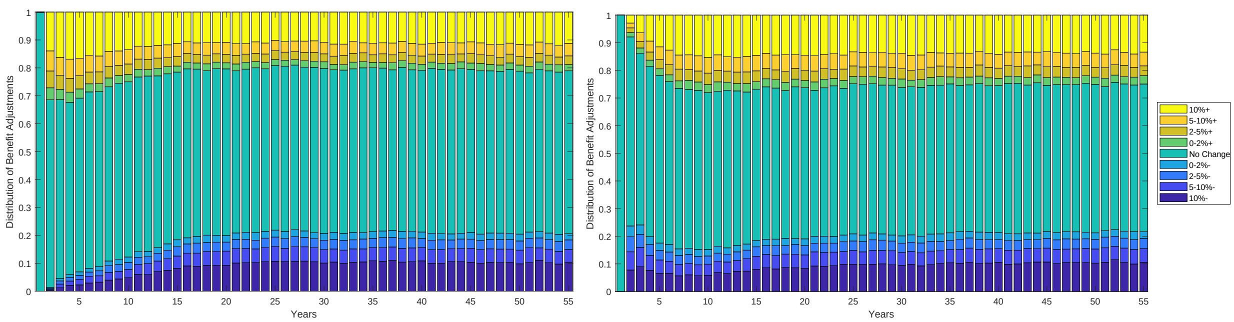

Figure 7: **Distribution of Benefit Adjustments by Valuation Year for Corridor Plans.**
See the description of Figure 1 for more details. The left panel reports the distributions of benefit adjustments for the symmetric corridor case. The right panel shows the distributions for the saving one.

The effectiveness of the no-action range in reducing the frequency and magnitude of benefit changes is demonstrated in the left panel of Figure 7: the benefits remain the same from one year to another with more than 50% probability. The probability of a large benefit adjustment (i.e., an increase or decrease greater than 10%) occurring in any given year is reduced to 20% from over 50%. In addition, negative adjustments are less likely during the first 15 years because the plan starts with a funded ratio that is above the midpoint of the no-action range.

The left panel of Figure 8 shows the value of the pension deal under the symmetric corridor plan design and the left panels of Figure 9 show the value transfers caused by adding a no-action range to the EROA plan, indicating that the no-action range centred at a funded ratio of 100% adds significant stability to the plan without creating hardly any additional value transfers.[9]

[9]These results depend on a stable demographic profile and the assumption that the simulation of the economic variables is started from the long-term equilibrium level. Given the current low interest rate environment, it may not be reasonable to assume that the economic and demographic conditions remain stationary over the next 50 years.

6.2 Triggers Centred at a Funded Ratio of 120%

Canadian regulations for CDC plans allow the use of no-action ranges, although not ones centred at 100%. Generally, the regulations require the lower end of the no-action range to start at 100%. In this section, we implement a plan with lower and upper trigger points set at funded ratios of 100% and 140%, respectively. Since the no-action range is deliberately biased to saving, we call this a saving corridor plan.

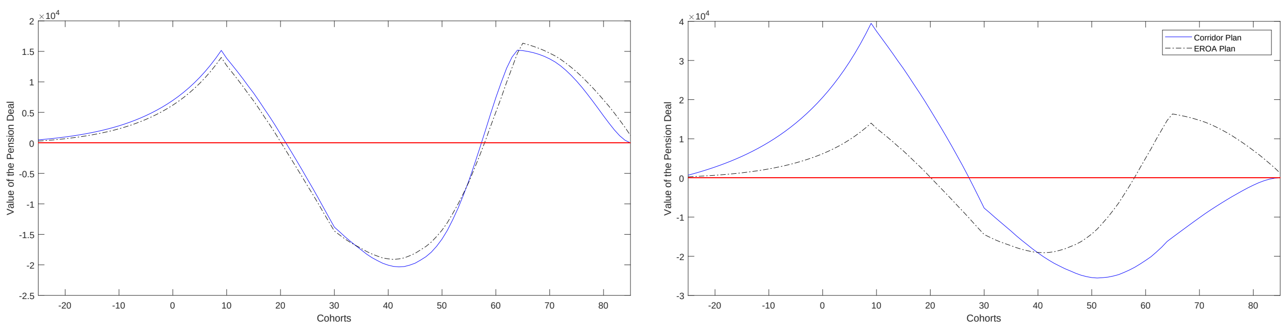

Figure 8: **Value of the Corridor Plans and the EROA Plan for All Cohorts.**
See the description of Figure 2. The left panel reports the value for the symmetric corridor case. The right panel shows the value for the saving one. The solid line represents the generational values under the symmetric (left) or saving (right) corridor plan and the dashed line represents the generational values under the EROA plan.

Regarding the dynamics of the funded ratio, the first 10 years form a transitional period during which time the distribution of the funded ratio widens and drifts upwards. The median funded ratio moves from a value of 110% at plan inception to the centre of the no-action range (i.e., 120%): instead of slowly eroding the initial surplus as the symmetric corridor plan does, the savings corridor plan is expected to build up additional buffers. As a result, the benefits are less likely to be adjusted upward and more likely to be adjusted downward than under the symmetric corridor plan, at least in the early years. This is confirmed by the right panel of Figure 7. After the first 10 years, the probability of the benefits remaining unchanged in a given valuation is around 50%, which is the same as under the symmetric corridor plan. This is because the no-action ranges have the same "size" or "span" in terms of the funded ratio. However, the probability of positive benefit adjustments is slightly higher and the probability of negative benefit adjustments is slightly lower under the saving corridor plan than under the symmetric corridor plan because positive returns on the surplus tend to spin off additional gains.

The right panel of Figure 8 illustrates the value of the pension deal under the saving corridor plan design. Under this plan design, all generations younger than age 30 have positive plan values because the plan is expected to build up a significant buffer and this buffer is distributed to the generations who remain in the plan when it is terminated. This is similar to distributing the savings accumulated under the pure CDC plan with a bond-based (conservative) valuation rate, except that the median size of the savings here is larger than under the plan in Section 5.1. The generations who are not entitled to the residual surplus at plan termination tend to be worse off under this design because a significant part of their benefit payments tend to be held back to build the buffer.

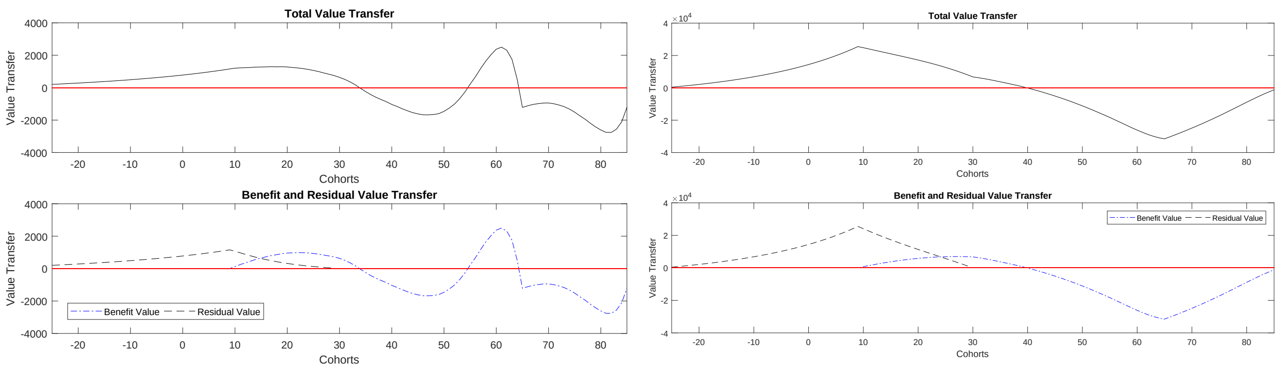

Figure 9: **Value Transfers Created by Adding a No-action Range.**
See the description of Figure 6 for more details. The left panels report the value transfers for the symmetric corridor case. The right panels show the value transfers for the saving one.

The right panels of Figure 9 report the value transfers resulting from adding a no-action range spanning funded ratios of 100% to 140% to the EROA plan from Section 5.2. There is a clear shift in value from older cohorts to younger cohorts. The oldest active cohort at plan termination, who is 9 years old at plan inception, is expected to have the largest gain under this plan design equal to nearly $30,000 in terms of market-consistent value. This generation has the highest actuarial liability at plan termination, thus is expected to share the largest portion of the extra savings in the residual assets. The generation retiring right after plan inception has the largest loss (around $30,000), equivalent to nearly 60% of their final year earnings. This is because a part of the surplus which would otherwise fund benefit increases is retained in the plan for building up a larger buffer. From these results, we can conclude that the stability under the saving corridor plan does come at someone's expense: in this case, at the expense of older generations.

7 Concluding Remarks

In this research, we study intergenerational equity under different CDC designs using the value-based ALM framework. We investigate the wealth redistribution effect of changing key plan design elements, such as changing valuation assumptions and adding explicit benefit stabilization mechanisms. The operation of four CDC plans are studied using Monte Carlo simulation over a 55-year time horizon. The cash flows of different generations are recorded separately using the generational accounting technique. The contingent retirement benefits paid under each plan are formulated as combinations of several embedded options written between the pension fund and the plan members. The value of the pension deal for each generation is evaluated using the risk-neutral asset pricing technique and compared across different plan designs.

To generate the economic scenarios on which the future evolution of pension funds depends, a real-world VAR-GARCH model is constructed. In addition, a corresponding risk-neutral model is derived and calibrated to help with the pricing of the embedded options within the pension contract.

The first design studied in this research is a pure CDC plan with a bond-based valuation rate. In the second plan, the bond-based valuation rate is replaced with the actuary's estimate of the expected return on assets. Lastly, we study two plans with benefit stabilizing designs, including the so-called symmetric corridor plan and the saving corridor plan.

The main contributions of this paper are threefold. First and foremost, we show that value transfers arise by simply joining a pure CDC scheme, i.e., without the inclusion of any explicit benefit smoothing mechanisms. Inflation and valuation rate risk are shared implicitly between generations, and explain, to some extent, these value transfers. Second, we show that changing the valuation rate to the EROA transfers part of the surplus of younger generations to the generations who already retired—or are close to retirement—but is not able to make the deal completely fair to all generations from an *ex ante* perspective. Third, adding a symmetric corridor is shown to reduce the volatility of retirement benefits without triggering significant additional value transfers, and plans with designs that are biased to saving tend to shift value from older generations to younger generations.

Another important finding of our study is that the generations aged between 30 to 55 at plan inception are expected to lose value under all plan designs investigated. This conclusion is somewhat troublesome from a policy perspective. The challenge is, therefore, to find a design that protects the pension value of these "middle cohorts," instead of transferring part of it to either younger or older members. This may involve ring-fencing part of the initial assets brought into the plan for the exclusive benefit of certain cohorts—so benefits are not paid out too fast—or releasing unused surplus upon the death of each member instead of retaining it in the plan—so benefits are not paid out too slowly. Ring-fencing is an interesting feature to include in pension designs and is left for future research.

References

Aitken, W. H. 1996. *A problem-solving approach to pension funding and valuation*. ACTEX Publications.

Alberta Reg 154/2014. 2014. Employment pension plans regulations.

Blommestein, H. J., N. Kortleve, and J. Yermo. 2009. Evaluating the design of private pension plans: Costs and benefits of risk-sharing. *Working paper*.

Bollerslev, T. 1986. Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics* 31:307–327.

Bovenberg, L., R. Koijen, T. Nijman, and C. Teulings. 2007. Saving and investing over the life cycle and the role of collective pension funds. *De Economist* 155:347–415.

CIA. 2015a. Determination of best estimate discount rates for going concern funding valuations. Tech. rep., Canadian Institute of Actuaries (CIA).

CIA. 2015b. Report of the task force on target benefit plans. Tech. rep., Canadian Institute of Actuaries (CIA).

Cui, J., F. De Jong, and E. Ponds. 2011. Intergenerational risk sharing within funded pension schemes. *Journal of Pension Economics and Finance* 10:1–29.

Engle, R. F., and C. W. Granger. 1987. Co-integration and error correction: Representation, estimation, and testing. *Econometrica* 55:251–276.

Gollier, C. 2008. Intergenerational risk-sharing and risk-taking of a pension fund. *Journal of Public Economics* 92:1463–1485.

Heston, S. L., and S. Nandi. 2000. A closed-form GARCH option valuation model. *Review of Financial Studies* 13:585–625.

Hoevenaars, R. P., and E. H. Ponds. 2008. Valuation of intergenerational transfers in funded collective pension schemes. *Insurance: Mathematics and Economics* 42:578–593.

Kocken, T. 2006. *Curious contracts*. Ph.D. thesis, Vrije University Amsterdam.

Lütkepohl, H. 2005. *New introduction to multiple time series analysis*. Springer Science & Business Media.

Ma, C.-M. G. 2018. Selecting discount rates for assessing funded status of target benefit plans. Tech. rep. Canadian Institute of Actuaries.

Sanders, B. 2016. Analysis of target benefit plan design options. Tech. rep., Society of Actuaries.

A More on the Model

The VAR-GARCH model is transformed into a corresponding risk-neutral model for pricing purposes. Specifically, in our study, we define our change of measure via a pricing kernel which is a function of the risk factors in the model:

$$\frac{M_{n+1}}{e^{-r_n^s}} = \frac{\frac{d\mathbb{Q}}{d\mathbb{P}}\Big|_{\mathcal{F}_{n+1}}}{\frac{d\mathbb{Q}}{d\mathbb{P}}\Big|_{\mathcal{F}_n}} = \frac{e^{-\boldsymbol{\Lambda}_{n+1}^\top \mathbf{y}_{n+1}}}{\mathbb{E}^{\mathbb{P}}\left[e^{-\boldsymbol{\Lambda}_{n+1}^\top \mathbf{y}_{n+1}} \,\middle|\, \mathcal{F}_n\right]},$$

where $\mathbf{y}_{n+1} = \boldsymbol{\Sigma}_{n+1}^{\frac{1}{2}} \boldsymbol{\varepsilon}_{n+1}$ is the time-n five-dimensional noise term under the real-world model, and $\boldsymbol{\Lambda}_{n+1}$ is a five-dimensional vector representing the time-varying market prices of risk.

As discussed in Section 2, $y_{i,n+1}$ follows a normal distributions with mean of 0 and variance of $\sigma_{i,n+1}^2$ under the physical measure. We can thus derive the moment generating function of $y_{i,n+1}$ under the measure $\mathbb{Q}$ as

$$\mathbb{E}^{\mathbb{Q}}\left[e^{uy_{i,n+1}} \,|\, \mathcal{F}_n\right] = \mathbb{E}^{\mathbb{P}}\left[\frac{e^{-\boldsymbol{\Lambda}_{n+1}^\top \mathbf{y}_{n+1}} \; e^{uy_{i,n+1}}}{\mathbb{E}^{\mathbb{P}}\left[e^{-\boldsymbol{\Lambda}_{n+1}^\top \mathbf{y}_{n+1}} \middle| \mathcal{F}_n\right]} \,\middle|\, \mathcal{F}_n\right] = \exp\left(\frac{u^2}{2}\sigma_{i,n+1}^2 - u\Lambda_{i,n+1}\sigma_{i,n+1}^2\right).$$

It is easy to show that the innovations are in fact normally distributed with a mean of $-\Lambda_{i,n+1}\sigma_{i,n+1}^2$ and a variance of $\sigma_{i,n+1}^2$, where $\Lambda_{i,n+1}$ is the i^{th} element of $\boldsymbol{\Lambda}_{t+1}$. Therefore, the elements of the $\mathbb{Q}$-measure noise terms $\mathbf{y}_{n+1}^*$ could be defined as a function of the $\mathbb{P}$-noise terms, i.e.,

$$y_{i,n+1}^* = y_{i,n+1} + \Lambda_{i,n+1}\sigma_{i,n+1}^2 \quad \text{for} \quad i = 1, ..., 5. \tag{13}$$

The elements of $\mathbf{y}_{n+1}^*$ are independently distributed, each with parameters $\mathcal{N}(0, \sigma_{i,n+1}^2)$. Therefore, vector $\mathbf{y}_{n+1}^*$ can also be written in the following form:

$$\mathbf{y}_{n+1}^* = \mathbf{y}_{n+1} + \boldsymbol{\Sigma}_{n+1}\boldsymbol{\Lambda}_{n+1}. \tag{14}$$

We then assume that the risk premium is affine in the state variables, i.e.,

$$\boldsymbol{\Sigma}_{n+1}\boldsymbol{\Lambda}_{n+1} = \boldsymbol{\lambda}_0 + \boldsymbol{\Sigma}_{n+1}\boldsymbol{\gamma}^* + \boldsymbol{\lambda}_1\mathbf{z}_n,$$

where $\boldsymbol{\lambda}_0$ is a five-dimensional vector that contains constant parameters while $\boldsymbol{\lambda}_1$ is a 5×5 matrix that accounts for the time-varying part in the risk premium. The parameter $\boldsymbol{\gamma}^*$ is given by

$$\boldsymbol{\gamma}^* = \begin{bmatrix} 0 & 0 & 0 & \gamma_4 & 0 \end{bmatrix}^\top,$$

and γ_4 addresses the equity risk premium of the excess stock return captured by the GARCH-in-mean model.

From Equations (13) and (14), we can establish the relationship between the noise terms under the real-world measure and those under the risk-neutral measure as

$$\mathbf{y}_{n+1} = \mathbf{y}^*_{n+1} - (\boldsymbol{\lambda}_0 + \boldsymbol{\Sigma}_{n+1}\boldsymbol{\gamma}^* + \boldsymbol{\lambda}_1\mathbf{z}_n).$$

Substituting these noise terms in Equations (2) and (4), we can get the corresponding risk-neutral dynamics of our VAR-GARCH model:

$$\mathbf{z}^*_{n+1} = (\boldsymbol{\nu} - \boldsymbol{\lambda}_0) + \boldsymbol{\Sigma}_{n+1}(\boldsymbol{\gamma} - \boldsymbol{\gamma}^*) + (\boldsymbol{\beta} - \boldsymbol{\lambda}_1)\mathbf{z}^*_n + \mathbf{y}^*_{n+1},$$

where

$$\mathbf{y}^*_{n+1} = \boldsymbol{\Sigma}^{\frac{1}{2}}_{n+1}\boldsymbol{\varepsilon}_{n+1}$$

and

$$\boldsymbol{\Sigma}_{n+1} = \boldsymbol{\omega} + \mathbf{A}\,\mathrm{diag}\left(\mathbf{y}^*_n - (\boldsymbol{\lambda}_0 + \boldsymbol{\Sigma}_n\boldsymbol{\gamma}^* + \boldsymbol{\lambda}_1\mathbf{z}_{n-1})\right)^2 + \mathbf{B}\,\boldsymbol{\Sigma}_n.$$

To achieve consistency between the real-world model and the risk-neutral model, the discounted value of the equity index should behave as a martingale.[10] Thus, the following conditions should be satisfied:

$$\lambda_{0,4} = \nu_4 \quad \text{and} \quad \boldsymbol{\lambda}_{1,4} = \boldsymbol{\beta}_4,$$

where $\lambda_{0,4}$ is the 4th element of vector $\boldsymbol{\lambda}_0$ and $\boldsymbol{\lambda}_{1,4}$ is the 4th row of matrix $\boldsymbol{\lambda}_1$. With these two conditions,

[10] A model is arbitrage-free if and only if there exists a change of measure under which the discounted processes for all traded assets are martingales—in our case, the equity index.

we can derive the risk-neutral dynamics of the excess stock return process as

$$z^*_{4,n+1} = -\frac{1}{2}\sigma^2_{4,n+1} + y^*_{4,n+1}.$$

The discounted value of the stock price S_{n+1} is, therefore, behaving as a martingale under the $\mathbb{Q}$ measure:

$$\mathbb{E}^{\mathbb{Q}}\left[S_{n+1}e^{-r^s_{n+1}}\,\middle|\,\mathcal{F}_n\right] = \mathbb{E}^{\mathbb{Q}}\left[S_n e^{z_{4,n+1}}\,|\,\mathcal{F}_n\right] = S_n\,\mathbb{E}^{\mathbb{Q}}\left[e^{-\frac{1}{2}\sigma^2_{4,n+1}+y^*_{4,n+1}}\,\middle|\,\mathcal{F}_n\right] = S_n,$$

because $y^*_{4,n+1} \sim \mathcal{N}(0, \sigma^2_{4,n+1})$.

B More on the Data

The yields of zero-coupon bonds with maturities ranging from three months to 30 years are available on the Bank of Canada's website. We use the yield on three-month Canadian Treasury bills and the 15-year bond yield observed on the first trading day of month n as proxies for the short-term bond yield, r^s_n, and long-term bond yield, r^l_n, respectively. The yields on bonds with one-year, three-year, five-year, seven-year, ten-year, 12-year, 14-year and 15-year maturities are used in the estimation of the risk premium parameters $\boldsymbol{\lambda}_0$ and $\boldsymbol{\lambda}_1$.

The values of the consumer price index (i.e., CPI_n) at the end of month n (with base year 2002) are taken from Statistics Canada's website. The month-n inflation rate i_n is calculated as

$$i_n = \log\left(\frac{\text{CPI}_n}{\text{CPI}_{n-1}}\right).$$

The closing price of the S&P/TSX Composite index (i.e., S_n) at the end of month n and its annualized dividend yield (i.e., D_n) are obtained from Statistics Canada's website. This index includes the stock prices of the largest companies traded on the Toronto Stock Exchange. The continuously compounded total stock return, s_n, and the dividend yield, d_n, can be calculated as

$$s_n = \log\left(\frac{S_n}{S_{n-1}} + \frac{D_n}{12}\right), \qquad d_n = \log\left(1 + \frac{D_n}{12}\right),$$

and the excess stock return r_n in excess of the risk free rate, proxied by the three-month Canadian Treasury bills yield, is defined as

$$r_n = \log\left(1 + \frac{D_n}{12}\right) - r^s_n.$$

C Estimation Results

Estimation Under the Physical Measure

First of all, the mean reversion level $\boldsymbol{\mu}$ of the VAR-GARCH process is set to the historical mean of the economic variables, except for the excess stock return. As explained in Section 2, the market risk premium for equities is driven by the time-varying uncertainty level. Therefore, we assume that the mean-reversion level of the excess stock return is proportional to its long-term variance. The other parameters in Equations (1) and (3) are estimated using maximum likelihood estimation.

The sequence $\{\mathbf{z}_n : n = 1, 2, .., N\}$ contains the N observed values of the five economic variables. Let $\boldsymbol{\Theta}$ be the vector that contains all the model parameters. Since $\boldsymbol{\varepsilon}_n$ is a white noise vector such that $\boldsymbol{\varepsilon}_n \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_5)$, the conditional distribution of $\{\mathbf{z}_n - \boldsymbol{\mu} : t = 2, ..., N\}$ given $\mathcal{F}_{n-1}$ is

$$\mathbf{z}_n - \boldsymbol{\mu} \mid \mathcal{F}_{n-1} \sim \mathcal{N}\left(\boldsymbol{\mu}_n(\boldsymbol{\Theta}), \boldsymbol{\Sigma}_n(\boldsymbol{\Theta})\right),$$

where

$$\boldsymbol{\mu}_n(\boldsymbol{\Theta}) = \mathbb{E}^{\mathbb{P}}\left[\mathbf{z}_n - \boldsymbol{\mu} \mid \mathcal{F}_{n-1}\right] = \boldsymbol{\Sigma}_n(\boldsymbol{\Theta})\boldsymbol{\gamma} + \boldsymbol{\beta}(\mathbf{z}_{n-1} - \boldsymbol{\mu})$$

$$\text{and } \boldsymbol{\Sigma}_n(\boldsymbol{\Theta}) = \text{Var}^{\mathbb{P}}\left[\mathbf{z}_n - \boldsymbol{\mu} \mid \mathcal{F}_{n-1}\right] = \boldsymbol{\omega} + \mathbf{A}\,\text{diag}(\mathbf{y}_{n-1}^2) + \mathbf{B}\,\boldsymbol{\Sigma}_{n-1}.$$

Let us assume that $\mathbf{y}_1 = \mathbf{z}_1 - \boldsymbol{\mu}$. The following residuals are thus given by

$$\mathbf{y}_n = \mathbf{z}_n - \boldsymbol{\mu} - \boldsymbol{\mu}_n(\boldsymbol{\Theta}), \quad \text{for } n = 2, ..., N.$$

Under the normality assumption, the conditional log-likelihood for the time t observations is given by

$$\mathcal{L}_n(\boldsymbol{\Theta}; \mathbf{z}_n) = -\frac{5}{2}\log(2\pi) - \frac{1}{2}\log|\boldsymbol{\Sigma}_n(\boldsymbol{\Theta})| - \frac{1}{2}\left(\mathbf{z}_n - \boldsymbol{\mu} - \boldsymbol{\mu}_n(\boldsymbol{\Theta})\right)^\top \boldsymbol{\Sigma}_n^{-1}\left(\mathbf{z}_n - \boldsymbol{\mu} - \boldsymbol{\mu}_n(\boldsymbol{\Theta})\right).$$

The vector $\boldsymbol{\Theta}$ is estimated by maximizing the total conditional log-likelihood function, i.e.,

$$\mathcal{L}(\boldsymbol{\Theta}) = \sum_{n=2}^{N} \mathcal{L}_n(\boldsymbol{\Theta}; \mathbf{z}_n).$$

Because the predictions done in this study are over a very long time horizon—more than 50 years—it requires that the ESG be stationary. As explained in Lütkepohl (2005), the stationarity condition for the VAR process requires that all eigenvalues of the autocorrelation matrix $\boldsymbol{\beta}$ have a modulus of less than 1. To make sure that the conditional variances are non-negative and stationary, the GARCH parameters a_i, b_i and ω_i also need to satisfy the following conditions, as defined in Bollerslev (1986):

$$a_i \geq 0, \quad b_i \geq 0, \quad \omega_i \geq 0, \quad \text{and} \quad a_i + b_i < 1, \quad \text{for } i = 1, ..., 5.$$

Calibration Under the Risk-Neutral Measure

We assume that the risk premium is zero for the inflation and dividend yield, and, as discussed in Appendix A, that the risk premium parameters of the excess stock returns are the same as the $\mathbb{P}$-parameters to achieve consistency between the real-world model and the risk-neutral model. By making these assumptions, some of the parameters in $\boldsymbol{\lambda}_0$ and $\boldsymbol{\lambda}_1$ are fixed and the number of parameters that need to be estimated is reduced to 12. Let $\mathbf{u}$ be the vector that contains all the free risk premium parameters and let $\boldsymbol{\lambda}_0$ and $\boldsymbol{\lambda}_1$ be two matrices are defined as

$$\boldsymbol{\lambda}_0 = \begin{bmatrix} u(1) \\ u(2) \\ 0 \\ 0.03472 \\ 0 \end{bmatrix}, \quad \boldsymbol{\lambda}_1 = \begin{bmatrix} u(3) & u(4) & u(5) & u(6) & u(7) \\ u(8) & u(9) & u(10) & u(11) & u(12) \\ 0 & 0 & 0 & 0 & 0 \\ -0.00340 & 0.06756 & -0.39160 & 0.10891 & 0.00217 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix}.$$

The risk-neutral parameters are calibrated by minimizing the sum of the squared differences between the historical yields of the zero-coupon bonds and the model bond yields. Under the $\mathbb{Q}$ measure, the bond yields with different maturities can be calculated by using Monte Carlo simulation. For specific values of $\boldsymbol{\lambda}_0$ and $\boldsymbol{\lambda}_1$, we generate $10,000$ risk-neutral scenarios for each month.[11] Thus, the price of a zero-coupon bond paying one dollar and maturing in T months at any given historical time point n_h can be approximated by

$$\widehat{P}_{n_h}^{(T)} = \mathbb{E}^{\mathbb{Q}}\left[\exp\left(-\sum_{m=1}^{T} r_{n_h+m}^{s}\right)\middle|\mathcal{F}_{n_h}\right] \approx \frac{1}{10,000}\sum_{i=1}^{10,000}\exp\left(-\sum_{m=1}^{T} r_{n_h+m}^{s,(i)}\right),$$

where $r_{n_h+m}^{s,(i)}$ denotes the i^{th} simulated short rate in the m^{th} month after the starting time n_h. Then, the model continuously compounded monthly yield for the T-month zero-coupon bond at time n_h is given

[11] At the beginning of each month, we start our simulation with the then current values of the economic variables and project the series forward for 180 months.

by

$$\widehat{r}_{n_h}^{(T)} = -\frac{\log\left(\widehat{P}_{n_h}^{(T)}\right)}{T}.$$

The vector **u** is then calibrated such that the sum of the squared difference between the model bond yields $\widehat{r}_{n_h}^{(T)}$ and the actual bond yields $r_{n_h}^{(T)}$ is minimized, i.e.,

$$\arg\min_{u} \sum_{n_h=1}^{N} \sum_{T} \left(\widehat{r}_{n_h}^{(T)} - r_{n_h}^{(T)}\right)^2, \tag{15}$$

where $T \in \{12, 36, 60, 84, 120, 144, 168, 180\}$.

Table 1: **Physical Parameter Estimates for the VAR-GARCH Model.**

Panel A: VAR Estimates.							
		$\log(r_n^s)$	$\log(r_n^l)$	i_n	r_n	$\log(d_n)$	R^2
$\boldsymbol{\beta}$	$\log(r_{n+1}^s)$	0.97274	0.05668	-0.61300	-0.29114	-0.03143	0.98138
		(0.00000)	(0.01104)	(0.72027)	(0.00852)	(0.00048)	
	$\log(r_{n+1}^l)$	0.00151	0.98779	0.14012	-0.11352	0.00138	0.98797
		(0.75795)	(0.00000)	(0.86633)	(0.03118)	(0.80196)	
	i_{n+1}	-0.00002	-0.00011	0.13351	0.00239	-0.00129	0.04721
		(0.95951)	(0.85406)	(0.00523)	(0.45318)	(0.00022)	
	r_{n+1}	-0.00239	0.00324	-0.41396	0.03639	0.00754	0.01058
		(0.61996)	(0.69250)	(0.50446)	(0.48594)	(0.03745)	
	$\log(d_{n+1})$	0.00045	-0.00844	-0.07619	-1.06080	0.99801	0.99446
		(0.58370)	(0.00000)	(0.59563)	(0.00000)	(0.00000)	
		$\log(r_n^s)$	$\log(r_n^l)$	i_n	r_n	$\log(d_n)$	
$\boldsymbol{\nu}$		-0.05370	-0.04947	-0.00751	0.05101	-0.05611	
$\boldsymbol{\gamma}$		-0.50000	-0.50000	0.00000	0.20391	-0.50000	
Panel B: GARCH Estimates.							
		$\log(r_n^s)$	$\log(r_n^l)$	i_n	r_n	$\log(d_n)$	
$\boldsymbol{\omega}$		0.00057	0.00029	0.00000	0.00006	0.00011	
		(0.35186)	(0.29881)	(0.79165)	(0.77700)	(0.00000)	
A		0.26859	0.16820	0.05688	0.10179	0.82056	
		(0.00028)	(0.00074)	(0.00006)	(0.0100)	(0.00000)	
B		0.72141	0.72434	0.93082	0.86577	0.16944	
		(0.00000)	(0.00001)	(0.00000)	(0.00000)	(0.00000)	

The VAR-GARCH parameters are estimated using maximum likelihood. Panel A reports the estimates of the mean process, including the autoregressive matrix $\boldsymbol{\beta}$, the constant vector $\boldsymbol{\nu}$, and the convexity correction and equity risk premium vector $\boldsymbol{\gamma}$. Panel B shows the estimates of the GARCH parameters.

Results

Panel A of Table 1 reports the estimated parameters (with *p*-values in brackets) for the VAR-GARCH model.[12] From the VAR estimate, we can observe that both the logarithmic transform of the short-term and the long-term bond yields are well-explained by their own lagged value and the inflation rate. The model explains a large proportion of the variations in the short-term and long-term bond yields with R^2 values as high as 98.13% and 98.79%, respectively (see the rightmost column of Panel A). Inflation rates appear to be weakly related to lagged values, with an R^2 of 4.72%. On the other hand, historical excess stock returns are consistent under the model with a white noise process, in which the current value and the lagged value have no significant autocorrelations with other economic variables. The relationship between the excess stock return and the dividend yields is significant. Dividend yields rise if stock prices decline and a higher dividend yield predicts an increase in the excess stock returns for the next period. The variability in the dividend yields is also well explained by the model with the R^2 equal to 99.44%.

Table 2: **Risk-Neutral Parameter Estimates for the VAR-GARCH Model.**

	$\boldsymbol{\lambda}_0$	$\boldsymbol{\lambda}_1$				
		$\log(r_n^s)$	$\log(r_n^l)$	i_n	r_n	$\log(d_n)$
$\log(r_n^s)$	-0.0010	0.0436	-0.0291	3.3698	-0.0475	-0.0110
$\log(r_n^l)$	-0.0612	0.0047	-0.0142	-0.3890	-0.1024	-0.0012
i_n	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
r_n	0.0510	-0.0024	0.0032	-0.4140	0.0364	0.0075
$\log(d_n)$	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000

The risk premium parameters $\boldsymbol{\lambda}_0$ and $\boldsymbol{\lambda}_1$ are calibrated by minimizing the total squared differences between the model zero-coupon bond yields and the actual zero-coupon bond yields.

Panel B of Table 1 shows the estimates of the GARCH parameters and the estimated unconditional variance levels for the five economic variables. In the GARCH model, a relatively high a_i and low b_i indicates a more volatile variance than those with relatively low a_i and relatively high b_i. From the GARCH parameter estimates, the dividend yield has the most volatile variances when compared to the other variables, while the inflation process has the most stable variance.

The estimates for the risk premium parameters are given in Table 2. With these estimates, the

[12]The *p*-values are testing whether the estimated parameters are significantly different from zero.

Table 3: **Statistics of the Historical Mispricing Term Structure.**

	$r^{(12)}$	$r^{(36)}$	$r^{(84)}$	$r^{(120)}$	$r^{(144)}$	$r^{(168)}$
Mean	0.00089	0.00038	-0.00045	-0.00034	-0.00011	0.00015
Standard Deviation	0.00288	0.00373	0.00256	0.00153	0.00097	0.00075

The difference between the model bond yield and the actual bond yield, $r^{(m)}$, is calculated for the 302 historical months available. The first row of the table reports the averages of the differences for bonds with different maturities m (given in months) and the second row reports the standard deviation of these differences.

value of the objective function in Equation (15) is 0.0133. Table 3 reports the mispricing statistics of the fitted bond yield term structure. The fit is worse on the short end of the yield curve: on average, the model overprices the one-year bond yield by 0.89%, which is the highest mispricing shown in the table.