How Much Can Collective Defined Contribution Plans Improve Risk-Sharing?

Deborah Lucas*
Daniel Smith**

This version: June 2020

* MIT Sloan School (corresponding author) dlucas@mit.edu

** MIT Golub Center for Finance and Policy (GCFP)

We would like to thank Zvi Bodie, Robert Merton and Robert Pozen for many insightful conversations and suggestions, and Jun-Mok Kim, who helped develop an earlier version of the CDC model presented here.

Abstract

Collective Defined Contribution (CDC) plans have been suggested as an attractive and sustainable alternative to public sector DB plans. A CDC plan is a hybrid structure, designed to provide more predictable retirement benefits than a traditional DC plan while operating at the lower cost of a DB plan. It does this by sharing investment risk across worker cohorts and centralizing asset management. We develop a model of an unsubsidized CDC plan, and use it to characterize the risk-sharing rules and investment policies that maximize a "scheduled benefit" for retirees that is almost always achieved or exceeded. We compare the outcomes under the CDC system with those from an otherwise similar options-augmented DC model, where participants have access to self-financing strategies that involve trading in one-year put and call options. The ability to effectively trade long-dated options in the CDC framework delivers a somewhat higher scheduled benefit than can be achieved by self-insuring in an options-augmented DC plan. However under current contribution policies, the scheduled benefit in the CDC plan falls short of what most would consider an adequate retirement income.

1. Introduction

The number of defined benefit (DB) pensions in the U.S. private sector has sharply declined in recent decades. The falloff has been attributed to a variety of factors that include high financial and regulatory costs to employers, tightened disclosure requirements, and a decline in unionization. By contrast, DB pensions are still prevalent in the U.S. public sector, where almost all states and cities sponsor such plans for their employees. However, many of those public plans are under financial stress and their sustainability is in doubt. Collectively their unfunded liabilities totaled over $4 trillion in 2017 on a market value basis (Rauh, 2019).

Defined contribution (DC) plans, such as the popular 401(k), are the leading alternative to DB. DC plans by construction never have an underfunding problem because future benefits are limited to what can be supported by accumulated assets. However, public-sector unions have resisted moving to DC. A switch would open the door to benefit cuts, and shift the risk of underfunding retirement from public employers to their employees. Participants, newly responsible for making their own investment decisions and often with limited financial knowledge, are at risk of making poor investment choices. Moreover, the administrative costs associated with individual accounts can be considerable.

A Collective Defined Contribution (CDC) plan is a hybrid structure, designed to provide more predictable benefits to retirees than a traditional DC plan, while operating at the lower costs of a DB plan. Variants have been adopted in several places including the Netherlands and the province of New Brunswick in Canada. A vital feature of a CDC plan is that all employer and employee contributions are invested and professionally managed in one collective pool. Benefits depend on investment performance, but uncertainty is mitigated by the sharing of investment risk across worker cohorts.

We consider the optimization problem facing a CDC plan manager whose objective is to offer employees the highest attainable "scheduled benefit," subject to the constraint that the realized benefit for vested workers falls below that scheduled benefit only rarely and by a limited amount. Managers control the asset allocation strategy and parameters of the mechanism for sharing risk across worker cohorts. We focus on steady-state outcomes of a closed system, without any subsidies from sponsors or taxpayers or cross-subsidies between cohorts.[1]

[1] A transition from an existing DB plan to a CDC plan involves additional considerations such as what to do about current underfunding of vested benefits, and how much should be credited to the notional personal accounts of existing workers at the time of the transition. Once these assumptions are specified, the model can be easily adapted to analyze outcomes along alternative transition paths.

The model is calibrated to match a typical age distribution of new public sector hires, typical age-specific separation rates, and average historical real wage growth. Monte Carlo simulation and grid search are used to solve for the optimal investment and risk-sharing policies.

We then turn to the central question posed, of whether the achievable outcomes in an optimized CDC structure are significantly better than for an "options-augmented DC plan"? An options-augmented DC plan incorporates a self-financing strategy of selling one-year calls and buying one-year puts into an otherwise standard DC framework. The comparison is salient because the risk-sharing mechanism in our CDC system implicitly trades puts and calls between worker cohorts. A theoretical advantage of the CDC structure is its ability to synthetically create the opportunity to trade long-dated options, instruments that are illiquid and often unavailable to individual investors in DC plans. The quantitative advantage of the CDC plan is evaluated by comparing its optimized outcomes with those obtainable in an optimized, options-augmented DC plan—one with the same goals, investment opportunities and contribution rates, and access to one-year put and call options that are priced using the approach of Black and Scholes (1973) and Merton (1973).

The analysis is motivated in part by the political economy concern that a government-sponsored CDC plan will tend to offer unsustainable benefits that eventually will lead to subsidies from taxpayers or in-plan transfers from younger to older cohorts, as has happened with DB plans. After all, a scheme that purports to produce fairly safe benefits for retirees but that partially funds those benefits with risky asset holdings without recourse to the plan sponsor is at risk of promising more than can be delivered. There is also always a tension between the goal of providing adequate retirement benefits and the cost of the contributions required to fund them.

The model may help address some of these concerns by providing a framework that provides realistic benchmarks for configurations of contribution rates, scheduled benefits, and investment policies that are in fact sustainable without subsidies. Furthermore, quantifying the benefits of a CDC system over a simpler options-augmented DC system can help all stakeholders--sponsors, unions, plan participants and taxpayers--evaluate whether the gains from greater risk-sharing and centralized asset management justify a CDC system's greater complexity. The model is detailed and general enough for plan sponsors to be able to adapt it to assess how the adoption of such a plan would perform based on the characteristics of their workforce, contribution rules, and other system parameters.

As anticipated, the unsubsidized CDC plan outperforms the corresponding options-augmented DC system: it is able to deliver a higher scheduled benefit to retirees because of the increased opportunities for risk-sharing across cohorts. However, the calibration suggests the attainable improvement is likely to be modest. Under either system, realized replacement rates are low relative to what is typically viewed as adequate for a comfortable retirement. Achieving a higher scheduled benefit requires either tolerating more shortfall risk and investing in riskier assets, or upping

contribution rates. For the CDC plan, the size of the scheduled benefit is also sensitive to the vesting rule because it determines how much is transferred from early separators to those that reach full vesting status.

The analysis here is related to several strands of the literature on retirement finance. The general topic of intergenerational risk-sharing via pension systems has received considerable attention, e.g., Gollier (2008), and Beetsma et al. (2009). Hoevenaars and Ponds (2008), and Bams et al. (2016) study risk-sharing in stylized CDC models. Previous research has recognized that minimum benefit guarantees and other conditional intra-fund transfers can be valued using options pricing techniques, and the idea has been applied to value minimum benefit guarantees on private accounts proposed to replace Social Security, e.g., by Smetters (2001), Biggs et al. (2006) and Sinclair et al. (2006).

The rest of the paper is organized as follows: Section 2 describes the objective function, and presents the CDC and options-augmented DC models. Section 3 describes the calibration of asset returns, contribution rates, payout rules, vesting rules, separation rates, and wages. Section 4 reports the results of a horse race between the optimized CDC and options-augmented DC models. Section 5 examines the sensitivity of the results to certain calibration and modeling choices. Section 6 concludes.

2. Model specifications

Here we describe the structure of the CDC and the options-augmented DC plans, along with the objective function against which both systems are evaluated.

2.1 Objective function

CDC plan managers are assigned the objective of maximizing a "scheduled benefit" for vested retirees, subject to the constraint that the realized benefit falls short of the scheduled benefit only rarely and by a limited amount. The objective is motivated by several considerations. For public sector workers that are accustomed to a DB environment, this will produce outcomes that are more familiar and predictable than would a switch to a standard DC plan. Certain behavioral considerations also suggest that this sort of objective may be optimal (Boes and Arjen, 2018). However, this objective function is considerably more risk averse than a standard utility specification, as no value is assigned to either the size or frequency of realized benefits that exceed the scheduled benefit.

Specifically, the objective is to choose investment and risk management policies that maximize the scheduled benefit for fully vested workers that retire at the normal retirement age. The scheduled benefit is subject to the constraint that the realized benefit only falls below the

scheduled benefit P_1 percent of the time, and that it only falls below a share λ of the scheduled benefit P_2 percent of the time.

The dual cutoff reflects two goals: (1) ensuring that the scheduled benefit is quite safe, and (2) ensuring that left tail outcomes are very rare. Making the latter operative requires setting $P_2 < P_1$. Lower values of P_1 and P_2, and higher values of λ, imply safer benefits, but they also preclude the higher returns from riskier investment policies. As always, there is a risk- return tradeoff.

In the analysis of the options-augmented DC model, we assume that individual participants or the investment professionals working on their behalf have this same objective function. Under that plan, all contributions are assumed to vest immediately.

2.2 CDC model

The necessary ingredients of any CDC system include: a centrally managed portfolio of assets; contribution rules that determine how much workers and employers pay into the fund each year; and the distribution rules that determine the payouts to retirees. There are many possible system designs that are consistent with this broad approach.

The implementation we consider is designed to focus on the fundamental question of whether a CDC plan can deliver an adequate and low-risk benefit to retirees through well-chosen investment policies and risk management practices, with typical DC contribution rates, and without subsidies or cross-subsidies. Relatedly, we use the model to explore how much the internal risk-sharing structure mitigates the risk of investing a portion of the collective portfolio in stocks and other risky assets.

The CDC system presented here can be understood in terms of the formal objective function for plan managers; the assumed dynamics of workforce demographics, compensation, and contribution rules; the account structures that implement risk-sharing and that track retirement claims as a function of cohort and vesting status at separation or retirement; and the investment instruments and risk management strategies available to plan managers.

The assumed vesting rule has a significant effect on the results. The less generous are the benefits for workers that separate before they fully vest, the higher is the attainable scheduled benefit. Many considerations go into a sponsor's choice of a vesting rule, and optimizing over that choice is outside the scope of this analysis. We focus on a particular vesting rule for the standalone CDC plan that favors vested retirees over early separators, as is typically the case for existing DB plans. We also report results for the case of immediate full vesting for comparability

to the options-augmented DC plan, and to be able to clearly distinguish between the gains from improved risk-sharing and those from redistribution via the vesting rule.

Numerical methods are used to solve for the investment and risk management policies that jointly maximize the manager's objective for a steady state distribution of worker ages and tenures. We focus on steady-state outcomes to establish what is achievable in a closed, mature system.[2]

The model incorporates the considerable heterogeneity within cohorts in starting age and eventual tenure, and the complicated value transfers that occur within and across cohorts arising from early separations and vesting rules. Because of that complexity, the model is more clearly described narratively than with equations. The R code that is used to produce the numerical solutions provides the most concise algebraic description of the system, and it is available in an online Appendix.

2.2.1 Risk management via the Reserve Fund

The "Reserve Fund," together with the rules governing its flows, is the mechanism by which investment risk is shared across worker cohorts. As described below, each active worker has several notional personal accounts that interact with the Reserve Fund, and which at separation or retirement determine the lump sum or benefit annuity paid out.

Specifically, the Reserve Fund has a notional claim on a portion of the collectively managed plan assets.[3] It is credited each year with investment returns that exceed a specified threshold, with the rest allocated to workers' personal accounts.

If a fully vested worker's account balance at retirement or separation is insufficient to support the "scheduled benefit" for that tenure, and if the Reserve Fund has a positive balance, then funds are transferred from the Reserve Fund to make up the shortfall. The focal scheduled benefit in this analysis is for retirees over a full working life. Because of the protection from the Reserve Fund, the scheduled benefit serves as a floor on retirement benefits for most vested retirees. However, the scheduled benefit is not guaranteed; following extended periods of unusually low returns that deplete the Reserve Fund, new retirees will receive a smaller annuity

[2] Optimal policies outside of the steady state are of practical interest, but what is optimal is likely to vary depending on the situation. For instance, analyzing a transition to a CDC plan from an underfunded DB system would require additional assumptions about which workers would receive benefits under the old and new systems, and how past accruals would be incorporated.

[3] Although here we assume the Reserve Fund has the same risk profile as the rest of the portfolio, plan managers could choose to invest the reserves in a safer or risker asset mix. The outcomes do not appear to be sensitive to this assumption.

payment. For interpretability, all benefits are expressed in terms of a replacement rate, i.e., the retirement annuity payment as a share of average annual wages over the last several years of employment.

An important choice variable is the target size of the Reserve Fund, which is expressed as a multiple of annual payouts. Targeting a higher average Reserve Fund balance reduces the risk of exhausting it, but also reduces the average amount of funds available to distribute to separators and retirees. When the Reserve Fund exceeds its target size, a portion of the excess is proportionally credited back to worker accounts as described below.

The Reserve Fund transfer rules generally have the side effect of transferring resources from unvested separators to vested retirees. All workers pay into the Reserve Fund as long as they are employed by the sponsor, but unvested separators receive back only a portion of those funds.

2.2.2 Account structures and payout rules

A "cohort of workers" refers to workers hired in a given year at a given age. Each cohort of workers has associated with it three notional accounts that are tracked over time: "Cohort Assets," "Worker Assets," and "Floor Assets." All active workers that belong to the same cohort share the same accounts, with each worker having a notional claim to his or her per capita share of the aggregate that may depend on vesting status. Participants have no control over the accounts; they exist only to keep track of contributions and investment returns, and as an input into the calculation of retirement or separation payouts.[4] The multiple accounts are necessary to keep track of the claims of workers with differing wage histories, tenures, and vesting status, and to track asset flows to and from the Reserve Fund.

The per capita balance of the Cohort Asset account for an active worker is based on the sum of cumulative contributions from workers, C_t^w, and employers, C_t^e, a pro rata share of any distributions from the Reserve Fund, E_t, and investment returns, R_t, that are capped at a ceiling rate of return RC, minus cumulative payouts to separated workers. Asset returns above the ceiling are credited to the Reserve Fund instead of to Cohort Assets. The ceiling effectively controls the strike price of call options written by participants, the value of which is collectivized in the Reserve Fund to provide insurance for vested workers whose average investment returns over their tenure are low. To summarize, the evolution of the Cohort Asset accounts follow:

[4] Plan communications to participants would not reference these accounts. Instead, sponsors would disclose statistics relevant to retirement planning, including scheduled benefits, and possibly also projections of realized benefits as a function of years until separation.

(1) $$A_t^{Cohort} = A_{t-1}^{Cohort} \times Min(RC, R_t) + E_t + C_t^w + C_t^e$$

For workers that separate prior to vesting, the difference between the amount in their Cohort Asset account and the lump sum payout they receive is transferred to the Reserve Fund, and their Cohort Asset account is zeroed out. That transfer to the Reserve Fund is generally positive but need not be.

The ceiling rate of return, *RC*, is an important choice variable for plan managers that affects the rate of accumulation into the Reserve Fund. We refer to transfers to the Reserve Fund through this channel as a sweep.

An alternative to the sweep explored in the sensitivity analysis is for the plan manager to sell call options to the market with a strike price equal to the same ceiling rate of return. Proceeds from the option sales are deposited into the Reserve Fund the payouts, which are identical to those with the sweep, are debited from Cohort Assets accounts. In the absence of transaction costs, both strategies result in transfers into the Reserve Fund of equivalent value, but the latter provides a more stable funding stream that in principle might support a higher scheduled benefit.

The Worker Assets account tracks the payout that a worker is eligible to receive when a separation occurs before full vesting. The balance depends on tenure, wage history, realized returns, and the details of the vesting rule. As noted earlier, any difference between the Cohort Assets and Worker Assets of separators is absorbed by the Reserve Fund in the year separation occurs.

The Floor Assets account tracks the value basis for scheduled benefits for each cohort of fully vested participants as a function of their years of service. The account includes all contributions plus accumulated interest, where the interest accrues at a fixed floor rate, *RF*. The floor rate translates into the focal scheduled benefit with 32 years of completed service. Floor Assets accounts evolve according to:

(2) $$A_t^{Floor} = A_{t-1}^{Floor} * RF + C_t^w + C_t^e$$

To summarize, payouts to separating and retiring workers are based on these account balances. Workers that separate prior to vesting receive a lump sum payout that equals their notional claim on the Worker Asset account in the year of separation. At retirement, vested participants receive a life annuity that in most cases is based on the maximum of their notional claim on the Cohort Asset account and on the Floor Asset account. However, if the Reserve Fund is depleted and

Floor Assets exceed Cohort Assets, the life annuity is based on their notional claim on the Cohort Asset account topped up with any available Reserve Fund assets. In the variant considered in the sensitivity analysis that also caps the realized benefits of vested retirees, the excess in their Cohort Asset account over what is needed to pay benefits at that cap is transferred to the Reserve Fund.

These payout and fund flow rules collectively ensure that the sum of the Cohort Asset accounts across all cohorts and the Reserve Fund is always equal to the balance of the collectively managed asset portfolio.

2.2.3 Workforce demographics, compensation, and contribution rules

The number of workers in each age bracket, from the minimum age for new hires to the retirement age, is initialized to match the steady-state worker distribution by age. The steady-state is a deterministic function of the assumed age distribution of new hires, separation rates by age, and the fixed retirement age. It can be computed by repeated forward iteration of the implied transition matrices. To maintain a stable age and tenure distribution for the workforce, the number of newly hired workers is set to match the number of separating and retiring workers in the steady-state. This also maintains a constant number of newly retiring and separating workers over time.

We abstract from several important sources of volatility in realized benefits. Real wages are assumed to grow at a constant rate that is independent of age or tenure, and inflation is held constant.

Workers and employers each contribute a fixed fraction of wages to the CDC plan annually. All workers participate in the plan from the date they are hired through separation or retirement. The reported scheduled benefit is the replacement rate for a vested worker that retires after 32 years of service.

2.2.4 Asset allocation and summary of control variables

Pension funds typically invest across a wide range of asset classes that include domestic and foreign stocks, bonds, cash, real estate, and alternative assets. We collapse all risky assets into a single category called “stock.” The asset portfolio is allocated between stock and a risk-free bond. Stock returns are stationary, and represent the only source of uncertainty in the steady state.

In principle the asset allocation strategy could be contingent on any function of the state variables of the steady state system, which include the balances in the Reserve Fund and Cohort Asset accounts. For most of the analysis, portfolio strategies are restricted to choosing the fixed share invested in stocks. In the sensitivity analysis we also consider allocations to stock that are a linear function of the Reserve Fund multiple.

A CDC system aims to economize on transactions costs and to minimize investment mistakes by centralizing and professionalizing asset management. The asset returns assumed here are net of average transactions costs. It would be straightforward to incorporate trading costs that depend on turnover, but because investment policies are limited to simple trading strategies we abstract from those effects.

To summarize the control variables, plan managers pursue the objective stated in section 2.1 through their choice of: (1) an asset allocation policy; (2) the ceiling *RC* on the rate of return credited to the Cohort Account; and (3) the target size of the Reserve Fund. Managers take as given the parameters governing demographics, compensation, contribution rates, vesting rules, and asset returns.

2.2.5 Monte Carlo Simulation

The time step in the model is annual, and outcomes are based on Monte Carlo simulations of sequences of risky asset returns. At the beginning of each Monte Carlo run, the age distribution of the active workforce corresponds to its steady-state distribution. All accounts and the Reserve Fund start with no assets, consistent with the absence of external subsidies to the system. The optimal policy is found by searching over the multi-dimensional grid of admissible choice variables for plan managers. The reported results are based on 10,000 Monte Carlo simulations.

In order to ensure that the system has reached its stochastic steady state, the determination of the scheduled benefit is based on the distribution of realized outcomes in the 35^{th} year of system operation. We verified that at this horizon the system has in fact stabilized by comparing the results to those obtained for longer simulation runs.

In each successive year, a fraction of each cohort of active workers separates or retires at a fixed rate that depends on age. An equal number of new hires replace the separating and retiring workers. The age distribution of new hires is also fixed, and skewed towards younger ages. The net effect of separations and new hires is to maintain a stable size and age distribution of the active workforce. The wages for all active workers grow at a deterministic rate.

Each year the realized return on stocks is determined by a random draw from the assumed distribution of returns. Those returns affect the balances of the collectively managed asset pool and the various funds according to the rules described above. The resources transferred to separating or retiring participants are decremented from the collective asset pools and the various accounts.

2.3 Options-augmented DC Model

The options-augmented DC model is designed to be as parallel as possible to the CDC model. Its purpose is to evaluate how the achievable benefit outcomes from the CDC model compare with what is attainable with a standalone DC account that is augmented with a self-financing portfolio of liquid stock options that can put a floor on risky asset returns. To that end, the employer and employee contribution rates and investment opportunities are the same for both models. In the version of the CDC plan used for comparison, all participants fully vest immediately. In the options-augmented DC system, the optimal investment and the self-financing options strategies are chosen to maximize the same objective function as for the CDC plan.

In place of the Reserve Fund, which implicitly trades both long and short-dated options between participants in different cohorts, an employee or delegated asset manager can limit risk by trading in a self-financing portfolio of one-year put and call options. To maintain the parallel with the sweep of returns above a ceiling into the Reserve Fund in the CDC model, we assume that each year, a call option is written on the entire asset portfolio. The strike price on the call options is a choice variable. The proceeds are used to buy put options. Possible strategies range from buying large numbers of deeply out-of-the-money puts, to buying small numbers of in-the-money puts. Option values are calculated using Black-Scholes-Merton.

As for the CDC model, portfolio choice involves choosing the constant percentage invested in stocks. The options strategies are limited to choosing a put and call strike price that are a constant multiple of current asset value. More complicated strategies that depend on tenure or accumulated balances might produce a higher scheduled benefit, but would require greater sophistication to execute and entail higher trading costs and are not explored.

In this setting a single account can be used to track accumulated savings based on contribution rates and asset returns, inclusive of the effects of the options. Separation rates and other demographic factors are irrelevant because the focus is on the achievable replacement rate for a worker that remains active until reaching the normal retirement age, and how that compares to what the same worker would obtain in a CDC plan.

The scheduled benefit for the options-augmented DC plan is identified from the simulated distribution of asset values at retirement for a worker that has worked for the requisite number of years. Specifically, it is possible to identify the point in the value distribution where a fraction P_1 of realizations are below it. One can then check whether, at that point, less than a fraction P_2 of realizations are less than a share λ of that point. If so, the first constraint is the binding one and that point is the scheduled benefit. If not, the second constraint binds, and the scheduled benefit equals the outcome such that realizations falls below a share λ of it exactly P_2 percent of the time.

3. Calibration

The parameters of the CDC and options-augmented DC models are described here. The two models are calibrated with identical parameters unless otherwise noted.

3.1 Contributions and Vesting Rules

Workers and employers each contribute 10 percent of wages to the plan each year, for a total 20 percent contribution rate. The contributions are invested as part of the centralized asset pool, and credited to the relevant accounts.

The main vesting rule used to evaluate the CDC model is as follows: Full vesting occurs with 10 years of service. The worker's portion of the contribution vests immediately, along with the investment income earned on those balances. The employer contributions and associated earnings cliff vest at 10 years of service. Separators receive a lump sum payout equal to the balance in the Worker Asset account, which is the sum of those two components.

We also evaluate the CDC model under the assumption of immediate full vesting in order to create a level comparison with the options-augmented DC model. Although DC plans also may not provide immediate vesting of an employer's contributions, the transfers from unvested separators would accrue to the sponsor rather than to vested retirees as in the CDC case. That difference in the incidence of transfers complicates comparisons between the two systems when vesting is not immediate. The effects of other vesting rules in the CDC model are also evaluated in the sensitivity analysis.

3.2 Assets and Return Distribution

Plan assets are allocated between risk-free bonds and risky stocks. The nominal risk-free rate is fixed at 2.5 percent. Stock returns are normally distributed, with an expected return of 7.5

percent and a standard deviation of 20 percent on a nominal basis. Inflation is assumed to be constant at 2 percent.

3.3 Demographics, Wages and Replacement Rates

The maintained assumptions about the workforce are summarized in Tables 1 to 4. Table 5 summarizes the implied steady-state age distribution of active workers, and Table 6 shows the implied distribution of tenures at separation or retirement. All simulations are populated initially from that steady-state. Parameters were chosen to be roughly in keeping with the demographics and dynamics of a typical public sector workforce.

Table 1: Fixed economic assumptions

Base wage in start year	$50,000
Real wage growth (annual)	0.50%
Nominal wage growth (annual)	2.50%
Annuitization factor	0.04

The annuitization factor is used to translate accumulated account assets into an annual benefit payment. The annuitization factor is set to .04, a typical value for sellers of life annuities. The replacement rate is calculated as the annual benefit payment divided by the average wage in the last five years of employment. The scheduled benefit is identified with the replacement rate for a vested retiree with a 32-year work history.

Table 2: Fixed plan parameters

Workforce size	100,000
Minimum age at hire	25
Retirement age	62
Service years to vest	32
Employer contribution rate	10%
Employee contribution rate	10%

Table 3: Age distribution of new hires

25 to 34	0.5
35 to 44	0.25
45 to 54	0.2
55 to 61	0.05

Table 4: Annual separation rates by current age

26 to 44	0.02
45 to 54	0.02835
55 to 61	0.0942

Table 5: Steady state worker age distribution

25 to 34	0.15
35 to 44	0.30
45 to 54	0.35
55 to 62	0.20

Table 6: Steady state distribution of tenure at separation or retirement

1 to 9	0.303
10 to 19	0.285
20 to 31	0.313
32 to 37	0.099

3.4 Parameterized Objective Functions

We consider two scenarios for risk tolerance. In the Lower-Risk Scenario, the chance of the realized benefit falling short of the scheduled benefit is less than 10 percent, and the chance of it falling short of 80 percent of the scheduled benefit is less than 2 percent. The Lower-Risk Scenario provides much of the safety of a traditional DB plan. In the Higher-Risk Scenario, the chance of a shortfall is less than 20 percent, and the chance of a shortfall in excess of 50 percent is less than 2 percent. The Higher-Risk Scenario produces a higher scheduled benefit, but at the cost of considerably more risk.

Plan participants are likely to differ in their personal risk tolerances. Those with Social Security, more wealth, or a spouse with good benefits may prefer to take more risk in exchange for a higher scheduled benefit, and those without those protections may be more risk averse. Sponsors of larger plans could accommodate that heterogeneity by allowing new participants a one-time choice to sign up for one of several different risk and scheduled benefit profiles. Effectively the sponsor would be running several separate plans under one umbrella, still with full asset pooling to minimize portfolio management expenses.

3.5 Reserve Fund rebates

The target size of the Reserve Fund is a choice variable, but we exogenously fix the rate at which balances in the Reserve Fund in excess of the target multiple are rebated back to worker accounts at 0.2.[5]

4. Results

The outcomes from two standard DC plans without options, parameterized as above, provide a useful reference point. We consider two benchmarks: (1) a risk-free bond portfolio and (2) a 60/40 stock-to-bond asset mix, both with immediate vesting. Recall that for a DC plan, the scheduled benefit is found by calculating the distribution of income replacement rates implied by a given investment policy, and identifying the scheduled benefit with the cutoff point that satisfies the optimization criteria in section 2.1.

A risk-free bond portfolio delivers a certain benefit of a 28 percent income replacement rate (henceforth described simply as a benefit of .28). That low replacement rate helps explain why DB plan sponsors that have promised higher benefits and do not want to increase contributions choose to hold riskier assets.

For the 60/40 stock-to-bond asset mix, the scheduled benefit depends on the risk tolerance in the objective function. For the Lower-Risk Scenario, the scheduled benefit is lower than for a risk-free portfolio, and for the Higher-Risk Scenario it is slightly higher (see Table 7). The average benefit of 0.46 is significantly higher. Although our objective function puts no weight on the average benefit, that statistic will also be of interest in assessing the attractiveness of a CDC plan.

Table 7: Scheduled and Average Benefits for a 60/40 DC plan

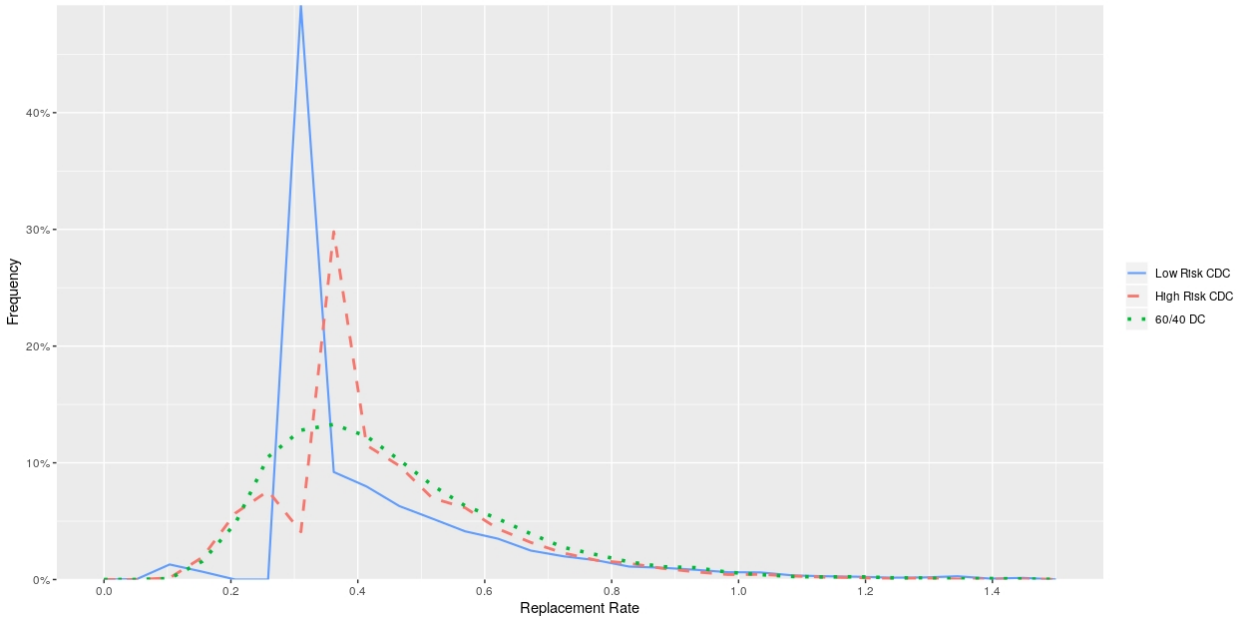
Scheduled Benefit, Lower-Risk Scenario	0.23
Scheduled Benefit, Higher-Risk Scenario	0.30
Average Benefit	0.46

[5] We fix this rate to avoid adding another dimension to the optimization. Changing it to 0.1 or 0.5 and re-optimizing had a negligible effect on the scheduled benefit.

4.1 Optimized CDC Model

The results for the optimized CDC model, calibrated under the base case parameters described in Section 3, are reported in Table 8. The scheduled benefit, average benefit, risky asset share, and Reserve Fund cap are shown for both the Lower-Risk and Higher-Risk scenarios.

Figure 1: Distribution of Realized Benefits for CDC Lower-Risk Scenario, CDC Higher-Risk Scenario, and for 60/40 Benchmark DC plans



The scheduled benefits are significantly higher for the CDC plan than for the benchmark 60/40 DC plan portfolio in Table 7. The increase is 9 percentage points for both scenarios. However, the Lower-Risk Scenario provides a scheduled benefit that is only 4 percentage points above the risk-free benchmark DC plan. Figure 1 compares the entire distribution of benefit realizations for each of the two scenarios with those of the benchmark 60/40 DC plan. It shows that for both scenarios the scheduled benefit is the modal benefit and realized benefits only rarely fall below it. There is also a long right tail of realized benefits that significantly exceed the scheduled benefit. That right tail account for the considerably higher average than scheduled benefit.

What is initially more surprising is that the Lower-Risk Scenario turns out to have a higher share invested in stocks and generates a slightly higher average benefit. The former happens because for the parameters assumed here, risk in the Lower-Risk Scenario is primarily controlled by targeting a much larger Reserve Fund, whereas risk in the Higher Risk Scenario is primarily controlled with a smaller allocation to stocks.[6] Relatedly, the return ceiling above which returns

[6] Because the objective function is piecewise linear with kinks, it can generate "bang-bang" policies that change significantly for small differences in parameter choices.

are credited to the Reserve Fund is much lower in the Lower-Risk Scenario. The differences also arise because the lower tail constraint is binding for the Lower-Risk Scenario, whereas the simple shortfall constraint is binding for the Higher-Risk Scenario. Average benefits do not enter into the objective function, and they may be larger or smaller than the scheduled benefit. However, for both scenarios it turns out that the average benefits are similar to those from the 60/40 benchmark portfolio. Hence, the CDC plan is able to deliver a similar average benefit with much more predictability. Finally, it is important to note that these outcomes can vary considerably depending on the parameters of the model and objective function.

An artifact of the vesting rule is that the Reserve Fund, and thereby vested retirees, are subsidized by unvested separators. This has the effect of making scheduled and average benefits for vested retirees higher when the target size of the Reserve Fund is larger. From the perspective of unvested separators, a further drawback of the vesting rule is that by accommodating a riskier investment policy, a larger Reserve Fund leaves them with a riskier distribution if they separate prior to vesting. To avoid these effects, below we also consider the results with immediate full vesting.

Table 8: Optimized CDC Model

Lower-Risk Scenario	
Scheduled Benefit	0.32
Average Benefit	0.47
Share Stocks	0.69
Reserve Fund Target Multiple	1.91
Ceiling Rate	4.60%
Higher-Risk Scenario	
Scheduled Benefit	0.39
Average Benefit	0.43
Share Stocks	0.47
Reserve Fund Target Multiple	0.15
Ceiling Rate	12.0%

4.1.2 Comparison of CDC and optimized options-augmented DC outcomes

Here we compare the outcomes in the CDC model with those of an options-augmented DC model where risk sharing possibilities are limited to trading in liquid short-term options. We first re-optimize the CDC model under the assumption that workers are immediately vested to avoid the transfers from unvested to vested workers that would not occur in a DC plan. Table 9 summarizes the CDC results with immediate vesting.

Table 9: Optimized CDC model with immediate vesting

Lower-Risk Scenario	
Scheduled Benefit	0.29
Average Benefit	0.43
Share Stocks	0.69
Target Reserve Fund Multiple	2.1
Ceiling Rate	3.25%

Higher-Risk Scenario	
Scheduled Benefit	0.36
Average Benefit	0.44
Share Stocks	0.59
Target Reserve Fund Multiple	0.25
Ceiling rate	15.0%

The scheduled benefits in Table 9 are lower than the corresponding ones in Table 8 because no transfers are coming from unvested workers that separate. Without transfers, the Lower-Risk Scenario has a scheduled benefit that is only one percentage point above that of the benchmark risk-free DC plan. However, it still has the advantage that the average benefit is considerably higher than for a risk-free investment portfolio. The Higher-Risk Scenario delivers a scheduled benefit that is eight percentage points above the risk-free DC plan, and an average benefit that is similar to the Lower-Risk Scenario.

Table 10 reports the outcomes for the options-augmented DC plan, which is intended as a simpler "do-it-yourself" version of a CDC risk-sharing scheme. The option strike prices are reported as a percentage of current plan asset values. With its collar strategy of buying puts financed by selling calls, workers participating in the plan can protect themselves against large negative returns and still retain some of the upside from investing in stocks. Consistent with the more risk-averse objective, we find that in the Lower-Risk Scenario a higher value of calls are sold (i.e., lower call strike price), and the purchased puts are less out-of-the-money.

Comparing Tables 9 and 10, the scheduled benefits are higher in the CDC plan than for the parallel options-augmented DC plan. However, the differences in the level of scheduled benefits are not large. In the Lower-Risk Scenario the CDC model adds only one percentage point to the scheduled benefit over the options-augmented DC model, and three percentage points for the Higher-Risk Scenario. The greater risk-sharing potential in the CDC model arises from the possibility of effectively trading longer dated options across cohorts through the Reserve Fund mechanism. At least for the specifications considered here, that difference generates only modest gains.

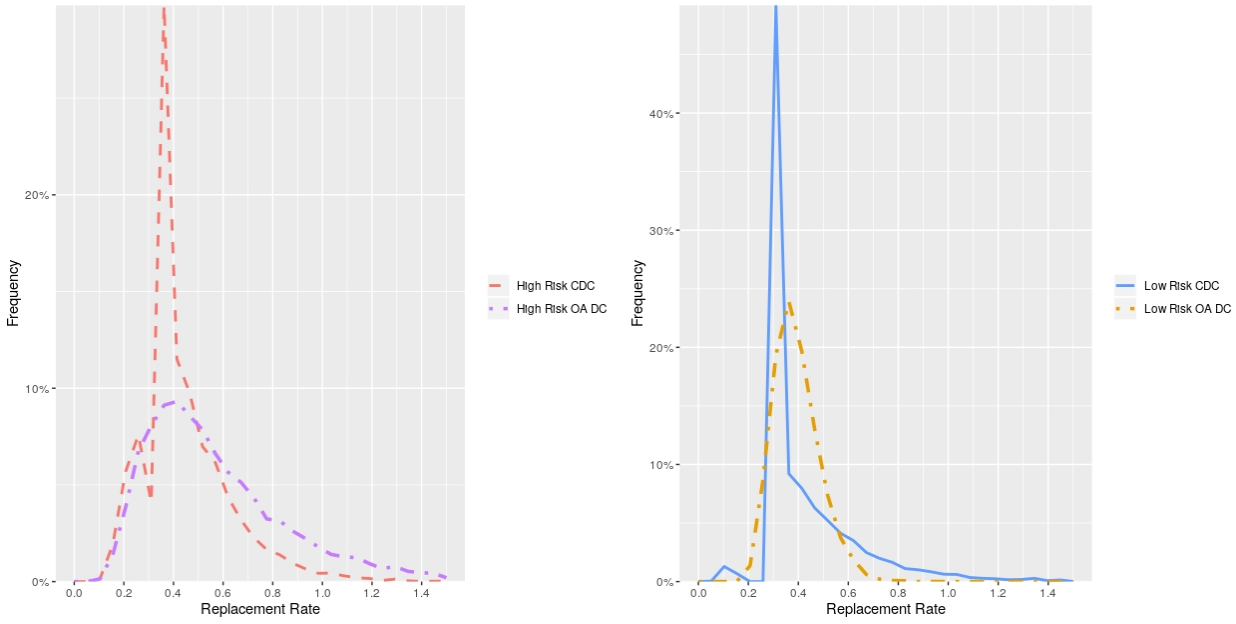
Figure 2 compares the distribution of realized benefits between the CDC and options-augmented DC systems for each scenario. It shows that the CDC model generates a much higher frequency of realized outcomes that are close to the scheduled benefit. That predictability is an additional advantage of the CDC system.

Table 10: Optimized options-augmented DC model

Lower-Risk Scenario	
Scheduled Benefit	0.28
Average Benefit	0.39
Share Stocks	0.60
Call Strike	119%
Put Strike	90%

Higher-Risk Scenario	
Scheduled Benefit	0.33
Average Benefit	0.59
Share Stocks	1.00
Call Strike	126%
Put Strike	85%

Figure 2: Comparing Distribution of Realized Benefits in CDC and Options-Augmented DC Models



5. Sensitivity analysis

The outcomes in the CDC model are sensitive to the details of how it is specified and parameterized. Here we briefly explore the effects of alternative vesting rules, changes in Reserve Fund risk management practices, allowing a more dynamic trading strategy, and altering contribution rates.

5.1 Vesting rules

The sensitivity of the scheduled benefit to the vesting rule is illustrated in Table 11, which shows the results for immediate, 5-year cliff, 10-year cliff, and 32-year cases. Unvested separators always receive a lump sum payout equal to the balance in their Worker Asset account.

Apart from the horizon, the 5-year cliff is identical to the 10-year cliff assumed in the base case analysis. In both cases workers' contributions and returns on them vest immediately. Under the 32-year rule, full vesting occurs only with 32 years of service. As always the worker's portion of the contribution vests immediately, along with the full amount of investment income earned on those balances. In this case the principal amount of the employer contributions also vests immediately, but earnings are only credited on those balances at the risk-free rate. That generates a higher scheduled benefit than in the other cases, but less than 10 percent of workers complete the 32 years necessary to receive it.

Table 11: Optimized CDC model, Alternative Vesting Rules

Lower-Risk Scenario	immediate	5-year	10-year	32-year
Scheduled Benefit	0.29	0.29	0.32	0.39
Average Benefit	0.43	0.49	0.47	0.6
Share Stocks	0.69	0.8	0.69	0.92
Reserve Fund Target Multiple	2.1	1.75	1.9	1.9

Higher-Risk Scenario	immediate	5-year	10-year	32-year
Scheduled Benefit	0.36	0.37	0.39	0.47
Average Benefit	0.44	0.44	0.43	0.44
Share Stocks	0.59	0.56	0.47	0.27
Reserve Fund Target Multiple	0.25	0.59	0.15	0.25

5.2 Alternative risk management strategies

We considered two variants of the CDC model aimed at increasing or stabilizing inflows into the Reserve Fund to potentially increase the scheduled benefit. The first was to cap the realized

replacement rate for vested retirees. That rule generates additional transfers from personal accounts to the Reserve Fund when the cap is hit, truncating the upper tail of realized benefits and potentially providing additional resources to increase scheduled benefits. It effectively introduces transactions involving long-term call options into the model. We considered caps of 0.5 and 0.7. The caps had a negligible effect on the scheduled benefits in either scenario. The intuition for why the additional transfers to the Reserve Fund made little difference is because they happened in good states of the world after long periods of above-average stock returns. At such times, the Reserve Fund is likely to be above its target holdings already and there will also be little need to use it.

The second variant was to replace the sweep of returns on Cohort Asset accounts above a ceiling with the sale of one-year call options with a strike price equal to the same ceiling and covered by the same returns on Cohort Asset accounts. The proceeds from the options sale are deposited into the Reserve Fund, producing a more stable stream of revenue than the sweeps. Unlike the cap on realized benefits, this reduces the correlation between stock returns and transfers into the Reserve Fund relative to the base case analysis. Holding all other policies as in Table 9, the scheduled benefit in the Lower-Risk Scenario increases by 2 percentage points In the Higher-Risk Scenario, which relies much less on the Reserve Fund, the change reduced the scheduled benefit by 1 percentage point.

We also looked at replacing the sweep with a sale of call options without constraining the other policy choices. That resulted in a very different optimal strategy of investing almost entirely in stocks and relying on the options sales to top up shortfalls in Cohort Accounts. It generated a significantly higher scheduled benefit for both scenarios, but it also increased the likelihood of very large shortfalls, and seemed inconsistent with the stability goals of a CDC system.

5.3 Dynamic portfolio allocation strategies

Allowing for the portfolio allocation to risky assets to vary with the state of the system in the CDC model has the potential to increase scheduled benefits over the static allocation strategy considered in the base case analysis. The potential for improvement arises because the objective function may cause effective risk aversion to vary with the Reserve Fund ratio.

To test this idea, we consider the case where the plan manager chooses an allocation rule that is a linear function of the Reserve Fund ratio. That limits the choice to two parameters, a slope, β, and an intercept, α, which keeps the size of the policy choice space tractable. Specifically, the share of stocks is given by

(3) $$\alpha \times S^* + \beta \times (\rho^* - \rho)$$

Here S^* is the optimal static share of stock, ρ^* is the optimal Reserve Fund target ratio, and ρ is the current Reserve Fund ratio. The share is also bounded to be between 0 and 1. Table 12 reports results, holding all other parameters as in Table 9.

Table 12: CDC fully vested model with dynamic portfolio rules

Lower-Risk Scenario	
Scheduled Benefit	0.31
Average Benefit	0.39
S^*, α, β	0.69, 1.15, 0.25
Target Reserve Fund Multiple	2.0
Ceiling rate	3.7%

Higher-Risk Scenario	
Scheduled Benefit	0.37
Average Benefit	0.42
S^*, α, β	0.59, 0.95, 0.02
Target Reserve Fund Multiple	0.5
Ceiling rate	15.75%

5.4 Asset management fees

A potential advantage of a CDC system over an options-augmented DC approach is the cost savings from more efficient asset management and avoiding the fees associated with individual accounts. In the base case we assumed expected returns were identical in either case. However, a CDC plan may be able to achieve higher average returns after expenses. To assess the potential effect of those differences, we considered a reduction of 25 basis points on all investment of the options-augmented DC plan. That has the effect of reducing scheduled and average benefits in that plan by about 2.5 percent relative to the Table 10 results.

5.5 Contribution rates

Scheduled and expected benefits are close to linear in the contribution rate, although not perfectly so because of the effects of vesting rules. Allowing workers to voluntarily increase their contribution rate above the required minimum could help those that want to save more achieve their retirement income goals. Offering that possibility could enhance retirement security and involve minimal additional expense for the plan.

6. Discussion and Conclusions

In this analysis we have developed a modeling framework to evaluate to what extent an unsubsidized CDC system might be able to offer a satisfactory alternative to public sector DB plans; and by how much a CDC system could be expected to improve risk-sharing relative to a much simpler options-augmented DC plan. Our findings suggest that an unsubsidized CDC system does in fact enhance risk-sharing by effectively allowing the trading of long-dated put and call options between worker cohorts, but that quantitatively the gains are not large relative to an options-augmented DC plan.

In the CDC model analyzed, the focus is entirely on managing investment risk. We abstract from the considerable risk to realized benefits arising from uncertain wages, inflation, longevity and other factors. Because other risks are absent, the analysis incorporates less risk-sharing across cohorts than is often envisioned for this sort of system. The potential of a CDC system to diversity additional risk factors could be analyzed using an extension of the model that explicitly incorporates some of them. For instance, partial insurance for additional risk factors could be achieved by using rebates of above-target balances in the Reserve Fund selectively rather than distributing them pro-rata as is done here. We also do not attempt to incorporate a more desirable alternative to a level nominal annuity in the de-accumulation phase (e.g., Muralidhar et al., 2016). Omitting such additional opportunities for welfare enhancement might suggest that our results are a conservative estimate of the potential risk-sharing gains from a CDC-type structure. However, a system with a more complex set of transfer rules and insurance goals may be more susceptible to unintended cross-subsidies and unanticipated shortfalls.

Whether a CDC plan offers participants sufficient choice is also a consideration. Certainly a typical DC plan allows individuals more control over the risk-return tradeoff, although there is disagreement about whether that is advantageous. In any case, a CDC plan can be structured to accommodate some heterogeneity in risk tolerance by allowing workers to choose from a menu of projected benefit distributions. For example, participants could be offered a one-time choice between something akin to the Lower-Risk and Higher-Risk scenarios presented here. The multiple sub-plans could be managed in a collective portfolio that would preserve economies of scale and expertise in asset management. However, an additional challenge would be to effectively communicate to workers the consequences to them of the available choices.

A further concern about CDC systems in general is that, similar to the DB plans they would replace, they tend to entail transfers from younger workers (through the diversion of contributions from investments to current payouts) or taxpayers (in the form of additional levies to fund the system) if Reserve Fund assets turn out to be insufficient to cover implicitly promised benefits. The particular CDC structure we analyze by construction precludes systematic transfers

between cohorts or subsidies from taxpayers. That restriction is essential for understanding what these systems could realistically offer without subsidies, and for comparison with alternative structures that also would not receive subsidies.

The scheduled benefits in the optimized calibrated CDC model may strike many as inadequate, particularly for those states and localities that do not participate in Social Security. Furthermore, in real terms the scheduled benefit declines over time; an inflation protected annuity of equivalent present value at retirement would entail a considerably lower initial replacement rate. The modest level of achievable scheduled benefits reflect that even with significant inter-cohort risk sharing, the amount of investment risk that is consistent with a fairly safe scheduled benefit is quite limited. The results also reflect the assumption of a lower average return environment than in the past. Of course, higher replacement rates always can be obtained by scaling up contribution rates.

This analysis focused on steady-state outcomes of the CDC model in order to examine the potential of a mature system. Switching from a DB to a CDC system involves significant transition issues. A useful feature of the model is that it can be easily adapted to evaluate the distributional consequences of alternative transition rules for an arbitrary set of initial conditions. We hope to use it in future work to shed light on how transitions between systems could be most efficiently and equitably managed.

Finally, although our primary focus was on what could be achieved in a CDC system, it was interesting to discover that the options-augmented DC model was able to produce outcomes that were considerably more predictable than the benefits in a traditional 60/40 DC plan, and fairly similar to those of the CDC model. Whether this sort of DC variant would be an appealing alternative or add-on for traditional DC plans seeking to offer workers more secure benefits is also a question for future research.

References

Bams, Dennis and Schotman, Peter C. and Tyagi, Mukul (2016), "Optimal Risk Sharing in a Collective Defined Contribution Pension System," Netspar Discussion Paper No. 03/2016-015.

Beetsma, R.M.W.J. and Lans Bovenberg, A.(2009), Pensions and Intergenerational Risk-sharing in General Equilibrium. Economica, 76: 364-386.

Biggs, Andrew Clark Burdick, and Kent Smetters (2006), "Pricing Personal Account Benefit Guarantees: A Simplified Approach," in Social Security Policy in a Changing Environment, editors Jeffrey Brown, Jeffrey Liebman and David Wise, University of Chicago Press

Black, Fisher, and Myron. J. Scholes. (1973) "The pricing of options and corporate liabilities," Journal of Political Economy 81, no. 3:637–54

Boes Mark-Jan and Arjen Siegmann (2018) "Intergenerational risk sharing under loss averse preferences," Journal of Banking and Finance, July 2018

Gollier, Christian. (2008). "Intergenerational risk-sharing and risk-taking of a pen- sion fund." Journal of Public Economics, 92(5-6): 1463–1485.

Hoevenaars, Roy and Eduard Ponds (2008), "Valuation of intergenerational transfers in funded collective pension schemes," Insurance: Mathematics and Economics, April 2008

Lachance, Marie-Eve, and Olivia Mitchell (2003). "Guaranteeing individual ac- counts," American Economic Review 93 (2): 257–60.

Merton, Robert. (1973). "Theory of rational option pricing." Bell Journal of Eco- nomics 4, no. 1 (spring): 141–83.

Muralidhar, A., K. Ohashi, and S. Shin. 2016. The Most Basic Missing Instrument in Financial Markets: The Case for Bonds for Financial Security, Journal of Investment Consulting, Vol. 16, No. 2, p. 34-47, 2016

Pennacchi, George G. (1999). "The value of guarantees on pension fund returns." Journal of Risk and Insurance 66 (2): 219–37.

Rauh, Joshua (2019), "Hidden Debt, Hidden Deficits, the 2019 Edition," Hoover Institution Essay

Sinclair, Sven, Deborah Lucas, Amy Rehder Harris, Michael Simpson, and Julie

Topoleski (2006). "Evaluating Benefit Guarantees in Social Security," Congressional Budget Office Background Paper

Smetters, Kent. (2001). The Effect of Pay-When-Needed Benefit Guarantees on the Impact of Social Security Privatization pp. 91—111, In: John Y. Campbell and Martin Feldstein, (Eds.) Risk Aspects of Investment Based Social Security Reform. NBER, Cambridge, MA