

Intergenerational Risk Sharing in a Defined Contribution Pension System: Analysis with Bayesian Optimization

An Chen¹†, Motonobu Kanagawa²‡ and Fangyuan Zhang²¶

¹Institute of Insurance Science, University of Ulm, Germany

²Data Science Department, Eurecom, France

Abstract

We study a fully funded, collective defined-contribution (DC) pension system with multiple overlapping generations. We investigate whether the welfare of participants can be improved by intergenerational risk sharing (IRS) implemented with a realistic investment strategy (e.g., no borrowing) and without an outside entity (e.g., share holders) that helps finance the pension fund. To implement IRS, the pension system uses an automatic adjustment rule for the indexation of individual accounts, which adapts to the notional funding ratio of the pension system. The pension system has two parameters that determine the investment strategy and the strength of the adjustment rule, which are optimized by expected utility maximization using Bayesian optimization. The volatility of the retirement benefits and that of the funding ratio are analyzed, and it is shown that the trade-off between them can be controlled by the optimal adjustment parameter to attain IRS. Compared with the optimal individual DC benchmark using the life-cycle strategy, the studied pension system with IRS is shown to improve the welfare of risk-averse participants, when the financial market is volatile.

Key words: Intergenerational risk sharing, Defined Contribution, Automatic Adjustment Rule, Bayesian Optimization

1. Introduction

Defined Contribution (DC) pension plans constitute key part of pension systems in many countries, such as the “401k plan” in the United States and the “personal pensions” in the United Kingdom. In a DC plan, each participant owns her account, pays fixed contributions to the account regularly, and the accumulated contributions are invested in a financial market. The amount of retirement benefits is determined by the market value of the individual account at the retirement date. By design, DC plans have advantages over traditional pension schemes such as Pay-As-You-Go (PAYG) and Defined Benefit (DB) pension plans in terms of the transparency, fairness, portability and sustainability, which encourage the prevalence of DC plans.

However, DC plans have major issues arising from that each participant bears the investment risk by herself. One issue is that the participant may not choose a good investment strategy due to the lack of financial expertise. Indeed, a typical DC plan participant in reality tends to perform the naive “ $1/n$ diversification,” equally dividing her contributions into the default options provided by the plan (Benartzi and Thaler

† an.chen@uni-ulm.de

‡ motonobu.kanagawa@eurecom.fr

¶ fangyuan.zhang@eurecom.fr (Corresponding)

2001). Optimal investment strategies studied in the literature (e.g., Merton 1971; Cairns 1996; Cairns et al. 2006; Boulier et al. 2001; Vigna and Haberman 2001; Chen and Delong 2015; Menoncin and Vigna 2017) can thus not be easily employed by such participants.

Another major issue of DC plans is the incapability of *intergenerational risk sharing* (IRS), which is to diversify non-diversifiable risks within one generation (e.g., those caused by economic shocks) across different generations. For example, a DC plan participant whose accumulation period overlaps an economic depression faces a financial risk that cannot be diversified by herself; even if she can choose an optimal investment strategy, she may not accumulate enough wealth for retirement. IRS enables diversifying such non-diversifiable risks by sharing the risks among different generations. By definition, however, IRS requires the pension scheme to be *collective*, and thus individual DC plans cannot implement IRS.

It is well known that carefully-designed collective schemes can implement IRS to improve the welfare of participants (e.g., Gordon and Varian 1988; Allen and Gale 1997; Shiller 1999). Gollier (2008) shows that a collective DC pension plan can improve the welfare of participants compared with an individual benchmark using the optimal life-cycle investment strategy. Similarly, Cui et al. (2011) study a collective DB-based hybrid pension scheme where both contributions and pension benefits may be adjusted, and show that it is welfare-improving compared with an optimal individual benchmark. Chen et al. (2016) consider a three-pillars pension system in which the second pillar is a collective hybrid plan, and show it is welfare improvement compared with a corresponding individual DC benchmark. See Barr and Diamond (2008, Chapter 7) and Beetsma and Romp (2016) for an overview of IRS and further references.

While the seminal work of Gollier (2008) shows that IRS in a collective DC pension system is welfare-improving, both his first-best and second-best strategies depend on rather strong assumptions. The first-best strategy attains IRS by enhancing the risk-taking ability of the pension fund, by treating the net present value of the contributions from all the future generations as part of the fund's *total* wealth. Consequently, the fund can invest more than the fund's *actual* wealth, i.e., the fund can perform *borrowing* for investment, similar to the life-cycle investment strategy for an individual investor (Merton 1971). However, borrowing is not realistic for pension funds in reality. On the other hand, his second-best strategy does not allow borrowing, but assumes the existence of an outside entity (shareholders) that helps finance the fund. Therefore, it is not clear whether the welfare-improvement by the second-best strategy can be attained without the shareholders.

Given these limitations of Gollier's analyses, one may ask: *Can the welfare-improvement by IRS be attained by a fully funded, collective DC pension system with a realistic investment strategy (e.g., no borrowing) and without an outside entity that helps finance the pension fund?* Previous related works do not exactly answer this question, as discussed later in detail. For example, while Chen et al. (2016) show that their collective scheme with IRS is welfare-improving compared with the corresponding individual DC scheme, it is assumed that both the collective and individual schemes use the *same* investment strategy; therefore, it is not clear whether their IRS is welfare-improving as compared with the *optimal* individual investment strategy.

Our main aim is to investigate the above question. To this end, we consider a stylized model for a fully funded, collective DC pension fund with multiple overlapping generations, which we call the *IRS-DC* model. As Gollier (2008), each participant pays a fixed annual contribution to the pension fund, and the fund makes investment on behalf of the participants. Different from the first-best strategy of Gollier (2008), however, the fund is not allowed to perform borrowing. Each participant has her own account in the fund,

which accumulates her contributions and is indexed to the fund's investment performance; the account value at the retirement date determines her pension benefit. The indexation rate of individual accounts is automatically adjusted to the (notional) funding ratio of the pension fund by the adjustment rule of Goecke (2013). This automatic adjustment rule is the device for implementing IRS in our model. In contrast to the second-best strategy of Gollier (2008), the fund is fully-funded and does not rely on any external entity to implement IRS.

We analyze how the automatic adjustment rule stabilizes the funding ratio and the benefits of participants to attain IRS. Analytic expressions are derived for the funding ratio and benefits. It is shown that there is a trade-off between the stability of the funding ratio and that of benefits, and that this trade-off is controlled by the strength of the automatic adjustment rule. That is, benefits can be made more stable by increasing the volatility of the funding ratio, and vice versa. IRS can be attained by balancing this trade-off.

Automatic adjustment rules in pension systems have been not only studied in the literature (e.g., Cui et al. 2011; Chen et al. 2016; Bams et al. 2016; Donnelly 2017) but also applied to real pension systems in such countries as Sweden and the Netherlands (OECD 2021, Chapter 2). They are used for improving the sustainability of a pension fund and for providing stable benefits to participants (e.g., Settergren 2001; Barr and Diamond 2011). However, a formal analysis is missing for justifying such use of automatic adjustment rules in collective pension systems. Our analysis thus provides a first step in this regard.

The IRS-DC model has two parameters, one for the investment strategy and the other for the automatic adjustment rule. For optimizing these parameters, we define an expected utility maximization problem that involves the benefits of all the generations including those in the future, following Gollier (2008). As this optimization problem cannot be solved analytically, we solve it numerically using *Bayesian optimization*, a machine learning approach to optimizing a black-box function (e.g., Shahriari et al. 2016). As discussed later, the use of Bayesian optimization is our computational contribution, in line with the recent deployments of machine learning in the insurance literature (e.g., Hainaut 2018; Gabrielli 2020; Wüthrich 2020; Scognamiglio 2022; Schnürch and Korn 2022).

To answer the question above, our main finding is that IRS can improve the welfare of participants without borrowing and shareholders, *if* the financial market is volatile and the participants are risk-averse; IRS may not be welfare-improving if this condition is not satisfied. We compare the welfare of the IRS-DC plan participants and the welfare of the corresponding individual DC plan participants, where the latter uses the optimal life-cycle strategy (Merton 1971). Several different settings of the financial market and the risk aversion of participants are investigated, and the above finding is obtained.

The paper proceeds as follows. Section 2 introduces the IRS-DC pension model, which is analyzed in Section 3. Section 4 explains the expected utility maximization problem, how to solve it with Bayesian optimization, and the setup for simulations. Section 5 presents numerical analyses and Section 6 concludes. The appendix contains a short tutorial on Bayesian optimization and the proof of a main analytical result.

2. Pension Model

This section describes the IRS-DC pension model. The pension fund contains multiple overlapping generations, where there are always incoming and outgoing generations. Each




Figure 1: Schematic Illustration of the pension model. The horizontal axis indicates time t , and the vertical axis generation identifiers. For each generation, the black and white circles indicate the time points when that generation joins and retires from the pension fund, respectively; the orange line indicates the duration of the participation in the fund. For example, the generation i joins the fund at time $t = i - N$ and retires at $t = i$. For each time point, the blue box indicates the current participants of the pension fund, and black diamonds indicate those participants (only shown for the generations displayed in the figure). For example, at time $t = i - 1$, the generations $i, i + 1, \dots, i + N - 1$ are participating in the fund.

generation pays fixed contributions annually to the pension fund. Before explaining the details, we summarize below the key features of the pension fund:

- (i) The pension fund collectively invests the contributions from different generalizations (participants) in a financial market.
- (ii) Each participant maintains her account in the pension fund that records her accumulated pension rights.
- (iii) The growth rate of individual accounts is automatically adjusted based on the fund's investment performance and a notional funding ratio, so that intergenerational risk sharing is implemented.

Section 2.1 describes the IRS-DC pension model in detail. Section 2.2 compares it with related pension models in the literature.

2.1. Description of the IRS-DC Pension Model

Figure 1 provides a schematic illustration of our pension model. The pension fund covers N overlapping working generations in each operating year (e.g., $N = 40$). The pension fund is fully-funded. For simplicity, we assume that each generation consists of one hypothetical participant. Let $t \geq 0$ denote the time, with the unit being one year. We assume that the fund starts at time $t = 0$ with N initial generations.

2.1.1. Generation Identifier

We use an integer $i \in \mathbb{N}$ as the identity of the generation who retires at time $t = i$ (see Figure 1). Namely, the generation i joins the fund at time $t = i - N$ and leaves the fund at time $t = i$; thus, this generation is in the fund for N years. Using this notation, we

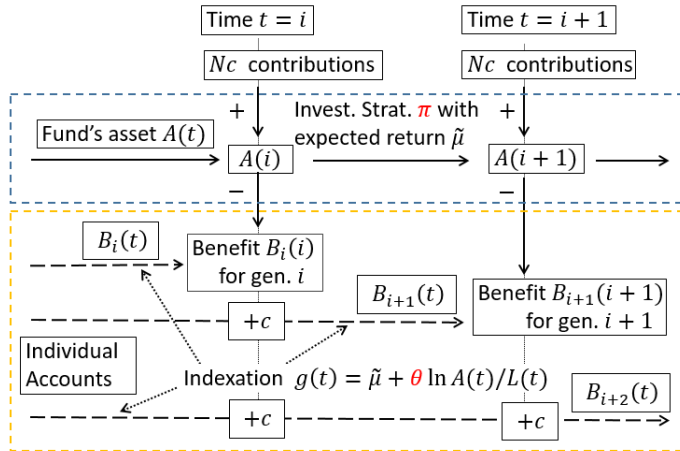


Figure 2: Schematic illustration of the cash flows in our pension model for two time points $t = i$ and $t = i + 1$. In the blue dotted box, $A(t)$ denotes the pension fund’s asset, whose dynamics follow (2.4) and (2.5). In the orange dotted box, $B_i(t)$, $B_{i+1}(t)$, and $B_{i+2}(t)$ denote the individual accounts of generations i , $i + 1$, and $i + 2$, respectively, whose dynamics follow (2.6) and (2.7). At time $t = i$, there is an incoming cash flow of Nc contributions from the N working generations, and an outgoing cash flow of the retirement benefit $B_i(i)$ for the generation i , who retires at time $t = i$; see (2.5). This retirement benefit is determined by the individual account $B_i(t)$, which accumulates the annual contribution c at the beginning of each year and develops continuously according to the indexation rate $g(t) = \tilde{\mu} + \theta \ln A(t)/L(t)$; see (2.7). On the other hand, the asset process $A(t)$ develops continuously according to the investment performance of the portfolio with the constant-mix strategy π ; see (2.4). Each “ $+c$ ” represents the accumulation of the annual contribution in the individual account: At time t , the amount c is added to each of $B_{i+1}(t)$ and $B_{i+2}(t)$.

can define the set of all the working generations in the fund at any time point $t \geq 0$ as

$$I_w(t) := \{[t] + 1, [t] + 2, \dots, [t] + N\},$$

where $[t]$ denotes the integer part of t (e.g., if $t = 3.4$ then $[t] = 3$; if $t = 3$ then $[t] = 3$).

2.1.2. Financial Market and Pension Asset Dynamics

We consider a financial market where there exist two investment opportunities: a risky asset (e.g., a stock) and a risk-free asset (e.g., a bond), denoted by $S(t)$ and $F(t)$, respectively. Specifically, we consider the Black-Scholes market, where $S(t)$ is driven by a diffusion process with constant drift $\mu > 0$ and volatility $\sigma > 0$, while $F(t)$ develops at a risk-free rate $r > 0$ such that $\mu > r$:

$$dS(t) = \mu S(t)dt + \sigma S(t)dZ(t), \quad S(0) = 1; \tag{2.1}$$

$$dF(t) = rF(t)dt, \quad F(0) = 1, \tag{2.2}$$

where $Z(t)$ is a standard Brownian motion under the real-world probability measure.

Let $A(t)$ denote the asset of the pension fund at time t , with $A(0) > 0$ being the initial asset. For simplicity, we assume that the fund invests a constant fraction $\pi \in [0, 1]$ of the pension asset $A(t)$ in the risky asset $S(t)$ and the rest $1 - \pi$ in the risk-free asset $F(t)$;

thus, we can write the dynamics of the pension asset as

$$dA(t) = \frac{\pi A(t)}{S(t)} dS(t) + \frac{(1-\pi)A(t)}{F(t)} dF(t), \quad \pi \in [0, 1]. \quad (2.3)$$

We assume that the fund is prohibited from borrowing ($\pi > 1$) and short selling ($\pi < 0$).

At the beginning of each year $t = 0, 1, 2, \dots$, each working generation pays a constant amount of contribution, $c > 0$, to the pension fund. Since there are N working generations, the fund thus receives a total of Nc contributions at the beginning of each year. At the same time, the fund pays a lump-sum benefit to the generation t , who retires at time t . With (2.3), the dynamics of the pension asset can thus be written as

$$dA(t) = A(t)(\pi(\mu - r) + r) dt + A(t)\pi\sigma dZ(t), \quad A(0) = A_0, \quad t > 0, \quad (2.4)$$

$$A(t)_+ = A(t) + Nc - B_t(t), \quad t = 0, 1, 2, \dots, \quad (2.5)$$

where $A(t)_+ := \lim_{\varepsilon \rightarrow +0} A(t + \varepsilon)$ denotes the right continuous limit, and $B_t(t)$ is the benefit paid to the generation t who has just retired and defined in (2.6)–(2.7) below (the double t notation of $B_t(t)$ is deliberate and its meaning will be clear shortly).

Therefore, the pension asset develops continuously over time $t > 0$, while there is a jump at each integer time $t = 0, 1, 2, \dots$ (i.e., at the beginning of each year) when there are incoming and outgoing cash flows of Nc and $B_t(t)$, respectively.

2.1.3. Individual Accounts and Retirement Benefits

Like pure DC and notional DC plans, each participant (generation) in our IRS-DC pension fund has her individual account, which keeps track of her pension rights. It records her annual contributions and grows according to the indexation rate (2.11) defined below. The terminal value of the account at the time of retirement becomes the lump-sum retirement benefit.

More formally, let $B_i(t)$ denote the individual account of generation $i \in \mathbb{N}$ at time $t \in [i - N, i]$, which starts from $B_i(i - N) = 0$ when this generation enters the fund at time $t = i - N$. Then we define its dynamics as

$$B_i(t)_+ = B_i(t) + c, \quad t = i - N, \quad i - N + 1, \dots, \quad i - 1, \quad (2.6)$$

$$dB_i(t) = B_i(t)g(t)dt, \quad i - N \leq t \leq i. \quad (2.7)$$

where $g(t)$ is the indexation rate at time t , defined in (2.11) below. Namely, the individual account $B_i(t)$ grows continuously from the entry time $t = i - N$ until retirement at $t = i$ according to the indexation rate $g(t)$ as (2.7), while the account accumulates the annual contribution c at the beginning of every year as (2.6). The account value at the time of retirement $t = i$, i.e., $B_i(i)$, is the retirement benefit of the generation i ; see (2.5).

2.1.4. Indexation Rate and Notional Liability

We now define the indexation rate $g(t)$ that determines the growth rate of individual accounts as in (2.7), which in turn affects the pension's asset dynamics in (2.5). To this end, following Bams et al. (2016) and Donnelly (2017), we first define a *notional liability* of the fund as

$$L(t) := \sum_{i \in I_w(t)} B_i(t), \quad t \geq 0. \quad (2.8)$$

That is, we define the notional liability $L(t)$ of the fund at time t as the sum of individual accounts $B_i(t)$ for the current working generations $i \in I_w(t)$. We call (2.8) *notional liability* because the fund uses it only for defining the indexation rate (see below).

If there is no cash flow in (2.5), the solution to the asset process (2.4) is given by

$$A(t) = A(0) \exp \left(\int_0^t \tilde{\mu} ds + \int_0^t \tilde{\sigma} dZ(s) \right), \quad (2.9)$$

where $\tilde{\mu} > 0$ and $\tilde{\sigma} > 0$ are constants defined as

$$\tilde{\mu} := \pi(\mu - r) + r - \frac{1}{2}\pi^2\sigma^2, \quad \tilde{\sigma} := \pi\sigma. \quad (2.10)$$

Following Goecke (2013, Eq. (5)),¹ we then define the indexation rate $g(t)$ as

$$g(t) = \tilde{\mu} + \theta \ln \left(\frac{A(t)}{L(t)} \right), \quad t \geq 0, \quad (2.11)$$

where $\theta > 0$ is a constant, $\tilde{\mu}$ is from (2.10), $A(t)$ is the pension fund's asset process (2.4) (2.5), and $L(t)$ is the notional liability (2.8).

2.1.5. The Role of the Indexation Rate

While we will present a more formal analysis in Section 3, we provide here an intuitive discussion of how the indexation rate (2.11) works. As defined in (2.7), the indexation rate $g(t)$ controls the growth rate of individual accounts $B_i(t)$. The first term $\tilde{\mu}$ in (2.11) is the expected annual log return using the same investment strategy π in the financial market *without* participating in the pension fund. In the second term, $A(t)/L(t)$ is the *notional funding ratio* that quantifies the balance between the asset $A(t)$ and the notional liability $L(t)$. We call the funding ratio $A(t)/L(t)$ *notional* because the liability $L(t)$ is notional; they are introduced only for defining the indexation rate $g(t)$. The second term adjusts the growth rate of individual accounts $B_i(t)$, and the parameter $\theta \geq 0$ specifies the strength of the adjustment.

If $\theta = 0$, then individual accounts grow *deterministically* at the rate $\tilde{\mu}$; therefore in this case, the retirement benefits are *ex-ante determined*, i.e., the pension plan becomes a defined-benefit plan, thus removing the investment risk of pension participants. However, it risks the sustainability of the pension fund, since the fund is of a defined-contribution type by design, and thus there is no way of adjusting the contributions when the fund is underfunding.

On the other hand, if one sets a large value of $\theta > 0$, then pension participants bear more investment risks to improve the sustainability of the pension fund. For example, suppose that the pension asset exceeds the notional liability, i.e., $A(t) > L(t) = \sum_{i \in I_w(t)} B_i(t)$. One can interpret this situation as that the fund yields a high return in the investment and thus there is a "surplus." Then the log notional funding ratio becomes positive, $\ln(A(t)/L(t)) > 0$, and the indexation rate $g(t)$ increases; therefore individual accounts $B_i(t)$ grow faster, reflecting the high investment return. On the other hand, if $A(t) < L(t) = \sum_{i \in I_w(t)} B_i(t)$, which happens when the fund yields a low return and thus there is a "deficit." In this case, we have $\ln(A(t)/L(t)) < 0$ and thus the indexation rate $g(t)$ decreases; therefore individual accounts $B_i(t)$ grow more slowly, reflecting the low investment return.

This argument implies that the adjustment parameter θ should be neither too small nor too large. One should choose θ appropriately to achieve a good trade-off between the risks of individual participants and the pension fund. We present a more formal analysis in Section 3.

2.2. Comparison with Related Pension Models

We compare the IRS-DC model with related pension models in the literature. Goecke (2013) studies the indexation rate (2.11) for the return smoothing in a self-financing pension plan. Goecke (2013)'s model consists only of one generation, and there exists no cash flow of contributions and payments. The earlier work by Baumann and Müller (2008) considers the indexation rate (2.11) where the risk-free rate r is used instead of $\tilde{\mu}$. Our model is a continuous-time version of the discrete-time model of Bams et al. (2016), which itself is an extension of the overlapping generations model of Gollier (2008). Bams et al. (2016) use the indexation rate in (2.11), but do not analytically study its use. Donnelly (2017) considers a funded collective DC pension plan. Donnelly (2017)'s model consists of *fixed* multiple overlapping generations, and there exists no new incoming generation. Donnelly (2017) uses an automatic adjustment rule similar to Goecke (2013)'s (and thus ours) but is different in its concrete form.

Cui et al. (2011) consider a funded DB-based hybrid pension system that can adjust both benefits and contributions. Their pension model does not have individual accounts, and thus there is no direct link between one's benefits and contributions. Their model is DB-based in this sense. They use automatic adjustment rules for contributions and retirement benefits based on the funding ratio. Chen et al. (2016) consider a hybrid pension plan as their second-pillar pension system, in which there exist individual accounts. They also use automatic adjustment rules for contributions and individual accounts' indexation rates. While the adjustment rules of Cui et al. (2011) and Chen et al. (2016) are conceptually similar to ours, they are different in their forms. For example, Chen et al. (2016) use the "tangent hyperbolic adjustment function," while our adjustment rule is based on the log notional funding ratio. Moreover, Cui et al. (2011) and Chen et al. (2016) define the liabilities in a DB manner, taking into account future retirement benefits, while we define our notional liability in a DC manner, i.e., as the sum of current individual account values. Again, our liability is notional since it is used only for defining the automatic adjustment rule.

Automatic adjustment mechanisms have been implemented in real pension systems; see OECD (2021, Chapter 2) for an overview. Notably, Sweden's first-pillar notional DC pension system uses an automatic adjustment rule for the indexation rate of individual accounts (Settergren 2001). This adjustment mechanism is based on a notional funding ratio² and is conceptually similar to other rules discussed here (See e.g., Hagen 2013, Eqs. (6.2) and (6.3)). The Swedish first-pillar notional DC pension system defines its notional liability essentially in the same way as (2.8) (Settergren 2001, Eq. (3)).³

3. Analysis

We present an analysis of the IRS-DC pension model in Section 2, focusing on the role of the indexation rate (2.11) for achieving IRS. In particular, we study how the adjustment parameter θ in the indexation rate impacts the notional funding ratio, which measures the stability of the pension fund, and the retirement benefits of individual participants.

In Section 3.1, we first study the dynamics of the log notional funding ratio. Based on this, we analyze the effects of the adjustment parameter on the dynamics of the notional funding ratio in Section 3.2, and on the retirement benefit of an individual participant in Section 3.3. In the latter, we obtain an analytic expression of the retirement benefit in terms of the log notional funding ratio and the adjustment parameter. Based on this expression, we compare the retirement benefits of the IRS-DC plan and the corresponding

pure DC plan in Section 3.4. This last analysis provides insights into how IRS works in the IRS-DC plan.

3.1. Dynamics of the Log Notional Funding Ratio

We start by analyzing the dynamics of the *log notional funding ratio* defined as

$$\rho(t) := \ln \left(\frac{A(t)}{L(t)} \right). \tag{3.1}$$

Goecke (2013, Proposition A.1) shows that $\rho(t)$ is an Ornstein-Uhlenbeck process under the assumption that there exists no cash flow. Baumann and Müller (2008, Section 3.1) obtain a similar result, but again assuming no cash flow. Since our model involves explicit cash flows as in (2.5), these earlier results are not directly applicable. Nevertheless, we show here that $\rho(t)$ in our model is also an Ornstein-Uhlenbeck process if the time t is between integer time points. (Recall that cash flows in our model occur only at integer time points; see (2.5)). This result, and intermediate derivations, are later used for deriving further results, so we present them here for completeness.

Let $t_0 \in \mathbb{N}$ be an arbitrary integer time point, which corresponds to the beginning of a year. Then by (2.4), with $A(t_0)$ being the initial value, the asset process $A(t)$ for $t_0 < t < t_0 + 1$ is written as

$$A(t) = A(t_0) \exp \left(\int_{t_0}^t \tilde{\mu} ds + \int_{t_0}^t \tilde{\sigma} dZ(s) \right).$$

Notice the difference from the previous expression (2.9), which starts from $t = 0$ and holds only under the assumption that there exists no cash flow. Similarly, by (2.7), (2.8) and (2.11), the notional liability $L(t)$ is given as

$$L(t) = L(t_0) \exp \left(\int_{t_0}^t \tilde{\mu} ds + \theta \int_{t_0}^t \rho(s) ds \right). \tag{3.2}$$

Then for $t_0 < t < t_0 + 1$, the log notional funding ratio $\rho(t)$ can be expanded as

$$\begin{aligned} \rho(t) &= \ln A(t) - \ln L(t) \\ &= \ln A(t_0) + \int_{t_0}^t \tilde{\mu} ds + \int_{t_0}^t \tilde{\sigma} dZ(s) - \ln L(t_0) - \int_{t_0}^t \tilde{\mu} ds - \theta \int_{t_0}^t \rho(s) ds \end{aligned} \tag{3.3}$$

$$= \rho(t_0) - \theta \int_{t_0}^t \rho(s) ds + \int_{t_0}^t \tilde{\sigma} dZ(s). \tag{3.4}$$

The last expression is obtained because the two identical terms $\int_{t_0}^t \tilde{\mu} ds$ in (3.3) are cancelled out. This is the result of the expected log return $\tilde{\mu}$ being used in defining the indexation rate (2.11), which in turn results in (3.2).

Equation (3.4) indicates that the log notional funding ratio $\rho(t)$ for $t_0 < t < t_0 + 1$ is an Ornstein-Uhlenbeck process with initial value $\rho(t_0)$ (e.g., Karatzas and Shreve 1998, p. 358), which can be written as

$$\rho(t) = e^{-\theta(t-t_0)} \rho(t_0)_+ + \tilde{\sigma} \int_{t_0}^t e^{-\theta(t-s)} dZ(s), \quad t_0 \in \mathbb{N}, \quad t_0 < t < t_0 + 1. \tag{3.5}$$

This expression shows that $\rho(t) = \ln(A(t)/L(t))$ is mean-reverting in the sense that, irrespective of the value of $\rho(t_0)$, it tends to 0 (in expectation) as t increases. In other words, the notional funding ratio $A(t)/L(t)$ tends to 1 as t increases.

3.2. Effects of the Adjustment Parameter on the Notional Funding Ratio

Based on the expression (3.5), we next study how the adjustment parameter θ affects the dynamics of the log notional funding ratio $\rho(t)$. The conditional expectation and conditional variance of $\rho(t)$ given $\rho(t_0)$ is, for $t_0 < t < t_0 + 1$,

$$\mathbb{E}[\rho(t) \mid \rho(t_0)_+] = e^{-\theta(t-t_0)} \rho(t_0)_+, \quad (3.6)$$

$$\mathbb{V}[\rho(t) \mid \rho(t_0)_+] = \tilde{\sigma}^2 \mathbb{E} \left[\left(\int_{t_0}^t e^{-\theta(t-s)} dZ(s) \right)^2 \right] = \tilde{\sigma}^2 \int_{t_0}^t e^{-2\theta(t-s)} ds. \quad (3.7)$$

where (3.7) follows from the Ito Isometry. From (3.6) and (3.7), we can make the following observations:

- If θ is large, the conditional expectation (3.6) approaches 0 more quickly as t increases, and the conditional variance (3.7) becomes smaller. In the limit $\theta \rightarrow \infty$, both (3.6) and (3.7) converge to 0 for $t_0 < t < t_0 + 1$. Therefore, as $\theta \rightarrow \infty$, $\rho(t)$ converges to 0 for $t_0 < t < t_0 + 1$.

- If θ is small, the conditional expectation (3.6) approaches 0 more slowly as t increases, and the conditional variance (3.7) becomes relatively larger. Indeed, as $\theta \rightarrow +0$, (3.6) converges to $\rho(t_0)_+$ and (3.7) converges to $\tilde{\sigma}^2(t - t_0)$. Therefore, in the limit $\theta \rightarrow +0$, the conditional expectation of $\rho(t)$ remains the same as the initial value $\rho(t_0)$ at time t_0 , and the conditional variance of $\rho(t)$ is determined by the variance of log returns $\tilde{\sigma}^2$ of the investment strategy π ; see (2.9) and (2.10).

These observations show how the adjustment parameter θ affects the notional funding ratio $A(t)/L(t)$ and thus the stability of the pension fund. A larger θ lets $A(t)/L(t)$ approach 1 more quickly and thus makes the fund more stable, while a smaller θ makes the fund more volatile. Recall that the value of θ determines how strong the adjustment in the indexation rate $g(t)$ works for the individual accounts $B_i(t)$; see (2.11). Therefore, a larger θ results in a stronger adjustment of the individual accounts $B_i(t)$, so that the notional liability $L(t) = \sum_{i \in I_w(t)} B_i(t)$ is adjusted more quickly to match the fund's asset $A(t)$; this is an intuitive explanation of how a large θ improves the stability of the pension fund.

While a larger θ may be beneficial for the fund's stability, it results in a stronger adjustment of the individual accounts $B_i(t)$, which may make the retirement benefits volatile. Therefore it is important to understand the effects of θ on the retirement benefits; we analyze this next.

3.3. Effects of the Adjustment Parameter on the Pension Benefits

We next study how the adjustment parameter θ affects the retirement benefit of each generation. To this end, we obtain an analytic expression of the retirement benefit in terms of the log notional funding ratio, as summarized in the following proposition. The proof can be found in Appendix A.

PROPOSITION 1. *Let $c > 0$ be the annual contribution, $\tilde{\mu} > 0$ and $\tilde{\sigma} > 0$ be the mean and the standard deviation of the annual log return in (2.10), $\rho(t) = \ln(A(t)/L(t))$ be the log notional funding ratio, and $\theta > 0$ be the adjustment parameter of the indexation rate*

$g(t)$ in (2.11). Then for generation $i \in \mathbb{N}$, the retirement benefit $B_i(i)$ is given by

$$B_i(i) = c \sum_{n=1}^N \exp \left(\underbrace{n\tilde{\mu}}_{(i)} + \underbrace{(1 - e^{-\theta}) \sum_{\ell=1}^n \rho(i - \ell)_+}_{(ii)} + \underbrace{\tilde{\sigma} \sum_{\ell=1}^n \int_{i-\ell}^{i-\ell+1} (1 - e^{-\theta(i-\ell+1-s)}) dZ(s)}_{(iii)} \right). \tag{3.8}$$

Proposition 1 enables studying the effects of the adjustment parameter θ on the retirement benefit $B_i(i)$ of the i -th generation, who retires at time $t = i$. The expression (3.8) consists of N terms, in which each term is indexed by $n = 1, \dots, N$. (Recall that N is the total number of years each generation contributes to the fund). One can understand the n -th term in (3.8) as corresponding to the contribution c made at time $t = i - n$, i.e., n years before the retirement at time $t = i$.

We can make the following observations for the exponent of the n -th term in (3.8):

- The term (i) corresponds to the deterministic growth term $\tilde{\mu}$ in the indexation rate $g(t)$; see (2.11).

- The term (ii) represents the effects of the fund’s “surplus” or “deficit” in the last n years before the retirement. One can understand that there is a “surplus” if $\sum_{\ell=1}^n \rho(i - \ell)_+ > 0$; in this case the retirement benefit increases accordingly, as a redistribution of the surplus. On the other hand, there is a “deficit” if $\sum_{\ell=1}^n \rho(i - \ell)_+ < 0$, and the retirement benefit decreases accordingly; one can understand this as risk sharing to make the fund sustainable. The adjustment parameter θ determines the strength of the effects of this term, as we have $\lim_{\theta \rightarrow +0} (1 - e^{-\theta}) = 0$ and $\lim_{\theta \rightarrow \infty} (1 - e^{-\theta}) = 1$.

- The term (iii) shows the effects of the volatility of the fund’s investment in the last n years before the retirement; recall the definition of $\tilde{\sigma} = \pi\sigma$ in (2.10). The adjustment parameter θ controls the influence of this volatility, as we have $\lim_{\theta \rightarrow +0} (iii) = 0$ and $\lim_{\theta \rightarrow \infty} (iii) = \tilde{\sigma}$.

From these observations, one can understand that the adjustment parameter θ determines how strongly the retirement benefit is linked to the fund’s actual investment performance. For a larger θ , the terms (ii) and (iii) become more significant, and the retirement benefit is more directly influenced by the fund’s investment performance. For a smaller θ , the terms (ii) and (iii) become less significant, and the retirement benefit is determined mainly by the deterministic growth term (i). This asymptotic analysis supports the informal discussion in Section 2.1.5 on the mechanism of the indexation rate.

One may conclude that a smaller θ may be more beneficial for individual participants, because it makes the retirement benefits less volatile. However, as discussed in Section 3.2, a smaller θ makes the fund’s operation more volatile, and thus a larger θ is more desirable for the fund’s sustainability. Therefore, θ should be neither too small nor too large. We will discuss how to select the adjustment parameter θ (and the investment strategy π) in Section 4.

3.4. Effects of Intergenerational Risk Sharing

Lastly, we discuss the effects of IRS, by comparing the pension benefits of the IRS-DC plan and the corresponding pure DC plan. Because our focus is to understand how IRS

works, we assume here that the pure DC plan uses the same investment strategy π as the IRS-DC plan. (Note that, in our numerical analysis in Section 5, we consider this setting as well as the setting where the pure DC plan uses the optimal investment strategy.)

Consider two hypothetical individuals from generation $i \in \mathbb{N}$, who retire in the year i . One individual participates in the IRS-DC plan, and receives the retirement benefit (3.9). The other participates in the pure DC plan using the investment strategy π , and receives the retirement benefit denoted by $A_i(i)$. It is easy to see that $A_i(i)$ is given by

$$A_i(i) = c \sum_{n=1}^N \exp \left(\underbrace{n\tilde{\mu}}_{(i')} + \underbrace{\tilde{\sigma} \int_{i-n}^i dZ(s)}_{(ii')} \right), \quad (3.9)$$

By comparing (3.8) and (3.9), we can make the following observations:

- The term (i) in (3.8) and the term (i') in (3.9) are the same.
- The term (ii) in (3.8), which represents the effects of the fund's surplus or deficit, does not exist in (3.9). This is reasonable, because there is no IRS in the pure DC plan.
- The term (iii) in (3.8), which shows the influence of the volatility of the investment, corresponds to the term (ii') in (3.9). However, one can see that the volatility term (iii) in (3.8) is smaller than the term (ii') in (3.9). The smaller volatility in (3.8) is the result of IRS, and controlled by the adjustment parameter θ .

This comparison describes how IRS works in the IRS-DC plan: *IRS reduces the volatility of investment returns (term (iii) in (3.8)), by letting the individuals share the fund's surplus or deficit (term (ii) in (3.8))*. This effect of IRS is particularly important for protecting individual participants when the market is turbulent. Our numerical analysis in Section 5 shows that IRS is beneficial in this way.

4. Optimizing the Investment Strategy and Adjustment Parameter

We describe how to optimize the parameters of the IRS-DC pension model, namely the investment strategy π and the adjustment parameter θ , so as to maximize the welfare of pension participants. In Section 4.1, we first introduce an expected utility maximization problem that involves the welfare of all the generations including those from the future. Since there is no analytical solution for this maximization problem, we next explain how to solve it numerically using *Bayesian optimization* in Section 4.2. We then describe the setting of simulations in Section 4.3, which will be used later in our numerical analysis.

4.1. Expected Utility Maximization Problem

We consider a hypothetical social planner (fund manager) who decides the investment strategy π and the adjustment parameter θ for the welfare of all the generations. To define the utility of this social planner, let $U_\gamma : (0, \infty) \rightarrow (-\infty, \infty)$ be the constant relative risk aversion (CRRA) utility function:

$$U_\gamma(x) := \frac{x^{1-\gamma}}{1-\gamma}, \quad \gamma > 0, \quad (\gamma \neq 1), \quad (4.1)$$

where $\gamma > 0$ is the level of relative risk aversion. We then define the utility of the social planner as the sum of discounted utilities of the retirement benefits for all the generations:

$$\sum_{i=1}^{\infty} \beta^i U_{\gamma}(B_i(i)), \tag{4.2}$$

where $\beta > 0$ is a discounting factor and $B_i(i)$ is the retirement benefit of the i -th generation who retires at time $t = i$; see Figure 2 and (2.6) (2.7).

Lastly, we define our expected utility maximization problem as

$$\max_{\pi, \theta} \mathbb{E} \left[\sum_{i=1}^{\infty} \beta^i U_{\gamma}(B_i(i)) \right] \quad \text{subject to} \quad 0 \leq \pi \leq 1, \quad \theta > 0, \tag{4.3}$$

where the expectation is with respect to the retirement benefits $B_i(i)$ for all generations $i \in \mathbb{N}$. Recall that $B_i(i)$ are path-dependent and depend on the investment strategy π and the adjustment parameter θ .

We numerically solve the maximization problem (4.3), since neither the expected utility in (4.3) nor the solution for π and θ are available in closed form. We approximate the expected utility in (4.3) by Monte Carlo simulations, and optimize π and θ using *Bayesian optimization*, as explained next.

4.2. Bayesian Optimization for Expected Utility Maximization

We briefly explain here how we use Bayesian optimization (BO) for solving the expected utility maximization problem (4.3). For details, see Appendix B and references therein. BO is a modern machine learning approach for globally optimizing a black-box objective function, and has been shown to be more efficient than traditional approaches such as grid search (Shahriari et al. 2016). It has been widely used in applications where the objective function is computationally expensive to evaluate, such as the optimization of hyper parameters of a large-scale AI model (Snoek et al. 2012). The current work is the first attempt to apply BO in optimizing a pension system.

The objective function in (4.3) takes π and θ as an input and outputs the expected utility:

$$(\pi, \theta) \mapsto f(\pi, \theta) := \mathbb{E} \left[\sum_{i=1}^{\infty} \beta^i U_{\gamma}(B_i(i)) \right],$$

where we note again that $B_i(i)$ depends on π and θ . The key idea of BO is to “learn” the landscape of the objective function $f(\pi, \theta)$ while searching for π and θ that maximizes the objective function. BO first evaluates the function values $f(\pi, \theta)$ for some initial candidates of π and θ , and obtains a rough estimate for the landscape of $f(\pi, \theta)$. In the next step, BO finds π and θ such that the function value $f(\pi, \theta)$ and its *uncertainty* are both high, so as to balance the so-called *exploitation-exploration trade-off*. BO then evaluates $f(\pi, \theta)$ for these π and θ , and updates the estimate of the landscape of $f(\pi, \theta)$. BO iterates this learning-optimization procedure. Estimates of the maximizers, $(\pi^*, \theta^*) = \arg \max f(\pi, \theta)$, are obtained after a sufficient number of iterations (Bull 2011).

The above procedure is called “Bayesian” because the learning of the objective function is done by a Bayesian nonparametric method (Rasmussen and Williams 2006). The Bayesian method is used because it can yield both an estimate of the landscape as well as its uncertainties, which are crucial for the exploitation-exploration trade-off and for gaining the optimization efficiency. For implementation, we use the R package *mlrMBO* (Bischl et al. 2017) in our numerical analysis.

4.3. Simulation Setting

We explain here how we approximate the expected utility in (4.3) by Monte Carlo simulations, which is necessary for applying Bayesian optimization. Moreover, we describe the problem setting for our numerical analysis in the next section. First of all, we set the number of working generations as $N = 40$, the discounting factor in (4.3) as $\beta = 0.98$, and the annual contribution as $c = 1$.

4.3.1. Financial Market

We consider three settings for the financial market that represent different market risks, to investigate when IRS in the IRS-DC model works most effectively. The market price of risk, a.k.a the Sharpe ratio, is defined by

$$\lambda := \frac{\mu - r}{\sigma} \quad (4.4)$$

where μ and σ are the rate and volatility of the risky asset, and r is the rate of the risk-free asset; see Section 2.1.2. The Sharpe ratio represents a baseline return for taking a risk in the market. It can be expected that the effectiveness of IRS depends on the market volatility σ , as discussed in Section 3. Hence, we consider the following three settings for the financial market, with different levels of the Sharpe ratio and the volatility:

$$\begin{aligned} \text{Market 1: } \lambda_1 &= \frac{\mu_1 - r_1}{\sigma_1} = \frac{0.065 - 0.02}{0.15} = 0.3; \\ \text{Market 2: } \lambda_2 &= \frac{\mu_2 - r_2}{\sigma_2} = \frac{0.065 - 0.01}{0.25} = 0.22; \\ \text{Market 3: } \lambda_3 &= \frac{\mu_3 - r_3}{\sigma_3} = \frac{0.065 - 0.01}{0.5} = 0.11. \end{aligned}$$

We may refer to Markets 1, 2, and 3 as M1, M2, and M3 for brevity. We call M1, M2, and M3 the markets with *high*, *intermediate*, and *low* Sharpe ratios, respectively.

4.3.2. Risk Aversion of the Social Planner

The relative risk aversion γ in the the CRRA utility function (4.1) represents the social planner's risk attitude: the social planner becomes more risk-averse if γ is larger. To study the impacts of γ on the optimal investment strategy π and adjustment parameter θ , we consider three settings: $\gamma = 3, 5, 10$.

4.3.3. Entry Cohorts

The generations with indicators $i = 1, \dots, 40$ are those who participate in the IRS-DC plan at time $t = 0$, and are called *entry cohorts*. For $t < 0$, i.e., before participating in the IRS-DC plan, we assume that the entry cohorts participate in a pure DC plan that applies the optimal *life-cycle investment strategy*, following Gollier (2008). Namely, the pure DC plan invests a constant fraction of the total wealth (including the discounted future incomes) into the risky asset, and the rest of the total wealth into the risk-free asset. To be more precise, for each generation $i = 1, \dots, 40$, let $B_i(t)$ be the individual account and $Y_i(t)$ be the net present value at time t of all the future contributions; then the optimal fraction to be invested in the risky asset is given by

$$\pi_i^{\text{ind}}(t) := \pi^c \frac{B_i(t) + Y_i(t)}{B_i(t)}, \quad (4.5)$$

where π^c is the so-called *Merton constant* (Merton 1971) defined as

$$\pi^c := \frac{\lambda}{\gamma\sigma}, \quad (4.6)$$




Figure 3: Examples of the optimal life-cycle investment strategy (4.5) for an individual who participate in the pure DC plan for her entire working period of 40 years. The left and right figures are for the Markets 1 and 3, respectively. We consider the three settings of relative risk aversion $\gamma = 3, 5, 10$ for the individual. The vertical axis represents (4.5) and the horizontal axis the time.

where λ is the Sharpe ratio in (4.4). Note that $Y_i(t)$ can be calculated straightforwardly here, as we do not consider the interest risk.

For an individual with the CRRA utility function, the life-cycle investment strategy (4.5) provides the highest expected utility (Merton 1971; Gollier 2008). Hence, the life-cycle investment strategy and its variants have been popular choices for pure DC plans (e.g., Booth and Yakoubov 2000; Haberman and Vigna 2002). Note that, when an individual is young, the discounted future income $Y_i(t)$ is much higher than her current wealth in the account $B_i(t)$, and thus π^{ind} in (4.5) is much larger than 1. Therefore, the life-cycle investment strategy (4.5) implies a high-leverage (i.e., borrowing) strategy when the individual is young (see Figure 3).

We also use the life-cycle investment strategy in our numerical analysis in Section 5.5 to make a comparison between the IRS-DC and the optimal pure DC plans.

4.3.4. Euler-Maruyama Approximation

For simulating the dynamics of the asset process (2.4), we use the Euler-Maruyama approximation. For a finite time horizon $T > 0$, we divide the interval $[0, T]$ into n equal time intervals:

$$[0, T] = [0, \Delta, 2\Delta, \dots, n\Delta],$$

where we set the step size as $\Delta = 1/12$, which corresponds to one month. Then we simulate the asset process as

$$\begin{aligned} A(t + \Delta) &= A(t) + (\pi(\mu - r) + r)\Delta + \pi\sigma\sqrt{\Delta}Z, \\ A(t)_+ &= A(t) + 40c - B_t(t), \quad \text{if } t \in \mathbb{N} \cup \{0\}, \end{aligned}$$

where Z is a standard normal random variable. The dynamics of the asset process in a pure DC plan is simulated in a similar way.

The step size $\Delta = 1/12$ implies that the indexation rate (2.11) is adjusted monthly according to the notional funding ratio $A(t)/L(t)$. This monthly update is more frequent than the cash flows of the fund, which are annual. This setting reflects the fact that the market values of individual accounts usually vary more frequently than cash flows in reality. We assume the initial value of the notional funding ratio is one: $A(0)/L(0) = 1$.

4.3.5. *Time Horizon T*

While the expected utility in (4.3) involves the infinite horizon, it is intractable for simulations and we need to use a finite horizon T . In our numerical analysis, we set the horizon as $T = 80$ years. For approximating the expected utility in (4.3), we then simulate 10,000 paths for the asset process until the horizon T and compute the Monte Carlo average. Note that, in this setting, the generations $i = 41, \dots, 80$ are those who spend their entire working periods in the IRS-DC plan.

4.3.6. *Upper Bound of the Adjustment Parameter θ*

While the adjustment parameter θ can take an arbitrarily large value in theory, for numerical optimization of θ we need to set its upper bound. We set the range of θ as $0 < \theta \leq 1$ in our numerical analysis.

5. Numerical Analysis

This section presents our numerical analysis of the IRS-DC pension model. In Section 5.1, we first discuss the optimal investment strategy and adjustment parameter obtained by Bayesian optimization. In Section 5.2, we then study the dynamics of the notional funding ratio, and discuss how the adjustment parameter θ affects its stability. The stability of the notional funding ratio can be understood as the stability of the pension fund's operation. In Sections 5.3, 5.4, and 5.5, we focus on the individual accounts in the IRS-DC fund. We first study the dynamics of individual accounts in Section 5.3, and then the distribution of retirement benefits in Section 5.4. Lastly, we study the welfare of pension participants in Section 5.5.

5.1. *Optimal Investment Strategy and Adjustment Parameter*

As explained Section 4.3, we consider 9 different settings for the numerical analysis, resulting from 3 different values for the relative risk aversion ($\gamma = 3, 5, 10$) of the social planner and 3 different values for the Sharpe ratio ($\lambda = 0.3, 0.22, 0.11$) of the financial market. In each setting, we find the optimal investment strategy π and the adjustment parameter θ by Bayesian optimization, as described in Section 4. We report the resulting optimal values of π and θ in Table 1. For comparison, we also report the Merton constant (4.6) for each setting in Table 1, which will be used in the experiment in Section 5.5.

For each value of the Sharpe ratio λ , the optimal π^* and θ^* tend to decrease as the risk aversion γ increases (with the exception of the case $\lambda = 0.3$, where θ^* remains 1). Regarding the optimal investment strategy π^* , this tendency can be anticipated by the same tendency in the Merton constant (4.6), which is inversely proportional to the risk aversion γ . Regarding the optimal adjustment parameter θ^* , this tendency can be expected from the analysis in Section 3.3, where it is shown that a smaller adjustment parameter θ lowers the volatility of the retirement benefits; thus a higher risk aversion γ leads to smaller θ^* .

For each value of the risk aversion γ , the optimal π^* and θ^* tend to be smaller as the Sharpe ratio λ becomes smaller (for the cases $\gamma = 3, 5$ and $\lambda = 0.22, 0.11$, the adjustment parameter θ^* is comparably small). One can understand this tendency in a similar way as the discussion in the above paragraph, since the Sharpe ratio represents the market price of risk. Note that θ^* is extremely small for $\gamma = 10$ and $\lambda = 0.01$, which implies the IRS-DC plan becomes similar to a DB plan, as discussed in Section 3.3; in this case π^* is also very small, meaning that the asset is mainly invested in the risk-free asset.

Table 1: Optimal investment strategy π^* and adjustment parameter θ^* obtained by Bayesian optimization for the 9 different π^* settings, resulting from the 3 different values of the Sharpe ratio λ and the 3 different values of the risk aversion γ . For comparison, the Merton constant is shown for each setting.

$\lambda = 0.3$	$\gamma = 3$	$\gamma = 5$	$\gamma = 10$
π^* :	0.832	0.519	0.267
θ^* :	1	1	1
Merton constant:	0.667	0.4	0.2
$\lambda = 0.22$:	$\gamma = 3$	$\gamma = 5$	$\gamma = 10$
π^* :	0.479	0.334	0.124
θ^* :	0.0651	0.0535	0.0237
Merton constant:	0.293	0.176	0.088
$\lambda = 0.11$:	$\gamma = 3$	$\gamma = 5$	$\gamma = 10$
π^* :	0.131	0.06	0.0544
θ^* :	0.0835	0.072	0.0000493
Merton constant:	0.073	0.044	0.022




Figure 4: The average and standard deviation (calculated from the 10,000 simulations) of the notional funding ratio process $A(t)/L(t)$ for the three market settings with risk aversion $\gamma = 10$.

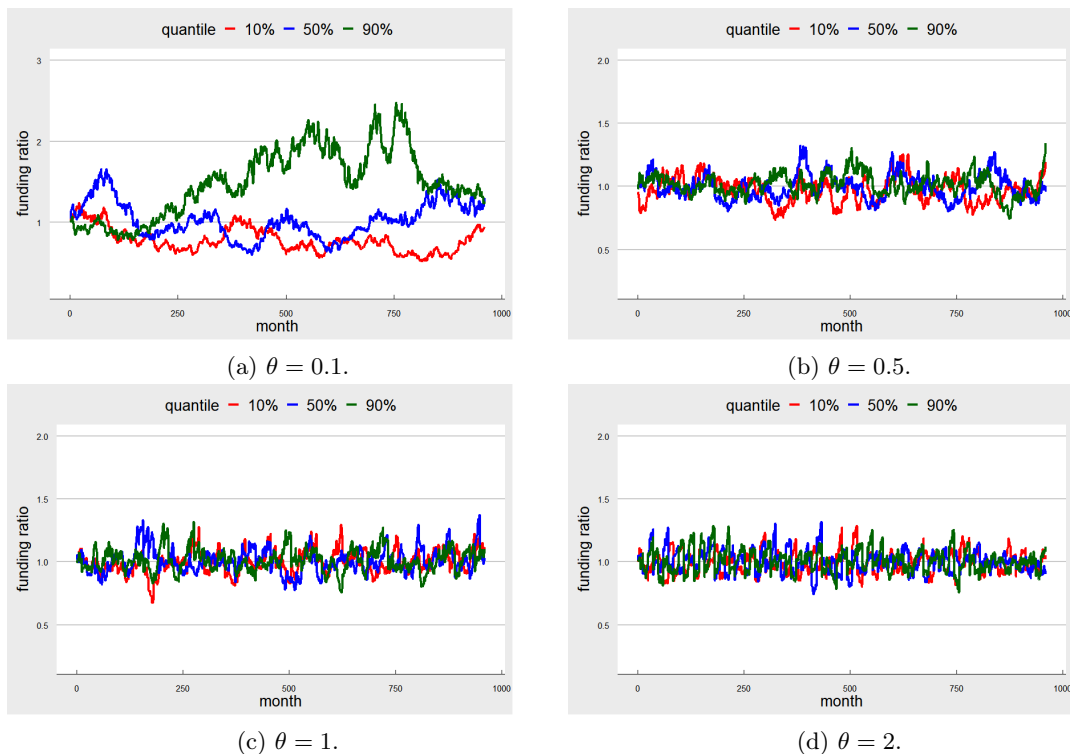


Figure 5: Notional funding ratio processes for “representative” scenarios in Market 1 with $\pi = 0.6$ and $\gamma = 3$.

5.2. Notional Funding-Ratio Process

We next study the dynamics of the notional funding ratio $A(t)/L(t)$. For each setting, we use the optimal parameters in Table 1 to obtain 10,000 simulation paths. Figure 4 shows the average and standard deviation of the notional funding ratio process (calculated from the 10,000 simulations) for 80 years, for the risk aversion $\gamma = 10$ and for the three markets M1 ($\lambda = 0.3$), M2 ($\lambda = 0.22$) and M3 ($\lambda = 0.11$); results for the other values of γ are presented in Appendix D.1. The zigzag shape of the notional funding ratio is because of the discrete cash flows at integer time points, which correspond to the beginnings of different years. For all the market settings, the average of the notional funding ratio process is stable and close to 1. On the other hand, the standard deviation significantly differs for the three settings. The analysis in Section 3.2 shows that a smaller adjustment parameter θ makes the notional funding ratio process more volatile. The smallest adjustment parameter θ^* is used in the setting M3 among the three settings of Figure 4 (see Table 1), which has the largest standard deviation; therefore Figure 4 is consistent with the analysis in Section 3.2. The same observation can be made for the additional results for $\gamma = 3, 5$ in Appendix D.1, where M2 has the smallest θ^* and the largest standard deviation.

To further investigate the effects of the adjustment parameter θ on the notional funding ratio process, we perform the following experiment, the results of which are shown in Figure 5. In Market 1 with $\gamma = 3$, we use the fixed investment strategy $\pi = 0.6$, and run 10,000 simulations for each of different values of the adjustment parameter $\theta \in \{0.1, 0.5, 1, 2\}$. (Note that here, as in theory, we allow for θ to take a value larger than

1). For each setting, we pick up the three scenarios from the 10,000 simulations that correspond to the top 10%, 50%, and 90% values of the utilities of the social planner (see (4.2)), and plot the corresponding notional funding ratio processes; these three scenarios can be interpreted as representing “good”, “medium” and “bad” realizations of the financial market. Figure 5 shows that, as θ increases, the variations among the three notional funding ratios become smaller. This result provides a further support for the the analysis in Section 3.2 that a larger adjustment parameter θ makes the notional funding ratio process more stable.

We have discussed how the adjustment parameter θ affects the stability of the notional funding ratio process, which can be understood as the stability of the pension fund’s operation. We next turn our attention to the stability of individual accounts in the IRS-DC fund.

5.3. Dynamics of Individual Accounts

We next study the dynamics of individual accounts in the IRS-DC plan. We focus on the generations $i = 41, \dots, 80$, as these are the generations who spend their entire working periods in the IRS-DC plan; see Section 4.3.5. As for the analysis in Section 3.4, to study the effects of IRS, each individual account in the IRS-DC plan is compared with the corresponding account in a pure DC plan that uses the same investment strategy π^* in Table 1. (A pure DC plan using the optimal life-cycle investment strategy is compared in Section 5.5.) Since IRS is absent, the dynamics of an individual account in the pure DC plan is given by the asset process yielding (3.9).

Figure 6 shows arbitrarily chosen paths of the individual accounts from the generation $i = 41$ in the IRS-DC and pure DC plans, for the three market settings and risk aversion $\gamma = 10$. For Market 1, for which the adjustment parameter θ^* is large (see Table 1), the paths of the IRS-DC and DC accounts are similar. In contrast, for Markets 2 and 3, for which the adjustment parameter θ^* is smaller, the IRS-DC account accumulates more stably than the pure DC account. This observation is consistent with the discussions in Sections 2.1.5, 3.3 and 3.4, where it is argued that a smaller adjustment parameter θ reduces the volatility of an IRS-DC account as a result of IRS.

We next quantify the effects of IRS on stabilizing the accumulation of an IRS-DC account. To this end, we calculate the *increment-ratio-based roughness (IR roughness)* (Bardet et al. 2011), a measure of roughness/smoothness of a stochastic process, for each path of the IRS-DC and pure DC accounts. The IR roughness takes a value between 0 and 1, and a larger value indicates that the path is smoother; see Appendix C for details. Table 2 shows the average of the IR roughness over the 10,000 simulations for each of the IRS-DC and pure DC accounts from the generation $i = 41$. For all the 9 settings considered, the IRS-DC account has a larger IR roughness than the pure DC account, which implies that the IRS-DC account is smoother. Therefore, IRS makes the accumulation of the IRS-DC account more stable than the pure DC account (Recall that the only difference between the IRS-DC and DC plans here is the existence of IRS).

5.4. Distribution of Retirement Benefits

We next study how IRS affects the distribution of retirement benefits. Figure 7 shows the histograms of retirement benefits (over the 10,000 simulations) of the IRS-DC and pure DC accounts for the generation $i = 41$, for the three market settings and risk aversion $\gamma = 10$. (Results for $\gamma = 3, 5$ are shown in Appendix D.2.) For Market 1, for which the adjustment parameter θ^* of the IRS-DC plan is large, the histograms of the IRS-DC and pure DC retirement benefits are almost identical. On the other hand, for

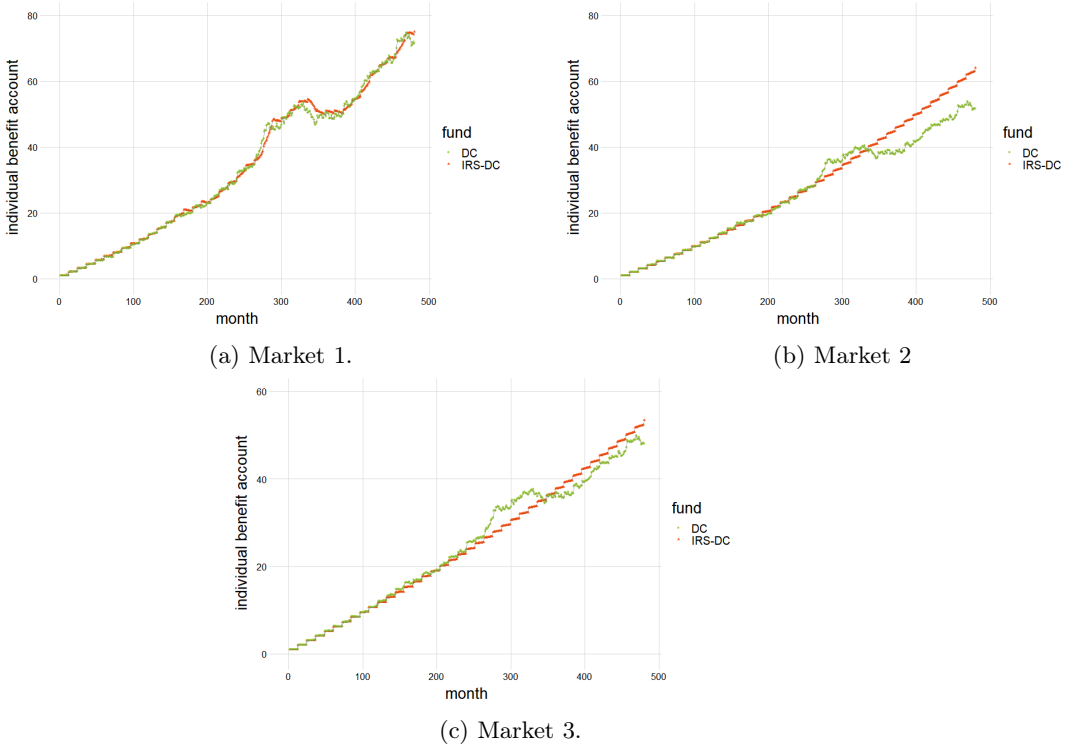


Figure 6: Examples of paths of the IRS-DC and pure DC accounts for the generation $i = 41$ in the three market settings with risk aversion $\gamma = 10$.

Table 2: Average IR roughness over 10,000 simulations calculated for each of the IR-DC and pure DC accounts, for the three market settings and the three levels of risk aversion.

		IR roughness	$\gamma = 3$	$\gamma = 5$	$\gamma = 10$
Market 1	IRS-DC		0.937	0.944	0.959
	DC		0.732	0.739	0.754
Market 2	IRS-DC		0.993	0.996	1.000
	DC		0.731	0.735	0.752
Market 3	IRS-DC		0.991	0.998	1.000
	DC		0.737	0.751	0.753

Markets 2 and 3, for which the adjustment parameter is smaller, the volatility of the IRS-DC benefits is smaller than the pure DC benefits. In particular, for Market 3, for which the adjustment parameter is close to 0, the volatility of the IRS-DC benefits is very small. These observations support the analysis in Sections 3.3 and 3.4 that a smaller adjustment parameter θ makes the retirement benefits less volatile by IRS. Moreover, our result is consistent with similar observations made by Bams et al. (2016) and Donnelly (2017) that a collective DC scheme can reduce the volatility of retirement benefits.

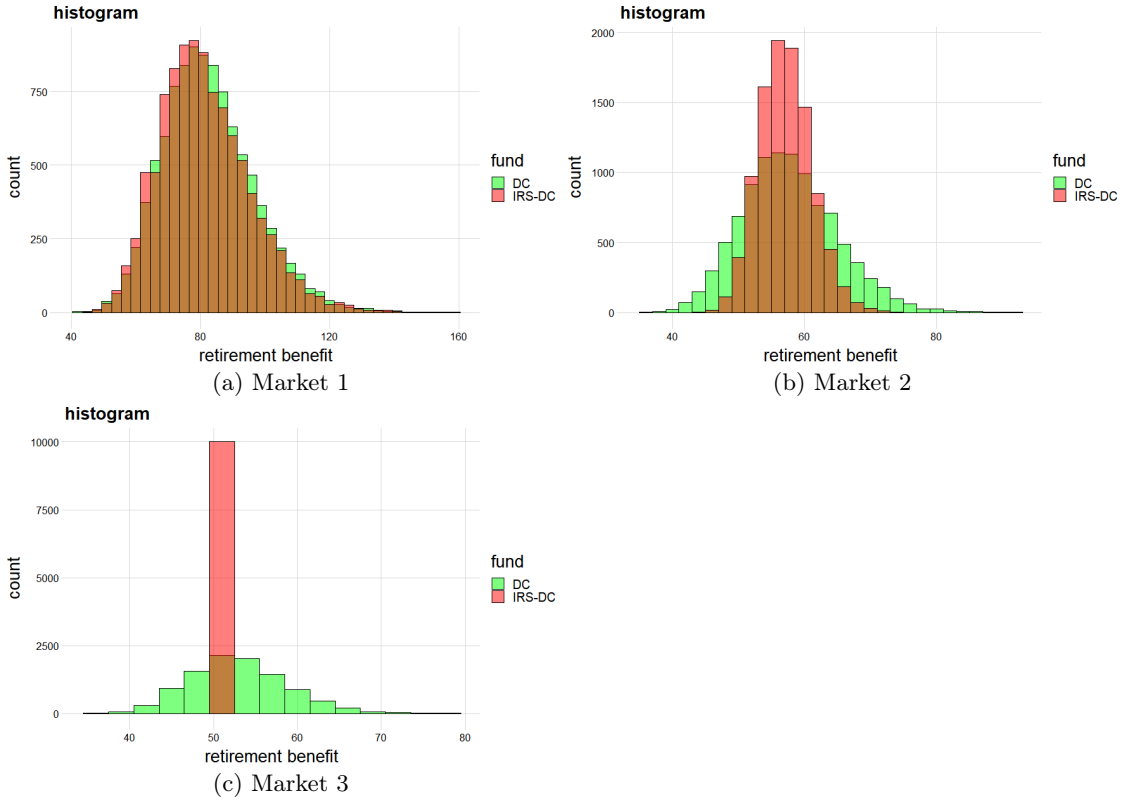


Figure 7: Histograms of the retirement benefits (obtained from the 10,000 simulations) of the IRS-DC and pure DC accounts for the generation $i = 41$, for the three market settings and risk aversion $\gamma = 10$. In each figure, the red and green histograms are those of the IRS-DC and pure DC benefits, respectively; the brown part shows the overlap between the two histograms.

5.5. Welfare of Participants

Lastly, we study how IRS can improve the welfare of the IRS-DC plan participants in terms of their expected utilities. To this end, we make a comparison with a pure DC plan that uses the optimal life-cycle investment strategy in (4.5), which yields the highest expected utility for an individual investor.

For simplicity, we assume that each participant in the IRS-DC plan has the same CRRA utility U_γ in (4.1) as the social planner. Similarly, to make a comparison straightforward, we assume that each participant in the pure DC plan has the same CRRA utility. To measure the welfare, we calculate the certainty equivalent (CE) for each participant. That is, for an IRS-DC participant from the generation i with retirement benefit $B_i(i)$, the CE is defined as the quantity $CE_i^{(\text{IRS-DC})} > 0$ satisfying

$$U_\gamma(CE_i^{(\text{IRS-DC})}) = \mathbb{E}[U_\gamma(B_i(i))], \quad (5.1)$$

where we approximate the expectation in the right hand side by the empirical average of 10,000 realizations of $B_i(i)$. The CE of each participant in the pure DC plan is calculated similarly. Note that, since the utility function is strictly concave, the expected utility is monotonically increasing with respect to the CE; a higher CE implies a higher expected

utility. We calculate the CEs of the participants in the IRS-DC plan and the pure DC plans for the generations $i = 41, \dots, 80$.

Figure 8 shows the CEs of the IRS-DC and pure DC participants for the generations $i = 41, \dots, 80$, for the three market settings and risk aversion $\gamma = 10$. (Results for $\gamma = 3, 5$ are shown in Appendix D.3.) For Market 1, where the Sharpe ratio is high, the pure DC participants obtain higher welfare than the IRS-DC participants. On the other hand, for Markets 2 and 3, where the Sharpe ratio is lower, the IRS-DC participants obtain higher welfare than the pure DC participants. This observation indicates that the IRS-DC plan can provide higher welfare than the optimal DC plan *when the market is more volatile* (in the sense of having a lower Share ratio). Therefore, IRS is expected to be particularly advantageous in protecting individual participants when the market is turbulent (e.g., when there is an economic shock).

While it has been generally known in the literature that IRS is welfare-improving, there are a few key differences in our contribution. To explain this, we make a comparison with closely related works. Gollier (2008) shows that IRS is welfare-improving over the optimal life-cycle investment strategy, but his analysis is based on the assumption that the pension fund can perform *borrowing* for investment (i.e., the investment strategy π can be larger 1); this assumption is not realistic for pension funds in reality. Moreover, his second-best strategy assumes the existence of a “shareholder” that helps finance the pension fund. Our result above shows that IRS can be welfare-improving even when borrowing is prohibited for the pension fund (i.e., $0 < \pi < 1$) and without a shareholder.

Cui et al. (2011) show that a hybrid pension plan with IRS can provide higher welfare than a pure DC plan with an “optimal” investment strategy. However, their “optimal” individual invest strategy is not allowed to perform borrowing, and therefore it is less optimal than the optimal life-cycle strategy (4.5), which performs borrowing. Moreover, Cui et al. (2011) optimize the parameters of the pension fund so as to maximize the expected utility of *one specific entry cohort*, not all the generations; they then compare this entry cohort’s welfare with a pure DC plan participant’s welfare. This way of optimizing the pension system is not appropriate as it ignores the other generations’ welfare. On the other hand, we show that the IRS-DC plan, which optimizes for all the generations’ utilities as in (4.3), can improve the welfare over the optimal life-cycle investment strategy, when the market is volatile.

Bams et al. (2016) consider a similar pension model as ours, but they do not show that their model can provide higher welfare than individual DC plans. Similarly, Donnelly (2017) considers a related collective DC plan, but does not compare it with individual DC plans. Chen et al. (2016) study a three-pillar model in which the second pillar is a collective DC, DB, or hybrid pension plan, and make a comparison with the corresponding three-pillar model with the second pillar being an individual DC plan. While they show that the former yields higher welfare than the latter, they assume that the both plans use the *same* investment strategy, with the fraction invested in the risky asset being $\pi = 0.5$; therefore their individual DC plan is not optimal. Different from these previous works, we make a comparison with the optimal life-cycle investment strategy. By doing so, we show that the volatility of the financial market is a key factor that determines whether IRS is welfare-improving over the optimal life-cycle investment strategy.

6. Concluding Remarks

We have shown that a fully funded collective DC pension system with intergenerational risk sharing (IRS) can improve the welfare of individual participants, as compared with individual DC benchmarks using the optimal life-cycle investment strategy, when the

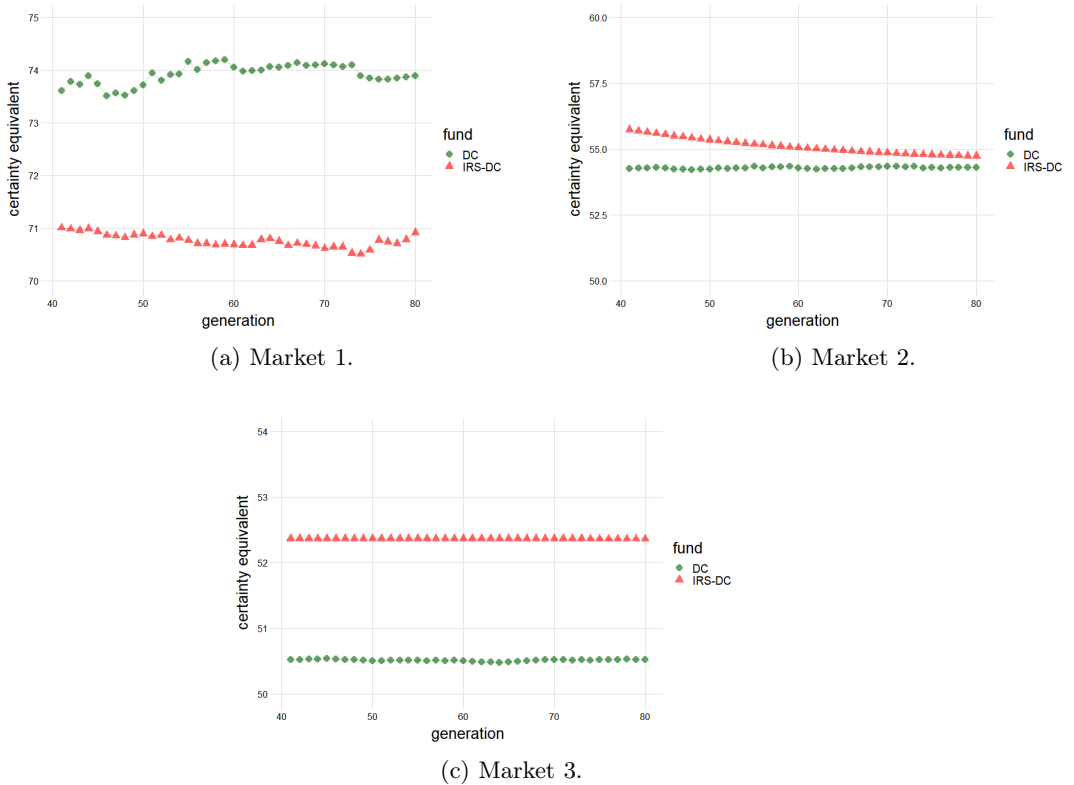


Figure 8: Certainty equivalents of the participants in the IRS-DC and pure DC plans for the generations $i = 41, \dots, 80$, for the three market settings and risk aversion $\gamma = 10$.

financial market is volatile. Key new findings to the literature include that i) the welfare improvement can be achieved without relying on borrowing and share holders, in contrast to, e.g., Gollier (2008), and that ii) whether IRS improves the welfare depends on the volatility of the financial market, as measured by the Sharpe ratio. These observations suggest that a fully funded pension system with a realistic investment strategy (i.e., without borrowing) can implement IRS and protect individual participants from a turbulent market.

Our investigation has been based on a stylized model, which we call the IRS-DC pension model, that uses an indexation rate of individual accounts as a device for IRS. This indexation rate, originally introduced by Goecke (2013), is automatically adjusted according to the notional funding ratio of the pension fund, so as to balance the welfare of different generations and the sustainability of the pension fund. We have analyzed the notional funding ratio process and retirement benefits in the IRS-DC model, and how their volatility is controlled by the adjustment parameter in the indexation rate. Moreover, we have shown how the adjustment parameter and the investment strategy can be optimized by using Bayesian optimization, a machine learning method for global optimization.

There are a number of possible future directions. First, as we have shown the effectiveness of indexation rate of Goecke (2013) as a means for IRS in a collective pension system, the same indexation rate may be applied to other collective schemes, such as hybrid pension systems (e.g., Cui et al. 2011; Chen et al. 2016) and notional DC pension

systems (e.g., Settergren 2001), where other forms of automatic adjustment rules are used for adjusting the individual accounts and/or contributions. This is worth investigation, as automatic adjustment rules have been used in real pension systems, such as the Dutch and Swedish pension systems (OECD 2021, Chapter 2).

Second, as Bayesian optimization provides an efficient way of optimizing the parameters of a pension system, it enables researchers to study optimal pension systems under more realistic but complex setups. For example, Bayesian optimization may be applied to optimize the three-pillars pension system of Chen et al. (2016), which involves a number of parameters, by expected utility maximization; this may enable showing that their collective scheme is welfare-improving over the optimal individual benchmark using the life-cycle investment strategy, as we have shown for our collective scheme.

Third, our finding that IRS is welfare-improving in a volatile market is worth further investigation in a more realistic setup of the financial market. While our setup of the Black-Scholes market (i.e., log-normally distributed risky asset returns) follows many of related works (e.g., Cui et al. 2011; Chen et al. 2016), it is known that this setup does not necessarily hold in reality (e.g., Cont 2001). In particular, log returns of real stocks are known to have heavy tails, which implies that real financial markets are more volatile than the Black-Scholes market. Therefore, IRS could be shown to be more effective for real financial markets than for the Black-Scholes market.

Fourth, the discontinuity risk may be discussed for the IRS-DC model. We have implicitly assumed the mandatory participation of individuals by modeling that the population in each generation remains the same, as in related works (e.g., Gollier 2008; Chen et al. 2016). One could relax this assumption by making the participation voluntary, and study individuals' preferences and how they impact the sustainability of the pension fund and the welfare of different generations (e.g., Beetsma et al. 2012). Because contributions are not adjusted in the IRS-DC plan by design, it may be anticipated that the IRS-DC plan is less prone to the discontinuity risk than DB-based pension plans. However, if the voluntary participation changes the populations of different generations, the effectiveness of IRS may be affected. It will be interesting to investigate whether the mandatory participation is necessary for the IRS-DC plan to maintain effective IRS.

Acknowledgements

This work in part has been supported by the French government, through the 3IA Cote d'Azur Investment in the Future Project managed by the National Research Agency (ANR) with the reference number ANR-19-P3IA-0002, and by Deutsche Forschungsgemeinschaft with Grant number 418318744 of the research project: "Zielrente: die Lösung zur alternden Gesellschaft in Deutschland".

Notes

¹Our indexation rate corresponds to Goecke (2013, Eq. (5)) with $\rho_{\text{target}} = 0$.

²The funding ratio in Sweden's notional DC pension system is notional since both "assets" and "liabilities" are notional, as it is a PAYG system.

³Since retirees leave the fund immediately after receiving lump-sum benefits, Eq. (5) in Settergren (2001) does not exist in our case and thus Eq. (3) in Settergren (2001) is equal to the notional liability in (2.8).

REFERENCES

- Allen, F. and Gale, D. (1997). Financial markets, intermediaries, and intertemporal smoothing. *Journal of Political Economy*, 105(3):523–546.
- Bams, D., Schotman, P. C., and Tyagi, M. (2016). Optimal risk sharing in a collective defined contribution pension system. *Netspar Discussion Paper*.
- Bardet, J.-M., Surgailis, D., et al. (2011). Measuring the roughness of random paths by increment ratios. *Bernoulli*, 17(2):749–780.
- Barr, N. and Diamond, P. (2008). *Reforming Pensions: Principles and Policy Choices*. Oxford University Press.
- Barr, N. and Diamond, P. (2011). Improving sweden’s automatic pension adjustment mechanism. *Center for Retirement Research at Boston College*, 11.
- Baumann, R. T. and Müller, H. H. (2008). Pension funds as institutions for intertemporal risk transfer. *Insurance: Mathematics and Economics*, 42(3):1000–1012.
- Beetsma, R. and Romp, W. (2016). Intergenerational risk sharing. In *Handbook of the Economics of Population Aging*, chapter 6, pages 311–380. Elsevier.
- Beetsma, R. M., Romp, W. E., and Vos, S. J. (2012). Voluntary participation and intergenerational risk sharing in a funded pension system. *European Economic Review*, 56(6):1310–1324.
- Benartzi, S. and Thaler, R. H. (2001). Naive diversification strategies in defined contribution saving plans. *American Economic Review*, 91(1):79–98.
- Bischl, B., Richter, J., Bossek, J., Horn, D., Thomas, J., and Lang, M. (2017). mlrmbo: A modular framework for model-based optimization of expensive black-box functions. *arXiv preprint arXiv:1703.03373*.
- Booth, P. and Yakoubov, Y. (2000). Investment policy for defined-contribution pension scheme members close to retirement: An analysis of the “lifestyle” concept. *North American Actuarial Journal*, 4(2):1–19.
- Boulier, J.-F., Huang, S., and Taillard, G. (2001). Optimal management under stochastic interest rates: the case of a protected defined contribution pension fund. *Insurance: Mathematics and Economics*, 28(2):173–189.
- Bull, A. D. (2011). Convergence rates of efficient global optimization algorithms. *Journal of Machine Learning Research*, 12:2879–2904.
- Cairns, A. J. (1996). Continuous-time pension-fund modelling. In *Proceedings of the 6th AFIR International Colloquium, Nuremberg*, pages 609–624. Citeseer.
- Cairns, A. J., Blake, D., and Dowd, K. (2006). Stochastic lifestyling: Optimal dynamic asset allocation for defined contribution pension plans. *Journal of Economic Dynamics and Control*, 30(5):843–877.
- Chen, A. and Delong, L. (2015). Optimal investment for a defined-contribution pension scheme under a regime switching model. *ASTIN Bulletin: The Journal of the IAA*, 45(2):397–419.
- Chen, D. H., Beetsma, R. M., Ponds, E. H., and Romp, W. E. (2016). Intergenerational risk-sharing through funded pensions and public debt. *Journal of Pension Economics & Finance*, 15(2):127–159.
- Cont, R. (2001). Empirical properties of asset returns: stylized facts and statistical issues. *Quantitative Finance*, 1(2):223.
- Cui, J., De Jong, F., and Ponds, E. (2011). Intergenerational risk sharing within funded pension schemes. *Journal of Pension Economics and Finance*, 10(1):1–29.
- Donnelly, C. (2017). A discussion of a risk-sharing pension plan. *Risks*, 5(1):12.
- Gabrielli, A. (2020). A neural network boosted double overdispersed poisson claims reserving model. *ASTIN Bulletin: The Journal of the IAA*, 50(1):25–60.
- Goecke, O. (2013). Pension saving schemes with return smoothing mechanism. *Insurance: Mathematics and Economics*, 53(3):678–689.
- Gollier, C. (2008). Intergenerational risk-sharing and risk-taking of a pension fund. *Journal of Public Economics*, 92(5-6):1463–1485.
- Gordon, R. H. and Varian, H. R. (1988). Intergenerational risk sharing. *Journal of Public Economics*, 37(2):185–202.
- Haberman, S. and Vigna, E. (2002). Optimal investment strategies and risk measures in defined contribution pension schemes. *Insurance: Mathematics and Economics*, 31(1):35–69.

- Hagen, J. (2013). A history of the Swedish pension system. Working Paper Series, Center for Fiscal Studies 2013:7, Uppsala University, Department of Economics.
- Hainaut, D. (2018). A neural-network analyzer for mortality forecast. *ASTIN Bulletin: The Journal of the IAA*, 48(2):481–508.
- Kanagawa, M., Hennig, P., Sejdinovic, D., and Sriperumbudur, B. K. (2018). Gaussian processes and kernel methods: A review on connections and equivalences. *arXiv preprint arXiv:1807.02582*.
- Karatzas, I. and Shreve, S. (1998). *Brownian Motion and Stochastic Calculus*. Springer Science & Business Media, 2nd edition.
- McKay, M. D., Beckman, R. J., and Conover, W. J. (2000). A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics*, 42(1):55–61.
- Menoncin, F. and Vigna, E. (2017). Mean–variance target-based optimisation for defined contribution pension schemes in a stochastic framework. *Insurance: Mathematics and Economics*, 76:172–184.
- Merton, R. C. (1971). Optimum consumption and portfolio rules in a continuous-time model. *Journal of Economic Theory*, 3(4):373–413.
- OECD (2021). Pensions at a Glance 2021: OECD and G20 Indicators.
- Rasmussen, C. and Williams, C. (2006). *Gaussian Processes for Machine Learning*. MIT Press.
- Schnürch, S. and Korn, R. (2022). Point and interval forecasts of death rates using neural networks. *ASTIN Bulletin: The Journal of the IAA*, 52(1):333–360.
- Scognamiglio, S. (2022). Calibrating the lee-carter and the poisson lee-carter models via neural networks. *ASTIN Bulletin: The Journal of the IAA*, 52(2):519–561.
- Settergren, O. (2001). The automatic balance mechanism of the Swedish pension system. *Wirtschaftspolitische Blätter*, 4.
- Shahriari, B., Swersky, K., Wang, Z., Adams, R. P., and De Freitas, N. (2016). Taking the human out of the loop: A review of Bayesian optimization. *Proceedings of the IEEE*, 104(1):148–175.
- Shiller, R. J. (1999). Social security and institutions for intergenerational, intragenerational, and international risk-sharing. In *Carnegie-Rochester Conference Series on Public Policy*, volume 50, pages 165–204. Elsevier.
- Snoek, J., Larochelle, H., and Adams, R. P. (2012). Practical bayesian optimization of machine learning algorithms. *Advances in Neural Information Processing Systems*, 25.
- Vigna, E. and Haberman, S. (2001). Optimal investment strategy for defined contribution pension schemes. *Insurance: Mathematics and Economics*, 28(2):233–262.
- Wüthrich, M. V. (2020). Bias regularization in neural network models for general insurance pricing. *European Actuarial Journal*, 10(1):179–202.

Appendix A. Proof of Proposition 1

Proof. Let $t_0 \in \mathbb{N}$ be such that $i - N \leq t_0 \leq i - 1$. Let $\rho(t) = \ln(A(t)/L(t))$ be the log notional funding ratio. Then we have

$$\begin{aligned}
 & B_i(t_0 + 1) \tag{A1} \\
 &= (B_i(t_0) + c) \exp\left(\tilde{\mu} + \theta \int_{t_0}^{t_0+1} \rho(s) ds\right), \\
 &\stackrel{(a)}{=} (B_i(t_0) + c) \exp\left(\tilde{\mu} + \rho(t_0)_+ - \rho(t_0 + 1) + \int_{t_0}^{t_0+1} \tilde{\sigma} dZ(s)\right), \\
 &\stackrel{(b)}{=} (B_i(t_0) + c) \exp\left(\tilde{\mu} + (1 - e^{-\theta})\rho(t_0)_+ - \int_{t_0}^{t_0+1} e^{-\theta(t_0+1-s)} \tilde{\sigma} dZ(s) + \int_{t_0}^{t_0+1} \tilde{\sigma} dZ(s)\right), \\
 &= (B_i(t_0) + c) \exp\left(\tilde{\mu} + (1 - e^{-\theta})\rho(t_0)_+ + \tilde{\sigma} \int_{t_0}^{t_0+1} (1 - e^{-\theta(t_0+1-s)}) dZ(s)\right), \tag{A2}
 \end{aligned}$$

where (a) follows from (3.4) and (b) follows from (3.5).

We show a proof by induction. Suppose that, for $m \in \mathbb{N}$ with $0 < m \leq N - 1$, we have

$$\begin{aligned}
 B_i(i - m) = c \sum_{n=m+1}^N \exp\left((n - m)\tilde{\mu} + (1 - e^{-\theta}) \sum_{\ell=m+1}^n \rho(i - \ell)_+ \right. \\
 \left. + \tilde{\sigma} \sum_{\ell=m+1}^n \int_{i-\ell}^{i-\ell+1} (1 - e^{-\theta(i-\ell+1-s)}) dZ(s)\right) \tag{A3}
 \end{aligned}$$

Note that the identity (A3) holds for $m = N - 1$, since we have by (A2) and $B_i(i - N) = 0$

$$B_i(i - N + 1) = c \exp\left(\tilde{\mu} + (1 - e^{-\theta})\rho(i - N)_+ + \tilde{\sigma} \int_{i-N}^{i-N+1} (1 - e^{-\theta(i-N+1-s)}) dZ(s)\right).$$

By using (A2) with $t_0 = i - m$, the assumption (A3) implies that

$$\begin{aligned}
 & B_i(i - m + 1) = \\
 & (B_i(i - m) + c) \exp\left(\tilde{\mu} + (1 - e^{-\theta})\rho(i - m)_+ + \tilde{\sigma} \int_{i-m}^{i-m+1} (1 - e^{-\theta(i-m+1-s)}) dZ(s)\right) \\
 &= \left[c \sum_{n=m+1}^N \exp\left((n - m)\tilde{\mu} + (1 - e^{-\theta}) \sum_{\ell=m+1}^n \rho(i - \ell)_+ \right. \right. \\
 & \quad \left. \left. + \tilde{\sigma} \sum_{\ell=m+1}^n \int_{i-\ell}^{i-\ell+1} (1 - e^{-\theta(i-\ell+1-s)}) dZ(s)\right) + c \right] \\
 & \quad \times \exp\left(\tilde{\mu} + (1 - e^{-\theta})\rho(i - m)_+ + \tilde{\sigma} \int_{i-m}^{i-m+1} (1 - e^{-\theta(i-m+1-s)}) dZ(s)\right) \\
 &= c \sum_{n=m}^N \exp\left((n - m + 1)\tilde{\mu} + (1 - e^{-\theta}) \sum_{\ell=m}^n \rho(i - \ell)_+ \right. \\
 & \quad \left. + \tilde{\sigma} \sum_{\ell=m}^n \int_{i-\ell}^{i-\ell+1} (1 - e^{-\theta(i-\ell+1-s)}) dZ(s)\right),
 \end{aligned}$$

which is the same expression as (A3) with m being replaced by $m - 1$. Therefore, by induction, (A3) holds with $m = 0$, which is (3.8). This completes the proof.

□

Appendix B. Tutorial on Bayesian optimization

We provide here a short tutorial on Bayesian optimization (BO). For further details and references, see e.g. Shahriari et al. (2016).

Let Ω be a parameter set and $f : \Omega \rightarrow \mathbb{R}$ be the objective function to be maximized. In our problem, this parameter set is $\Omega = [0, 1] \times [0, 1]$ and each $x := (\pi, \theta) \in \Omega$ represents a pair of the investment strategy π and adjustment parameter θ . We define the objective function as the certainty equivalent (CE) of the expected utility in (4.3) with input parameters $x = (\pi, \theta)$:

$$f(x) := f(\pi, \theta) := CE(\pi, \theta), \quad (\text{B1})$$

$$\text{where } CE(\pi, \theta) \geq 0 \text{ is such that } U_\gamma(CE(\pi, \theta)) = \mathbb{E} \left[\sum_{t=1}^{\infty} \beta^t U_\gamma(B_t(t)) \right].$$

Note that the expected utility is a function of $x = (\pi, \theta)$, as the payment $B_t(t)$ depends on π and θ . Since the CRRA utility function U_γ is strictly monotonically increasing with respect to its argument, the maximizer of the certainty equivalent is the same as the maximizer of the expected utility:

$$\arg \max_{(\pi, \theta) \in \Omega} CE(\pi, \theta) = \arg \max_{(\pi, \theta) \in \Omega} \mathbb{E} \left[\sum_{t=1}^{\infty} \beta^t U_\gamma(B_t(t)) \right].$$

Thus, the maximization of the expected utility can be equivalently formulated as the maximization of the objective function (B1).

In our study, the expected utility is approximated by the Monte Carlo average of 10,000 simulations of the asset process $A(t)$ (and thus the resulting $B_t(t)$) for $t = 1, \dots, T := 80$. Therefore, each evaluation of $f(x)$ for a given $x = (\pi, \theta)$ involves 10,000 simulations over 80 years on monthly basis, which is computationally expensive. If $A(t) \leq 0$ happens at any $t > 0$ for any of 10,000 simulations of the financial market, we set the objective function value to its minimum: i.e., $f(x) = 0$.

B.1. Procedure of Bayesian optimization.

First, we generate initial design points $x_1, \dots, x_{n_{\text{init}}}$ for some $n_{\text{init}} \in \mathbb{N}$, and evaluate the function values $f(x_1), \dots, f(x_{n_{\text{init}}})$ on these points. One can generate these initial points randomly (e.g. uniform sampling on Ω) or deterministically (e.g., grid points). In our study, we use the design given by Latin hypercube sampling (McKay et al. 2000) on Ω with $n_{\text{init}} = 10$.

Below we use the notation $D_n := \{(x_i, f(x_i))\}_{i=1}^n \subset \Omega \times \mathbb{R}$ to write the collection of points x_1, \dots, x_n and the resulting function values $f(x_1), \dots, f(x_n)$. D_n can be understood as “data” or “observations” about f after n -time evaluations of the function. We also denote by $\alpha(x; D_n)$ the *acquisition function*, whose concrete form will be introduced later in Section B.3. The acquisition function $\alpha(x; D_n)$ is a function of $x \in \Omega$ and defined from D_n .

BO iterates the following procedure for $n = n_{\text{init}} + 1, n_{\text{init}} + 2, \dots, N$, where N is the total number of function evaluations.

- (i) Compute $x_{n+1} \in \arg \max_{x \in \Omega} \alpha(x; D_n)$,
- (ii) Simulate $f(x_{n+1})$, and augment the data $D_{n+1} := D_n \cup \{(x_{n+1}, f(x_{n+1}))\}$.

An estimate of the optimal parameters is then given as the maximizer from the evaluated inputs x_1, \dots, x_N :

$$x^* \in \arg \max \{f(x) \mid x \in \{x_1, \dots, x_N\}\}$$

The acquisition function $\alpha(x; D_n)$ determines the next point x_{n+1} to evaluate the objective function f . Note that the computational cost of solving $\max_{x \in \Omega} \alpha(x; D_n)$ is negligible compared to the computational cost of evaluating $f(x_{n+1})$, as $\alpha(x; D_n)$ can be evaluated cheaply.

The acquisition function is designed so as to balance the *exploitation* and *exploration*. Exploitation is a strategy to search for in a region near the current maximizer in $x_n^* := \arg \max \{f(x) \mid x \in \{x_1, \dots, x_n\}\}$; exploration is to search for in a region far from the evaluated points x_1, \dots, x_n . This *exploration-exploitation trade-off* is enabled by the learning and uncertainty quantification of the response surface of f from the data D_n . This is done by Gaussian process regression, which we will explain next.

B.2. Gaussian process regression

Gaussian process regression (Rasmussen and Williams 2006) is a Bayesian non-parametric method for learning (or approximating) an unknown function $f : \Omega \rightarrow \mathbb{R}$ from its finite observations (data) $D_n = \{(x_i, f(x_i))\}_{i=1}^n$. Recall that Bayesian inference in general proceeds as follows: a) define a *prior distribution* for the quantity of interest, b) collect observations (data) related to that quantity, and c) update the prior distribution to the *posterior distribution* using the observed data, applying Bayes' rule. In Gaussian process regression, the quantity of interest is the unknown *function* f , and a') one defines a prior distribution of f as a Gaussian process (or Gaussian random field), b') collects data $D_n = \{(x_i, f(x_i))\}_{i=1}^n$, and c') updates the prior Gaussian process to the *posterior* Gaussian process, applying Bayes' rule. See Figure 9 for illustrations of Gaussian process regression.

B.2.1. Prior Gaussian Process

A Gaussian process is completely specified by its *mean function* $m : \Omega \rightarrow \mathbb{R}$ and *covariance function* $k : \Omega \times \Omega$. We write $f \sim \mathcal{GP}(m, k)$ to mean that f is a sample path of the Gaussian process with mean function μ and covariance function k . Then we have $\mu(x) = \mathbb{E}[f(x)]$, $x \in \Omega$ and $k(x, x') = \mathbb{E}[(f(x) - \mu(x))(f(x') - \mu(x'))]$, $x, x' \in \Omega$. By specifying μ and k , we implicitly specify the corresponding Gaussian process.

For simplicity, we consider a Gaussian process with the zero-mean function (i.e., $\mu(x) = 0, \forall x \in \Omega$) for our prior distribution of the objective function f :

$$f \sim \mathcal{GP}(0, k). \tag{B2}$$

What we need is to specify the covariance function k . By doing so, we can express our assumption or knowledge regarding key properties of the objective function f , such as its smoothness and structure.

Popular choices of covariance kernels include square-exponential kernel $k(x, x') = \exp(-\|x - x'\|^2/h)$ with $h > 0$ and *Matérn kernels*. In our study we use the so-called Matérn-5/2 kernel of the form

$$k(x, x') = \left(1 + \frac{\sqrt{5}\|x - x'\|}{h} + \frac{5\|x - x'\|^2}{3h^2}\right) \exp\left(-\frac{\sqrt{5}\|x - x'\|}{h}\right) \tag{B3}$$

where $h > 0$ is a scale parameter.⁴ Roughly, this kernel leads to $f \sim \mathcal{GP}(0, k)$ that is almost surely twice differentiable (e.g., Kanagawa et al. 2018, Section 4.4). Thus, with




Figure 9: Illustrations of Gaussian process regression, adapted from Kanagawa et al. (2018, Figure 1). **Left:** The green curves are 5 sample paths from the prior Gaussian process (B2) using the Matérn kernel (B3) with $h = 1$. The thick black line is the prior mean function $m(x) = 0$. **Right:** the three points are noise-perturbed observations $(x_i, y_i)_{i=1}^3$, where $y_i = f^*(x_i) + \varepsilon_i$ with f^* being the ground-truth function and ε_i being an independent zero-mean Gaussian noise with variance $\sigma^2 := 0.01$. The thick black curve is the posterior mean function $m_n(x)$ in (B5) and the two thin black curves are the posterior standard deviation function $\sigma_n(x) = \sqrt{k_n(x, x)}$ in (B7), in which \mathbf{K}_n^{-1} is replaced by $(\mathbf{K}_n + \sigma^2 \mathbf{I})^{-1}$ and $\mathbf{f}_n := (y_1, y_2, y_3)^\top$. (These modifications are theoretically justified; see e.g. Kanagawa et al. (2018) and references therein for details.) The green curves are 5 sample paths from the posterior Gaussian process (B4)

this kernel we essentially assume this degree of smoothness for the objective function, and this is our prior assumption.

B.2.2. Posterior Gaussian Process

The use of a Gaussian process as a prior leads to an *analytic* expression of the resulting posterior distribution. Given data $D_n = \{(x_i, f(x_i))\}_{i=1}^n$, the posterior distribution of f is also given as a Gaussian process

$$f|D_n \sim \mathcal{GP}(m_n, k_n), \quad (\text{B4})$$

where $\mu_n : \Omega \rightarrow \mathbb{R}$ is the *posterior mean function* and $k_n : \Omega \times \Omega \rightarrow \mathbb{R}$ is the *posterior covariance function*, given by

$$m_n(x) = \mathbb{E}[f(x)|D_n] = \mathbf{f}_n^\top \mathbf{K}_n^{-1} \mathbf{k}_n(x), \quad x \in \Omega, \quad (\text{B5})$$

$$\begin{aligned} k_n(x, x') &= \mathbb{E}[(f(x) - m_n(x))(f(x') - m_n(x))|D_n] \\ &= k(x, x') - \mathbf{k}_n(x)^\top \mathbf{K}_n^{-1} \mathbf{k}_n(x'), \quad x, x' \in \Omega, \end{aligned} \quad (\text{B6})$$

where $\mathbf{f}_n := (f(x_1), \dots, f(x_n))^\top$, $\mathbf{k}_n(x) := (k(x, x_1), \dots, k(x, x_n))^\top \in \mathbb{R}^n$ and $\mathbf{K}_n := (k(x_i, x_j))_{i,j=1}^n \in \mathbb{R}^{n \times n}$. For the detail of the above derivation, see Rasmussen and Williams (2006).

The posterior mean function m_n in (B5) is an approximation of the objective function f based on the data D_n . It works as a computationally cheaper surrogate model of f . On the other hand, the *posterior standard deviation*

$$\sigma_n(x) := \sqrt{k_n(x, x)} = \sqrt{\mathbb{E}[(f(x) - \mu_n(x))^2|D_n]} \quad (\text{B7})$$

quantifies the uncertainty about the unknown function value $f(x)$. These μ_n and σ_n are the building blocks of the acquisition function, as we will see next.

B.3. Acquisition function

We now introduce the concrete form acquisition function $\alpha(x; D_n)$. There are many acquisition functions proposed in the literature; see Shahriari et al. (2016, Section IV). Most popular ones include the *EI* (Expected Improvement), *GP-UCB* (Gaussian Process Upper Confidence Bound), and *ES* (Entropy Search). In this paper, we use the EI acquisition function, which is standard and theoretically well studied (Bull 2011). Let

$$f_n^* := \max_{i=1, \dots, n} f(x_i), \quad x_n^* \in \arg \max\{f(x) \mid x \in \{x_1, \dots, x_n\}\}.$$

be the maximum and the maximizer of the objective function $f(x)$ over the currently evaluated inputs x_1, \dots, x_n . The EI acquisition function $\alpha(x; D_n)$ at x is defined as the expected improvement of the function value $f(x)$ over the current maximum f_n^* , where the expectation is with respect to the posterior Gaussian process (B4):

$$\begin{aligned} \alpha(x; D_n) &:= \mathbb{E}_{f \sim \mathcal{GP}(\mu_n, k_n)}[\max(f(x) - f_n^*, 0)] \\ &= \underbrace{\sigma_n(x) \phi\left(\frac{\mu_n(x) - f_n^*}{\sigma_n(x)}\right)}_{\text{Exploration}} + \underbrace{(\mu_n(x) - f_n^*) \Phi\left(\frac{\mu_n(x) - f_n^*}{\sigma_n(x)}\right)}_{\text{Exploitation}}, \end{aligned} \quad (\text{B8})$$

where $\phi: \mathbb{R} \rightarrow [0, \infty)$ is the probability density function of a standard Gaussian random variable, and $\Phi: \mathbb{R} \rightarrow [0, 1]$ is its cumulative distribution function: $\Phi(y) := \int_{-\infty}^y \phi(s) ds$, $y \in \mathbb{R}$.

The first term in (B8) represents the *exploration*, as it becomes large when $\sigma_n(x)$, which represents the uncertainty about the function value $f(x)$, is large. This is typically the case when x is far from already evaluated locations x_1, \dots, x_n . The second in (B8) represents the *exploitation*, as it becomes large when $\mu_n(x) - f_n^*$ is large and $\sigma_n(x)$ is small. This is typically the case when x is near the current maximizer x_n^* . Thus, the EI acquisition function naturally balances the exploration-exploitation trade-off, and the next point $x_{n+1} \in \arg \max_{x \in \Omega} \alpha(x; D_n)$ achieves such a balance.

B.4. Demonstration

Figure 10 shows an example of points $x = (\pi, \theta)$ evaluated by BO for $\gamma = 10$. The green points are the $n_{\text{init}} = 10$ initial design points $x_1, \dots, x_{n_{\text{init}}}$ generated by Latin hypercube sampling. The total number of evaluated points is $N = 100$. The red point is the maximizer x_N^* , and the blue points are the 10 other second-best parameters (largely overlapping with the red point). For a comparison, we show 10×10 grid points.

Appendix C. IR-roughness Measure

To describe the IR-roughness measure (Bardet et al. 2011), we suppose that the path of each individual account is represented by a function $h: [0, \tilde{T}] \rightarrow \mathbb{R}$, where $\tilde{T} = 40$ is its terminal time. Note that \tilde{T} is the terminal time for one generation but is different from the terminal time for the operation of the pension fund T . Discretizing the domain to $n - 1 \in \mathbb{N}$ intervals, the first-order IR-roughness is defined as

$$R^{1,n}(h) := \frac{1}{n-1} \sum_{j=0}^{n-2} \frac{\left| h\left(\frac{j+1}{n}\right) - h\left(\frac{j}{n}\right) + h\left(\frac{j+2}{n}\right) - h\left(\frac{j+1}{n}\right) \right|}{\left| h\left(\frac{j+1}{n}\right) - h\left(\frac{j}{n}\right) \right| + \left| h\left(\frac{j+2}{n}\right) - h\left(\frac{j+1}{n}\right) \right|}. \quad (\text{C1})$$




Figure 10: *Demonstration of Bayesian optimization with $\gamma = 10$ (see Section B.4). The points represent the values of (π, θ) evaluated by Bayesian optimization. The green points are the initial 10 points given by Latin hypercube sampling. The red point is the maximizer found by Bayesian optimization after 100 evaluations, and the blue points (largely overlapping the red point) are the 10 second-best points.*

By the triangle inequality, the numerator in the sum is less than or equal to the denominator, and thus $R^{1,n}(h)$ takes values between 0 and 1. When the signs of the two increments $h(\tilde{T}(j+1)/n) - h(\tilde{T}j/n)$ and $h(\tilde{T}(j+2)/n) - h(\tilde{T}(j+1)/n)$ are the same, the numerator equals the denominator. While when those signs are different, the numerator is smaller than the denominator. As such, $R^{1,n}(h)$ reflects the sign changes of the function h and thus quantifies its roughness. Intuitively, $R^{1,n}(h)$ is close to 0 when h is rough, and is close to 1 when h is smooth. In fact, Bardet et al. (2011) show that, for a sufficiently smooth h , $R^{1,n}(h)$ converges to 1 as $n \rightarrow \infty$.

Appendix D. Additional Numerical Results

We show here additional numerical results not included in the main body of the paper.

D.1. Notional Funding Ratio Process

Figures 11 and 12 show additional results for the notional funding ratio process; see Section 5.2 for details.

D.2. Distribution of Retirement Benefits

Figure 13 shows the histograms of the retirement benefits of the IRS-DC and pure DC plan participants for two settings of the risk aversion, $\gamma = 3, 5$; see Section 5.4 for details.

D.3. Welfare of Participants

Figure 14 shows the certainty equivalents of the IRS-DC and pure DC participants for the generations $i = 41, \dots, 80$ for two settings of the risk aversion, $\gamma = 3, 5$; see Section 5.5 for details.




Figure 11: Notional funding ratio processes for the risk aversion $\gamma = 3$ and for the three market settings M1, M2 and M3. For each setting, the average and standard deviation of the notional funding ratio processes over 10,000 simulations are shown.




Figure 12: Notional funding ratio processes for the risk aversion $\gamma = 5$ and for the three market settings M1, M2 and M3. For each setting, the average and standard deviation of the notional funding ratio processes over 10,000 simulations are shown.

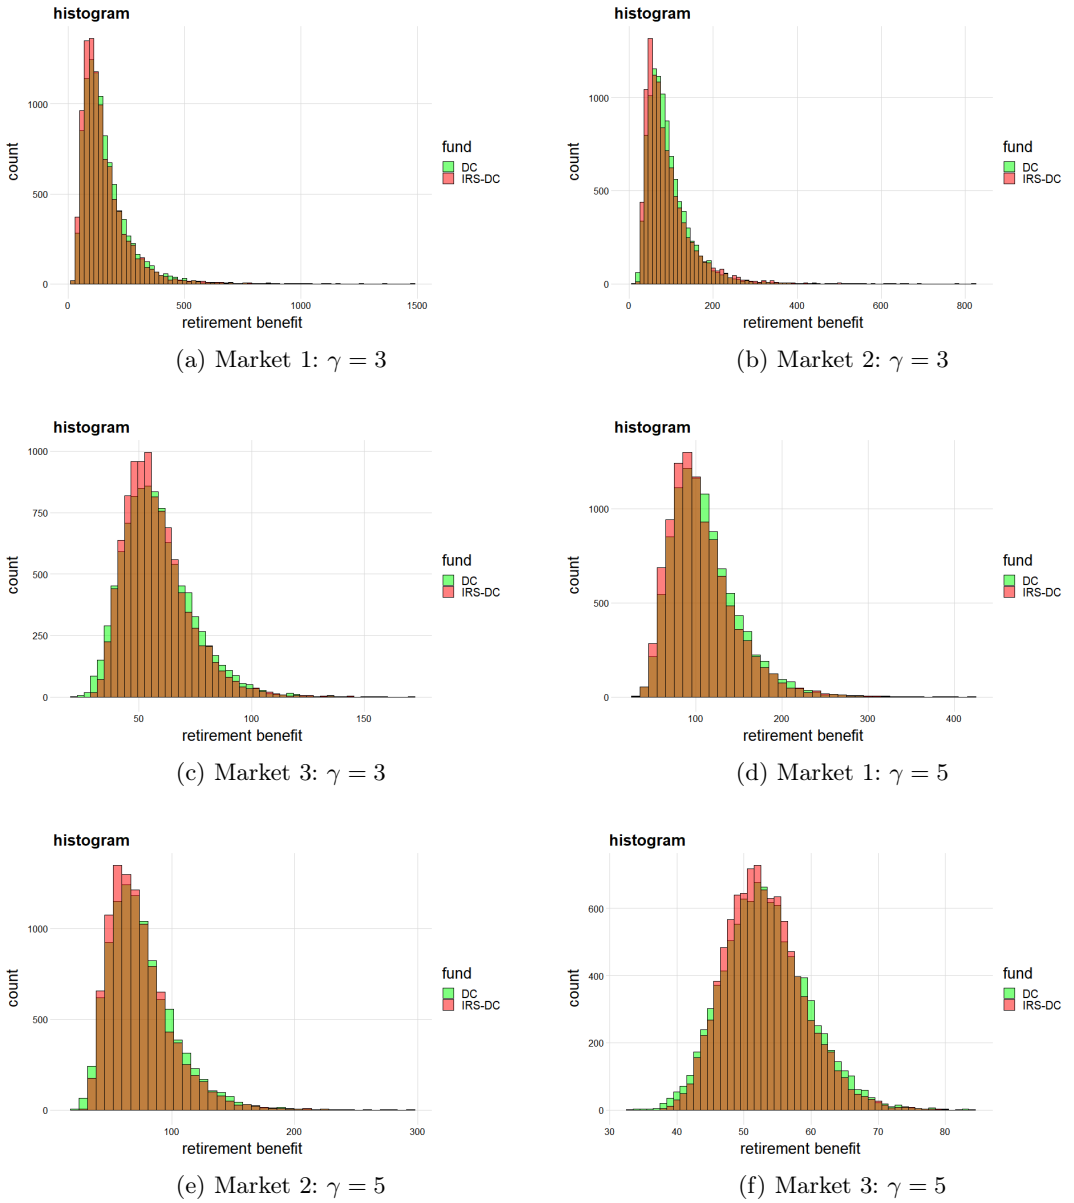





Figure 13: Histograms of the retirement benefits (obtained from the 10,000 simulations) of the IRS-DC and pure DC accounts for the generation $i = 41$, for the three market settings and risk aversion $\gamma = 3, 5$. In each figure, the red and green histograms are those of the IRS-DC and pure DC benefits, respectively; the brown part shows the overlap between the two histograms.




(a) Market 1: $\gamma = 3$




(b) Market 2: $\gamma = 3$




(c) Market 3: $\gamma = 3$



(d) Market 1: $\gamma = 5$



(e) Market 2: $\gamma = 5$



(f) Market 3: $\gamma = 5$

Figure 14: Certainty equivalents of the participants in the IRS-DC and pure DC plans for the generations $i = 41, \dots, 80$, for the three market settings and risk aversion $\gamma = 3, 5$.