Cooperation in Target Benefit Plans: A Game Theoretical Perspective

by

Jing Wang

B.Sc., Simon Fraser University, 2016

Project Submitted in Partial Fulfillment of the
Requirements for the Degree of
Master of Science

in the
Department of Statistics and Actuarial Science
Faculty of Science

SIMON FRASER UNIVERSITY
Fall 2018

Approval

Name: **Jing Wang**

Degree: **Master of Science (Actuarial Science)**

Title: **Cooperation in Target Benefit Plans: A Game Theoretical Perspective**

Examining Committee: **Chair:** Jinko Graham
Professor

Barbara Sanders
Senior Supervisor
Associate Professor

Bégin, Jean-François
Supervisor
Assistant Professor

Cary Chi-Liang Tsai
External Examiner
Professor

Date Defended: **December 12, 2018**

Abstract

Many occupational pension plans rely on intergenerational cooperation to deliver stable retirement benefits; however, this cooperation has natural limits and exceeding these limits can threaten the sustainability of the plan. In this project, we cast the problem of intergenerational cooperation within funded pension plans in a game theoretic framework that incorporates overlapping generations and uncertainty in the cost of cooperation. Employing the concept of a subgame perfect equilibrium, we determine the threshold above which cooperation should not be enforced. Using two different processes for the stochastic cost of cooperation, we illustrate the combination of parameters that allow for the existence of a reasonable threshold, and study how the level of prefunding and the stochastic process parameters affect both the threshold and the probability of sanctioned non-cooperation.

Keywords: Non-cooperative game theory; Overlapping generations; Stochastic processes; Pension plan sustainability; Intergenerational cooperation

Dedication

To my parents who give me endless love and instill the passion of learning in me.

Acknowledgements

First and foremost, I would like to thank my senior supervisor, Professor Barbara Sanders, for originally getting me interested in research and motivating me to pursue masters' degree when I was an undergraduate research assistant. I deeply appreciate her enlightening guidance, constant patience, and enormous encouragement both in my coursework and in my research project. I feel that I have been blessed to have Barbara as my supervisor. I am even more blessed to learn from all the amazing professors in our department who guided, inspired and supported me so much.

Huge thanks go to my supervisor, Professor Jean-François Bégin, for his thorough review of my project and insightful feedback. I am grateful to Professor Cary Chi-Liang Tsai for serving on my examining committee and providing valuable suggestions. Special thanks to Professor Jinko Graham for chairing the examining committee. I am indebted to Professors Gary Parker, Cary Chi-Liang Tsai and Yi Lu for their equal measure of support, inspiration and kind motivation throughout my undergraduate and graduate studies. My gratitude extends to all the faculty members, staffs, and fellow students in the Department of Statistics and Actuarial Science for making my journey here so enjoyable.

Last but not the least, I wish to thank my family for their everlasting love and encouragement. Without their unconditional support, I would not have achieved this.

Table of Contents

List of Tables

List of Figures

Chapter 1

Introduction

1.1 Background and Motivation

This project explores cooperation in target benefit plans from the perspective of stochastic game theory. A target benefit plan (TBP) is a hybrid pension plan, which combines elements of both defined contribution (DC) plans and defined benefit (DB) plans. It looks like a DC plan in that the employer's contribution is usually fixed. A TBP also resembles a DB plan since a target benefit is defined at plan inception; it differs, however, in that this target is not guaranteed. The actual benefit received by members depends on plan experience and might be different from the target. Unlike a DB plan, the employer bears almost none of the investment, interest rate or demographic risks. Instead, these risks are borne by the employees as a group. How exactly these risks are shared is a key aspect of TBP design. Some TBPs have very little risk-sharing, resulting in volatile benefits that are adjusted frequently in response to emerging plan experience. Other TBPs aim to create more stability in pensions in pay by either adjusting active members' contributions or their future benefit accruals. In this case, stability for older generations comes at the expense of younger generations. Whatever the risk sharing arrangement looks like, a key question is whether it is sustainable.

In theory, pension plans are institutions that create enforceable, or binding, contracts; however, many pension contracts turned out over time not to be nearly as binding as expected. There are many instances of DB plan failures where employers promised to fund the pension plan but ultimately failed to do so. For example, the United Airlines pension plans failed in 2005 with a shortfall of \$10.2 billion, of which \$6.8 billion was funded by the Pension Benefit Guaranty Corporation, a U.S. government agency that insures pensions (Walsh, 2005). Even so, United Airlines' 130,000 employees and pensioners had lost \$3.4 billion worth of benefits (Walsh, 2005). More recently, the Sears bankruptcy left about 34,000 plan members without full benefits, including about 18,000 pensioners (Morgan, 2018). Therefore, even though a pension plan is binding in theory, in practice it may not be.

A particular feature of target benefit plans that may threaten sustainability is their reliance on intergenerational solidarity, especially where retired members' risks are underwritten by younger members. Solidarity can break down when there is intergenerational inequity; that is, the risk-sharing deal is not fair ex-ante so that there is a persistent one-way transfer of wealth. The ex-ante fairness of TBPs has already been explored by others, including Gollier (2008) and Cui et al. (2011). This is not our primary focus. Instead, we are interested in exploring the limits on intergenerational solidarity from an ex-post perspective.

Due to the uncertainty in investments and other plan experience, the transfers among generations tend to be different under each state of nature, and different over time. However, sometimes adverse conditions may persist. If, ex-post, one generation is asked to subsidize other generations for an extended period, they might lose patience. Similarly, if the size of the subsidy required in any given period is too large, then there is a risk that the generation who is asked to provide this subsidy (often the younger members) will not uphold their part of the deal any more, even if the deal was fair ex-ante. Hence, from an ex-post perspective, the size of transfers may be significant and thus can reduce the willingness of younger cohorts to participate.

The decision that the current young generation is making today regarding whether or not to break the deal depends on two things: the contribution they are required to pay today, which is known, and the benefit they will receive when they become old, which is uncertain since it depends on the continued cooperation of the new young generation that comes after them. The level of trust between these two generations will impact the decision that the current young will make: if there is no trust, the current young might want to be cautious and break their promise earlier, lest they be caught as the unlucky generation who is old when the deal is broken. So, the absence of trust is likely to make the pension contract fall apart earlier than in the case where complete trust is assumed to exist.

In this project we attempt to solve the problem of trust by making the contract self-enforceable. In other words, we are looking for a sustainable intergenerational "contract"; one that continues to be satisfactory or valuable to all stakeholders over time. One way to achieve this would be to put in place rules that do not ask generations to deviate from what would be in their self-interest in any given situation. We explore this idea in detail using the concepts and tools of game theory.

1.2 Cooperative versus Non-cooperative Frameworks

There are two possible frameworks in game theory to explore cooperation among agents.[1]

The first framework (Framework 1) is cooperative game theory. Cooperative game theory generally assumes that a binding agreement among agents can exist. Agents have an

[1] Agents are the decision makers in game theory.

incentive to form coalitions because the collective payoff of the coalition is greater than the sum of the individual payoffs would be in the absence of a coalition (Bilodeau, 1998, p. 177). For example, in the context of a TBP, Gollier (2008) has shown that pooling investment risk across generations improves certainty-equivalent returns. Cooperative game theory then focuses on the behaviour of coalitions and potential subcoalitions, and tries to allocate the collective payoff among agents fairly. Allocation rules must ensure that each agent's and each subcoalition's allocated payoff is at least as great as it would be in the absence of the coalition. In this case, forming a "collective" makes rational economic sense.

The second framework (Framework 2) is the non-cooperative game theoretic approach. Non-cooperative game theory is interested in individual agents' decision making and aims to maximize individual payoffs. As mentioned in Bilodeau (1998), every agent knows what he can achieve under all possible circumstances and takes into account other agents' possible decisions. In this framework, cooperation is achieved as an equilibrium such that no individual would deviate from the equilibrium path, out of his own interest to maximize payoffs.

It is worth noting that cooperation can be studied by using either the cooperative or non-cooperative game theoretic framework. The main difference is that the existence of the coalition is presumed in cooperative game theory, but is not assumed in non-cooperative game theory.

To cast TBPs in a game theoretic framework, we need three elements.

- We need a concept of self-enforceability over time to make sure that agents will not have an incentive to break the deal later, not just at plan inception. This should solve the absence of trust issue described above.
- We want uncertainty in payoffs to reflect the stochastic investment environment and specifically the impact of random investment experience on either the pension amount or on the size of the subsidy required to keep pension stable.
- We need an overlapping generation (OLG) structure because generations keep aging: the young become older and as a new generation joins, the old are replaced.

Framework 1 can handle the first two elements (self-enforceability and uncertainty in payoffs), but it cannot handle the third one (OLG structure). All of the cooperative game theoretic frameworks we explored assume that the objectives and characteristics of the agents in the game stay the same; for example, in a game played by a company and a worker, the workers are either assumed to be ageless and infinitely lived, or they are replaced by workers with the same characteristics and objectives every time. This is not true when the deal is among generations.

By contrast, Framework 2 can incorporate all three elements. Specifically, to make cooperation self-enforceable under this framework, the absence of intergenerational trust can be solved by making it costly to deviate from the agreed upon rules, potentially due to punishment by subsequent generations. In this way, the self-enforceability element can be implicitly captured by the OLG structure. So, we choose to use Framework 2 in this project.

1.3 Key Concepts in Non-cooperative Game Theory

In this section, the key concepts of non-cooperative game theory are introduced, based on the textbook by Fudenberg and Tirole (1991). In non-cooperative game theory, a game is represented by three elements: the set of agents who play the game, the information and strategies available to each agent at each decision point, and utility functions that give each agent's payoffs, or utilities, of each strategy profile. A strategy of an agent is a full contingency plan, which should completely describe the actions to be taken at any decision point. It specifies a probability distribution over an agent's actions at any decision point. A collection of each agent's strategy is called a strategy profile. Complete information of a game means that each agent knows everyone's payoff from each outcome. Perfect information of a game means that one agent acts at a time, and every agent knows all previous actions while making his decision.

A game can be represented by a game tree, which specifies the order of the moves, i.e., which agent moves when, as shown in Figure 1.1. A tree is a finite collection of ordered nodes containing all possible moves from each position. Each node fully describes all events (or the path) that preceded it, not just the "physical location" at a decision point. Besides the initial node, each node is required to have exactly one immediate predecessor, so there is only one path through the tree to a given node. The nodes that do not have any successors are called terminal nodes, which completely determine a path through the tree.

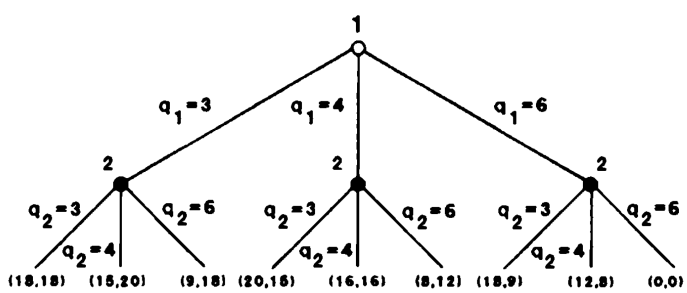

Figure 1.1: An example of the game tree where there are two agents, each choosing from the three possible outcomes: 3, 4 and 6, with payoffs being displayed at the end of the tree for agent 1 and 2, respectively. (Fudenberg and Tirole, 1991, p. 79).

A strategy is a "best response" for an agent if, given what other agents are doing, no other strategy yields a strictly higher expected payoff for this agent. In terms of predicting the outcome of a game, the Nash equilibrium is one solution, which specifies a strategy profile where every agent's strategy is a best response to all other agents' strategies. Basically, in a

Nash equilibrium, no agent wants to alter his strategy unilaterally, with each agent trying to maximize their own payoffs.

A repeated game involves repetition of the same game (the "stage game") to deal with situations where an agent will act based on other agents' actions in the previous round (i.e., in the previous "stage game"). A subgame is a part of the game that can be studied as a game on its own. For games with perfect information, as in repeated games, a subgame is any part of the game tree that starts at a non-terminal node, and that includes everything following that node up to the end of the game tree. A subgame can stand alone as a game.

A strategy profile of a game is a subgame perfect equilibrium if it is a Nash equilibrium for every proper subgame of the original whole game. This concept of the subgame perfect equilibrium supports cooperation among agents over time because if the agreed upon strategy profile is a subgame perfect equilibrium, then at any subgame, there is no incentive for any agent to deviate from the strategy profile unilaterally.

In repeated games, trigger strategies are considered. An agent using a trigger strategy starts with cooperating and will punish any deviators who defect. The severity of punishments and the level of defection to trigger punishment can vary. One popular trigger strategy is the grim strategy. The agents using the grim strategy cooperate until one agent defects and cooperation will break down forever if any defect ever happens. This strategy requires punishing the deviator forever even with one single deviation from the agreed upon strategy.

In repeated games, the concept of discounting future payoffs matters greatly. No discounting means that the agent's subjective time preference rate, β, is zero so that money in the future is valued the same as money today. When there is discounting ($\beta > 0$), money in the future is worth less than money today, which can result in myopic decision making. No discounting implies that future rewards are more valuable and future punishments are more severe than in the case with discounting.

In an infinitely repeated game, if agents are patient enough (i.e., there is no discounting), then the strategy profile with all agents following the same grim strategy is a subgame perfect equilibrium (Fudenberg and Tirole, 1991, p. 145). In repeated games, more cooperation can happen than in games played just once because agents consider the effects of their current actions on other agents' future actions in succeeding games. Cooperation is made even more likely by using the grim strategy because of agents' fear of punishment from subsequent generations. The infinite time horizon is also helpful because there will always be a succeeding generation to punish the previous deviator, which discourages the agents from deviating and leads to more cooperation. Therefore, intuitively, a sufficiently patient agent will care more about the long-run impact of punishment than the short-term gain from defection (Fudenberg and Tirole, 1991, p. 145).

1.4 Deterministic Models with Overlapping Generations

Suppose we replace the infinitely lived agents with finitely lived overlapping generations in the above formulation. We then have an infinitely repeated discrete time game without discounting with overlapping generations of agents each using the grim strategy. Cremer (1986) illustrates that cooperation can still be supported as a subgame perfect equilibrium in this case. The use of the grim strategy, the absence of discounting, and the infinite time horizon support cooperation in the same way as was discussed in the previous section. In the model of Cremer (1986), a fourth feature that also enhances cooperation is that the oldest agents are not required to put in effort in their last period (just as, for example, pensioners receive benefits without any further contributions). This can be viewed as a reward in the future to motivate cooperation.

Cremer's model particularly interests us as we try to adapt a framework to the pension context. First, the infinite time horizon helps us study the long-term interaction among finitely lived generations. Here the infinite time does not necessarily mean pension plans last forever, it just means the end day is indefinite. Second, the OLG structure captures the succession of finitely lived agents. Third, it makes sense to have agents put in effort while young and enjoy its benefits while old.[2]

Similar to Cremer's setup, Kandori (1992) studies cooperative behaviour between patient agents in infinitely repeated games. First, Kandori (1992) notes that in repeated games theory, Folk's theorem implies that "any mutually beneficial outcome can be sustained in an ongoing relationship between infinitely-lived agents" (p. 81). In the original theorem, the identities of infinitely-lived agents are the same. Kandori (1992) also extends this theorem to a more realistic setup which is the infinitely repeated game with finitely lived agents and an OLG structure. Using the concept of subgame perfect equilibrium, it is shown that "any mutually beneficial outcome can approximately be sustained if the agent's life span and the overlapping periods are long enough" when there is no discounting (Kandori, 1992, p. 81).

As described in Esteban and Sákovics (1993), intergenerational transfers can also be made self-enforceable by establishing a costly institution[3] to facilitate transfers from the young to the old, in an infinite discrete time OLG model with two-period living agents. This is because the agents who choose to defect will need to pay a cost to establish a new

[2]From an economic perspective, young agents have the highest marginal productivity of labour (Cremer, 1986). From a game theory perspective, only young agents are threatened by successors: since old agents will not survive until the next period for further consequences, they do not care about the threat in the next period. So, the young are expected to put in effort to form cooperation, however we cannot expect the old to put in much effort.

[3]Institutions are "social artifacts" that "crystallize agreements by means of built-in rules" (Esteban and Sákovics, 1993, p. 190). An institution is costly if there is a significant price to be paid for creating or changing it. A pension plan is an example of a costly institution.

institution, which discourages them from defection. This is different from Cremer (1986) and Kandori (1992) where cooperation is enforced by subsequent generations.

Bhaskar (1998) studies how the agents' knowledge regarding past events affects the sustainability of intergenerational transfers in an OLG model with two-period living agents. This limited knowledge is different from Cremer (1986) and Kandori (1992), as both those studies consider agents that know all past events. Specifically, when agents only know information from the past m periods (imperfectly informed about past events), different equilibria can arise.

As Bhaskar (1998) points out, Hammond (1975) has already found that if each agent knows all previous agents' actions, there exist subgame perfect equilibria where transfers could be sustained. Bhaskar (1998) also sets out two examples of cooperation enforcement or subgame perfect equilibrium strategy profiles where each agent makes the transfer decision based on the actions of previous agents. The first one is the grim strategy mentioned previously, which is non-forgiving in the sense that once a deviation happens, cooperation breaks down forever. The second one is the so-called resilient strategy in which agents are punished only if they make no transfer when they should. In this case, even if someone deviates, after some finite deviation periods, cooperation can still be restored. This equilibrium strategy improves the sustainability of cooperation and is an ideal circumstance we hope to achieve in the context of intergenerational cooperation within pension plans.

1.5 Stochastic OLG Models

Having addressed the OLG structure and the issue of self-enforceability over time, the uncertainty aspect still needs to be considered. This notion has been studied in two papers: Messner and Polborn (2003) and Miyazaki (2014).

Messner and Polborn (2003) introduce uncertainty to a model similar to that of Cremer (1986), adding random shocks to the cost of cooperation in an infinitely repeated OLG game where agents are two-period living. In their setup, it is the young agents' decision to cooperate or not with the old agents. If they decide to cooperate, the young agent will pay a cost of cooperation, and the current old agent will receive \$1. There is a threshold for the cost of cooperation for the young to make cooperation decisions; that is, young agents will cooperate only when the cost is not above the threshold. Otherwise the young will choose not to cooperate this time, i.e., the young pay nothing and the old receive nothing. Messner and Polborn (2003) focus on showing the parameter space for which cooperation can be sustained in a subgame perfect equilibrium. They show that by adding even a very small uncertainty to the cost of cooperation, the parameter space where cooperation can be sustained as a subgame perfect equilibrium decreases dramatically compared to the deterministic case.

The model setup in Messner and Polborn (2003) is very similar to that in Cremer (1986), but with two major extensions. One is the uncertainty in payoffs: the cost of cooperation is assumed to follow a stochastically monotone Markov process. The other one is defining a "correct behaviour" strategy profile, which is essentially a resilient strategy, as opposed to the grim strategy used by Cremer (1986). A more relaxed strategy is needed because, due to the uncertainty in the cost of cooperation, the grim strategy would make cooperation too easy to break without the opportunity to restore later, which seems too harsh. As mentioned previously, a resilient equilibrium strategy makes sense in the context of intergenerational cooperation in pensions. In pension plans, it should be reasonable for the young generation to refuse cooperation when the cost of underwriting the old generation is too expensive, and later return to cooperation when the cost is more manageable. We want the equilibrium strategy to be cooperation reverting after some finite periods of deviation. From this angle, the framework proposed by Messner and Polborn (2003) is desirable in terms of studying sustainability of intergenerational cooperation.

Miyazaki (2014) considers a slightly different model where both the young and the old are given some endowments at each discrete time point. These endowments could be interpreted as investment returns or other plan experience (e.g., surplus resulting from favorable demographic changes). Randomness comes from endowment shocks that affect both the young and the old agents. In this framework, the central question becomes how to allocate the total endowment between the young and the old agents in order to maintain subgame perfect equilibrium. Miyazaki (2014) focuses on golden-rule type allocations[4] and their characteristics.

The main differences between Miyazaki (2014) and Messner and Polborn (2003) are as follows:

- Miyazaki (2014) assumes that both the young and the old agents are given some endowments at each time, which are then shared via transfers. Specifically, transfers from the old to the young are possible, which are not considered in Messner and Polborn (2003).
- Miyazaki (2014) and Messner and Polborn (2003) establish the stochastic payoffs in different ways, which may or may not be consistent. Miyazaki (2014) assumes an independent and identically distributed stochastic process for endowment shocks, where the shock at each time is from a given finite set with a known probability. Messner and Polborn (2003) assume that the cost of cooperation follows a stochastically monotone Markov process.

Even though some aspects of Miyazaki (2014) are conceptually attractive, such as the existence of specific rules for allocating the endowments to ensure subgame perfect equilib-

[4]"An allocation is a golden-rule type allocation if the allocation maximizes the weighted sum of the young agent's conditional expected lifetime utility" (Miyazaki, 2014, p. 500).

rium, ultimately we found this framework difficult to adapt to our interest. Therefore, we chose the Messner and Polborn (2003) framework which captures all the aspects we need. Using this framework, we are not trying to allocate surpluses or endowments, or figure out how much money each generation should get. Instead, we focus on the interaction or cooperative behaviour within the pension plan and try to determine the maximum amount of money that the young generations can be expected to give to the old generations.

In what follows, we abstract the target benefit plan into a simple non-cooperative framework by treating each generation as an agent who makes decisions to maximize individual utilities or payoffs. In this case, cooperation can happen only if it is within the agent's self interest. We extend the framework developed by Messner and Polborn (2003) by introducing a fixed cost in addition to the random cost of cooperation, and explore pension sustainability by asking what is the maximum threshold for cooperation and how often cooperation can be achieved. Chapter 2 describes our model in detail. We show the parameter space for when cooperation can be sustained in a subgame perfect equilibrium and the behaviour of the threshold under various assumptions in Chapters 3 and 4. Chapter 5 concludes and makes suggestions for future extensions.

Chapter 2

Model and Assumptions

2.1 Model Setup

The original framework of Messner and Polborn (2003) has no storage capacity for future consumption and the only explicit source of the old generation's benefit is the subsidy provided by the young generation. By contrast, TBPs normally operate with prefunding, which means that there is some money set aside ahead of time to pay for each generation's future retirement benefits. So, in the model with prefunding the old generation has some money saved for themselves in advance and is only partially funded by the young generation. Hence, we expand the original model to include contributions to cover some or all of the estimated cost of benefits as well as storage capacity (for the investment). The uncertainty in investment returns then gives rise to the uncertainty in each young generation's actual costs.

2.1.1 Extension with Prefunding

Consider an OLG structure in an infinitely repeated game, where agents are two-period living. At every discrete time point, an agent enters the pension plan as a young generation and makes a contribution to the plan. One period later, he becomes the old generation and collects his pension payment. As agents keep being replaced (i.e., a new young agent enters the plan; the original young becomes old; and the original old dies), at every time point there exists a young agent and an old agent. An infinite time horizon is used because a pension plan is ongoing, and the end day is indefinite. The single agent in each generation can be generalized to multiple agents with the same preference who will take the same action. For simplicity, we use one agent per generation.

We introduce a fixed cost m paid by the young generation. This cost is payable by the young regardless of whether they decide to cooperate or not, and is calculated as the present value of a \$1 benefit receivable one period later using a valuation interest rate of i_v. Then,

$$0 < m = \frac{1}{1+i_v} < 1, \text{ and}$$

$$d = 1 - m = \frac{i_v}{1 + i_v}.$$

Note that, with two-period living agents, each period is actually as long as 40 years. As a result, i_v is a 40-year rate. Some reasonable values of i_v and corresponding values of m are shown in Table 2.1.

Annual Effective Rate	i_v	m
2%	121%	0.5
4%	380%	0.2
5%	604%	0.1

Table 2.1: Values of i_v and corresponding values of m.

The case with $m = 0$ is equivalent to the pay-as-you-go pension system with no prefunding. We denote the investment return during the period $[t, t + 1)$ by i_{t+1}. Suppose an agent enters the plan at time t. The fixed contribution m made by this agent is assumed to earn investment return i_{t+1} while the agent is young, and the accumulated value of this contribution one period later, when the agent is old, is $m(1 + i_{t+1})$. The targeted benefit of the old is \$1, but it may be less if cooperation breaks down.

At each time point, the young agent gets to choose whether to cooperate or not. The young agent who decides to cooperate at t will pay the fixed cost m mentioned above, plus a stochastic cost of cooperation γ_t required in order for the current old agent to receive their fixed \$1 benefit. Basically, the young agent pays the fixed entrance cost m to prefund his own retirement, which will be invested and paid back to him when he is old. On the other hand, the stochastic cost of cooperation γ_t funds the deficit of the current old generation. From the perspective of the old agent at time t, the \$1 benefit payable to him was partially funded by himself when he was young: this is the m contributed at time $t - 1$ accumulating with interest, $m(1 + i_t)$. The rest comes from the next generation's variable cost of cooperation, γ_t, where

$$\gamma_t = 1 - m(1 + i_t).$$

If the young agent chooses to defect at time t, he only pays m and the current old only receives the self-financing portion $m(1 + i_t)$ without the young agent making up the deficit.

Intuitively, when the cost of funding the old agent's deficit is too high, the young agent would rather be on his own. So there will be a threshold for the cost of cooperation, γ^T, above which a young agent will not be willing to cooperate.

In Messner and Polborn (2003), in order to achieve subgame perfect equilibrium, a resilient type of strategy profile is defined to set the rule for the initial path and punishment paths. It is called the "correct behaviour" strategy profile, where young agents' decisions are made based on comparing the cost of cooperation with the threshold, γ^T, and taking into account the actions of previous agents. We assume that all agents follow this "correct behaviour" strategy profile to achieve equilibrium.

Definition 1. (reproduced from Messner and Polborn, 2003, p. 155) *The "correct behaviour" for the young agent in the first period i.e., $t = 1$, is to cooperate if and only if $\gamma_1 \leq \gamma^T$. The correct behaviour for the young agent at time $t > 1$ is to cooperate if and only if $\gamma_t \leq \gamma^T$ and the young agent in the previous period $t-1$ behaved correctly.*

As implied by the correct behaviour, an agent will be punished by the next generation if he defects when $\gamma \leq \gamma^T$, or if he fails to punish his predecessor who did not behave correctly. In this project, we focus on the maximum threshold γ^*, i.e., the greatest value of γ^T such that the correct behaviour strategy profile is still a subgame perfect equilibrium. This maximum threshold will result in more cooperation than any other choice of γ^T.

Under the original model, the result of non-cooperation is loss of all benefits: the generation who did not follow the correct behaviour when they were young gets nothing when old, which is quite a severe punishment. When a fixed cost m is added, non-cooperation at time t results in the loss of $\gamma_{t+1} = 1 - m(1 + i_{t+1})$ when the noncompliant agent becomes old. Depending on the size of m, this punishment may not be very severe. In fact, it may not be serious enough to create a subgame perfect equilibrium. To motivate cooperation, we assume an additional punishment. If an agent does not follow the correct behaviour, one period later he will receive only m, the money he originally put in. The deviator is punished by being deprived not only of the next generation's support but also his own investment returns.

Overall, this setup is a crude but appropriate model of TBPs. The variables m and γ_t can be interpreted as contributions, where m is the fixed part that each generation puts into the pension plan, and γ_t is the variable part of the contribution paid by the young generation in order to partially underwrite the benefit of \$1 promised to the old generation. The pension plan features vary depending on the decisions made by the young. On the surface, if cooperation is sustained, the plan looks like a DB plan with fixed benefits and variable contributions, except that the employees are responsible for these variable contributions instead of the employer. The variability in the benefit appears when cooperation breaks down. In this case, the old generation has to take care of itself. This way we can have both contributions and benefits variable, as in a TBP.

We then want to characterize the "sustainability" of the TBP in question: given a target benefit level (\$1), a fixed cost (m) and some process that governs the cost of cooperation (or implicitly, the investment return), what would be the maximum "sustainability threshold" (γ^*) from the perspective of the younger generation and how often would this threshold be breached?

2.1.2 Assumptions

We make the following important assumptions. The first three assumptions are the same as in Messner and Polborn (2003), but the fourth one is different.

- Agents are risk-neutral.

- There is no discounting, which means agents are not myopic. In other words, the concern over the long-term punishment exceeds the short-term gain from noncompliance.
- Every agent plays the correct behaviour strategy.
- The stochastic cost of cooperation, γ_t, follows a process that satisfies Condition 1 below.

In Messner and Polborn (2003), the stochastic cost of cooperation, γ_t, is required to satisfy the stochastically monotone Markov property, which is defined as:

$$Pr(\gamma_t > \overline{\gamma} \,|\, \gamma_{t-1} = \overline{\gamma}) \geq Pr(\gamma_t > \overline{\gamma} \,|\, \gamma_{t-1} = \underline{\gamma}), \quad \forall\, \overline{\gamma} > \underline{\gamma}. \tag{2.1}$$

Messner and Polborn (2003) require that the probability of γ_t exceeding some fixed value $\overline{\gamma}$ be greater if this value $\overline{\gamma}$ is already reached at $t-1$ than if starting from a smaller value $\underline{\gamma}$ at $t-1$. However, with the extra prefunding, m, this condition is no longer enough to guarantee the existence of a subgame perfect equilibrium. We impose a stronger condition defined as follows, which is a sufficient but perhaps not necessary condition.

Condition 1. *Let $\pi(y\,|\,x)$ be the one-step transition density from x to y. The stochastic process is required to be a Markov process that satisfies:*

$$\pi(\overline{\gamma} \,|\, \widetilde{\gamma}) \geq \pi(\overline{\gamma} \,|\, \underline{\gamma}), \quad \forall\, \overline{\gamma} \geq \widetilde{\gamma} > \underline{\gamma}. \tag{2.2}$$

In our model, the probability of γ_t increasing from time $t-1$ to t (that is, $\gamma_t \geq \gamma_{t-1}$), and reaching any fixed value $\overline{\gamma}$ is required to be higher if γ_{t-1} is higher.

We introduce the following notation for the probabilities that cooperation breaks down at time t given last period's cooperation cost:

$$p = Pr(\gamma_t > \gamma^* \,|\, \gamma_{t-1} = \gamma^*), \text{ and}$$

$$p' = Pr(\gamma_t > \gamma^* \,|\, \gamma_{t-1} < \gamma^*).$$

Here, γ^* is the maximum threshold defined above; that is, the greatest value of γ^T such that the correct behaviour strategy is still a subgame perfect equilibrium.

Similarly, we introduce the following notation for the expected cost of cooperation at time t given that the cost exceeds the threshold and conditional on the cost of cooperation in the previous period:

$$E = E[\gamma_t \,|\, \gamma_t > \gamma^*,\ \gamma_{t-1} = \gamma^*], \text{ and}$$

$$E' = E[\gamma_t \,|\, \gamma_t > \gamma^*,\ \gamma_{t-1} < \gamma^*].$$

This is also the expected amount lost by old agents if cooperation breaks down at time t. It is shown in Appendix A that (2.2) implies

$$pE \geq p'E' \tag{2.3}$$

and that a process that satisfies Condition 1 also has the stochastic monotone Markov property in (2.1).

2.1.3 Processes that Satisfy Condition 1

Messner and Polborn (2003) use two different binomial random walk processes and a general AR(1) process to illustrate their model, which have the stochastically monotone Markov property. However these processes do not work in our model because they do not satisfy Condition 1. In our case, the Gaussian random walk and the AR(1) process with $\mu = 0$ are examples of processes that satisfy Condition 1.

Proposition 1. *An AR(1) process with $\mu = 0$ and $0 < \alpha < 1$ satisfies Condition 1.*

Proof. Suppose the cost of cooperation, γ_t, follows an AR(1) process with $\mu = 0$, $0 < \alpha < 1$ and $\sigma > 0$, where μ is the long-term mean of the process, α is the drift term that controls how fast the process returns to the long-term mean, and σ is the local volatility parameter for the noise term, ε_t. The evolution of γ_t can be described by

$$\gamma_t = \alpha\gamma_{t-1} + \varepsilon_t, \text{ where } \varepsilon_t \overset{iid}{\sim} \mathcal{N}(0, \sigma^2),\ \ 0 < \alpha < 1,\ \ \sigma > 0. \tag{2.4}$$

It follows that the conditional density of the cost of cooperation at time t given the cost of cooperation at the previous time point is normally distributed. Specifically,

$$\gamma_t | \gamma_{t-1} \sim \mathcal{N}(\alpha\gamma_{t-1}, \sigma^2).$$

Then, for all $\overline{\gamma} \geq \widetilde{\gamma} > \underline{\gamma}$,

$$\pi(\overline{\gamma} \,|\, \widetilde{\gamma}) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{\left(\overline{\gamma} - \alpha\widetilde{\gamma}\right)^2}{2\sigma^2}} \geq \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{\left(\overline{\gamma} - \alpha\underline{\gamma}\right)^2}{2\sigma^2}} = \pi(\overline{\gamma} \,|\, \underline{\gamma}). \tag{2.5}$$

□

Corollary 1. *A Gaussian random walk process also satisfies Condition 1.*

Proof. If the cost of cooperation γ_t follows a Gaussian random walk with parameter σ, such that

$$\gamma_t = \gamma_{t-1} + \varepsilon_t, \text{ where } \varepsilon_t \overset{iid}{\sim} \mathcal{N}(0, \sigma^2), \tag{2.6}$$

then, it follows that

$$\gamma_t | \gamma_{t-1} \sim \mathcal{N}(\gamma_{t-1}, \sigma^2).$$

The proof is thus the same as above with $\alpha = 1$.

$\square$

For both processes described above, there are relationships between the distributions of γ_t and i_t. For the AR(1) process with $\mu = 0$, from Equation (2.4), it follows that

$$1 - m(1 + i_t) = \alpha[1 - m(1 + i_{t-1})] + \varepsilon_t,$$

which implies

$$i_t = \alpha i_{t-1} + \frac{(1-m)(1-\alpha)}{m} - \frac{\varepsilon_t}{m} = \alpha^* i_{t-1} + \mu^* + \varepsilon_t^*,$$

where $\varepsilon_t^* = -\frac{\varepsilon_t}{m} \overset{iid}{\sim} \mathcal{N}(0, \frac{\sigma^2}{m^2})$. So, i_t also follows an AR(1) process with parameters

$$\alpha^* = \alpha, \text{ and}$$

$$\mu^* = \frac{(1-m)(1-\alpha)}{m}.$$

For example, a (reasonable) value of $m = 0.3$ and $\alpha = 0.4$ produce a long-term compound mean return $\mu^* = 1.4$ over a 40-year horizon, corresponding to an annualized return of about 2.2%.

Similarly, when γ_t follows a Gaussian random walk process, i_t also follows a Gaussian random walk process. This can be seen by setting $\alpha = 1$ in the above equations.

As noted above, in our model each time period could possibly represent 40 years. It would be interesting to know if the two processes that we considered are realistic. Specifically, is it reasonable to assume autocorrelation and mean reversion in 40-years returns? Unfortunately, there is not enough data for us to test whether these properties hold. Thus, we cannot confirm whether a Gaussian random walk or an AR(1) process with $\mu = 0$ are reasonable choices.

2.2 Expected Net Payoff of Threshold Cooperators

We borrow the concept of the "threshold cooperator" from Messner and Polborn (2003). The threshold cooperator can be defined as the young agent whose cost of cooperation is at the threshold, γ^*.

A threshold cooperator who enters at $t-1$ with $\gamma_{t-1} = \gamma^*$ is supposed to cooperate based on the correct behaviour strategy described above. If he behaves correctly and chooses to cooperate, his total cost at time $t-1$ is $m + \gamma^*$. One period later, when he becomes old at t, he can either receive \$1 if cooperation continues ($\gamma_t \leq \gamma^*$) or receive $m(1+i_t)$ if cooperation breaks down ($\gamma_t > \gamma^*$). On the other hand, if he deviates from the correct behaviour strategy by choosing not to cooperate at $t-1$, he can only receive the original m he contributed, without any investment returns.

Figure 2.1 illustrates the decision faced by the threshold cooperator at $t-1$ ($\gamma_{t-1} = \gamma^*$) and the consequences of that decision.

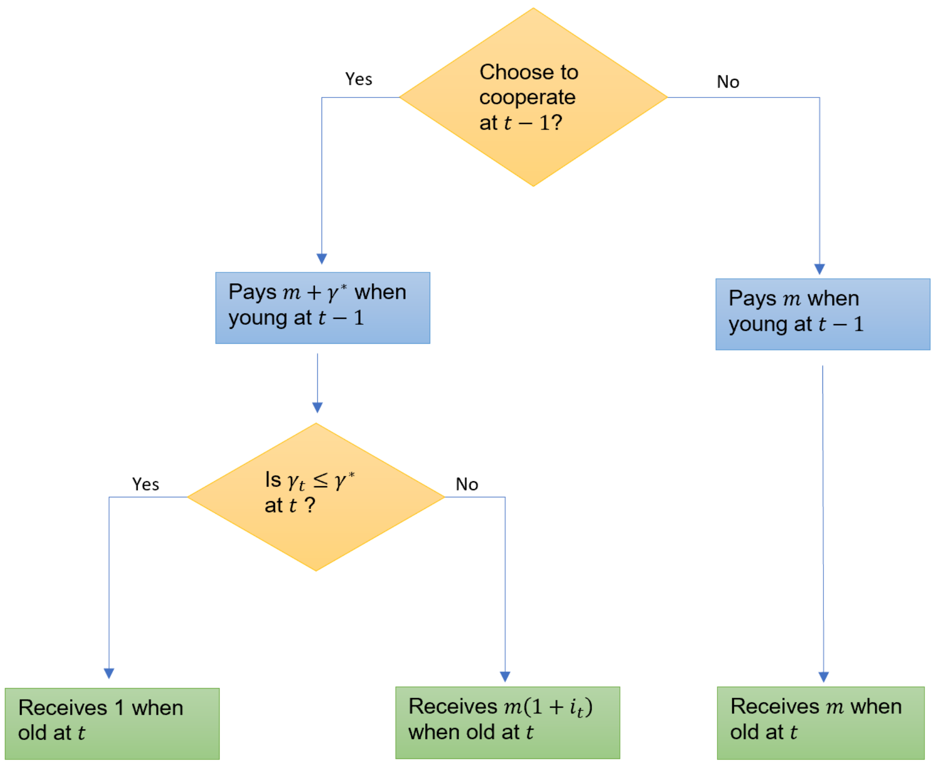

Figure 2.1: Flow chart for the threshold cooperator at $t-1$ (i.e., the agent with $\gamma_{t-1} = \gamma^*$).

The expected net payoffs under the various options facing the threshold cooperator are derived as follows. First, we look at his expected net payoff if he chooses to follow the correct behaviour strategy.

$$
\begin{aligned}
&E[\text{net payoff of a threshold cooperator who follows the correct behaviour strategy at } t-1] \\
&= E[\text{net payoff of a threshold cooperator who cooperates at } t-1] \\
&= E[\text{benefit received at } t \text{ by a person who was threshold cooperator at } t-1] \\
&\quad - \text{cost of cooperation at } t-1 \\
&= E[\text{benefit received at } t \mid \text{cooperation at } t]\ Pr[\text{cooperation at } t] \\
&\quad + E[\text{benefit received at } t \mid \text{no cooperation at } t]\ Pr[\text{no cooperation at } t] - m - \gamma^* \\
&= 1 \times (1-p) + E\left[m(1+i_t) \mid \gamma_t > \gamma^*,\ \gamma_{t-1} = \gamma^*\right] \times p - m - \gamma^* \\
&= (1-p) + E\left[1-\gamma_t \mid \gamma_t > \gamma^*,\ \gamma_{t-1} = \gamma^*\right] \times p - m - \gamma^* \\
&= (1-p) + \left\{1 - E\left[\gamma_t \mid \gamma_t > \gamma^*,\ \gamma_{t-1} = \gamma^*\right]\right\} \times p - m - \gamma^* \\
&= (1-p) + (1-E) \times p - m - \gamma^* \\
&= 1 - pE - m - \gamma^*
\end{aligned}
\tag{2.7}
$$

Next, we consider the expected net payoff of a threshold cooperator who deviates from the correct behaviour strategy at $t-1$:

$$
\begin{aligned}
&E[\text{net payoff of a threshold cooperator who deviates from the correct behaviour} \\
&\quad \text{strategy at } t-1] \\
&= E[\text{net payoff of a threshold cooperator who does not cooperate at } t-1] \\
&= E[\text{benefit received at } t \text{ by a person who was threshold cooperator at } t-1] \\
&\quad - \text{cost of cooperation at } t-1 \\
&= E[m \mid \gamma_{t-1} = \gamma^*] - m \\
&= 0
\end{aligned}
\tag{2.8}
$$

When the cost of cooperation follows a stochastically monotone Markov process, threshold cooperators are the agents who are least likely to be subsidized by the next generation. For example, at $t-1$, a young agent enters the plan and faces a cost of cooperation $\gamma_{t-1} = \gamma^*$; if the cost increases the next period, so that $\gamma_t > \gamma^*$, the next young agent will not cooperate at time t. In this case, the original threshold cooperator's deficit cannot be made up at t. To keep agents at the threshold cooperating, their payoff from cooperation must be high enough. Messner and Polborn (2003, p. 153) clarified that this is true when cooperation is efficient, which in our framework corresponds to $m + \gamma^* < 1$.

2.3 Expected Net Payoff of Agents Below the Threshold

The agent whose cost of cooperation is below the threshold ($\gamma_{t-1} < \gamma^*$) faces decisions that are very similar to those of a threshold cooperator. He is also supposed to cooperate at $t-1$ according to the correct behaviour strategy and pay $m + \gamma_{t-1}$, so that when he becomes old at t, he can either receive \$1 if cooperation continues ($\gamma_t \leq \gamma^*$); or receive $m(1+i_t)$ if cooperation breaks down ($\gamma_t > \gamma^*$). But if he chooses to deviate from the strategy by not cooperating at $t-1$, he can only receive the original m he put in and thus lose the investment return.

The options faced by the agent below the threshold at $t-1$ (i.e., $\gamma_{t-1} < \gamma^*$) are as shown in Figure 2.2.

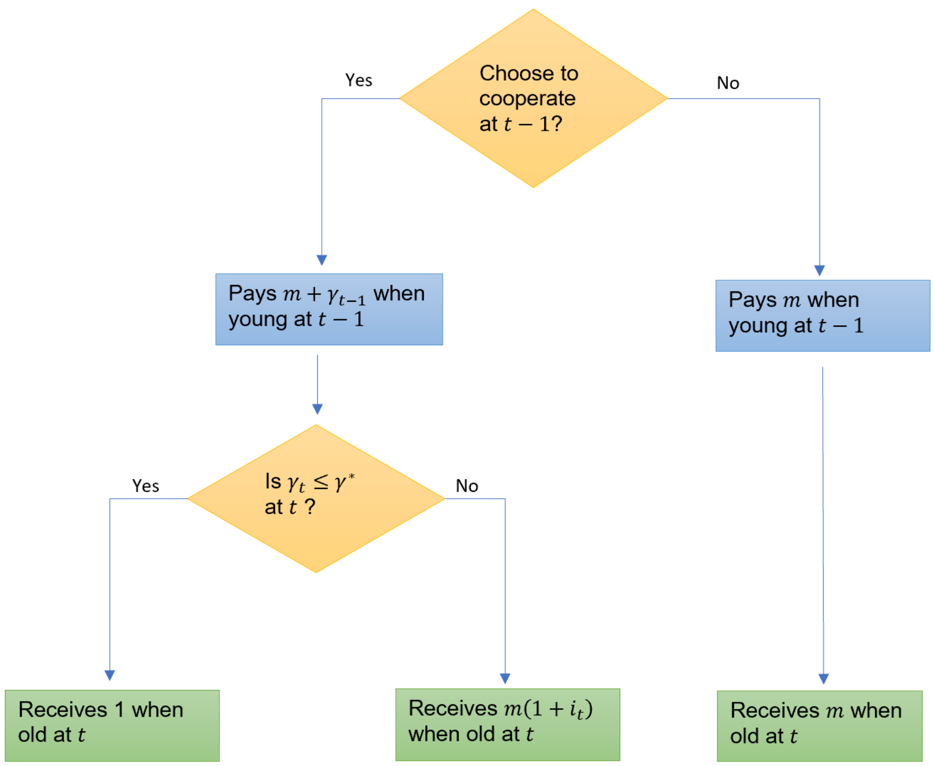

Figure 2.2: Flow chart for the agent below the threshold at $t-1$ (i.e., the agent with $\gamma_{t-1} < \gamma^*$).

The expected net payoffs under the various options facing the agent below the threshold at $t-1$ are derived as follows. First, we consider the expected net payoff in case of following the agreed-upon strategy.

$$\begin{aligned}
&E[\text{net payoff of a "below threshold" agent who follows the correct behaviour strategy}\\
&\quad \text{at } t-1]\\
&= E[\text{net payoff of anyone below the threshold who cooperates at } t-1]\\
&= E[\text{benefit received at } t \text{ by anyone below the threshold at } t-1]\\
&\quad - \text{cost of cooperation at } t-1\\
&= E[\text{benefit received at } t \mid \text{cooperation at } t]\ Pr[\text{cooperation at } t]\\
&\quad + E[\text{benefit received at } t \mid \text{no cooperation at } t]\ Pr[\text{no cooperation at } t] - m - \gamma_{t-1}\\
&= 1 \times (1-p') + E[m(1+i_t)|\ \gamma_t > \gamma^*,\ \gamma_{t-1} < \gamma^*] \times p' - m - \gamma_{t-1}\\
&= (1-p') + E[1-\gamma_t|\ \gamma_t > \gamma^*,\ \gamma_{t-1} < \gamma^*] \times p' - m - \gamma_{t-1}\\
&= (1-p') + \{1 - E[\gamma_t|\ \gamma_t > \gamma^*,\ \gamma_{t-1} < \gamma^*]\} \times p' - m - \gamma_{t-1}\\
&= (1-p') + (1-E') \times p' - m - \gamma_{t-1}\\
&= 1 - p'E' - m - \gamma_{t-1}
\end{aligned}$$

Next, we derive the expected net payoff of a "below threshold" agent who deviates from the correct behaviour strategy at $t-1$:

$$\begin{aligned}
&E[\text{net payoff of a "below threshold" agent who deviates from the correct behaviour}\\
&\quad \text{strategy at } t-1]\\
&= E[\text{net payoff of anyone below the threshold who does not cooperate at } t-1]\\
&= E[\text{benefit received at } t \text{ by anyone below the threshold at } t-1]\\
&\quad - \text{cost of cooperation at } t-1\\
&= E[m|\ \gamma_{t-1} < \gamma^*] - m\\
&= 0
\end{aligned}$$

As mentioned above, threshold cooperators are the agents whose deficits are least likely to be funded when they are old. Any young agent whose cost of cooperation is below the threshold is in a better situation. For example, suppose a young agent enters at $t-1$ when $\gamma_{t-1} < \gamma^*$; even if the cost of cooperation increases in the next period, it may still be below the threshold γ^*, and the next young agent may still cooperate at time t.

2.4 Subgame Perfectness

As discussed earlier, a resilient strategy profile leads to a subgame perfect equilibrium.

Proposition 2. *Suppose that the cost of cooperation, γ_t, develops according to a process that satisfies Condition 1. Then, the most cooperative subgame perfect equilibrium is defined implicitly by the greatest solution of*

$$1 - pE - m - \gamma^* = 0. \tag{2.9}$$

Proof. Intuitively, a strategy produces a subgame perfect equilibrium if it is not in the agents' interest to deviate from the strategy, i.e., nobody is better off by deviating. So we need to show that the expected net payoff for behaving correctly is always at least as much as the expected net payoff for not behaving correctly.[1] As threshold cooperators are the ones who are least likely to behave correctly (because they are most likely not to be subsidized when old), they must be the ones who produce the lowest expected payoffs from cooperation. The below-threshold agents should always have higher expected net payoffs than threshold cooperators if both of them choose to behave correctly. To make threshold cooperators indifferent between their two choices (that is, behaving correctly and not behaving correctly) their expected net payoffs should be equal under both cases. This gives rise to Equation (2.9). The rest of the proof is divided into two steps, which make sure that a below threshold agent always has incentive to behave correctly in a subgame perfect equilibrium.

Step 1. Show that for any young agent who is below the threshold, their expected net payoff when cooperating at $t-1$ is no less than their expected net payoff when they are not cooperating at $t-1$.

$$\begin{aligned}
&E[\text{net payoff of agent below the threshold if cooperating at } t-1] \\
&= 1 - p'E' - m - \gamma_{t-1} \\
&> 1 - p'E' - m - \gamma^* \\
&\geq 1 - pE - m - \gamma^*, \text{ by (2.3)} \\
&= 0, \text{ by (2.9)}.
\end{aligned} \tag{2.10}$$

Step 2. Show that, for any young agent who is below the threshold, their expected net payoff is no less than the threshold cooperator's expected net payoff when both of them choose to behave correctly.

$$\begin{aligned}
&E[\text{net payoff of agent below the threshold if cooperating at } t-1] \\
&> 0, \text{ by (2.10)} \\
&= E[\text{net payoff of a threshold cooperator if cooperating at } t-1]
\end{aligned} \tag{2.11}$$

On the other hand, if both agents, the threshold cooperator and the below-threshold agent, choose to breach at $t-1$, their expected net payoffs are both equal to 0. $\square$

[1] Because we assume risk neutrality, individuals' utility of payoffs is proportional to the payoffs.

Proposition 3. *Let* $\gamma^* > 0$ *and* $Pr(m + \gamma_t > 1 | \ \gamma_0) > 0$ *for all* γ_0 *and at least some* t*, together with Condition 1, then the greatest solution of Equation* (2.9)*,* $1 - pE = m + \gamma^*$*, satisfies* $\gamma^* < 1 - m$*.*

Proof. The requirement that $Pr(m + \gamma_t > 1 | \ \gamma_0)$ be positive for all γ_0 and at least some t means that it is possible for the cost $m + \gamma_t$ to exceed \$1 at some point in time. It implies that for any realization of the process $\{\gamma_t\}_{t=0}^{\infty}$ with $\gamma_0 < 1$, there exists a $\tau \in \mathbb{N}$ such that $m + \gamma_\tau \leq 1$ and $m + \gamma_{\tau+1} > 1$.

Then it follows that

$$Pr(m + \gamma_t > 1 \,|\, m + \gamma_{t-1} \leq 1) > 0, \text{ for some } t \geq 0,$$

which is equivalent to

$$Pr(\gamma_t > 1 - m \,|\, \gamma_{t-1} \leq 1 - m) > 0, \text{ for some } t \geq 0.$$

Therefore, by Condition 1,

$$Pr(\gamma_t > 1 - m \,|\, \gamma_{t-1} = 1 - m) \geq Pr(\gamma_t > 1 - m \,|\, \gamma_{t-1} \leq 1 - m) > 0, \tag{2.12}$$

for some $t \geq 0$.

Since $\pi(x \,|\, \gamma^*) \geq 0$, and $\gamma^* > 0$, $pE = \int_{\gamma^*}^{\infty} x \cdot \pi(x \,|\, \gamma^*) dx \geq 0$. We show $pE > 0$ by contradiction. Suppose $pE = 0$. This implies that $\pi(x \,|\, 1 - m) = 0$ for all $x > 1 - m$. But then we would have

$$\begin{aligned} & Pr(\gamma_t > 1 - m \,|\, \gamma_{t-1} = 1 - m) \\ & = \int_{1-m}^{\infty} \pi(x \,|\, 1 - m) dx \\ & = 0, \end{aligned} \tag{2.13}$$

which conflicts with Equation (2.12), so $pE \neq 0$. Therefore, we must have $pE > 0$, and, by Equation (2.9), we have that

$$1 - pE = m + \gamma^* < 1,$$

and

$$\gamma^* < 1 - m.$$

□

It is reasonable to require γ^* be positive because otherwise young agents would cooperate only when they pay nothing or receive some compensation from the old. In this case, cooperation offers no benefit to the old generation.

Definition 2. *The feasible region is the set of all parameters for which* $\gamma^* > 0$*.*

Note that even though $\gamma^* > 0$, γ_t can still be negative, in which case retirees transfer wealth to the young in good times. So in our setup, the young benefit from the upside potential in exchange for bearing the downside risk for retirees. However, there is an asymmetry between the upside potential and downside risk in the sense that the cost of cooperation is capped from above but not from below. Since the young are the only ones who make cooperation decisions, they can grasp all the upside potential from the old but only absorb the downside risk up to a reasonable threshold, as there is no lower threshold when γ_t is negative.

Chapter 3

Application to the Gaussian Random Walk

In this chapter, we assume that the cost of cooperation, γ_t, follows a Gaussian random walk process with parameter σ. The evolution of γ_t can be described by Equation (2.6). Then, it follows that

$$\gamma_t \,|\, \gamma_{t-1} \sim \mathcal{N}(\gamma_{t-1}, \sigma^2).$$

3.1 Constraints in Parameters

We require the following two conditions of Proposition 3 to be true:

1. $Pr(m + \gamma_t > 1 \,|\, \gamma_0) > 0$ for all γ_0 and at least some t;
2. $\gamma^* > 0$.

We first show that the first condition is satisfied when γ_t follows a Gaussian random walk. We use the fact that $\gamma_t|\gamma_0 = \gamma_0 + \sum_{i=1}^{t} \varepsilon_t \sim \mathcal{N}(\gamma_0, \sigma^2 t)$:

$$\begin{aligned}
Pr(m + \gamma_t > 1 \,|\, \gamma_0) &= Pr(\gamma_t > 1 - m \,|\, \gamma_0) \\
&= Pr\left(\frac{\gamma_t - \gamma_0}{\sqrt{\sigma^2 t}} > \frac{1 - m - \gamma_0}{\sqrt{\sigma^2 t}} \,\middle|\, \gamma_0\right) \\
&= 1 - \Phi\left(\frac{1 - m - \gamma_0}{\sqrt{\sigma^2 t}}\right) \\
&> 0,
\end{aligned}$$

where $\Phi(\cdot)$ is the cumulative distribution function (cdf) of a standard normal random variable.

Next, we find out the constraints on the parameters of the Gaussian random walk that will make the second condition satisfied. We use the fact that $\gamma_t \,|\, \gamma_{t-1} \sim \mathcal{N}(\gamma_{t-1}, \sigma^2)$, and Equation (A.2) to write

$$pE = \int_{\gamma^*}^{\infty} x \cdot \pi(x \mid \gamma^*) dx = \int_{\gamma^*}^{\infty} x \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(x-\gamma^*)^2}{2\sigma^2}} dx.$$

Let $y = \frac{x-\gamma^*}{\sigma}$. Applying this change of variable yields the following:

$$\begin{aligned} pE &= \int_0^{\infty} (\gamma^* + \sigma y) \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{y^2}{2}} \sigma dy \\ &= \int_0^{\infty} \gamma^* \frac{1}{\sqrt{2\pi}} e^{-\frac{y^2}{2}} dy + \int_0^{\infty} \sigma y \frac{1}{\sqrt{2\pi}} e^{-\frac{y^2}{2}} dy \\ &= \frac{1}{2}\gamma^* + \sigma \int_0^{\infty} y\phi(y) dy \\ &= \frac{1}{2}\gamma^* + \sigma\phi(0) \\ &= \frac{1}{2}\gamma^* + \frac{1}{\sqrt{2\pi}}\sigma. \end{aligned}$$

Substituting this result into Equation (2.9), we have

$$1 - \frac{3}{2}\gamma^* - \frac{1}{\sqrt{2\pi}}\sigma - m = 0. \tag{3.1}$$

This yields an explicit solution for the cooperation threshold:

$$\gamma^* = \frac{2}{3}(1 - \frac{\sigma}{\sqrt{2\pi}} - m). \tag{3.2}$$

Finally, requiring $\gamma^* > 0$ then corresponds to the following constraint:

$$m + \frac{\sigma}{\sqrt{2\pi}} - 1 < 0. \tag{3.3}$$

The feasible region when γ_t follows a Gaussian random walk is the set of all combinations of σ and m that satisfy $m \in (0,1)$, $\sigma \in (0, \sqrt{2\pi})$ and Equation (3.3).

This corresponds to the grey area in Figure 3.1. Note that, although the σ's in the plot are feasible values, they are not all realistic. In fact, only small values of σ correspond to values of i_t that make practical sense. Figure 3.1 also shows some level sets of γ^*: these are values of m and σ that result in the same value of γ^*. Since γ^* is linear in both σ and m without any interaction terms, the level sets are simply straight lines.

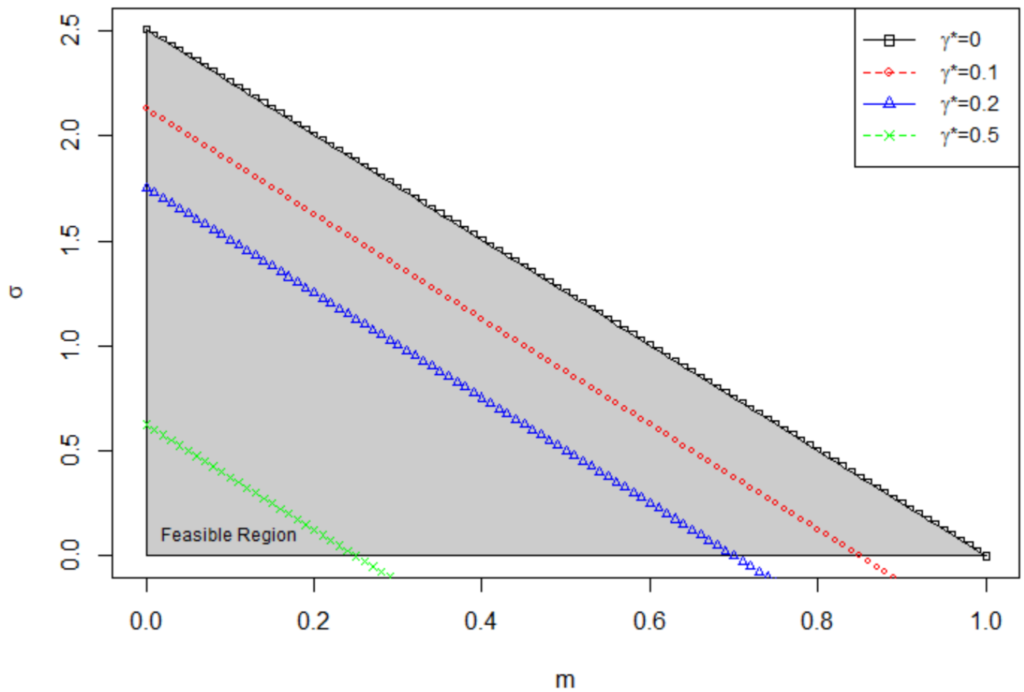

Figure 3.1: Feasible region and level curves when the cost of cooperation follows a Gaussian random walk

From Equation (3.2), we see that γ^* becomes smaller as m or σ increases. Intuitively, when m is large (i.e., the young is self-financing a large portion) he will be willing to pay less to receive a top-up to the \$1 retirement benefit. On the other hand, larger σ means more risk: the young agents are more likely to have to pay for the old, but might not get paid by the next generation. So as σ increases, the young agents are less willing to pay.[1] Therefore, γ^* decreases because agents want to be more cautious, and be committed to paying less often.

3.2 Probability of Non-cooperation

Under the assumption of a Gaussian random walk,

$$\gamma_t|\gamma_0 = \gamma_0 + \sum_{i=1}^{t} \varepsilon_t \sim \mathcal{N}(\gamma_0, \sigma^2 t). \tag{3.4}$$

[1]It is important to note that this sensitivity to volatility (which can be interpreted as a form of risk-aversion) arises from the pattern of payoffs, not from concavity in the utility function. The net payoff is a concave function of γ_t. Therefore, even though the utility function is linear, the utility of net payoffs is a concave function.

The probability of sanctioned non-cooperation at each time point t, in other words the t^{th} generation's probability of allowed cooperation breakdown, is equal to

$$
\begin{aligned}
&Pr(\gamma_t > \gamma^* |\ \gamma_0) \\
&= 1 - \Phi\left(\frac{\gamma^* - \gamma_0}{\sigma\sqrt{t}}\right) \\
&= 1 - \Phi\left(\frac{\frac{2}{3}(1 - \frac{\sigma}{\sqrt{2\pi}} - m) - \gamma_0}{\sigma\sqrt{t}}\right), \text{ where } \gamma_0 = 1 - m(1 + i_0).
\end{aligned}
\tag{3.5}
$$

If we set $i_0 = i_v$, then $\gamma_0 = 0$. As time goes by, the probability of allowed cooperation breakdown increases slowly for reasonable values of σ. For example, with $m = 0.21$ and $\sigma = 0.1$, we have $\gamma^* = 0.5$ and $Pr(\gamma_1 > \gamma^* | \gamma_0) \approx 0$ whereas $Pr(\gamma_5 > \gamma^* | \gamma_0) = 1.3\%$. So the 5^{th} generation still has a very small probability of cooperation breakdown.

σ												
2.6												
2.5	50%											
2.4	50%											
2.3	50%											
2.2	49%	50%										
2.1	49%	50%										
2	49%	49%	50%									
1.9	48%	49%	50%									
1.8	48%	49%	49%									
1.7	48%	48%	49%	50%								
1.6	47%	48%	49%	50%								
1.5	47%	48%	48%	49%	50%							
1.4	46%	47%	48%	49%	50%							
1.3	46%	47%	47%	48%	49%							
1.2	45%	46%	47%	48%	49%	50%						
1.1	44%	45%	46%	47%	48%	49%						
1	43%	44%	45%	46%	48%	49%	50%					
0.9	42%	43%	44%	46%	47%	48%	49%					
0.8	40%	41%	43%	44%	46%	47%	49%					
0.7	38%	40%	41%	43%	45%	46%	48%	50%				
0.6	35%	37%	39%	41%	43%	45%	47%	49%				
0.5	32%	34%	36%	38%	41%	43%	45%	48%	50%			
0.4	27%	29%	32%	34%	37%	40%	43%	46%	49%			
0.3	19%	22%	25%	28%	32%	35%	39%	43%	47%			
0.2	9%	11%	14%	18%	22%	27%	32%	37%	43%	49%		
0.1	0%	1%	1%	2%	5%	9%	14%	22%	32%	43%		
0												
	0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1	m

Figure 3.2: A heatmap for the probability of allowed non-cooperation for the 5^{th} generation when γ_t follows Gaussian random walk with $\gamma_0 = 0$.

Looking at Figure 3.2, for a given m, as σ becomes larger cooperation is less likely because it is more likely for the cost to jump above the threshold. Additionally, for a fixed σ, as m gets larger, the probability of allowed non-cooperation goes up because the threshold becomes lower as m increases.

From Figure 3.2, the most sustainable system (the one with the lowest probability of non-cooperation) appears to be the one with the smallest value of m, that is, the pay-as-you-go system (described at the beginning of Chapter 2); in this case, there is a higher cooperation threshold and only a small probability of non-cooperation. In addition, the total maximum possible cost $m+\gamma^*$ is also small. However, this type of arrangement creates another problem: if there is no prefunding, then the old will lose everything when the young choose not to cooperate. When there is some prefunding, say $m = 0.3$ and $\sigma = 0.3$, there is still about a 28% chance that the young will not cooperate in the next period, but in case of that non-cooperation, the old will not lose everything, as they still receive the prefunded portion of the benefit. When m is larger, retirees do not need as much help from the young generation, since they already self-financed a large portion of their benefit, so the amount at stake (i.e., what the old stand to lose in case of non-cooperation by the young) is smaller. We may be willing to accept a larger likelihood of that loss, because it is not a catastrophic loss. From the perspective of a government social planner, it may be better to prefund the pension plan, so the retirees' benefit at risk is not large.

To explore this idea further, we calculated the expected loss for the old at $t = 5$ when there is no cooperation. Let $\pi_t(y \,|\, x)$ be the t-step transition density from x at time 0 to y at time t.

$$
\begin{aligned}
&E[\text{benefit loss of the old at } t \mid \text{no cooperation at } t] \\
&= E(\gamma_t|\gamma_t > \gamma^*, \gamma_0) \\
&= \frac{\int_{\gamma^*}^{\infty} x \cdot \pi_t(x \,|\, \gamma_0)dx}{Pr(\gamma_t > \gamma^*|\ \gamma_0)} \\
&= \frac{\int_{\gamma^*}^{\infty} x \frac{1}{\sqrt{2\pi\sigma^2 t}} e^{-\frac{(x-\gamma_0)^2}{2\sigma^2 t}} dx}{Pr(\gamma_t > \gamma^*|\ \gamma_0)}, \text{ by (3.4).}
\end{aligned}
$$

Focusing on the numerator only, we let $y = \frac{x-\gamma_0}{\sigma\sqrt{t}}$. By changing the variable of integration from x to y, we obtain:

$$
\begin{aligned}
&\int_{\frac{\gamma^*-\gamma_0}{\sigma\sqrt{t}}}^{\infty} (\gamma_0 + \sigma\sqrt{t}y) \frac{1}{\sqrt{2\pi\sigma^2 t}} e^{-\frac{y^2}{2}} \sigma\sqrt{t} dy \\
&= \gamma_0 \int_{\frac{\gamma^*-\gamma_0}{\sigma\sqrt{t}}}^{\infty} \phi(y)dy + \sigma\sqrt{t} \int_{\frac{\gamma^*-\gamma_0}{\sigma\sqrt{t}}}^{\infty} y\phi(y)dy \\
&= \gamma_0 \left[1 - \Phi\left(\frac{\gamma^*-\gamma_0}{\sigma\sqrt{t}}\right)\right] + \sigma\sqrt{t}\ \phi\left(\frac{\gamma^*-\gamma_0}{\sigma\sqrt{t}}\right).
\end{aligned}
$$

Combining the numerator and denominator, we have

$$
\begin{aligned}
&E[\text{benefit loss of the old at } t \mid \text{no cooperation at } t] \\
&= \frac{\gamma_0 \left[1 - \Phi\left(\frac{\gamma^* - \gamma_0}{\sigma\sqrt{t}}\right)\right] + \sigma\sqrt{t}\, \phi\left(\frac{\gamma^* - \gamma_0}{\sigma\sqrt{t}}\right)}{Pr(\gamma_t > \gamma^* \mid \gamma_0)} \\
&= \frac{\gamma_0 \left[1 - \Phi\left(\frac{\gamma^* - \gamma_0}{\sigma\sqrt{t}}\right)\right] + \sigma\sqrt{t}\, \phi\left(\frac{\gamma^* - \gamma_0}{\sigma\sqrt{t}}\right)}{1 - \Phi\left(\frac{\gamma^* - \gamma_0}{\sigma\sqrt{t}}\right)}, \text{ by Equation (3.5).} \\
&= \gamma_0 + \sigma\sqrt{t} \left[\frac{\phi\left(\frac{\gamma^* - \gamma_0}{\sigma\sqrt{t}}\right)}{1 - \Phi\left(\frac{\gamma^* - \gamma_0}{\sigma\sqrt{t}}\right)}\right].
\end{aligned}
$$

The results are summarized in Figure 3.3: when m is small, the severity (expressed as the conditional expected benefit loss) is worse. The social planner then has to balance sustainability against severity of loss.

σ												
2.6												
2.5	4.46											
2.4	4.30											
2.3	4.14											
2.2	3.98	3.93										
2.1	3.82	3.77										
2	3.65	3.61	3.57									
1.9	3.49	3.45	3.41									
1.8	3.33	3.29	3.25									
1.7	3.17	3.13	3.08	3.04								
1.6	3.01	2.97	2.92	2.88								
1.5	2.85	2.81	2.76	2.72	2.68							
1.4	2.69	2.64	2.60	2.56	2.52							
1.3	2.53	2.48	2.44	2.40	2.35							
1.2	2.37	2.32	2.28	2.24	2.19	2.15						
1.1	2.21	2.16	2.12	2.07	2.03	1.99						
1	2.05	2.00	1.96	1.91	1.87	1.83	1.78					
0.9	1.89	1.84	1.80	1.75	1.71	1.67	1.62					
0.8	1.73	1.68	1.64	1.59	1.55	1.50	1.46					
0.7	1.57	1.52	1.48	1.43	1.39	1.34	1.30	1.26				
0.6	1.41	1.37	1.32	1.27	1.23	1.18	1.14	1.10				
0.5	1.26	1.21	1.16	1.11	1.07	1.02	0.98	0.94	0.89			
0.4	1.10	1.05	1.01	0.96	0.91	0.86	0.82	0.77	0.73			
0.3	0.96	0.90	0.85	0.80	0.75	0.71	0.66	0.61	0.57			
0.2	0.82	0.76	0.71	0.66	0.60	0.55	0.50	0.46	0.41	0.37		
0.1	0.71	0.64	0.58	0.52	0.47	0.41	0.35	0.30	0.25	0.20		
0												
	0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1	m

Figure 3.3: Conditional expected dollar loss for the old generation at $t = 5$ given no cooperation when γ_t follows a Gaussian random walk with $\gamma_0 = 0$.

3.3 Decision Process

In this section, we give a detailed example of the decision process outlined in Chapter 2. Suppose the cost of cooperation follows a Gaussian random walk process with parameters $m = 0.21$, $\gamma_0 = 0$ and $\sigma = 0.1$. By Equation (3.2), the maximum cooperation threshold based on the subgame perfect equilibrium requirement, γ^*, is 0.5. In Table 3.1, we show each element of the decision made by six consecutive generations facing six specific values of realized γ_t from the Gaussian random walk process described above. Cooperation is broken twice (at times 4 and 6) but all agents follow the correct behaviour.

	(1)	(2)	(3)	(4)	(5) = (3) - (4)	(6)
t	γ_t	Correct behaviour	Expected net payoff under the correct behaviour	Expected net payoff of deviating	Incentive to behave correctly	Actual behaviour
1	0.35	Cooperate	0.4	-	0.4	Cooperate (correct)
2	0.42	Cooperate	0.3	-	0.3	Cooperate (correct)
3	0.5	Cooperate	0.0	-	0.0	Cooperate (correct)
4	0.55	Do not cooperate	0.4	-0.2	0.6	Do not cooperate (correct)
5	0.48	Cooperate	0.1	-	0.1	Cooperate (correct)
6	0.6	Do not cooperate	0.3	-0.3	0.6	Do not cooperate (correct)

Table 3.1: Example of cooperation decisions using the threshold $\gamma^* = 0.5$, for a Gaussian random walk process with $\gamma_0 = 0,\ m = 0.21,\ \sigma = 0.1$.

By contrast, Table 3.2 describes the decision process when the cooperation threshold is set at a level above γ^*, i.e., at $\widetilde{\gamma} = 0.55$. In Table 3.2, at time 4, the higher threshold gives the young agent incentive to deviate from the correct behaviour, because following the correct behaviour would lead to a negative expected net payoff, which is worse than deviating. So the 4$^{\text{th}}$ generation will not follow the rule. As a consequence, the 5$^{\text{th}}$ generation will also not cooperate, because he needs to punish his predecessor who deviates from the rule. So looking at the 6 periods overall, a threshold that is set to be higher than γ^* leads to less cooperation over time. This is consistent with the claim that the most cooperative solution is achieved when the threshold is set based on Equation (2.9).

	(1)	(2)	(3)	(4)	(5) = (3) - (4)	(6)
t	γ_t	Suggested behaviour using the threshold $\widetilde{\gamma}$ instead of γ^*	Expected net payoff under the correct behaviour	Expected net payoff of deviating	Incentive to behave correctly	Actual behaviour
1	0.35	Cooperate	0.4	-	0.4	Cooperate (correct)
2	0.42	Cooperate	0.3	-	0.3	Cooperate (correct)
3	0.5	Cooperate	0.1	-	0.1	Cooperate (correct)
4	0.55	Cooperate	-0.1	-	-0.1	Do not cooperate (deviate)
5	0.48	Cooperate	0.2	-	0.2	Do not cooperate (correct, punishment)
6	0.6	Do not cooperate	0.3	-0.3	0.6	Do not cooperate (correct)

Table 3.2: Example of cooperation decisions using the threshold $\widetilde{\gamma} = 0.55 > \gamma^*$, for the same Gaussian random walk process with $\gamma_0 = 0,\ m = 0.21,\ \sigma = 0.1$.

Chapter 4

Application to AR(1) with $\mu = 0$

In terms of the actual setup of a pension plan, the purpose of the fixed contribution m is to prefund the \$1 benefit, whereas γ_t is the intergenerational cost. In an AR(1) process, $\mu = 0$ means that the long-term mean of the process is 0. If the cost of cooperation follows an AR(1) process with zero mean, then the pension plan's goal must be full prefunding, since in the long term, the expected size of intergenerational transfers is zero. By contrast, $\mu > 0$ means that the pension plan is only partially prefunded, and we expect the young generation to put up the rest of the cost. Finally, $\mu < 0$ means each generation is expected to overfund itself which the next generation can expect to benefit from. In this chapter, we focus on $\mu = 0$, because most TBPs are expected to operate with full prefunding.

Assume that the cost of cooperation, γ_t, follows an AR(1) process with parameters $\mu = 0,\ 0 < \alpha < 1$, and $\sigma > 0$, so the evolution of γ_t can be described by Equation (2.4) and

$$\gamma_t \,|\, \gamma_{t-1} \sim \mathcal{N}(\alpha\gamma_{t-1}, \sigma^2).$$

Let us define the predicted value of γ_t given γ_{t-1} as

$$\hat{\gamma}_t = \alpha\gamma_{t-1}.$$

The realized value, γ_t, is this predicted value, plus a random term:

$$\gamma_t = \hat{\gamma}_t + \varepsilon_t, \text{ where } \varepsilon_t \sim \mathcal{N}(0, \sigma^2).$$

Here, α is the speed of return to the long-term mean of $\mu = 0$. The cost of cooperation today (at time t) is equal to a proportion of the cost last time (at time $t-1$) plus some randomness. Small α promotes fast reversion to the mean $\mu = 0$. For example, $\alpha = 0.2$ means that the predicted cost today is only 20% of the realized cost in the last period. Large α means slow reversion to the mean, so the predicted cost today is almost equal to the realized cost in the last period.

4.1 Constraints in Parameters

As before, we require the two conditions of Proposition 3 to be true:

1. $Pr(m+\gamma_t > 1|\ \gamma_0) > 0$ for all γ_0 and at least some t;
2. $\gamma^* > 0$.

We first show that the first condition is satisfied when γ_t follows an AR(1) process with $\mu = 0$.

Proof. Suppose the process starts with some value γ_0. Then, for $t = 1, 2, ...$, we have that

$$\begin{aligned}
\gamma_t &= \alpha\gamma_{t-1} + \varepsilon_t \\
&= \alpha\left(\alpha\gamma_{t-2} + \varepsilon_{t-1}\right) + \varepsilon_t \\
&= ... \\
&= \alpha^t\gamma_0 + \sum_{i=1}^{t} \varepsilon_i \alpha^{t-i}.
\end{aligned} \tag{4.1}$$

Therefore, the conditional distribution of γ_t given γ_0 is normal with $E(\gamma_t|\gamma_0) = \alpha^t\gamma_0$ and $Var(\gamma_t|\gamma_0) = \frac{(1-\alpha^{2t})\sigma^2}{1-\alpha^2}$. Consequently,

$$\begin{aligned}
Pr(m+\gamma_t > 1|\ \gamma_0) &= Pr(\gamma_t > 1-m|\ \gamma_0) \\
&= 1 - \Phi\left(\frac{1-m-E(\gamma_t|\gamma_0)}{\sqrt{Var(\gamma_t|\gamma_0)}}\right) \\
&> 0.
\end{aligned} \tag{4.2}$$

□

Next, we find out constraints on the parameters of the AR(1) process when $\mu = 0$, so that the second condition is satisfied. By Equation (A.2),

$$\begin{aligned}
pE &= \int_{\gamma^*}^{\infty} x \cdot \pi(x\,|\,\gamma^*)dx \\
&= \int_{\gamma^*}^{\infty} x \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(x-\alpha\gamma^*)^2}{2\sigma^2}} dx.
\end{aligned}$$

Let $y = \frac{x-\alpha\gamma^*}{\sigma}$. Using a change of variable in the previous integral (i.e., from x to y), we have that

$$
\begin{aligned}
pE &= \int_{\frac{\gamma^* - \alpha\gamma^*}{\sigma}}^{\infty} (\alpha\gamma^* + \sigma y)\frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{y^2}{2}} \sigma dy \\
&= \int_{\frac{(1-\alpha)\gamma^*}{\sigma}}^{\infty} (\alpha\gamma^*)\frac{1}{\sqrt{2\pi}} e^{-\frac{y^2}{2}} dy + \int_{\frac{(1-\alpha)\gamma^*}{\sigma}}^{\infty} \sigma y \frac{1}{\sqrt{2\pi}} e^{-\frac{y^2}{2}} dy \\
&= \alpha\gamma^* \left[1 - \Phi\left(\frac{(1-\alpha)\gamma^*}{\sigma}\right)\right] + \sigma \int_{\frac{(1-\alpha)\gamma^*}{\sigma}}^{\infty} y\phi(y) dy \\
&= \alpha\gamma^* \left[1 - \Phi\left(\frac{(1-\alpha)\gamma^*}{\sigma}\right)\right] + \sigma\phi\left(\frac{(1-\alpha)\gamma^*}{\sigma}\right).
\end{aligned}
\tag{4.3}
$$

Substituting this result into Equation (2.9), we have

$$
1 - \alpha\gamma^* \left[1 - \Phi\left(\frac{(1-\alpha)\gamma^*}{\sigma}\right)\right] - \sigma\phi\left(\frac{(1-\alpha)\gamma^*}{\sigma}\right) - m - \gamma^* = 0. \tag{4.4}
$$

There is no explicit solution for γ^*, but we can study the behaviour of γ^* in terms of the parameters m, α and σ. The feasible region when γ_t follows an AR(1) process with $\mu = 0$ is the set of combinations of m, α and σ, which satisfies $\gamma^* > 0$.

4.2 The Behaviour of γ^*

Let $z = f(\gamma, \alpha, \sigma) = \frac{(1-\alpha)\gamma}{\sigma}$, so

$$
\frac{\partial z}{\partial \gamma} = \frac{(1-\alpha)}{\sigma}, \quad \frac{\partial z}{\partial \alpha} = \frac{-\gamma}{\sigma}, \quad \frac{\partial z}{\partial \sigma} = \frac{-z}{\sigma}.
$$

The cooperation threshold, γ^*, is the largest root of

$$
g(\gamma, \alpha, \sigma, m) = 1 - \alpha\gamma\left[1 - \Phi(z)\right] - \sigma\phi(z) - m - \gamma,
$$

where $\phi(s) = \frac{1}{\sqrt{2\pi}} e^{-\frac{s^2}{2}}$, and $\Phi(s) = \int_{-\infty}^{s} \phi(u) du$.

The function g is continuous. We know that $\gamma^* > 0$ so we are only interested in the behaviour of g on the range $\gamma > 0$, which is equivalent to $z > 0$. In Appendix B, we show that the function $g(\gamma, \alpha, \sigma, m)$ is decreasing in γ for $\gamma > 0$. We also show that the partial derivatives of g with respect to α, m and σ are negative whenever $\gamma > 0$. This means that increases in m, α and σ shift the function $g(\gamma, \alpha, \sigma, m)$ downwards, at least when $\gamma > 0$. Since the function $g(\gamma, \alpha, \sigma, m)$ is decreasing in γ, the largest root γ^* decreases in m, α and σ. Hence, the threshold decreases as the prefunding level and volatility increases, and as the speed of reversion becomes slower.

4.3 Probability of Non-cooperation

When γ_t follows an AR(1) process with $\mu = 0$, the conditional distribution of γ_t given γ_0 is normal with $E(\gamma_t|\gamma_0) = \alpha^t\gamma_0$ and $Var(\gamma_t|\gamma_0) = \frac{(1-\alpha^{2t})\sigma^2}{1-\alpha^2}$.

The probability of sanctioned non-cooperation at each time point t, or the t^{th} generation's probability of allowed cooperation breakdown, is equal to

$$\begin{aligned}
&Pr(\gamma_t > \gamma^*|\gamma_0) \\
&= 1 - \Phi\left(\frac{\gamma^* - E(\gamma_t|\gamma_0)}{\sqrt{Var(\gamma_t|\gamma_0)}}\right) \\
&= 1 - \Phi\left(\frac{\gamma^* - \alpha^t\gamma_0}{\sqrt{\frac{(1-\alpha^{2t})\sigma^2}{1-\alpha^2}}}\right),
\end{aligned} \tag{4.5}$$

where $\gamma_0 = 1 - m(1 + i_0)$. If we set $i_0 = i_v$, then $\gamma_0 = 0$.

For example, with $m = 0.21$, $\alpha = 0.9$ and $\sigma = 0.1$, we have that $\gamma^* = 0.61$, $Pr(\gamma_1 > \gamma^*|\gamma_0) \approx 0$ and $Pr(\gamma_5 > \gamma^*|\gamma_0) = 0.051\%$, which is lower than in the Gaussian random walk case.

4.4 Interactions among Parameters

To study how the level of prefunding and the process parameters α and σ jointly affect γ^*, we construct three-dimensional plots. First, for fixed α, the effect of m and σ on γ^* is studied.

Looking at Figure 4.1, when m is close to 1 (almost full prefunding), we cannot find a subgame perfect equilibrium with $\gamma^* > 0$. By contrast, when m is close to zero (almost pay-as-you-go) and σ is very small (very little uncertainty), γ^* is close to 1. As m or σ increase individually from zero, the threshold usually decreases as shown in the blue lower region, similar to the Gaussian random walk case.

However, the flat part of the surface in the upper region suggests an interaction between m and σ. When σ is relatively small, increases in σ do not affect γ^* much at first. But, further increases in σ result in decreasing γ^*. So, at first, the threshold is insensitive to the increasing risk, but after a certain point the threshold becomes sensitive. The size of the insensitive region (the flat part) decreases as m increases.

The decrease in γ^* means that agents are worried about the possibility that as they become old in the next period, γ_{t+1} will be over the threshold and they will not get paid by the next generation; therefore, they will lower the threshold today in anticipation. The argument is that if the next generation is quite likely to go beyond the threshold, then the current young generation is not willing to obey that same threshold, thus γ^* drops. How worried agents are about the possibility of the next generation not honouring the deal is a combination of α and the variability in the random term.

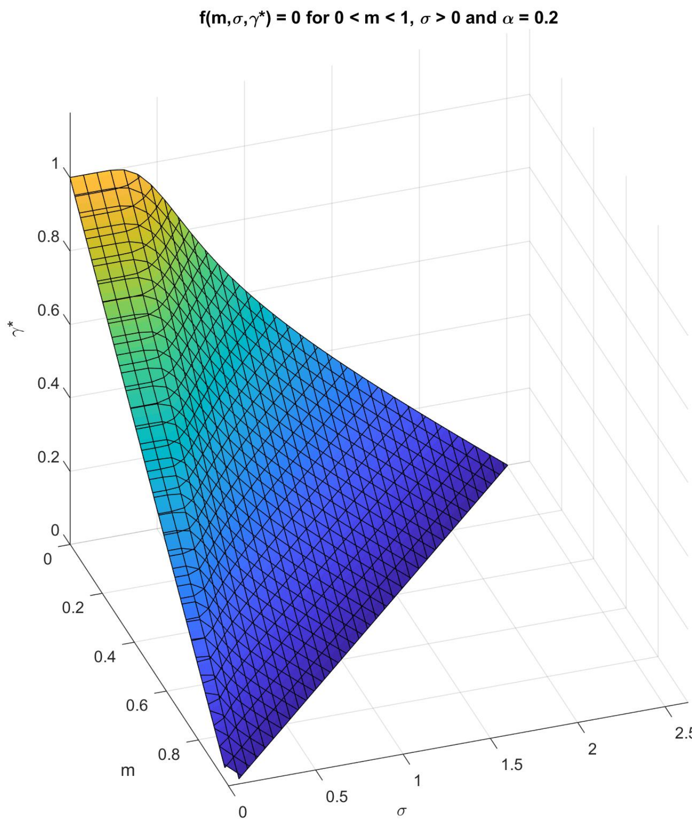

Figure 4.1: A three-dimensional plot of Equation (4.4), where $\alpha = 0.2$ for $0 < m < 1$, and $\sigma > 0$.

For a fixed α at 0.2, the small drift parameter suggests fast mean-reversion, which implies that the realized cost of cooperation in the last period, γ_{t-1}, does not have much influence on the predicted cost today. Instead, what mostly influences the actual cost today is σ. In this case, when σ is very small, it does not have much influence on γ^*, the capacity to cooperate. In fact, agents do not have much to be afraid of until σ becomes significant compared with $1 - m$ (i.e., the upper limit for γ^*). The point at which agents start to be concerned about volatility depends on the level of prefunding, m. The larger m is, the smaller the region that agents are insensitive to the increasing volatility. This is because as m increases σ is larger compared with $1 - m$.

Next, we look at the case with $\alpha = 0.8$, shown in Figure 4.2.

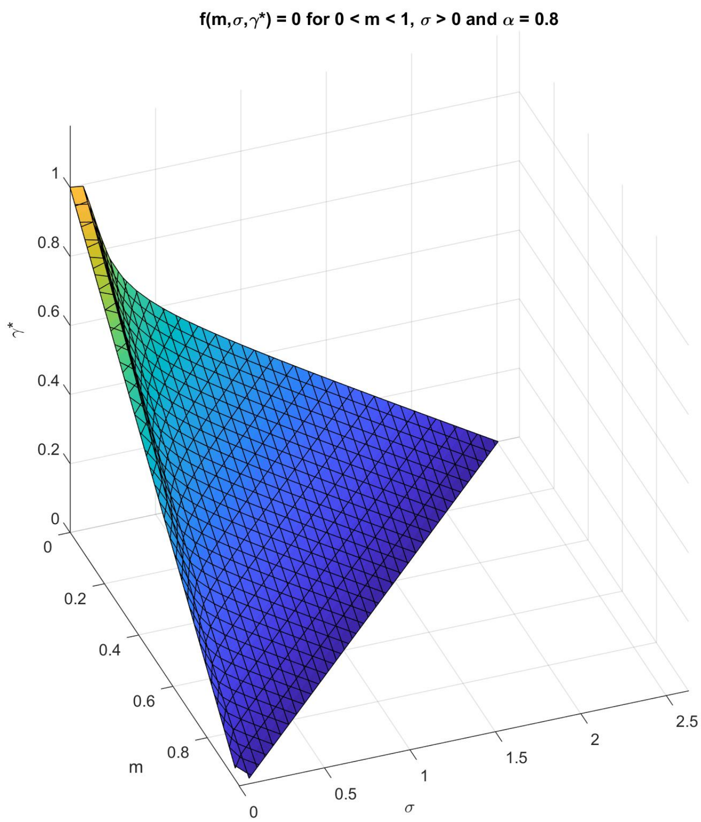

Figure 4.2: A three-dimensional plot of Equation (4.4), where $\alpha = 0.8$ for $0 < m < 1$, and $\sigma > 0$.

With large α, the predicted cost $\hat{\gamma}_t$ will be close to γ_{t-1} and might be high. Then, the cost is more likely to go over the threshold in the next period, so agents will be concerned about the increasing randomness much earlier. As illustrated in Figure 4.2, the plot looks similar to the Gaussian random walk case: the flat region where agents are insensitive to increasing risk is quite thin. To summarize, agents become concerned about the increasing σ much earlier when α is large (slow mean-reversion) and much later when α is small (fast mean-reversion).

Next, we study the effect of α and σ on γ^* for a fixed value of m.

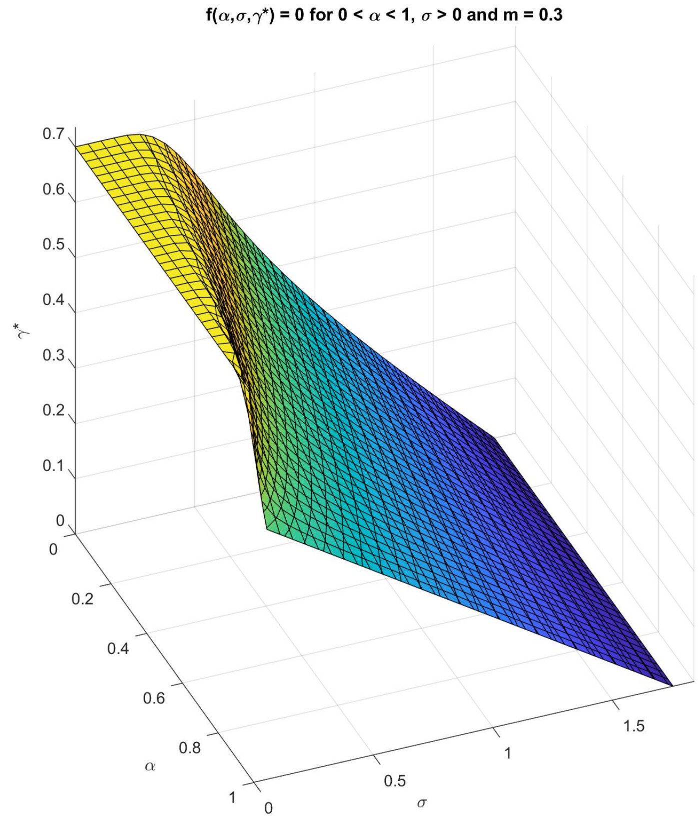

Figure 4.3: A three-dimensional plot of Equation (4.4), where $m = 0.3$ for $0 < \alpha < 1$, and $\sigma > 0$.

When $m = 0.3$, the upper bound for γ^* is 0.7. Figure 4.3 also illustrates the previously discussed insensitive region. Another thing to notice is the rate of change in γ^*. In Figure 4.3, when α is large, agents are very concerned about increasing volatility at first: there is a sharp drop in γ^* for small values of σ. Here, an increase in σ makes a large difference in how worried agents are about going beyond the threshold in the next time period and this affects agents' willingness to cooperate now. Once volatility reaches a certain point, agents appear to be less concerned about further increases in volatility: the rate of change in γ^* is slower. The opposite can be observed for small values of α where, as σ increases, the rate of change in γ^* is slow when σ is small but then becomes fast for sufficiently large values of σ.

To take a closer look, we redraw the surface from Figure 4.3 as a two-dimensional contour plot with varying α and σ.

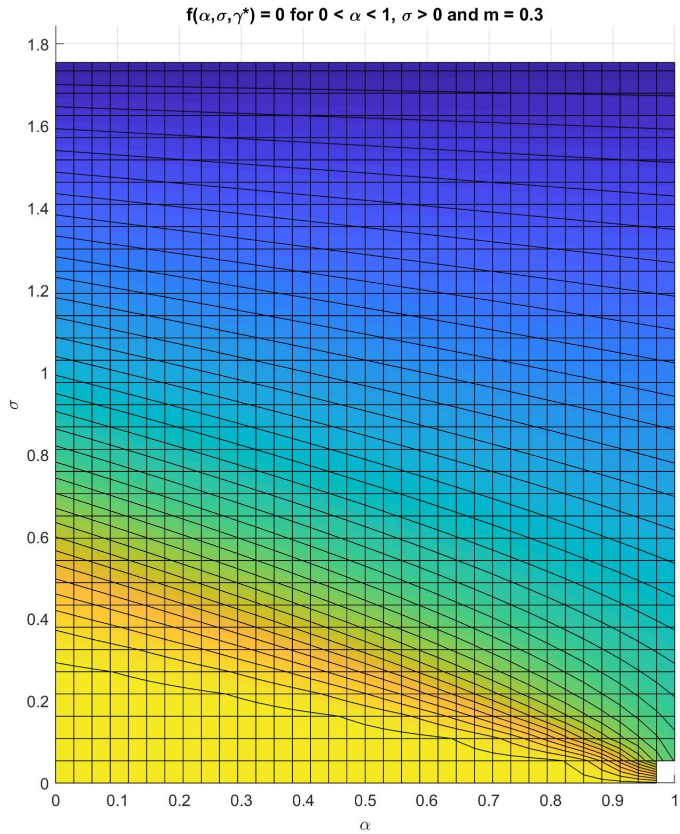

Figure 4.4: A contour plot of γ^* where $m = 0.3$ for $0 < \alpha < 1$, and $\sigma > 0$.

For large values of α, the contour lines are very dense in the small σ range, corresponding to a steep decline in γ^*. The lines are more sparse in the large σ range, corresponding to a slower decline in γ^*. In the small σ range (for example, when σ goes from 0.1 to 0.2), agents are not very concerned about increasing volatility if α is small. When $\alpha = 0.1$ and $\sigma = 0.1$, γ^* is quite large at 0.7. Recall that the realized value of γ_t is driven by the speed of mean reversion and that σ controls the volatility of the random innovation term. With small α, $\hat{\gamma}_t$ is pulled back to 0 quickly. In this case, adding an error term with a small σ does not matter because it is unlikely to bring γ_t near the threshold. Changing σ from 0.1 to 0.2 does not change this favourable situation much, and agents are able to sustain similar thresholds in both cases. As a result, the rate of change in γ^* is small.

When $\alpha = 0.8$ and $\sigma = 0.1$, γ^* is smaller at 0.63. The predicted cost of cooperation $\hat{\gamma}_t$ is close to γ_{t-1}, which may be much higher than the long-term mean $\mu = 0$. Then, there is a higher probability that γ_t will end up above the threshold and the σ that agents can tolerate is much smaller compared to that in the small α case. Increasing σ might make a large difference in terms of γ_t going beyond the threshold or not. As a consequence, γ^* decreases faster in the small range of σ, when α is large.

In the large σ range (for example, when σ goes from 0.8 to 0.9), agents are more concerned about increasing volatility when α is small. When $\alpha = 0.1$ and $\sigma = 0.8$, γ^* is quite small at 0.4. Changing σ from 0.8 to 0.9 will make things noticeably worse in terms of the possibility of γ_t jumping over the already low threshold. As a result, γ^* will decrease substantially. When $\alpha = 0.8$ and $\sigma = 0.8$, γ^* is even smaller at 0.28. This combination of large σ and large α is already disastrous, i.e., there is already hardly any trust because the probability of γ_t exceeding the threshold is so high. An even larger σ does not change the situation by much: the threshold is already so low that it cannot go much lower. So the rate of change in γ^* is not as fast.

Normally, a pension plan would know what the prefunding level m is, but would not know the process parameters. We produce contour plots with varying α and σ for $m = 0.3$, 0.5, and 0.8, which represent very unfunded, median funded and mostly funded plans, respectively.

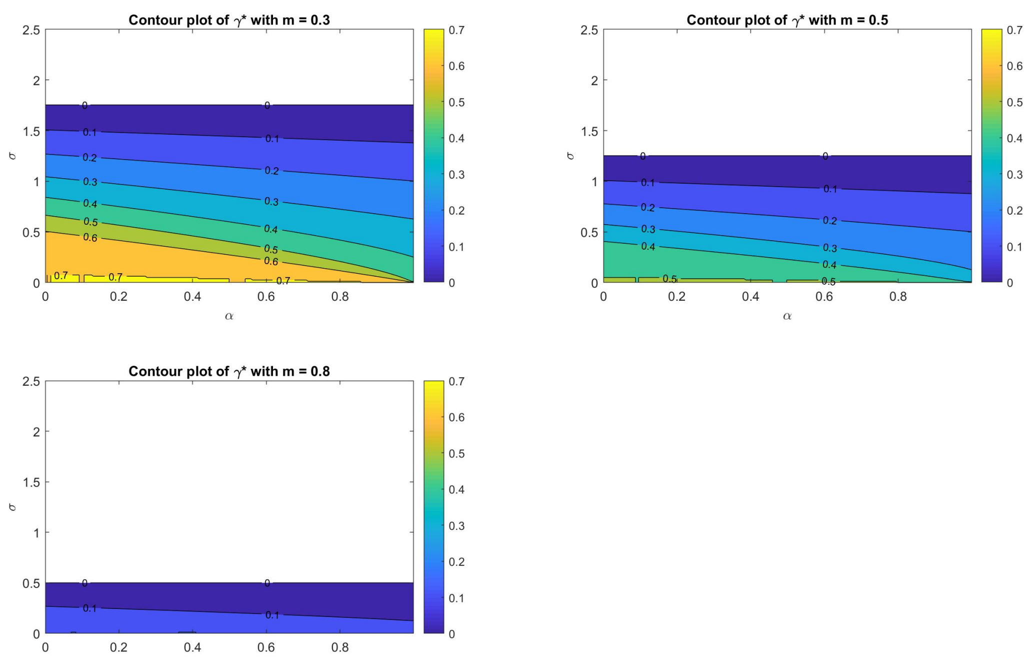

Figure 4.5: Contour plots of γ^* where $m = 0.3$, 0.5, and 0.8 for $0 < \alpha < 1$, and $\sigma > 0$.

In Figure 4.5, as we vary the level of prefunding the general pattern for the rate of change in γ^* still persists; i.e., for a large value of α, γ^* decreases faster at first and slower later, as σ increases, and vice versa for a small value of α. Moreover, as m increases, the largest value that γ^* can take is smaller.

Chapter 5

Conclusions

In this project, a game theoretic framework is employed to study intergenerational solidarity and sustainability in target benefit plans. We start from the model in Messner and Polborn (2003), where overlapping generations are combined with uncertain payoffs and cooperation is studied using non-cooperative game theory. We extend their original model by adding a fixed contribution paid by the young generation and a slightly different benefit received by the old generation. Then overall, the young generation's contribution to the plan consists of two parts, the fixed part that is mandatory for all the young generations and that grows with investment returns, and a variable cost of cooperation that is paid only when choosing to cooperate. Accordingly, the old generation's benefit, which is targeted at \$1, is composed of a self-financed portion that varies with investment returns and a subsidy from the next generation contingent on the cooperation decision.

Through the definition of subgame perfect equilibrium, the threshold for the cost of cooperation, which leads to the most cooperative outcome, can be calculated. The Gaussian random walk and the AR(1) process centred around zero are assumed for the stochastic cost of cooperation. We study the feasible region for the parameters, and how these parameters affect the threshold.

When the cost of cooperation follows a Gaussian random walk process, the threshold decreases linearly as either the level of prefunding or the volatility of the cost of cooperation increases. At first glance, a pension plan that is close to the pay-as-you-go system seems to be more sustainable than a prefunded pension plan in terms of the probability of cooperation breakdown, but we find that this small probability event corresponds to much more severe losses. When the cost of cooperation follows an AR(1) process, the parameters have much more complex influences on the threshold than that described in Messner and Polborn (2003).

Our model is a very basic framework that allows us to study cooperation in TBPs, and it has some significant limitations that could be addressed in future work:

- The current model assumes risk neutrality. Risk aversion can be added by incorporating a risk averse utility function and redefining the subgame perfect equilibrium.

Specifically, let $u(\cdot)$ be a concave utility function. Then, the expression in Equation (2.7) would be modified as follows:

$$E[u(\text{net payoff of a threshold cooperator who follows the correct behaviour strategy at } t-1)]$$

Similarly, the expression in Equation (2.8) would become:

$$E[u(\text{net payoff of a threshold cooperator who deviates from the correct behaviour strategy at } t-1)]$$

Subgame perfect equilibrium can be defined by setting the above two expressions to be equal. With this modification, we expect that γ^* would be even more sensitive to σ and cooperation will be even less likely, because the utility of the net payoff will be even more concave than in the case with risk neutrality.

- In Chapter 2, the asymmetry between the upside potential and downside risk borne by the young generation was mentioned: the cost of cooperation is capped from above but not from below. With regard to this limitation, it may be of interest to introduce a lower threshold, γ'. When γ_t is negative, the retirees might pass on the investment gains to the young generation only when $0 > \gamma_t \geq \gamma'$. This modification would mean that, at each time point, both the young and the old generation could make a choice about cooperation: the young may be willing to subsidize the retirees only when the cost is affordable; and the retirees may be willing to give up the upside potential only up to a limit. Hence, the risk transfers among generations, no matter which way they go, will happen only when they are acceptable to all, ex-post.
- Our model only allows for either full cooperation or no cooperation, and the retirees will receive their full benefit only if cooperation is at 100%. Messner and Polborn (2003) describe an extension of their basic model, which allows partial cooperation by defining a cooperation function, through which the relationship between retirees' benefits and the level of cooperation is specified. This feature could be added to our model.
- We assume two-period living agents where at each time point there are only two agents: the young and the old. It would be interesting to explore multiple-period living agents, as in Cremer (1986).
- The investment return that is forfeited by the generation that deviates from the correct behaviour is not given to anyone in the current model. As a further improvement, the forfeited amount could be used to benefit the current old, or could be given to the next young generation as a reward for behaving correctly. Forfeiting the entire investment return is sufficient to make the subgame perfect equilibrium exist. However, smaller

penalties might also be sufficient. It would be interesting to explore the existence of subgame perfect equilibrium under smaller penalties.

- Other stochastic processes that satisfy Condition 1 can be investigated. Alternatively, a less strict condition could be explored that would still guarantee the existence of a subgame perfect equilibrium.
- Our model assumes no discounting, which makes cooperation more likely to happen but might not be very realistic. It would be interesting to explore whether subgame perfect equilibrium can exist when there is discounting.
- It would also be interesting to consider γ_t, m and γ^* in the context of option pricing theory. We would then compare our results to Cui et al. (2011).

Bibliography

Abreu, D. (1988). On the theory of infinitely repeated games with discounting. *Econometrica*, *56*(2), 383–396.

Bhaskar, V. (1998). Informational constraints and the overlapping generations model: Folk and anti-folk theorems. *Review of Economic Studies*, *65*(1), 135–149.

Bilodeau, C. (1998). The ownership of the pension plan surplus using cooperative game theory. *Actuarial Research Clearing House*, *1*, 173–187.

Chatain, O. (2016). Cooperative and non-cooperative game theory. In M. Augier & D. J. Teece (Eds.), *The palgrave encyclopedia of strategic management* (pp. 1–3). London: Palgrave Macmillan UK.

Cremer, J. (1986). Cooperation in ongoing organizations. *Quarterly Journal of Economics*, *101*(1), 33–49.

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

Esteban, J. M. & Sákovics, J. (1993). Intertemporal transfer institutions. *Journal of Economic Theory*, *61*(2), 189–205.

Fudenberg, D. & Tirole, J. (1991). *Game theory* (1st ed.). The MIT Press.

Gollier, C. (2008). Intergenerational risk-sharing and risk-taking of a pension fund. *Journal of Public Economics*, *92*(5), 1463–1485.

Hammond, P. (1975). Charity: Altruism or cooperative egoism? In E. S. Phelps (Ed.), *Altruism, morality and economic theory* (pp. 115–131). New York: Russell Sage Foundation.

Kandori, M. (1992). Repeated games played by overlapping generations of players. *Review of Economic Studies*, *59*(1), 81–92.

Messner, M. & Polborn, M. K. (2003). Cooperation in stochastic OLG games. *Journal of Economic Theory*, *108*(1), 152–168.

Miyazaki, K. (2014). Efficiency and lack of commitment in an overlapping generations model with endowment shocks. *Japanese Economic Review*, *65*(4), 499–520.

Morgan, G. (2018, February 16). Why are such supposedly compassionate Liberals letting Sears rip off its workers? *Financial Post*. Retrieved from https://business.financialpost.com/opinion/gwyn-morgan-why-are-such-supposedly-compassionate-liberals-letting-sears-rip-off-its-workers

Walsh, M. W. (2005, July 31). How Wall Street wrecked United's pension. *The New York Times*. Retrieved from https://www.nytimes.com/2005/07/31/business/yourmoney/how-wall-street-wrecked-uniteds-pension.html

Appendix A

Implications of Condition 1

This section shows that Condition 1 implies $pE \geq p'E'$ and the stochastic monotone Markov property.

$$\begin{aligned} E &= E[\gamma_t \,|\, \gamma_t > \gamma^*,\ \gamma_{t-1} = \gamma^*] \\ &= \frac{\int_{\gamma^*}^{\infty} x \cdot \pi(x \,|\, \gamma^*)dx}{\int_{\gamma^*}^{\infty} \pi(x \,|\, \gamma^*)dx} \qquad \text{(A.1)} \\ &= \frac{\int_{\gamma^*}^{\infty} x \cdot \pi(x \,|\, \gamma^*)dx}{p} \\ \implies pE &= \int_{\gamma^*}^{\infty} x \cdot \pi(x \,|\, \gamma^*)dx \qquad \text{(A.2)} \end{aligned}$$

Similarly, if an agent at $t-1$ is below the threshold such that $\gamma_{t-1} = \underline{\gamma} < \gamma^*$,

$$\begin{aligned} E' &= E[\gamma_t |\ \gamma_t > \gamma^*,\ \gamma_{t-1} = \underline{\gamma}] \\ &= \frac{\int_{\gamma^*}^{\infty} x \cdot \pi(x \,|\, \underline{\gamma})dx}{\int_{\gamma^*}^{\infty} \pi(x \,|\, \underline{\gamma})dx} \qquad \text{(A.3)} \\ &= \frac{\int_{\gamma^*}^{\infty} x \cdot \pi(x \,|\, \underline{\gamma})dx}{p'} \\ \implies p'E' &= \int_{\gamma^*}^{\infty} x \cdot \pi(x \,|\, \underline{\gamma})dx \qquad \text{(A.4)} \end{aligned}$$

If Condition 1 is satisfied, i.e., $\pi(\overline{\gamma} \,|\, \widetilde{\gamma}) \geq \pi(\overline{\gamma} \,|\, \underline{\gamma}),\ \forall\ \overline{\gamma} > \widetilde{\gamma} > \underline{\gamma}$, then

$$\int_{\gamma^*}^{\infty} x \cdot \pi(x \,|\, \gamma^*)dx \geq \int_{\gamma^*}^{\infty} x \cdot \pi(x \,|\, \underline{\gamma})dx,$$

and Equation (2.3), $pE \geq p'E'$, also holds.

Furthermore, Condition 1 implies the stochastic monotone Markov property in Messner and Polborn (2003), because

$$Pr(\gamma_t > \overline{\gamma}|\gamma_{t-1} = \overline{\gamma}) = \int_{\overline{\gamma}}^{\infty} \pi(x\,|\,\overline{\gamma})dx \geq \int_{\overline{\gamma}}^{\infty} \pi(x\,|\,\underline{\gamma})dx = Pr(\gamma_t > \overline{\gamma}|\gamma_{t-1} = \underline{\gamma}),\ \forall\ \overline{\gamma} > \underline{\gamma}.$$

Appendix B

Proofs for the Behaviour of γ^*

This section first shows that $\frac{\partial g}{\partial \gamma} = -\alpha\left[1 - \Phi(z)\right] + z\phi(z) - 1$ is negative for all $z > 0$.

Proof. The first term $-\alpha\left[1 - \Phi(z)\right]$ is between -1 and 0, since $\Phi(z) \in (0, 1)$ for $z > 0$, and $\alpha \in (0, 1)$.

The second term $z\phi(z)$ is between 0 and $\frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}} \approx 0.242$. On the one hand, as we require $z > 0$, $z\phi(z) > 0$. On the other hand,

$$\begin{aligned}
\frac{d}{dz}\Big[z\phi(z)\Big] &= \phi(z) + z\frac{d}{dz}\phi(z) \\
&= \phi(z) - z^2\phi(z) \\
&= (1 - z^2)\phi(z).
\end{aligned}$$

We can therefore see that $\frac{d}{dz}\Big[z\phi(z)\Big]$ is positive for $-1 < z < 1$, and negative otherwise. Combining with $z > 0$, the second term, $z\phi(z)$, increases when $0 < z < 1$, and decreases when $z > 1$. The maximum value of $z\phi(z)$ on the domain $z > 0$ will then occur at $z = 1$. The maximum value is $1 \cdot \phi(1) = \frac{1}{\sqrt{2\pi}} e^{-\frac{1^2}{2}} \approx 0.242$.

Combining all three terms, $-2 < \frac{\partial g}{\partial \gamma} < -0.758$. Therefore, $\frac{\partial g}{\partial \gamma}$ is negative for all $z > 0$.

$\square$

We also show that the partial derivatives of g with respect to α, m and σ are negative whenever $\gamma > 0$. First, we look at how the function $g(\gamma, \alpha, \sigma, m)$ changes with respect to γ.

Since $\phi'(s) = -s\phi(s)$, and $\Phi'(s) = \phi(s)$, we have that

$$\begin{aligned}
\frac{\partial g}{\partial \gamma} &= -(\alpha\gamma)\left[-\frac{\partial}{\partial \gamma}\Phi(z)\right] - \alpha\left[1 - \Phi(z)\right] - \sigma\frac{\partial}{\partial \gamma}\phi(z) - 1 \\
&= (\alpha\gamma)\phi(z)\frac{\partial z}{\partial \gamma} - \alpha\left[1 - \Phi(z)\right] + \sigma z\phi(z)\frac{\partial z}{\partial \gamma} - 1 \\
&= -\alpha\left[1 - \Phi(z)\right] + (\alpha\gamma + \sigma z)\phi(z)\frac{\partial z}{\partial \gamma} - 1 \\
&= -\alpha\left[1 - \Phi(z)\right] + \gamma\left(\frac{1-\alpha}{\sigma}\right)\phi(z) - 1 \\
&= -\alpha\left[1 - \Phi(z)\right] + z\phi(z) - 1.
\end{aligned} \tag{B.1}$$

This derivative is always negative for $z > 0$. In terms of α, we have

$$\begin{aligned}
\frac{\partial g}{\partial \alpha} &= (\alpha\gamma)\frac{\partial}{\partial \alpha}\Phi(z) - \gamma\left[1 - \Phi(z)\right] - \sigma\frac{\partial}{\partial \alpha}\phi(z) \\
&= (\alpha\gamma)\phi(z)\frac{\partial z}{\partial \alpha} - \gamma\left[1 - \Phi(z)\right] + \sigma z\phi(z)\frac{\partial z}{\partial \alpha} \\
&= -\gamma\left[1 - \Phi(z)\right] + (\alpha\gamma + \sigma z)\phi(z)\frac{\partial z}{\partial \alpha} \\
&= \gamma\left\{ -\left[1 - \Phi\left(\frac{(1-\alpha)\gamma}{\sigma}\right)\right] - \left(\frac{\gamma}{\sigma}\right)\phi\left(\frac{(1-\alpha)\gamma}{\sigma}\right)\right\},
\end{aligned} \tag{B.2}$$

which is negative for any $\gamma > 0$. We also have

$$\frac{\partial g}{\partial m} = -1 < 0.$$

Finally, with respect to σ, we have

$$\begin{aligned}
\frac{\partial g}{\partial \sigma} &= (\alpha\gamma)\frac{\partial}{\partial \sigma}\Phi(z) - \phi(z) - \sigma\frac{\partial}{\partial \sigma}\phi(z) \\
&= (\alpha\gamma)\phi(z)\frac{\partial z}{\partial \sigma} - \phi(z) + \sigma z\phi(z)\frac{\partial z}{\partial \sigma} \\
&= -\phi(z)\left[-(\alpha\gamma)\frac{\partial z}{\partial \sigma} + 1 - \sigma z\frac{\partial z}{\partial \sigma}\right] \\
&= -\phi(z)\left[(\alpha\gamma)\frac{z}{\sigma} + 1 + \sigma z\frac{z}{\sigma}\right] \\
&= -\phi\left(\frac{(1-\alpha)\gamma}{\sigma}\right)\frac{(1-\alpha)\gamma^2 + \sigma^2}{\sigma^2} < 0.
\end{aligned} \tag{B.3}$$