Understanding Heterogeneity in Lifetime Pension Pools: Implications for Practice

JULY | 2026

Understanding Heterogeneity in Lifetime Pension Pools: Implications for Practice

Author

Jean-François Bégin, FSA, FCIA, PhD
Associate Professor
Simon Fraser University

Isaac Marchand, BSc
Master's Student
Simon Fraser University

Barbara Sanders, FSA, FCIA, MSc
Associate Professor
Simon Fraser University

Sponsors

Aging and Retirement Strategic Research Program

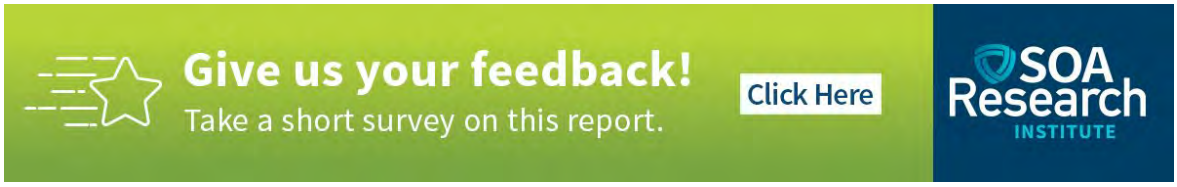

Caveat and Disclaimer

The opinions expressed and conclusions reached by the authors are their own and do not represent any official position or opinion of the Society of Actuaries Research Institute, Society of Actuaries, or its members. The Society of Actuaries Research Institute makes no representation or warranty to the accuracy of the information.

Copyright © 2026 by the Society of Actuaries Research Institute. All rights reserved.

CONTENTS

Executive Summary **4**

Section 1 Introduction and Literature Review **6**

Section 2 The Most Common Lifetime Pension Pool Design **10**
2.1 Pool Membership 10
2.2 Initial Benefits 10
2.3 Whole-Pool Adjustments 12
2.4 Example of the Whole-Pool Adjustment Calculation 13

Section 3 Conceptual Insights on the Impact of Heterogeneity on Stability **16**
3.1 Homogeneous Pool 16
3.1.1 General Setup 16
3.1.2 Standard Deviation of the Adjustment 17
3.2 Heterogeneous Pool 19
3.2.1 General Setup 19
3.2.2 Standard Deviation of the Adjustment 20
3.2.3 Impact of Age 21
3.2.4 Impact of the Initial Benefit 22
3.2.5 Impact of Both Dimensions 23
3.3 Implied Number of Homogeneous Pool Members 27

Section 4 A Deeper Look at Stability **30**
4.1 Introducing Another Measure of Stability: The Stable Income Period 30
4.2 Computation of the Stable Income Period Measure 31
4.3 Homogeneous Pool 32
4.4 Heterogeneous Pool 33
4.4.1 Impact of Age 33
4.4.2 Impact of the Initial Benefit 34
4.4.3 Impact of Both Dimensions 34
4.5 Impact of Investment Experience on Stable Income Period and Heterogeneity 39
4.6 Comparison with the Standard Deviation of the Adjustment 40

Section 5 Mitigation of Heterogeneity **42**
5.1 Considering Stability Only 43
5.2 Risk–Reward Consideration 43

Section 6 Concluding Remarks and Future Research **48**

Section 7 Acknowledgments **50**

Appendix A Proofs **51**
A.1 Proof of Proposition 1 51
A.2 Proof of Proposition 2 52

Appendix B Alternative Setting of a Fixed-Size Pool **55**
B.1 Standard Deviation of the Adjustment 55
B.2 Stable Income Period 56

Appendix C More Details on Interpreting Age for Mortality Differentials **58**
C.1 Age as a Proxy for Mortality Heterogeneity 58
C.2 Effect of Mortality Dispersion on the Stable Income Period 58

References **60**

About The Society of Actuaries Research Institute **63**

Copyright © 2026 Society of Actuaries Research Institute

Understanding Heterogeneity in Lifetime Pension Pools: Implications for Practice

Executive Summary

This report examines how heterogeneity within lifetime pension pools affects the stability of retirement income. While lifetime pension pools are often developed under the assumption that *larger pools are more stable*, the analysis shows that pool composition matters at least as much as pool size. Using two complementary measures of stability, the standard deviation of benefit adjustments and the stable income period (SIP), the report demonstrates that heterogeneity can influence benefit outcomes in systematic, predictable, and operationally relevant ways.

The results highlight a fundamental asymmetry across sources of heterogeneity. On the one hand, mortality heterogeneity can be either stabilizing or destabilizing depending on how it enters the pool. Adding younger members generally improves stability, often substantially, by extending the period over which income remains predictable. Adding older members, by contrast, yields limited diversification benefits and initially reduces stability unless their number is sufficiently large. These findings are robust across different stability measures.

On the other hand, the results illustrate that heterogeneity, in terms of the amount invested in the pool (proxied by differences in initial benefit levels), is consistently destabilizing. Pools that include groups with different benefit levels tend to exhibit greater variability and shorter SIPs than homogeneous pools of the same size, particularly when a small group has disproportionately high or low initial benefits. While increasing the number of members eventually improves stability through scale, benefit heterogeneity dilutes these diversification gains rather than enhancing them.

When the multiple dimensions of heterogeneity interact, stability outcomes depend on their joint configuration. Younger members can sustain higher initial benefit levels without undermining stability, whereas older members improve stability only when their initial benefits are lower.

The report also shows that the investment environment matters. When investment risk is high, return volatility becomes the dominant driver of benefit adjustments, compressing differences in stability across pool compositions. In such settings, heterogeneity remains relevant but plays a secondary role relative to asset allocation.

Overall, these findings illustrate that additional consideration is warranted when assuming that *more pooling is better*. Larger, more heterogeneous pools can deliver meaningful gains, but only when heterogeneity is well balanced. Poorly balanced heterogeneity can reduce stability, complicate operations, and undermine member confidence.

The analysis also shows that adverse effects of heterogeneity can be mitigated through design choices. In particular, adjusting pricing at entry can rebalance heterogeneous pools so that participation remains attractive (that is, stability is maintained or enhanced) relative to operating independently. These adjustments necessarily involve trade-offs between initial benefit levels and stability: improving stability for all groups may require some groups to accept lower initial benefits. Realistic approaches to pricing adjustments explicitly take these trade-offs into account, balancing stability and benefit adequacy rather

Copyright © 2026 Society of Actuaries Research Institute

than targeting stability alone. Importantly, mitigation should be viewed as an operational and design tool, not as a substitute for sound pool composition.

The findings of this report can help inform the design, governance, and evaluation of lifetime pension pools. They show that heterogeneity affects income stability in systematic ways that can be anticipated and managed but not fully eliminated. The observations below provide practical considerations for pool operators on how to assess heterogeneity, when pooling is likely to enhance stability, and when alternative design choices could be considered.

1. **Heterogeneity can interact with design decisions in determining outcomes.** The results show that pool size alone is not sufficient to ensure stable outcomes, and poorly balanced heterogeneity can reduce stability even in large pools. Consequently, when considering creating, expanding, merging, or consolidating lifetime pension pools, practitioners can benefit from explicitly assessing who is being added, not just the number of members. The specific impact of age composition and initial benefit level dispersion can be evaluated jointly using stability metrics such as the SIP.

2. ***Extreme* heterogeneity can undermine stability.** When large differences in age and initial benefit level coexist within a pool (e.g., when older members who have high initial benefits are commingled with younger members who have low initial benefits), income stability is reduced relative to more balanced configurations. Segmentation, entry pricing adjustments, or benefit caps can help mitigate these effects and support more stable outcomes.

3. **Instability arising from heterogeneity can be mitigated by tools such as entry pricing adjustments.** When lifetime pension pools include groups that differ materially, adjusting entry prices can be an effective way to mitigate instability. By rebalancing entry prices across groups, pricing adjustments can offset the destabilizing effects of extreme heterogeneity while maintaining overall fairness at the pool level. This finding suggests that entry pricing adjustments can be a design tool, alongside benefit limits, rather than relying solely on pool size or uniform benefit rules to manage heterogeneity.

4. **Investment risk can have a greater impact than heterogeneity in volatile portfolios**. When lifetime pension pools are exposed to volatile investment returns, investment risk becomes the dominant driver of benefit variability, and the incremental impact of heterogeneity in age or initial benefit level on income stability is substantially reduced. In such settings, the results show that differences in pool composition generate only modest changes in stability relative to the variability induced by asset returns. From a practitioner's perspective, this implies that heterogeneity may be considered as a second-order concern in pools with moderate to highly volatile investment strategies.

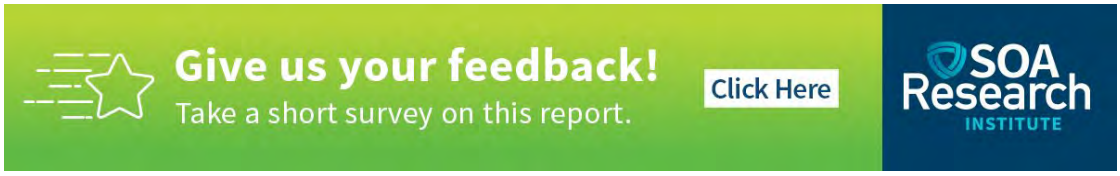

Copyright © 2026 Society of Actuaries Research Institute

Section 1 Introduction and Literature Review

With the decline in guaranteed pension arrangements worldwide, retirement systems are increasingly turning toward more flexible decumulation solutions. Among these, lifetime pension pools are expected to play an increasingly important role. At a high level, a lifetime pension pool allows retirees to convert a lump sum of retirement savings into income paid for life. The defining feature of these arrangements is that longevity and investment risks are shared collectively across participants, rather than being borne individually or transferred to a sponsor or insurer. This pooling mechanism directly addresses one of the central challenges of retirement planning: the risk of outliving one's savings, which remains a significant concern under traditional income drawdown strategies.

The literature suggests that pooling retirement assets in this way can lead to better outcomes than alternative lifetime income solutions. In particular, lifetime pension pools can outperform insured annuities on both financial and utility grounds, largely because they do not require risk capital to support guarantees (see, e.g., Bernhardt and Donnelly, 2019; Chen and Rach, 2023). Instead of fixing benefit levels in advance, benefits adjust over time based on the experience of the pool. This flexibility allows lifetime pension pools to deliver lifelong income.

Although the term lifetime pension pool is increasingly used as an umbrella label, arrangements based on this idea appear under a variety of names. These include group self-annuitization (GSA) schemes (Piggott et al., 2005; Qiao and Sherris, 2013), pooled annuity funds (Stamos, 2008; Donnelly et al., 2013), annuity overlay funds (Donnelly et al., 2014), retirement tontines (Milevsky and Salisbury, 2015, 2016; Fullmer, 2019; Chen et al., 2021; Winter and Planchet, 2022), variable annuities (Balter and Werker, 2020; Balter et al, 2020) variable payout annuities (Horneff et al., 2010), and variable payment life annuity (VPLA) funds (Sanders and Bégin, 2025).

Importantly, lifetime pension pools are not merely theoretical constructions. Long-standing and well-established examples exist. One of the earliest is the College Retirement Equities Fund (CREF), operated by the Teachers Insurance and Annuity Association of America in the US since 1952. In Canada, VPLAs have been paid by the Faculty Pension Plan at the University of British Columbia since 1967. More recently, Australia introduced similar arrangements through QSuper (now part of the Australian Retirement Trust) with the launch of its Lifetime Pension in 2021. Two Canadian financial services firms followed suit in 2021 and 2022, respectively: Purpose Investments with their Longevity Pension Fund, and Guardian Capital with the GuardPath Modern Tontine Trust.[1]

Despite their growing adoption, lifetime pension pools raise several important challenges. Because benefits are not guaranteed, clear communication is essential so that members understand that income levels may vary over time and can meaningfully compare these arrangements with other retirement income options (see, e.g., Bégin and Sanders, 2023a). Lifetime pension pools must also operate within a regulatory framework, be it pension or securities regulation, creating additional design and governance considerations.

Beyond these issues, the design of a lifetime pension pool involves several key choices. These include how the investment policy is specified and how hurdle rates or payout adjustment mechanisms are defined (see, e.g., Li et al., 2022; Olivieri et al., 2022; Bégin and Sanders, 2024; Bégin et al., 2026a), whether the pool is open to new entrants or closed (see, e.g., Bégin and Sanders, 2023b), and whether benefits are

[1] Note that Guardian announced it will terminate the GuardPath Modern Tontine Trust in early 2025 as part of a wind-up process.

Copyright © 2026 Society of Actuaries Research Institute

determined within a single collective group or across multiple groups. This report focuses on the latter, which is a particularly important and relatively underexplored aspect of design. Specifically, the influence of heterogeneity within a pool on benefit variability and stability is investigated. Understanding this relationship is central to assessing the practical performance of lifetime pension pools.

From a retiree's perspective, stability is often as important as the expected level of income itself. Retirement budgets are typically fixed and inflexible, making it difficult to absorb unexpected income reductions. Evidence from behavioral economics also suggests that individuals place a disproportionately high weight on losses relative to gains, meaning that a decrease in benefits is felt more acutely than an equivalent increase (Kahneman and Tversky, 1979). As a result, two retirement income arrangements with the same expected benefits may be perceived very differently depending on how stable those benefits are over time. For lifetime pension pools, in particular, stability plays a central role in member confidence, communication, and long-term acceptance.

Despite its importance, there is no single way to measure benefit stability. In practice, stability is often discussed qualitatively or assessed using simple indicators, such as the frequency or magnitude of benefit adjustments. Academic literature has proposed a range of quantitative measures, including dispersion-based metrics, downside-risk measures, and path-dependent indicators that capture the persistence of benefit reductions (see, e.g., Donnelly, 2017; Chen et al., 2023; Bégin et al., 2026b). Each of these measures emphasizes different aspects of the benefit experience, and their implications can differ materially, especially when the pool membership is heterogeneous.[2]

This report provides some insights into ways to measure benefit stability. Stability, however, is not the only dimension that may be relevant when evaluating lifetime pension pools. Two other concepts are often invoked in this context: sustainability and fairness.

1. Sustainability, despite being an important notion in pension, is used to describe several distinct concepts in practice (Gros and Sanders, 2019). In the context of a lifetime pension pool, sustainability may simply refer to the ability of the pool to continue paying benefits at all. The devastation measure introduced by Donnelly (2017) provides a formal way to quantify the risk of such an event.

   For the class of lifetime pension pools considered in this report (with immediate, whole-pool benefit adjustments), this form of sustainability is not a binding concern provided benefits are priced reasonably accurately at inception.[3] As shown in Bégin and Sanders (2023b), the adjustment mechanism is self-correcting: deviations from expected experience are absorbed through benefit changes rather than accumulating as funding shortfalls. Although mispricing can create redistribution and fairness concerns, it does not threaten sustainability in the narrow sense of complete asset depletion, even in heterogeneous pools. Therefore, this aspect of sustainability is not pursued further.

[2] Several quantitative measures of benefit stability have been proposed in the literature. These include downside-risk indicators, such as benefit at risk (Bégin and Sanders, 2023a), time-based measures that quantify how long benefits remain within specified bounds (Bernhardt and Donnelly, 2021), and related extensions that assess the impact of heterogeneity by comparing stability to an equivalent homogeneous pool (Donnelly, 2023; Bernhardt and Qu, 2024).

[3] In this report, arrangements that provide a uniform adjustment across the whole membership are considered. These are referred to as whole-pool or uniform adjustments. In the literature, this type of adjustment is often referred to as a group-based adjustment because a pool and a group are considered synonymous. Since the framework considered here allows the pool to consist of multiple distinct groups with heterogeneous characteristics, these adjustments are instead referred to as whole-pool adjustments.

Copyright © 2026 Society of Actuaries Research Institute

2. Fairness and, in particular, actuarial fairness, is another commonly cited objective. At its core, actuarial fairness requires that the present value of benefits equal the value of contributions. This condition holds at the aggregate pool level but need not hold for individual members or groups. At these levels, the strongest feasible criterion is ex ante fairness, defined in terms of expected present values, since realized outcomes inevitably differ in any arrangement that pools idiosyncratic mortality risk.

   In homogeneous pools, the uniform benefit adjustment rule used in this report delivers benefits that are actuarially fair in expectation at the individual level. This property generally breaks down under heterogeneity; because a single adjustment factor applies to all members, persistent transfers between groups can arise. More refined sharing rules can mitigate these effects, but at the cost of added complexity and reduced transparency. Moreover, fairness assessments depend on assumed models of future investment returns and mortality, which are themselves uncertain. As a result, while actuarial fairness remains a useful benchmark, it is a limited lens through which to study heterogeneity.

Against this backdrop, benefit stability provides a more operational and member-relevant criterion, so this dimension is the focus of the report. The analysis shows that the impact of heterogeneity on benefit stability depends on both the measure used to assess stability and the source of heterogeneity.

A simple, one-period measure based on the standard deviation of benefit adjustments is first used to highlight clear conceptual mechanisms. This analysis shows that differences in survival probabilities can either increase or decrease the variability of benefits depending on how new members compare to the existing pool; adding younger members generally improves stability through diversification, whereas adding older members initially increases variability unless their number is sufficiently large. By contrast, heterogeneity in the amount invested in the pool at inception (proxied by the initial level of benefits) is unambiguously destabilizing relative to a homogeneous benchmark, particularly when a small group has benefits that differ sharply from the rest of the pool. When age and initial benefit heterogeneity interact, stability outcomes depend on their joint configuration, and some configurations improve stability for one group but not for both simultaneously, highlighting the presence of trade-offs across members.

The analysis is then refined using a more comprehensive stability metric, the stable income period, or SIP. Specifically, this measure captures the persistence of stable benefits over time, leading to more nuanced but consistent conclusions. Age heterogeneity remains the dominant driver of stability outcomes; pools that expand by adding younger members experience substantially longer periods of stable income, whereas adding older members yields limited or even negative gains. Dispersion in initial benefit levels consistently shortens the period of stable income, regardless of direction, reinforcing the conclusion that benefit heterogeneity undermines stability. When multiple sources of heterogeneity are present simultaneously, the analysis reveals clear trade-offs among age, initial benefit levels, and group size. Overall, the results show that, while pooling can improve stability relative to remaining independent, heterogeneity matters in systematic ways, and stability gains depend on balance rather than pool size alone.

This second stability metric also makes it possible to examine practical approaches to mitigating the loss of stability that can arise in heterogeneous lifetime pension pools. Rather than proposing a single optimal design, the analysis evaluates how design levers, particularly adjustments to entry pricing, can be used to rebalance heterogeneous pools so that participation remains attractive relative to operating independently. The analysis highlights the trade-offs between income stability and initial benefit levels and shows how setting appropriate objectives can help identify configurations that are acceptable to all groups, while preserving actuarial fairness at the pool level.

Copyright © 2026 Society of Actuaries Research Institute

Note that several important topics fall outside the scope of this report. The report did not propose underwriting or eligibility rules, address issues related to selection or disclosure policies, or advocate for a specific lifetime pension pool design. These issues are highly relevant in practice, but they raise distinct regulatory, legal, and policy considerations that merit separate analysis. Throughout the report, initial pricing is assumed to be appropriate and consistent with member characteristics, and the relevant sources of heterogeneity are assumed to be observable.

The remainder of the report is organized as follows. Section 2 introduces the mechanics of the lifetime pension pool considered in this study. Sections 3 and 4 provide the measures used to assess benefit stability: the standard deviation of the adjustment and the SIP, respectively. These chapters are intended to build intuition and provide the tools needed to assess heterogeneity. Section 5 brings the analysis together by presenting a discussion of potential mitigation strategies for the impact of heterogeneity on stability. Section 6 provides the conclusion.

Copyright © 2026 Society of Actuaries Research Institute

Section 2 The Most Common Lifetime Pension Pool Design

As noted earlier, various lifetime pension pool designs have already been proposed in the literature. This section describes some assumptions regarding the pool membership and most common benefit-update mechanism used in lifetime pension pools. Specifically, this approach relies on uniform, whole-pool benefit adjustments across all members, drawing inspiration from earlier work on GSA schemes by Piggott et al. (2005), Valdez et al. (2006), Qiao and Sherris (2013), Hanewald et al. (2013), and Bégin and Sanders (2024). Similar methods are currently employed by institutions such as the College Retirement Equities Fund (CREF, 2025) and the University of British Columbia's Faculty Pension Plan.

2.1 POOL MEMBERSHIP

This section begins by formally defining the membership, as this will be particularly helpful later when addressing pool heterogeneity. Members of the pool do not need to be identical; they can differ significantly in terms of age, sex, socioeconomic status, and initial benefit level. It is assumed that the membership could be divided into S groups denoted by $G \in \{1, \dots, S\}$. Let $N_t^{(G)}$ represent the number of survivors in group G at time t. The total number of members N_t is obtained by summing up the size of all groups:

$$N_t = \sum_{G=1}^{S} N_t^{(G)}.$$

2.2 INITIAL BENEFITS

For simplicity, and unless stated otherwise, this report assumes a simple mortality law that is representative of systematic mortality in the pool. Specifically, the Gompertz law is used for illustrative purposes, where the probability that a member aged x survives at least t years is given by

$${}_tp_x = \exp\left(-\exp\left(\frac{x-m}{b}\right)\left(\exp\left(\frac{t}{b}\right)-1\right)\right),$$

where m represents the modal value of the future lifetime distribution and b is the dispersion coefficient. The Gompertz law makes it possible to model mortality heterogeneity among groups as differences in age, providing a transparent and practical way to capture variation in mortality risk without introducing additional complexity.

This law could be different from one group to the next, but it is assumed to be known here.[4] In practice, it would be possible to incorporate stochastic mortality or generational life tables. However, for simplicity, these more advanced approaches were not considered here. This simplifying assumption is not expected to qualitatively impact the results.

[4] Appendix C shows that differences in mortality risk beyond age can be interpreted, to a large extent, as differences in implied age. From a stability perspective, this implies that the effects of mortality heterogeneity are largely captured by age-based analyses, which motivates the focus of the main text.

Copyright © 2026 Society of Actuaries Research Institute

The term structure of interest rates used for discounting is flat and denoted by the constant (continuously compounded) interest rate h, which is the so-called hurdle rate.[5]

Let $\ddot{a}_x^{(G)}$ denote the price of an annuity due for a life who is assumed to be aged x. The superscript G indicates the group to which the individual belongs. Here, a group consists of members sharing the same characteristics such as sex, age, socioeconomic status, and initial benefit level. The price of this annuity due, based on the assumptions mentioned above, is given by

$$\ddot{a}_x^{(G)} = \sum_{t=0}^{\infty} {}_tp_x^{(G)} \exp(-ht), \tag{1}$$

where ${}_tp_x^{(G)}$ is the probability that a member in group G aged x survives at least t years. The annuity price in Equation (1) can be used to derive the following recursive relationship that is quite useful in the context of lifetime pension pool adjustment calculations:

$$\ddot{a}_{x+t-1}^{(G)} - 1 = p_{x+t-1}^{(G)} \ddot{a}_{x+t}^{(G)} \exp(-h). \tag{2}$$

Assume that, at time 0, each member in group G has an investment amount of $A_0^{(G)}$, so that the total asset value at inception is given by

$$A_0 = \sum_{G=1}^{S} N_0^{(G)} A_0^{(G)}.$$

A common approach in the literature is to determine each member's entry benefit by dividing their investment amount by an annuity factor that reflects their individual risk profile. Specifically, for each member in group G, the benefit $B_0^{(G)}$ at time 0 can be obtained from

$$B_0^{(G)} = \frac{A_0^{(G)}}{\ddot{a}_{x^{(G)}}^{(G)}},$$

where $x^{(G)}$ is the age at inception of the members in group G.

It is important to note that a challenge in designing pools with heterogeneous members is achieving fairness in pricing at entry. Milevsky and Salisbury (2016) introduced the concept of equitable share prices, whereby each cohort purchases shares in the pool at a cost that equalizes the expected present value of payouts per dollar invested across all participants. Building on this idea, Chen et al. (2021) derived optimal withdrawal rates for mixed cohorts and then adjusted share prices numerically to approximate fairness across groups. In both approaches, strict actuarial fairness cannot be achieved when mixing cohorts in a pool, since differences in survival prospects and time horizons prevent equal present values across all groups. So, fairness at entry is not absolute but can be operationalized through the calibration of entry prices, which plays a role in balancing contributions and expected benefits among diverse cohorts.

[5] Continuously compounded rates are used for mathematical convenience in this report. Some practitioners might be more familiar with effective annual rates. In this case, one can simply replace the discounting factors. These substitutions should not change the results presented in this report, assuming that continuously compounded rates are appropriately converted into effective annual rates.

Copyright © 2026 Society of Actuaries Research Institute

2.3 WHOLE-POOL ADJUSTMENTS

After the first payment is made at time 0, the operation of the pool under study is similar to that explained in Piggott et al. (2005) and Qiao and Sherris (2013) in the context of GSA schemes. They describe how the amount left by decedents is shared among the surviving members, among other things.

Assume that at time $t-1$, each member in group G has an investment amount of $A_{t-1}^{(G)}$, so that the total asset value is given by[6]

$$A_{t-1} = \sum_{G=1}^{S} N_{t-1}^{(G)} A_{t-1}^{(G)}.$$

This amount is sufficient to support the benefit $B_{t-1}^{(G)}$ for each group:

$$A_{t-1} = \sum_{G=1}^{S} N_{t-1}^{(G)} B_{t-1}^{(G)} \ddot{a}_{x^{(G)}+t-1}^{(G)}. \tag{3}$$

At time t, the total asset value becomes

$$A_t = \left(A_{t-1} - \sum_{G=1}^{S} N_{t-1}^{(G)} B_{t-1}^{(G)}\right) \exp(r_t^{\mathrm{PF}}) = \sum_{G=1}^{S} N_{t-1}^{(G)} B_{t-1}^{(G)} \left(\ddot{a}_{x^{(G)}+t-1}^{(G)} - 1\right) \exp(r_t^{\mathrm{PF}}),$$

where r_t^{PF} is the time-t continuously compounded rate of return realized on the asset portfolio.[7] Unless otherwise specified, it is assumed that the portfolio was invested in the risk-free asset.

The equation above is obtained retrospectively: it is a roll-forward of assets using actual benefit payments and actual investment returns. One can also get a value for A_t, prospectively, similar to Equation (3):

$$A_t = \sum_{G=1}^{S} N_t^{(G)} B_t^{(G)} \ddot{a}_{x^{(G)}+t}^{(G)} = \sum_{G=1}^{S} N_t^{(G)} \left(\delta_t B_{t-1}^{(G)}\right) \ddot{a}_{x^{(G)}+t}^{(G)},$$

where δ_t is the multiplicative adjustment factor applied at time t to the benefits from the prior period.

Simply put, the value of the adjusted future benefits should equal the accumulated value of the assets. Equating the retrospective and prospective values for A_t, and assuming the same adjustment is applied to all surviving members' benefits at time t, gives

$$\delta_t = \left(\frac{\sum_{G=1}^{S} N_{t-1}^{(G)} B_{t-1}^{(G)} \left(\ddot{a}_{x^{(G)}+t-1}^{(G)} - 1\right)}{\sum_{G=1}^{S} N_t^{(G)} B_{t-1}^{(G)} \ddot{a}_{x^{(G)}+t}^{(G)}}\right) \exp(r_t^{\mathrm{PF}}).$$

Applying the recursion from Equation (2) to the annuity factors yields the following expression:

---

[6] These amounts are notional and are used solely for the purpose of calculating benefit adjustments. Members do not have a direct entitlement to these amounts; upon a member's death, the entire amount is forfeited and redistributed among the surviving members.

[7] The factor $\exp(r_t^{\mathrm{PF}})$ could be replaced with $(1+\tilde{r}_t^{\mathrm{PF}})$, where $\tilde{r}_t^{\mathrm{PF}} = \exp(r_t^{\mathrm{PF}}) - 1$ is the effective annual rate equivalent to the continuously compounded rate r_t^{PF}.

Copyright © 2026 Society of Actuaries Research Institute

$$\begin{aligned}\delta_t &= \left(\frac{\sum_{G=1}^{S} N_{t-1}^{(G)}\ B_{t-1}^{(G)}\ p_{x^{(G)}+t-1}^{(G)} \ddot{a}_{x^{(G)}+t}^{(G)}}{\sum_{G=1}^{S} N_t^{(G)} B_{t-1}^{(G)}\ \ddot{a}_{x^{(G)}+t}^{(G)}}\right) \exp(r_t^{\mathrm{PF}} - h) \\ &= \mathrm{MEA}_t \times \mathrm{IEA}_t,\end{aligned} \tag{4}$$

where

$$\mathrm{MEA}_t = \frac{\sum_{G=1}^{S} N_{t-1^+}^{(G)}\ B_{t-1}^{(G)}\ p_{x^{(G)}+t-1}^{(G)} \ddot{a}_{x^{(G)}+t}^{(G)}}{\sum_{G=1}^{S} N_t^{(G)} B_{t-1}^{(G)}\ \ddot{a}_{x^{(G)}+t}^{(G)}}$$

is the time-t mortality experience adjustment factor and

$$\mathrm{IEA}_t = \exp(r_t^{\mathrm{PF}} - h)$$

is the time-t investment experience adjustment factor.

This version of the mortality experience adjustment includes sums of annuity factors for different individuals across various groups. Note that these mortality and investment adjustment factors are the same for each member. Ultimately, this means that every member's benefit changes by the same factor, δ_t, regardless of their group G.

An important feature of these adjustment factors is their *immediate* nature. Benefits are recalibrated at the end of each period based on realized mortality and investment outcomes, so that any deviation from expectations is reflected in the following year's payments. This immediacy preserves the solvency of the pool, since payouts are always aligned with available assets, and ensures that members do not accumulate unfunded promises that would need to be corrected in later periods. The adjustment mechanism, therefore, acts as a self-correcting tool; as soon as the experience of the pool diverges from the actuarial assumptions, benefits adjust for all surviving members.

This property stands in contrast with other potential designs that might rely on delayed or smoothed adjustments (see, e.g., Guillén et al., 2006; Cui et al., 2011; Goecke, 2013; Bégin, 2020; Chen et al., 2023; Bégin et al., 2026b). For example, one could imagine spreading the impact of an unfavorable mortality shock or investment shortfall over several future periods, thereby reducing the volatility of benefits in any single year. Such smoothing mechanisms may be attractive from a behavioral perspective, as they dampen year-to-year fluctuations. However, smoothing introduces intertemporal transfers among groups and raises questions of fairness, since some members may die before the full adjustment is applied.

Another alternative would be to base adjustments on cohort-specific experience, rather than pooling across the whole pool (see, e.g., Fullmer, 2019; Bégin and Sanders, 2023b). In this case, younger and older groups would see their benefits evolve according to their own survival outcomes, potentially reducing perceived cross-subsidies. Yet, this approach sacrifices the diversification benefits of larger, combined pools and could expose small cohorts to significant volatility if a few unexpected deaths occur. By contrast, the immediate, whole-pool adjustment mechanism described above strikes a balance; it treats all members uniformly within the pool, ensures financial sustainability at every step, and avoids the complexities of cohort-specific corrections.

2.4 EXAMPLE OF THE WHOLE-POOL ADJUSTMENT CALCULATION

A simple numerical example with two groups is then considered to illustrate the operation of this stylized lifetime pension pool. In practice, one would generally expect more members and potentially more groups, too; the full demographic model used in the analysis is presented later in this report.

Copyright © 2026 Society of Actuaries Research Institute

Assume a pool of 100 female members, half of whom are aged 65 years old at time 0 and are high-income members; the other half of the members are aged 75 years old and are low-income. The former group is denoted by $G = 1$ and the latter by $G = 2$. Assume that the hurdle rate h is set to 4%. The actuarial present values of the annuities due associated with these groups are assumed to be $\ddot{a}_{65}^{(1)} = 15$ and $\ddot{a}_{75}^{(2)} = 10$.

High-income members deposit \$1,000,000 in the pool, whereas low-income members deposit \$500,000, so that the total assets at inception are \$75,000,000. These initial deposits allow members in the first group to receive a benefit of

$$B_0^{(1)} = \frac{1{,}000{,}000}{15} = 66{,}667, \text{ for } G = 1,$$

and members in the second group to receive a benefit of

$$B_0^{(2)} = \frac{500{,}000}{10} = 50{,}000, \text{ for } G = 2$$

at inception. After paying these benefits, the total asset value drops to

$$75{,}000{,}000 - \sum_{G=1}^{2} N_0^{(G)} B_0^{(G)} = 75{,}000{,}000 - (50)(66{,}667) - (50)(50{,}000) = 69{,}166{,}667.$$

This amount is then invested for a year in a portfolio earning an uncertain rate of return. Suppose that the continuously compounded rate of return realized in the first year is $r_1^{\text{PF}} = 3\%$; this leads to a time-1 asset value of

$$A_1 = 69{,}166{,}667 \times \exp(0.03) = 71{,}273{,}105.$$

In the adjustment factor calculation, it is assumed that five members die during the first year: two younger high-income members and three older low-income members. The time-1 post-redistribution benefit $B_1^{(G)}$ for each surviving member is given by

$$B_0^{(G)} \frac{A_1}{N_1^{(1)} B_0^{(1)} \ddot{a}_{66}^{(1)} + N_1^{(2)} B_0^{(2)} \ddot{a}_{76}^{(2)}} = B_0^{(G)} \frac{71{,}273{,}105}{(50-2)(66{,}667)(14.7185) + (50-3)(50{,}000)(9.8603)}$$
$$= \begin{cases} 67{,}617 & \text{if } G = 1 \\ 50{,}713 & \text{if } G = 2 \end{cases},$$

where $\ddot{a}_{66}^{(1)} = 14.7185$ and $\ddot{a}_{76}^{(2)} = 9.8603$ if $p_{65}^{(1)} = 0.99$ and $p_{75}^{(2)} = 0.95$. In other words, the time-1 benefits are the time-1 asset values divided across the survivors (while accounting for their updated annuity factors).

Identical adjustments are obtained when using the formula developed in Section 2.3:

$$\text{MEA}_1 = \left(\frac{N_0^{(1)} B_0^{(1)} p_{65}^{(1)} \ddot{a}_{66}^{(1)} + N_0^{(2)} B_0^{(2)} p_{75}^{(2)} \ddot{a}_{76}^{(2)}}{N_1^{(1)} B_0^{(1)} \ddot{a}_{66}^{(1)} + N_1^{(2)} B_0^{(2)} \ddot{a}_{76}^{(2)}} \right) = \left(\frac{71{,}989{,}412}{70{,}271{,}050} \right) = 1.0245,$$
$$\text{IEA}_1 = \exp(r_1^{\text{PF}} - h) = \exp(0.03 - 0.04) = 0.9901,$$

which led to a total adjustment factor of

$$\delta_1 = \text{MEA}(1) \times \text{IEA}(1) = 1.0143,$$

Copyright © 2026 Society of Actuaries Research Institute

and a benefit of

$$B_1^{(G)} = B_0^{(G)}\,\delta_1 = \begin{cases} 67{,}617 & \text{if } G = 1 \\ 50{,}713 & \text{if } G = 2 \end{cases}.$$

Even though investment returns were below the hurdle rate—which should have caused a decrease in the time-1 benefits—more members than expected died during the first year. This ultimately led to an adjustment factor that was higher than one, meaning that the time-1 benefits were larger than those paid at time 0. By applying the same logic recursively, one can obtain the benefit for each year and member once the realized mortality and rate of return are observed.

Copyright © 2026 Society of Actuaries Research Institute

Section 3 Conceptual Insights on the Impact of Heterogeneity on Stability

This section examines a one-period lifetime pension pool following the framework introduced in Section 2. The goal of the section is to build intuition for how participant heterogeneity affects the variability of the benefit adjustment factor within such pools. Differences in characteristics, such as survival probabilities (via age) and initial benefit levels, are considered, and their effects on the pool's overall performance are analyzed. To capture benefit stability in a tractable way, the standard deviation of the adjustment—available via an approximation—is used as a proxy measure, and changes in its magnitude are studied under varying degrees of heterogeneity.

3.1 HOMOGENEOUS POOL

The standard deviation of benefit adjustments is first computed in the context of a homogeneous pool. This benchmark case provides a setting in which the behavior of the adjustment mechanism can be fully understood and serves as a point of comparison for subsequent analyses incorporating heterogeneity.

3.1.1 GENERAL SETUP

A closed pool with only one group such that $S = 1$ is considered. Consider a pool of $N_0^{(1)}$ individuals, all aged $x^{(1)}$, who enter a one-period lifetime pension pool. Each participant receives \$1 upon joining the pool (i.e., at time 0). A year later, each surviving member is targeted to receive a payment of \$1 if they are still alive (i.e., at time 1). The actual payment at time 1 will depend on the number of survivors. For simplicity, it is assumed that the pool invested in a risk-free asset that yielded a (continuously compounded) return of r over the next year, and that the hurdle rate h was set to the risk-free rate.

Assume that all individuals are in the same group, so that they are all identical copies of one another, with a survival probability of $p_{x^{(1)}}^{(1)}$ over the next year.[8] Each members pays

$$\ddot{a}_x^{(1)} = 1 + p_{x^{(1)}}^{(1)} \exp(-h)$$

at inception to join the pool. They then receive a payment of \$1 right away, and the rest is invested in the risk-free asset, yielding a fund value of

$$A_1 = N_0^{(1)} p_{x^{(1)}}^{(1)} \exp(-h) \exp(r) = N_0^{(1)} p_{x^{(1)}}^{(1)} \tag{5}$$

at time 1. This amount is then split among the survivors at time 1. The number of survivors, $N_1^{(1)}$, is assumed to be distributed as a binomial random variable in this example:

$$N_1^{(1)} \sim \text{Bin}\left(N_0^{(1)}, p_{x^{(1)}}^{(1)}\right).$$

As long as there are still individuals in the pool at time 1, the prospective value of the fund should be given by

[8] In this section, the survival probabilities are given by a Gompertz law with $m = 85$ and $b = 10$. This represents a modal value of 85 years and a dispersion coefficient of 10 years. The qualitative conclusions are robust to the choice of these parameters.

Copyright © 2026 Society of Actuaries Research Institute

$$A_1 = N_1^{(1)}\left(\delta_1 B_0^{(1)}\right) = N_1^{(1)}\delta_1, \tag{6}$$

where δ_1 is the adjustment factor at time 1. Because Equations (5) and (6) need to be equal, the formula is adjusted to:

$$\delta_1 = \frac{N_0^{(1)} p_{x^{(1)}}^{(1)}}{N_1^{(1)}}$$

at time 1, which is also the amount received by each survivor. This is defined if $N_1^{(1)} > 0$. If all members of the pool have died, however, it is assumed that the adjustment was $\delta_1 = 0$, as no further benefits were paid.

3.1.2 STANDARD DEVIATION OF THE ADJUSTMENT

One could proxy the pool's stability by calculating the standard deviation of the adjustment δ_1 (or its variance by taking the square of the standard deviation). Yet this could be very complicated, as it involves the variance of $1/N_1^{(1)}$. Therefore, an approximation for the standard deviation of δ_1, $\mathrm{SD}[\delta_1]$, is developed based on a second-order Taylor series expansion, which is used throughout Section 3.

Proposition 1. *The standard deviation of the adjustment* $\mathrm{SD}[\delta_1]$ *can be approximated by taking the square root of the quantity*

$$\left(N_0^{(1)} p_{x^{(1)}}^{(1)}\right)^2 q\left(\frac{1}{\mu^2} + \frac{3\sigma^2}{\mu^4}\right) - \left(N_0^{(1)} p_{x^{(1)}}^{(1)}\right)^2 q^2 \left(\frac{1}{\mu} + \frac{\sigma^2}{\mu^3}\right)^2,$$

where

$$\begin{aligned} q &= 1 - p_{N_1^{(1)}}(0), \\ \mu &= \frac{N_0^{(1)} p_{x^{(1)}}^{(1)}}{q}, \\ \sigma^2 &= \frac{N_0^{(1)} p_{x^{(1)}}^{(1)}\left(1 - p_{x^{(1)}}^{(1)}\right) + \left(N_0^{(1)} p_{x^{(1)}}^{(1)}\right)^2}{q} - \mu^2, \end{aligned}$$

and $p_{N_1^{(1)}}$ *is the probability mass function (pmf) of* $N_1^{(1)}$ *such that* $p_{N_1^{(1)}}(0) = \left(1 - p_{x^{(1)}}^{(1)}\right)^{N_0^{(1)}}$.

Proof. See Appendix A.1.

This approximation highlights that stability depends jointly on the pool size $N_0^{(1)}$ and the survival probability $p_{x^{(1)}}^{(1)}$. Intuitively, larger pools with higher survival probabilities reduce the likelihood of extreme adjustments, thereby lowering $\mathrm{SD}[\delta_1]$. However, the metric remains imperfect; it captures average variability but not the full distributional risk, particularly in the tails where rare events (such as complete depletion of survivors) may produce abrupt changes in payouts. Hence, while useful as a first step, this measure should only be interpreted as an initial proxy for stability rather than a definitive assessment.

Figure 1 shows the standard deviation of the adjustment as a function of the number of members $N_0^{(1)}$ and the age of the members $x^{(1)}$, which captures different mortality profiles. As expected, the uncertainty in the adjustment decreases as a function of the number of members and increases as a function of age;

Copyright © 2026 Society of Actuaries Research Institute

having more members or having younger members with better expected mortality experience decreases the standard deviation of the adjustment.[9]

Figure 1

STANDARD DEVIATION AS A FUNCTION OF THE NUMBER OF MEMBERS AND THEIR AGE FOR HOMOGENEOUS POOLS

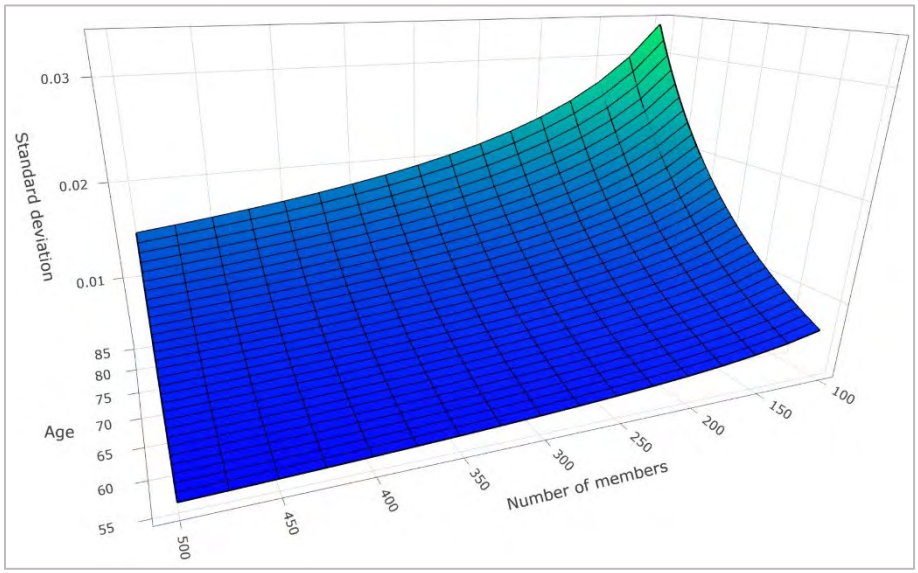

Notes: This figure illustrates the standard deviation of the adjustment as a function of the number of members and their age. Homogeneous pools are considered in this figure.

An important feature of homogeneous pools is the effect of *diversification* through numbers, which can be directly linked to the law of large numbers. When the pool size increases by adding identical participants, individual survival outcomes remain random, but the aggregate proportion of survivors becomes more predictable. In other words, idiosyncratic mortality risk is averaged out across a larger cohort. This reduces the variability of the adjustment factor δ_1, thereby improving stability.

From a mathematical perspective, as $N_0^{(1)}$ becomes very large, the realized survival rate $N_1^{(1)}/N_0^{(1)}$ converges to the expected survival probability $p_{x^{(1)}}^{(1)}$. Consequently, the adjustment δ_1 converges to 1, eliminating fluctuations that arise from sample variation. This highlights how stability is fundamentally a by-product of scale; small pools are exposed to significant randomness in outcomes, while large pools are less prone to those shocks.

Figure 2 reports similar results to Figure 1, but in a slightly different format: the standard deviation of the adjustment is showed for three representative entry ages—60, 65, and 70—as a function of the pool size. The results clearly illustrate the benefit of diversification; as the number of participants increases, the standard deviation of the adjustment declines, reflecting improved stability of the benefit distribution. The decline is most pronounced when moving from very small pools to moderately sized ones, with the steepest drop occurring as membership grows from just a handful of individuals to roughly 100. Beyond

[9] Note that the initial benefit amount is set to $B_0^{(1)} = 1$ without loss of generality here: having a higher or lower value does not impact the adjustment uncertainty. Changing its value would not change the conclusion of this section.

Copyright © 2026 Society of Actuaries Research Institute

that point, the marginal gains from adding further members become much smaller, and the curve gradually flattens out.

Figure 2
STANDARD DEVIATION AS A FUNCTION OF THE NUMBER OF MEMBERS FOR THREE AGES FOR HOMOGENEOUS POOLS

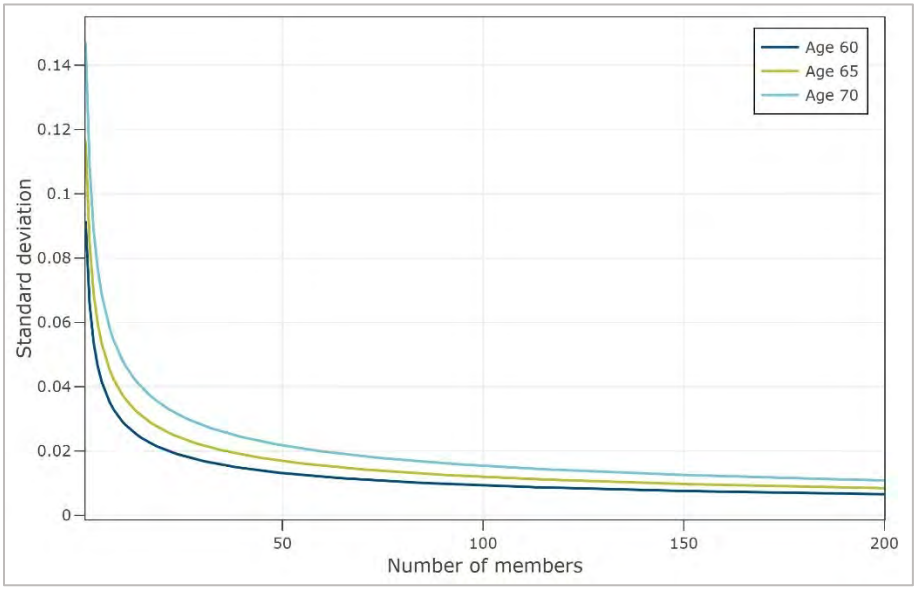

Notes: This figure illustrates the standard deviation of the adjustment as a function of the number of members for three ages: 60, 65, and 70. Homogeneous pools are considered in this figure.

This diversification effect provides a natural benchmark for understanding the role of heterogeneity in later sections. In homogeneous pools, stability improves mechanically with size. In heterogeneous pools, however, differences in survival probabilities and initial benefits mean that diversification is not perfect, and stability must instead be evaluated in terms of how those differences interact within the adjustment mechanism.

3.2 HETEROGENEOUS POOL

The analysis of stability is then extended by considering heterogeneous pools composed of groups that differed in characteristics such as age or initial benefit level. This setting introduces asymmetries in how risks are shared across members, making the adjustment mechanism more complex. Studying heterogeneity makes it possible to assess how differences between participants influence the pool's overall stability.

3.2.1 GENERAL SETUP

To compare and contrast the results of Section 3.1, a slightly different closed pool with two distinct groups of members is considered; that is, $S = 2$.

Consider $N_0^{(2)}$ individuals aged $x^{(2)}$ of a second group joining the pool at inception. They have a different survival probability; that is, $p_{x^{(2)}}^{(2)}$. Instead of having an initial benefit of \$1, the second group starts with a benefit of $\$y$ at time 0, such that $B_0^{(2)} = y$. These members, therefore, pay

$$y\,\ddot{a}_{x^{(2)}}^{(2)} = y\left(1 + p_{x^{(2)}}^{(2)} \exp(-h)\right)$$

Copyright © 2026 Society of Actuaries Research Institute

to join the pool. After paying \$1 to each member of the first group and $\$y$ to each member of the second group, the pool operator invests the remaining funds in the risk-free asset, yielding a fund value of

$$A_1 = \left(N_0^{(1)} p_{x^{(1)}}^{(1)} \exp(-h) + y\, N_0^{(2)} p_{x^{(2)}}^{(2)} \exp(-h)\right) \exp(r) = N_0^{(1)} p_{x^{(1)}}^{(1)} + y\, N_0^{(2)} p_{x^{(2)}}^{(2)}$$

at time 1. This amount is again split among the survivors; $N_1^{(1)}$ surviving members in the first group and $N_1^{(2)}$ in the second group, distributed as

$$N_1^{(2)} \sim \text{Bin}\left(N_0^{(2)}, p_{x^{(2)}}^{(2)}\right).$$

To obtain the benefit adjustment, the prospective value of the time-1 fund is considered:

$$A_1 = N_1^{(1)}\left(\delta_1 B_0^{(1)}\right) + N_1^{(2)}\left(\delta_1 B_0^{(2)}\right) = \delta_1\left(N_1^{(1)} + y\, N_1^{(2)}\right),$$

where δ_1 was the adjustment. Setting this equal to the retrospective value obtained above yielded

$$\delta_1 = \frac{N_0^{(1)} p_{x^{(1)}}^{(1)} + y\, N_0^{(2)} p_{x^{(2)}}^{(2)}}{N_1^{(1)} + y\, N_1^{(2)}},$$

so that the benefit paid to the first group of individuals was

$$\frac{N_0^{(1)} p_{x^{(1)}}^{(1)} + y\, N_0^{(2)} p_{x^{(2)}}^{(2)}}{N_1^{(1)} + y\, N_1^{(2)}}$$

and that paid to the second group was

$$y\,\frac{N_0^{(1)} p_{x^{(1)}}^{(1)} + y\, N_0^{(2)} p_{x^{(2)}}^{(2)}}{N_1^{(1)} + y\, N_1^{(2)}}$$

at time 1.

3.2.2 STANDARD DEVIATION OF THE ADJUSTMENT

The same strategy as in Section 3.1.1 is used to find an approximation of the adjustment's standard deviation.

Proposition 2. *The standard deviation of the adjustment* $\text{SD}[\delta_1]$ *can be approximated by taking the square root of the quantity*

$$\left(N_0^{(1)} p_{x^{(1)}}^{(1)} + y\, N_0^{(2)} p_{x^{(2)}}^{(2)}\right)^2 q\left(\frac{1}{\mu^2} + \frac{3\sigma^2}{\mu^4}\right) - \left(N_0^{(1)} p_{x^{(1)}}^{(1)} + y\, N_0^{(2)} p_{x^{(2)}}^{(2)}\right)^2 q^2 \left(\frac{1}{\mu} + \frac{\sigma^2}{\mu^3}\right)^2,$$

where

$$q = \left(1 - p_{N_1^{(1)}}(0)\right) + \left(1 - p_{N_1^{(2)}}(0)\right) - \left(1 - p_{N_1^{(1)}}(0)\right)\left(1 - p_{N_1^{(2)}}(0)\right),$$

$$\mu = \frac{N_0^{(1)} p_{x^{(1)}}^{(1)} + y\, N_0^{(2)} p_{x^{(2)}}^{(2)}}{q},$$

Copyright © 2026 Society of Actuaries Research Institute

$$\sigma^2 = \frac{\begin{pmatrix} N_0^{(1)} p_{x^{(1)}}^{(1)} \left(1 - p_{x^{(1)}}^{(1)}\right) + \left(N_0^{(1)} p_{x^{(1)}}^{(1)}\right)^2 + 2y N_0^{(1)} p_{x^{(1)}}^{(1)} N_0^{(2)} p_{x^{(2)}}^{(2)} \\ + y^2 N_0^{(2)} p_{x^{(2)}}^{(2)} \left(1 - p_{x^{(2)}}^{(2)}\right) + y^2 \left(N_0^{(2)} p_{x^{(2)}}^{(2)}\right)^2 \end{pmatrix}}{q} - \mu^2,$$

and $p_{N_1^{(2)}}$ *is the pmf of* $N_1^{(2)}$.

Proof. See Appendix A.2.

3.2.3 IMPACT OF AGE

The impact of introducing groups with different ages is investigated, which corresponded to different survival probabilities. This subsection focuses on cases where additional members were added to the pool, rather than fixing the total pool size and examining its composition. For comparison, Appendix B.1 analyzes the alternative setting of a fixed-size pool divided into two groups.

Figure 3

STANDARD DEVIATION AS A FUNCTION OF THE NUMBER OF MEMBERS IN THE SECOND GROUP AND THEIR AGES

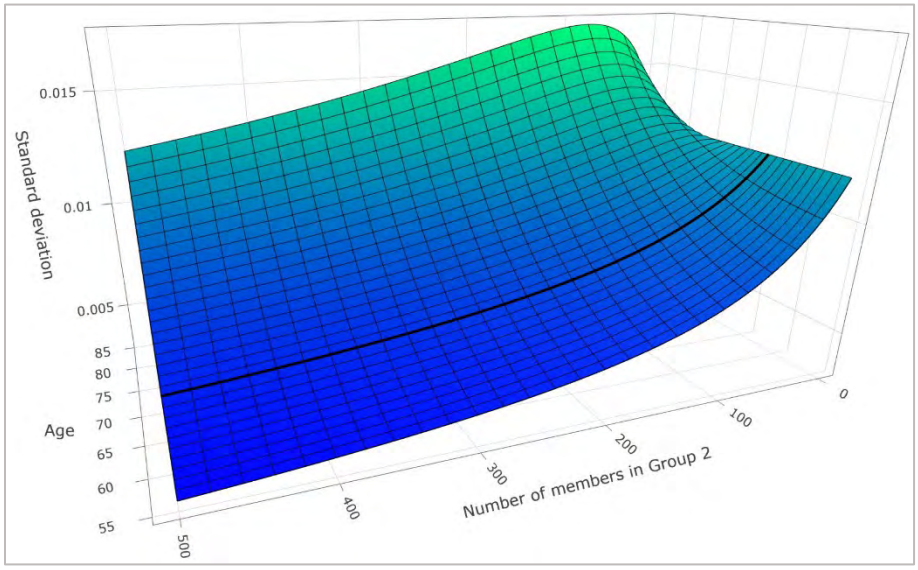

Notes: This figure shows the standard deviation of the adjustment as a function of the number of members in the second group $N_0^{(2)}$ and their age $x^{(2)}$, assuming $N_0^{(1)} = 100$, $x^{(1)} = 65$, and $y = 1$.

Figure 3 shows the standard deviation of the adjustment as a function of the number of individuals in the second group $N_0^{(2)}$ and their age $x^{(2)}$, assuming $N_0^{(1)} = 100$, $x^{(1)} = 65$, and $y = 1$. This figure assumed that the size of the pool could increase by the inclusion of additional members from the second group. The black line represents the standard deviation when the members in the second group are the same age as those in the first group, which is consistent with the results in Section 3.1. When a few individuals who differed from those in the first group were added, smaller standard deviations were observed if the individuals in the second group had better or similar mortality—younger than 65 years old—and larger standard deviations if they were older than 65. As more and more individuals were added, the standard deviation eventually began to decline, even if the new members had worse mortality than the original members. The number of individuals needed in order to start seeing a decrease in the standard deviation of the adjustment depends on how different the survival probability is between the two groups.

Copyright © 2026 Society of Actuaries Research Institute

Figure 4 presents results analogous to those in Figure 3, but now for five specific ages: 55, 60, 65, 70, and 75. Introducing members with better mortality prospects than the 65-year-olds in the original pool enhances stability; the dark blue and green curves decline more rapidly than the light blue curve, which represents the homogeneous benchmark. By contrast, incorporating members with worse mortality increases variability; the yellow and red curves generally remain above the homogeneous case. The figure also allows for a direct comparison across pools, highlighting the role of heterogeneity. For example, a homogeneous pool of 230 members (100 initial members plus 130 additional ones) achieves a stability level comparable to that obtained by adding 500 members aged 75.

Figure 4

STANDARD DEVIATION AS A FUNCTION OF THE NUMBER OF MEMBERS IN THE SECOND GROUP AND THEIR AGES

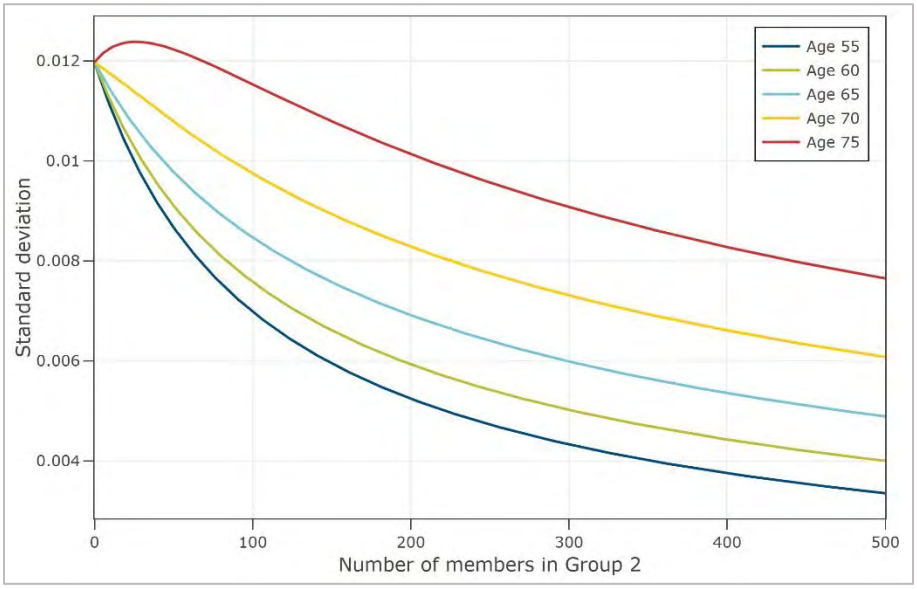

Notes: This figure shows the standard deviation of the adjustment as a function of the number of members in the second group $N_0^{(2)}$ and their age $x^{(2)}$, assuming $N_0^{(1)} = 100$, $x^{(1)} = 65$, and $y = 1$. The lines are calculated for different ages for the second group, $x^{(2)}$.

3.2.4 IMPACT OF THE INITIAL BENEFIT

The impact of heterogeneity in an individual's investment in a lifetime pension pool has been examined before in the literature. Bernhardt and Qu (2024) studied how differences in members' initial savings affected the income stability of participants in a closed lifetime pension pool. Distinct from much of the prior work that assumed homogeneous members, their analysis isolated the role of unequal savings by holding all other characteristics constant. Donnelly (2023) also addressed the issue of heterogeneity in the initial benefit levels, among other sources of variation, and introduced a heterogeneity measure that captured the degree of dispersion among participants. Building on this line of research, this subsection explores the impact of heterogeneity but does so using an approximation to the standard deviation of adjustments.

Figure 5 reports the standard deviation of the adjustment as a function of the number of individuals in the second group $N_0^{(2)}$ and their initial benefit y, assuming $N_0^{(1)} = 100$, and that members in both groups are age 65; that is, $x^{(1)} = 65$, and $x^{(2)} = 65$. The black line depicts the standard deviation when the additional members are homogeneous in terms of the initial benefit, serving as a benchmark for comparison with Section 3.1. Introducing a small number of heterogeneous members—those whose initial benefits differ

Copyright © 2026 Society of Actuaries Research Institute

from Group 1—initially increases the standard deviation, regardless of whether their benefits are lower or higher. As the number of such members grows, however, the standard deviation declines, reflecting a law-of-large-numbers effect, even though the added members are not identical copies of the originals. The point at which this decline sets in depends on the initial benefit level y of Group 2.

Figure 6, completing the information shown in Figure 5, underscores an important distinction in how different forms of heterogeneity affect stability. The most favorable outcome is obtained when additional members are identical to the existing ones, thereby reinforcing homogeneity. In this case, stability improves monotonically with pool size, consistent with the law of large numbers. By contrast, introducing heterogeneity in the form of different initial benefit amounts reduces stability relative to the homogeneous benchmark. No matter whether the second group has higher or lower benefits, the resulting pool performs worse in terms of standard deviation than one composed solely of homogeneous members.

Figure 5

STANDARD DEVIATION AS A FUNCTION OF THE NUMBER OF MEMBERS IN THE SECOND GROUP AND THEIR INITIAL BENEFIT LEVEL

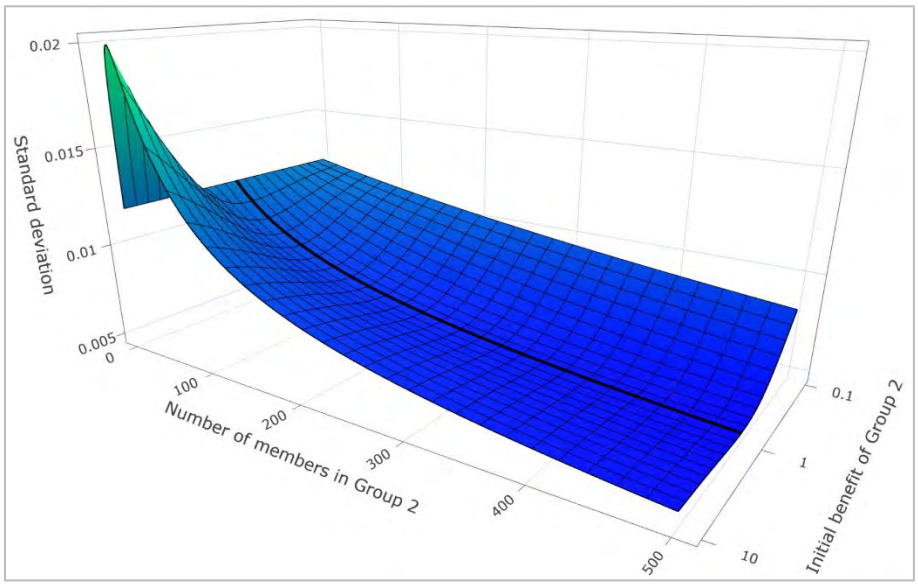

Notes: This figure shows the standard deviation of the adjustment as a function of the number of members in the second group $N_0^{(2)}$ and their initial benefit level y, assuming $N_0^{(1)} = 100$ and $x^{(1)} = x^{(2)} = 65$.

This finding stands in sharp contrast to the effect of heterogeneity in survival probability or age considered earlier. In that setting, the interaction between groups created diversification benefits; carefully mixing cohorts could produce outcomes that are superior to those of a homogeneous pool. When heterogeneity stems from benefit amounts, however, such beneficial diversification is absent—heterogeneity only dilutes stability. This comparison highlights that the nature of heterogeneity matters; differences in mortality can sometimes be advantageous, but differences in payouts are always detrimental when measured against the homogeneous benchmark.

3.2.5 IMPACT OF BOTH DIMENSIONS

The analysis is then extended by considering heterogeneity along two dimensions simultaneously: age, or survival probability, and the initial benefit level. By varying both factors simultaneously, a more complete picture is obtained of how heterogeneous membership influences outcomes, beyond the effect of each source of variation taken in isolation, as in Sections 3.2.3 and 3.2.4.

Copyright © 2026 Society of Actuaries Research Institute

Figure 6

STANDARD DEVIATION AS A FUNCTION OF THE NUMBER OF MEMBERS IN THE SECOND GROUP AND THEIR INITIAL BENEFIT LEVEL

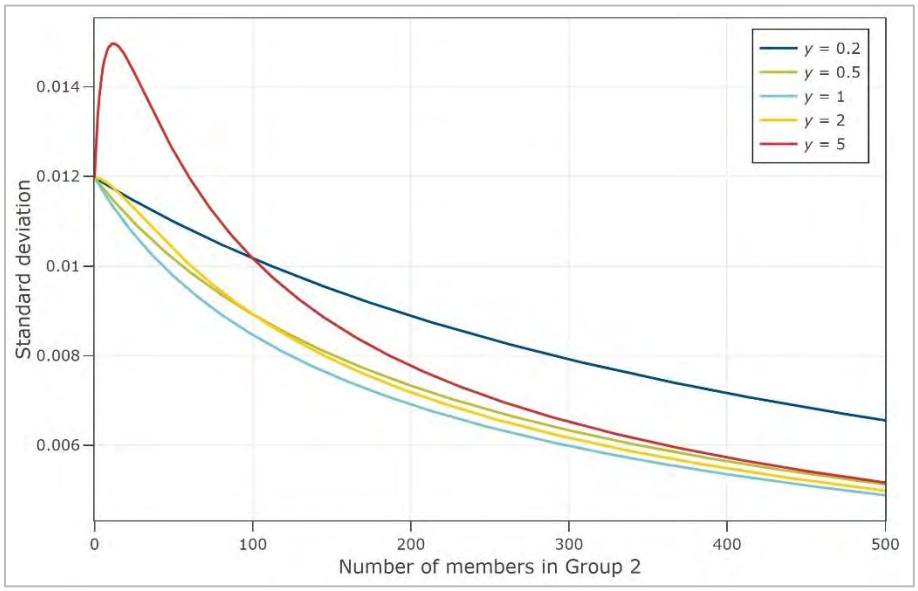

Notes: This figure shows the standard deviation of the adjustment as a function of the number of members in the second group $N_0^{(2)}$ and their initial benefit level y, assuming $N_0^{(1)} = 100$ and $x^{(1)} = x^{(2)} = 65$. The lines are calculated for different initial benefit levels for the second group, y.

Figure 7 reports the standard deviation of the adjustment as a function of the initial benefit level of Group 2, y, for three different ages of Group 2 (60, 65, and 70). The figure shows that the point of minimum standard deviation (marked by vertical dotted lines) differs across these cases: for 60-year-olds, the lowest variability occurs when their initial benefit level exceeds that of Group 1, whereas for 70-year-olds, the minimum is attained when y is below one. This suggests that younger entrants should contribute more capital to the pool—generally translating into higher benefits—which, in turn, improves stability. For reasonable values of y, the pool as a whole benefit from expansion (i.e., moving from 100 to 200 members). Relative to a homogeneous pool of 200 members, however, outcomes depend on the specific age–benefit profile: in some cases, the heterogeneous pool performs better (e.g., 60 years old with an initial benefit of $y = 2$), while in others it performs worse (e.g., 70 years old with an initial benefit of $y = 0.7$).

Figure 8 examines the reverse perspective of Figure 7, focusing on how the age of members in Group 2 affects the standard deviation for different values of the initial benefit y. When the initial benefit of Group 2 is small, the effect of age is limited; the dark blue and green lines remain relatively flat, indicating little sensitivity to changes in the age $x^{(2)}$. By contrast, when the initial benefit is larger, age becomes a critical factor, with substantial variation in the standard deviation across different $x^{(2)}$ values; the yellow and red curves report such increases. This pattern suggests that pool operators need only be concerned with the age composition of an additional group when its benefit level is disproportionately large.

Having established the overall behavior of the adjustment's standard deviation, contour plots of the standard deviation surface are then used to provide a convenient way to visualize how stability varies jointly with the two characteristics of Group 2: age on the x-axis and the initial benefit level on the y-axis. In Figure 9, contour lines trace combinations of age and initial benefit that yield the same standard deviation, thereby mapping the surface of stability outcomes. The color shading highlights regions of particular interest; red areas correspond to pool compositions that perform worse than the first group alone,

Copyright © 2026 Society of Actuaries Research Institute

indicating reduced stability, while green areas identify compositions that perform better than a fully homogeneous pool of the same size (i.e., with 200 members). Thus, these plots illustrate that heterogeneity does not always erode stability; under certain combinations of ages and initial benefits, it can even enhance it for one of the two groups (i.e., Group 1 in this case). More generally, any composition outside the red region can be interpreted as a favorable pooling arrangement from Group 1's perspective.

Figure 7

STANDARD DEVIATION AS A FUNCTION OF THE INITIAL BENEFIT LEVEL OF MEMBERS IN THE SECOND GROUP FOR THREE AGE GROUPS

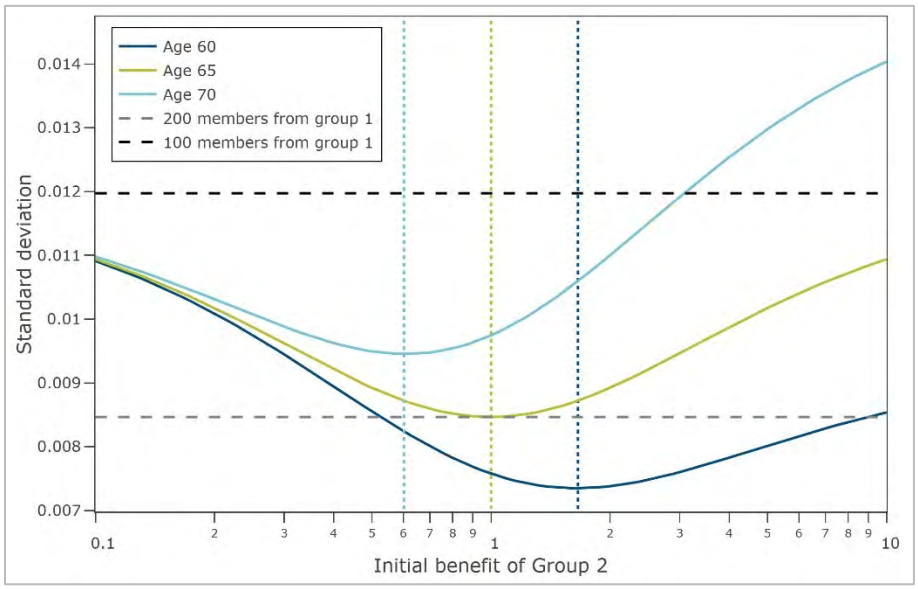

Notes: This figure shows the standard deviation of the adjustment as a function of the initial benefit level of members in the second group y, assuming $N_0^{(1)} = N_0^{(2)} = 100$ and $x^{(1)} = 65$. The lines are calculated for different ages for the second group, $x^{(2)}$. Horizontal dashed lines represent the standard deviation for a homogeneous group of 100 and 200 members in black and gray, respectively; these lines are used for comparison purposes.

It is equally important to consider how the second group perceives heterogeneous pools. In Figure 9, the analysis has focused on the perspective of the first group, but stability gains are meaningful only if both groups benefit. Figure 10 presents the same contour plot as in Figure 9, but here the red areas mark pool compositions that are worse than remaining independent for at least one group. As expected, the upper-right corner remains red, since Group 1 already judged those pools as unfavorable. In addition, the bottom-left region is also red, indicating cases where Group 2 is younger and has lower initial benefits than Group 1. From their perspective, this implies pooling with an older group with higher initial benefits—an arrangement that has previously been shown to yield worse outcomes.

Notably, no green region appears in this plot. For a heterogeneous pool to be preferable to a homogeneous one of the same size, one group must include younger members. Since one group must always be older than the other, it is impossible for both groups to perceive such an advantage simultaneously. It is, therefore, futile to search for group combinations that would make both groups better off than in the homogeneous case; some degree of cross-subsidization between groups is inevitable. This does not imply that combining the groups is undesirable—in fact, both may still enjoy improved stability relative to remaining on their own. Rather, the point is that no heterogeneous arrangement can simultaneously outperform the homogeneous benchmark from the perspective of both groups.

Copyright © 2026 Society of Actuaries Research Institute

Figure 8

STANDARD DEVIATION AS A FUNCTION OF AGE OF MEMBERS IN THE SECOND GROUP FOR FIVE INITIAL BENEFIT LEVELS

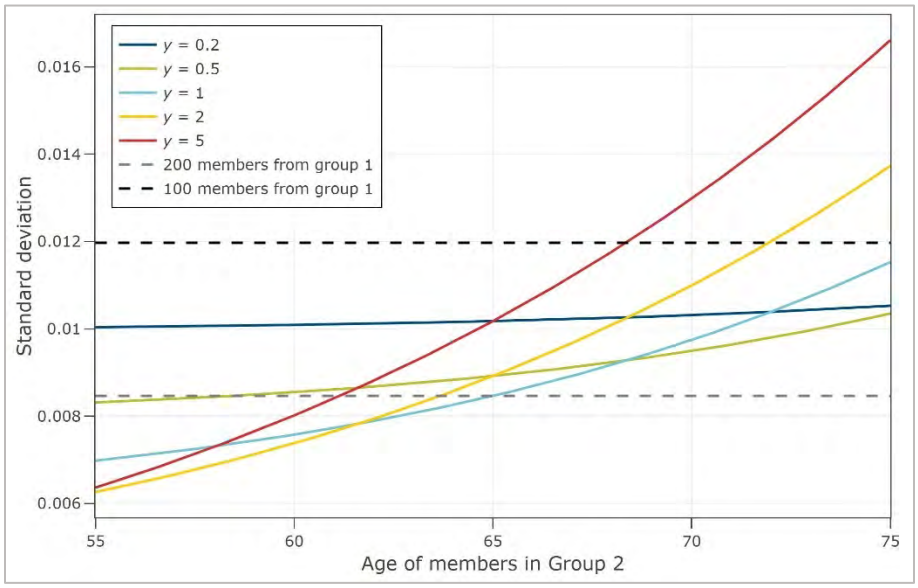

Notes: This figure shows the standard deviation of the adjustment as a function of the age of members in the second group $x^{(2)}$, assuming $N_0^{(1)} = N_0^{(2)} = 100$ and $x^{(1)} = 65$. The lines are calculated for different levels of initial benefit for the second group, y. Horizontal dashed lines represent the standard deviation for a homogeneous group of 100 and 200 members in black and gray, respectively; these lines are used for comparison purposes.

Figure 9

CONTOUR PLOT OF STANDARD DEVIATION SURFACE AND PREFERENCE REGIONS FOR THE FIRST GROUP

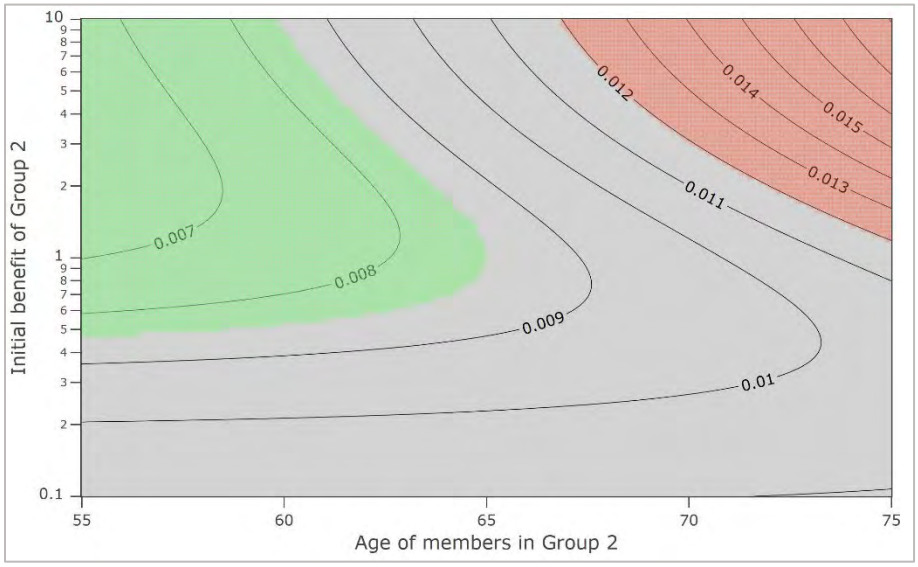

Notes: This figure shows the contour plot of the surface of standard deviation of the adjustment as a function of the age of members in the second group $x^{(2)}$ and their initial benefit y, assuming $N_0^{(1)} = N_0^{(2)} = 100$ and $x^{(1)} = 65$. The contour lines represent combinations of group age and initial benefit that yield the same level of standard deviation, thereby mapping the stability surface of the pool. The red region corresponds to outcomes with a higher standard deviation than group 1 on its own, indicating reduced stability. The green region highlights outcomes with a lower standard deviation than in a fully homogeneous pool, showing where heterogeneity can be beneficial.

Copyright © 2026 Society of Actuaries Research Institute

3.3 IMPLIED NUMBER OF HOMOGENEOUS POOL MEMBERS

So far, the report has assessed stability by reporting the standard deviation of the adjustment. While this raw measure captures variability, the numbers are not immediately interpretable on their own. To provide more intuition, an approach similar to that of Bernhardt and Qu (2024) was followed, who introduced the concept of an implied number of homogeneous members. Specifically, the standard deviation observed in a heterogeneous pool is translated into the size of a homogeneous pool from the first group that would deliver the same level of stability. This transformation provides a natural benchmark; it makes it possible to evaluate heterogeneous pools in terms of their homogeneous equivalents, thereby quantifying directly how much stability is gained or lost through heterogeneity. Figure 11 mirrors the analysis of Figure 9, whereas Figure 12 provides the corresponding counterpart to Figure 10. Although the contour levels differ across figures (i.e., comparing Figure 9 and Figure 11 or Figure 10 and Figure 12), the location and shape of the regions remain unchanged. This consistency means that the underlying results are not affected, only the scale at which they are displayed. Adjusting the contour levels in this way enhances readability and makes the numerical values more easily interpretable for practitioners, while preserving the essential structure of the red and green regions that inform the conclusions.

Figure 10

CONTOUR PLOT OF STANDARD DEVIATION SURFACE AND PREFERENCE REGIONS FOR THE FIRST GROUP

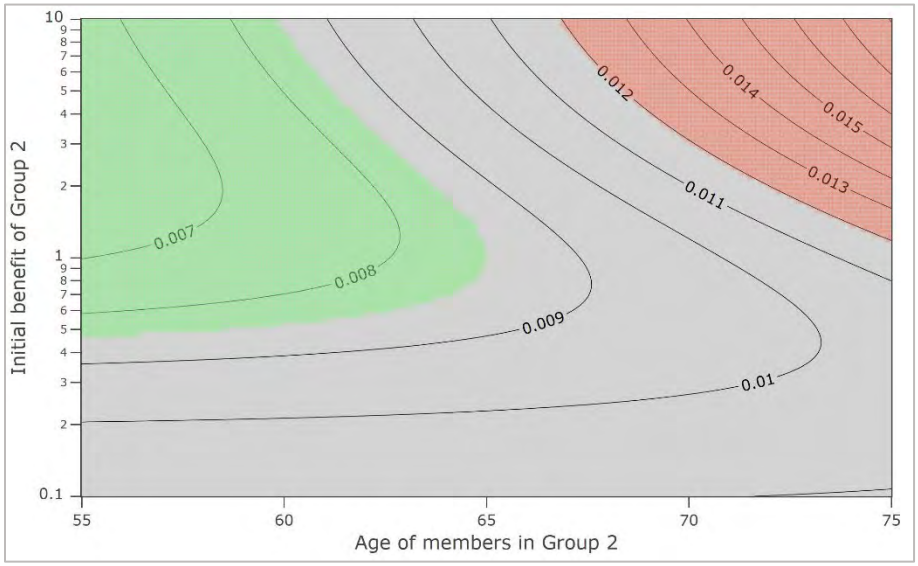

Notes: This figure shows the contour plot of the surface of standard deviation of the adjustment as a function of the age of members in the second group $x^{(2)}$ and their initial benefit y, assuming $N_0^{(1)} = N_0^{(2)} = 100$ and $x^{(1)} = 65$. The contour lines represent combinations of group age and initial benefit that yield the same level of standard deviation, thereby mapping the stability surface of the pool. The red region corresponds to outcomes with a higher standard deviation than group 1 on its own, indicating reduced stability. The green region highlights outcomes with a lower standard deviation than in a fully homogeneous pool, showing where heterogeneity can be beneficial.

Copyright © 2026 Society of Actuaries Research Institute

Figure 11

CONTOUR PLOT OF STANDARD DEVIATION SURFACE AND PREFERENCE REGIONS FOR BOTH GROUPS

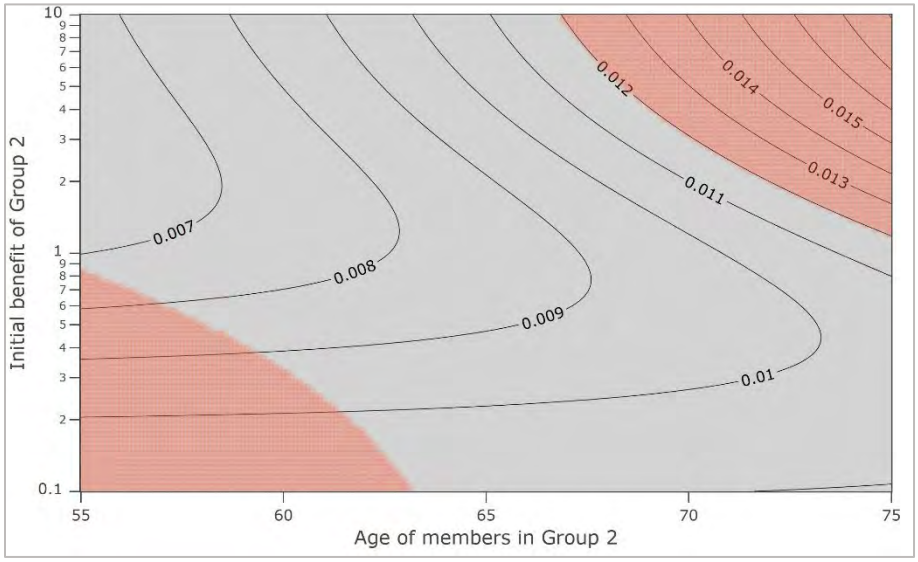

Notes: This figure shows the contour plot of the surface of standard deviation of the adjustment as a function of the age of members in the second group $x^{(2)}$ and their initial benefit y, assuming $N_0^{(1)} = N_0^{(2)} = 100$ and $x^{(1)} = 65$. The contour lines represent combinations of group age and initial benefit that yield the same level of standard deviation, thereby mapping the stability surface of the pool. The red regions correspond to outcomes where at least one of the two groups experiences a higher standard deviation than if it were on its own, indicating reduced stability.

Figure 12

CONTOUR PLOT OF STANDARD DEVIATION SURFACE EXPRESSED AS THE IMPLIED NUMBER OF HOMOGENEOUS MEMBERS AND PREFERENCE REGIONS FOR THE FIRST GROUP

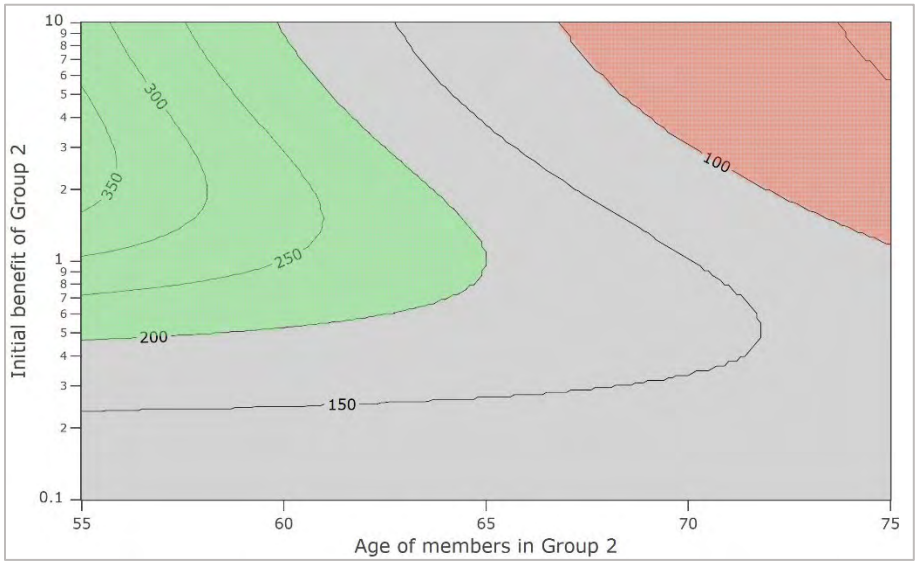

Notes: This figure shows the contour plot of the surface of standard deviation of the adjustment expressed as the implied number of homogeneous members as a function of the age of members in the second group $x^{(2)}$ and their initial benefit y, assuming $N_0^{(1)} = N_0^{(2)} = 100$ and $x^{(1)} = 65$. The contour lines represent combinations of group age and initial benefit that yield the same level of implied number of homogeneous members, thereby mapping the stability surface of the pool. See Figure 9's caption for more details.

Copyright © 2026 Society of Actuaries Research Institute

Figure 13

CONTOUR PLOT OF STANDARD DEVIATION SURFACE EXPRESSED AS THE IMPLIED NUMBER OF HOMOGENEOUS MEMBERS AND PREFERENCE REGIONS FOR BOTH GROUPS

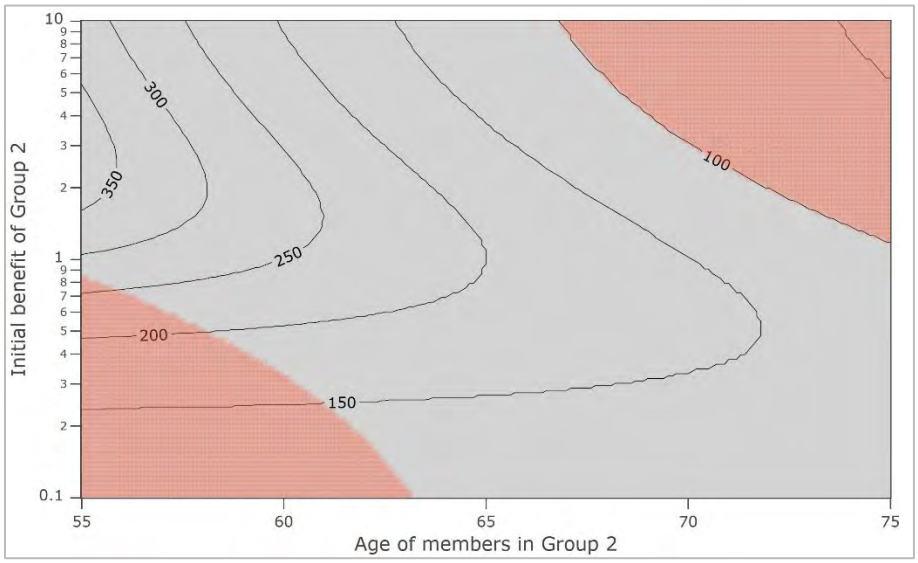

Notes: This figure shows the contour plot of the surface of standard deviation of the adjustment expressed as the implied number of homogeneous members as a function of the age of members in the second group $x^{(2)}$ and their initial benefit y, assuming $N_0^{(1)} = N_0^{(2)} = 100$ and $x^{(1)} = 65$. The contour lines represent combinations of group age and initial benefit that yield the same level of implied number of homogeneous members, thereby mapping the stability surface of the pool. See Figure 10's caption for more details.

Copyright © 2026 Society of Actuaries Research Institute

Section 4 A Deeper Look at Stability

As discussed in Section 3, using the standard deviation of the (one-year) adjustment provides an intuitive and convenient way to describe stability in a lifetime pension pool. The standard deviation can be easily approximated and even expressed as the implied number of homogeneous members that would generate the same stability. However, this measure has limitations; it reflects only short-term fluctuations in benefits, lacks standalone interpretability over the entire lifetime of the pool, and cannot capture the tail risk of extreme adjustments since the standard deviation summarizes only typical, average variation.

4.1 INTRODUCING ANOTHER MEASURE OF STABILITY: THE STABLE INCOME PERIOD

To evaluate income stability more comprehensively, the SIP is introduced, first proposed by Bernhardt and Donnelly (2021). This time-based proxy asks: for how long can we be confident that a member's income will remain within specified bounds around its initial level? Unlike short-term volatility metrics, such as the standard deviation considered in Section 3, the SIP captures the persistence of stability over the life of the pool. Members and other stakeholders can, therefore, interpret it directly as: "income is stable for at least 20 years with 90% confidence."

Formally, for tolerances $\varepsilon_1 \in (0,1)$ and $\varepsilon_2 \in (0,\infty)$, and confidence level $\beta \in (0,1)$, the SIP is defined as

$$\mathrm{SIP}(\varepsilon_1, \varepsilon_2, \beta) = \max\left[t \in \mathbb{N} : \mathbb{P}\left[B_0^{(G)}(1-\varepsilon_1) \le B_k^{(G)} \le B_0^{(G)}(1+\varepsilon_2)\,,\ \forall\, k \le t\right] \ge \beta\right], \qquad (7)$$

In words, the SIP is the maximum number of years T such that, with probability at least β, the member's income remains between the lower bound $B_0^{(G)}(1-\varepsilon_1)$ and the upper bound $B_0^{(G)}(1+\varepsilon_2)$ for every period up to T. This SIP definition is very flexible: ε_1, ε_2, and β can be tailored to reflect different profiles of risk tolerances. Note that if ε_2 is set to infinity, the measure focuses exclusively on downward deviations, disregarding any upside fluctuations.

This formulation was first introduced by Bernhardt and Donnelly (2021) in the homogeneous case and extended by Bernhardt and Qu (2024) and Donnelly (2023) to heterogeneous pooled funds. Its strength lies in combining an explicit probabilistic guarantee with an interpretable time horizon, thereby capturing both the magnitude and the duration of stability in income streams.[10]

Table 1
COMPARISON OF STANDARD DEVIATION AND SIP

Feature	Standard Deviation of the Adjustment	Stable Income Period
Type of measure	One-period, variance-based	Time-horizon, probability-based
What it captures	Typical short-term fluctuations	Maximum period that income remains within bounds with high probability
Interpretability	Expressed in units of variability; requires translation into an implicit number of homogeneous members	Directly interpretable as stable for T years with 100β% confidence
Focus	Normal variation	Persistence of stability and protection against adverse outcomes
Tail risk sensitivity	Limited	Explicitly incorporates the probability of breaching specified bounds

[10] Bernhardt and Donnelly (2021) and Bernhardt and Qu (2024) proposed an analytical approximation of the SIP under certain restrictions. They did so because finding a value is often too complex analytically, even under a simple sharing rule like those of Section 2.

Copyright © 2026 Society of Actuaries Research Institute

This formulation is more complicated than the simpler stability measure introduced in Section 3. Whereas the earlier approach provided a straightforward variance-based snapshot of short-term fluctuations, the SIP introduces an explicitly time-based perspective. This added complexity reflects the richer information conveyed by the measure, making it possible to assess not only the variability of income at a given point, but also the duration over which income stability can be expected with high confidence. Table 1 summarizes the differences between the two measures.

4.2 COMPUTATION OF THE STABLE INCOME PERIOD MEASURE

The SIP calculation relies on the joint distribution of the cumulative adjustments because the probability in Equation (7) can be rewritten as

$$\begin{aligned}&\mathbb{P}\left[B_0^{(G)}(1-\varepsilon_1) \le B_k^{(G)} \le B_0^{(G)}(1+\varepsilon_2)\,,\ \ \forall\, k \le t\right]\\ &= \mathbb{P}\left[(1-\varepsilon_1) \le \prod_{i=1}^{k} \delta_i \le (1+\varepsilon_2)\,,\ \ \forall\, k \le t\right],\end{aligned}$$

where the adjustment δ_i is defined in Equation (4). Put simply, this calculation yields the probability that the cumulative adjustments up to time t stay within $[1-\varepsilon_1, 1+\varepsilon_2]$. The SIP is determined by reference to the first time when this probability falls below the confidence level β. Accordingly, it suffices to compute the measure only up to this point, since once the probability drops below β, it cannot rise above it again.

Computing the exact value of these probabilities, even for a single t, is extremely complex and computationally expensive. For this reason, Monte Carlo simulations are used to evaluate the distribution (see Marchand et al., 2026, for more details). Specifically, realizations of the cumulative adjustment factors are simulated. This yields an empirical probability of stable adjustments. The SIP, as defined in Equation (7), is then obtained by taking the β^{th} quantile of this distribution. An illustration of the distribution is provided in Figure 13, where the blue vertical dotted line indicates the SIP for $\beta = 0.95$.

Figure 14
EMPIRICAL PROBABILITY OF STABLE ADJUSTMENTS

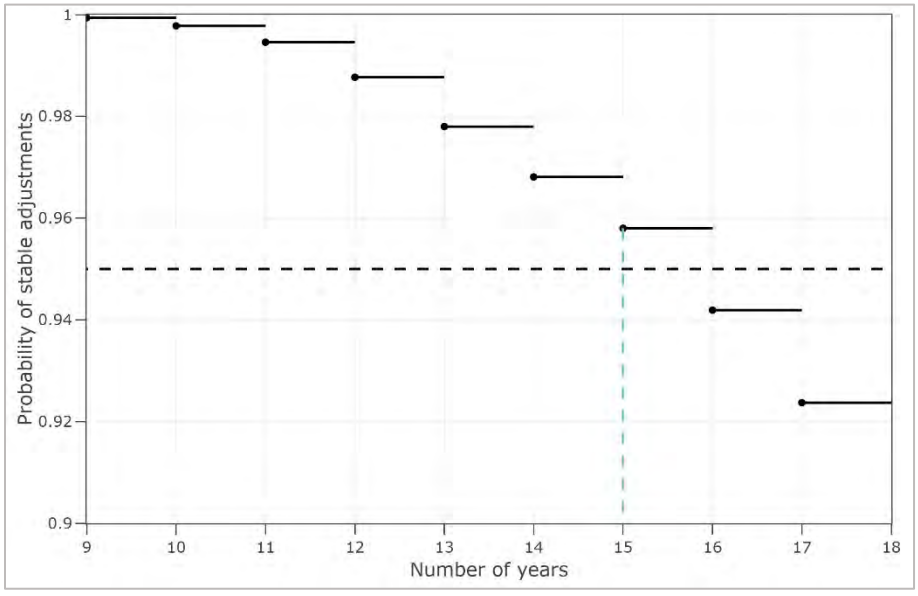

Notes: The figure reports the empirical probability of stable adjustments. The black horizontal dashed line represents the 95$^{\text{th}}$ quantile of this distribution (i.e., β^{th} quantile), whereas the light blue vertical dashed line shows the SIP for this example.

Copyright © 2026 Society of Actuaries Research Institute

4.3 HOMOGENEOUS POOL

The SIP is applied to homogeneous pools as done in Section 3.1. This allows building intuition for how the SIP behaves under simple, controlled settings before turning into more complex heterogeneous cases. By considering identical participants of a given entry age, the effects of pool size and age profile on stability can be isolated, making it easier to understand how the SIP reflects the underlying mortality dynamics. In this section, unless otherwise stated, the values $\varepsilon_1 = 0.1$, $\varepsilon_2 = \infty$, and $\beta = 0.95$ are used.

Figure 14 shows the SIP as a function of the number of members in a homogeneous pool and the entry age of the members for three ages: 60, 65, and 70. The SIP provides a complementary perspective to the standard deviation analysis; rather than focusing on the variability of the first adjustment factor, it measures how long members can expect their income stream to remain stable before mortality shocks cause significant adjustments.

The results illustrate two key effects. First, as the number of members increases, the SIP lengthens substantially, reflecting the diversification benefits of larger pools. This again highlights the law of large numbers; with more members, individual mortality outcomes average out, and the collective income path remains smoother for longer. Second, entry age has a clear influence on stability. Pools of younger members (e.g., age 60) exhibit longer SIPs compared to pools of older members (e.g., age 70), since younger cohorts have higher survival probabilities and less variability in early mortality outcomes.

The shape of the curves is also informative. The SIP rises sharply when moving from very small pools to moderately sized ones, meaning that even a modest increase in membership significantly extends stability. However, beyond a few hundred members, the marginal benefit of adding more participants diminishes; the curves flatten, showing that very large pools deliver relatively smaller additional gains in stability.

Figure 15

SIP AS A FUNCTION OF THE NUMBER OF MEMBERS FOR THREE AGES.

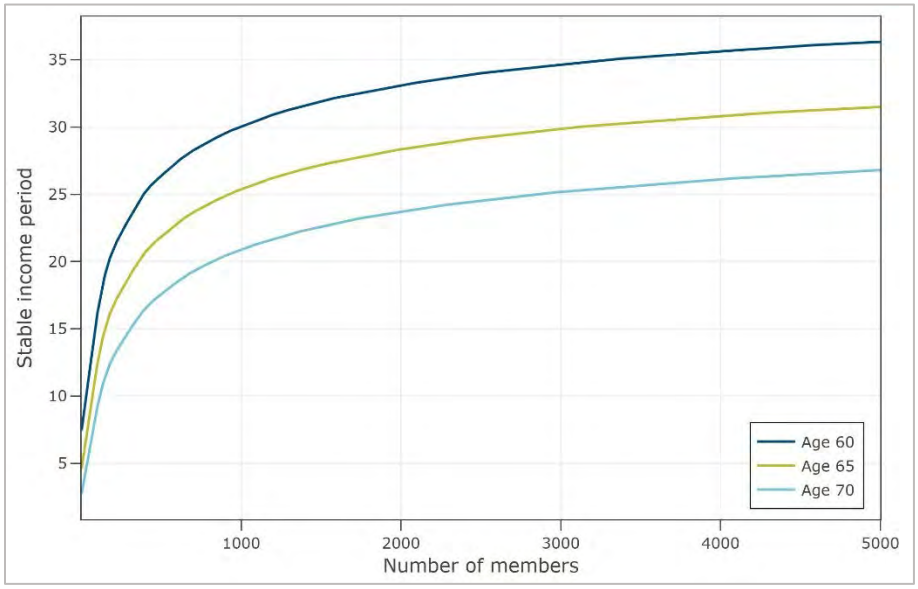

Notes: This figure illustrates the stable income period as a function of the number of members for three ages: 60, 65, and 70.

Taken together, these results reinforce the interpretation from the standard deviation measure while framing it in more operational terms; the SIP provides a direct, intuitive measure of how long benefits are expected to remain unchanged. The next subsection will extend this analysis to heterogeneous pools,

Copyright © 2026 Society of Actuaries Research Institute

where the trade-offs between diversification, age structure, and initial benefit distribution must be considered jointly.

4.4 HETEROGENEOUS POOL

Again, this subsection focuses on cases where additional members were added to the pool, rather than fixing the total pool size and examining its composition. For comparison, Appendix B.2 analyzes the alternative setting of a fixed-size pool in the context of the SIP.

4.4.1 IMPACT OF AGE

Figure 15 shows the SIP as a function of the number of individuals in the second group $N_0^{(2)}$ and their age $x^{(2)}$, assuming $N_0^{(1)} = 100$, $x^{(1)} = 65$, and $y = 1$. In this setup, the size of the pool increases by adding members from a second group with potentially different mortality characteristics. The light blue line represents the SIP for the homogeneous benchmark case in which additional members are also 65 years old.

The figure highlights that the effect of enlarging the pool depends critically on the mortality profile of the group being added. When the additional members are younger than 65, the SIP rises more rapidly, reflecting longer periods of stable benefits. By contrast, when the additional members are older than 65, the SIP increases only modestly, indicating that stability gains are much more limited. In fact, adding members aged 75 produces curves that lie consistently below the homogeneous benchmark, reflecting the destabilizing influence of higher mortality.

Figure 16
SIP AS A FUNCTION OF THE NUMBER OF MEMBERS IN THE SECOND GROUP AND THEIR AGES

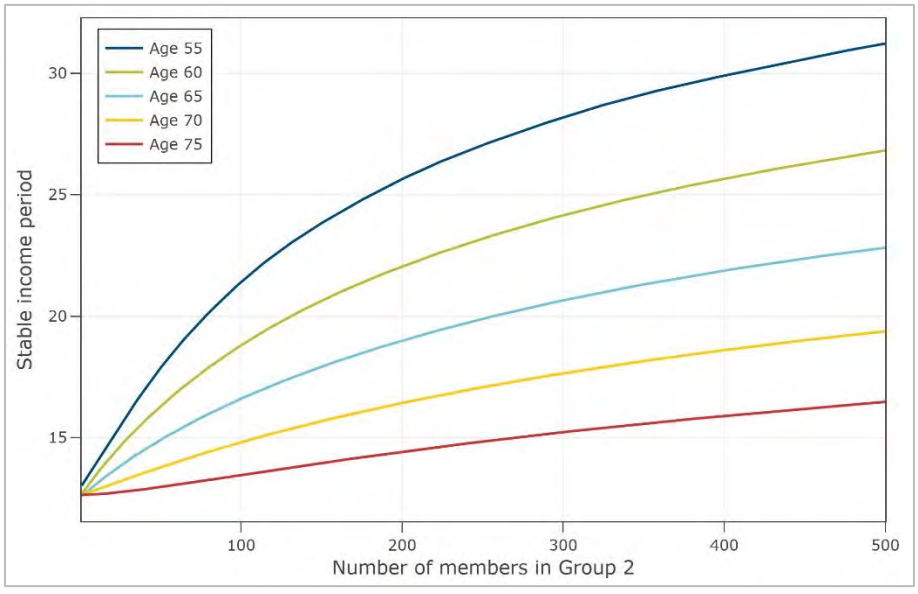

Notes: This figure shows the stable income period as a function of the number of members in the second group $N_0^{(2)}$ and their age $x^{(2)}$, assuming $N_0^{(1)} = 100$, $x^{(1)} = 65$, and $y = 1$. The lines are calculated for different ages for the second group, $x^{(2)}$.

As more individuals are added, the SIP increases in all cases, but the rate of improvement differs markedly by age. Pools with younger groups enjoy disproportionately larger stability gains, while those with older groups experience slower growth. The point at which the SIP begins to meaningfully diverge from the homogeneous case depends on the survival gap between the two groups.

Copyright © 2026 Society of Actuaries Research Institute

4.4.2 IMPACT OF THE INITIAL BENEFIT

Figure 16 shows the SIP as a function of the number of individuals in the second group $N_0^{(2)}$ and their initial benefit level y, assuming $N_0^{(1)} = 100$ and $x^{(1)} = x^{(2)} = 65$. The light blue line corresponds to the homogeneous benchmark where additional members have the same benefit level as the original group.

Figure 17

SIP AS A FUNCTION OF THE NUMBER OF MEMBERS IN THE SECOND GROUP AND THEIR INITIAL BENEFIT LEVEL

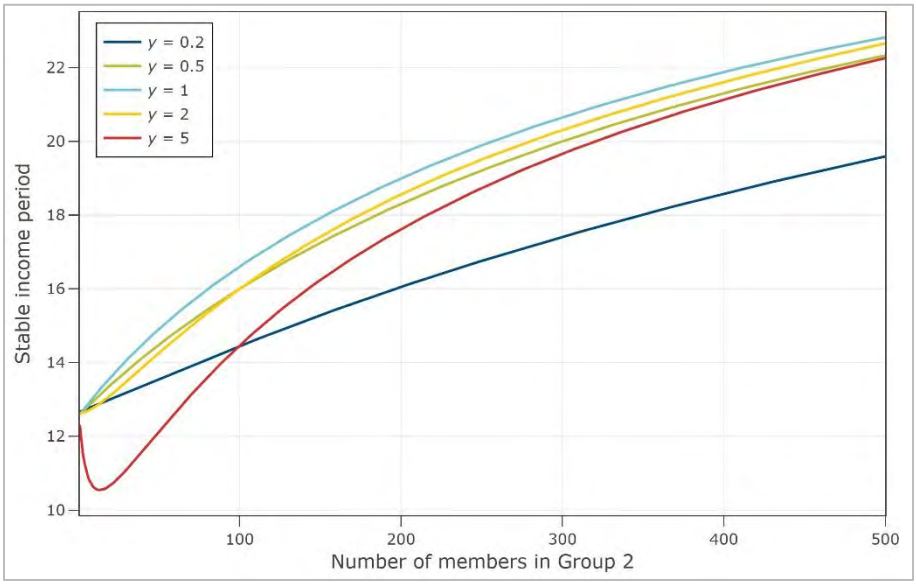

Notes: This figure shows the stable income period as a function of the number of members in the second group $N_0^{(2)}$ and their initial benefit level y, assuming $N_0^{(1)} = 100$ and $x^{(1)} = x^{(2)} = 65$. The lines are calculated for different initial benefit levels for the second group, y.

The results highlight that introducing heterogeneity in initial benefits has a markedly different effect on stability compared to heterogeneity in age. When additional members have initial benefits that differ from the benchmark, the SIP is consistently lower than in the homogeneous case (i.e., the light blue line is above all other curves in Figure 16). This reduction in SIP occurs regardless of whether the new group has smaller or larger benefits. The red curve with $y = 5$ is particularly illustrative: stability initially decreases when a small number of high-benefit members are added and, although the SIP eventually rises as more individuals are included, it never catches up to the homogeneous case.

The figure underscores that, while increasing the number of members always leads to some eventual improvement in stability through diversification, heterogeneity in benefit levels dilutes this effect. Unlike the case of heterogeneity in mortality, where mixing younger and older groups could sometimes enhance stability, benefit heterogeneity produces no such gains. Instead, it systematically shortens the SIP relative to a pool of otherwise identical members.

4.4.3 IMPACT OF BOTH DIMENSIONS

Figure 17 illustrates how age and initial benefit levels jointly shape the SIP. The downward slope of all curves highlights the strong influence of mortality: as the added group becomes older, the SIP shortens, regardless of the initial benefit level.

Copyright © 2026 Society of Actuaries Research Institute

Figure 18

SIP AS A FUNCTION OF AGE OF MEMBERS IN THE SECOND GROUP FOR FIVE INITIAL BENEFIT LEVELS

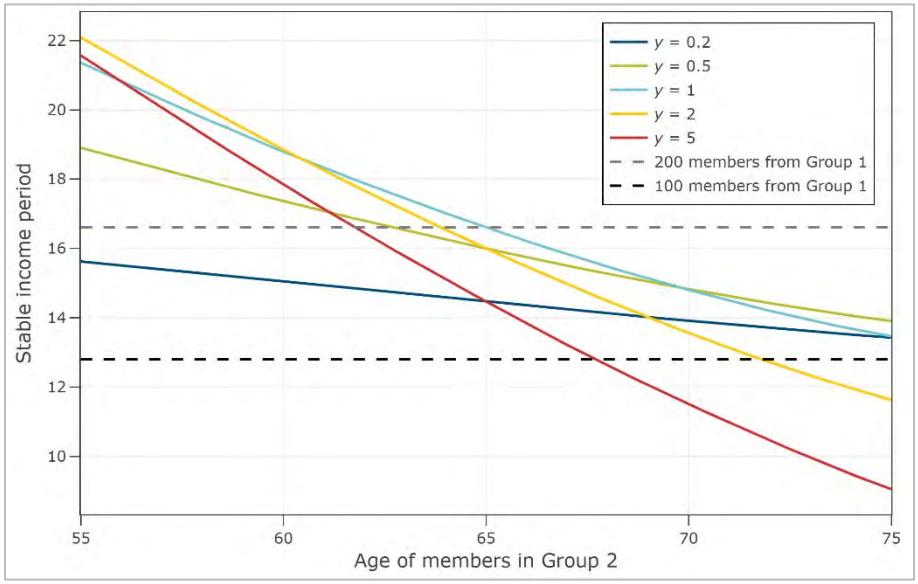

Notes: This figure shows the stable income period as a function of the age of members in the second group $x^{(2)}$, assuming $N_0^{(1)} = N_0^{(2)} = 100$ and $x^{(1)} = 65$. The lines are calculated for different levels of initial benefit for the second group, y. Horizontal dashed lines represent the stable income period for a homogeneous group of 100 and 200 members in black and gray, respectively; these lines are used for comparison purposes.

Initial benefits interact with age in important ways. For younger groups, higher initial benefits (e.g., $y = 2$ or $y = 5$) help extend the SIP when compared to lower initial benefits. For older groups, however, high benefits accelerate instability, while lower benefits soften the effect of higher mortality. In other words, the *adequate* benefit level depends on the age composition of the pool.

The dashed lines, representing homogeneous benchmarks of 100 and 200 members, provide useful reference points. Heterogeneous pools often improve upon the smaller homogeneous case, showing that diversity can add stability when the balance of age and benefits is favorable. Yet compared with the larger homogeneous benchmark, outcomes are mixed; some combinations (such as age 60 with $y = 2$) yield greater stability, while others (such as age 70 with $y = 5$) perform worse.

Figure 18 shows the SIP as a function of the initial benefit level of Group 2, y, for three different ages (60, 65, and 70). The figure reveals that the point of maximum stability varies with age. For age 60, the SIP peaks when the initial benefit level is above that of Group 1 while, for age 70, the longest SIP occurs when y is below one. This indicates that younger members can contribute more capital—and thereby receive higher benefits—without undermining stability, whereas older members improve stability only when their initial benefits are smaller.

For reasonable values of y, expanding the pool generally increases stability relative to the smaller 100-member homogeneous benchmark. However, relative to the 200-member homogeneous benchmark, outcomes are more nuanced. Certain age–benefit combinations, such as age 60 with $y = 2$, deliver a SIP longer than that of the homogeneous case, while others, such as age 70 with $y = 0.7$, fall short.

Taken together, the figure underscores that stability cannot be judged by age or benefits in isolation. Instead, it is their interaction that matters: younger, higher-benefit members strengthen stability, while older, higher-benefit members erode it. The SIP makes these trade-offs particularly transparent, as it

Copyright © 2026 Society of Actuaries Research Institute

translates complex heterogeneity into a direct measure of how long income streams can be expected to remain stable.

Figure 19

SIP AS A FUNCTION OF THE INITIAL BENEFIT LEVEL OF MEMBERS IN THE SECOND GROUP FOR THREE AGE GROUPS

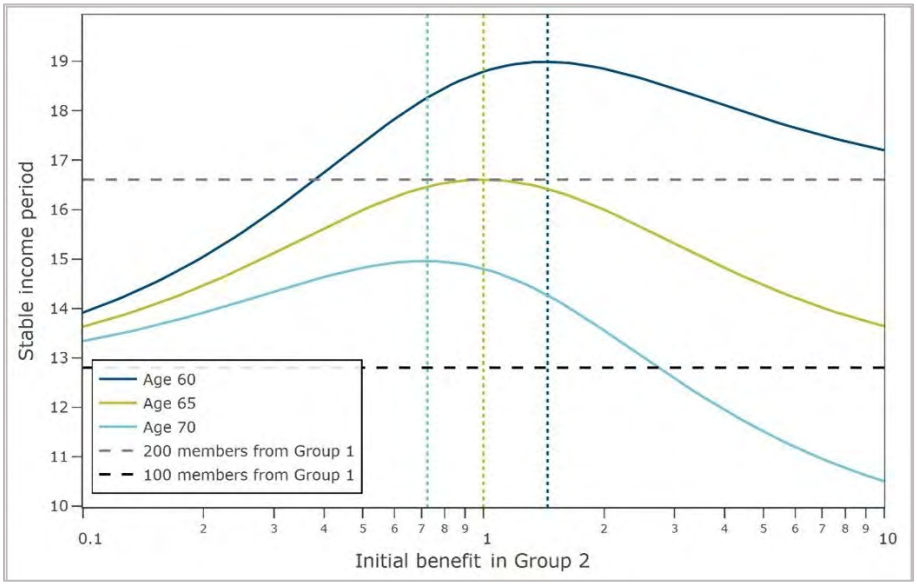

Notes: This figure shows the stable income period as a function of the initial benefit level of members in the second group y, assuming $N_0^{(1)} = N_0^{(2)} = 100$ and $x^{(1)} = 65$. The lines are calculated for different ages for the second group, $x^{(2)}$. Horizontal dashed lines represent the stable income period for a homogeneous group of 100 and 200 members in black and gray, respectively; these lines are used for comparison purposes.

Having established the main properties of the SIP in simpler settings, contour plots are used to show how stability evolves when age and benefit heterogeneity are varied jointly. These plots provide a global view of stability surfaces, highlighting how different group compositions translate into longer or shorter periods of stable payouts.

In Figure 19, contour lines trace combinations of age and initial benefit that yield the same SIP, while green and red shading emphasizes which pools are more or less stable than the benchmarks, respectively. As with the standard deviation, regions exist where heterogeneity erodes stability, particularly when older members are combined with disproportionately high benefits. At the same time, the SIP highlights more clearly the ranges where pooling extends stability, since it translates variability into a measure of expected duration. In this sense, the SIP gives a more intuitive picture of how long stability lasts under different group arrangements.

Figure 20 expands the perspective by asking whether both groups simultaneously benefit from pooling. The red areas mark cases where at least one group would be worse off than if it remained independent. Consistent with earlier results, strongly asymmetric arrangements, where one group was both younger and had lower initial benefits, tend to generate dissatisfaction, since the other group captured most of the stability gains. Importantly, no region emerged where both groups strictly outperformed their homogeneous benchmarks. The SIP, therefore, reinforced the point where some degree of cross-subsidization is unavoidable; while pooling can improve stability for both groups, it cannot do so in a way that makes both simultaneously better than in the homogeneous case.

Copyright © 2026 Society of Actuaries Research Institute

Figure 20

CONTOUR PLOT OF SIP SURFACE AND PREFERENCE REGIONS FOR THE FIRST GROUP

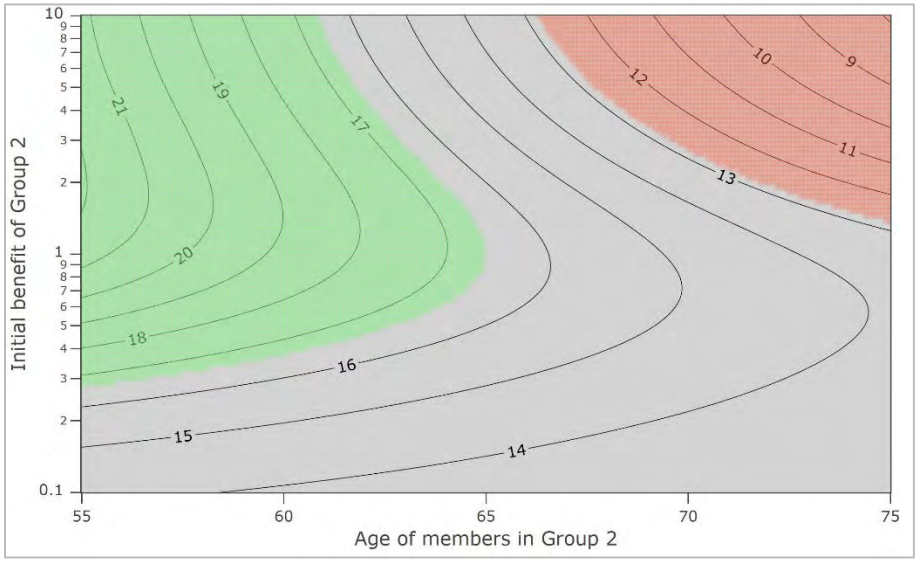

Notes: This figure shows the contour plot of the surface of stable income period as a function of the age of members in the second group $x^{(2)}$ and their initial benefit y, assuming $N_0^{(1)} = N_0^{(2)} = 100$ and $x^{(1)} = 65$. The contour lines represent combinations of group age and initial benefit that yield the same level of stable income period, thereby mapping the stability surface of the pool. The red region corresponds to outcomes with a lower stable income period than Group 1 on its own, indicating reduced stability. The green region highlights outcomes with a higher stable income period than in a fully homogeneous pool, showing where heterogeneity can be beneficial.

Figure 21

CONTOUR PLOT OF SIP SURFACE AND PREFERENCE REGIONS FOR BOTH GROUPS ASSUMING 100 MEMBERS IN THE SECOND GROUP

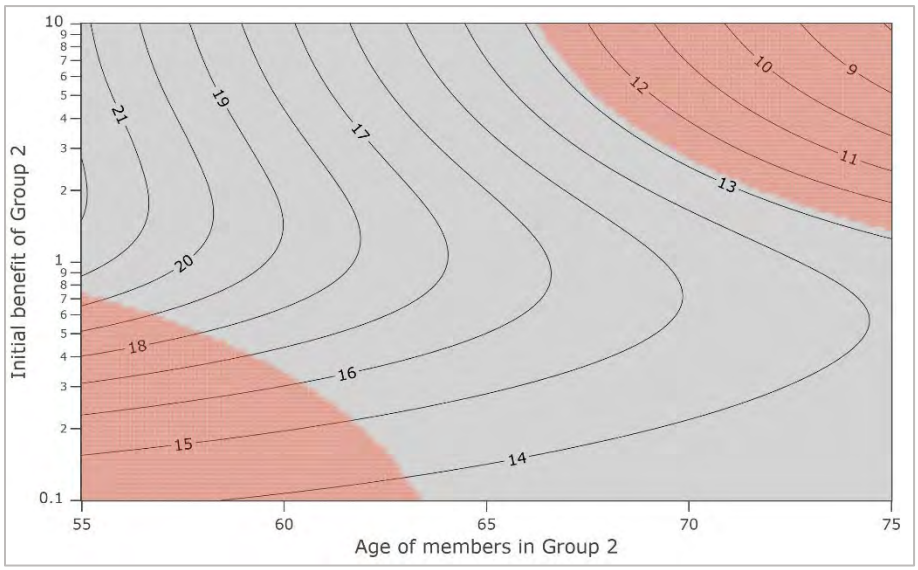

Notes: This figure shows the contour plot of the stable income period as a function of the age of members in the second group $x^{(2)}$ and their initial benefit y, assuming $N_0^{(1)} = N_0^{(2)} = 100$ and $x^{(1)} = 65$. The contour lines represent combinations of group age and initial benefit that yield the same level of stable income period, thereby mapping the stability surface of the pool. The red regions correspond to outcomes where at least one of the two groups experiences a lower stable income period than if it were on its own, indicating reduced stability.

Copyright © 2026 Society of Actuaries Research Institute

Figure 22

CONTOUR PLOT OF SIP SURFACE AND PREFERENCE REGIONS FOR BOTH GROUPS ASSUMING 50 MEMBERS IN THE SECOND GROUP

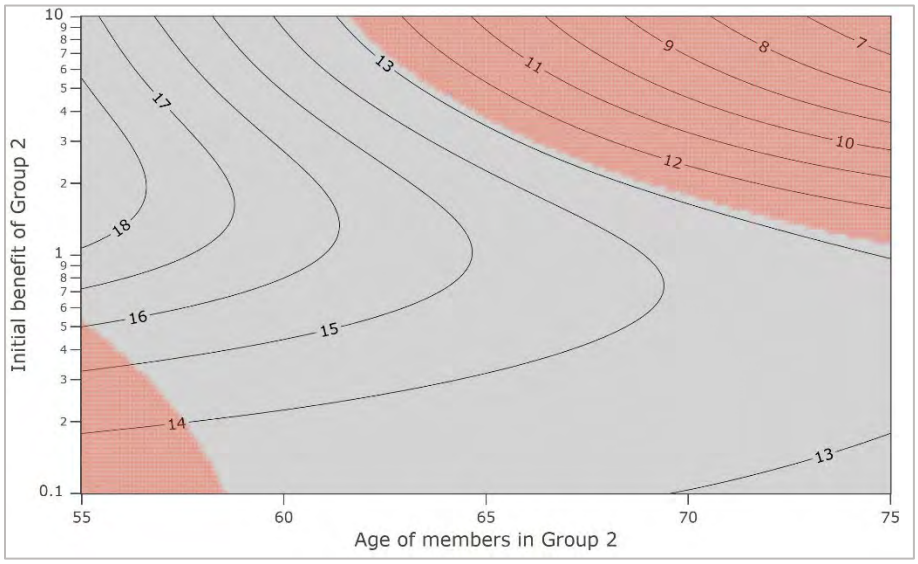

Notes: This figure shows the contour plot of the stable income period as a function of the age of members in the second group $\boldsymbol{x}^{(2)}$ and their initial benefit $\boldsymbol{y}$, assuming $\boldsymbol{N}_{\boldsymbol{0}}^{(1)} = 100$, $\boldsymbol{N}_{\boldsymbol{0}}^{(2)} = 50$, and $\boldsymbol{x}^{(1)} = 65$. See Figure 20's caption for more details.

Figure 23

CONTOUR PLOT OF SIP SURFACE AND PREFERENCE REGIONS FOR BOTH GROUPS ASSUMING 200 MEMBERS IN THE SECOND GROUP

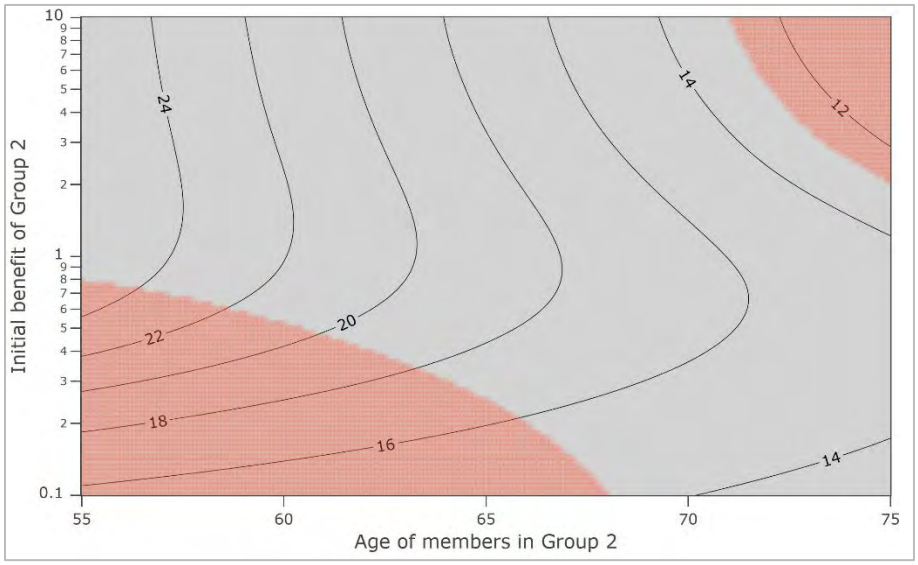

Notes: This figure shows the contour plot of the stable income period as a function of the age of members in the second group $\boldsymbol{x}^{(2)}$ and their initial benefit $\boldsymbol{y}$, assuming $\boldsymbol{N}_{\boldsymbol{0}}^{(1)} = 100$, $\boldsymbol{N}_{\boldsymbol{0}}^{(2)} = 200$, and $\boldsymbol{x}^{(1)} = 65$. See Figure 20's caption for more details.

Figure 19 and Figure 20 report contour plots for a second group of 100 members. Different sizes for the second group are considered: 50 members in Figure 21 and 200 members in Figure 22.

When Group 2 is smaller than Group 1 (Figure 21), the SIP surface changes; from Group 2's perspective, joining the larger pool produces strong diversification gains, so heterogeneity must be quite unfavorable before stability falls below what they could achieve on their own. For Group 1, however, the pooling effect

Copyright © 2026 Society of Actuaries Research Institute

is weaker, and even moderate levels of heterogeneity are enough to reduce the SIP below their standalone benchmark. This explains why the red region in the upper-right corner is larger, while the red region in the lower-left corner is smaller.

When Group 2 is larger than Group 1 (Figure 22), the pattern reverses. The larger group now perceives only modest benefits from adding a smaller, heterogeneous partner, and the SIP more quickly drops below the standalone alternative. In contrast, the smaller group enjoys stronger diversification gains, and its outcomes remain favorable across a wider range of age–benefit combinations.

These results emphasize that the relative weight of each group matters as much as their characteristics; stability outcomes depend not only on age and initial benefits, but also on which group dominates the pool's composition.

4.5 IMPACT OF INVESTMENT EXPERIENCE ON STABLE INCOME PERIOD AND HETEROGENEITY

So far, this section has focused on scenarios where asset returns were fixed at the risk-free rate, thereby isolating the impact of mortality experience on stability. This subsection extends the analysis to a mixed portfolio that allocates 50% to a risk-free asset, yielding 2% and 50% to a risky asset with normally distributed log-returns, with a mean of 7% and a volatility of 15%, such that

$$r_t^{\mathrm{PF}} = \log\left(\frac{1}{2}\, e^{0.02} + \frac{1}{2}\, e^{y_t}\right),$$

where $y_t \sim \mathcal{N}(0.07, 0.15^2)$. These parameters were chosen to be broadly consistent with the historical performance of major stock indices. This setting introduces investment risk on top of mortality heterogeneity, thereby providing a richer picture of stability when asset returns are uncertain.

Figure 23 reports the SIP surface when the pool invests in the mixed portfolio. The SIP is computed with a lower bound ε_1 of 0.3 and a confidence level β of 0.95; it indicates the time horizon over which the benefit remains above 70% of its original value with at least 95% probability. The results show that the overall structure of the SIP surface is preserved. Green and red regions appear in the same places as in the fixed-return case, indicating that the qualitative effects of heterogeneity in age and benefit remain the same; younger entrants with lower initial benefit levels tend to extend stability, while older entrants with higher benefit levels reduce it. However, the range of variation in SIP across the surface is much smaller. In this setting with random asset returns, the difference between the favorable and unfavorable zones is less than 0.2 years of stability, compared with more than 3.5 years in the fixed-return benchmark.

The reason is straightforward: when returns are random, investment risk becomes the dominant source of variation in income adjustments. This masks much of the effect of group composition, compressing the SIP surface toward a narrower band. Here, contour levels span only from about 6.0 to 6.5 years of stability whereas, in the fixed-return case, the range extended from 6 to 16 years. In other words, mortality and benefit heterogeneity play a secondary role once investment risk is introduced because the random fluctuation in asset returns already drive much of the instability.

Figure 24 presents the SIP surface when the evaluation accounts for the perspective of both groups simultaneously. In contrast to the earlier plot that highlighted green regions where at least one group could benefit, here only red areas appear, marking pool compositions that are worse than remaining independent for at least one group. This reinforces the point that, while pooling can generate gains relative to standalone participation, such gains cannot be distributed symmetrically. Even under investment risk, stability improvements cannot make both groups simultaneously better off compared to homogeneous benchmarks.

Copyright © 2026 Society of Actuaries Research Institute

Figure 24

CONTOUR PLOT OF SIP SURFACE AND PREFERENCE REGIONS FOR THE FIRST GROUP WHEN ASSET RETURNS ARE RANDOM

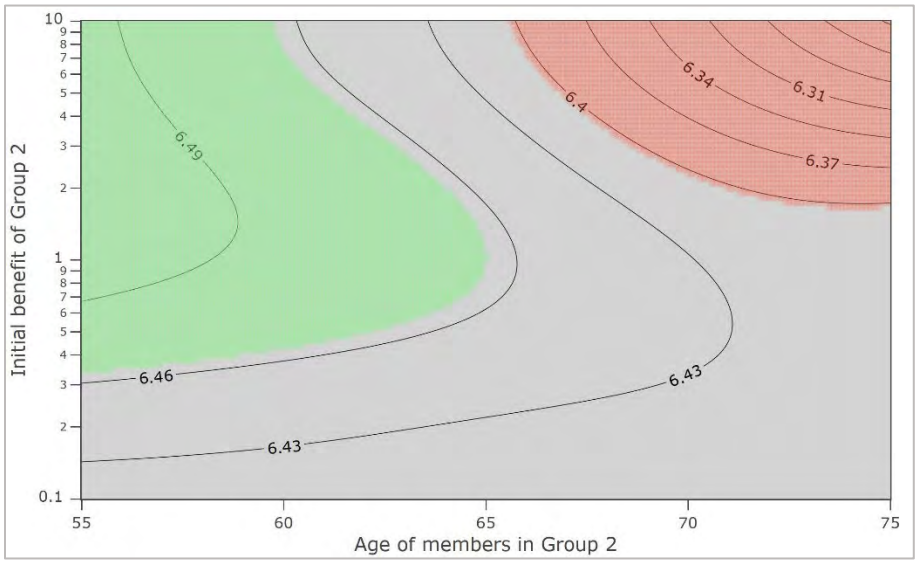

Notes: This figure shows the contour plot of the surface of stable income period as a function of the age of members in the second group $x^{(2)}$ and their initial benefit y, assuming $N_0^{(1)} = N_0^{(2)} = 100$ and $x^{(1)} = 65$. The asset returns are random, based on a mixed portfolio that allocates 50% to a risk-free asset, yielding 2% and 50% to a risky asset with normally distributed returns, with a mean of 7% and a volatility of 15%. See Figure 19's caption for more details.

These findings illustrate that the importance of heterogeneity depends on the investment environment. In a fixed-return setting, differences in group age and benefits have a pronounced impact on stability, creating wide gaps across the SIP surface. In a risky-return environment, by contrast, these effects remain present but are overshadowed by the noise introduced by asset returns. For practical purposes, this suggests that, while group composition still matters, its influence on stability may be much less visible in portfolios exposed to substantial investment risk.

4.6 COMPARISON WITH THE STANDARD DEVIATION OF THE ADJUSTMENT

The standard deviation of the adjustment and the SIP provide closely related information about the stability of lifetime pools. Both measures highlight the role of diversification through numbers, the impact of mortality differences across groups, and the destabilizing effect of heterogeneity in initial benefits. The patterns observed in the previous figures confirm that the two metrics tell essentially the same story.

That said, the measures differ in scope and interpretation. The standard deviation of the adjustment is straightforward to compute and can be expressed in terms of an implied number of homogeneous members, which provides a useful benchmark. However, it is inherently a one-period measure, reflecting variability from year to year rather than the persistence of stability over time. This makes it an excellent proxy, although a somewhat narrow one.

By contrast, the SIP is computationally more demanding, yet it captures the entire trajectory of income adjustments across multiple periods. Its interpretation is also more direct; it measures the expected length of time over which benefits remain stable before significant adjustments occur. In doing so, it accounts not only for normal year-to-year variation, but also for extreme cases where stability breaks down more abruptly. This feature makes the SIP a more comprehensive and practically meaningful metric, especially when assessing the long-term robustness of collective arrangements.

Copyright © 2026 Society of Actuaries Research Institute

Figure 25

CONTOUR PLOT OF SIP SURFACE AND PREFERENCE REGIONS FOR BOTH GROUPS ASSUMING 100 MEMBERS IN THE SECOND GROUP WHEN ASSET RETURNS ARE RANDOM

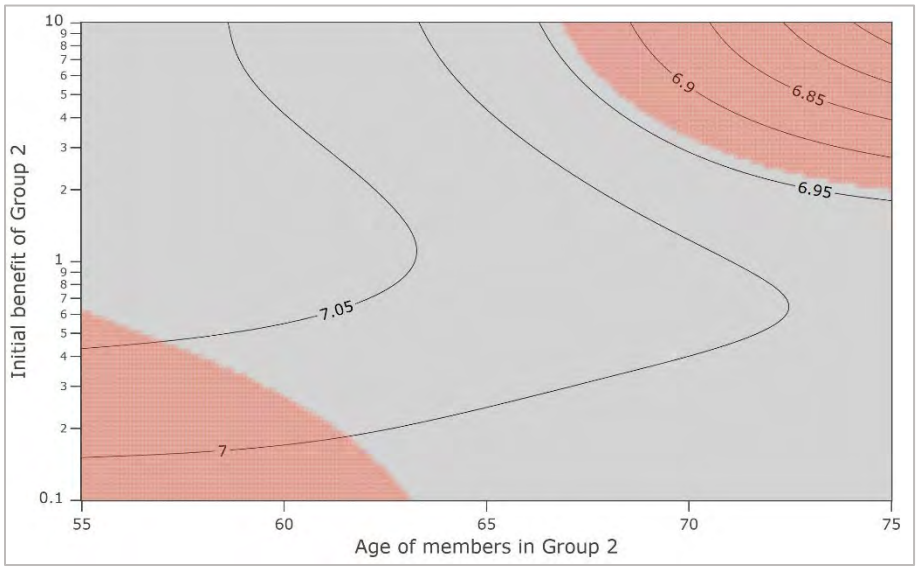

Notes: This figure shows the contour plot of the stable income period as a function of the age of members in the second group $\boldsymbol{x}^{(2)}$ and their initial benefit $\boldsymbol{y}$, assuming $\boldsymbol{N}_{\boldsymbol{0}}^{(\boldsymbol{1})} = \boldsymbol{N}_{\boldsymbol{0}}^{(\boldsymbol{2})} = 100$ and $\boldsymbol{x}^{(\boldsymbol{1})} = 65$. The asset returns are random, based on a mixed portfolio that allocates 50% to a risk-free asset, yielding 2% and 50% to a risky asset with normally distributed returns, with a mean of 7% and a volatility of 15%. See Figure 20's caption for more details.

Copyright © 2026 Society of Actuaries Research Institute

Section 5 Mitigation of Heterogeneity

Section 4 demonstrated how the SIP provides a practical and intuitive way to quantify heterogeneity within a lifetime pension pool. In particular, it offers a clear criterion for assessing whether a specific group would be better off operating on its own rather than joining a pool with multiple, heterogeneous groups. That said, even when such gains are identified, practical considerations, such as administrative simplicity, governance constraints, regulatory requirements, or communication objectives, may make it preferable to operate a single, unified pool instead of creating multiple smaller and more homogeneous ones.

This section examines practical approaches to reducing the loss of stability that some groups may experience when participating in a heterogeneous pool. For each group, operating as a standalone pool is used as a natural benchmark for comparison. The objective is to design arrangements that make participation in the combined pool at least as attractive as this benchmark for every group involved. The analysis, therefore, concentrated on cases where heterogeneity led to a reduction in stability for one or more groups (i.e., any pool composition in the red regions in the figures of Sections 3 and 4).

Even when all groups benefit relative to their standalone benchmark (i.e., any pool composition in the gray regions), the size of these gains could vary across groups. No attempt was made to equalize these gains. Since every group is already better off by participating in the lifetime pension pool, addressing differences in the magnitude of benefits would add unnecessary design and governance complexity without a clear practical justification.

The approach to mitigating the loss of stability experienced by some groups was inspired by Milevsky and Salisbury (2016) and Chen and Rach (2023). In these articles, the authors adjusted the *price per share* and the *payout rate* of a tontine to make it optimal for the members. Milevsky and Salisbury (2016) defined the optimal tontine as the one that maximizes the expected utility of each group individually, but they found that no single arrangement was optimal for every group simultaneously. Chen and Rach (2023) instead aimed for a socially optimal arrangement that maximized a weighted average of the individual expected utilities.

Within this report's lifetime pension pool framework, the payout rate was dictated by the sharing rule; nonetheless, the price per share could still be adjusted, which is defined as the entry price for one dollar of initial benefit. Also, instead of maximizing the expected utility of each group, which cannot be done simultaneously, the objective was to make each group better off than its benchmark under a given objective function, subject to the constraint that the pool as a whole remained actuarially fair. This *pooling constraint* means that the present value of the adjusted benefits must equal the total assets invested; that is,

$$A(0) = \sum_{G=1}^{S} N_0^{(G)} \frac{A_0^{(G)}}{\pi_G} \ddot{a}_{x^{(G)}}^{(G)}, \tag{8}$$

where π_G is the price per share for group G. Note that if $\pi_G = \ddot{a}_{x^{(G)}}^{(G)}$ is used, then $A_0^{(G)}/\pi_G = B_0^{(G)}$, which is referred to as the natural pricing; that is, the rule as introduced in Section 2.

Choosing a set of price-per-share values that differs from those implied by the natural pricing rule (i.e., $\pi_G = \ddot{a}_{x^{(G)}}^{(G)}$) amounts to a reallocation of the pool's assets across groups. To understand the effect of this reallocation, an objective function based solely on stability is first considered, assuming that each member cared only about the SIP. Then, a second, more realistic objective function is introduced that incorporated both the SIP and the impact of alternative entry pricing, reflecting a trade-off between stability and the initial benefit level.

Copyright © 2026 Society of Actuaries Research Institute

For simplicity, attention is restricted to a pool composed of two groups, although the method could readily be extended to more groups using more sophisticated optimization algorithms.

5.1 CONSIDERING STABILITY ONLY

If each group's only objective is to maximize stability, as measured by its SIP, the objective function for group G in a heterogeneous pool consisting of S groups can be written as

$$f_G(\pi_1, \ldots, \pi_S) = \text{SIP}, \tag{9}$$

where $\text{SIP} \equiv \text{SIP}(\varepsilon_1, \varepsilon_2, \beta)$ is the SIP of the heterogeneous pool, evaluated using the corresponding set of price-per-share values $\{\pi_G\}_{G=1}^{S}$. If there exists a price-per-share configuration, satisfying Equation (8), for which the value of this objective function for every group in the heterogeneous pool exceeds that of its respective standalone benchmark, then participation in the larger heterogeneous pool is preferable for all groups.

When the objective function is simply the SIP, the set of price-per-share values that maximize the objective function for each group is the same, since the SIP is the same for every member in a heterogeneous pool. Therefore, for this specific objective function, the pool operator could simply choose the set of price-per-share values that maximizes the SIP subject to Equation (8).

Consider a pool that consists of a first group of 100 members aged 65, pooled with a second group of 100 members aged 70, with $\mathbf{y} = 5$ under natural pricing. As illustrated in Figure 25, this configuration lies in the red region (see blue dot), indicating that at least one group would achieve a higher SIP by operating independently. To address this, the price-per-share values for both groups is adjusted, subject to the pooling constraint, seeking to maximize the objective function in Equation (9). The resulting configuration, shown by the green dot in Figure 25, moves the pool out of the red region, implying that both groups experience a higher SIP by participating in the heterogeneous pool.

This revised pool configuration satisfies all groups when the only objective is stability, as expressed through the SIP. However, this improvement, while respecting the pooling constraint in Equation (8), requires a meaningful trade-off in initial benefit levels. Table 2 reports this trade-off. Under natural pricing (first row of Panel A), Group 1 attains a higher SIP in its standalone benchmark than in the heterogeneous pool. When prices per share are adjusted to maximize the SIP (second row of both panels), the heterogeneous pool now delivers a higher SIP than the benchmark.

That said, Panel B of Table 2 shows that this gain in stability comes at a substantial cost for Group 2; for the same initial investment, its initial benefit level falls by more than half, driven by a significantly higher price per share, π_2. As a result, despite the improvement in stability, participation in the heterogeneous pool may not be attractive for Group 2 once initial benefit levels are considered.

5.2 RISK–REWARD CONSIDERATION

The previous subsection demonstrated that it is important to consider both the benefit level and the pool's stability when adjusting the price per share. A family of objective functions that accounts for both of these aspects is proposed, inspired by the assumed behavior of individual investors in Sharpe (1964). It is assumed that each group would decide whether to join the heterogenous lifetime pension pool by considering the distribution of future returns on investment. Groups were assumed to base their decisions on two characteristics of this distribution: the initial yield, $Y_G = 1/\pi_G$, which could be viewed as a proxy for the expected annual returns, and the SIP, which represents a dispersion measure of the distribution that can be parameterized to put greater weight on the lower tail.

Copyright © 2026 Society of Actuaries Research Institute

Figure 26

CONTOUR PLOT OF THE STABLE INCOME PERIOD SURFACE WITH AN EXAMPLE OF MITIGATION CONSIDERING STABILITY ONLY

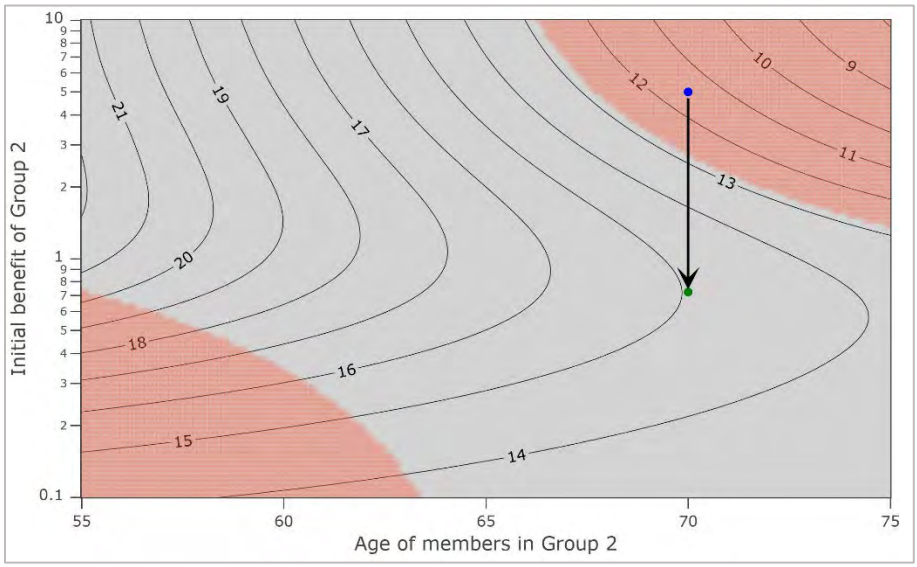

Notes: The blue and green dots represent an example pool before and after changing the natural price per share to maximize the objective function of Equation (9). See Figure 20's caption for more details.

Table 2

EXAMPLE OF MITIGATION CONSIDERING STABILITY ONLY

Panel A: Impact of Different Pricing Rules on Group 1						
Price per share	π_1	**Initial Benefit for Group 1**	**Heterogeneous SIP**	**Heterogeneous Objective for Group 1**	**Benchmark SIP for Group 1**	**Benchmark Objective for Group 1**
Natural pricing	14.95	1.00	11.5	11.5	12.8	12.8
Optimal pricing	4.65	3.22	15.0	15.0	12.8	12.8
Panel B: Impact of Different Pricing Rules on Group 2						
Price per share	π_2	**Initial Benefit for Group 2**	**Heterogeneous SIP**	**Heterogeneous Objective for Group 2**	**Benchmark SIP for Group 2**	**Benchmark Objective for Group 2**
Natural pricing	12.42	5.00	11.5	11.5	9.4	9.4
Optimal pricing	26.65	2.33	15.0	15.0	9.4	9.4

Notes: This table shows the prices per share, the initial benefits, the SIP and objective function values for Groups 1 and 2 pooled together, as well as the benchmark SIP and objective function values for each group using the natural price per share and the optimal pricing as per Equation (9). It is based on a pool comprised of 100 members aged 65 (Group 1) and 100 members aged 70 (Group 2).

When deciding whether a pool is acceptable,it is assumed that all members within the same group shared the same level of risk aversion. The objective function f_G may take various forms and parameterizations depending on the members' risk aversion. In general, members would prefer a higher initial yield, such that the objective increased for higher yields; that is,

$$\frac{\partial f_G}{\partial y_G} > 0.$$

It is further assumed that the groups are stability seekers, so that the objective function increases when the SIP is larger; that is,

$$\frac{\partial f_G}{\partial \text{SIP}} > 0.$$

Copyright © 2026 Society of Actuaries Research Institute

With this type of objective function, there is generally no single set of price-per-share values that simultaneously maximizes the objective for all groups. Instead, the objective function can be used as a screening tool to identify price-per-share combinations under which participation in the heterogeneous pool is preferable to operating independently for every group. A heterogeneous pool configuration is defined as acceptable if each group's value of the objective function exceeded that of its respective benchmark pool.

The pool from the previous subsection is revisited, represented by the blue dot in Figure 26, and used it to illustrate how this framework could be applied in practice. Specifically, the following objective function is considered:

$$f_G = \mathcal{Y}_G - \frac{\alpha_G}{2} \times \mathrm{SIP}^{-2}, \tag{10}$$

with $\alpha_G = 2$ for all groups. The parameter α_G is tied to risk aversion: a positive value of α_G corresponds to a risk-averse group, $\alpha_G = 0$ represents risk neutrality, and a negative value corresponds to a risk-seeking group. Note that this objective function is inspired by the expected quadratic utility from a traditional Gaussian mean–variance framework; Williams (1977) showed the derivation of this formula. In this context, the SIP acts as a proxy for variance, and its inverse is used in order to preserve positive convexity in the indifference curves.

Using the objective function defined in Equation (10), the set of price-per-share combinations for which the heterogeneous pool delivered a higher objective function value for both groups than those obtained under separate, homogeneous pools is identified. These combinations are shown via the green line in Figure 26. At first glance, this line lies within the red region; however, it is important to note that Figure 26 does not display the price-per-share dimension for each group and, therefore, does not fully capture the underlying trade-offs implied by these combinations. Table 3 reports the pool characteristics at the two extremities of the range of acceptable price-per-share values, evaluated from the perspective of each group. Under natural pricing, while Group 2 achieved a higher objective function value in the heterogeneous pool than in its standalone benchmark, the opposite held for Group 1. In that case, Group 1 would have preferred to operate independently rather than participate in the pooled arrangement.

To address this issue for Group 1, its price per share can be reduced, thereby increasing its initial benefit. To maintain the constraint in Equation (8), this adjustment must be offset by an increase in the price per share for Group 2, which lowers its initial benefit. This rebalancing raises the SIP of the heterogeneous pool by narrowing the gap in initial benefits across groups. The second rows of Table 3 illustrate the highest price-per-share level (i.e., corresponding to the lowest benefit level) that Group 1 is willing to accept. At this point, Group 1 is indifferent between participating in the heterogeneous pool and operating independently, as the value of its objective function matches that of its benchmark. Group 2, meanwhile, continues to prefer the heterogeneous pool, indicating that this pricing configuration remains acceptable to both groups.

If the price per share of Group 1 were decreased further while that of Group 2 were increased, additional acceptable pricing combinations would be observed until an upper pricing limit for Group 2 was reached.. As shown in the third rows of each panel of Table 3, when Group 2 reaches this upper limit, the value of its objective function in the heterogeneous pool has decreased to the point where it equals the value of its benchmark pool. This limit indicates the lowest initial benefit that Group 2 is willing to accept in exchange for the extra stability offered by the heterogeneous pool relative to its benchmark. These upper limits for both groups dictate the range of the green line in Figure 26. Any pricing between these two limits that satisfies Equation (8) would create an acceptable heterogeneous pool arrangement for both groups.

Copyright © 2026 Society of Actuaries Research Institute

Figure 27

CONTOUR PLOT OF THE STABLE INCOME PERIOD SURFACE WITH AN EXAMPLE OF MITIGATION CONSIDERING BOTH INITIAL YIELD AND THE STABLE INCOME PERIOD

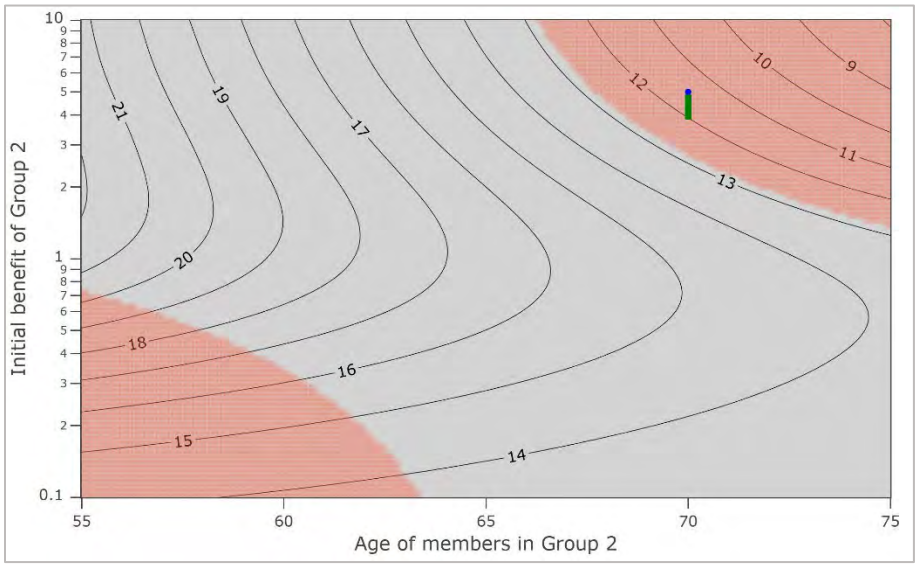

Notes: The blue point represents the example pool using the natural price per share, and the green line represents the range of acceptable price-per-share combinations based on the objective function of Equation (10). See Figure 20's caption for more details.

Table 3

EXAMPLE OF MITIGATION CONSIDERING BOTH INITIAL YIELD AND THE STABLE INCOME PERIOD

Panel A: Impact of Different Pricing Rules on Group 1						
Price-per-share	π_1	**Initial Benefit for Group 1**	**Heterogeneous SIP**	**Heterogeneous Objective for Group 1**	**Benchmark SIP for Group 1**	**Benchmark Objective for Group 1**
Natural pricing	14.95	1.00	11.5	0.059	12.8	0.061
Maximum pricing for Group 1	14.66	1.02	11.6	0.061	12.8	0.061
Maximum pricing for Group 2	12.14	1.23	12.1	0.075	12.8	0.061
Panel B: Impact of Different Pricing Rules on Group 2						
Price-per-share	π_2	**Initial Benefit for Group 2**	**Heterogeneous SIP**	**Heterogeneous Objective for Group 2**	**Benchmark SIP for Group 2**	**Benchmark Objective for Group 2**
Natural pricing	12.42	5.00	11.5	0.073	9.4	0.069
Maximum pricing for Group 1	12.48	4.97	11.6	0.072	9.4	0.069
Maximum pricing for Group 2	13.15	4.72	12.1	0.069	9.4	0.069

Notes: See Table 2's caption for more details.

It is important to emphasize that the analysis presented in this section is illustrative rather than prescriptive. The specific objective function and parameter values used here are chosen to demonstrate how the framework can be applied in practice and how trade-offs between stability and initial benefit levels can be made explicit. Different pools, member profiles, and institutional contexts will naturally call for different objective functions and parameterizations. As a result, the numerical outcomes and acceptable ranges of price-per-share values should not be interpreted as universal recommendations, but rather as examples of how the approach can inform design and pricing decisions.

From a practical perspective, the key consideration is that it is helpful for pool operators and plan sponsors to develop a good understanding of their members' preferences, both across and within groups, before applying this framework. By calibrating an objective function that reflects how members value stability relative to initial benefit levels, pool operators can use the methodology to identify ranges of price-per-share values that support voluntary participation. In this sense, the framework provides a structured

Copyright © 2026 Society of Actuaries Research Institute

decision-support tool; it does not dictate a single *optimal* design but instead offers guidance on how to balance competing objectives when structuring heterogeneous lifetime pension pools.

Copyright © 2026 Society of Actuaries Research Institute

Section 6 Concluding Remarks and Future Research

This report examined how heterogeneity within lifetime pension pools affects benefit stability, with a particular focus on mortality and initial benefit level differences across groups. Using two complementary measures of stability—the standard deviation of benefit adjustments and the SIP—it is shown that heterogeneity matters in systematic and predictable ways. Pool size alone is not sufficient to guarantee stable outcomes: the composition of the pool and the balance between its constituent groups play a central role.

Several clear conclusions emerged from the analysis. First, heterogeneity in mortality risk (proxied by age in this study) can either enhance or undermine stability depending on how it is introduced. Adding younger members generally improves stability, often substantially so, by extending the period over which income remains predictable. Adding older members, by contrast, yields limited diversification benefits and can initially reduce stability unless their number is sufficiently large. These effects persist across both short- and long-term measures of stability and remain qualitatively robust when investment risk is introduced.

Second, heterogeneity in the amount of funds brought to the pool at inception (proxied by initial benefits) is destabilizing. Pools that include small groups with disproportionately high or low initial benefits experience greater variability and shorter stable income periods than otherwise comparable homogeneous pools. From a design perspective, this highlights benefit dispersion as a key source of fragility in lifetime pension pools.

Third, when multiple dimensions of heterogeneity interact, stability outcomes depend on their joint configuration. Younger members can sustain higher benefits without undermining stability, whereas older members improve stability only when their benefits are lower.

The analysis showed that the presence of investment risk affects the relative importance of the results. When asset returns are uncertain, investment volatility becomes a dominant driver of benefit adjustments, compressing differences in stability across pool compositions. As a result, the impact of heterogeneity is virtually absent in these settings. In such environments, stability outcomes are driven primarily by the investment policy, with heterogeneity acting as a secondary factor. This finding suggests that the role of pool composition must be interpreted in light of the underlying investment strategy; heterogeneity matters most when investment risk is limited whereas, in riskier portfolios, its effects are partially masked by return volatility.

Building on these insights, the report explored practical approaches to mitigating the adverse effects of heterogeneity through adjustments to entry pricing. By reallocating assets across groups (subject to an overall actuarial fairness constraint), it is possible to design heterogeneous pools that are acceptable to all groups when evaluated against their standalone benchmarks. However, these adjustments necessarily involve trade-offs between stability and initial benefit levels. Designs that maximize stability alone may require reductions in initial benefits that some groups are unwilling to accept. More realistic objective functions that balance stability and initial yield, therefore, provide a more appropriate framework for assessing participation incentives.

From a design and operational perspective, the results underscore that *more pooling* is not always better. Larger, more heterogeneous pools can deliver important diversification benefits, but only when heterogeneity is well balanced. In some cases, maintaining separate pools, closing pools to new entrants, or introducing limited differentiation across groups may be preferable. These decisions depend not only on actuarial considerations, but also on administrative feasibility, communication objectives, and regulatory constraints. The tools developed in this report are intended to support such decisions, not to dictate a single optimal design.

Copyright © 2026 Society of Actuaries Research Institute

An important dimension that needs further testing is the role of open heterogeneous lifetime pension pools on stability. Throughout this report, much of the analysis has focused on closed pools or on the addition of a second group in a controlled setting. In practice, however, many lifetime pension pools operate on an open basis, with new entrants joining over time. An open pool can materially alter the impact of heterogeneity. On the one hand, admitting new members—particularly younger cohorts—can mitigate instability by continuously refreshing the mortality profile of the pool and extending the effective diversification horizon. In this sense, openness may function as a stabilizer, especially when entry conditions are designed to align contributions with prevailing pricing assumptions. On the other hand, open pools may also amplify heterogeneity if successive entry cohorts differ substantially in age, socioeconomic profile, or initial benefit levels. Understanding how openness interacts with heterogeneity over time is essential for evaluating lifetime pension pool designs.

Several additional important avenues for future research remain. First, the analysis assumes that relevant sources of heterogeneity are observable and correctly priced at entry. Extending the framework to account for selection effects, mispricing, or unobservable heterogeneity would be a valuable next step. Second, while this report focused on immediate, whole-pool adjustments, alternative adjustment mechanisms deserve further study, particularly considering their implications for fairness. Third, incorporating richer models of investment risk and stochastic mortality would help assess the robustness of the conclusions under more realistic conditions.

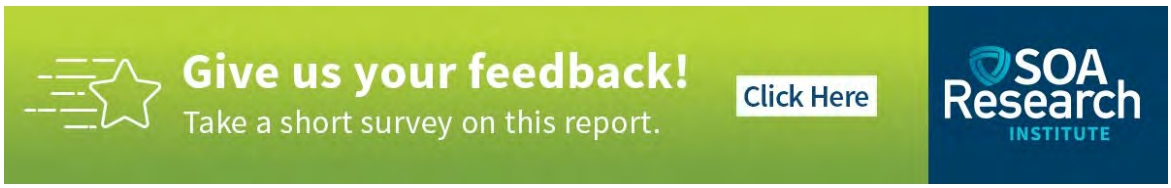

Copyright © 2026 Society of Actuaries Research Institute

Section 7 Acknowledgments

The authors' deepest gratitude goes to those without whose efforts this project could not have come to fruition: the volunteers who generously shared their wisdom, insights, advice, guidance, and arm's-length review of this study prior to publication. Any opinions expressed may not reflect their opinions nor those of their employers. Any errors belong to the authors alone.

Project Oversight Group members:

Gavin S. Benjamin, FSA, FCIA
David R. Cantor, ASA, EA, CFA, FRM
Douglas P. Chandler, FSA, FCIA
Seungjae Lee, ASA, MAAA
Grant T. Martin, FSA, EA, CERA, FCA
Lawrence I. Pollack, FSA, MAAA, FCA
Mark Shemtob, FSA, MAAA, EA, MSPA, FCA

At the Society of Actuaries Research Institute:

Barbara Scott, Senior Research Administrator
Steven C. Siegel, ASA, MAAA, SOA Senior Practice Research Actuary

Copyright © 2026 Society of Actuaries Research Institute

Appendix A Proofs

This appendix contains the proofs of Propositions 1 and 2 from Section 3.

A.1 PROOF OF PROPOSITION 1

Proof. The variance of δ_1 is decomposed by using

$$\mathrm{Var}[\delta_1] = \mathbb{E}[\delta_1^2] - \mathbb{E}[\delta_1]^2.$$

The first moment of δ_1 was given by

$$\begin{aligned}\mathbb{E}[\delta_1] &= \mathbb{E}\left[\frac{N_0^{(1)}\, p_{x^{(1)}}^{(1)}}{N_1^{(1)}}\middle| N_1^{(1)} > 0\right]\mathbb{P}\left[N_1^{(1)} > 0\right] + 0\ \mathbb{P}\left[N_1^{(1)} = 0\right] \\ &= N_0^{(1)}\, p_{x^{(1)}}^{(1)}\,\mathbb{E}\left[\frac{1}{N_1^{(1)}}\middle| N_1^{(1)} > 0\right]\ \mathbb{P}\left[N_1^{(1)} = 0\right],\end{aligned}$$

whereas the second moment of δ_1 was

$$\begin{aligned}\mathbb{E}[\delta_1^2] &= \mathbb{E}\left[\left(\frac{N_0^{(1)}\, p_{x^{(1)}}^{(1)}}{N_1^{(1)}}\right)^2\middle| N_1^{(1)} > 0\right]\mathbb{P}\left[N_1^{(1)} > 0\right] + 0^2\ \mathbb{P}\left[N_1^{(1)} = 0\right] \\ &= \left(N_0^{(1)}\, p_{x^{(1)}}^{(1)}\right)^2\,\mathbb{E}\left[\left(\frac{1}{N_1^{(1)}}\right)^2\middle| N_1^{(1)} > 0\right]\ \mathbb{P}\left[N_1^{(1)} = 0\right],\end{aligned}$$

which led to

$$\begin{aligned}\mathrm{Var}[\delta_1] = \left(N_0^{(1)}\, p_{x^{(1)}}^{(1)}\right)^2\,\mathbb{E}\left[\left(\frac{1}{N_1^{(1)}}\right)^2\middle| N_1^{(1)} > 0\right]\ \mathbb{P}[N_1 > 0] \\ - \left(N_0^{(1)}\, p_{x^{(1)}}^{(1)}\right)^2\,\mathbb{E}\left[\frac{1}{N_1^{(1)}}\middle| N_1^{(1)} > 0\right]^2\ \mathbb{P}\left[N_1^{(1)} = 0\right]^2.\end{aligned}$$

Using a second-order Taylor series expansion of $f\left(N_1^{(1)}\right) = \frac{1}{N_1^{(1)}}$ around

$$\mu \equiv \mathbb{E}\left[N_1^{(1)}\,\middle|\, N_1^{(1)} > 0\right] = \frac{N_0^{(1)}\, p_{x^{(1)}}^{(1)}}{\mathbb{P}\left[N_1^{(1)} > 0\right]} = \frac{N_0^{(1)}\, p_{x^{(1)}}^{(1)}}{\left(1 - p_{N_1^{(1)}}(0)\right)},$$

where $p_{N_1^{(1)}}$ was the pmf of $N_1^{(1)}$, then

$$\frac{1}{N_1^{(1)}} \approx f(\mu) + f'(\mu)\left(N_1^{(1)} - \mu\right) + \frac{1}{2} f''(\mu)\left(N_1^{(1)} - \mu\right)^2.$$

Since

$$f'(\mu) = -\frac{1}{\mu^2}, \qquad f''(\mu) = \frac{2}{\mu^3},$$

It is easy to show that

Copyright © 2026 Society of Actuaries Research Institute

$$\frac{1}{N_1^{(1)}} \approx \frac{1}{\mu} - \frac{1}{\mu^2}\left(N_1^{(1)} - \mu\right) + \frac{1}{\mu^3}\left(N_1^{(1)} - \mu\right)^2.$$

From this expansion, an approximation of the expected value of $\frac{1}{N_1^{(1)}}$ is obtained:

$$\mathbb{E}\left[\frac{1}{N_1^{(1)}} \middle| N_1^{(1)} > 0\right] \approx \frac{1}{\mu} + \frac{\sigma^2}{\mu^3},$$

where

$$\begin{aligned}\sigma^2 \equiv \mathrm{Var}\left[N_1^{(1)} \mid N_1^{(1)} > 0\right] &= \frac{N_0^{(1)} p_{x^{(1)}}^{(1)}\left(1 - p_{x^{(1)}}^{(1)}\right) + \left(N_0^{(1)} p_{x^{(1)}}^{(1)}\right)^2}{\mathbb{P}\left[N_1^{(1)} > 0\right]} - \frac{\left(N_0^{(1)} p_{x^{(1)}}^{(1)}\right)^2}{\mathbb{P}\left[N_1^{(1)} > 0\right]^2} \\ &= \frac{N_0^{(1)} p_{x^{(1)}}^{(1)}\left(1 - p_{x^{(1)}}^{(1)}\right) + \left(N_0^{(1)} p_{x^{(1)}}^{(1)}\right)^2}{\left(1 - p_{N_1^{(1)}}(0)\right)} - \frac{\left(N_0^{(1)} p_{x^{(1)}}^{(1)}\right)^2}{\left(1 - p_{N_1^{(1)}}(0)\right)^2}.\end{aligned}$$

The expression $g(N_1) = \left(\frac{1}{N_1^{(1)}}\right)^2$ is approximated with a similar expansion:

$$g'(\mu) = -\frac{2}{\mu^3}, \qquad g''(\mu) = \frac{6}{\mu^4},$$

which led to

$$\mathbb{E}\left[\left(\frac{1}{N_1^{(1)}}\right)^2 \middle| N_1^{(1)} > 0\right] \approx \frac{1}{\mu^2} + \frac{3\sigma^2}{\mu^4}.$$

Simplifying and putting everything back yielded

$$\mathrm{Var}[\delta_1] \approx \left(N_0^{(1)} p_{x^{(1)}}^{(1)}\right)^2 \left(1 - p_{N_1^{(1)}}(0)\right)\left(\frac{1}{\mu^2} + \frac{3\sigma^2}{\mu^4}\right) - \left(N_0^{(1)} p_{x^{(1)}}^{(1)}\right)^2 \left(1 - p_{N_1^{(1)}}(0)\right)^2 \left(\frac{1}{\mu} + \frac{\sigma^2}{\mu^3}\right)^2.$$

Finally, once the variance is obtained, the standard deviation could be easily obtained by taking the square root: $\mathrm{SD}[\delta_1] = \sqrt{\mathrm{Var}[\delta_1]}$.

A.2 PROOF OF PROPOSITION 2

Proof. The variance of δ_1 is decomposed by using

$$\mathrm{Var}[\delta_1] = \mathbb{E}[\delta_1^2] - \mathbb{E}[\delta_1]^2.$$

The first moment of δ_1 was given by

$$\begin{aligned}\mathbb{E}[\delta_1] &= \mathbb{E}\left[\frac{N_0^{(1)} p_{x^{(1)}}^{(1)} + y\, N_0^{(2)} p_{x^{(2)}}^{(2)}}{N_1^{(1)} + y\, N_1^{(2)}} \middle| N_1^{(1)} + N_1^{(2)} > 0\right] \mathbb{P}\left[N_1^{(1)} + N_1^{(2)} > 0\right] + 0\, \mathbb{P}\left[N_1^{(1)} + N_1^{(2)} = 0\right] \\ &= \left(N_0^{(1)} p_{x^{(1)}}^{(1)} + y\, N_0^{(2)} p_{x^{(2)}}^{(2)}\right) \mathbb{E}\left[\frac{1}{N_1^{(1)} + y\, N_1^{(2)}} \middle| N_1^{(1)} + N_1^{(2)} > 0\right] \mathbb{P}\left[N_1^{(1)} + N_1^{(2)} = 0\right],\end{aligned}$$

Copyright © 2026 Society of Actuaries Research Institute

whereas the second moment of δ_1 was

$$\mathbb{E}[\delta_1^2] = \mathbb{E}\left[\left(\frac{N_0^{(1)} p_{x^{(1)}}^{(1)} + y\, N_0^{(2)} p_{x^{(2)}}^{(2)}}{N_1^{(1)} + y\, N_1^{(2)}}\right)^2 \middle| N_1^{(1)} + N_1^{(2)} > 0\right] \mathbb{P}\left[N_1^{(1)} + N_1^{(2)} > 0\right] + 0^2\, \mathbb{P}\left[N_1^{(1)} + N_1^{(2)} = 0\right]$$
$$= \left(N_0^{(1)} p_{x^{(1)}}^{(1)} + y\, N_0^{(2)} p_{x^{(2)}}^{(2)}\right)^2 \mathbb{E}\left[\left(\frac{1}{N_1^{(1)} + y\, N_1^{(2)}}\right)^2 \middle| N_1^{(1)} + N_1^{(2)} > 0\right] \mathbb{P}\left[N_1^{(1)} + N_1^{(2)} > 0\right],$$

which led to

$$\mathrm{Var}[\delta_1] = \left(N_0^{(1)} p_{x^{(1)}}^{(1)} + y\, N_0^{(2)} p_{x^{(2)}}^{(2)}\right)^2 \mathbb{E}\left[\left(\frac{1}{N_1^{(1)} + y\, N_1^{(2)}}\right)^2 \middle| N_1^{(1)} + N_1^{(2)} > 0\right] \mathbb{P}\left[N_1 + N_1^{(2)} > 0\right]$$
$$- \left(N_0^{(1)} p_{x^{(1)}}^{(1)} + y\, N_0^{(2)} p_{x^{(2)}}^{(2)}\right)^2 \mathbb{E}\left[\frac{1}{N_1^{(1)} + y\, N_1^{(2)}} \middle| N_1^{(1)} + N_1^{(2)} > 0\right]^2 \mathbb{P}\left[N_1^{(1)} + N_1^{(2)} = 0\right]^2.$$

Using a second-order Taylor series expansion of $f\left(N_1^{(1)} + y\, N_1^{(2)}\right) = \frac{1}{N_1^{(1)} + y\, N_1^{(2)}}$ around

$$\mu \equiv \mathbb{E}\left[N_1^{(1)} + y\, N_1^{(2)} \,\middle|\, N_1^{(1)} + N_1^{(2)} > 0\right] = \frac{N_0^{(1)} p_{x^{(1)}}^{(1)} + y\, N_0^{(2)} p_{x^{(2)}}^{(2)}}{q},$$

where $p_{N_2^{(2)}}$ was the pmf of $N_2^{(2)}$ and

$q \equiv \mathbb{P}\left[N_1^{(1)} + N_1^{(2)} > 0\right] = \left(1 - p_{N_1^{(1)}}(0)\right) + \left(1 - p_{N_1^{(2)}}(0)\right) - \left(1 - p_{N_1^{(1)}}(0)\right)\left(1 - p_{N_1^{(2)}}(0)\right)$, it can be shown that

$$\frac{1}{N_1^{(1)} + y\, N_1^{(2)}} \approx f(\mu) + f'(\mu)\left(N_1^{(1)} + y\, N_1^{(2)} - \mu\right) + \frac{1}{2} f''(\mu)\left(N_1^{(1)} + y\, N_1^{(2)} - \mu\right)^2.$$

Since

$$f'(\mu) = -\frac{1}{\mu^2}, \qquad f''(\mu) = \frac{2}{\mu^3},$$

the following expression was obtained:

$$\frac{1}{N_1^{(1)} + y\, N_1^{(2)}} \approx \frac{1}{\mu} - \frac{1}{\mu^2}\left(N_1^{(1)} + y\, N_1^{(2)} - \mu\right) + \frac{1}{\mu^3}\left(N_1^{(1)} + y\, N_1^{(2)} - \mu\right)^2.$$

From this expansion, an approximation of the expected value of $\frac{1}{N_1^{(1)} + y\, N_1^{(2)}}$ is obtained:

$$\mathbb{E}\left[\frac{1}{N_1^{(1)} + y\, N_1^{(2)}} \middle| N_1^{(1)} + N_1^{(2)} > 0\right] \approx \frac{1}{\mu} + \frac{\sigma^2}{\mu^3},$$

where

Copyright © 2026 Society of Actuaries Research Institute

$$\sigma^2 \equiv \mathrm{Var}\left[N_1^{(1)} + y\, N_1^{(2)} \,\middle|\, N_1^{(1)} + N_1^{(2)} > 0 \right]$$
$$= \frac{\begin{pmatrix} N_0^{(1)} p_{x^{(1)}}^{(1)} \left(1 - p_{x^{(1)}}^{(1)}\right) + \left(N_0^{(1)} p_{x^{(1)}}^{(1)}\right)^2 + 2y N_0^{(1)} p_{x^{(1)}}^{(1)} N_0^{(2)} p_{x^{(2)}}^{(2)} \\ + y^2 N_0^{(2)} p_{x^{(2)}}^{(2)} \left(1 - p_{x^{(2)}}^{(2)}\right) + y^2 \left(N_0^{(2)} p_{x^{(2)}}^{(2)}\right)^2 \end{pmatrix}}{q} - \mu^2.$$

The expression $g(N_1) = \left(\frac{1}{N_1^{(1)} + y\, N_1^{(2)}}\right)^2$ is obtained with a similar expansion:

$$g'(\mu) = -\frac{2}{\mu^3}, \qquad g''(\mu) = \frac{6}{\mu^4},$$

which led to

$$\mathbb{E}\left[\left(\frac{1}{N_1^{(1)} + y\, N_1^{(2)}} \right)^2 \middle| N_1^{(1)} + N_1^{(2)} > 0 \right] \approx \frac{1}{\mu^2} + \frac{3\sigma^2}{\mu^4}.$$

Simplifying and putting everything back yielded

$$\mathrm{Var}[\delta_1] \approx \left(N_0^{(1)} p_{x^{(1)}}^{(1)} + y\, N_0^{(2)} p_{x^{(2)}}^{(2)}\right)^2 q \left(\frac{1}{\mu^2} + \frac{3\sigma^2}{\mu^4}\right) - \left(N_0^{(1)} p_{x^{(1)}}^{(1)} + y\, N_0^{(2)} p_{x^{(2)}}^{(2)}\right)^2 q^2 \left(\frac{1}{\mu} + \frac{\sigma^2}{\mu^3}\right)^2.$$

Finally, once the variance was obtained, the standard deviation could be easily obtained by taking the square root: $\mathrm{SD}[\delta_1] = \sqrt{\mathrm{Var}[\delta_1]}$.

Copyright © 2026 Society of Actuaries Research Institute

Appendix B Alternative Setting of a Fixed-Size Pool

B.1 STANDARD DEVIATION OF THE ADJUSTMENT

Figure 27 reports the standard deviation of the adjustment for a 100-member pool. The first group was assumed to be 65 years old, but its size varied from 0 to 100 members. The age of the second group is also varied from 55 to 75, as well as its size from 100 to 0 members, so that the total number of members was 100 for each case.

Figure 28

STANDARD DEVIATION AS A FUNCTION OF THE NUMBER OF MEMBERS IN THE SECOND GROUP AND THEIR AGE

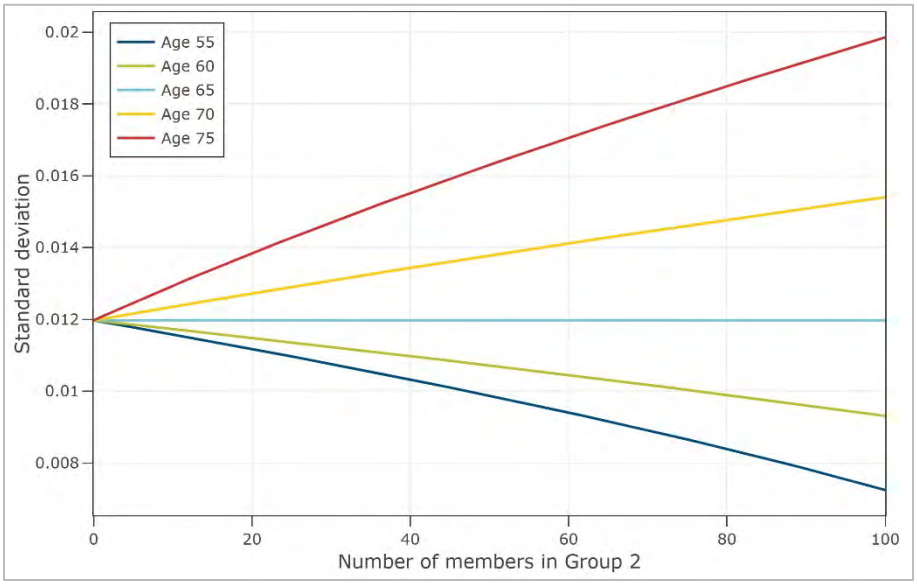

Notes: This figure shows the standard deviation of the adjustment as a function of the number of members $N_0^{(2)}$, assuming a total of 100 members, $x^{(1)} = 65$, and $y = 1$. The lines are calculated for different ages for the second group, $x^{(2)}$.

Removing younger members who have high survival probabilities and replacing them with older individuals who have low survival probabilities increases the standard deviation of the adjustment, as expected. Conversely, adding members with high survival probabilities to a pool consisting mainly of individuals with low survival probabilities reduces this standard deviation. Thus, there exists an implicit trade-off between different groups: individuals with higher survival probabilities are disadvantaged by the inclusion of individuals with significantly lower survival probabilities, and vice versa. However, if the survival probabilities across groups are similar, the difference in standard deviation remains relatively small.

Overall, pools composed of younger individuals with higher survival probabilities exhibit lower adjustment uncertainty. Therefore, for a given pool size, selecting individuals with high survival probabilities minimizes the standard deviation of adjustments.

Figure 28 illustrates the standard deviation of the adjustment for, again, a 100-member pool. Both groups are assumed to be age 65, but they differ in their initial benefit level y. As in Figure 27, the size of the first group varies from 0 to 100 members, with the second group making up the remainder. Interestingly, the highest standard deviations arise when a small group has an initial benefit level that differs sharply from the rest of the pool. This suggests a general guideline: when there is a large discrepancy between the initial

Copyright © 2026 Society of Actuaries Research Institute

benefits of two groups, stability is improved if the group with the higher benefit level is also the one with more members.

An important feature of Figure 28 is its symmetry; what ultimately drives the adjustment is not the absolute level of the benefits but their relative size across groups. Whether the first group receives higher or lower initial benefits than the second, the resulting standard deviation depends only on how imbalanced the payouts are relative to one another. In other words, interchanging the roles of the two groups yields the same outcome: for instance, having 10 members with $y = 10$ and 90 members with $y = 1$ is equivalent to having 10 members with $y = 1$ and 90 members with $y = 0.1$.

Figure 29

STANDARD DEVIATION OF THE ADJUSTMENT AS A FUNCTION OF THE NUMBER OF MEMBERS IN THE SECOND GROUP AND THEIR INITIAL BENEFIT LEVEL

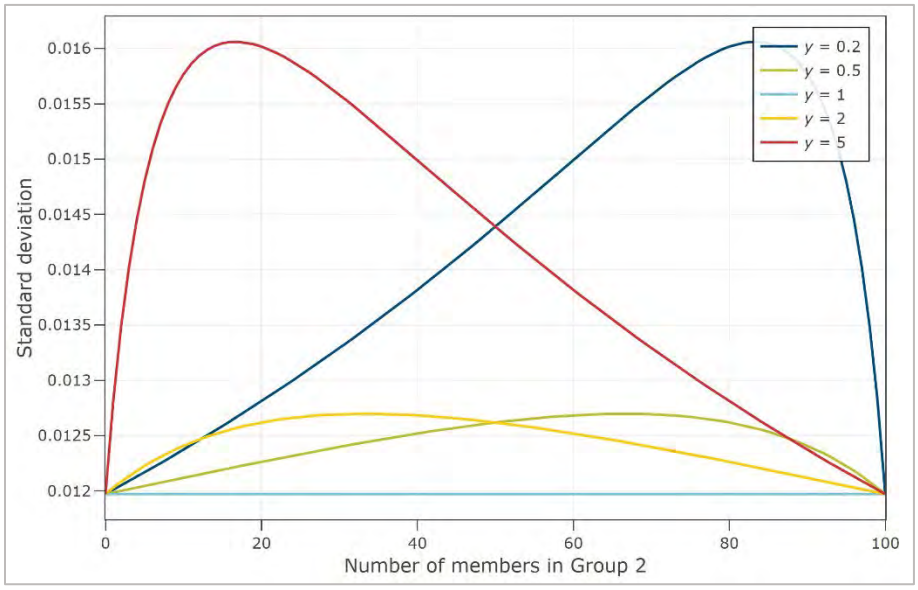

Notes: This figure shows the standard deviation of the adjustment as a function of the number of members $N_0^{(2)}$, assuming a total of 100 members and $x^{(1)} = x^{(2)} = 65$. The lines are calculated for different initial benefit levels for the second group, y.

B.2 STABLE INCOME PERIOD

Figure 29 reports the SIP for a 100-member pool composed of two groups. Similar to Appendix B.1, the first group is age 65, while the second varies in age, with the total pool size always equal to 100. The results show that adding younger participants extends stability, as reflected in the rising curves for ages 55 and 60. In contrast, introducing older participants reduces the SIP: the lines for ages 70 and 75 slope downward, indicating shorter periods of stable payouts. The homogeneous benchmark of 65-year-olds lies in between.

These patterns underline the trade-offs inherent in mixing groups. Stability improves when the pool tilts toward members with higher survival probabilities, while it erodes when older, riskier participants dominate. The overall message is clear: for a fixed pool size, age composition is decisive in determining how long benefits remain stable.

Figure 30 reports the SIP for a 100-member pool where both groups are age 65 but differ in their initial benefit level y. The size of the first group varies from 0 to 100 members, with the second group making up the balance so that the total pool size remains fixed at 100. Similar to the conclusion reached in Section 4, the results show that the worst (lowest) SIP values occur when a small group has initial benefits that differ sharply from the rest of the pool. For example, a handful of members with much higher or lower benefits

Copyright © 2026 Society of Actuaries Research Institute

can markedly reduce stability, even when the majority is homogeneous. This points to a general principle: large discrepancies in initial benefits undermine stability most when the group with the extreme payout is relatively small.

Figure 30

SIP AS A FUNCTION OF THE NUMBER OF MEMBERS IN THE SECOND GROUP AND THEIR AGE

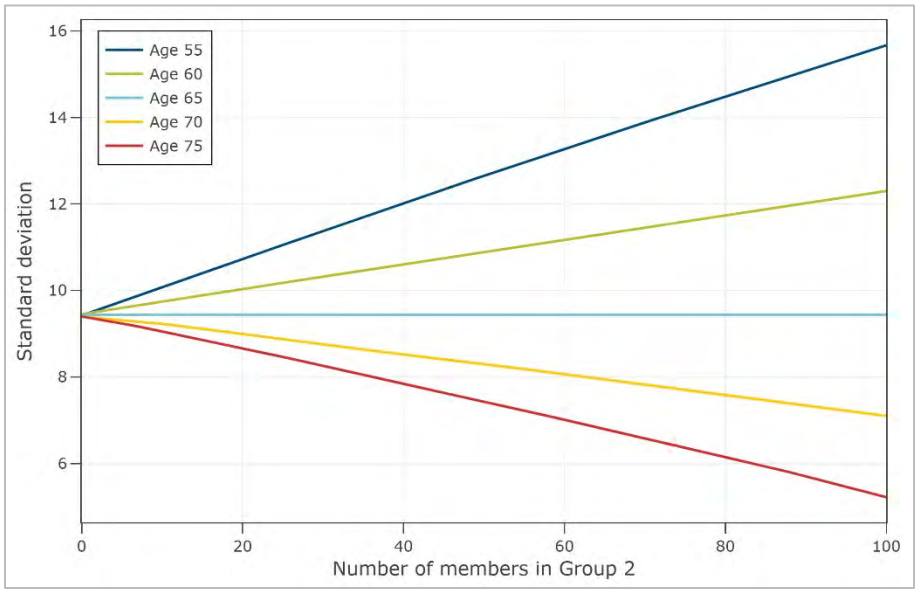

Notes: This figure shows the stable income period as a function of the number of members $N_0^{(2)}$, assuming a total of 100 members, $x^{(1)} = 65$, and $y = 1$. The lines are calculated for different ages for the second group, $x^{(2)}$.

Figure 31

SIP AS A FUNCTION OF THE NUMBER OF MEMBERS IN THE SECOND GROUP AND THEIR INITIAL BENEFIT LEVEL

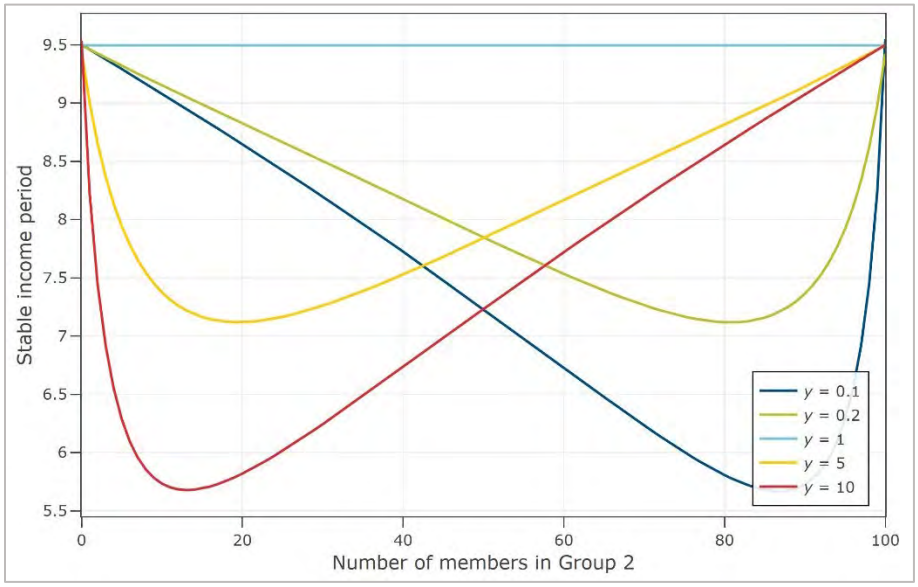

Notes: This figure shows the stable income period as a function of the number of members $N_0^{(2)}$, assuming a total of 100 members and $x^{(1)} = x^{(2)} = 65$. The lines are calculated for different initial benefit levels for the second group, y.

Copyright © 2026 Society of Actuaries Research Institute

Appendix C More Details on Interpreting Age for Mortality Differentials

In the main body of this report, mortality heterogeneity across pool members is primarily represented through differences in age at entry. This modeling choice is intentional and reflects a balance between realism and tractability. Age is a readily observable characteristic, is well understood by members and regulators, and captures a large share of systematic differences in mortality risk. Nevertheless, in practice, individuals of the same age may face different mortality risks due to differences in sex, socioeconomic status, occupation, lifestyle, or health status.

The purpose of this appendix is to explain how such sources of mortality heterogeneity can be incorporated into the analytical framework used in this report, and to clarify why the age-based analysis presented in the main sections remains appropriate for assessing benefit stability.

C.1 AGE AS A PROXY FOR MORTALITY HETEROGENEITY

The analysis relies on a reference mortality model to describe survival probabilities for pool members. Relative to this reference, differences in mortality across groups can be represented by adjusting age. Conceptually, a group with a higher mortality risk can be treated as being effectively older under the reference mortality law, while a group with a lower mortality risk can be treated as being effectively younger. For example, it is common to approximate differences in mortality between males and females of the same chronological age by assigning an age offset that reflects their relative mortality experience. Under such an approximation, pooling 65-year-old males with 65-year-old females can be analyzed by treating the males as if they were several years older than the females under a common reference mortality table.

If this proxy is reasonable, then the results presented in the sections analyzing mortality and initial benefit heterogeneity extend naturally to mortality differences arising from characteristics other than age, provided that an appropriate implied age can be identified for each group.

To understand why the age-based proxy is appropriate, it is helpful to consider the structure of the mortality model used in this report. Survival probabilities are modeled using a Gompertz law presented in Section 2.2, which is characterized by two key components:

1. A parameter governing the overall level of mortality; that is, m.
2. A parameter governing how concentrated deaths are around that level; that is, b.

Differences in the first component correspond directly to shifts in age. When two groups differ only in this dimension, their mortality distributions are identical up to a translation along the age axis. In such cases, representing mortality heterogeneity through age differences is exact and introduces no approximation error. Differences in the second component affect the spread of mortality over time. A lower dispersion implies that deaths are more concentrated around the modal age, resulting in higher survival probabilities at younger ages and lower survival probabilities at older ages. A higher dispersion has the opposite effect, spreading deaths more evenly over time.

C.2 EFFECT OF MORTALITY DISPERSION ON THE STABLE INCOME PERIOD

Allowing the dispersion component of the mortality model to differ across groups has a measurable but limited effect on the SIP. Importantly, these effects do not alter the qualitative relationship between age and income stability identified in the main analysis. In particular, when dispersion differs across groups, the SIP remains decreasing in age across all specifications considered.

Copyright © 2026 Society of Actuaries Research Institute

Figure 31 illustrates how differences in mortality dispersion affect the SIP as a function of members' ages in the second group. The figure mirrors the analysis in the main body of the report but allows the dispersion parameter of the mortality model to vary across groups, while holding other characteristics constant.

Figure 32

SIP AS A FUNCTION OF THE NUMBER OF MEMBERS IN THE SECOND SUBGROUP AND THEIR MORTALITY DISPERSION PARAMETER

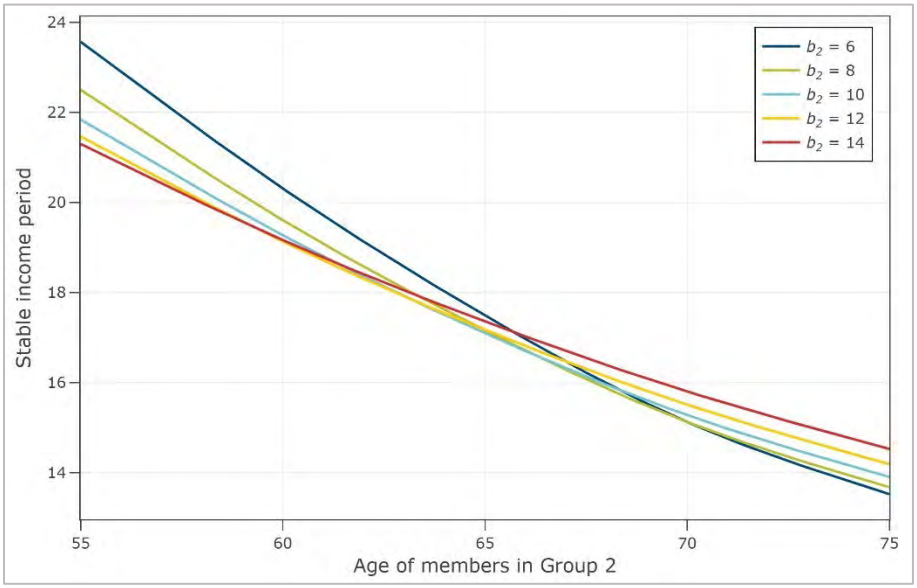

Notes: This figure shows the stable income period as a function of the age of members in the second group $x^{(2)}$, assuming $N_0^{(1)} = N_0^{(2)} = 100$, $x^{(1)} = 65$, and $b_1 = 10$. The lines are calculated for different mortality dispersion parameters b_2 for the second group, assuming that the modal parameter is the same for both groups; that is, $m_1 = m_2 = 85$.

The figure reports that, for younger members, lower dispersion is associated with slightly higher SIP values, reflecting more predictable survival outcomes during the early part of retirement. For older members, lower dispersion leads to slightly lower SIP values, as mortality risk increases more sharply at advanced ages. Across all cases, however, the SIP declines smoothly with age and preserves the same overall shape. This reinforces the conclusion that allowing for realistic differences in mortality dispersion changes the level of benefit stability only marginally and does not affect the qualitative conclusions drawn from the age-based analysis.

From a practical perspective, these results support the use of a single reference mortality model combined with age-based adjustments to account for mortality heterogeneity. Differences in mortality related to sex, socioeconomic status, or other characteristics can be incorporated by assigning an implied age that reflects relative mortality risk.

As a result, the conclusions drawn elsewhere in this report regarding the impact of mortality heterogeneity on income stability remain valid even when mortality differences extend beyond age alone. Age provides an effective and practical proxy for mortality heterogeneity when assessing income stability in lifetime pension pools.

Copyright © 2026 Society of Actuaries Research Institute

References

Balter, A. G. and Werker, B.J. (2020). The effect of the assumed interest rate and smoothing on variable annuities. *ASTIN Bulletin*, *50*(1), 131–154.

Balter, A.G., Kallestrup-Lamb, M., and Rangvid, J. (2020). Variability in pension products: A comparison study between The Netherlands and Denmark. *Annals of Actuarial Science*, *14*(2), 338–357.

Bégin, J. F. (2020). Levelling the playing field: A VIX-linked structure for funded pension schemes. *Insurance: Mathematics and Economics*, *94*, 58–78.

Bégin, J.-F. and Sanders, B. (2023a). *Benefit at Risk in Lifetime Pension Pools*. Society of Actuaries and Canadian Institute of Actuaries.

Bégin, J.-F. and Sanders, B. (2023b). *Exploration of Lifetime Pension Pool Design Elements*. Society of Actuaries and Canadian Institute of Actuaries.

Bégin, J.-F. and Sanders, B. (2024). Benefit volatility-targeting strategies in lifetime pension pools. *Insurance: Mathematics and Economics*, *118*, 72–94.

Bégin, J.-F., Sanders, B., and Sun, Y. (2026a). Optimal hurdle rate and investment policy in lifetime pension pools. *Working Paper*.

Bégin, J.-F., Desai, V., Donnelly, C., and Sanders, B. (2026b). Capital preservation or pension smoothing? A general framework for collective pension plan designs. *Working Paper*.

Bernhardt, T. and Donnelly, C. (2019). Modern tontine with bequest: Innovation in pooled annuity products. *Insurance: Mathematics and Economics*, *86*, 168–188.

Bernhardt, T. and Donnelly, C. (2021). Quantifying the trade-off between income stability and the number of members in a pooled annuity fund. *ASTIN Bulletin*, *51*(1), 101–130.

Bernhardt, T. and Qu, G. (2024). Wealth heterogeneity in a closed pooled annuity fund. *Scandinavian Actuarial Journal*, *2024*(3), 199–226.

Chen, A., Qian, L., and Yang, Z. (2021). Tontines with mixed cohorts. *Scandinavian Actuarial Journal*, *2021*(5), 437–455.

Chen, A., Kanagawa, M., and Zhang, F. (2023). Intergenerational risk sharing in a defined contribution pension system: analysis with Bayesian optimization. *ASTIN Bulletin*, *53*(3), 515–544.

Chen, A. and Rach, M. (2023). Actuarial fairness and social welfare in mixed-cohort tontines. Insurance: Mathematics and Economics, 111, 214–229.

CREF. (2025). Prospectus. College Retirement Equities Fund. Available online at https://fundcompli.rightprospectus.com/documents/TIAACREFVA/CREF_ar.pdf. Accessed on April 30, 2026.

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

Donnelly, C. (2017). A discussion of a risk-sharing pension plan. *Risks*, *5*(1), 12.

Donnelly, C. (2023). Heterogeneity among new entrants to a pooled annuity fund. *Working Paper*.

Copyright © 2026 Society of Actuaries Research Institute

Donnelly, C., Guillén, M., and Nielsen, J. P. (2013). Exchanging uncertain mortality for a cost. *Insurance: Mathematics and Economics*, *52*(1), 65–76.

Donnelly, C., Guillén, M., and Nielsen, J. P. (2014). Bringing cost transparency to the life annuity market. *Insurance: Mathematics and Economics*, *56*, 14–27.

Fullmer, R. K. (2019). *Tontines: A Practitioner's Guide to Mortality-Pooled Investments*. CFA Institute Research Foundation.

Goecke, O. (2013). Pension saving schemes with return smoothing mechanism. *Insurance: Mathematics and Economics*, *53*(3), 678–689.

Gros, B. and Sanders, B. (2019). *The Quest for Sustainability in Contingent Pension Plans*. C. D. Howe Institute Commentary 553, September 2019.

Guillén, M., Jørgensen, P. L., and Nielsen, J. P. (2006). Return smoothing mechanisms in life and pension insurance: Path-dependent contingent claims. *Insurance: Mathematics and Economics*, *38*(2), 229–252.

Hanewald, K., Piggott, J., and Sherris, M. (2013). Individual post-retirement longevity risk management under systematic mortality risk. *Insurance: Mathematics and Economics*, *52*(1), 87–97.

Horneff, W. J., Maurer, R. H., Mitchell, O. S., and Stamos, M. Z. (2010). Variable payout annuities and dynamic portfolio choice in retirement. *Journal of Pension Economics & Finance*, 9(2), 163–183.

Kahneman, D. and Tversky, A. (1979). Prospect Theory: An Analysis of Decision under Risk. *Econometrica*, 47(4), 263–291.

Li, S., Labit Hardy, H., Sherris, M., and Villegas, A. M. (2022). A managed volatility investment strategy for pooled annuity products. *Risks*, *10*(6), 121.

Marchand, I., Bégin, J.-F., and Sanders, B. (2026). When heterogeneity helps: Stability in lifetime pension pools. *Working Paper*.

Milevsky, M. A. and Salisbury, T. S. (2015). Optimal retirement income tontines. *Insurance: Mathematics and Economics*, *64*, 91–105.

Milevsky, M. A. and Salisbury, T. S. (2016). Equitable retirement income tontines: Mixing cohorts without discriminating. *ASTIN Bulletin*, *46*(3), 571–604.

Olivieri, A., Thirurajah, S., and Ziveyi, J. (2022). Target volatility strategies for group self-annuity portfolios. *ASTIN Bulletin*, *52*(2), 591–617

Piggott, J., Valdez, E. A., and Detzel, B. (2005). The simple analytics of a pooled annuity fund. *Journal of Risk and Insurance*, *72*(3), 497–520.

Qiao, C. and Sherris, M. (2013). Managing systematic mortality risk with group self-pooling and annuitization schemes. *Journal of Risk and Insurance*, *80*(4), 949–974.

Sanders, B. and Bégin, J.-F. (2025). *Comparative Analysis of Variable Payment Life Annuities*. Global Risk Institute.

Sharpe, W. F. (1964). Capital asset prices: A theory of market equilibrium under conditions of risk. *Journal of Finance*, *19*(3), 425–442.

Copyright © 2026 Society of Actuaries Research Institute

Stamos, M. Z. (2008). Optimal consumption and portfolio choice for pooled annuity funds. I*nsurance: Mathematics and Economics*, *43*(1), 56–68

Valdez, E. A., Piggott, J., and Wang, L. (2006). Demand and adverse selection in a pooled annuity fund. *Insurance: Mathematics and Economics*, *39*(2), 251–266

Williams, J. T. (1977). A note on indifference curves in the mean-variance model. *Journal of Financial and Quantitative Analysis*, *12*(1), 121–126.

Winter, P. and Planchet, F. (2022). Modern tontines as a pension solution: A practical overview. *European Actuarial Journal*, *12*(1), 3-32.

Copyright © 2026 Society of Actuaries Research Institute

About The Society of Actuaries Research Institute

Serving as the research arm of the Society of Actuaries (SOA), the SOA Research Institute provides objective, data-driven research bringing together tried and true practices and future-focused approaches to address societal challenges and your business needs. The Institute provides trusted knowledge, extensive experience and new technologies to help effectively identify, predict and manage risks.

Representing the thousands of actuaries who help conduct critical research, the SOA Research Institute provides clarity and solutions on risks and societal challenges. The Institute connects actuaries, academics, employers, the insurance industry, regulators, research partners, foundations and research institutions, sponsors and non-governmental organizations, building an effective network which provides support, knowledge and expertise regarding the management of risk to benefit the industry and the public.

Managed by experienced actuaries and research experts from a broad range of industries, the SOA Research Institute creates, funds, develops and distributes research to elevate actuaries as leaders in measuring and managing risk. These efforts include studies, essay collections, webcasts, research papers, survey reports, and original research on topics impacting society.

Harnessing its peer-reviewed research, leading-edge technologies, new data tools and innovative practices, the Institute seeks to understand the underlying causes of risk and the possible outcomes. The Institute develops objective research spanning a variety of topics with its strategic research programs: aging and retirement; actuarial innovation and technology; mortality and longevity; diversity, equity and inclusion; health care cost trends; and catastrophe and climate risk. The Institute has a large volume of topical research available, including an expanding collection of international and market-specific research, experience studies, models and timely research.

Society of Actuaries Research Institute
8770 W Bryn Mawr Ave, Suite 1000
Chicago, IL 60631
www.soa.org

Copyright © 2026 Society of Actuaries Research Institute