Graphon Games for Optimal Investment in a Competitive Market

Ananya Parashar

Advisor: Professor Ludovic Tangpi

Submitted in partial fulfillment

of the requirements for the degree of

Bachelor of Science in Engineering

Department of Operations Research and Financial Engineering

Princeton University

May 2024

I hereby declare that I am the sole author of this thesis.

I authorize Princeton University to lend this thesis to other institutions or individuals for the purpose of scholarly research.

Ananya Parashar

I further authorize Princeton University to reproduce this thesis by photocopying or by other means, in total or in part, at the request of other institutions or individuals for the purpose of scholarly research.

Ananya Parashar

Abstract

In this paper, we study the portfolio optimization problem formulated by Lacker and Zariphopoulou [5]. They formulate a finite time horizon model that allows agents to be competitive, measuring their CRRA utility not only by their absolute wealth but also relative performance compared to the average of other agents. We extend this model to include an individual weight that an agent may place on each other agent should they want to compete. We find the optimal control of this problem by deriving the HJB equation then prove it is a Nash Equilibrium by showing it is a fixed point. To find the optimal control in the graphon case, we first restrict the graphon to be continuous then approximate the discontinuous graphon with continuous ones to use this result. To conclude, we analyze the behavior of the Nash Equilibrium investment strategy through simulations and compare with prior results.

Acknowledgements

First and foremost, my deepest gratitude goes to my advisor Professor Ludovic Tangpi. Had he not responded to my cold email, I would not have done research during the duration of my junior year and continued on to write a theoretical thesis this year. These research experiences under Professor Tangpi's mentorship transformed my academic aspirations, turning the once distant thought of graduate studies into a reality. His patience and generosity in sharing his time and knowledge have not only shaped my intellectual growth but also modeled the kind of person and scholar I aspire to be.

Not a day goes by where I am not thankful for my family, whose unwavering support and love have been my bedrock throughout my four years at Princeton. Every phone call, every word of encouragement, and every visit home has been a source of strength for me.

To the person who pushed me to step out of my comfort zone by doing research and taking graduate level classes, I thank you from the bottom of my heart.

Last and certainly not least, I extend my thanks to my friends. To those who have been with me from the beginning, who shared moments in Spoon and through First College, and to those I've grown close to in recent times in Dod 401 – Princeton would not be the same without you. Thank you for supporting me in ways I thought not possible and for the memories I will never forget.

To Dad, Mom, Aarushi, and Atharv

Contents

List of Figures

Chapter 1

Introduction

Today, as financial markets become more complex and interconnected, the pursuit of optimal investment strategies has never become more important. Whether it be individual investors or institutional investors, who manage important funds such as pension funds or insurances, each seek to maximize their reward. Typically, we see that this is not so straightforward a problem to solve as maximizing reward comes at the cost of incurring high volatility. That is, while we can maximize our reward, we also have a very high downside, thus what most investors pursue is maximal reward while minimizing risk. This balance between risk and reward is achieved differently for various investors: those that are more risk tolerant can see higher expected reward but also a higher downside risk in the case that their investments rapidly decline.

Clearly, this optimal investment problem is rooted in decision making, a fundamental and important aspect of life. This type of decision making motivates an entire class of problems known as optimal control problems. Such problems vary from minimizing fuel consumption for a rocket landing in space to determining the dynamic price for a new product being introduced into the market. The first of these problems has historical importance: it was widely studied during the race to the moon in the 50s and 60s [2]. To formulate each of these problems, given the dynamics of a system, we pick variables called *controls* at each time step. We choose these decision variables based on the information that is available in order to optimize some objective. In the rocket problem, given the dynamics of the rocket moving through space, we may pick $\alpha(t)$ at each time step t, where this represents the thrust at time t achieved by burning fuel. By formulating this optimization problem as in [2], we find that it is optimal to use either the full thrust or not use thrust at all.

Overall, it is clear that this class of problems allows for powerful modeling in decision making problems; therefore, we use this to model the optimal investment problem. Here, the controls we define are the proportion we invest in risky assets at each time step. The goal thus of our research is to develop a model to provide this optimal strategy for investors, who are individually competitive with one another and, in general, have varying attitudes about competition. Doing so not only provides important results in an economic sense but contributes to the field of graphon games.

Chapter 2

Literature Review

In this section, we begin by reviewing studies of the single investor optimal investment problem, classically known as Merton's problem. We expand the setting by considering n investors who seek to optimize investments as well as outperform other investors, then finally look at a continuum of investors who seek to do the same.

2.1 Merton's Portfolio Problem

In 1969, Robert Merton produced his groundbreaking paper "Lifetime Portfolio Selection under Uncertainty: The Continuous-Time Case" which formulated the single investor optimal investment problem as well as produced a solution [7]. Note that in his paper, Merton uses different notation; however, for consistency, we write the equations using the same notation as in the remainder of this paper.

We first delve into the model itself as gaining intuition for this is essential to understanding the dynamics in subsequent sections. Merton's problem investigates a singular investor who seeks to maximize their terminal expected utility by investing in a risky asset and a riskless asset. The riskless asset, normally thought of as a bond, has price S_t^0 that evolves, with fixed interest rate r, according to the following:

$$dS_t^0 = rS_t^0 dt. \tag{2.1}$$

The risky asset, traditionally thought of as a stock, has price evolution according to the following stochastic differential equation:

$$dS_t = \mu S_t dt + \sigma S_t dW_t, \tag{2.2}$$

where $\mu > 0$ is a constant notating the drift rate of the price process, $\sigma > 0$ is also a constant representing the volatility of the price, and W is a standard Brownian motion on a filtered probability space. Now define c_t as the investor's consumption per unit of wealth at time t, where we assume $c_t \in [0, 1]$. Moreover, let π_t be the proportion of wealth (potentially negative) invested in the stock at time t, taking values in some closed, convex subset of $\mathbb{R}$, say A. Consequently, $1 - \pi_t$ is the proportion invested in bonds at time t, thus we have the investor's wealth process as:

$$dX_t = \frac{X_t \pi_t}{S_t} dS_t + \frac{X_t(1 - \pi_t)}{S_t^0} dS_t^0 - X_t c_t dt.$$

Intuitively, the investor buys $X_t \pi_t / S_t$ shares of stocks with $X_t \pi_t$ of their wealth given the price S_t per share. The same reasoning holds true for the number of bonds the investor buys. Now, substituting Equation (2.1) and Equation (2.2) for the bond and stock price processes respectively then rearranging, we have:

$$dX_t = X_t(\pi_t \mu + (1 - \pi_t)r - c_t)dt + X_t \pi_t \sigma dW_t.$$

With this wealth process, Merton formulates the optimization objective of an agent, given some terminal time T, utility function U, and discount rate ρ as:

$$\max_{(\pi,c)\in\mathcal{A}\times\mathcal{C}} \mathbb{E}\left[\int_t^T e^{-\rho s}U(c_s X_s)ds + B(X_T, T)\right]$$

where $\mathcal{A}$ and $\mathcal{C}$ are the sets where the strategies take value in. Moreover, B is defined by Merton as the bequest valuation function written as $\epsilon^{1-\gamma}e^{-\rho T}U(X_T)$.[1] Here, the parameter $\gamma \in [0,1)$ is the risk aversion and the parameter ϵ allows us to scale the utility of the bequest in comparison to the utility from consumption. Since Merton specifies that $0 < \epsilon \ll 1$, we know that the agent derives significantly less utility from wealth left as a bequest than from wealth consumed.

To find a solution to this optimization problem, first consider an agent with constant relative risk aversion (CRRA) utility, where we may see that $1-\gamma = 1/\delta$ is the Pratt measure of relative risk aversion:[2]

$$U(x;\gamma) = \begin{cases} \frac{1}{\gamma}x^\gamma, & \text{if } \gamma \neq 0 \\ \log(x), & \text{if } \gamma = 0. \end{cases}$$

The solution to the problem over a finite time horizon (i.e. T is finite) or the optimal fraction invested in stocks $\pi^*(t)$ as well as the optimal consumption rate $c^*(t)$ was derived using Bellman's Principle of Optimality.[3] The solution is as follows:

$$\pi^* = \frac{\mu - r}{\sigma^2(1-\gamma)} \quad \text{and} \quad c^*(t) = \begin{cases} \frac{1}{1+(v\epsilon-1)e^{-v(T-t)}}, & \text{for } v \neq 0 \\ \frac{1}{T-t+\epsilon}, & \text{for } v = 0 \end{cases}$$

where $v \equiv \frac{\mu}{1-\gamma}$. The fraction invested in stocks is surprisingly independent of time

[1]We may alter how we value terminal wealth by using a separate utility function from $U(X_T)$.

[2]Throughout this paper, we will see that δ is the inverse of the risk aversion measure, and we will eventually define this as risk tolerance of an agent.

[3]Bellman's Principle of Optimality is the discrete time analogue of the Hamilton-Jacobi-Bellman equation which we use for our continuous time setting.

and has thus acquired the name "Merton's fraction." The solution in the infinite time horizon case is almost identical.

Merton also derives a solution using similar techniques for the Constant Absolute Risk Aversion (CARA) utility but only in the infinite time horizon case with the CARA utility defined as:

$$U(x) = -e^{-\eta X}/\eta,$$

where $\eta > 0$ is Pratt's measure of absolute risk aversion. The resulting solution is

$$\pi^* = \frac{a-r}{\eta r \sigma^2 X(t)} \qquad \text{and} \qquad c^*(t) = r + \frac{\rho - r + (a-r)^2/2\sigma^2}{X(t)\eta r}.$$

Since Merton's paper has been published, this investment problem has become a fundamental, widely studied problem in portfolio theory. As we discussed, Merton assumes investment in a singular stock and a singular bond which we may expand to multiple stocks. Note also that Merton employs several simplifying assumptions, such as constant interest rate as well as constant parameters in the stock price process (i.e. constant μ and constant σ). While the topic of expanding these parameters has been addressed in papers beyond Merton, we will see that these assumptions are useful in the tractability of our problem, thus we only consider literature which employs these simplifications.[4] Nevertheless, in Merton's model, it is clear that a singular agent investing in the market is not a realistic model and thus we extend this to n players.

2.2 Competitive N-Agent Models

As mentioned before, analysis is typically extended to the n-player setting to obtain a more realistic model of the market and hence this model is more complex. With n players, we may study how these players interact, or how the actions of other players

[4] Aside from the graphon game where literature is not available for constant drift and volatility.

affect a particular player. Here, we seek to find the Nash Equilibrium of players – where no one player would be better off should they unilaterally deviate strategies.

In particular, consider n investors each seeking to achieve the same goal as previously stated. That is, they hope to maximize their expected utility while consuming their wealth at intermediate periods. Without interaction between agents, this problem boils down to the single investor problem, thus hasn't been widely studied. Rather, the extension of Merton's problem to n-players with a specific interaction was first introduced by Espinosa in 2010 in his PhD thesis, but more formally presented in his paper "Optimal Investment Under Relative Performance Concerns" [1] which was first published in 2011. As in Merton's problem, he considers the market to contain one riskless bond but now d risky assets. Here, he introduces a setting where the agents interact by competing, where each agent seeks to relatively outperform his competitors.

Using Espinosa's model for the CARA utility case, we consider the arithmetic mean as the metric agents compete against. Excluding agent i, this average may be written as $Y_t^i := \frac{1}{n}\sum_{k\neq i} X_t^k$.[5] Then, we consider the following objective, with $\mathcal{A}$ being the set of admissible portfolios, which we will define later for our purposes, and θ_i being agent i's competition weight on the average of other agents' wealths:

$$\begin{aligned}
&\sup_{\pi^i\in\mathcal{A}} \mathbb{E}\left[U_i\left((1-\theta_i)X_T^i + \theta_i(X_T^i - Y_T^i)\right)\right] \\
&= \sup_{\pi^i\in\mathcal{A}} \mathbb{E}\left[U_i\left(X_T^i - \theta_i Y_T^i\right)\right] \\
&= \sup_{\pi^i\in\mathcal{A}} \mathbb{E}\left[-\exp\left(-\frac{1}{\delta_i}\left(\left(1-\frac{\theta_i}{n}\right)X_T^i - \theta_i Y_T^i\right)\right)\right].
\end{aligned} \tag{2.3}$$

Clearly, the first supremum above shows how each agent takes into account a convex combination of their absolute wealth as well as their relative wealth (i.e. how much

[5] Note that we will use $k \neq i$ throughout this paper to denote all agents k from 1 through n excluding agent i.

more wealth they have than the average). Moreover, we write U_i as a general utility function for each agent i, but Espinosa specifically considers CARA utility parameterized by δ_i, the ith agent's risk tolerance: $U_i(x;\delta_i) = -e^{-\frac{1}{\delta_i}x}$. Thus the last step reflects this choice of utility function.

Espinosa proceeds to provide optimal solutions under various portfolio constraints by using equilibrium pricing; however, these techniques lead to n-dimensional BSDE (Backward Stochastic Differential Equation) systems and thus are hard to solve as noted in [5]. Because this paper finds difficulty with the well-posedness of the problems and solutions, Espinosa mainly proves the existence of the Nash Equilibrium and thus we look to [5] which solves this competitive n-agent model using stochastic optimal control techniques as was done in Merton's paper.

In [5], Daniel Lacker and Thaleia Zariphopoulou explicitly solve out the Nash Equilibrium for the CARA utility case, using the objective introduced by Espinosa in Equation (2.3). They make a key assumption for tractability that investment strategies are chosen at time 0, and are thus constants (i.e. $\pi_t = \pi$). Moreover, they assume that the market consists now of n stocks, where S^i for stock $i \in [n]$ is agent i's individual stock.[6] The market also consists of a common riskless bond (with 0 interest rate). Thus we have the stock price is:

$$dS^i_t = \mu_i S^i_t dt + \nu_i S^i_t dW^i_t + \sigma_i S^i_t dB_t \tag{2.4}$$

Here, $\mu_i > 0$, $\nu_i, \sigma_i \geq 0$ are all constants with μ_i is the drift rate, ν_i is the volatility of the prices for stock i (idiosyncratic volatility), and σ_i is the common noise that affects all stocks equally (systematic volatility). They also enforce that $\nu_i + \sigma_i > 0$. W^i is a standard Brownian motion on the filtered probability space $(\Omega, \mathcal{F}, \mathbb{F} = (\mathcal{F}_t)_{t\in[0,T]}, \mathbb{P})$. We will use the same probability space for subsequent sections.

[6]Note that we will use the notation $[n]$ to denote the set $\{1, \ldots, n\}$ for the remainder of the paper.

Enforcing some measurability conditions for the π proportion the agent invests and taking $\mathcal{A}$ to be the set of admissible portfolios, Lacker and Zariphopoulou then write the ith agent's wealth process as, using the same process as we introduced in Merton:

$$dX_t^i = X_t^i \pi_t^i (\mu_i dt + \nu_i dW_t^i + \sigma_i dB_t) \tag{2.5}$$

with initial condition $X_0^i = x_0^i$.[7] The solution to this problem with objective in Equation (2.3), where other players have fixed strategies, is derived through the Hamilton-Jacobi-Bellman (HJB) equation in [5], and is as follows, for constants defined by

$$\varphi_n := \frac{1}{n}\sum_{k=1}^{n} \delta_k \frac{\mu_k \sigma_k}{\sigma_k^2 + \nu_k^2(1-\theta_k/n)} \qquad \text{and} \qquad \psi_n := \frac{1}{n}\sum_{k=1}^{n} \theta_k \frac{\sigma_k^2}{\sigma_k^2 + \nu_k^2(1-\theta_k/n)}.$$

If $\psi_n < 1$, the Nash equilibrium is given by:

$$\pi^{i,*} = \delta_i \frac{\mu_i}{\sigma_i^2 + \nu_i^2(1-\theta_i/n)} + \theta_i \frac{\sigma_i}{\sigma_i^2 + \nu_i^2(1-\theta_i/n)} \frac{\varphi_n}{1-\psi_n}, \tag{2.6}$$

else if $\psi_n = 1$, there does not exist an equilibrium.

They also introduce, for the first time, modeling this problem under CRRA utility. Here, each agent compares themselves with the geometric mean of agents[8] defined by $Y_t = \left(\prod_{k \neq i} X_t^k\right)^{\frac{1}{n}}$. Again, we consider our utility function parameterized by δ_i:

$$U_i(x;\delta_i) = \begin{cases} \frac{1}{1-\delta_i} x^{1-1/\delta_i}, & \text{if } \delta_i \neq 1 \\ \log(x), & \text{if } \delta_i = 1. \end{cases} \tag{2.7}$$

Then, we have the following objective, with $\overline{X}_T$ being the geometric mean of all agents

[7]This is true since $dX_t^i = \frac{X_t^i \pi_t^i}{S_t^i} dS_t^i$ then we substitute Equation (2.4) for the stock price process.

[8]This is due to the difference in the form of the utility and leads to tractability of the problem. The interested reader may read [5] for more discussion on this topic.

(including i) at terminal time T:

$$\sup_{\pi^i \in \mathcal{A}} \mathbb{E}\left[U_i\left(X_T^i \overline{X}_T^{-\theta_i}\right)\right] = \sup_{\pi^i \in \mathcal{A}} \mathbb{E}\left[U_i\left(X_T^{i^{1-\theta_i/n}} Y_T^{i^{\theta_i}}\right)\right].$$

The authors prove that the constant Nash equilibrium always exists, and thus for constants

$$\varphi_n := \frac{1}{n}\sum_{k=1}^{n} \delta_k \frac{\mu_k \sigma_k}{\sigma_k^2 + \nu_k^2(1 + (\delta_k - 1)\theta_k/n)},$$
$$\psi_n := \frac{1}{n}\sum_{k=1}^{n} \theta_k(\delta_k - 1)\frac{\sigma_k^2}{\sigma_k^2 + \nu_k^2(1 + (\delta_k - 1)\theta_k/n)}$$

the constant equilibrium is given by:

$$\pi^{i,*} = \delta_i \frac{\mu_i}{\sigma_i^2 + \nu_i^2(1 + (\delta_i - 1)\theta_i/n)} - \theta_i(\delta_i - 1)\frac{\sigma_i}{\sigma_i^2 + \nu_i^2(1 + (\delta_i - 1)\theta_i/n)} \frac{\varphi_n}{1 - \psi_n} \quad (2.8)$$

In 2019, Daniel Lacker and Agathe Soret introduced consumption into this competitive n-agent problem, noting that it was the first paper to do so [3]. However, they only introduce consumption per unit wealth c_t^i for the CRRA utility problem and not for CARA utility. Under consumption, the new wealth process for agent i is similar to Equation (2.5) with the intermediate consumption term:

$$dX_t^i = X_t^i \pi_t^i(\mu_i dt + \nu_i dW_t^i + \sigma_i dB_t) - c_t^i X_t^i \quad (2.9)$$

and thus the optimization objective for agent i is given by, where U is the CRRA utility function defined in Equation (2.7),

$$\sup_{(\pi^i, c^i) \in \mathcal{A} \times \mathcal{C}} \mathbb{E}\left[\int_0^T U\left((c_t^i X_t^i)^{1-\theta_i/n}(\overline{c}_{-i}(t) Y_t)^{-\theta_i}; \delta_i\right) + \epsilon_i U\left((X_T^i)^{1-\theta_i/n} Y_T^{-\theta_i}; \delta_i\right)\right]. \quad (2.10)$$

Here, $\overline{c}_{-i}(t) = (\prod_{k \neq i} c_k(t))^{\frac{1}{n}}$. We may see that the optimization problem now includes an integral since we care about intermediate consumption in addition to the final wealth. Furthermore, we use $\epsilon_i > 0$ to scale the terminal wealth in accordance with the importance the agent places on intermediate consumption compared to terminal wealth. The solution to this $\pi^{i,*}$ is the same as in the non-consumption case given by Equation (2.8). For the optimal consumption, they write:

$$c_t^{i,*} = \begin{cases} \left(\frac{1}{\beta_i} + \left(\frac{1}{\lambda_i} - \frac{1}{\beta_i}\right) e^{-\beta_i(T-t)}\right)^{-1} & \text{if } \beta_i \neq 0 \\ (T - t + \lambda_i^{-1})^{-1} & \text{if } \beta_i = 0 \end{cases} \tag{2.11}$$

where β_i and λ_i are constants defined in [3]. While we can find solutions in the CRRA case, these n-player problems are limited in tractability. As noted in [3], the CARA utility function with consumption only potentially has a solution. Thus, finding a solution using a different model could be beneficial, as a more tractable framework extends itself to more complex models. Thus we look at Mean Field Games (MFGs) to solve this problem.

2.3 Mean Field Games for Competitive Agents

The first application of MFG theory to portfolio optimization is in [5] by Lacker and Zariphopoulou looking at both CRRA and CARA utility functions. Moreover, the introduction of consumption into the MFG for CRRA utility is done in [3]. The MFG is a framework where, rather than n agents, we assess the continuum of agents (i.e. as $n \to \infty$). This analysis is useful to analyze a large population of players as analyzing the n-player game directly can become intractable due to the combinatorial explosion of possible states and strategies. This framework allows the model to become more tractable by understanding the distribution of states or strategies in the population rather than tracking every individual interaction.

We look at Lacker and Soret's analysis of the MFG in [3] with consumption for the

CRRA utility case as that is the most general. To analyze the MFG, we first consider the n-player game where each player $i \in [n]$ has type vector $\zeta_i := (\xi_i, \delta_i, \theta_i, \mu_i, \nu_i, \sigma_i)$. Here, the agent's type vector is their specified initial wealth (ξ), individual preference parameters (δ, θ), and market parameters (μ, ν, σ). The equilibrium strategies (for investment and consumption) in the n-player game, as in Equations (2.8) and (2.11), only depend on this type vector as well as the type distribution of these type vectors, written as $m_n = \frac{1}{n}\sum_{k=1}^{n} \delta_{\zeta_k}$. Thus as $n \to \infty$, by Law of Large Numbers, the type distribution converges weakly to m, a limiting probability measure and thus the equilibrium strategy converges as well. Rather than doing this analysis in context of the n-player game, we represent the MFG as its own game for a continuum of agents with type distribution m.[9]

Thus, to analyze the MFG, the authors pick a representative agent at random from the continuum of agents and assign them a random type vector $\zeta := (\xi, \delta, \theta, \mu, \nu, \sigma)$ then solve. They solve by first considering

$$\overline{X}_t := \exp \mathbb{E}[\log X_t | \mathcal{F}_t^B] \quad \text{and} \quad \overline{\Gamma}_t := \exp \mathbb{E}[\log c_t | \mathcal{F}_t^B],$$

the continuous analog of geometric mean.[10] The objective is

$$\sup_{(\pi,c)\in\mathcal{A}_{MF}} \mathbb{E}\left[\int_0^T U\left(c_t - \theta\overline{c}_t; \delta\right) dt + \epsilon U\left(X_T - \theta Y_T; \delta\right)\right]. \tag{2.12}$$

Using the HJB equation to solve, the optimal solution (π^*, c^*) is (the unique strong

[9]This explanation is adapted from [3] and [5]. For more careful analysis, we defer the reader to either of these papers.

[10]Here, we also condition on the filtration generated by the common noise for consistency. That is, we want $\exp \mathbb{E}[\log X_t | \mathcal{F}_t^B] = \overline{X}$.

equilibrium):

$$\pi^* = \frac{\delta\mu}{\sigma^2+\nu^2} - \frac{\theta(\delta-1)\sigma}{\sigma^2+\nu^2}\frac{\phi}{1+\psi}$$
$$c_t^* = \begin{cases} \left(\frac{1}{\beta} + \left(\frac{1}{\lambda} - \frac{1}{\beta}\right) e^{-\beta(T-t)}\right)^{-1}, & \text{if } \beta \neq 0 \\ (T-t+\lambda^{-1})^{-1}, & \text{if } \beta = 0 \end{cases}$$

which is exactly the same as what was derived in the n player game rather with the limit of the original constants.

In this model, we may see that these competition weights θ_i are quite restrictive, since agent i may only weigh the average of the other agents. It does not take into account whether agent i cares more about agent j as opposed to agent k ($j, k \neq i \in [n]$). Thus, we keep this θ_i to measure how much the agent i cares about their relative wealth compared to their absolute wealth and introduce a new parameter. The rest of this paper focuses on this model with this new parameter both in the setting of n players as well as in the setting of a continuum of players should that make the problem more tractable.

2.4 Graphon Games for Competitive Agents

Since these weights are individual, the MFG approach cannot be extended to this model as it can only be used for homogeneous agents (up to their type vector). This weight, in fact, breaks the symmetry required for the MFG and thus we look at graphon games to analyze this more heterogeneous problem.

The introduction of the weight λ_{ik} for which player i weighs player k with was done both in the n-player game and in the limit in [9] by Ludovic Tangpi and Xuchen Zhou. Here, Tangpi and Zhou consider a slightly different model than the ones described above where the stock evolution for each player $u \in I$ where $I = [0, 1]$ is as follows:

$$dS_t^u = \text{diag}(S_t^u)(\mu_t^u dt + \sigma_t^u dW_t^u + \sigma_t^{*u} dW_t^*)$$

where we now have a continuum of stocks indexed by u. Here, we consider a riskless bond with zero interest rate as we did before, but now the drift term and volatility terms are not constants rather they are dependent on time. Moreover, as opposed to Lacker and Zariphopoulou's model, each agent in the continuum of agents invest in all stocks as opposed to being restricted to their own individual stock. Thus agent u's wealth process with strategy π^u follows the dynamics below, where we assume the agent does not consume wealth at intermediate periods:

$$dX_t^u = \pi_t^u \cdot (\Sigma_t^u \theta_t^u dt + \sigma_t^u dW_t^u + \sigma_t^{*u} dW_t^*), \quad X_0^u = \xi^u$$

where $\Sigma_t^u = (\sigma_t^u, \sigma_t^{*u})$ and $\theta_t^u = \Sigma_t^{u\top} \left(\Sigma_t^u \Sigma_t^{u\top}\right)^{-1} \mu_t^u$. Tangpi and Zhou then fix graphon $G : I \times I \to I$ which is a symmetric and measurable function. The utility maximization problem for agent u is the following under CARA utility:

$$V_0^{u,G} = V_0^{u,G}\left((\pi^v)_{v \neq u}\right) \tag{2.13}$$

$$:= \sup_{\pi^u \in \mathcal{A}^G} \mathbb{E}\left[-\exp\left(-\frac{1}{\eta_u}\left(X_T^{u,\pi^u} - \mathbb{E}\left[\rho \int_I X_T^{v,\pi^v} G(u,v)dv | \mathcal{F}_T^*\right]\right)\right)\right] \tag{2.14}$$

where $\mathbb{F}^* := (\mathcal{F}_t^*)_{t \in [0,T]}$ is the $\mathbb{P}$-completion of the filtration generated by W^*. Here, since the agent uses the CARA utility function, the model subtracts the arithmetic average of other agents' wealth from the agent's wealth. The main result of this paper is firstly that the graphon game exists and secondly that the n player game with discrete weights converges to the graphon game with a continuum of weights. Tangpi and Zhou prove these results extensively both in the common noise case and without. In our paper, we consider a similar problem but with the CRRA utility under constant drift and volatility terms as well as only investment in individualized stocks. However, we approach the graphon game without using the BSDE methods introduced in this paper, rather we use aforementioned HJB equations.

Chapter 3

Weighted N Player Game Model

In this section, we introduce a model, similar to that introduced in [3] for CRRA utility, but where we introduce a new parameter to weigh each individual competitor. This new problem can be formulated as a question: in a game of n agents, how do we maximize the expected terminal utility of each agent who have an individualized competition weight with respect to each competitor, as well as a relative performance weight to determine how much they care about competition? We proceed to solve this new problem by deriving the Hamilton-Jacobi-Bellman equation, guessing a suitable ansatz, then obtaining a closed form solution for the optimal control. To extend this best response solution to the Nash Equilibrium, we prove that there exists a unique fixed point of the best response using the Banach Fixed Point Theorem.

3.1 Introduction

To start, we formulate the optimization objective for the CRRA utility function in the n-player game with individual weights. This model looks similar to (2.10) that we introduced in the n-player competitive agent model.

Throughout the remainder of this paper, we will work on the probability space $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$ where the filtration $\mathbb{F} = (\mathcal{F})_{t\in[0,t]}$ is the natural filtration generated by the Brownian motion of each stock as well as the common noise. Again, we have that each agent's wealth evolves as in (2.5), assuming a zero interest rate, with each agent trading their individual stock i at each time step t over a common finite investment horizon $[0,T]$:

$$dX_t^i = X_t^i \pi_t^i (\mu_i dt + \nu_i dW_t^i + \sigma_i dB_t)$$

with initial condition $x_0^i \in \mathbb{R}$. Each individual stock i here has drift rate $\mu_i > 0$, volatility $\nu_i \geq 0$, and common noise coefficient $\sigma_i \geq 0$ where $\sigma_i + \nu_i > 0$. Note that the Brownian motions $B_1, \ldots, B_n, W$ are independent. Each agent uses the CRRA utility function as below, with personal risk tolerance parameter $\delta_i > 0$:

$$U_i(x;\delta_i) = \begin{cases} \left(1 - \frac{1}{\delta_i}\right) x^{1-\frac{1}{\delta_i}}, & \text{if } \delta_i \neq 1 \\ \log(x), & \text{if } \delta_i = 1. \end{cases}$$

Specifically, each agent i seeks to maximize their expected utility under this CRRA utility. They do so by picking the optimal control π_t^i at each time step $t \in [0,T]$ where this represents the fraction of wealth invested in stocks. π_t^i is an admissible strategy if it satisfies the following definition.

Definition 3.1.1 (Admissible Strategy)**.** *A portfolio strategy π is admissible if it is an $\mathbb{R}$ valued, $\mathbb{F}$ progressively measurable process satisfying the integrability condition $\mathbb{E}[\int_0^T |\pi_t|^2 dt] < \infty$.*

Thus, if each agent i picks admissible strategy π^i for all $i \in [n]$, agent i's expected payoff J_i is the following:

$$J_i(\pi^1, \ldots, \pi^n) := \mathbb{E}\left[U\left(X_T^i \left(Y_T^i\right)^{-\theta_i}, \delta_i\right)\right], \text{where } Y_T^i = \left(\prod_{k=1}^n (X_t^k)^{\tilde{\lambda}_{ik}}\right)^{\frac{1}{\sum_{k=1}^n \tilde{\lambda}_{ik}}}. \quad (3.1)$$

Expected utility here is a function of the competition parameter which captures the tradeoff between absolute and relative wealth: with θ_i close to 0, agent i gains utility from a higher absolute terminal wealth whereas θ_i close to 1 signifies that agent i gains utility from outperforming their competition, the weighted geometric mean of other agents' terminal wealth.

Again, we see the similarities to Lacker and Soret's model introduced in [3], the only difference being the added parameter weighing each individual agent $\widetilde{\lambda}_{ik}$. That is, $\widetilde{\lambda}_{ik} \in [0,1]$ is the weight that agent i gives each other agent k when competing with them. We make the modeling decision to write Y_T^i as the above with $\widetilde{\lambda}_{ik}$ in exponent since this is equivalent to

$$\log Y_T^i = \frac{1}{n} \sum_{k=1}^{n} \widetilde{\lambda}_{ik} \log(X_T^k),$$

which is the weighted arithmetic average.

$\widetilde{\lambda}_{ik}$ is similar to θ_i which weighs the whole population but rather is individualized: if $\widetilde{\lambda}_{ik}$ is 1, then agent i cares about agent k's wealth whereas if $\widetilde{\lambda}_{ik}$ is 0, then agent i is indifferent to agent k's wealth. We take $\widetilde{\lambda}_{ii} = 0$ as we don't want agent i to consider itself in its competition. We define these weights more formally as an indicator below:

$$\widetilde{\lambda}_{ij} = \begin{cases} 1, & \text{if agent } i \text{ cares about agent } j\text{'s performance} \\ 0, & \text{otherwise.} \end{cases}$$

Thus, the weight parameter and the competition parameter θ_i work in conjunction: while the weight of each agent that agent i cares about is 1, we know that these weights can be magnified by θ_i close to 1 or diminished by θ_i being close to 0. Now, we may write as before

$$Y_t^i = \left(\prod_{k \neq i} (X_t^k)^{\widetilde{\lambda}_{ik}} \right)^{\frac{1}{\sum_{k=1}^{n} \widetilde{\lambda}_{ik}}} = \prod_{k \neq i} (X_t^k)^{\lambda_{ik}}$$

where $\lambda_{ik} = \frac{\widetilde{\lambda}_{ik}}{\sum_{k=1}^{n} \lambda_{ik}}$.[1]

Given this model, we seek to find a Nash equilibrium of the n players defined in [5] as the following:

Definition 3.1.2 (N-Player Nash Equilibrium)**.** *A vector $(\pi^{1,*}, \ldots, \pi^{n,*})$ of admissible strategies is a Nash Equilibrium if, for every player $i = 1, \ldots, n$, the following inequality holds true for all admissible strategies $\pi^i \in \mathcal{A}$:*

$$J_i(\pi^{1,*}, \ldots, \pi^{i,*}, \ldots, \pi^{n,*}) \geq J_i(\pi^{1,*}, \ldots, \pi^{i-1,*}, \pi^i, \pi^{i+1,*}, \ldots, \pi^{n,*}). \tag{3.2}$$

In addition to the above condition, if, for all players i, their corresponding strategy $\pi^{i,*}$ remain constant over the finite horizon $[0, T]$ (formally, $\pi_t^{i,*} = \pi_0^{i,*}\ \forall t \in [0, T]$), this is a constant Nash Equilibrium.

We prove the existence of a Nash Equilibrium by fixing other players strategies to be $(\pi^1, \ldots, \pi^{i-1}, \pi^{i+1}, \ldots, \pi^n)$ then computing the best response of player i and finding a fixed point. This is done by maximizing the payoff of player i in (3.1) over π^i, where we consider the set of admissible strategies $\mathcal{A}$:

$$\sup_{\pi^i \in \mathcal{A}} \mathbb{E}\left[U\left(X_T^{i^{1-\theta_i/n}} {Y_T}^{i^{-\theta_i}}; \delta_i\right)\right]. \tag{3.3}$$

To solve this optimization problem, we must derive the Hamilton-Jacobi-Bellman (HJB) equation and find the solution.

3.2 HJB Equation

Each agent must now solve the same optimization problem as in (2.10). To solve, we will need to derive the HJB equation. While typically, a problem with n agents lends itself to an n-dimensional HJB equation, we rather treat the geometric mean Y_t as an

[1]Each λ_{ik} is thus fractional and retains the property of $\lambda_{ii} = 0$.

uncontrolled state variable. This allows us to derive a single HJB equation. Before deriving the HJB equation, we first proceed in a similar manner as the competitive n-player game, deriving the dynamic of this new state variable Y_t.

3.2.1 Derivation of Geometric Mean Process

We start off by fixing agent i as well as constant (time-independent) investment strategies for agents $k \neq i$ (i.e. we may write $\pi_t^k = \pi^k$). Thus, each X_t^k solves (2.10) with initial solution $X_0^k = x_0^k$, that is:

$$dX_t^k = X_t^k \pi^k (\mu_k dt + \nu_k dW_t^k + \sigma_k dB_t). \tag{3.4}$$

The application of Itô's formula allows us to write:

$$d(\log X_t^k) = \left(\mu_k \pi^k - \frac{1}{2}(\sigma_k^2 + \nu_k^2){\pi^k}^2 \right) dt + \nu_k \pi^k dW_t^k + \sigma_k \pi^k dB_t. \tag{3.5}$$

We may see that $d(\log Y_t^i) = \frac{1}{n}\sum_{k\neq i} \lambda_{ik} d(\log X_t^k)$ and thus, using the above (3.5),

$$d(\log Y_t^i) = \left(\widehat{\lambda_i \mu \pi}_{-i} - \frac{1}{2}\widehat{\lambda_i \Sigma \pi^2}_{-i} \right) dt + \frac{1}{n}\sum_{k\neq i} \lambda_{ik}\nu_k \pi^k dW_t^k + \widehat{\lambda_i \sigma \pi}_{-i} dB_t$$

where $\widehat{\lambda_i \mu \pi}_{-i} = \frac{1}{n}\sum_{k\neq i}\lambda_{ik}\mu_k\pi^k$, $\Sigma_k = \sigma_k^2 + \nu_k^2$ and thus $\widehat{\lambda_i \Sigma \pi^2}_{-i} = \frac{1}{n}\sum_{k\neq i}\lambda_{ik}\Sigma_k {\pi^k}^2$. Moreover, we have that $\widehat{\lambda_i \sigma \pi}_{-i} = \frac{1}{n}\sum_{k\neq i}\lambda_{ik}\sigma_k\pi^k$ and $\widehat{\lambda_i \nu \pi}_{-i} = \frac{1}{n}\sum_{k\neq i}\lambda_{ik}\nu_k\pi^k$. Using Itô's formula again to find $d(Y_t^i)$, we find that:

$$dY_t^i = Y_t^i(\eta_i)dt + Y_t^i \widehat{\lambda_i \nu \pi}_{-i} dW_t^k + Y_t^i \widehat{\lambda_i \sigma \pi}_{-i} dB_t \tag{3.6}$$

where $\eta_i = \widehat{\lambda_i \mu \pi}_{-i} - \frac{1}{2}\left(\widehat{\lambda_i \Sigma \pi^2}_{-i} - \widehat{\lambda_i \sigma \pi}_{-i}^{\,2} - \frac{1}{n}\widehat{(\lambda_i \nu \pi)^2}_{-i} \right)$ defined in terms of $\widehat{\lambda_i \mu \pi}_{-i}$, $\widehat{\lambda_i \Sigma \pi^2}_{-i}$ and $\widehat{\lambda_i \sigma \pi}_{-i}$ which we defined previously, and $\widehat{(\lambda_i \nu \pi)^2}_{-i} = \frac{1}{n}\sum_{k\neq i}\lambda_{ik}^2\nu_k^2{\pi^k}^2$.

3.2.2 Derivation of Coupled HJB Equation

To find the best response of player i, we derive the HJB equation, with Y_t^i being the solution to Equation (3.6) and X_t^i being the solution to Equation (3.4). However for simplicity, we write the SDEs as, where they are valued in $\mathbb{R}^n$:

$$dX_t^i = b_i^{(x)}(t, X_t^i)dt + \nu_i^{(x)}(t, X_t^i)dW_t^i + \sigma_i^{(x)}(t, X_t^i)dB_t$$
$$dY_t^i = b_i^{(y)}(t, Y_t^i)dt + \nu_i^{(y)}(t, Y_t^i)dW_t^i + \sigma_i^{(y)}(t, Y_t^i)dB_t$$

with initial conditions $X_0^i = x_0^i$ and $Y_0^i = y_0^i$. We couple these as one process $Z^i = (X^i, Y^i)$. We use notation $Z^{i,t,x,y}$ for the process Z^i with initial condition $(t, x, y) \in [0, T] \times \mathbb{R}^n \times \mathbb{R}^n$.

Theorem 3.2.1 (HJB Equation)**.** *The coupled HJB equation for the above problem is the following for constant control $a^i \in \mathbb{R}$ and value function φ^i:*

$$\varphi_t^i + \sup_{a^i \in \mathbb{R}} \Bigg[\varphi_x^i b_i^{(x)} + \varphi_y^i b_i^{(y)} + \varphi_{xy}^i (\nu_i^{(x)} \nu_i^{(y)} + \sigma_i^{(x)} \sigma_i^{(y)}) + \frac{1}{2} \varphi_{xx}^i (\nu_i^{(x)^2} + \sigma_i^{(x)^2})$$
$$+ \frac{1}{2} \varphi_{yy}^i (\nu_i^{(y)^2} + \sigma_i^{(y)^2}) \Bigg] = 0.$$

To derive this, as in [8], we first state the Dynamic Programming Principle (DPP).

Theorem 3.2.2 (Dynamic Programming Principle)**.** *For any $\mathbb{F}$ measurable stopping time $\theta \in \mathcal{T}_{t,T}$, where $\mathcal{T}_{t,T}$ is the set of $[t, T]$ valued stopping times:*

$$\varphi^i(t, x, y) = \sup_{\alpha^i \in \mathcal{A}} \mathbb{E}\left[\varphi^i(\theta, Z_\theta^{i,t,x,y})\right]$$

for control $\alpha^i = (\alpha_s^i)_{s \in [0,T]}$, a progressively measurable processes valued in $\mathcal{A} \subseteq \mathbb{R}^n$ and $(t, x, y) \in [0, T] \times \mathbb{R}^n \times \mathbb{R}^n$. Moreover, f_i is a measurable function such that $f_i : [0, T] \times \mathbb{R}^n \times \mathcal{A} \times \mathcal{C} \to \mathbb{R}$ for each $i \in [n]$. Finally, we define φ^i to be the value function such that $\varphi^i(x) = \sup_{\alpha^i \in \mathcal{A}} J_i(\pi^1, \pi^2, \ldots, \pi^{i-1}, \alpha^i, \pi^{i+1}, \ldots, \pi^n)$ where J_i is the *expected payoff as in (3.1).*

Prooof of Theorem 3.2.1. Consider the stopping time $\theta = t + h$ where h is a small value, and t is some time. Moreover, we fix constant controls $\alpha_s^i = a^i$ for some arbitrary $a^i \in \mathbb{R}$. By the DPP in Theorem 3.2.2, we know that:

$$\varphi^i(t,x,y) \geq \mathbb{E}\left[\varphi^i(t+h, Z_{t+h}^{i,t,x,y})\right]. \tag{3.7}$$

We start by looking at the right hand side, specifically $\varphi^i(t+h, Z_{t+h}^{i,t,x,y})$ where we assume $\varphi^i \in C^{1,2}$ in the interval $[t, t+h]$.[2] We may then apply Itô's Formula:

$$\begin{aligned}
\varphi^i(t+h, Z_{t+h}^{i,t,x,y}) = \varphi^i(t,x,y) &+ \int_t^{t+h} \frac{\partial \varphi^i}{\partial t} ds + \int_t^{t+h} \left(\frac{\partial \varphi^i}{\partial x} b_i^{(x)} + \frac{\partial \varphi^i}{\partial y} b_i^{(y)} \right) ds \\
&+ \int_t^{t+h} \frac{\partial^2 \varphi^i}{\partial x \partial y} (\nu_i^{(x)} \nu_i^{(y)} + \sigma_i^{(x)} \sigma_i^{(y)}) ds + \int_t^{t+h} \frac{1}{2} \frac{\partial^2 \varphi^i}{\partial x^2} (\nu_i^{(x)^2} + \sigma_i^{(x)^2}) ds \\
&+ \int_t^{t+h} \frac{1}{2} \frac{\partial^2 \varphi^i}{\partial y^2} (\nu_i^{(y)^2} + \sigma_i^{(y)^2}) ds + \int_t^{t+h} \left(\frac{\partial \varphi^i}{\partial x} \nu_i^{(x)} + \frac{\partial \varphi^i}{\partial y} \nu_i^{(y)} \right) dW_s^i \\
&+ \int_t^{t+h} \left(\frac{\partial \varphi^i}{\partial x} \sigma_i^{(x)} + \frac{\partial \varphi^i}{\partial y} \sigma_i(y) \right) dB_s.
\end{aligned}$$

Taking the expectation of this, we know that the Brownian motions W_s^i and B_s are independent for all i by prior assumption and thus $\mathbb{E}[dW_s^i dB_s] = \mathbb{E}[dW_s^i]\mathbb{E}[dB_s]$. Moreover, we know that the stochastic integrals (i.e. integrals with respect to dW_s^i or dB_s) have 0 expectation since they are martingales vanishing in 0.[3] We obtain:

$$\begin{aligned}
\mathbb{E}[\varphi^i(t+h, Z_{t+h}^{i,t,x,y})] = \varphi^i(t,x,y) + \mathbb{E}\Bigg[&\int_t^{t+h} \frac{\partial \varphi^i}{\partial t} ds + \int_t^{t+h} \left(\frac{\partial \varphi^i}{\partial x} b_i^{(x)} + \frac{\partial \varphi^i}{\partial y} b_i^{(y)} \right) ds \\
&+ \int_t^{t+h} \frac{\partial^2 \varphi^i}{\partial x \partial y} (\nu_i^{(x)} \nu_i^{(y)} + \sigma_i^{(x)} \sigma_i^{(y)}) ds + \int_t^{t+h} \frac{1}{2} \frac{\partial^2 \varphi^i}{\partial x^2} (\nu_i^{(x)^2} + \sigma_i^{(x)^2}) ds \\
&+ \int_t^{t+h} \frac{1}{2} \frac{\partial^2 \varphi^i}{\partial y^2} (\nu_i^{(y)^2} + \sigma_i^{(y)^2}) ds \Bigg].
\end{aligned}$$

[2] To the unfamiliar reader, this denotes the set of continuous functions differentiable in time t and twice differentiable in x as well as y.

[3] Typically, such stochastic integrals would only result in a local martingale. However, since we impose appropriate integrability conditions on the integrand, we recover a true martingale.

By Equation (3.7), since $\varphi^i(t,x,y) \geq \mathbb{E}\left[\varphi^i(t+h, Z_{t+h}^{i,t,x,y})\right]$, we may substitute the above expression into this equation:

$$\varphi^i(t,x,y) + \mathbb{E}\Bigg[\int_t^{t+h} \frac{\partial \varphi^i}{\partial t} ds + \int_t^{t+h} \left(\frac{\partial \varphi^i}{\partial x} b_i^{(x)} + \frac{\partial \varphi^i}{\partial y} b_i^{(y)}\right) ds \\ + \int_t^{t+h} \frac{\partial^2 \varphi^i}{\partial x \partial y} (\nu_i^{(x)} \nu_i^{(y)} + \sigma_i^{(x)} \sigma_i^{(y)}) ds + \int_t^{t+h} \frac{1}{2} \frac{\partial^2 \varphi^i}{\partial x^2} (\nu_i^{(x)^2} + \sigma_i^{(x)^2}) ds \\ + \int_t^{t+h} \frac{1}{2} \frac{\partial^2 \varphi^i}{\partial y^2} (\nu_i^{(y)^2} + \sigma_i^{(y)^2}) ds\Bigg] \leq \varphi^i(t,x,y).$$

Clearly, this is exactly

$$\mathbb{E}\Bigg[\int_t^{t+h} \frac{\partial \varphi^i}{\partial t} ds + \int_t^{t+h} \left(\frac{\partial \varphi^i}{\partial x} b_i^{(x)} + \frac{\partial \varphi^i}{\partial y} b_i^{(y)}\right) ds + \int_t^{t+h} \frac{\partial^2 \varphi^i}{\partial x \partial y} (\nu_i^{(x)} \nu_i^{(y)} + \sigma_i^{(x)} \sigma_i^{(y)}) ds \\ + \int_t^{t+h} \frac{1}{2} \frac{\partial^2 \varphi^i}{\partial x^2} (\nu_i^{(x)^2} + \sigma_i^{(x)^2}) ds + \int_t^{t+h} \frac{1}{2} \frac{\partial^2 \varphi^i}{\partial y^2} (\nu_i^{(y)^2} + \sigma_i^{(y)^2}) ds\Bigg] \leq 0.$$

Now we multiply by $\frac{1}{h}$ and take the limit as $h \to 0$. We also rewrite φ^i as $\varphi^i(s, Z_s^i)$ to represent that it is a function of s and Z_s^i,

$$\lim_{h \to 0} \frac{1}{h} \mathbb{E}\Bigg[\int_t^{t+h} \Bigg(\frac{\partial \varphi^i(s,Z_s^i)}{\partial t} + \frac{\partial \varphi^i(s,Z_s^i)}{\partial x} b_i^{(x)} + \frac{\partial \varphi^i(s,Z_s^i)}{\partial y} b_i^{(y)} + \frac{\partial^2 \varphi^i(s,Z_s^i)}{\partial x \partial y} (\nu_i^{(x)} \nu_i^{(y)} \\ + \sigma_i^{(x)} \sigma_i^{(y)}) + \frac{1}{2} \frac{\partial^2 \varphi^i(s,Z_s^i)}{\partial x^2} (\nu_i^{(x)^2} + \sigma_i^{(x)^2}) + \frac{1}{2} \frac{\partial^2 \varphi^i(s,Z_s^i)}{\partial y^2} (\nu_i^{(y)^2} + \sigma_i^{(y)^2})\Bigg) ds\Bigg] \leq 0.$$

Here, we may use the Fundamental Theorem of Calculus, since taking the limit as $h \to 0$ is equivalent to taking the derivative, thus by evaluating at $s = t$, we obtain

$$\mathbb{E}\Bigg[\frac{\partial \varphi^i(t,Z_t^i)}{\partial t} + \frac{\partial \varphi^i(t,Z_t^i)}{\partial x} b_i^{(x)} + \frac{\partial \varphi^i(t,Z_t^i)}{\partial y} b_i^{(y)} + \frac{\partial^2 \varphi^i(t,Z_t^i)}{\partial x \partial y} (\nu_i^{(x)} \nu_i^{(y)} + \sigma_i^{(x)} \sigma_i^{(y)}) \\ + \frac{1}{2} \frac{\partial^2 \varphi^i(t,Z_t^i)}{\partial x^2} (\nu_i^{(x)^2} + \sigma_i^{(x)^2}) + \frac{1}{2} \frac{\partial^2 \varphi^i(t,Z_t^i)}{\partial y^2} (\nu_i^{(y)^2} + \sigma_i^{(y)^2})\Bigg] \leq 0.$$

We know that $Z_t^i = (x, y)$ and so we may get rid of the expectation as well as simplify notation, where $\varphi^i_t = \frac{\partial \varphi^i(t,Z^i_t)}{\partial t}$ and other partial derivatives are notated similarly:

$$\varphi^i_t + \varphi^i_x b^{(x)} + \varphi^i_y b^{(y)} + \varphi^i_{xy}(\nu_i^{(x)}\nu_i^{(y)} + \sigma_i^{(x)}\sigma_i^{(y)}) + \frac{1}{2}\varphi^i_{xx}(\nu_i^{(x)^2} + \sigma_i^{(x)^2}) + \frac{1}{2}\varphi^i_{yy}(\nu_i^{(y)^2} + \sigma_i^{(y)^2}) \leq 0.$$

This above inequality is true for all a^i. Thus we may write that the inequality holds for the suprema, where a^i takes values in $\mathbb{R}$:

$$\varphi^i_t + \sup_{a^i\in\mathbb{R}}\Big[\varphi^i_x b_i^{(x)} + \varphi^i_y b_i^{(y)} + \varphi^i_{xy}(\nu_i^{(x)}\nu_i^{(y)} + \sigma_i^{(x)}\sigma_i^{(y)}) + \frac{1}{2}\varphi^i_{xx}(\nu_i^{(x)^2} + \sigma_i^{(x)^2}) + \frac{1}{2}\varphi^i_{yy}(\nu_i^{(y)^2} + \sigma_i^{(y)^2})\Big] \leq 0,$$

where φ^i_t is not in the supremum as it is a function of only t, x, y. For the other inequality, we follow the same steps instead taking $\alpha^{i,*}$ and $c^{i,*}$ to be the optimal controls in the DPP. That is, $\varphi^i(t,x) = \mathbb{E}[\varphi^i(t+h, Z^{i,t,x,y,\alpha^{i,*},c^{i,*}}_{t+h})]$. We apply Itô's formula to $\varphi^i(t+h, Z^{i,t,x,y,\alpha^{i,*},c^{i,*}}_{t+h})$ as we did before and follow the same steps applying h leading to:

$$\mathbb{E}\Big[\frac{\partial \varphi^i(t,Z^i_t)}{\partial t} + \frac{\partial \varphi^i(t,Z^i_t)}{\partial x} b_i^{(x)} + \frac{\partial \varphi^i(t,Z^i_t)}{\partial y} b_i^{(y)} + \frac{\partial^2 \varphi^i(t,Z^i_t)}{\partial x \partial y}(\nu_i^{(x)}\nu_i^{(y)} + \sigma_i^{(x)}\sigma_i^{(y)}) + \frac{1}{2}\frac{\partial^2 \varphi^i(t,Z^i_t)}{\partial x^2}(\nu_i^{(x)^2} + \sigma_i^{(x)^2}) + \frac{1}{2}\frac{\partial^2 \varphi^i(t,Z^i_t)}{\partial y^2}(\nu_i^{(y)^2} + \sigma_i^{(y)^2})\Big] = 0.$$

We take the supremum and realize that this is at least 0 by the above:

$$\mathbb{E}\Big[\frac{\partial \varphi^i(t,Z^i_t)}{\partial t} + \sup_{a^i\in\mathbb{R}}\Big(\frac{\partial \varphi^i(t,Z^i_t)}{\partial x} b_i^{(x)} + \frac{\partial \varphi^i(t,Z^i_t)}{\partial y} b_i^{(y)} + \frac{\partial^2 \varphi^i(t,Z^i_t)}{\partial x \partial y}(\nu_i^{(x)}\nu_i^{(y)} + \sigma_i^{(x)}\sigma_i^{(y)}) + \frac{1}{2}\frac{\partial^2 \varphi^i(t,Z^i_t)}{\partial x^2}(\nu_i^{(x)^2} + \sigma_i^{(x)^2}) + \frac{1}{2}\frac{\partial^2 \varphi^i(t,Z^i_t)}{\partial y^2}(\nu_i^{(y)^2} + \sigma_i^{(y)^2})\Big)\Big] \geq 0.$$

We know once more that $Z^i_t = (x, y)$ and thus we may get rid of the expectation as well as simplify notation:

$$
\begin{aligned}
\varphi_t^i + \sup_{a^i\in\mathbb{R}} \Big[&\varphi_x^i b_i^{(x)} + \varphi_y^i b_i^{(y)} + \varphi_{xy}^i(\nu_i^{(x)}\nu_i^{(y)} + \sigma_i^{(x)}\sigma_i^{(y)}) + \frac{1}{2}\varphi_{xx}^i(\nu_i^{(x)^2} + \sigma_i^{(x)^2}) \\
&+ \frac{1}{2}\varphi_{yy}^i(\nu^{(y)^2} + \sigma_i^{(y)^2})\Big] \geq 0
\end{aligned}
$$

We may combine this with the previous inequality to get the coupled HJB equation with equality:

$$
\begin{aligned}
\varphi_t^i + \sup_{a^i\in\mathbb{R}} \Big[&\varphi_x^i b_i^{(x)} + \varphi_y^i b_i^{(y)} + \varphi_{xy}^i(\nu_i^{(x)}\nu_i^{(y)} + \sigma_i^{(x)}\sigma_i^{(y)}) + \frac{1}{2}\varphi_{xx}^i(\nu_i^{(x)^2} + \sigma_i^{(x)^2}) \\
&+ \frac{1}{2}\varphi_{yy}^i(\nu_i^{(y)^2} + \sigma_i^{(y)^2})\Big] = 0
\end{aligned}
$$

□

3.3 Solution to Weighted N Player Game

3.3.1 Best Response Strategy

We first state the HJB equation derived in Theorem 3.2.1, plugging in the original coefficients of the X Itô process and Y Itô process as in Equations (3.4) and (3.6) where $\varphi^i(x,y,t)$ is the solution:

$$
\begin{aligned}
\varphi_t^i + \sup_{\pi^i\in\mathbb{R}} \Big[&\varphi_x^i(\pi^i\mu_i)x + \varphi_y^i(\eta_i)y + \varphi_{xy}^i(\pi^i x\nu_i\cdot 0 + \pi^i x\sigma_i y\widehat{\lambda_i\sigma\pi}_{-i}) + \frac{1}{2}\varphi_{xx}^i((\pi^i x\nu_i)^2 \\
&+ (\pi^i x\sigma_i)^2) + \frac{1}{2}\varphi_{yy}^i y^2(\frac{1}{n}(\widehat{\lambda_i\nu\pi})^2{}_{-i} + \widehat{\lambda_i\sigma\pi}^2_{-i})\Big] = 0
\end{aligned}
$$

with terminal condition $\varphi^i(x,y,T) = \epsilon_i U(x^{1-\theta_i/n} y^{-\theta_i}; \delta_i)$.[4] Rearranging and splitting apart the suprema:

$$\varphi_t^i + \sup_{\pi^i \in \mathbb{R}} \Big[\underbrace{\pi^i(\varphi_x^i \mu_i x + \varphi_{xy}^i xy\sigma_i \widehat{\lambda_i \sigma \pi}_{-i}) + \frac{1}{2}\varphi_{xx}^i {\pi^i}^2 x^2 \Sigma_i}_{\Lambda} \Big] + \varphi_y^i y(\eta_i)$$
$$+ \frac{1}{2}\varphi_{yy}^i y^2 (\frac{1}{n}(\widehat{\lambda_i \nu \pi})^2{}_{-i} + \widehat{\lambda_i \sigma \pi}^2_{-i}) = 0.$$

Now we may take the derivative of the first suprema to find the optimal π:

$$\frac{d\Lambda}{d\pi^i} = \varphi_x^i \mu_i x + \varphi_{xy}^i xy\sigma_i \widehat{\lambda_i \sigma \pi}_{-i} + \varphi_{xx}^i \pi^i x^2 \Sigma_i = 0.$$

This is solved by

$$\pi^{i,*} = \frac{-\mu_i x \varphi_x^i - \sigma_i \widehat{\lambda_i \sigma \pi}_{-i} xy \varphi_{xy}^i}{\Sigma_i x^2 \varphi_{xx}^i}. \tag{3.8}$$

We may plug this back in and the equation becomes:

$$\varphi_t^i - \frac{(\mu_i x \varphi_x^i + \sigma_i \widehat{\lambda_i \sigma \pi}_{-i} xy \varphi_{xy}^i)^2}{2\Sigma_i x^2 \varphi_{xx}^i} + \varphi_y y(\eta_i) + \frac{1}{2}\varphi_{yy}^i y^2 (\frac{1}{n}(\widehat{\lambda_i \nu \pi})^2{}_{-i} + \widehat{\lambda_i \sigma \pi}^2_{-i}) = 0. \tag{3.9}$$

For each of the two cases for the CRRA utility function, we make separate ansatzes and solve.

$\boldsymbol{\delta_i \neq 1}$: We make the following ansatz for the solution φ^i, for some function of t denoted by $f_i(t)$. Since the terminal condition at time T must be the expression as in Equation (3.3.1), we put $f_i(T) = 1$ as well:

$$\varphi^i(x,y,t) = \epsilon_i \left(1 - \frac{1}{\delta_i}\right)^{-1} x^{(1-\theta_i/n)(1-1/\delta_i)} y^{-\theta_i(1-1/\delta_i)} f_i(t). \tag{3.10}$$

[4]This suprema is finite if the function is concave which is clearly true.

We plug this back in and divide by $\epsilon_i \left(1-\frac{1}{\delta_i}\right)^{-1} x^{(1-\theta_i/n)(1-1/\delta_i)} y^{-\theta_i(1-1/\delta_i)}$. Note that we may do this since $x\varphi^i_x$, $y\varphi^i_y$, $x^2\varphi^i_{xx}$ and $y^2\varphi^i_{yy}$ all yield this expression with this constant in front of it (as well as an $f_i(t)$ which we do not divide out as it does not appear in φ_t). Thus we may rewrite the equation, replacing the constant expression with ρ_i defined as:

$$\rho_i = \left(1-\frac{1}{\delta_i}\right)\Bigg(\frac{(1-\theta_i/n)\left(\mu_i - \sigma_i\theta_i(1-1/\delta_i)\widehat{\lambda_i\sigma\pi}_{-i}\right)^2}{2\Sigma_i(1-(1-\theta_i/n)(1-1/\delta_i))} + \frac{1}{2}\Big((\widehat{\lambda_i\sigma\pi}_{-i})^2 + \widehat{\lambda_i\nu\pi}^2_{-i}\Big)\theta_i^2(1-1/\delta_i) - \theta_i\widehat{\lambda_i\mu\pi}_{-i} + \frac{\theta_i}{2}\widehat{\lambda_i\Sigma\pi^2}_{-i}\Bigg).$$

With this, we have the following equation:

$$0 = \left(1-\frac{1}{\delta_i}\right)^{-1} f_i'(t) + \rho_i f_i(t). \tag{3.11}$$

This differential equation is easily solved by $f_i(t) = e^{\rho_i\left(1-\frac{1}{\delta_i}\right)(T-t)}$. Plugging in and using the verification theorem, we have that the optimal control or best response of player i is:[5]

$$\pi^{i,*} = \frac{\delta_i\mu_i - \sigma_i\widehat{\lambda_i\sigma\pi}_{-i}\theta_i(\delta_i-1)}{(\sigma_i^2+\nu_i^2)(\delta_i-(1-\theta_i/n)(\delta_i-1))}. \tag{3.12}$$

$\boldsymbol{\delta_i = 1}$: With $\delta_i = 1$, the utility can be written as:

$$U\left((X^i_t)^{1-\theta_i/n}(Y^i_t)^{-\theta_i/n};\delta_i\right) = \log\left(x^{1-\theta_i/n}y^{-\theta_i/n}\right)$$

Now we make an ansatz with value function φ^i for some function of t denoted by $f_i(t)$, where we put $f_i(T) = 0$ to fulfill the boundary condition:

$$\varphi^i(x,y,t) = U\left(x^{1-\theta_i/n}y^{-\theta_i/n};\delta_i\right) + f_i(t) = \left(1-\frac{\theta_i}{n}\right)\log x - \theta_i \log y + f_i(t).$$

[5]This is a classical result in optimal control theory that guarantees that the function φ is the optimal value function for the problem and π is an optimal control if φ satisfies the HJB equations and meets the boundary conditions as well as if π optimizes the Hamiltonian.

Plugging this into the HJB equation in Equation (3.9), we get

$$f_i'(t) + \rho_i = 0$$

with

$$\rho_i := \frac{\mu_i^2(1 - \frac{\theta_i}{n})}{2(\sigma_i^2 + \nu_i^2)} - \theta_i\eta + \frac{1}{2}\theta_i\left(\widehat{\lambda_i\sigma\pi}^2_{-i} + \frac{1}{n}\widehat{(\lambda_i\nu\pi)^2}_{-i}\right).$$

Then clearly $f_i(t)$ is

$$f_i(t) = \rho_i(T - t)$$

and plugging φ^i back in, we get

$$\pi^{i,*} = \frac{\mu_i}{\Sigma_i}.$$

Again, we use the verification theorem to prove this is the optimal control and thus player i's best response.

3.3.2 Existence of Nash Equilibrium

To prove that the best response we found in Equation (3.12) is the Nash equilibrium $(\pi^{1,*}, \dots, \pi^{n,*})$, we now show that a fixed point of this mapping exists and is unique. We begin by stating the Banach Fixed Point Theorem, the main tool we will use to prove this result.

Theorem 3.3.1 (Banach Fixed Point Theorem)**.** *Let $\Psi : X \to X$ be a function on the complete metric space (X, d). If Ψ is a contraction mapping, that is, we have that $\forall x, x' \in X$,*

$$d(\Psi(x), \Psi(x')) \leq \epsilon d(x, x')$$

for $\epsilon < 1$, then Ψ admits a unique fixed point.

To use this theorem, we first define the complete metric space as $(\mathbb{R}^n, d)$ where for $x^{(1)}, x^{(2)} \in \mathbb{R}^n$, the Euclidean distance d is written as

$$d(x^{(1)}, x^{(2)}) = \sqrt{\sum_{i=1}^{n} \left(x_i^{(2)} - x_i^{(1)}\right)^2}.$$

Define the vector $\boldsymbol{\pi} = (\pi^1, \ldots, \pi^n) \in \mathbb{R}^n$. Then the function $\Psi : \mathbb{R}^n \to \mathbb{R}^n$ is the following: $\Psi(\boldsymbol{\pi}) = (\Psi(\boldsymbol{\pi})^1, \ldots, \Psi(\boldsymbol{\pi})^n)$ where each element is the best response derived in Equation (3.12):

$$\Psi(\boldsymbol{\pi})^i = \pi^i = \frac{\delta_i \mu_i - \frac{\sigma_i}{n} \sum_{j=1}^{n} (\lambda_{ij} \sigma_j \pi^j) \theta_i (\delta_i - 1)}{(\sigma_i^2 + \nu_i^2)(\delta_i - (1 - \theta_i/n)(\delta_i - 1))}.$$

Now, to show there exists a fixed point $\Psi(\boldsymbol{\pi}) = \boldsymbol{\pi}$, we prove that Ψ is a contraction mapping then invoke the Banach Fixed Point Theorem.

Theorem 3.3.2. *The function Ψ defined as above is a contraction mapping for sufficiently large n.*

Proof. We prove this by computing for any $\boldsymbol{\pi}, \boldsymbol{\pi} \in \mathbb{R}^n$:

$$d(\Psi(\boldsymbol{\pi}) - \Psi(\boldsymbol{\pi}'))^2 = \sum_{i=1}^{n} |\pi^i - \pi^{i\prime}|^2.$$

Plugging in, we have:

$$= \sum_{i=1}^{n} \left| \frac{\delta_i \mu_i - \frac{\sigma_i}{n} \sum_{j=1}^{n} (\lambda_{ij} \sigma_j \pi^j) \theta_i (\delta_i - 1)}{(\sigma_i^2 + \nu_i^2)(\delta_i - (1 - \theta_i/n)(\delta_i - 1))} - \frac{\delta_i \mu_i - \frac{\sigma_i}{n} \sum_{j=1}^{n} (\lambda_{ij} \sigma_j \pi'^j) \theta_i (\delta_i - 1)}{(\sigma_i^2 + \nu_i^2)(\delta_i - (1 - \theta_i/n)(\delta_i - 1))} \right|^2.$$

We may rearrange terms, noting that the only term that differs is that with π^j and $\pi^{j\prime}$:

$$= \sum_{i=1}^{n} \left(\left(\frac{\sigma_i \theta_i (\delta_i - 1)}{n(\sigma_i^2 + \nu_i^2)(\delta_i - (1 - \theta_i/n)(\delta_i - 1))} \right)^2 \left| \sum_{j=1}^{n} \lambda_{ij} \sigma_j (\pi^{j\prime} - \pi^j) \right|^2 \right).$$

By Cauchy Schwarz, we can bound the second term of the product:

$$= \sum_{i=1}^{n} \left(\left(\frac{\sigma_i \theta_i (\delta_i - 1)}{n(\sigma_i^2 + \nu_i^2)(\delta_i - (1 - \theta_i/n)(\delta_i - 1))} \right)^2 \sum_{j=1}^{n} \lambda_{ij}^2 \sigma_j^2 (\pi^{j\prime} - \pi^j)^2 \right).$$

Now note we have that $\lambda_{ij}^2 \sigma_j^2 \leq (\max_{j \in [n]} \sigma_j)^2$ since $0 \leq \lambda_{ij} \leq 1$. Denote C to be $(\max_{j \in [n]} \sigma_j)$. Then we know that, since $d(\boldsymbol{\pi}, \boldsymbol{\pi'})^2 = \sum_{j=1}^{n} (\pi^{j\prime} - \pi^j)^2$,

$$= d(\boldsymbol{\pi}, \boldsymbol{\pi'})^2 \sum_{i=1}^{n} \left(\frac{C \sigma_i \theta_i (\delta_i - 1)}{n(\sigma_i^2 + \nu_i^2)(\delta_i - (1 - \theta_i/n)(\delta_i - 1))} \right)^2.$$

Clearly, we may bound the above as required if

$$\epsilon = \sqrt{\frac{\left(\frac{1}{n} \sum_{i=1}^{n} \left(\frac{C \sigma_i \theta_i (\delta_i - 1)}{(\sigma_i^2 + \nu_i^2)(\delta_i - (1 - \theta_i/n)(\delta_i - 1))} \right)^2 \right)}{n}},$$

where the numerator is an average over bounded parameters for each agent and the denominator is $\sqrt{n}$. Thus as n grows, we see ϵ becomes small. For sufficiently large n, $\epsilon < 1$ and, hence, Ψ is a contraction mapping on $\mathbb{R}^n$. □

As a result of the above theorem, we may invoke Banach Fixed Point Theorem and say that for enough players in the n-player game, Ψ admits a unique fixed point, thereby proving the existence of such a Nash equilibrium.[6]

3.4 Discussion

In the logarithmic case, we see that the new investment strategy is exactly the same. This is an artifact of the logarithmic utility function which causes agents with such utility to not be competitive. In the general CRRA utility, the investment strategy is similar to that proved in [3] for the unweighted case presented earlier in Equation

[6]Note that this equilibrium is a constant Nash equilibrium since π is time independent.

(2.11) with the exception of the definitions of our constants. This is expected as the only change in our formulation is the inclusion of individualized weights which doesn't change the problem significantly and thus the optimal strategy would be similar in competing with weighted averages rather than the original averages. We discuss more about the properties of the optimal investment strategy in Section 5.3 since the n-player solution has similar structure to the graphon solution.

Chapter 4

Graphon Games

Now consider taking the number of players n to the infinite limit. We use a graphon game to represent this network of interactions. To do so, we first build out relevant theory for graphon games then proceed to find the Nash equilibrium under CRRA utility.

4.1 Background

We build out this theory, taking concepts from [9]. Rather than indexing as we did before, where we studied agent $i \in \{1, 2, \ldots, n\}$, we now represent this as a continuum with agent $u \in [0, 1]$. Hence, each strategy, rather than being represented as π^i, is represented as π^u. Moreover, for the market model, we have now a continuum of stocks for each agent and therefore a continuum of Brownian motions $(W^u)_{u\in[0,1]}$. We also have B that we had established previously as the Brownian motion for common noise. Thus, we write the following for each stock:

$$\frac{dS^u}{S^u} = \mu_u dt + \nu_u dW_t^u + \sigma_u dB_t$$

with constants μ_u, ν_u, and σ_u defined previously as drift rate and volatilities respectively for stock u. Player u's wealth is:

$$dX_t^u = \pi^u X_t^u(\mu_u dt + \nu_u dW_t^u + \sigma_u dB_t)$$

similar to what we defined before. Here, the difference lies in the total utility. Originally, the CRRA utility, without consumption, was defined as in Equation (3.3), rather now we define it as follows:

$$\sup_{\pi^u \in \mathcal{A}} \mathbb{E}\left[U\left(X_T^{x,\pi} Y_T^{x,\pi-\theta_u}\right)\right] \tag{4.1}$$

where

$$Y_t^u = \exp\left(\mathbb{E}\left[\int_0^1 G(u,v)\log(X_t^v)dv|\mathcal{F}_t^B\right]\right)$$

is the weighted geometric mean of the continuum of agents. Moreover, $G(u,v)$ is the continuous analogue to λ_{uv} or the weight player u places on player v. More generally, G is a graphon or a bounded, measurable, and symmetric function $G : [0,1]^2 \to [0,1]$ which measures the interaction among a continuum of agents. Now we proceed to define a graphon Nash equilibria.

Definition 4.1.1 (Graphon Nash Equilibria)**.** *A family of strategy profiles* $(\hat{\pi}^u)_{u\in I}$ *such that for* $u \in [0,1]$*:*

$$\mathbb{E}\left[U\left(X_T^{x,\hat{\pi}^u}\left(Y_T^{x,\hat{\pi}^v}\right)^{-\theta_u}\right)\right] = \max_{\pi^u \in \mathcal{A}} \mathbb{E}\left[U\left(X_T^{x,\pi^u}\left(Y_T^{x,\hat{\pi}^v}\right)^{-\theta_u}\right)\right].$$

This definition essentially means that should we fix a strategy profile for all players $v \neq u$, then the strategy player u picks is the one that maximizes their expected utility. In other words, there is no incentive for any player u to deviate from π^u as it yields the maximum expected reward. Now that we have defined all relevant definitions, we proceed to solve for the graphon Nash equilibrium.

4.2 Solution

4.2.1 Derivation of Ansatz

To approach this problem, we first derive the Itô process for Y_t^u. We know that Y_t^u is defined as the following:

$$Y_t^u = \exp\left(\mathbb{E}\left[\int_0^1 G(u,v)\log(X_t^v)dv \middle| \mathcal{F}_t^B\right]\right).$$

Thus we begin with the Itô process for X_t^v:

$$dX_t^v = \pi^v X_t^v(\mu_v dt + \nu_v dW_t^v + \sigma_v dB_t).$$

Using Itô's formula, we see that:

$$d\log(X_t^v) = \frac{1}{X_t^v}dX_t^v - \frac{1}{2}\pi^{v^2}(\nu_v^2 + \sigma_v^2)dt.$$

Computing $\int_0^1 G(u,v)d\log(X_t^v)dv$, we have

$$\int_0^1 G(u,v)d\log(X_t^v)dv = \int_0^1 G(u,v)\left(\frac{1}{X_t^v}dX_t^v - \frac{1}{2}\pi^{v^2}(\nu_v^2 + \sigma_v^2)dt\right)dv.$$

Expanding this, we get:

$$\begin{aligned}\int_0^1 G(u,v)d\log(X_t^v)dv = &\int_0^1 G(u,v)\left(\pi^v\mu_v - \frac{1}{2}\pi^{v^2}(\nu_v^2 + \sigma_v^2)\right)dtdv \\ &+ \int_0^1 G(u,v)\pi^v\nu_v dW_t^v dv + \int_0^1 G(u,v)\pi^v\sigma_v dB_t dv.\end{aligned}$$

Now taking the expectation conditioned on the filtration generated by B and finding $d\log Y_t^u = \mathbb{E}\left[\int_0^1 G(u,v)d\log(X_t^v)dv|\mathcal{F}_t^B\right]$, we have

$$\begin{aligned} d\log Y_t^u = \mathbb{E}\Bigg[& \int_0^1 G(u,v)\left(\pi^v\mu_v - \frac{1}{2}\pi^{v^2}(\nu_v^2+\sigma_v^2)\right)dtdv + \int_0^1 G(u,v)\pi^v\nu_v dW_t^v dv \\ & + \int_0^1 G(u,v)\pi^v\sigma_v dB_t dv\Bigg|\mathcal{F}_t^B\Bigg]. \end{aligned}$$

Here, the dependence on $\mathcal{F}_t^B$ arises from the investment strategy π^v. All other terms with individual parameters and W_t^v are independent of this filtration. Thus, the conditional expectation of the dW_t^v term is just the regular expectation which is 0 since the stochastic integral is a martingale vanishing in 0. Thus we have:

$$\begin{aligned} d\log Y_t^u = \mathbb{E}\Bigg[& \int_0^1 G(u,v)\left(\pi^v\mu_v - \frac{1}{2}\pi^{v^2}(\nu_v^2+\sigma_v^2)\right)dv\Bigg|\mathcal{F}_t^B\Bigg]dt \\ & + \mathbb{E}\Bigg[\int_0^1 G(u,v)\pi^v\sigma_v dv\Bigg|\mathcal{F}_t^B\Bigg]dB_t. \end{aligned}$$

Taking $\overline{\pi\mu}_u := \int_0^1 G(u,v)\mu_v\mathbb{E}[\pi^v|\mathcal{F}_t^B]dv$, $\overline{\pi\sigma}_u := \int_0^1 G(u,v)\sigma_v\mathbb{E}[\pi^v|\mathcal{F}_t^B]dv$, $\overline{\pi^2\nu^2}_u = \int_0^1 G(u,v)\nu_v^2\mathbb{E}[\pi^{v^2}|\mathcal{F}_t^B]dv$, and $\overline{\pi^2\sigma^2}_u = \int_0^1 G(u,v)\sigma_v^2\mathbb{E}[\pi^{v^2}|\mathcal{F}_t^B]$, we may rewrite this as:

$$d\log Y_t^u = (\overline{\pi\mu}_u - \frac{1}{2}(\overline{\pi^2\nu^2}_u + \overline{\pi^2\sigma^2}_u))dt + \overline{\pi\sigma}_u dB_t.$$

Using Itô's formula again to transform this process to the desired dynamics with the exponential function, we get the following:

$$dY_t^u = Y_t^u(\gamma_u dt + \overline{\pi\sigma}_u dB_t),$$

where $\gamma_u = \overline{\pi\mu}_u - \frac{1}{2}(\overline{\pi^2\nu^2}_u + \overline{\pi^2\sigma^2}_u)$. Now we treat (X_t^u, Y_t^u) as a state process and use the coupled HJB equations derived in Theorem 3.2.1. While the value function is $\varphi^u(t,x,y)$ for each player u, we drop this superscript and (t,x,y) for notational simplicity and we also use the notation $\varphi_t = \frac{\partial \varphi^u}{\partial t}$ for all partial derivatives:

$$0 = \varphi_t + \sup_{\pi^u \in \mathbb{R}} \Big[\varphi_x \pi^u x \mu_u + \varphi_y y (\overline{\pi\mu}_u - \frac{1}{2}(\overline{\pi^2\nu^2}_u + \overline{\pi^2\sigma^2}_u) + \frac{1}{2}\overline{\pi\sigma}_u^2) \\ + xy\pi^u \varphi_{xy}(\sigma_u \overline{\pi\sigma}_u) + x^2 \varphi_{xx} \pi^{u^2} \frac{1}{2}\left(\nu_u^2 + \sigma_u^2\right) + y^2 \varphi_{yy} \frac{1}{2}\left(\overline{\pi\sigma}_u^2\right) \Big].$$

In turn, we may make the following ansatz for the solution φ^u, for some function of t denoted by $f_u(t)$. Since the terminal condition at time T must be $\varphi^u(x, y, T) = U_u\left((xy)^{-\theta_u}; \delta_u\right)$, we put $f_u(T) = 1$ as well:

$$\varphi^u(x, y, t) = \left(1 - \frac{1}{\delta_u}\right)^{-1} x^{(1-1/\delta_u)} y^{-\theta_u(1-1/\delta_u)} f_u(t). \tag{4.2}$$

We plug this ansatz in (dropping the superscript u for notational simplicity) and divide by $\left(1 - \frac{1}{\delta_u}\right)^{-1} x^{(1-1/\delta_u)} y^{-\theta_u(1-1/\delta_u)}$ to get the following:

$$0 = f_u'(t) + f_u(t) \Bigg(\sup_{\pi^u \in \mathbb{R}} \Big[\pi^u (1 - \frac{1}{\delta_u})(\mu_u - \theta_u(1 - \frac{1}{\delta_u})(\sigma_u \overline{\pi\sigma}_u)) \\ - \frac{1}{2\delta_u}(1 - \frac{1}{\delta_u})\pi^{u^2}\left(\nu_u^2 + \sigma_u^2\right) - \theta_u(1 - \frac{1}{\delta_u})(\overline{\pi\mu}_u - \frac{1}{2}(\overline{\pi^2\nu^2}_u + \overline{\pi^2\sigma^2}_u) \\ + \frac{1}{2}\overline{\pi\sigma}_u^2) - \frac{1}{2}\theta_u(1 - \frac{1}{\delta_u})(-\theta_u + \frac{\theta_u}{\delta_u} - 1)\left(\overline{\pi\sigma}_u^2\right) \Big].$$

Then clearly the solution to this is:

$$f_u(t) = \exp\Bigg(- \sup_{\pi^u \in \mathbb{R}} \Big[\pi^u (1 - \frac{1}{\delta_u})(\mu_u - \theta_u(1 - \frac{1}{\delta_u})(\sigma_u \overline{\pi\sigma}_u)) \\ - \frac{1}{2\delta_u}(1 - \frac{1}{\delta_u})\pi^{u^2}\left(\nu_u^2 + \sigma_u^2\right) - \theta_u(1 - \frac{1}{\delta_u})(\overline{\pi\mu}_u - \frac{1}{2}(\overline{\pi^2\nu^2}_u + \overline{\pi^2\sigma^2}_u) \\ + \frac{1}{2}\overline{\pi\sigma}_u^2) - \frac{1}{2}\theta_u(1 - \frac{1}{\delta_u})(-\theta_u + \frac{\theta_u}{\delta_u} - 1)\left(\overline{\pi\sigma}_u^2\right) \Big](T - t)\Bigg). \tag{4.3}$$

4.2.2 Continuous Graphon Solution

Note that we cannot directly take the derivative with respect to π^u in the above expression as each of the integrals contain π^u. As a result, we assume Riemann integrability and write each integral as the limit of a Riemann sum. This assumption is true under the conditions outlined in the following theorem.

Theorem 4.2.1 (Riemann-Lebesgue Theorem)**.** *A function f is Riemann integrable if and only if f is bounded and the set of discontinuities of f has measure* 0.

This holds as we are given G is bounded, and we can bound the deterministic constants $\{\mu_u\}_{u\in[0,1]}$, $\{\sigma_u\}_{u\in[0,1]}$ and $\{\nu_u\}_{u\in[0,1]}$. Further, should we restrict G to be a graphon that is continuous a.e., then the integrals in Equation (4.3) would be Riemann integrable.[1]

Then, we may write Equation (4.3) as the following, where $\widehat{\pi\sigma}_u^N = \frac{1}{N}\sum_{i=1}^N G(u,v_i)$ $\mathbb{E}[\pi^{v_i}|\mathcal{F}_t^B]\sigma_{v_i}$ and we define other Riemann sums using similar notation with v_i as the partition point:

$$\begin{aligned} f_u(t) = \exp\Bigg(&- (1-\frac{1}{\delta_u}) \sup_{\pi^u\in\mathbb{R}} \lim_{N\to\infty} \Bigg[\pi^u(\mu_u - \theta_u(1-\frac{1}{\delta_u})(\sigma_u \widehat{\pi\sigma}_u^N)) \\ &- \frac{1}{2\delta_u}\pi^{u^2}\left(\nu_u^2+\sigma_u^2\right) - \theta_u(\widehat{\pi\mu}_u^N - \frac{1}{2}(\widehat{\pi^2\nu^2}_u^N + \widehat{\pi^2\sigma^2}_u^N) \\ &+ \frac{1}{2}\widehat{\pi\sigma}_u^{N^2}) - \frac{1}{2}\theta_u(-\theta_u + \frac{\theta_u}{\delta_u} - 1)\left(\widehat{\pi\sigma}_u^{N^2}\right)(T-t)\Bigg]\Bigg). \end{aligned}$$

In order to derive an explicit form for the solution, we must swap the supremum and limit which can be done should the strategies be restricted to a compact set since the other parameters are bounded. Thus we restrict π^u such that $\pi^u \in \mathcal{A}$ where $\mathcal{A} \subseteq \mathbb{R}$

[1] This is somewhat of a restrictive assumption on the weights players can place on one another. While we could consider many nontrivial graphon examples like $G(u,v) = uv$ or $\max(u,v)$, we work to extend this in Section 4.2.3.

is a compact set. Then we may write the following:

$$f_u(t) = \exp\Bigg(-(1-\frac{1}{\delta_u}) \lim_{N\to\infty} \sup_{\pi^u \in \mathbb{R}} \Bigg[\pi^u(\mu_u - \theta_u(1-\frac{1}{\delta_u})(\sigma_u \widehat{\pi\sigma}_u^N)) - \frac{1}{2\delta_u}\pi^{u^2}\left(\nu_u^2+\sigma_u^2\right) - \theta_u(\widehat{\pi\mu}_u^N - \frac{1}{2}\widehat{\pi^2\nu^2}_u^N + \widehat{\pi^2\sigma^2}_u^N + \frac{1}{2}\widehat{\pi\sigma}_u^{N^2}) - \frac{1}{2}\theta_u(-\theta_u + \frac{\theta_u}{\delta_u} - 1)\left(\widehat{\pi\sigma}_u^{N^2}\right) \Bigg](T-t)\Bigg).$$

Now, to take the supremum, note that for all of the above Riemann sums, we can take the derivative with respect to π^u (with the respective deterministic functions multiplied into this sum):

$$\begin{aligned} \frac{d}{d\pi^u}\left[\frac{1}{N}\sum_{i=1}^N G(u,v_i)\mathbb{E}[\pi^{v_i}|\mathcal{F}_t^B]\right] &= \frac{d}{d\pi^u}\left[\frac{1}{N}\sum_{i=1}^N G(u,v_i)\pi^{v_i}\right] \\ &= \frac{1}{N}\frac{d}{d\pi^u}\left(\sum_{i=1}^N G(u,v_i)\pi^{v_i}\mathbb{1}_{v^i=u}\right) \\ &= \frac{1}{N}\frac{d}{d\pi^u}G(u,u)\pi^u = 0. \end{aligned}$$

The first equality follows from that we pick π^{v_i} to be deterministic when taking the supremum and thus $\mathbb{E}[\pi^{v_i}|\mathcal{F}_t^B] = \pi^{v_i}$.[2] Moreover, the last equality follows from the assumption that player u puts 0 weight on themselves or that $G(u,u) = 0$. Taking the derivative with respect to π^u of the supremum expression and finding critical points yields:

$$\mu_u - \theta_u(1-\frac{1}{\delta_u})(\sigma_u\widehat{\pi\sigma}_u^N) - \frac{1}{\delta_u}\pi^u\left(\nu_u^2+\sigma_u^2\right) = 0.$$

Thus the π^u that satisfies such is:

$$\pi^u = \frac{\delta_u\mu_u - \theta_u(\delta_u - 1)(\sigma_u\widehat{\pi\sigma}_u^N)}{\nu_u^2+\sigma_u^2}. \tag{4.4}$$

[2]From here on out, thus, the constants defined earlier are deterministic. One example of this being $\overline{\pi\mu}_u := \int_0^1 G(u,v)\mu_v\mathbb{E}[\pi^v|\mathcal{F}_t^B]dv$ which is now exactly equal to $\int_0^1 G(u,v)\mu_v\pi^v dv$.

Using this expression to find an explicit form of the ansatz in Equation (4.3), we see

$$f_u(t) = \exp\Bigg(-(1-\frac{1}{\delta_u}) \lim_{N\to\infty} \Bigg(\frac{\Big(\mu_u - \theta_u(\sigma_u \widehat{\pi\sigma}_u^N)\Big)^2 \delta_u}{-2\left(\nu_u^2+\sigma_u^2\right)} - \frac{\theta_u}{2} \Big((2\widehat{\pi\mu}_u^N - \widehat{\pi^2\nu^2}_u^N - \widehat{\pi^2\sigma^2}_u^N \\ + \widehat{\pi\sigma}_u^{N^2}) + (-\theta_u + \frac{\theta_u}{\delta_u} - 1)\left(\widehat{\pi\sigma}_u^{N^2}\right)\Big)\Bigg).$$

From here, we see that all Riemann sums have been preserved thus we take the limit and see convergence to the respective Riemann integral. Further, we know that the Riemann integral and Lebesgue integral coincide here due to our prior assumptions; therefore using our shorthand for the Lebesgue integrals, we obtain

$$\xi := \Bigg(\frac{\Big(\mu_u - \theta_u(1-\frac{1}{\delta_u})(\sigma_u\overline{\pi\sigma}_u)\Big)^2 (\delta_u - 1)}{-2\left(\nu_u^2+\sigma_u^2\right)} - \frac{\theta_u}{2}(1-\frac{1}{\delta_u}) \Big((2\overline{\pi\mu}_u - \overline{\pi^2\nu^2}_u - \overline{\pi^2\sigma^2}_u \\ + \overline{\pi\sigma}_u^2) + (-\theta_u + \frac{\theta_u}{\delta_u} - 1)\left(\overline{\pi\sigma}_u^2\right)\Big)\Bigg),$$

and hence the final function is

$$f_u(t) = e^{\xi(T-t)}.$$

This corresponds to $\pi^{u^*} = \lim_{N\to\infty} \pi^u$ where π^u is as in (4.4). This is exactly:

$$\pi^{u^*} = \frac{\delta_u\mu_u - \theta_u(\delta_u - 1)(\sigma_u\overline{\pi\sigma}_u)}{\nu_u^2+\sigma_u^2}. \tag{4.5}$$

Normally, we would use the verification theorem to claim that the π^{u^*} we solved for is the optimal control; however, this is not so obvious with the convergence argument we use. Thus, we prove that this is the optimal control alternatively by first denoting the Hamiltonian as

$$\begin{aligned}H(\pi^u) :=&\pi^u(1-\frac{1}{\delta_u})(\mu_u - \theta_u(1-\frac{1}{\delta_u})(\sigma_u\overline{\pi\sigma}_u)) \\ &-\frac{1}{2\delta_u}(1-\frac{1}{\delta_u})\pi^{u2}\left(\nu_u^2+\sigma_u^2\right) - \theta_u(1-\frac{1}{\delta_u})(\overline{\pi\mu}_u - \frac{1}{2}(\overline{\pi^2\nu^2}_u + \overline{\pi^2\sigma^2}_u) \\ &+\frac{1}{2}\overline{\pi\sigma}_u^2) - \frac{1}{2}\theta_u(1-\frac{1}{\delta_u})(-\theta_u + \frac{\theta_u}{\delta_u} - 1)\left(\overline{\pi\sigma}_u^2\right).\end{aligned}$$

For any other strategy $\widehat{\pi^u} \in \mathcal{A}$, we derive:

$$H(\widehat{\pi^u}) = \lim_{N\to\infty} H^N(\widehat{\pi^u})$$

where H^N is the Riemann sum approximation of the integrals. Moreover, we claim that since π^u optimizes H^N

$$\begin{aligned}\lim_{N\to\infty} H^N(\widehat{\pi^u}) &\leq \lim_{N\to\infty} H^N(\pi^u) \\ &= H(\pi^{u^*}).\end{aligned}$$

Therefore, we have that $H(\widehat{\pi^u}) \leq H(\pi^u)$, $\forall \widehat{\pi^u} \in \mathcal{A}$ and thus is the optimal control.

4.2.3 General Graphon Solution

In this section, we derive a solution for the general graphon case to broaden the prior result for continuous graphons. The solution we find should intuitively match up with that for the continuous case. To derive this general result, we approximate the graphon with continuous functions then use the same Riemann approach for each continuous function. Once we have strategies for each continuous function, we analyze convergence and prove that this final strategy is the optimal control. In order to approximate G with continuous functions, we must first understand the following well-known result.

Theorem 4.2.2 (Lusin's Theorem)**.** *Let $(X, \mathcal{M}, \mu)$ be a measure space, and let $G : X \to \mathbb{R}$ be a bounded, measurable function. For every $\epsilon > 0$, there exists a compact set $E \subseteq X$ and a continuous function $F : X \to \mathbb{R}$ such that*

1. *$G(x) = F(x)$ for all $x \in E$,*

2. *$\mu(X \setminus E) < \epsilon$.*

This theorem implies that within any given level of approximation $\epsilon > 0$, one can find a continuous function F that approximates the measurable function G closely on a large portion of the domain, with the measure of the set where G and F differ being less than ϵ.

For each $k \in \mathbb{N}$, we choose $\epsilon_k = 1/k$ and then apply Lusin's Theorem to obtain the continuous function G_k and a compact set E_k with the above properties, mainly that $G = G_k$ on E_k. Note that as $k \to \infty$, $\mu(X \setminus E) < 1/k$ decreases to 0 thus the set where G_k and G differ decreases to being of measure 0. Therefore, as $k \to \infty$, we have $G_k(x) \to G(x)$ almost everywhere.

We analyze the function $f_u(t)$ as in Equation (4.3) for each G_k. Denote $\overline{\pi\sigma}_u^k := \int_0^1 G_k(u,v)\pi^v\sigma_v dv$ and we define the other approximations using similar notation. We know then the approximate ansatz for each G_k is the following:

$$\begin{aligned} f_u^k(t) = \exp\Bigg(- \sup_{\pi^u \in \mathbb{R}} \Bigg[& \pi^u(1 - \frac{1}{\delta_u})(\mu_u - \theta_u(1 - \frac{1}{\delta_u})(\sigma_u \overline{\pi\sigma}_u^k)) \\ & - \frac{1}{2\delta_u}(1 - \frac{1}{\delta_u})\pi^{u2}\left(\nu_u^2 + \sigma_u^2\right) - \theta_u(1 - \frac{1}{\delta_u})(\overline{\pi\mu}_u^k - \frac{1}{2}(\overline{\pi^2\nu^2}_u^k + \overline{\pi^2\sigma^2}_u^k) \\ & + \frac{1}{2}\overline{\pi\sigma}_u^{k^2}) - \frac{1}{2}\theta_u(1 - \frac{1}{\delta_u})(-\theta_u + \frac{\theta_u}{\delta_u} - 1)\left(\overline{\pi\sigma}_u^{k^2}\right)\Bigg](T - t)\Bigg). \end{aligned} \tag{4.6}$$

Since each G_k is continuous, we may proceed as in Section 4.2.2, by first writing the integral as a Riemann sum, then finding the optimal control. We use that we may interchange G_k with G as the continuous result we had derived holds for any graphon G and, hence, the optimal control corresponding to G_k is as in Equation

(4.4), which we denote as π_k^u in terms of Riemann sums. As we saw previously, these Riemann sums converge to the Riemann integrals (and hence Lebesgue integrals by our assumption) thus we derive Equation (4.5), which we write in terms of G_k here:

$$\pi_k^{u^*} = \frac{\delta_u \mu_u - \theta_u(\delta_u - 1)(\sigma_u \overline{\pi\sigma}_u^k)}{\nu_u^2 + \sigma_u^2}.$$

This converges in limit as $k \to \infty$ to the below since $G_k(x) \to G(x)$ a.e..

$$\pi^{u^*} = \frac{\delta_u \mu_u - \theta_u(\delta_u - 1)(\sigma_u \overline{\pi\sigma}_u)}{\nu_u^2 + \sigma_u^2}.$$

Note that this is the same expression as the continuous case. We must simply show now that this is indeed the optimal control. We prove such by first denoting the Hamiltonian as we did above:

$$\begin{aligned} H(\pi^u) :=& \pi^u(1 - \frac{1}{\delta_u})(\mu_u - \theta_u(1 - \frac{1}{\delta_u})(+\sigma_u \overline{\pi\sigma}_u)) \\ & - \frac{1}{2\delta_u}(1 - \frac{1}{\delta_u})\pi^{u2}\left(\nu_u^2 + \sigma_u^2\right) - \theta_u(1 - \frac{1}{\delta_u})(\overline{\pi\mu}_u - \frac{1}{2}(\overline{\pi^2\nu^2}_u + \overline{\pi^2\sigma^2}_u) \\ & + \frac{1}{2}\overline{\pi\sigma}_u^2) - \frac{1}{2}\theta_u(1 - \frac{1}{\delta_u})(-\theta_u + \frac{\theta_u}{\delta_u} - 1)\left(\overline{\pi\sigma}_u^2\right). \end{aligned}$$

For any other strategy $\widehat{\pi^u} \in \mathcal{A}$, we know that:

$$H(\widehat{\pi^u}) = \lim_{k\to\infty} \lim_{N\to\infty} H_k^N(\widehat{\pi^u})$$

where H_k^N is the Riemann representation of the continuous approximation (i.e. with functions G_k) of the Hamiltonian. Moreover, we claim that since π_k^u optimizes H_k^N:

$$\begin{aligned} \lim_{k\to\infty} \lim_{N\to\infty} H_k^N(\widehat{\pi^u}) &\leq \lim_{k\to\infty} \lim_{N\to\infty} H_k^N(\pi_k^u) \\ &= H(\pi^{u^*}). \end{aligned}$$

Therefore we have that $H(\widehat{\pi^u}) \leq H(\pi^{u^*})$, for all $\widehat{\pi^u} \in \mathcal{A}$ and we conclude that π^{u^*} is the optimal control. Rewriting the optimal control without shorthand notation:

$$\pi^{u^*} = \frac{\delta_u \mu_u - \theta_u(\delta_u - 1)(\sigma_u \int_0^1 G(u,v)\pi^v \sigma_v dv)}{\nu_u^2 + \sigma_u^2}. \tag{4.7}$$

4.2.4 Existence of a Nash Equilibrium

To prove that the best response we found in Equation (4.7) is the Nash equilibrium, we show, as we did in the n-player game, that a fixed point of this mapping exists and is unique. We again use the Banach Fixed Point Theorem stated in Theorem 3.3.1, first proving the necessary conditions.

We first define a complete metric space $(\mathcal{L}^2([0,1]), d)$ where $\mathcal{L}^2[0,1]$ denotes the set of square-integrable functions on $[0,1]$. Moreover, for $f, g \in \mathcal{L}^2([0,1])$, the ℓ_2 distance d is written as

$$d(f,g) = \sqrt{\int_0^1 (f(x) - g(x))^2 \, dx}.$$

Now let $\pi : [0,1] \to \mathbb{R}$ be a function in $\mathcal{L}^2[0,1]$. We know the function $\Psi : \mathcal{L}^2([0,1]) \to \mathcal{L}^2([0,1])$ is the following for each $u \in [0,1]$ due to the best response derived in Equation (4.7):

$$\Psi(\boldsymbol{\pi})^u = \pi^u = \frac{\delta_u \mu_u - \theta_u(\delta_u - 1)(\sigma_u \int_0^1 G(u,v)\pi^v \sigma_v dv)}{\nu_u^2 + \sigma_u^2}.$$

To show there exists a fixed point $\Psi(\boldsymbol{\pi}) = \boldsymbol{\pi}$, we prove that Ψ is a contraction mapping then invoke the Banach Fixed Point Theorem.

Theorem 4.2.3. *The function Ψ defined as above is a contraction mapping for*

$$\sqrt{\int_0^1 \left(\left(\frac{\theta_u(\delta_u - 1)}{\nu_u^2 + \sigma_u^2} \right)^2 \left(\int_0^1 G(u,v)^2 (\sigma_u \sigma_v)^2 dv \right) \right) du} < 1.$$

Proof. We prove this by computing, for any $\boldsymbol{\pi}, \boldsymbol{\pi}' \in \mathcal{L}^2[0,1]$:

$$d(\Psi(\boldsymbol{\pi}), \Psi(\boldsymbol{\pi}'))^2 = \int_0^1 |\pi^u - \pi^{u\prime}|^2 du.$$

We see that the $\nu_u \mu_u$ terms cancel out and we may factor out the common terms to get:

$$= \int_0^1 \left(\frac{\theta_u(\delta_u - 1)}{\nu_u^2 + \sigma_u^2} \right)^2 \left| \sigma_u \int_0^1 G(u,v)(\pi^v - \pi^{v\prime}) \sigma_v dv \right|^2 du.$$

Factoring again, we have

$$= \int_0^1 \left(\frac{\theta_u(\delta_u - 1)}{\nu_u^2 + \sigma_u^2} \right)^2 \left| \int_0^1 G(u,v)(\pi^v - \pi^{v\prime})(\sigma_u \sigma_v) dv \right|^2 du.$$

Now, we may use Cauchy Schwarz on the integral squared in order to get:

$$= \int_0^1 \left(\frac{\theta_u(\delta_u - 1)}{\nu_u^2 + \sigma_u^2} \right)^2 \left(\int_0^1 G(u,v)^2 (\sigma_u \sigma_v)^2 dv \right) \left(\int_0^1 (\pi^v - \pi^{v\prime})^2 dv \right) du.$$

By definition of the ℓ_2 distance, this is

$$= d(\boldsymbol{\pi}, \boldsymbol{\pi}')^2 \int_0^1 \left(\frac{\theta_u(\delta_u - 1)}{\nu_u^2 + \sigma_u^2} \right)^2 \left(\int_0^1 G(u,v)^2 (\sigma_u \sigma_v)^2 dv \right) du.$$

Hence, we deduce that we must have

$$\epsilon = \sqrt{\int_0^1 \left(\left(\frac{\theta_u(\delta_u - 1)}{\nu_u^2 + \sigma_u^2} \right)^2 \left(\int_0^1 G(u,v)^2 (\sigma_u \sigma_v)^2 dv \right) \right) du} < 1$$

in order for Ψ to be a contraction mapping. □

As a result of the above theorem, we may invoke Banach Fixed Point Theorem and conclude that for $\epsilon < 1$, Ψ admits a unique fixed point, thereby proving the existence of a Nash equilibrium.

4.3 Stability Analysis

Informally, stability means that if the graphons differ by a small amount then the resulting strategies differ by a small amount in expectation. To show stability of the solution, we want to first define such a metric that captures the distance between two graphons.

Definition 4.3.1 (Cut Norm). *For graphon G, the cut norm is defined as*

$$\|G\|_{\square} := \sup_{I} \left| \int_{I \times I} G(u,v)\, du\, dv \right|.$$

Moreover, we define the following norm which is equivalent to the cut norm:

Definition 4.3.2 ($L^{\infty} \to L^{1}$ operator norm)**.**

$$\|G\|_{\infty \to 1} := \sup\{\|G\varphi\|_{L^1[0,1]} : \varphi \in L^{\infty}[0,1], |\varphi| \leq 1\},$$

where $G\varphi(u) := \int_0^1 G(u,v)\varphi(v)\, dv$.

Furthermore, we know that the following relation holds by [4]:

$$\|g\|_{\square} \leq \|g\|_{\infty \to 1} \leq 4\|g\|_{\square} \tag{4.8}$$

which we will use later in our proof. Now that we have the above facts, we proceed to prove stability of the solution, defined formally in the theorem below.

Theorem 4.3.3. *The equilibrium strategy is stable. That is, if for some $\epsilon > 0$, where $\|G_1 - G_2\|_{\square} \leq \epsilon$ then for the corresponding π_1, π_2, there exists $\delta_{bound} > 0$ such that*

$$\mathbb{E}\|\pi^1 - \pi^2\|_2^2 \leq \delta_{bound}$$

where

$$\pi^1 = \frac{\delta_u \mu_u - \theta_u(\delta_u - 1)(\sigma_u \int_0^1 G_1(u,v)\pi^v \sigma_v dv)}{\nu_u^2 + \sigma_u^2}$$

and

$$\pi^2 = \frac{\delta_u \mu_u - \theta_u(\delta_u - 1)(\sigma_u \int_0^1 G_2(u,v)\pi^v \sigma_v dv)}{\nu_u^2 + \sigma_u^2}.$$

Proof. By simple math,

$$\mathbb{E}\|\pi^1 - \pi^2\|_2^2 = \left\| \frac{(-\theta_u(\delta_u - 1))}{\nu_u^2 + \sigma_u^2} \left(\sigma_u \int_0^1 (G_1(u,v) - G_2(u,v))\pi^v \sigma_v dv \right) \right\|_2^2.$$

Let $\Delta I = \sigma_u \int_0^1 (G_1(u,v) - G_2(u,v))\pi^v \sigma_v dv$. Then,

$$\mathbb{E}\|\pi^1 - \pi^2\|_2^2 = \left(\frac{(-\theta_u(\delta_u - 1))}{\nu_u^2 + \sigma_u^2} \right)^2 \|\Delta I\|_2^2. \tag{4.9}$$

We proceed to bound the ΔI by first bounding the absolute value:

$$|\Delta I| = \left| \sigma_u \int_0^1 (G_1(u,v) - G_2(u,v))\pi^v \sigma_v dv \right|$$

which is the following since $\sigma_u \geq 0$:

$$|\Delta I| = \sigma_u \left| \int_0^1 (G_1(u,v) - G_2(u,v))\pi^v \sigma_v dv \right|.$$

We may normalize $\pi^v \sigma_v$ by $\int_0^1 \pi^v \sigma_v dv$ since π^v and σ_v are nonnegative:

$$|\Delta I| = \sigma_u \int_0^1 \pi^v \sigma_v dv \left| \int_0^1 (G_1(u,v) - G_2(u,v)) \frac{\pi^v \sigma_v}{\int_0^1 \pi^v \sigma_v dv} dv \right|.$$

Now we may define the function $\varphi_1(v) := \frac{\pi^v \sigma_v}{\int_0^1 \pi^v \sigma_v dv}$. Note that $\varphi_1 \in L^\infty[0,1]$ since π^v and σ_v are bounded by prior assumption and $|\varphi_1| \leq 1$ through our normalization.

Then, we may upper bound the above quantity by the suprema over these:

$$|\Delta I| \leq \sigma_u \left(\int_0^1 \pi^v \sigma_v dv \right) \sup_{\varphi_2} \left| \int_0^1 (G_1(u,v) - G_2(u,v))\varphi_1(v) dv \right|.$$

This is equivalent to the $L^\infty \to L^1$ operator norm we defined in Equation (4.3.2). Hence, we rewrite as such:

$$|\Delta I| \leq \sigma_u \left(\int_0^1 \pi^v \sigma_v dv \right) \|G_1 - G_2\|_{\infty \to 1}.$$

We group terms together and bound this quantity by the inequality relating the cut norm to the $L^\infty \to L^1$ operator norm in Equation (4.8):

$$|\Delta I| \leq (4\|G_1 - G_2\|_\square) \left(\sigma_u \left(\int_0^1 \pi^v \sigma_v dv \right) \right).$$

We know that both π^v and σ_v are nonnegative and pointwise bounded thus finite at every point. Since we are integrating over a finite region, the integral of the product of any two of these is bounded. Therefore, let $C = \sigma_u \left(\int_0^1 \pi^v \sigma_v dv \right)$ for constant $C > 0$. We also know that by assumption $\|G_1 - G_2\|_\square \leq \epsilon$ and so we have

$$|\Delta I| \leq 4\epsilon C.$$

We now use this to bound $\|\Delta I\|_2^2$. In order to bound this term, first note that by definition of the L^1 norm, we have:

$$\|\Delta I\|_1 = \int_0^1 |\Delta I| du$$

and so we may compute a bound for the L^1 norm as:

$$\|\Delta I\|_1 \leq \int_0^1 4\epsilon C du = 4\epsilon C.$$

We also know that $\|\Delta I\|_2 \leq \|\Delta I\|_1$.[3] Therefore, we have

$$\|\Delta I\|_2 \leq 4\epsilon C.$$

Plugging this into Equation (4.9):

$$\mathbb{E}\|\pi^1 - \pi^2\|_2^2 \leq \left(\frac{(-\theta_u(\delta_u - 1))}{\nu_u^2 + \sigma_u^2}\right)^2 (4\epsilon C)$$

where we may set $\delta_{\text{bound}} := \left(\frac{(-\theta_u(\delta_u-1))}{\nu_u^2+\sigma_u^2}\right)^2 (4\epsilon C)$. Hence, we conclude our solution is stable. $\square$

[3] This is easy to prove by Cauchy-Schwarz.

Chapter 5

Simulations

To understand the behavior of our Nash Equilibrium strategy, we simulate by fixing parameters while changing others. We do this specifically in the graphon case as this has a smooth equilibrium over the players $u \in [0, 1]$ as opposed to the discrete strategies posed by the n-player game.

5.1 Methodology

In order to simulate, we must first find π^u which is not obvious from Equation (4.7). This is because we did not find an explicit form rather showed there was a fixed point to this mapping under some condition. To represent π^u, we would thus need to numerically find a fixed point of the mapping. The algorithm to find such is in Algorithm 1 where we use the formula we found for $\pi^{u,*}$, plugging in the old value of π^u into the right hand side of the equation and solving to find a new value of π^u. We keep iterating until there is convergence.

Finding the expected terminal wealth is more complex – we generate 5000 paths of the process X_t and average the terminal wealths X_T to get $\mathbb{E}[X_T]$. To generate these paths, we first set dt to be the number of time steps in a year, here 1/250 for the number of trading days per year. We also set N to be the total horizon T divided

Algorithm 1 Fixed Point Calculation

```
function FixedPoint(μ_u, θ_u, δ_u, σ_u, ν_u, σ_v, ν_v)
    pi_old ← 1.0, max_it ← 1000, eps ← 10^-3
    for k ← 1 to max_it do
        int ← ∫_0^1 pi_old × G(u, v) × σ_v dv
        pi_new ← (δ_u × μ_u − θ_u × (δ_u − 1) × (σ_u × int))/(ν_u^2 + σ_u^2)
        if |pi_old − pi_new| < eps then
            return pi_new
        end if
        pi_old ← pi_new
    end for
    return pi_old
end function
```

by this dt and thus N is the total number of time steps we simulate over. Now, to generate a path for X_t, for each of these N time steps, we generate two random samples, dW_t and dB_t, from the normal distribution with mean 0 and variance dt. We then update X_t by adding $X_t\pi_u(\mu_u dt + \nu_u dW_t + \sigma_u dB_t)$ where we found π_u in our fixed point calculations above.

5.2 Mean Field Game Simulation

To simulate the MFG, we know that each agent interacts with other agents the same, i.e. $G(u, v) = 1$. The symmetry of the game also means that we do not have to find a fixed point numerically, rather we may compute the solution analytically by first using Equation (4.7) with $G(u, v) = 1$:

$$\pi^u = \frac{\delta_u\mu_u - \theta_u(\delta_u - 1)(\sigma_u \int_0^1 \pi^v \sigma_v dv)}{\nu_u^2 + \sigma_u^2}. \tag{5.1}$$

We multiply this expression by σ_u and taking the integral over u and get

$$\int_0^1 \pi^u \sigma_u du = \int_0^1 \frac{\delta_u \sigma_u \mu_u - \sigma_u \theta_u (\delta_u - 1)(\sigma_u \int_0^1 \pi^v \sigma_v dv)}{\nu_u^2 + \sigma_u^2} du.$$

Let $\int_0^1 \pi^u \sigma_u du = \overline{\pi\sigma}$. Then this is exactly:

$$\overline{\pi\sigma} = \int_0^1 \frac{\delta_u \sigma_u \mu_u - \sigma_u \theta_u (\delta_u - 1)(\sigma_u \overline{\pi\sigma})}{\nu_u^2 + \sigma_u^2} du.$$

Rearranging, we get:

$$\overline{\pi\sigma} \left(1 + \int_0^1 \frac{\sigma_u^2 \theta_u (\delta_u - 1)}{\nu_u^2 + \sigma_u^2} \right) = \int_0^1 \frac{\delta_u \sigma_u \mu_u}{\nu_u^2 + \sigma_u^2} du.$$

Then, we have

$$\overline{\pi\sigma} = \frac{\gamma}{1 + \psi}, \tag{5.2}$$

where $\gamma = \int_0^1 \frac{\delta_u \sigma_u \mu_u}{\nu_u^2 + \sigma_u^2} du$ and $\psi = \int_0^1 \frac{\sigma_u^2 \theta_u (\delta_u - 1)}{\nu_u^2 + \sigma_u^2} du$. Note that this is equivalent to what was derived in [3] where instead of expectations over the type parameters such as $\Psi = \mathbb{E}[\frac{\sigma^2 \theta(\delta - 1)}{\nu^2 + \sigma^2}]$, we have these averages explicitly in terms of integrals.[1] Recall that an average of a function $f(u)$ over $[0, 1]$ is $\frac{1}{1-0} \int_0^1 f(u) du$ which is exactly what we have. Plugging Equation (5.2) back into the expression for π^u given by Equation (5.1), the optimal strategy becomes

$$\pi^u = \frac{\delta_u \mu_u - \theta_u (\delta_u - 1)(\frac{\sigma_u \gamma}{1 + \psi})}{\nu_u^2 + \sigma_u^2}.$$

With this explicit form, we may set the following parameters for the graph: $\mu_u = 0.5$, $\theta_u = 0.9$, $\sigma_u = 0.9$, $\nu_u = 0.6$, $\gamma = 1$, $\psi = 2$ and our time horizon $T = 1$. This results in the graph in Figure 5.1. In the MFG setting, that is without player

[1] In [3], they use similar notation ψ but φ instead to write our γ.

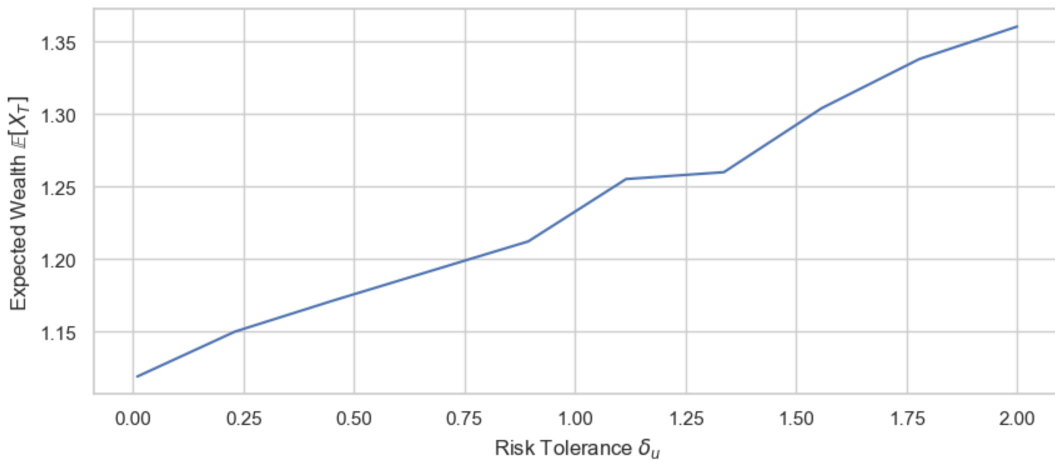

Figure 5.1: Expected Wealth versus Risk Tolerance in the MFG

interactions on the individual level, this graph highlights that as risk tolerance grows, our expected wealth grows as well. This agrees with our intuition that the more risky an investment is, the higher expected return.

5.3 Bilinear Graphon Simulation

For the case where $G(u, v) = uv$, we do not have an explicit form for π^u and thus we use Algorithm 1 to find a fixed point by setting parameters conveniently. We set similar parameters as before: $\mu_u = 0.5$, $\theta_u = 0.9$, $\sigma_u = \sigma_v = 0.9$, $\nu_u = \nu_v = 0.6$ for all players $v \neq u$ and time horizon $T = 1$. This allows us to write the integral as

$$\int_0^1 G(u, v)\sigma_v dv = \int_0^1 0.9uvdv$$

which integrates to $0.9u/2$. The graph is as in Figure 5.2.

Here, due to our graphon, we know that the index coincides with the amount of interaction agent u has with other agents. That is, if u is larger, player u interacts more and vice versa. In the Figure 5.2 thus, we see the same pattern as we saw in the MFG for the bilinear graphon in that as risk tolerance increases, expected wealth

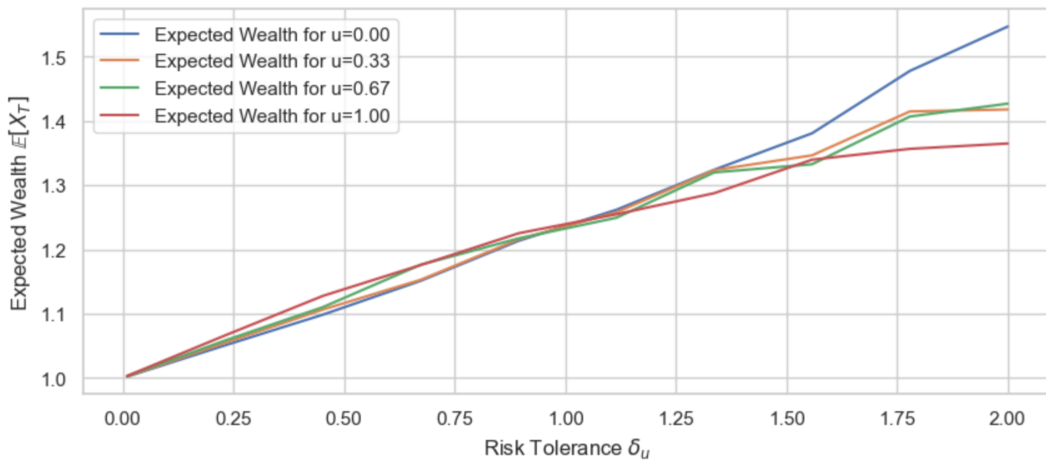

Figure 5.2: Expected Wealth versus Risk Tolerance for Different Interaction Levels

increases (regardless of the amount of interaction).

Moreover, we see that an agent that interacts less in this model tends to outperform agents that interact more if the agent's risk tolerance is greater than 1. This coincides with the solution in the CARA utility case in [6] since agents that interacted more had lower expected terminal wealth than those that interacted less. However, if risk tolerance is lower, it is not clear from Figure 5.2 exactly what happens. Thus, to more closely analyze these trends, we look at the graph in Figure 5.3.

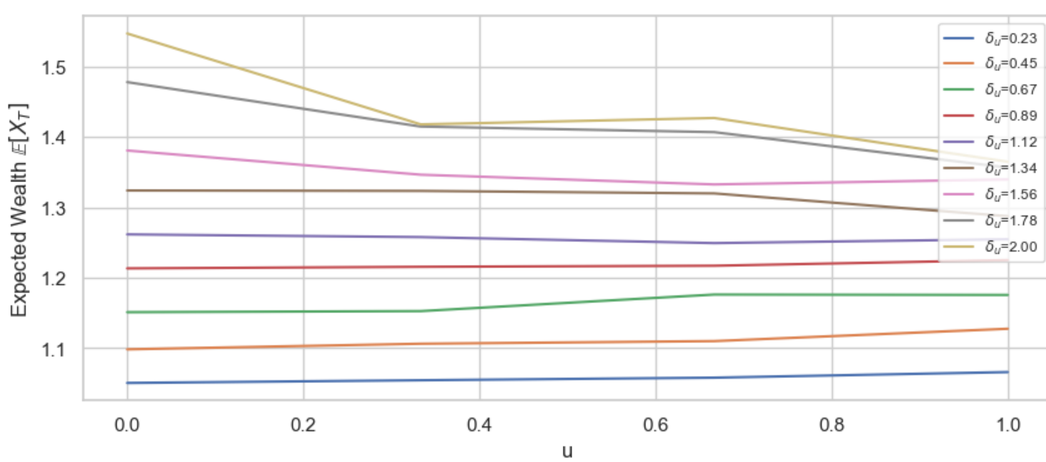

Figure 5.3: Expected Wealth versus Interaction Level for Different Risk Tolerance

In Figure 5.3, it is clear that for a smaller risk tolerance, an agent that interacts with a greater number of agents has a higher terminal expected wealth (albeit not very clear for $\delta_u = 0.23$). That is, the graph of expected wealths is increasing in u when $\delta_u < 1$. This agrees with a similar finding from [3] where they measure interactions through solely the parameter θ_u. Here, they find that, for $\delta_u < 1$, agents invest more for a higher θ_u, leading to a higher terminal expected wealth. This paper also says that as θ_u decreases and $\delta_u > 1$, agents invest less. In our setting, we may similarly see that only when risk tolerance is high enough ($\delta_u > 1$) does an increased number of interactions lead to a lower terminal expected wealth.

Chapter 6

Conclusion

In this work, we introduced an individualized competition weight into the model in [5] by Lacker and Zariphopoulou then derived the Nash Equilibrium investment strategy both in the n-player game and the graphon game. This weighted n-player and graphon models is somewhat similar to that of [9] introduced by Tangpi and Zhou, although we analyze CRRA utility rather than CARA utility with constant drift and volatility terms as well as investment in individual stocks. Moreover, our approach to finding the optimal control for this model is similar to [5] in deriving the HJB equation for the n-player game and graphon game. However, the work done to prove the existence of a Nash Equilibrium, approximation techniques to find the optimal control, and stability analysis performed in the graphon case is entirely novel.

The resulting solution coincides with the MFG when the graphon is constant and thus exhibits similar behavior in this case. We see that increases in risk tolerance leads to increases in expected wealth linearly. Moreover, the investment strategy we derived, regardless of the graphon we set, behaves similarly to that for the MFG. We compare these solutions by using θ_u as a proxy for $G(u, v)$ in the original model in [5] since both measure the amount of interactions.

Future work would attempt to find an explicit form for the Nash equilibrium π^u

rather than simply show the existence of a fixed point in both the n-player game and graphon game. While this explicit form is found in other papers such as [3] and [5], the averaging technique these papers use does not apply in this work as the parameter λ_{ij} or the graphon $G(u, v)$ is a function of two players. Moroever, future work here may try to gain a more general condition for the existence of a fixed point in the graphon case since the condition we gain is not easily interpretable nor easily fulfilled as opposed to the n-player game.

Furthermore, an easy way to extend this would be to introduce intermediate consumption into the problem. We expect the n-player strategy to look the same as the one we derived. We would also expect consumption to look similar to that in model without the weights; that is, as in [3]. However, extending consumption to the setting of the graphon game would be more difficult due to the approximation techniques we used, although it should follow similarly.

Finally, a different direction in which this model could be extended is by allowing agents to invest in multiple stocks rather than just their own stock. This idea was explored in [9] where each player is allowed to invest in the same vector of stocks.

Bibliography

[1] G.-E. Espinosa and N. Touzi. Optimal investment under relative performance concerns. *Mathematical Finance*, 25(2):221–257, 2015.

[2] E. Hubert. Optimal learning: Lecture notes for orf 418, 2022.

[3] D. Lacker and A. Soret. Many-player games of optimal consumption and investment under relative performance criteria, 2019.

[4] D. Lacker and A. Soret. A label-state formulation of stochastic graphon games and approximate equilibria on large networks, 2022.

[5] D. Lacker and T. Zariphopoulou. Mean field and n-agent games for optimal investment under relative performance criteria, 2018.

[6] M. Laurière, L. Tangpi, and X. Zhou. A deep learning method for optimal investment under relative performance criteria among heterogeneous agents, 2024.

[7] R. C. Merton. Lifetime portfolio selection under uncertainty: The continuous-time case. *The Review of Economics and Statistics*, 51(3):247–257, 1969.

[8] A. Parashar. Mean field games for optimal investment under relative performance criteria, 2023.

[9] L. Tangpi and X. Zhou. Optimal investment in a large population of competitive and heterogeneous agents, 2023.