UNIVERSITY OF GOTHENBURG
SCHOOL OF BUSINESS, ECONOMICS AND LAW

Evaluating the Adequacy of Skin in the Game in Central Counterparties

Investigating the Theory-Practice Gap

Edea Reci and Ilona Fedorova

Graduate School
Master of Science in Finance
Supervisor: Alexander Herbertsson

Evaluating the Adequacy of Skin in the Game in Central Counterparties
Investigating the Theory-Practice Gap
Edea Reci
Ilona Fedorova
Graduate School
School of Business, Economics and Law
University of Gothenburg

Abstract

Central counterparties (CCPs) are increasingly important in modern financial markets. They lower the counterparty credit risk for their members, usually a bank or a brokerage firm, by acting as an intermediary between the buyers and sellers to secure trade settlement. "Skin in the game" (SITG), the own capital of a CCP that can be used when one of its members defaults, acts as a safety mechanism by encouraging the CCP to manage risk diligently. By ensuring the CCP itself carries some of the financial burden in extreme loss situations, such as financial shortfalls from member defaults, SITG thus seeks to reduce moral hazard, where the CCP might be less motivated to manage risk given the absence of potential losses. However, the sufficient amount of SITG needed remains a subject of ongoing discussion. Theoretical models, such as Cont and Ghamami (2025), often propose more SITG than is generally found in practice.

Focusing on Cont and Ghamami (2025)'s "Cover 2" resilience standard, which requires CCPs to withstand simultaneous defaults of the two largest clearing members, we investigate this gap between theory and reality using data from major international CCPs for the period from 2019 to 2024. To determine whether the observed SITG levels align with the Cont and Ghamami (2025)'s incentive compatibility assumptions on tail risk, we conduct sensitivity analysis and examine key model parameters implied by market data. Our findings indicate that actual SITG contributions are frequently inconsistent with the underlying framework and often fall short of what the theory implies. These results suggest that CCPs may not have enough capital for severe but uncommon events, and raise questions about whether their incentives align with those of members. We conclude that continuous and focused efforts by regulators and the industry are necessary for strengthening the financial system from extreme shocks.

Keywords: central counterparty (CCP), Cover 2, default waterfall, incentive compatibility, moral hazard, skin in the game (SITG), systemic risk, tail risk

Acknowledgements

We would like to thank our supervisor, Alexander Herbertsson, for all the help and advice he gave us while we were writing our thesis. We also want to thank our families for always being there and encouraging us.

Edea Reci and Ilona Fedorova, Gothenburg, 2025-06-02

Contents

List of Figures

List of Tables

List of Abbreviations

CCP	central counterparty
CPMI	Committee on Payments and Market Infrastructures
CVaR	conditional value-at-risk
DF	default fund
EMIR	European Market Infrastructure Regulation
EVT	extreme value theory
FSB	Financial Stability Board
GPD	generalized Pareto distribution
ICC	incentive compatibility constraint
IM	initial margin
IOSCO	International Organization of Securities Commissions
OTC	over-the-counter
PFMI	Principles for Financial Market Infrastructures
PQD	public quantitative disclosure
SITG	skin in the game
VaR	value-at-risk

1

Introduction

1.1 Background and Problem Statement

Following the 2008 Global Financial Crisis (GFC), global financial markets underwent significant restructuring, and central clearing for standardised over-the-counter (OTC) derivatives was widely adopted (Group of Twenty (G20) 2009). As a result, central counterparties (CCPs) have become critical financial institutions, acting as intermediaries between buyers and sellers to absorb counterparty risk (CPMI-IOSCO 2020; Duffie and Zhu 2011). CCPs collect margins and default funds from their members (such as regulated financial institutions like banks and brokerage firms that directly participate in the CCP) to reduce the direct risk exposures in the market, which leads to risk being concentrated on their own balance sheets (Cont and Kokholm 2014; Tucker 2014).

Fundamental to a CCPs resilience is its "default waterfall", a sequence of financial buffers designed to cover the losses that arise if one or more clearing members default. Among this sequence is "skin in the game" (SITG), the CCP's own financial contribution. SITG ensures that the CCP has a direct financial interest in prudent risk management (CPMI-IOSCO 2020; Cont and Ghamami 2025). The CCP may underinvest in risk controls if SITG is too low, depending instead on surviving members' contributions to cover the losses (Haene and Sturm 2009; Cont and Ghamami 2025). Setting it too high, on the other hand, could discourage participation and increase clearing costs.

1.2 Problem and Research Gap

While it is widely accepted that CCPs should have sufficient SITG, there is no consensus on the optimal amount or the best method for determining it (Avalos and Mamatzakis 2018). Jurisdictions handle SITG in different ways. For instance, the European Market Infrastructure Regulation (EMIR) restricts SITG to 25% of the CCP's regulatory capital (European Commission 2013), while others like the US are more flexible (United States 2010; Commodity Futures Trading Commission 2011; SEC 2016). Research, such as by Cont and Ghamami (2025), suggests that meeting robust incentive goals might require 15-20% of the default fund in SITG, which is significantly larger than what many CCPs actually hold.

Based on empirical data, it appears that SITG for CCPs often ranges from 2% to 5% of the default fund (Thiruchelvam 2022; Walker 2023), which is far below the theoretical targets proposed by regulators and researchers. This gap raises concerns regarding CCPs' real incentives to restrain excessive risk-taking, and whether the financial system is shielded from systemic fallout from potential CCP members' defaults.

Using quarterly public quantitative disclosure (PQD) data from major CCPs, this thesis aims to address the gap between SITG levels observed in practice, and those suggested by theory. PQDs enable a comparative analysis using transparent, consistent data on metrics such as initial margin (IM), default fund (DF), SITG, member concentration, and stress-test results related to key resilience benchmarks, including "Cover 1" (the minimum resilience level of a CCP, which requires it to have the necessary resources to survive the default of its largest clearing member), and "Cover 2" (which ensures the CCP has sufficient resources to withstand the simultaneous defaults of the two largest members).

1.3 Research Objectives

This thesis aims to bridge theoretical insights with practice by applying the Cont and Ghamami (2025) SITG framework to real market data from CCPs, and examines how closely actual SITG levels align with the model's predictions. This thesis pursues several key objectives:

1. Provide a broad overview of CCP risk management practices, focusing mainly on the default waterfall, and the rationale for SITG, viewed within the current regulations.
2. Explain the theoretical model of the Cont and Ghamami (2025) incentive compatibility framework for SITG under the "Cover 2" standard.
3. Investigate trends of key central counterparty financial metrics, such as IM, DF, and SITG, alongside changes in membership concentration, to assess the extent of risk concentration in a CCP's largest member, using PQD data.
4. Use the Cont and Ghamami (2025) model to estimate SITG that is required under different parameters (e.g., tail exponents, concentration ratios) and compare these findings to the actual SITG reported in PQD by the CCPs.
5. Explore the "implied alpha" concept, in which we solve for the Pareto tail exponent used in Cont and Ghamami (2025) to align a CCP's observed SITG with the theoretical SITG formula, and interpret any related risk measures.
6. Reflect on how these findings might influence a CCP's risk management practices, SITG-related regulations, and overall financial stability.

1.4 Methodological Approach

This thesis combines theoretical modelling with empirical data analysis. The theoretical basis in our thesis is the SITG incentive framework developed by Cont and Ghamami (2025), which uses extreme value theory (EVT), specifically the Pareto tail distribution, to model losses that exceed initial margin. On the empirical side, we use quarterly public quantitative disclosures released by nine large CCPs from

Q1 2019 to Q4 2024 (where quarters correspond to calendar quarters and are denoted as Q1-Q4), thus capturing a range of markets and currencies. Empirical data were sourced from official regulatory disclosure reports published by each respective CCP, typically found on their corporate websites. We focus on SITG, IM, DF, stress estimates (Cover 1 and Cover 2), and membership concentration metrics; all data have been cross-currency converted into USD.

Our analysis examines CCP data to highlight how CCP resources and concentration measures evolved over time, applies Cont and Ghamami (2025)'s SITG formulas for the Cover 2 scenario to compare predictions with actual SITG values, conducts a sensitivity analysis of the Cont and Ghamami (2025) framework, particularly for the Pareto tail exponent and membership concentration, and estimates an implied alpha by inverting the SITG formula to find the Pareto tail exponent that can reproduce each CCP's observed SITG, followed by an examination of related tail-risk metrics.

1.5 Structure of the Thesis

By evaluating the proposed SITG model by Cont and Ghamami (2025) with real CCP disclosures, this thesis contributes to showing how much actual SITG levels differ from theoretically estimated SITG. In addition, it explores how realistic the Pareto-based tail assumptions might be.

The thesis is structured as follows: First, Chapter 2 provides an overview of central counterparties and presents the concepts of the default waterfall, SITG, and relevant regulatory developments. In Chapter 3, the theoretical model developed by Cont and Ghamami (2025) is presented, with a focus on its application to SITG incentives under the Cover 2 standard. The empirical analysis then examines CCP disclosures between 2019-2024. Chapter 4 focuses on trends in initial margin, default funds, SITG contributions, and concentration of clearing members. Chapter 5 then applies the Cont and Ghamami (2025) model and evaluates the sensitivity of SITG estimates with respect to various parameter assumptions and contrasts the model outputs with CCP reports. The concept of implied alpha is further examined in Chapter 6, where the parameter is derived from observed SITG in order to assess underlying tail risk assumptions. Chapter 7 summarises the thesis's main findings, limitations, and potential research avenues.

2

Literature Review

This literature review begins by outlining the critical role of central counterparties (CCPs) and their growing importance in the aftermath of the Global Financial Crisis (Section 2.1). It then examines several core dimensions of CCPs: their key operational functions (Section 2.2.1), the structure and function of the default waterfall (Section 2.2.2), and the main features of the regulatory framework in which they operate (Section 2.2.3). This basic understanding then allows for a more detailed understanding of "skin in the game" (SITG) and what it means for policy (Section 2.3). The chapter then examines theoretical models for SITG, focusing on the Cont and Ghamami (2025) framework (Section 2.4), and concludes with public quantitative disclosure (PQD) data as a source for empirical analysis (Section 2.5).

2.1 Central Counterparties (CCPs) in Post-Crisis Regulation

The financial landscape after the 2008 Global Financial Crisis (GFC) changed significantly. Before then, a significant portion of the global financial market involved over-the-counter (OTC) derivatives, which are financial contracts made directly between two parties, in contrast to trades in a public exchange (BIS 2008). Although OTC derivatives are useful for hedging risks, their bilateral structure led to interconnectedness, which played a large part in the severity of the crisis (FSB 2022). Counterparty credit risk, the risk that one of the parties in a deal might not meet its obligations, causing losses for the other party, is the primary threat of the OTC derivatives market (Duffie and Singleton 2003).

The default of Lehman Brothers, a major institution, led to a chain reaction where counterparties faced unexpected losses and uncertainty about their exposures (BIS 2009). In response to this lack of transparency in the OTC market, which made it difficult to assess concentrated risk, the G20 leaders' Pittsburgh Summit in 2009 mandated reforms. A key reform was the strengthening of the OTC derivatives market infrastructure with mandatory central clearing (Group of Twenty (G20) 2009). This agenda led to CCPs as critical pillars of the financial market.

Central Counterparties

A CCP is a specialised Financial Market Infrastructure that steps in between the trading parties in a financial contract. According to the Basel Committee on Banking Supervision, a CCP is "a clearing house that interposes itself between counterparties, becoming the buyer to every seller and the seller to every buyer" (BIS 2019). This interposition is achieved through Novation, the core legal and operational mechanism. Through this mechanism, the original bilateral agreement is replaced by two new contracts: between the original buyer and the CCP, and between the original seller and the CCP (Gregory 2014). Thus, through novation, the CCP bears all the counterparty credit risk that the initial participants would otherwise have faced (Duffie and Zhu 2011). According to CPMI-IOSCO Principle 18, only Clearing Members, typically regulated financial institutions that participate in the CCP, have direct interaction and can hold positions on the CCP, allowing the remaining market participants to access CCP services through them (CPMI-IOSCO 2012).

The primary economic reason for creating and using CCPs is to reduce and manage counterparty credit risk while improving transparency (Pirrong 2011). Central counterparties reduce counterparty credit risk by establishing contract performances, preventing defaults from cascading through markets (Pirrong 2011; Duffie and Zhu 2011). According to Hull (2023), the CCP acts as a circuit breaker by taking on the defaulting member's loss and ensuring that the obligations of non-defaulting members are met, thus reducing the risk of financial contagion.

The ownership structure of CCPs can be divided into two main types: Investor-Owned CCPs and Member-Owned CCPs. Investor-Owned are CCPs held by external shareholders, similar to typical corporations, and can be publicly traded companies. The main goal of investor-owned CCPs is profit maximisation, usually made from clearing fees or investment income on collateral. In contrast, Member-Owned CCPs are collectively owned by the clearing members themselves, the banks and brokers who use the CCP's services, and operate like a utility or a cooperative. The primary goal is to provide reliable and affordable clearing services to their members (Pirrong 2011).

Challenge of Risk Concentration

The increasing use of central clearing has led to a new and complex risk profile: the concentration of risk within the CCPs themselves (Tucker 2014). Major CCPs have become critical nodes in the financial system as they act as intermediaries for vast trade volumes. In Europe, LCH, formerly known as London Clearing House and owned by the London Stock Exchange Group (LSEG), handles most of the clearing for interest rate swaps globally, while Eurex Clearing AG, an investor-owned and leading CCP in Europe, handles mainly the European equity index derivatives. In North America, the Chicago Mercantile Exchange (CME) Group, a publicly traded investor-owned company, and a major clearer of futures, options, and swaps across various asset classes. ICE Clearing Credit, part of Intercontinental Exchange (ICE), is also investor-owned, and the leading clearer for credit default swaps. Additionally, the Depository Trust & Clearing Corporation (DTCC), a subsidiary entity that is member-owned, is crucial in the clearing of US equities and fixed income market.

Considering their size and importance, if a systemically important player were to fail, the consequences would be enormous and possibly freeze critical markets and trigger chain reactions internationally (Cont and Ghamami 2025).

The concentration of risk means CCPs must maintain exceptionally high levels of resilience. This risk arises because the CCP, by acting as both the buyer and the seller for multiple trades, concentrates the counterparty credit risk that was previously spread across numerous bilateral relationships. Therefore, the risk management frameworks used by CCPs must be robust to withstand severe market stress and even the default of one or more large clearing members. CCPs achieve this resilience through their 'default waterfall', a sequence of prefunded resources designed to reduce risk and cover losses from a member's potential default (Cont and Kokholm 2014; Murphy, Vasios and Vause 2014).

2.2 CCP Functions and Risk Structure

This section explores the specific functions of the operation of CCPs, and how they manage the risks they handle. Understanding these functions, along with the CCP's central role (discussed earlier), is key to analysing the incentive structures and quantitative models discussed in later chapters.

2.2.1 Core Functions

A CCP performs several core functions, such as Novation, Netting, and Collateral Management, that alter risk in clearing markets.

Novation

As introduced earlier, novation is the legal process where the CCP steps into the middle of the trade. The original contract is extinguished and replaced by two new ones, as it is illustrated in Figure 2.1.

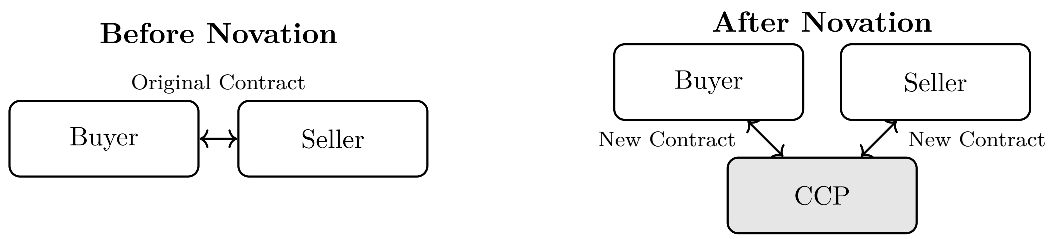

Figure 2.1: Illustration of Novation: CCP interposition creates two new contracts replacing the original bilateral one.

Netting

For each Clearing Member, the CCP conducts multilateral netting, which involves calculating the net obligation across all the cleared contracts within a certain class or market. More specifically, complex bilateral arrangements are replaced with a single net position of the member with the CCP for each product. This improves operational efficiency and reduces liquidity requirements under normal market conditions (Cont and Kokholm 2014; Heller and Vause 2012).

Collateral Management

To ensure that the novated contract obligations of the clearing member are met, the CCPs manage collateral (Gregory 2014). The CCP estimates, collects, and manages the financial assets that make up the collateral (also called margin), from their clearing members. These collaterals create a buffer against potential losses when a member defaults and are mainly of two types: initial margin (IM) and variation margin. The initial margin is the collateral acquired by each member at the start of the trade, and covers potential CCP losses if that member defaults while the CCP tries to close out the defaulter's positions. IM is a key concept for clearing, as it is the margin that should cover the largest expected losses from transactions. Variation margin, on the other hand, involves daily or more frequent regular payments made between the CCP and its clearing members. These payments reflect the current market value of each member's positions, with funds collected from losing members and transferred to those who are winning (Heller and Vause 2012).

2.2.2 The CCP Default Waterfall

Usually, the CCP starts by auctioning off a clearing member's portfolio to other non-defaulting participants if that member cannot meet its commitments. This action is intended to rapidly restore the balance of the CCP's positions, a so-called "matched book", and prevent larger shocks from spreading across the market. The CCP also encourages members to submit more competitive bids by holding these auctions before using its own resources (SITG), thereby controlling the systematic risks that can come from large, concentrated exposures (CPMI-IOSCO 2020).

Should the auction process not cover all losses, the CCP uses its default waterfall, a layered defence mechanism that the CCP uses to cover losses resulting from a clearing member's default or inability to meet its commitments. It is the cornerstone of a CCP's risk management framework, and it ensures that even the default of large members can be handled without the CCP itself defaulting. Its main purpose is to ensure that the CCP can continue to meet its obligations to surviving clearing members and to keep market stability (Gregory 2014; CPMI-IOSCO 2012).

The hierarchy of loss allocation is predefined, making it clear and predictable for all members what risks they might face if another member defaults. Although there are specific variations in the established hierarchy between CCPs and jurisdictions, the standard structure described in academic literature includes the following lines in a so-called "multilayered" default waterfall and is summarised below, also see in Figure 2.2.

Defaulter's Initial Margin (IM) and Variation Margin

The first line of defence against a potential loss is the defaulting member's own collateral that is meant to cover the costs of hedging the defaulter's portfolio under normal or moderately stressed market conditions (Rahman 2015). As noted above, IM is estimated based on the member's specific portfolio. Although it is usually determined quantitatively, the methodology varies among CCPs and can be complex. According to Gregory (2014), a common approach is the calculation of value-at-risk (VaR) or expected shortfall (ES), as it measures the risk of a loss distribution with a high confidence level (e.g., 99%) that losses will not exceed VaR or ES. In their paper

"Skin in the Game: Risk Analysis of Central Counterparties", Cont and Ghamami (2025) state that IM is usually determined from a VaR of the CCP's gross exposure to that member. The CCP also values the variation margin, as the counterparty of all transactions.

Defaulter's Default Fund Contribution

The second line of defence is the defaulter's pre-funded contribution into the default fund of the CCP, and it is used when the defaulter's IM is not enough (Haene and Sturm 2009). The DF is a pool of financial resources contributed by all clearing members ex-ante. The main objective of this particular fund is to absorb losses that are larger than the defaulted member's margin. Stress testing scenarios are usually used to calculate the size of the DF, where again various CCPs decide on different sizes, but aim to have enough funds to cover the default of their largest members.

CCP's Own Capital (First-Layer Skin in the Game)

The CCP's own capital is a layer used after the defaulter's resources are exhausted but before the contributions of non-defaulting members. This layer, commonly known as skin in the game (SITG), will be discussed in detail below.

Non-Defaulters' DF Contributions

If losses exceed the defaulter's resources and the CCP's own capital, surviving members' contributions to the DF are used, usually in proportion to their contributions (Heath et al. 2016). This last line mutualises the remaining losses among the clearing members.

Additional CCP Resources (Second-Layer Skin in the Game)

In some CCP structures, particularly under regulations such as the EU's European Parliament and Council (2021), a second-layer SITG appears after the contributions from non-defaulting members but before other recovery tools (Cont and Ghamami 2025).

Members' Unfunded DF Contributions (Assessments)

If the usual margins and DF contributions are not enough, the CCP can ask non-defaulting members for extra funds, sometimes called "cash calls". These additional contributions go beyond each member's initial DF commitments and are a way of further sharing the losses among the members who did not default. Regulations usually cap how high these cash calls can go, and the details (like maximum amounts and triggers) are specified in each CCP's rulebook (CPMI-IOSCO 2017; Cont 2015).

Other Recovery Measures

If even those additional assessments are not enough to cover the losses, CCPs may take more drastic steps, including variation margin gains haircuts (Cont and Ghamami 2025)

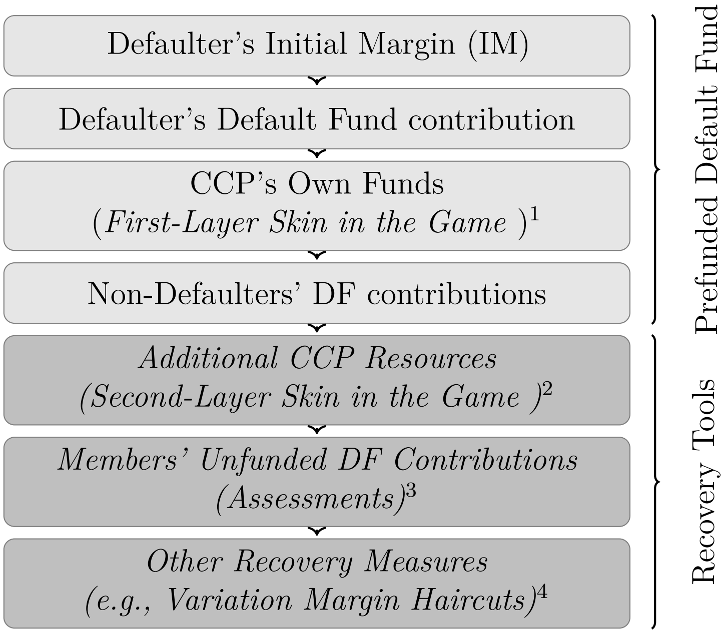

Figure 2.2: A simplified view of a CCP's multilayered default waterfall. The top boxes show prefunded default fund resources; the bottom boxes show additional tools used if the prefunded resources run out. Jurisdictional rules differ (EU, UK, US etc.), as detailed next.

While the sequence described above represents a common waterfall design, variations in the structure exist. Cont and Ghamami (2025) specifically point out the difference between monolayer and multilayer default waterfalls. The main difference between these default waterfall structures is how member collateral is organised and used.

A multilayer setup, which is common in clearing derivatives CCPs, has a clear separation of IM, a mutualised DF, and one or more SITG layers placed at certain points in the waterfall. Derivatives CCPs like LCH and CME follow this model. In contrast, monolayer CCPs, which for example Cont and Ghamami (2025) state that are commonly used by "systemically important securities CCPs in the US", do not have a separate DF, instead use a single pool of member collateral. SITG is placed

[1] *First-Layer SITG*: In the EU, 25% of a CCP's total regulatory capital must be placed ahead of non-defaulters' contributions (European Commission 2013); the UK mirrors this requirement in its onshored version of EMIR (United Kingdom 2018); US regulations under Dodd-Frank (United States 2010) afford CCPs some flexibility, subject to regulatory approval.

[2] *Second-Layer SITG*: EU rules (10–25% of risk-based capital) require these additional CCP resources to be used after non-defaulters' DF contributions (European Parliament and Council 2021); the UK is considering a similar structure (HM Treasury 2021); in the US, no explicit rule mandates a second SITG (United States 2010).

[3] *Cash Calls*: Resolution authorities in the EU and the UK can demand extra member contributions (up to 2x DF in the EU; up to 2–3x DF in the UK, depending on whether losses stem from defaults or other events) (European Parliament and Council 2021; HM Treasury 2021). In the US, Dodd-Frank allows CCP rulebooks to define such assessments (United States 2010).

[4] *Other Recovery Measures*: EU and UK regulations explicitly permit such measures (European Parliament and Council 2021; HM Treasury 2021); US CCPs include them in their regulator-approved recovery or wind-down plans (United States 2010; SEC 2016).

at a certain level in that single pool. Because monolayer structures can concentrate a large amount of mutualised collateral, some argue that SITG must be relatively higher (as a fraction of total resources) to counter potential moral hazard (Cont and Ghamami 2025).

2.2.3 Regulatory Framework

CCP risk management is regulated by a framework primarily developed after the GFC. The CPMI-IOSCO's Principles for Financial Market Infrastructures (PFMIs) serve as international benchmarks (CPMI-IOSCO 2012). Among the principles for financial resilience, Principle 4 (Credit Exposures) requires CCP to maintain sufficient financial resources, with "Cover 1" (largest member default) serving as the minimum baseline, and "Cover 2" (top two members default) usually required for more systemically important CCPs. Furthermore, Principle 15 (General Business Risk) requires that CCPs maintain sufficient capital (i.e., equity) to address non-default risks, including operational losses. Although PFMIs are fundamental, their implementation is based on principles, permitting jurisdictions to adopt varying SITG standards, as no universal formula exists.

European Union: EMIR and CCPRRR

The EU's EMIR requires CCPs to commit at least 25% of their capital as first-layer SITG (European Commission 2013). The CCP Recovery and Resolution Regulation (CCPRRR), (EU) 2021/23, requires a second-layer SITG from CCPs' "pre-funded dedicated own resources" for extreme losses (European Parliament and Council 2021, Art. 9(14)). Commission Delegated Regulation (EU) 2023/840 states that this second-layer (10-25% of risk-based capital) should reflect the CCP's specific activities, structure, and incentives (European Commission 2023, Art. 1 & Annex). EMIR 2.2 (European Parliament and Council 2019) also introduced stricter tiering for systemically important third-country CCPs.

United States: Dodd-Frank, CFTC, and SEC

The US aligns with PFMI standards through the Dodd-Frank Act and the rules of the CFTC (Commodity Futures Trading Commission), and SEC (Securities and Exchange Commission), but does not specify a required SITG amount (United States 2010; Commodity Futures Trading Commission 2011; SEC 2016). US CCPs set SITG policies which regulators must approve, with emphasis on CCP-defined recovery and wind-down plans.

United Kingdom: Post-Brexit UK EMIR and FSMA 2023

UK EMIR maintains the 25% first-layer SITG rule (United Kingdom 2018). The Financial Services and Markets Act 2023 (FSMA 2023) broadened the Bank of England's CCP resolution powers, including tools like statutory cash calls, variation margin gains haircuts, and contract tear-ups (HM Treasury 2021). The UK also considers an EU-style second-layer SITG.

Global Harmonization

Other nations (e.g., Japan, Australia) also align SITG policies with PFMI, adapting them locally (EACH 2021). Despite this, fully harmonising CCP oversight globally remains challenging, especially for cross-border products, due to different views on regulation, governance, and risk-calibration methods. International groups like the Financial Stability Board (FSB), CPMI, and IOSCO are working to improve cross-border coordination and promote consistent CCP resilience measures, including recovery and resolution standards. Agreeing on complex issues such as SITG amounts and resolution tools is important for financial stability and avoiding regulatory arbitrage.

2.3 Skin in the Game (SITG): Rationale and Policy Context

Skin in the game (SITG), the CCP's own capital contribution to its default waterfall, is a crucial element. This layer of the defence mechanism is used to cover losses from a member's default. The primary justification for this requirement is to align the incentives between the CCP and its clearing members (Cont and Ghamami 2025; Kuong and Maurin 2024). In this context, incentive alignment involves designing the loss allocation process so that the CCP's own interests encourage actions that protect its members and the stability of the clearing system.

This requires the CCP to place its own financial resources at risk before the mutualised DF contributions of the surviving members are fully used, aiming to ensure that the CCP has a direct financial stake in its risk management. In essence, SITG forces the CCP to bear some of the costs of its own poor risk management decisions and shields the surviving members from initial losses, reinforcing the principle that the defaulter and the CCP bear the losses first (CPMI-IOSCO 2012).

Several proponents argue that a CCP, specifically an investor-owned one, without a sufficient SITG may have less reason to manage risk carefully if most losses are paid by surviving members. This practice creates moral hazard and agency problems (Haene and Sturm 2009). In addition, the CCP showing it has financial commitment, gives its clearing members more confidence that it can manage risk and survive (Menkveld 2017). According to Cont (2017), having its own capital at risk encourages the CCP to maintain robust risk management practices, such as strict membership criteria, a conservative margin model, and monitoring members' exposure.

Although the principle of SITG is widely accepted, determining the appropriate level and structure of SITG is still debated by industry groups, regulators, and academics, and depends on factors such as the specific default waterfall design and the CCP's ownership structure (Cont and Ghamami 2025). In member-owned CCPs, clearing participants collectively fund and govern the facility, and, in theory, have fewer conflicts between shareholder returns and default risk control. However, collective decisions made among heterogeneous members can still lead to suboptimal capital contributions (Hart 1996; Hansmann 2014). Investor-owned CCPs, on the other hand, face a clearer tension between delivering returns to shareholders and committing enough capital for systemic resilience. In both structures, misaligned incentives can result in CCP equity being not enough at critical moments in the waterfall (Avalos and Mamatzakis 2018).

The debate over the levels of SITG includes many different views. Some clearing members, such as ABN-AMRO, Allianz, Barclays, and BlackRock (2020), argue for a minimum Cover 2 standard and for a significant second-layer SITG, to ensure the CCP's incentives are properly aligned (ABN-AMRO, Allianz, Barclays, and BlackRock 2020). Others, including Clancy (2025), have reported on analyses such as that of Cont and Ghamami (2025), highlighting that academics have called for a tenfold increase in CCP capital requirements to substantially enhance systemic resilience. Industry organisations, including EACH (European Association of CCP Clearing Houses) maintain that current capital rules, which include the 25% requirement (EACH 2021), are a reasonable compromise. Academics such as Cont and Ghamami (2025) propose quantitative metrics that tie CCP capital more directly to the risks from its largest one or two members to improve incentives and lower moral hazard. Nonetheless, the path to global convergence is not yet complete, with many calling for more studies based on data about CCP resources.

Cont and Ghamami (2025) argue that SITG should be linked to the largest member exposures (e.g., the top one or two exposures from their clearing members) rather than the entire DF. Echoing these concerns over adequacy, reports such as Clancy (2025) underscore the potential need for substantial increases in SITG, particularly for certain CCP structures, given the scale of systemic risks involved. Others highlight the need to scale SITG with the overall risk profile of the CCP, given evidence that the CCP's own capital contributions have remained flat even as member exposures have increased (Wilkes 2020; Walker 2023; EACH 2021). Empirical studies also suggest that CCPs with higher SITG tend to use more careful risk management models, but consensus on a global standard remains elusive. This debate highlights the need for quantitative calibration methods based on actual risk exposure.

2.4 Theoretical Models on SITG

Surprisingly few explicit models in academic literature quantitatively calculate CCP SITG levels. Baseline minimum SITG requirements are provided by regulatory guidelines such as EMIR and CCPRRR (see Section 2.2.3). However, these regulations often use fixed percentages of CCPs' capital without taking into account a direct relationship to the CCP's actual risk exposures or incentives. Models aiming to quantify risk or determine capital requirements under specific assumptions (such as concentration or tail behaviour) can enhance transparency regarding parameter impacts. Researchers have explored netting benefits and moral hazard in CCPs (Duffie and Zhu 2011; Biais, Heider and Hoerova 2016; Glasserman and Young 2016). However, many models do not provide explicit formulas for SITG amounts.

This thesis builds on the framework developed by Cont and Ghamami (2025), which directly addresses SITG sizing through an incentive compatibility approach. This approach seeks to address the agency problem where an insufficient capitalised CCP might lead to under-investment in risk controls, leaving members to cover extreme losses. The model proposed by Cont and Ghamami (2025) includes quantitative methods for calculating SITG layers, a first layer and a second layer, so that the CCP's incentives align with member preferences and financial stability. By using the extreme value theory (EVT), which assumes Pareto tails for losses exceeding margins

(McNeil, Frey and Embrechts 2015), Cont and Ghamami (2025) derives relationships between SITG, tail probabilities, trigger probabilities, the tail exponent, and relevant concentration metrics.

While the Cont and Ghamami (2025) model offers analytical closed-form expressions, it also presents limitations common to many static frameworks. For instance, the reliance on a particular distribution to model extreme losses, such as the Pareto tail distribution, may help its quantitative implementation, but may not capture certain dynamic feedback effects or complex dependencies as well as simulation-based approaches. Such trade-offs are common in analytical models.

Nonetheless, the conceptual strength of the Cont and Ghamami (2025) model lies in its proposed incentive-based formulation. This offers a theoretically grounded and risk-sensitive alternative to fixed capital ratios set by regulators. The model of Cont and Ghamami (2025) forms the core analytical basis for subsequent chapters of this thesis.

2.5 Data Source: Public Quantitative Disclosure (PQD)

Recognising the need for greater transparency about CCP risks and resources, regulators introduced PQDs. These disclosures provide standardised information on core metrics like initial margin levels, default fund sizes, concentration of top clearing members, and other capital details. Although PQD data are available at a quarterly frequency from major CCPs, the level of detail and consistency across CCPs can differ, which limits cross-comparisons.

Nonetheless, PQD data have become an important input for empirical studies looking at default fund adequacy, SITG calibrations, and how CCP exposures change over time. Later chapters in this thesis will use these data to explore current SITG practices across different CCPs, test hypotheses derived from the Cont and Ghamami (2025) framework, and evaluate the fit of various modelling approaches under real-world conditions.

This literature review showcases the CCPs pivotal role in the post-crisis financial landscape. Their concentrated risk exposures suggest they need a clear, robust institutional model and transparent regulation. SITG emerges as a key yet often imprecisely defined or calibrated dimension of CCP resilience, reflecting complex agency problems and governance trade-offs. While global and regional regulations have advanced significantly, debates persist on whether current frameworks do enough to align CCP incentives with broader financial stability goals.

The next chapter builds upon the SITG background by examining additional theoretical foundations of default fund structures and SITG placement. This sets the stage for a deeper look at the Cont and Ghamami (2025) model specifications in Chapter 3, followed by empirical analysis using PQD data in Chapter 4, Chapter 5, and Chapter 6.

3

Theoretical Framework: The Cont and Ghamami (2025) SITG Model

Building upon the foundational understanding of central counterparties (CCPs) and their default waterfall structures (as described in Chapter 2, Section 2.2.2), this chapter provides a detailed explanation of the skin in the game (SITG) framework proposed by Cont and Ghamami (2025). The quantitative, closed-form SITG formulas derived from the model are a key part of this thesis. These formulas are derived from a theoretical framework addressing economic challenges in central clearing, particularly the incentive alignment between a CCP and its members during rare, severe tail events.

This chapter combines the main ideas of the Cont and Ghamami (2025) framework, providing a clear overview before applying it in the chapter that follows. We begin with the specification of the default waterfall under the relevant "Cover 2" standard, Section 3.1, and the economic rationale for SITG, focusing on incentive alignment and moral hazard, Section 3.2. The discussion will then cover the use of extreme value theory (EVT) to model the tail risks that drive these incentive problems, including the reasoning behind the Pareto distribution assumption and the calculation of key parameters, Section 3.3.

The SITG formulas for the Cover 2 case are subsequently analysed, with a detailed derivation provided in Appendix A. These formulas are explicitly linked to the incentive constraints and EVT assumptions, Section 3.4. We also include a discussion of the model's key assumptions and scope, Section 3.5, followed by the derivation of additional risk metrics derived from the EVT framework, Section 3.6.

By presenting the reasons behind the model inCont and Ghamami (2025), along with its derivations, this chapter aims to provide a comprehensive understanding of how and why the Cont and Ghamami (2025) SITG formulas are developed and used in this thesis.

3.1 Cover 2 Default Waterfall Specification

The SITG formulas discussed in this chapter apply specifically to a **Cover 2** default scenario. This scenario requires the CCP to be resilient against the simultaneous default of the two clearing members that have the largest credit exposures. The relevant layers of the default waterfall, as defined by Cont and Ghamami (2025) for this scenario, are:

Let M_i, D_i, and E_i be the initial margin, default fund contribution, and tail exposure (net of initial margin) for a general member i ($i = 1, \ldots, N$), respectively. We denote the k-th largest value from these sets as $M_{(k)}$, $D_{(k)}$, and $E_{(k)}$, such that $M_{(1)} \geq M_{(2)} \geq \ldots$, $D_{(1)} \geq D_{(2)} \geq \ldots$, and $E_{(1)} \geq E_{(2)} \geq \ldots$. In the context of defaults, member (1) is typically the one with the largest tail exposure $E_{(1)}$, and $M_{(1)}$ and $D_{(1)}$ are its associated resources.

The Cover 2 default waterfall, as defined by Cont and Ghamami (2025) for this scenario, applies these resources in a specific sequence. Initially, losses are covered by the largest defaulter's (member (1)) own resources: its initial margin $M_{(1)}$ and its specific default fund contribution $D_{(1)}$. If these prove insufficient, the CCP's own capital is used, in the form of first-layer SITG (S_2), where the subscript '2' indicates calibration for a Cover 2 scenario. Should losses persist, the mutualised default fund (DF) is drawn upon. For Cover 2, this DF is sized at $D_{s,2} = E_{(1)} + E_{(2)}$, reflecting the sum of tail exposures from the two largest members. Each $E_{(k)}$ denotes the tail exposure of the k-th largest member, with its corresponding scale parameter $\kappa_{(k)}$ derived as per Eq. (3.5), and is funded primarily by surviving members. After $D_{s,2}$ and S_2 are depleted, the second-layer SITG ($\tilde{S}_2$) contributed by the CCP can be used. Beyond these resources, further recovery and resolution tools, such as member assessments can be employed, although these are not explicitly modelled within the SITG sizing framework. This specific ordering is important for the ICCs used to determine SITG.

3.2 Model Assumptions and Incentive Design

Central to the Cont and Ghamami (2025) framework is its approach to the economic problem of ensuring careful CCP risk management and quantifying relevant extreme risks.

The Incentive Problem and Compatibility Constraints

Moral Hazard and Agency Problems as Motivation

A main challenge in CCP resilience is determining the adequate amount of the CCP's own capital commitment, or SITG. The primary economic motivation for requiring SITG comes from the possibility of moral hazard and agency problems (Haene and Sturm 2009; Cont and Ghamami 2025). If a CCP operator (whether investor-owned or member-owned) perceives that extreme member defaults will primarily affect members who did not default (through the mutualisation of losses via the DF), it may have reduced incentives to invest enough in risk management practices (e.g., conservative margin setting, strict membership criteria, and monitoring positions). This misalignment, where the CCP risk manager does not fully internalise the impact of the risks taken, can lead to capital inadequacy and pose systemic risks (Duffie, Scheicher and Vuillemey 2015; Cont and Ghamami 2025).

Solution Approach: Incentive Compatibility

To formally address this, Cont and Ghamami (2025) use an incentive compatibility approach. This involves designing the loss allocation mechanism, specifically, the amount and placement of SITG, so that the CCP's own financial interests align with careful risk management that benefits the entire clearing system (Cont and Ghamami 2025). The core idea is to structure SITG so the CCP appropriately "shares the pain" of extreme losses, thereby taking on some of the risk it handles for its members.

The alignment between CCP and its members, is achieved through incentive compatibility constraints (ICCs), which are mathematical requirements on the loss probabilities faced by different parties. A key ICC focuses on protecting members' prefunded DF contributions. ICC requires that the probability of a surviving member i (where $i \neq j$) experiencing any loss on their default fund contribution ($L_i^j > 0$) due to the default of member j must not exceed a threshold q_D. This loss event occurs if the defaulter's (member j's) net loss ($U_j - M_j$) exceeds their own default fund contribution D_j plus the CCP's first SITG layer S. Formally, the probability ($L_i^j > 0$) is given by

$$P(L_i^j > 0) = P\Big(\underbrace{U_j - M_j}_{\text{Net Loss from member } j} > \underbrace{D_j + S}_{\text{Defaulter } j\text{'s Resources + SITG}} \Big) \leq q_D \tag{3.1}$$

where, for the defaulting member j, U_j is the CCP's exposure, M_j is their initial margin, and D_j their default fund contribution. Furthermore, S is the first SITG layer (denoted S_2 in the Cover 2 context). The threshold q_D is the unconditional probability $P((U_j - M_j)^+ > E_j)$, where E_j is the tail exposure of member j, defined as the value-at-risk (VaR) of this excess loss at the tail probability level q_D, formally $E_j = \mathrm{VaR}_{q_D}((U_j - M_j)^+)$.

The motivation for the constraint (3.1), as highlighted by Cont and Ghamami (2025), is that in the absence of SITG ($S = 0$), the probability of a surviving clearing member i suffering a loss is generally higher than the target protection level q_D. This is because the threshold for member loss (D_j) is usually less than the tail exposure E_j, used to define q_D. Specifically, $P(L_i^j > 0|S = 0) = P(U_j - M_j > D_j)$. Since D_j (often proportional to E_j) is usually less than E_j, and the loss distribution has a decreasing tail, $P(U_j - M_j > D_j) \geq P(U_j - M_j > E_j) = q_D$ due to the definition of E_j via q_D. Thus, $P(L_i^j > 0|S = 0) \geq q_D$. This inequality (Cont and Ghamami (2025, Eq. (3.6))) shows a mismatch without SITG. Requiring $P(L_i^j > 0) \leq q_D$ means there must be a positive SITG ($S > 0$) to absorb the initial part of the tail loss, thereby fixing the incentive imbalance and protecting member contributions up to the desired level.

To further strengthen alignment, particularly for the critical scenario involving the default of the member with the largest tail exposure (member (1), associated with $U_{(1)}$, $M_{(1)}$), Cont and Ghamami (2025) introduce an incentive compatibility approach (ICC). This ICC requires the loss probability for any surviving member i (where $i \neq (1)$) to equal the CCP's own probability of loss beyond the entire Cover

2 DF ($D_{s,2}$):

$$P(L_i^{(1)} > 0) = P(L_0^{(1)} > 0), \tag{3.2}$$

where $L_0^{(1)} = (U_{(1)} - M_{(1)} - D_{s,2})^+$ denotes the CCP's loss after exhausting the Cover 2 DF due to the default of member (1). By making these probabilities equal for the most impactful default, the framework ensures the CCP faces the same likelihood of loss as its members precisely when the mutualised resources are fully used up, strongly motive the prevention of such extreme events (Cont and Ghamami 2025, Section 3.2).

The Cont and Ghamami (2025) framework further sizes SITG using specific target loss probabilities for the CCP's capital layers. For the Cover 2 scenario, these are π_2 for the first SITG layer (S_2) and $\tilde{\pi}_2$ for the second layer ($\tilde{S}_2$).

The first target probability, π_2, represents the target probability that losses from the largest defaulter (member (1)) exceed its own resources ($D_{(1)}$) plus the first SITG layer (S_2). It is a policy choice, generally set such that $\pi_2 \leq q_D$, where q_D is the baseline default fund protection level.

For the second layer, $\tilde{\pi}_2$, is a stricter target probability ($\tilde{\pi}_2 < \pi_2$) for losses from member (1) breaching all pre-funded resources up to and including the mutualised default fund ($D_{s,2}$), S_2, and also $\tilde{S}_2$. This reflects greater risk aversion for more extreme losses.

These target probabilities, π_2 and $\tilde{\pi}_2$, are therefore crucial inputs for the SITG formulas, converting risk tolerance and incentive alignment goals into quantifiable CCP capital requirements. The detailed formulation of S_2 and $\tilde{S}_2$ using these targets is presented in Section 3.4.

3.3 Tail Risk Modeling via Extreme Value Theory (EVT)

Calculating the small probabilities ($P(\cdot)$ in Eqs. (3.1)-(3.2)) associated with large losses requires a methodology that properly captures tail behaviour. Standard risk models relying on assumptions like normality often perform poorly in the tails, underestimating the frequency and magnitude of extreme financial events (McNeil, Frey and Embrechts 2015). This underestimation is particularly dangerous for CCPs, which exist precisely to manage rare but very severe counterparty defaults. The following subsections explain how EVT is used to model these extreme risks properly.

Approach: Extreme Value Theory and Peaks Over Threshold (POT)

To address this, Cont and Ghamami (2025) use extreme value theory (EVT), a branch of statistics created to model the extremes of probability distributions. They adopt the peaks-over-threshold (POT) methodology (McNeil, Frey and Embrechts 2015, Ch. 5; Tsay 2010, Ch. 7), which directly models the distribution of losses that exceed a predetermined high threshold. In the CCP context, the natural threshold for a member i is their initial margin M_i, as losses beyond this point, $L_i = (U_i - M_i)^+$, are those that can potentially use up the layers in the default waterfall. The exposure variable U_i measures the CCP's underlying, pre-collateral market risk exposure to its member i. It represents the total potential loss the CCP estimates it could face if member i were to default, caused by adverse changes in the market value of that member's entire portfolio positions cleared by CCP. This potential loss is assessed over the margin period of risk (MPOR) - the time horizon the CCP expects it needs to successfully manage and neutralise the risk from the defaulting member's portfolio (Cont and Ghamami 2025, Section 2.1). As an example of how U_i might be constructed, particularly for centrally cleared derivatives, Cont and Ghamami (2025, Section 2.1) note that it may be taken as the sum of potential losses across all individual positions member i holds that are cleared by the CCP.

Theoretical Justification: GPD Convergence

A key result from EVT, the Pickands-Balkema-de Haan theorem, gives the theoretical reason for the POT approach. It states that for a broad class of underlying loss distributions (from which U_i is drawn), the distribution of excesses L_i over a sufficiently high threshold M_i approaches a generalised Pareto distribution (GPD) (McNeil, Frey and Embrechts 2015, Thm 5.20). This theorem gives a mathematical basis to use GPDs when modelling very large losses.

Assumption: Standard Pareto Tail

Financial data frequently exhibits "heavy tails", where extreme outcomes are more likely than predicted by, for example, a normal distribution. This heavy-tailed behaviour is described mathematically by a GPD that has a positive shape parameter. Based on this theoretical foundation (GPD convergence) and empirical observations of heavy tails in finance (often finding α between 2 and 4 for relevant asset classes (Cont 2001; Cont and Kan 2011)), Cont and Ghamami (2025) make a reasoned assumption that these excess losses follow a standard Pareto distribution (Cont and Ghamami (2025, Assumption 4.1)), a specific, analytical type of GPD suitable for such situations.

The Conditional Pareto Distribution and Key Parameters

The Conditional Pareto Law

Let $L_i = (U_i - M_i)^+$ be an excess loss for member i. The core assumption from Cont and Ghamami (2025, Assumption 4.1) is that the conditional distribution of L_i given $U_i \geq M_i$ (i.e., given $L_i > 0$) follows a Pareto law with a member-specific scale $\kappa_i > 0$ and a common (Pareto) tail exponent $\alpha > 1$:

$$P\Big(L_i > x \,\Big|\, L_i > 0\Big) = P\Big(U_i - M_i > x \,\Big|\, U_i \geq M_i\Big) = \left(\frac{\kappa_i + x}{\kappa_i}\right)^{-\alpha}, \quad x \geq 0. \tag{3.3}$$

This implies that the probability of exceeding x gets smaller with ($\sim x^{-\alpha}$) for large x. The parameter α controls this decay (lower α means heavier tail), while κ_i sets the scale for member i.

Unconditional Probability and Derivation of κ_i

Let $q = P(U_i \geq M_i)$ be the probability of member i breaching its initial margin M_i. The unconditional probability of the excess loss L_i for member i exceeding x is then found by conditions:

$$\begin{aligned} P(L_i > x) &= P(L_i > x \mid L_i > 0)P(L_i > 0) \\ &= P\Big(U_i - M_i > x \,\Big|\, U_i \geq M_i\Big)P(U_i \geq M_i) \\ &= \left(\frac{\kappa_i + x}{\kappa_i}\right)^{-\alpha} q \\ &= q\left(1 + \frac{x}{\kappa_i}\right)^{-\alpha}. \end{aligned} \tag{3.4}$$

The scale parameter κ_i is crucial but not directly observable. Cont and Ghamami (2025) link it to the tail exposure E_i of member i. E_i is defined as the value-at-risk (VaR) of the excess loss L_i at an unconditional probability level q_D, where:

$$E_i = \mathrm{VaR}_{q_D}(L_i) \quad \text{meaning} \quad P(L_i > E_i) = q_D\,.$$

Here, q_D is a small tail probability (e.g., 1%, 0.5%, where $q_D < q$). To find κ_i, we set $x = E_i$ in Eq. (3.4), that is:

$$q_D = q\left(1 + \frac{E_i}{\kappa_i}\right)^{-\alpha}.$$

Solving for κ_i gives the practical formula used in Cont and Ghamami (2025, Eq. (4.2)):

$$\kappa_i = \frac{E_i}{\left(\frac{q}{q_D}\right)^{1/\alpha} - 1}\,. \tag{3.5}$$

Note that Equation (3.5) is crucial as it allows the unobservable κ_i to be determined from E_i (a tail risk measure CCPs often compute for member i), q (margin breach probability), q_D (unconditional tail probability for E_i), and the tail exponent α.

3.4 SITG Formulas for the Cover 2 Case

We now present the explicit formulas for the SITG-quantities, S_2 and $\tilde{S}_2$ in the Cover 2 case, taken from the framework of Cont and Ghamami (2025). These formulas combine the EVT-Pareto model, Section 3.3, with the ICCs, Section 3.2, and the Cover 2 waterfall in Section 3.1. Further detailed derivations are also provided in Appendix A.

3.4.1 Concentration Ratios and Parameters

The SITG amounts derived from the Cont and Ghamami (2025) framework are sensitive to how member risk is concentrated. We define the necessary terms and relationships for the Cover 2 scenario.

Let E_i be the tail exposure of member i (net of initial margin), and $E_{(k)}$ be the k-th largest among these tail exposures. The sum of all members' tail exposures is $E_{Total} = \sum_{i=1}^{N} E_i$. The concentration ratios for the top two members are then defined as $c_k = E_{(k)}/E_{Total}$ for $k = 1, 2$. The Cover 2 default fund size, denoted $D_{s,2}$, is sized to cover the two largest tail exposures: $D_{s,2} = E_{(1)} + E_{(2)}$.

From these definitions, the largest tail exposure, $E_{(1)}$, can be expressed in terms of $D_{s,2}$ and the concentration ratios. Since $c_1 + c_2 = (E_{(1)} + E_{(2)})/E_{Total} = D_{s,2}/E_{Total}$, it follows that $E_{Total} = D_{s,2}/(c_1 + c_2)$. Substituting this into the definition of c_1, we obtain the key relationship for $E_{(1)}$:

$$E_{(1)} = c_1 E_{Total} = c_1 \frac{D_{s,2}}{c_1 + c_2} = \frac{c_1}{c_1 + c_2} D_{s,2} \,. \tag{3.6}$$

Furthermore, a key component in the allocation of financial resources during a default event is the defaulting member's own contribution to the default fund. To determine this for the largest member (denoted as member (1), corresponding to exposure $E_{(1)}$) within the Cover 2 framework, we apply the principle of proportional default fund contributions. This principle, as stated by Cont and Ghamami (2025, Eq. 2.4), asserts that a member i's contribution to the default fund, D_i, is proportional to its share of the total exposure: $D_i = D \cdot \frac{E_i}{E_{\text{Total}}}$, where D is the total size of the default fund.

When deriving SITG components that are scaled by the Cover 2 fund size ($D_{s,2}$), we apply this principle by taking the relevant total default fund as $D = D_{s,2}$. The largest member's share of total exposure is its concentration ratio $c_1 = \frac{E_{(1)}}{E_{\text{Total}}}$. Therefore, the default fund contribution of the largest member, $D_{(1)}$, is:

$$D_{(1)} = D_{s,2} \frac{E_{(1)}}{E_{\text{Total}}} = D_{s,2}\, c_1 \,. \tag{3.7}$$

These relationships for $E_{(1)}$ and $D_{(1)}$ in terms of $D_{s,2}$ and concentration ratios are central to formulating the SITG amounts for the Cover 2 case.

3.4.2 First-Layer SITG

The first-layer SITG, S_2, is calculated using an ICC which says that the probability of excess losses from the largest defaulter (member (1)), $L_{(1)} = (U_{(1)} - M_{(1)})^+$, exceeding that member's own resources plus S_2 equals a target probability $\pi_2 \leq q_D$: $P(L_{(1)} > D_{(1)} + S_2) = \pi_2$. Applying the Pareto tail model and the defined parameters, the formula for S_2 is taken from the general S_n formula in Cont and Ghamami (2025, Appendix Eq. (A.16)) and adjusted for the Cover 2 case ($n = 2$) as:

$$S_2 = \left[\left(\frac{(q/\pi_2)^{1/\alpha} - 1}{(q/q_D)^{1/\alpha} - 1} \right) \left(\frac{c_1}{c_1 + c_2} \right) - c_1 \right] D_{s,2} \,. \tag{3.8}$$

The detailed derivation steps are shown in Appendix A.1 of the thesis.

3.4.3 Second-Layer SITG

The Second-Layer SITG, $\tilde{S}_2$, provides protection after the mutualised DF $D_{s,2}$ and S_2 are used up, based on a stricter target probability $\tilde{\pi}_2 \leq \pi_2$. The formula for

$\tilde{S}_2$ is adapted from the general $\tilde{S}_n$ formula in Cont and Ghamami (2025, Appendix Eq. (A.24)) for the Cover 2 case ($n = 2$) as:

$$\tilde{S}_2 = \left[\left(\frac{(q/\tilde{\pi}_2)^{1/\alpha} - (q/\pi_2)^{1/\alpha}}{(q/q_D)^{1/\alpha} - 1} \right) \frac{c_1}{c_1 + c_2} + c_1 - 1 \right] D_{s,2} \,. \tag{3.9}$$

A detailed derivation of Eq. (3.9) is shown in Appendix A.2 of the thesis.

3.4.4 Boundary Probability for the Second Layer

For $\tilde{S}_2$ to be positive when S_2 is sized with $\pi_2 = q_D$, the chosen $\tilde{\pi}_2$ must be less than a boundary value $\tilde{\pi}_{0,2}$. This boundary is the probability $P(\tilde{L}^{(1)}_{i,2} > D_{i,2})$ calculated when $\tilde{S}_2 = 0$ and $\pi_2 = q_D$. The formula for $\tilde{\pi}_{0,2}$ is given by

$$\tilde{\pi}_{0,2} = q\left[1 + \left(\left(\tfrac{q}{q_D}\right)^{1/\alpha} - 1\right)\left(1 + (c_1 + c_2)\left(\tfrac{1-c_1}{c_1}\right)\right)\right]^{-\alpha}. \tag{3.10}$$

A derivation of Eq. (3.10), which uses the Pareto tail assumption with these boundary conditions, is given in Appendix A.3.

3.4.5 Total SITG Contribution

The total required SITG under the Cover 2 standard is $S_2 + \tilde{S}_2$ and the formula (based on $S_n + \tilde{S}_n$ formula in Cont and Ghamami (2025, Appendix Eq. (A.25)) for the Cover 2 case ($n = 2$)) is given by

$$S_2 + \tilde{S}_2 = \left[\left(\frac{(q/\tilde{\pi}_2)^{1/\alpha} - 1}{(q/q_D)^{1/\alpha} - 1} \right) \frac{c_1}{c_1 + c_2} - 1 \right] D_{s,2} \,. \tag{3.11}$$

The formula in Eq. (3.11) depends on $\tilde{\pi}_2$, EVT parameters, concentration, and the Cover 2 default fund size. The algebraic derivation of Eq. (3.11) is shown in Appendix A.4 of this thesis.

3.5 Assumptions and Scope of the Model

The usefulness of the Cont and Ghamami (2025) framework depends on several key assumptions, which also define its scope and limitations. First, the model's results rely on the assumption that conditional excess losses $(U_i - M_i)^+$ follow the specific Pareto distribution given in Eq. (3.3) (Cont and Ghamami 2025, Assumption 4.1). While supported by EVT and empirical observations, real-world loss distributions differ, potentially reducing the accuracy of calculated SITG amounts.

The framework requires specifying several parameter inputs, which include the tail exponent α, the probability of breaching margin q, the unconditional tail probability q_D, the target probabilities for incentive compatibility π_2 and $\tilde{\pi}_2$, and the concentration ratios c_1, c_2. The calculated SITG levels are sensitive to the chosen values for these inputs, which need careful calibration and sensitivity analysis (as explored in Chapter 5).

A further limitation is the model's static nature; it calculates SITG requirements based on estimates made beforehand of exposures and probabilities. It does not inherently incorporate dynamic feedback loops, liquidity effects, contagion, or fire-sale dynamics that can greatly change losses during a real-world crisis.

Finally, the framework primarily focuses on determining the size of the CCP's pre-funded SITG layers (S_2 and $\tilde{S}_2$) within the default waterfall and does not consider subsequent recovery tools. Its strength lies in ensuring adequate CCP capital commitment before relying on subsequent recovery tools.

Despite these limitations, the framework offers a logical and analytical method for integrating EVT-based tail risk modelling with incentive compatibility constraints, thereby providing insights into appropriate SITG levels.

3.6 Additional Risk Metrics

While the original Cont and Ghamami (2025) framework provides SITG formulas, additional risk measures can enhance the understanding of CCP tail risks. In this section, we introduce complementary risk measures, namely, value-at-risk (VaR), conditional value-at-risk (CVaR), and conditional expected loss. These metrics are not specifically discussed in the original paper, but we introduce them to add more to our study.

The Pareto tail assumption, seen in Eq. (3.3), for a generic member i's excess loss $L_i = (U_i - M_i)^+$ (with its specific scale parameter κ_i and common tail exponent α) not only helps define κ_i but also lets us calculate other important risk measures. In these calculations, we use this member-specific κ_i and the associated tail exposure $E_i = \mathrm{VaR}_{q_D}(L_i)$ (where $q_D = P(L_i > E_i)$).

3.6.1 Value-at-Risk (VaR)

Value-at-Risk (VaR) is a widely used risk measure. McNeil, Frey and Embrechts (2015) define VaR in terms of a confidence level. For a given confidence level, denoted here as $p \in (0, 1)$ (to avoid confusion with the Pareto tail exponent α used throughout this thesis), the VaR of a portfolio with loss L at this confidence level p is "given by the smallest number l such that the probability that the loss L exceeds l is no larger than $1 - p$" (McNeil, Frey and Embrechts 2015, Eq. (5.20)). Formally, this is expressed as

$$\mathrm{VaR}_p(L) = \inf\{l \in \mathbb{R} : P(L > l) \leq 1 - p\} = \inf\{l \in \mathbb{R} : F_L(l) \geq p\} .$$

This definition relies on the concept of the quantile function. As detailed by McNeil, Frey and Embrechts (2015, Definition 2.9), for a given distribution function F_L, its generalized inverse, $F_L^{\leftarrow}(p)$, is called the quantile function of F_L. For $p \in (0, 1)$, the p-quantile of F_L is given by $q_p(F_L) := F_L^{\leftarrow}(p) = \inf\{l \in \mathbb{R} : F_L(l) \geq p\}$. Thus, $\mathrm{VaR}_p(L)$ is precisely the p-quantile of the loss distribution L, denoted $q_p(L)$ (McNeil, Frey and Embrechts 2015). In this formulation, $1 - p$ represents the tail probability, which is the probability that the loss L exceeds its $\mathrm{VaR}_p(L)$.

In the Cont and Ghamami (2025) setup, and consistently within this thesis, VaR is often characterised directly by this tail probability. The tail exposure E_i for a generic member i is defined as the VaR of the excess loss $L_i = (U_i - M_i)^+$ at an unconditional probability level q_D. Here, q_D directly represents the tail probability (i.e., $q_D = 1-p$ if using the notation from McNeil, Frey and Embrechts (2015) above). Thus, E_i is the value such that the probability of the excess loss L_i exceeding E_i is precisely q_D:

$$E_i = \mathrm{VaR}_{q_D}(L_i), \quad \text{such that} \quad P(L_i > E_i) = q_D\,.$$

This notation $\mathrm{VaR}_{q_D}(L_i)$ means that q_D is the probability with likelihood of surpassing a threshold.

The relationship between this E_i (which is $\mathrm{VaR}_{q_D}(L_i)$) and the member's scale parameter κ_i was given in Section 3.3, seen Eq. (3.5)

$$E_i = \kappa_i \left[\left(\frac{q}{q_D} \right)^{1/\alpha} - 1 \right]. \tag{3.12}$$

This equation shows how the size of the tail exposure E_i is determined by its scale parameter κ_i, the Pareto tail exponent α, and the probabilities q and q_D.

3.6.2 Conditional Value-at-Risk (CVaR)

Conditional value-at-risk (CVaR), also known as expected shortfall (ES), measures the expected value of a loss given that this loss L_i exceeds a certain VaR threshold. McNeil, Frey and Embrechts (2015, Eq. (2.23)), define expected shortfall at a confidence level p as $\mathrm{ES}_p(L) = \mathbb{E}[L \mid L > \mathrm{VaR}_p(L)]$.

In the context of this thesis, we are interested in the CVaR of the excess loss $L_i = (U_i - M_i)^+$. Its CVaR at the unconditional probability level q_D is defined relative to its VaR, $E_i = \mathrm{VaR}_{q_D}(L_i)$, which was introduced in Section 3.6.1, seen in Eq. (3.12). Formally, the CVaR is

$$\mathrm{CVaR}_{q_D}(L_i) = \mathbb{E}[L_i \mid L_i > \mathrm{VaR}_{q_D}(L_i)] = \mathbb{E}[L_i \mid L_i > E_i]$$

since $E_i = \mathrm{VaR}_{q_D}(L_i)$. The conditioning event $L_i > E_i$ implies $L_i > 0$, as E_i represents a VaR of positive excess losses.

Within the (Cont and Ghamami 2025) framework adopted in this thesis, which assumes that the excess loss L_i follows a distribution characterised by a Pareto tail with exponent $\alpha > 1$ (detailed in Section 3.3), the CVaR can be expressed in a specific relationship to its VaR, E_i. For such distributions, the CVaR is given by

$$\mathrm{CVaR}_{q_D}(L_i) = E_i \, \frac{\alpha}{\alpha - 1} \quad \text{where} \quad \alpha > 1\,. \tag{3.13}$$

Here, the α is the same Pareto tail exponent that characterises the assumed distribution of L_i. Further details on the properties and applications of Pareto distributions in risk management contexts can be further found in McNeil, Frey and Embrechts (2015).

The formula in Equation (3.13) indicates that the average excess loss L_i, conditional on it being in the q_D-tail (i.e., exceeding its VaR, E_i), is E_i scaled up by the factor $\frac{\alpha}{\alpha-1}$.

This factor, which is always greater than 1 for $\alpha > 1$, depends only on the tail exponent α from the underlying Pareto assumption. Recalling the definition of E_i from Eq. (3.12),

$$E_i = \kappa_i \left[\left(\frac{q}{q_D} \right)^{1/\alpha} - 1 \right],$$

the CVaR is determined by the parameters κ_i, q, q_D, and α of the tail model adopted from Cont and Ghamami (2025).

3.6.3 Conditional Expected Loss Beyond Threshold

The conditional expected loss beyond threshold estimates the average size of total losses $L_i = (U_i - M_i)^+$ for a member i given a specific general non-negative threshold u. The derivation relies on the mean excess function, defined as $e_i(x) = \mathbb{E}[L_i - x \mid L_i > x]$ for $x \geq 0$. For the conditional Pareto distribution defined by Eq. (3.3) (which is conditional on $L_i > 0$), the mean excess function is a result from EVT (McNeil, Frey and Embrechts 2015, Ch. 5), that is

$$e_i(x) = \frac{x + \kappa_i}{\alpha - 1}, \quad \text{for } \alpha > 1\,.$$

The conditional expectation $\mathbb{E}[L_i \mid L_i > u]$ (implicitly, $L_i > u, L_i > 0$) is the threshold u plus the mean excess above u, so that

$$\begin{aligned} \mathbb{E}[L_i \mid L_i > u] &= u + e_i(u) \\ &= u + \frac{u + \kappa_i}{\alpha - 1} \end{aligned}$$

which simplifies to

$$\mathbb{E}[L_i \mid L_i > u] = \frac{\alpha u + \kappa_i}{\alpha - 1}\,. \tag{3.14}$$

Equation (3.14) provides a way to estimate expected losses conditional on breaching past a certain level u. For instance, as explored in Chapter 6, u can be set to an observed SITG level (S_{Observed}) to assess potential further losses.

4

Trends in CCP Financial Resources

This chapter reviews the evolution of major CCPs in our sample, based on publicly disclosed PQD data between 2019 and 2024. The analysis focuses on three main aspects. First, Section 4.1 discusses the data source and the method used. Next, Section 4.2 explores overall trends in key CCP resources IM, DF, and the CCP's own SITG and examines how they respond to market stress. Then, Section 4.3 looks for notable patterns in how concentrated members are and how these resources relate to each other across different CCPs. The chapter concludes with a brief summary of the main observations in Section 4.4.

4.1 Data Source and Preparation

This section explains how we collected and cleaned the PQD data, outlines the sample used in this study, and notes any limitations or assumptions made.

The primary data source for this thesis were public quantitative disclosures from global CCPs. PQD data were sourced from the main disclosure pages of key CCP groups including: Eurex Clearing AG (Eurex Clearing AG 2025), CME Group (for CME Clearing services) (CME Group 2025), LSEG (for LCH Ltd and LCH SA services) (London Stock Exchange Group 2025), Nasdaq Clearing AB (Nasdaq Clearing 2025), and Intercontinental Exchange (ICE) (for ICE Clear Europe and ICE Clear Credit services) (Intercontinental Exchange 2025). The dataset for this analysis covers the period from Q1 2019 to Q4 2024. This dataset also serves as the empirical foundation for the analyses presented in Chapter 5 and Chapter 6.

The analysis focuses mainly on data derived from this sources for nine major CCP entities or their key clearing services, selected for their systemic relevance. These include Eurex Clearing AG (ECAG, Germany - offering clearing for a diverse range of products including derivatives), ICE Clear Europe (UK - a key clearer of Futures & Options), CME Clearing (USA - covering its Interest Rate Swaps clearing services), CME Clearing (USA - covering its Futures & Options clearing services), LCH Ltd (UK - covering its Equities clearing services), LCH Ltd (UK - as a leading clearer of Interest Rate Swaps, notably via SwapClear), Nasdaq Clearing AB (Sweden - primarily for its Commodities clearing services), ICE Clear Credit (USA - as a leading clearer of Credit Default Swaps), and LCH SA (France - for its Over-the-Counter Credit Default Swaps clearing service, notably via CDSClear).

While Regulation (EU) 2021/23 formally requires a second-layer SITG requirement only on EU-domiciled CCPs in the sample, namely Eurex Clearing AG, LCH SA, and Nasdaq Clearing AB, this analysis nonetheless includes all CCPs in its theoretical framework to be thorough. Non-EU CCPs, CME in the United States or ICE Clear Europe in the United Kingdom are not legally bound by the EU requirement.

As the sample includes CCPs reporting in multiple currencies (USD, EUR, and GBP), all monetary variables are converted into USD to ensure we can compare their actual amounts consistently. Historical quarterly end-of-period exchange rates were used from Yahoo Finance.

This analysis focuses on a set of PQD fields that reflect the CCPs' major risk-management resources. The main variables taken from the PQD templates, using Python, for the analyses in this thesis are defined in Table 4.1. These fields are drawn from the CPMI-IOSCO template and can be further clarified by guidance such as the CCP Global PQD FAQ Guide (CCP Global 2024).

Table 4.1: Key PQD Fields and Their Mapping to Thesis Variables

PQD Field(s)[a]	**Thesis Variable**	**Description**
6.1.1	Required initial margin (IM)	Total IM required from all participants, including any add-ons, as reported for the clearing service or CCP.
4.1.4	Required default fund (DF)	Total DF requirement for participants, explicitly excluding the CCP's own SITG contribution.
4.3.15	Held DF (DF_{held})	Market value of DF resources received from participants, reported pre-/post-haircut.
4.1.1+4.1.3[b]	Skin in the game (SITG)	Total prefunded CCP own capital in the default waterfall, measured as the sum of PQD fields 4.1.1 and 4.1.3. This sum is consistently used as the SITG, also referred to as "Actual SITG" or "SR_{Observed}" in the analysis.
18.3.2	Top 5 IM Concentration (%)	Fraction of total IM posted by the top five clearing members (house + client).
18.4.2	Top 5 DF Concentration (%)	Fraction of total DF contributed by the top five participants.
4.4.3[c]	E_1	Estimated largest stress loss from default of a single participant and affiliates (excess over IM). Referred to as Cover 1 Stress Loss and used as a proxy for the largest tail exposure, $E_{(1)}$.
4.4.7[c]	$D_{s,2}$	Estimated largest stress loss from default of two largest participants and affiliates (excess over IM). Referred to as Cover 2 Stress Loss and used as the empirical measure for the Cover 2 default fund size ($E_{(1)} + E_{(2)}$).

[a] Field numbers based on typical CPMI-IOSCO PQD templates CPMI-IOSCO (2015). Exact references can vary by CCP and jurisdiction. See also CCP Global (2024).

[b] PQD field 4.1.1 refers to 'Prefunded Own Capital Before' (a first layer of SITG), and 4.1.3 refers to 'Prefunded Own Capital After' (second layer of SITG)

[c] Based on standard CCP stress testing disclosures, this study uses 'Peak Day' over a trailing 12-month window.

4.2 CCP Resource Trends and Ratios

This section examines how these key resources and their ratios change over time, and whether they vary during times of market stress.

First, we examine trends in the core resource pools. Figure 4.1 shows how the median IM, DF (both required and held), and SITG change from 2019 to 2024 across the sample of our 9 CCPs. By observing Figure 4.1 we can see a fairly steady rise in IM and DF, especially during volatile markets, such as early 2020 and parts of 2022. Meanwhile, SITG levels barely change, at least when looking at the median. This implies that CCPs adjust margins and DFs when market conditions change, but their own capital share seems to stay the same.

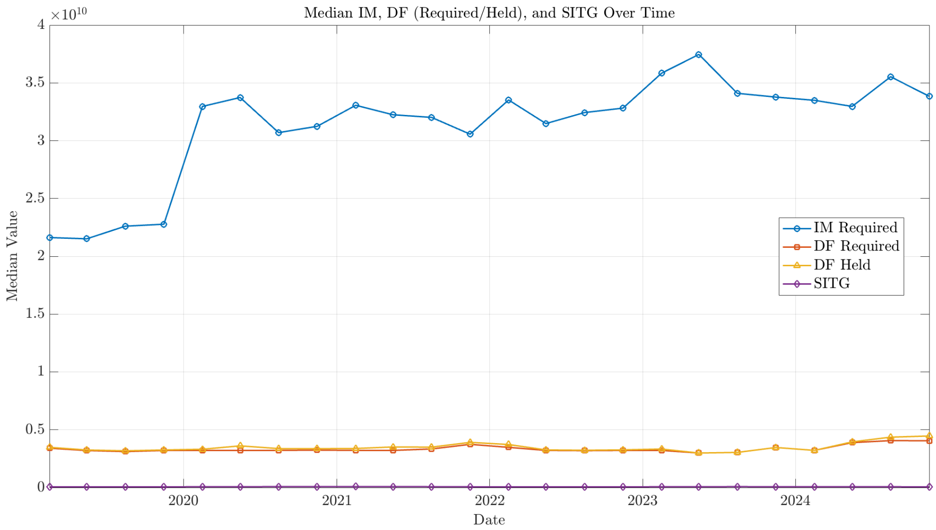

Figure 4.1: Median IM Required, DF Required, DF Held, and SITG (2019-2024)

To explore how member-funded resources compare to CCP-funded resources, we define a measure of "member pre-funded resources net of SITG", given as: $\text{IM}_{\text{required}} + \text{DF}_{\text{held}} - \text{SITG}$. This metric quantifies the extent to which pre-funded resources contributed by members (IM and held DF) exceed the CCP's own SITG contribution. Figure 4.2 shows the median of this measure over time. It indicates that member pre-funded resources net of SITG grows over time, suggesting that resources funded by members increase more rapidly than the CCP's own capital commitment.

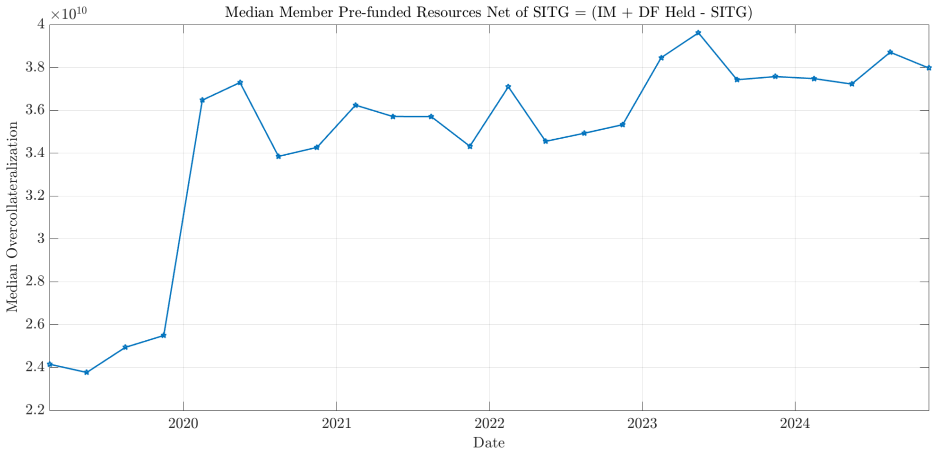

Figure 4.2: Median Member Pre-funded Resources Net of SITG (IM Required + DF Held - SITG) (2019-2024)

To further examine potential shifts in risk-bearing, we next compute two key ratios often discussed in the literature.

The first ratio, DF/SITG shows how large the DF is relative to the CCP's own stake, is displayed over the study period in Figure 4.3. It shows that the ratio rises from about 36x in 2019 to over 80x by the end of 2024. In other words, the member-funded share of the waterfall grows much more compared to the CCP's share. This finding is consistent with critiques that CCPs keep raising member requirements without also increasing their own SITG by a similar amount (Cont and Ghamami 2025; Pirrong 2011; FSB 2023).

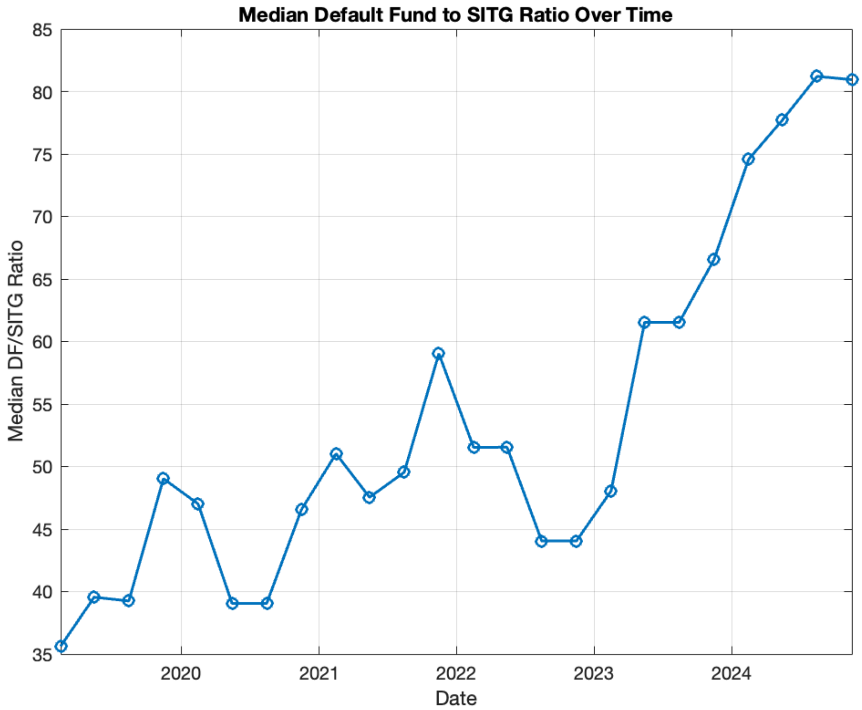

Figure 4.3: Median DF/SITG Ratio (2019-2024)

The second ratio, median IM/DF, illustrates how heavily the CCP relies on individual margining versus mutualised default protection and is displayed in Figure 4.4. It peaks around 2022 at about 18, then declines. While IM is almost always bigger than the DF, the spike around 2022 might reflect increased market volatility that led to larger margin calls (Duffie, Scheicher and Vuillemey 2015).

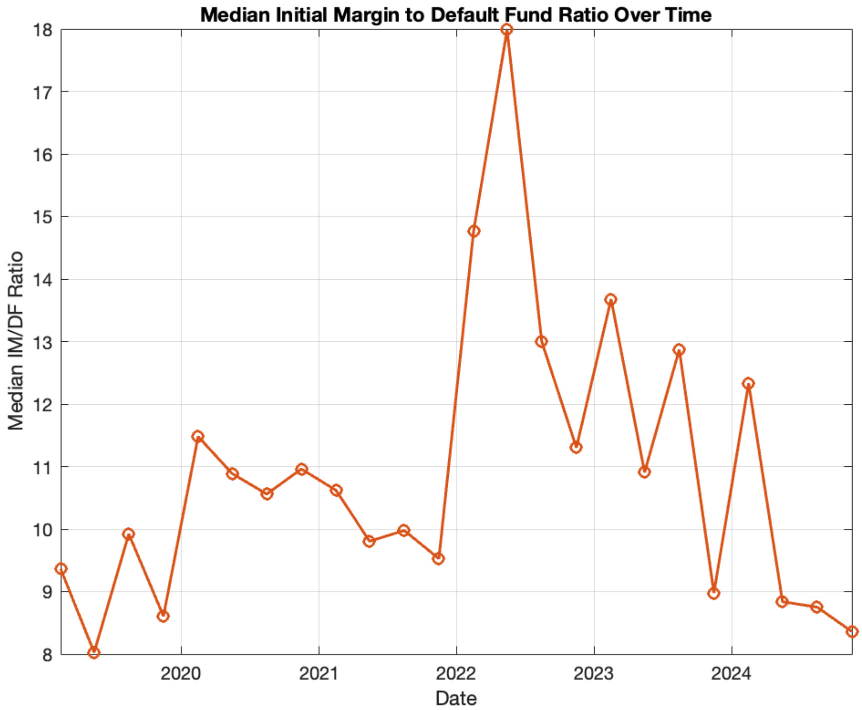

Figure 4.4: Median IM/DF Ratio (2019-2024)

4.3 Member Concentration and Exposure Patterns

This section examines more closely whether higher IM or higher participant concentration relates to larger SITG. The analysis mainly uses scatter plots of pooled quarterly data, drawn from the same dataset used in earlier sections and detailed in Section 4.1. Full regression outputs for the models discussed in this section are available in Appendix B.

Figure 4.5 shows a log-log plot of each CCP's reported SITG against its required initial margin. This scale is used to better visualize percentage changes and relationships that may follow exponential or power-law dynamics. The plot shows a positive trend, suggesting that CCPs with higher IM also tend to hold more SITG. However, a simple log-log regression ($\log(\text{SITG}) = \beta_0 + \beta_1 \log(\text{IM})$) suggests $\beta_1 \approx 0.49$ (p-value $< 10^{-20}$), which means SITG grows about half as fast as IM. The statistic measure that shows how well the regression line approximates our actual data, is $R^2 \approx 0.34$, and indicates other factors besides IM also play a role (see Appendix B.1 for detailed regression output).

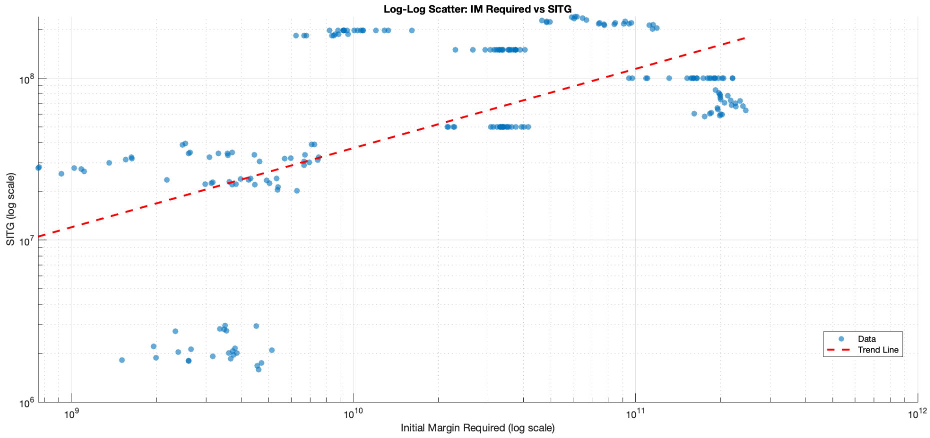

Figure 4.5: SITG vs. Required Initial Margin (Log-Log Scale, 2019-2024)

Next, Figure 4.6 plots the required DF against SITG, also on a log-log scale. Here, the estimated slope from a similar regression is about 0.89 (p-value $\approx 4\times10^{-47}$), with $R^2 \approx 0.62$ (detailed output in Appendix B.2). This is a stronger relationship than found with IM, yet is still sub-linear, as the slope is less than 1.0. Thus, while SITG increases as the DF grows, it does not increase proportionally.

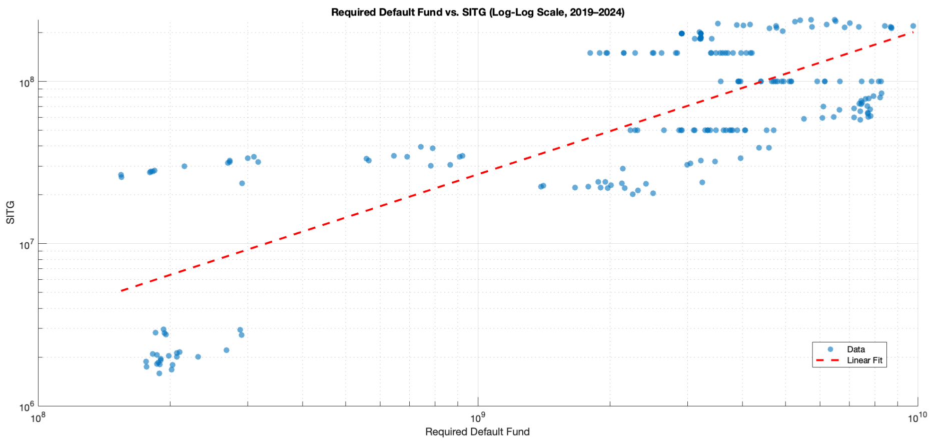

Figure 4.6: SITG vs. Required DF (Log-Log Scale, 2019-2024)

We also examine if CCPs with more concentrated membership, where the top few members hold a large fraction of IM or DF, tend to have higher SITG. Figure 4.7 compares SITG to the share of IM contributed by the top five members. A statistically significant but weak negative trend is observed (slope for Top5IM% $\approx -1.87 \times 10^8$, p-value for slope $\approx 1.5 \times 10^{-9}$, $R^2 \approx 0.157$ shown in detail in Appendix B.3, suggesting that higher IM concentration does not always mean CCPs put up more of their own capital.

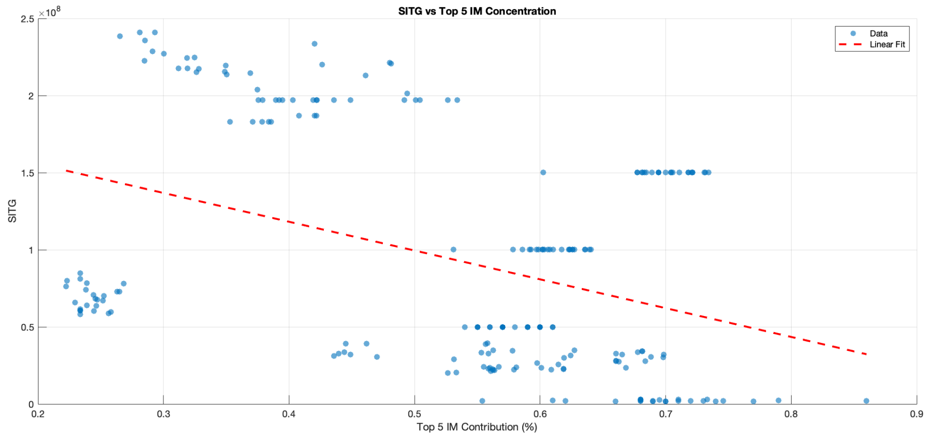

Figure 4.7: SITG vs. Top 5 IM Concentration (%), 2019-2024

Similarly, Figure 4.8 shows SITG compared to the fraction of the DF contributed by the top five participants. The slope in a simple linear fit is negative (slope for Top5DF% $\approx -1.78 \times 10^8$, p-value for slope $\approx 2.0 \times 10^{-9}$) and the overall $R^2 \approx 0.155$ is not very high (see Appendix B.4 for full details). This suggests that CCPs with especially concentrated DF membership do not increase their SITG in response, at least on average.

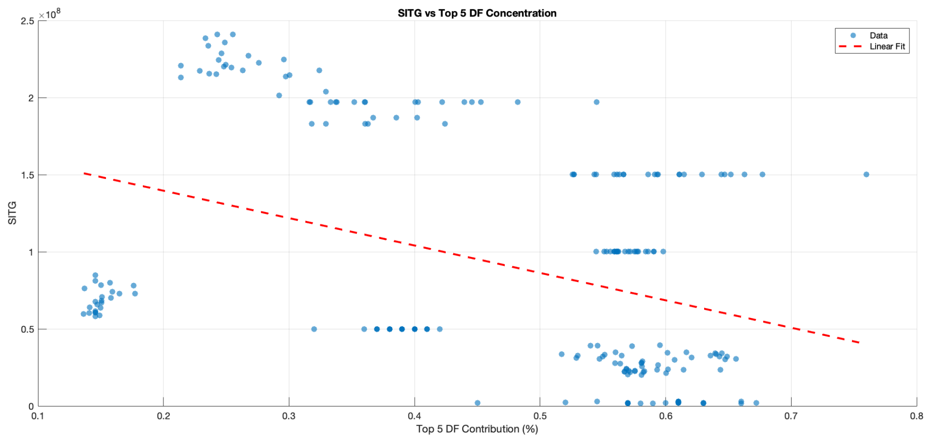

Figure 4.8: SITG vs. Top 5 DF Concentration (%), 2019-2024

Of course, each CCP has unique products and member bases, so the interpretation of these broad relationships should be made with caution. However, the data suggest that CCPs do not always hold more SITG when member exposures are heavily concentrated among a few participants.

4.4 Key CCP Resource Trends

From our empirical analysis of the 2019-2024 public quantitative disclosure (PQD) data, detailed in Sections 4.2 and 4.3, we make the following key observations.

Specifically, we observe relatively flat SITG levels in stress periods, as seen in Figure 4.1, even while IM and DF tend to spike when market volatility increases the SITG often remains generally unchanged. This trend, demonstrated by the growth in member pre-funded resources net of SITG, shown in Figure 4.2, also reflected in the rising DF/SITG ratio, in Figure 4.3, suggests that as member-funded resources increase during stress periods, the CCP's own capital staying mostly the same. Therefore, members end up covering a bigger portion of the risk during a crisis. Second, based on log-log regressions, SITG grows more slowly than either required IM or DF contributions, as illustrated in Figures 4.5 and 4.6, which raises questions on whether CCPs' own capital is keeping up as member-funded resources grow. Third, higher IM or DF concentration in the top five members does not seem to correlate with larger CCP capital contributions, seen in Figures 4.7 and 4.8. This observation makes us question how CCPs calibrate their own stake when fewer members share the risk.

The inherent limitations of PQD data, including inconsistent reporting intervals, currency conversions, and aggregated concentration metrics, are important to consider. The Cont and Ghamami (2025) model is applied, and its implications are further investigated in the subsequent chapters, but these initial observations establish the foundation for a more in-depth exploration of the theoretical arguments surrounding skin in the game.

5

The Cont and Ghamami (2025) SITG Model: Sensitivity and Empirical Implementation

This chapter applies the incentive compatibility framework from Cont and Ghamami (2025) to a typical Cover 2 default management setup used by CCPs. Our main goal is to determine the theoretical SITG requirements that the model in Cont and Ghamami (2025) implies, and then examine how those theoretical requirements compare to real reported data from CCPs between 2019-2024. This comparison helps us understand whether CCPs' actual capital stakes align align with overall financial stability needs.

We focus on three key model parameters, the tail exponent α, the largest-member concentration c_1, and the total top-two concentration r (where $r = c_1 + c_2$). In Section 5.1, we present the equations and discuss the assumptions of the parameters. Section 5.2 presents the results of our sensitivity analysis. In Section 5.3, we compare the theoretical SITG levels to those observed in PQD data. Finally, Section 5.4 summarises the main conclusions and limitations of this chapter.

5.1 Methodology and Parameter Specification

This section explains how we use the Cont and Ghamami (2025) model to estimate SITG for a Cover 2 scenario. We start with recalling the core formulas, describe key parameters, and note important assumptions, especially about the distribution of losses.

5.1.1 Model Formulas

Cont and Ghamami (2025) separate SITG into two layers within a CCP's default waterfall, in order to align the CCP's incentives with those of its clearing members. The first layer, denoted as S_2, is applied after the defaulting member's own resources (initial margin and default fund contribution) are gone, but importantly before the DF contributions of surviving members. The second layer, $\tilde{S}_2$, would be used only after surviving members' DF is used, but before more severe measures, such as unfunded assessments. In total, theoretical SITG is the sum of $S_2 + \tilde{S}_2$. These layers are designed to ensure the CCP has the right incentives, particularly in case of extreme losses that could impact the DF.

First we restate the Cover 2 setup (i.e., $n = 2$) given by Equations (3.8) and (3.9) from Chapter 3,

$$S_2 = \left[\left(\frac{(q/\pi_2)^{1/\alpha} - 1}{(q/q_D)^{1/\alpha} - 1} \right) \cdot \left(\tfrac{c_1}{c_1+c_2} \right) \; - \; c_1 \right] D_{s,2}$$

$$\tilde{S}_2 = \left[\left(\tfrac{(q/\tilde{\pi}_2)^{1/\alpha} - (q/\pi_2)^{1/\alpha}}{(q/q_D)^{1/\alpha} - 1} \right) \cdot \left(\tfrac{c_1}{c_1+c_2} \right) \; + \; c_1 \; - \; 1 \right] D_{s,2}$$

and the combined total theoretical SITG is therefore given by Equation (3.11), that is

$$S_2 + \tilde{S}_2 = \left[\left(\frac{(q/\tilde{\pi}_2)^{1/\alpha} - 1}{(q/q_D)^{1/\alpha} - 1} \right) \left(\frac{c_1}{c_1 + c_2} \right) - 1 \right] D_{s,2} \, .$$

Where $D_{s,2}$ is the Cover 2 default fund size, the sum of the two largest members' tail exposures ($E_{(1)} + E_{(2)}$), as defined in Chapter 3 (Section 3.4.1). Recall from Chapter 3, that the terms c_1 and c_2 are the exposure concentration ratios for the two largest members, and π_2 and $\tilde{\pi}_2$ are the target trigger probabilities for activating each SITG layer, which is also discussed below. Lastly, α is the tail exponent of the Pareto distribution assumed for excess losses $L_i = (U_i - M_i)^+$. The theoretical basis for this model, including the Pareto assumption, was discussed in Chapter 3 and the selection of α for this analysis is detailed in Section 5.1.2. We often analyse SITG as a ratio to $D_{s,2}$ (e.g., $S_2/D_{s,2}$, $\tilde{S}_2/D_{s,2}$) to assess its relative size.[1]

5.1.2 Parameter Selection

For the parameter values in Equations (3.8) and (3.9), we followed the guidance in Cont and Ghamami (2025) and typical CCP risk management practices.

Regarding the *Quantile Levels* (q, q_D), a CCP's IM is normally set to cover losses at a high confidence level, which often is around 99%, so $q = 0.01$. Meanwhile, DF is used for more extreme events, such as a 99.9% scenario, $q_D = 0.001$. This implies $q > q_D$, which is reasonable as the DF is there for bigger tail risks than the IM. In our baseline for the sensitivity analysis (Section 5.2), we set $q = 0.01$ and $q_D = 0.001$. Later, when we compare theory with real PQD data (Section 5.3), we also test $q_D = 0.005$.

The framework also requires "trigger probabilities" that determine when the first and second-layer SITG start, denoted as *Loss Probability Targets* $(\pi_2, \tilde{\pi}_2)$ *for Sensitivity Analysis*. In line with Cont and Ghamami (2025) we have $q_D \geq \pi_2 \geq \tilde{\pi}_2$. For consistency in our sensitivity analysis, we align the first SITG layer trigger (π_2) with the DF sizing probability (q_D), so $\pi_2 = q_D$. This ensures the CCP's first SITG is used under a similar degree of stress as when the DF is used.

For the second-layer SITG in the sensitivity analysis, $\tilde{\pi}_2$ must remain below a boundary $\tilde{\pi}_{0,2}$ to ensure model accuracy, such as non-negative $\tilde{S}_2$. This boundary, derived from the Cover 2 scenario where $\pi_2 = q_D$, is given by Eq. (3.10), the term

[1] This analysis focuses on standard multilayer default waterfalls in derivatives CCPs. However, Cont and Ghamami (2025) also examine monolayer setups (often found in some securities CCPs). Their results suggest such monolayer CCPs may require higher SITG to achieve similar incentive alignment.

$\left(2 - c_1 + \frac{c_2}{c_1} - c_2\right)$ used here is equivalent to the form $\left(1 + (c_1 + c_2)\left(\frac{1-c_1}{c_1}\right)\right)$, so that

$$\tilde{\pi}_{0,2} = q \times \left[1 + \left(\left(\frac{q}{q_D}\right)^{\frac{1}{\alpha}} - 1\right) \times \left(2 - c_1 + \frac{c_2}{c_1} - c_2\right)\right]^{-\alpha}.$$

In our sensitivity analysis in Section 5.2, we set $\tilde{\pi}_2 = 0.9\,\tilde{\pi}_{0,2}$, ensuring $\tilde{\pi}_2$ is safely below this theoretical limit. This maintains the order $\tilde{\pi}_2 < \tilde{\pi}_{0,2} \leq \pi_2 = q_D < q$. In is worth noting that for the empirical comparison in Section 5.3, a different, direct assumption for $\tilde{\pi}_2$ is used, as detailed there.

The *Tail Exponent (α)* is a critical element of this model, as it deals with extreme losses beyond the IM threshold. Cont and Ghamami (2025) assume these tail exposures follow a Pareto distribution, based on extreme value theory (EVT) and empirical observations of heavy tails in finance (McNeil, Frey and Embrechts 2015). A smaller α means fatter tails and a higher probability of extreme outcomes. We explore α between the values of 2.0 and 4.0 for our sensitivity analysis, a range commonly cited for financial data (Cont and Kan 2011). As noted in Chapter 4, PQD data is generally not detailed enough to reliably calculate α for each CCP directly. Thus, following Cont and Ghamami (2025)'s approach of assuming a Pareto tail, and examine SITG sensitivity across a plausible range for α.

For the *Concentration Ratios (c_1, c_2) and the Total Top-Two Concentration* r, which are key factors in determining risk, show the concentration of exposures among the largest members. The model incorporates $c_1$and c_2, largest and second largest members' share of E_{Total} respectively. The total concentration of the top two members is $r = c_1 + c_2$. Cont and Ghamami (2025) highlight that the relationship between concentration and SITG is not straightforward, as higher concentration does not always increase total SITG due to large members contribute more resources. The allocation between S_2 and $\tilde{S}_2$ is particularly sensitive to the relative sizes of c_1 and c_2. For our sensitivity analysis simulations, we vary c_1 testing values between $\{0.1, 0.3, 0.5\}$ and r, then derive $c_2 = r - c_1$, ensuring $c_1 \geq c_2 > 0$, as c_2 must be positive for a Cover 2 scenario with two distinct large defaulters, and that $r < 1$. This lets us look at different ways concentration can be structured, like when the top members are similar in size ($c_1 \approx c_2$) or when one is much larger.

Overall, these parameter choices for the sensitivity analysis reflect both standard industry practices and the theoretical constraints of the Cont and Ghamami (2025)'s model, allowing us to systematically check SITG requirements under different conditions.

5.1.3 Empirical Data Mapping

Chapter 4 , and more specifically Section 4.1, explains how we obtained actual SITG and operational estimates for stress losses that are used as E_1 and $D_{s,2}$ from PQD data. Briefly, we used the largest single- and two-member stress loss estimates, item 4.4.3 and 4.4.7 in PQD respectively, to obtain our empirical E_1 and $D_{s,2}$ values, and we estimated CCP SITG from PQD items 4.1.1 - 4.1.3.

5.1.4 Implementation Steps

In the sensitivity analysis detailed in Section 5.2, the procedure for each combination of input parameters (α, c_1, r) begins with the calculation of $c_2 = r - c_1$. The concentration parameters are then verified to ensure that $c_1 \geq c_2 > 0$ and $r < 1$ are plausible. Subsequently, the boundary probability $\tilde{\pi}_{0,2}$ is calculated using Eq. (5.1.2) with the specified parameters α, c_1, c_2, q, and q_D. Thereafter, the second-layer trigger probability is set as $\tilde{\pi}_2 = 0.9\,\tilde{\pi}_{0,2}$. Our selection of probability satisfies model constraints, such as $\tilde{\pi}_2 > 0$, and the condition $\pi_2 (= q_D) \geq \tilde{\pi}_{0,2}$, verifying that $\tilde{\pi}_{0,2}$ is valid and does not surpass q_D.

Provided all additional constraints are fulfilled, Eqs. (3.8) and (3.9) are solved for S_2 and $\tilde{S}_2$, ensuring non-negative results by applying $\max(0, \cdot)$ for each layer. Ultimately, the ratios $S_2/D_{s,2}$, $\tilde{S}_2/D_{s,2}$, and their sum, Total SITG / $D_{s,2}$, are saved for table and plot outputs. Only parameter combinations that produce valid solutions and satisfy all constraints are taken into account in the sensitivity results.

5.2 Sensitivity Analysis Results

This section examines how the theoretical estimated SITG levels, derived from the Cont and Ghamami (2025) model, are affected by changes in the parameters. Specifically, we analyse the impact of the tail exponent α and the total top-two member concentration $r = c_1 + c_2$. For this example analysis, the largest member's concentration is held constant at $c_1 = 0.3$, while r is varied, implying changes in $c_2 = r - c_1$ but ensuring that $c_2 > 0$ and $c_1 \geq c_2$. The tail exponent α ranges from 2.0 to 4.0, heavier to lighter tails respectively. Other parameters are held constant at $q = 0.01$, $q_D = 0.001$, and the target probability for the first SITG layer $\pi_2 = q_D$. The target probability for the second layer, $\tilde{\pi}_2$, is set to $0.9\tilde{\pi}_{0,2}$, as previously mentioned. It is important to note that the calculation of $\tilde{\pi}_{0,2}$, $\tilde{\pi}_2$ and $\tilde{S}_2$ layer in this setup, requires clear assumptions for individual concentrations c_1 and c_2.

Figure 5.1 presents these sensitivities through a series of line plots, showing each SITG component as a percentage of the Cover 2 default fund size, $D_{s,2}$.

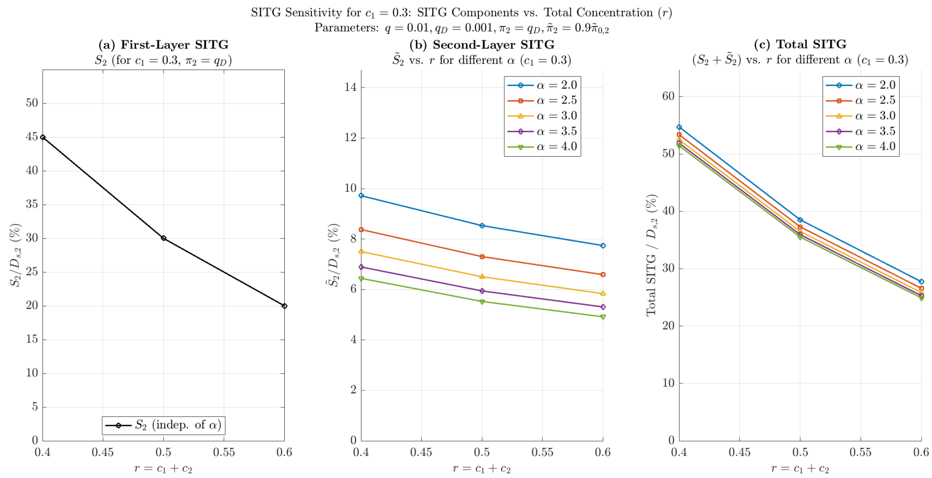

Figure 5.1: Theoretical SITG components (as % of $D_{s,2}$) for a fixed largest member concentration $c_1 = 0.3$, illustrating sensitivity to total top-two concentration r and tail exponent α. Parameters: $q = 0.01, q_D = 0.001, \pi_2 = q_D, \tilde{\pi}_2 = 0.9\tilde{\pi}_{0,2}$. **Panel (a)**: First-layer SITG (S_2) vs. r. Due to $\pi_2 = q_D$, S_2 is independent of α. **Panel (b)**: Second-layer SITG ($\tilde{S}_2$) vs. r, with distinct lines for different α values. **Panel (c)**: Total SITG ($S_2 + \tilde{S}_2$) vs. r, with distinct lines for different α values.

Panel (a) of Figure 5.1 shows the first-layer SITG, S_2. A key insight from Equation (3.8) is that when $\pi_2 = q_D$, the term $\frac{(q/\pi_2)^{1/\alpha}-1}{(q/q_D)^{1/\alpha}-1}$ simplifies to 1. This makes S_2 independent of the tail exponent α. As a result, S_2 is determined by the concentration parameters c_1, which is fixed at 0.3 in this specific setup, and r. The plot clearly shows that for a fixed c_1, the $S_2/D_{s,2}$ ratio decreases as the total top-two concentration r increases. Specifically, S_2 is approximately 45% of $D_{s,2}$ when $r = 0.4$, falling to 30% when $r = 0.5$, and to 20% when $r = 0.6$. This suggests that with higher overall concentration r, meaning c_2 is larger relative to c_1, the model requires a smaller first SITG layer from the CCP, as the defaulting members' own resources, weighted by c_1 in the S_2 formula, become relatively more important.

Panel (b) shows how sensitive the second-layer SITG, $\tilde{S}_2$, is to both α and r. Unlike S_2, $\tilde{S}_2$ is more affected by the tail exponent α. For any given level of r, $\tilde{S}_2/D_{s,2}$ is much higher for smaller α values, where tails are heavier, and decreases as α increases, when tails are lighter. For instance, when $r = 0.5$, $\tilde{S}_2$ can be over 8% of $D_{s,2}$ for $\alpha = 2.0$, but it drops to around 5% for $\alpha = 4.0$. This indicates the core principle that heavier tails imply a greater likelihood of extreme losses, that could exhaust the first layers of the default waterfall, thus meaning a larger $\tilde{S}_2$ is needed from the CCP to maintain incentive alignment for these more rare events. Evident by the plotted range is also the dependence on r, with $\tilde{S}_2$ generally lacking a monotonic relationship with r for fixed α, affected by the way it interacts with $\tilde{\pi}_{0,2}$.

The behaviour of Total SITG, shown in Panel (c), is a combination of its two layers. Since $\tilde{S}_2$ is sensitive to α, Total SITG also decreases as α increases, meaning that tails become lighter. For example, at $r = 0.4$, Total SITG is approximately 55% of $D_{s,2}$ for $\alpha = 2.0$, but drops to about 51% for $\alpha = 4.0$. The impact of r is also significant, as for a given α at 2.5, Total SITG decreases from roughly 53% at $r = 0.4$ to about 26% at $r = 0.6$. This shows that the overall capital the CCP needs to commit is influenced by tail risk perception (α) and the way its largest members (c_1, r) are concentrated.

As outlined, the sensitivity analysis performed in this chapter, determines $\tilde{\pi}_2$ by linking it to $\tilde{\pi}_{0,2}$, which is the boundary probability ensuring that $\tilde{S}_2 \geq 0$. The calculation of $\tilde{\pi}_{0,2}$, seen on Eq. (5.1.2), always requires specific values for c_1 and c_2. Thus, the values of S_2 and $\tilde{S}_2$ presented here depend on the concentration assumptions, $c_1 = 0.3$ and c_2 satisfies the plausibility conditions. Estimating total SITG with an exogenously set $\tilde{\pi}_2$, meaning not tied to $\tilde{\pi}_{0,2}$ which depends on individual concentrations, would change the determination of $\tilde{\pi}_2$ depending on the assumptions. For example, the total SITG formula, Eq. (3.11), primarily features the ratio $\frac{c_1}{c_1+c_2}$. Therefore, by changing the setting of $\tilde{\pi}_2$, changes how specific the assumptions need to be. In our empirical comparison in Section 5.3, $\tilde{\pi}_2$ is set differently, which affects how concentration inputs are used to estimate total SITG.

For a broader understanding, Table 5.1 presents results for various c_1 values beyond 0.3, offering more detailed numerical comparisons across different α, c_1, and r. Several entries illustrate how both S_2 and $\tilde{S}_2$ change under heavier or lighter tails. For instance, at $\alpha = 2.0$ and $c_1 = 0.1$, one row shows that if $r = 0.2$, the first-layer SITG S_2 is around 40%, whereas if r increases to 0.4, that same α produces $S_2 = 45\%$, and the total SITG rises above 50%. Such numeric details help clarify that even small changes in concentration or tail exponent can move total SITG by 5-10 percentage points. Overall, testing different α and c_1 combinations confirms the Cont and Ghamami (2025) model's sensitivity, that a higher top member ratio tends to reduce S_2 but may still result in a significant $\tilde{S}_2$.

Table 5.1: General Theoretical SITG Sensitivity to α, c_1, and r

α	r	c_1	c_2	$\tilde{\pi}_{0,2}$	$\tilde{\pi}_2$	$S_2(\%)$	$\tilde{S}_2(\%)$	Total SITG (%)
2.0	0.2	0.1	0.1	0.00020	0.00018	40.00	8.82	48.82
2.0	0.4	0.3	0.1	0.00037	0.00034	45.00	9.72	54.72
2.0	0.5	0.3	0.2	0.00031	0.00028	30.00	8.53	38.53
2.0	0.6	0.3	0.3	0.00026	0.00023	20.00	7.74	27.74
2.0	0.6	0.5	0.1	0.00050	0.00045	33.33	9.30	42.63
2.0	0.7	0.5	0.2	0.00046	0.00041	21.43	8.36	29.78
2.0	0.8	0.5	0.3	0.00042	0.00038	12.50	7.65	20.15
2.0	0.9	0.5	0.4	0.00038	0.00034	5.56	7.10	12.66
2.5	0.2	0.1	0.1	0.00016	0.00014	40.00	7.45	47.45
2.5	0.4	0.3	0.1	0.00033	0.00030	45.00	8.38	53.38
2.5	0.5	0.3	0.2	0.00026	0.00024	30.00	7.30	37.30
2.5	0.6	0.3	0.3	0.00022	0.00020	20.00	6.59	26.59
2.5	0.6	0.5	0.1	0.00046	0.00042	33.33	8.11	41.45
2.5	0.7	0.5	0.2	0.00042	0.00037	21.43	7.26	28.69
2.5	0.8	0.5	0.3	0.00037	0.00034	12.50	6.62	19.12
2.5	0.9	0.5	0.4	0.00034	0.00030	5.56	6.13	11.68
3.0	0.2	0.1	0.1	0.00013	0.00012	40.00	6.55	46.55
3.0	0.4	0.3	0.1	0.00030	0.00027	45.00	7.51	52.51
3.0	0.5	0.3	0.2	0.00023	0.00021	30.00	6.50	36.50
3.0	0.6	0.3	0.3	0.00019	0.00017	20.00	5.84	25.84
3.0	0.6	0.5	0.1	0.00043	0.00039	33.33	7.35	40.68
3.0	0.7	0.5	0.2	0.00038	0.00035	21.43	6.55	27.98
3.0	0.8	0.5	0.3	0.00034	0.00031	12.50	5.96	18.46
3.0	0.9	0.5	0.4	0.00031	0.00028	5.56	5.49	11.05
3.5	0.2	0.1	0.1	0.00011	0.00010	40.00	5.92	45.92
3.5	0.4	0.3	0.1	0.00027	0.00025	45.00	6.89	51.89
3.5	0.5	0.3	0.2	0.00021	0.00019	30.00	5.94	35.94
3.5	0.6	0.3	0.3	0.00016	0.00015	20.00	5.31	25.31
3.5	0.6	0.5	0.1	0.00041	0.00037	33.33	6.81	40.14
3.5	0.7	0.5	0.2	0.00036	0.00033	21.43	6.06	27.48
3.5	0.8	0.5	0.3	0.00032	0.00029	12.50	5.49	17.99
3.5	0.9	0.5	0.4	0.00028	0.00025	5.56	5.05	10.61
4.0	0.2	0.1	0.1	0.00010	0.00009	40.00	5.45	45.45
4.0	0.4	0.3	0.1	0.00025	0.00023	45.00	6.44	51.44
4.0	0.5	0.3	0.2	0.00019	0.00017	30.00	5.53	35.53
4.0	0.6	0.3	0.3	0.00015	0.00013	20.00	4.92	24.92
4.0	0.6	0.5	0.1	0.00039	0.00035	33.33	6.42	39.75
4.0	0.7	0.5	0.2	0.00034	0.00031	21.43	5.69	27.12
4.0	0.8	0.5	0.3	0.00030	0.00027	12.50	5.15	17.65
4.0	0.9	0.5	0.4	0.00026	0.00024	5.56	4.72	10.28

Notes: Table shows theoretical SITG percentages calculated for various valid combinations of α, $r = c_1 + c_2$, and c_1, enforcing $c_1 \geq c_2 > 0$. Fixed parameters: $q = 0.01$, $q_D = 0.001$, $\pi_2 = q_D$. $\tilde{\pi}_2$ is set to $0.9\tilde{\pi}_{0,2}$. Rows with $c_2 = 0$ are excluded.

5.3 Comparing Theoretical and Observed SITG

To investigate how these SITG levels compare with actual CCP's SITG, we used quarterly data from 2019 Q1 to 2024 Q4 for the sample of major CCPs described in Chapter 4. Specifically, for each CCP and each quarter, we obtained its reported "Actual SITG" or "Observed SITG" , its Cover 2 default fund size, and the largest single-member stress loss estimate. We then computed Actual SITG as a percentage of $D_{s,2}$ for each CCP quarter.

Meanwhile, to estimate total theoretical SITG ($S_2 + \tilde{S}_2$) for each CCP for this empirical comparison, we used Equation (3.11). The parameter choices used are that the tail exponent is fixed at $\alpha = 3.0$, the standard quantile level for initial margin as $q = 0.01$, and two scenarios for the DF sizing quantile q_D: first, $q_D = 0.001$ (99.9% confidence, green bars in Figure 5.2), and second, $q_D = 0.005$ (99.5% confidence, red bars). For the purpose of this empirical comparison, the trigger probability for the Second-Layer SITG, $\tilde{\pi}_2$, was set to $\tilde{\pi}_2 = q_D/5$ for each respective q_D scenario. The concentration ratio term $\frac{c_1}{c_1+c_2}$ required in Eq. (3.11) was represented for each CCP in each quarter by $\frac{E_1}{D_{s,2}}$, using their reported largest single-member stress loss, E_1, relative to its reported Cover 2 stress loss $D_{s,2}$.[2] With these parameters, the total theoretical SITG ($S_2 + \tilde{S}_2$) was calculated as a percentage of the CCP's empirical $D_{s,2}$ for that quarter using Eq. (3.11). It is worth noting that this way of calculating total SITG does not need estimation of individual c_1 and c_2 values, only their ratio $E_1/D_{s,2}$. The resulting theoretical SITG values were then averaged across all quarters for each of the 9 CCPs used in our data set and the outcome of this comparison is shown in Figure 5.2.

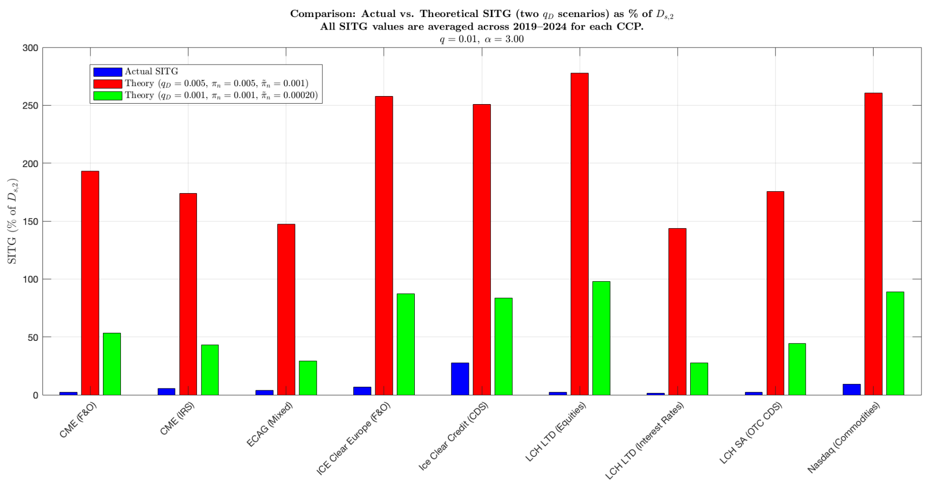

Figure 5.2: Average Actual SITG vs. Theoretical SITG (as % of $D_{s,2}$) for selected CCPs, 2019-2024. Theoretical based on Cont and Ghamami (2025) with $\alpha = 3.0$, $q = 0.01$, and two scenarios for q_D.

At first glance in Figure 5.2, most CCPs' actual SITG, shown by the blue bars, is

[2] This proxy assumes $E_1 \approx E_{(1)}$ (the largest member's tail exposure relevant for the model) and $D_{s,2} \approx E_{(1)} + E_{(2)}$ (the sum of the top two tail exposures) from actual stress test disclosures. Thus, $E_1/D_{s,2} \approx E_{(1)}/(E_{(1)} + E_{(2)})$, which, if $E_{(k)} = c_k E_{Total}$, equals $c_1/(c_1 + c_2)$.

much lower than both theoretical estimates. Meanwhile, the red bars corresponding to $q_D = 0.005$ are even higher than the green bars, which represent $q_D = 0.001$. The way the total SITG formula and the assumed $\tilde{\pi}_2 = q_D/5$ interact with other fixed parameters, such as $\alpha = 3.0$ and $q = 0.01$, leads to CCPs would need more capital with these specific settings. Indeed, many CCPs have red bars that go over 100-200% of $D_{s,2}$, while actual SITG is often under 20%.

The observed gap matches the concerns raised in Cont and Ghamami (2025) that if a CCP invests only a small fraction of its own capital, it may rely heavily on resources funded by members and have less reason to deal with extreme risks. As already pointed out in Chapter 4, some CCPs disclose SITG of about 1%, much lower than a heavy-tailed loss might need (Thiruchelvam 2022; Walker 2023). Moreover, because real CCP rules may tie SITG to a static percentage of regulatory capital (European Commission 2013), SITG can remain unchanged even as exposures grow or the risk of extreme losses increases.

From a policy standpoint, the gap between actual and theoretical SITG raises questions about moral hazard and whether incentives are properly aligned (Cont and Ghamami 2025). It also suggests that banks' exposures to CCP DFs might be treated as less risky than they are under current Basel standards (BIS 2023). As noted in Cont and Ghamami (2025), if one adopts a monolayer waterfall instead of a separate DF, the model indicates even higher SITG might be needed to keep incentives aligned.

Overall as seen in Figure 5.2, whether $q_D = 0.005$ or $q_D = 0.001$, actual CCP SITG frequently falls short of the model's suggestions by a wide margin. CCPs' own 'skin' may be too small to ensure robust incentives in the face of heavy-tailed risk. In addition, the lack of improvement over time for the studied period, as seen in Chapter 4, suggests a structural issue rather than a momentary data problem. This could lead to more debates among regulators, especially since some regulations link CCP SITG to a fixed percentage of regulatory capital, not to the actual risk exposures.

5.4 Key Model Insights of Cont and Ghamami (2025) SITG Framework

This chapter applied the Cont and Ghamami (2025) theoretical framework to a Cover 2 default management structure. We explored how the model's SITG requirements vary with the tail exponent α, the largest-member concentration ratio c_1, and different choices of the DF quantile q_D. We then compared these theoretical SITG levels to actual CCP data from 2019-2024. The analysis suggests four main points.

Firstly, a significant gap often exists between theory and practice, as many CCPs hold far less SITG than the model recommends under the tested parameters, often by factors of 5 to 10 or more. There was little indication that this gap narrowed over the 2019-2024 period, as discussed in Chapter 4, pointing to a structural rather than temporary shortfall.

Secondly, the model demonstrates high sensitivity to tail risk, where heavier-tailed assumptions (smaller α) generally produce significantly higher SITG requirements. This is particularly evident for the second layer $\tilde{S}_2$ when its trigger $\tilde{\pi}_2$ is determined via the $\tilde{\pi}_{0,2}$ boundary mechanism, as seen Eq. (5.1.2), which is sensitive to α. This fits with the model's goal of encouraging the CCP to consider extreme losses that go beyond what members' DF contributions cover.

Thirdly, there are complexities to how concentration is measured. The distribution of exposures, specifically the comparison of c_1 to c_2, influences how the total SITG is split between the first layer S_2 and the second layer $\tilde{S}_2$, as described by Eqs. (3.8) and (3.9). A very large top member might post substantial IM/DF which can reduce the CCP's first-layer SITG requirement (S_2), but the framework may still call for a significant second layer ($\tilde{S}_2$) if the overall DF is at risk of depletion during a severe event.

Finally, the interaction of model parameters can lead to non-intuitive outcomes. For example, since $\pi_2 = q_D$ directly influences the first-layer SITG (S_2) primarily via the concentration terms, as the probability ratio term in Eq. (3.8) simplifies to 1 when $\pi_2 = q_D$, changing q_D has a bigger effect on the second layer ($\tilde{S}_2$). In our empirical comparison calibrations, where $\tilde{\pi}_2$ was set directly to $q_D/5$, setting $q_D = 0.005$ can yield higher total SITG requirements than $q_D = 0.001$. This outcome comes from how the term $(q/\tilde{\pi}_2)^{1/\alpha}$ in the total SITG formula, Eq. (3.11), responds to q_D when $\tilde{\pi}_2$ is a fixed fraction (1/5) of q_D. This contrasts with the sensitivity analysis where $\tilde{\pi}_2$ is calculated using the more complex $\tilde{\pi}_{0,2}$, as shown in Eq. (5.1.2). Thus, a seemingly 'less conservative' DF sizing quantile does not automatically translate to lower SITG under all parts of this model's specific incentive-based structure and chosen parameters for the empirical comparison.

The above model insights demonstrate its sensitivity to various parameters and the clear difference from CCP's observed practices. The specific limitations about how the model was used in this chapter, such as the necessary parameter assumptions and data proxies used for comparison, are important to note. Broader implications of these findings, the overall study limitations, and potential avenues for future research are discussed in detail in Chapter 7.

6

SITG and the Implied Tail Risk

Building on the foundational understanding of CCPs in Chapter 2 and the SITG framework in Chapter 3, this chapter further investigates the implications of the Cont and Ghamami (2025) SITG framework, using the Pareto-based tail exponent (α). While previous Chapter 5, considered α as a fixed parameter, the current chapter tries to identify the implied tail exponent, α_{Implied}, that aligns the observed SITG levels to the model's predictions from Cont and Ghamami (2025).

This perspective, which we define as "implied alpha", aims to show how heavy or light tailed a CCP's loss distribution would need to be, according to the Cont and Ghamami (2025) model, to justify its actual SITG. A higher α_{Implied} suggests thinner tails, meaning less extreme risk, whereas a lower α_{Implied} indicates fatter tails, greater extreme risk.

We organise this chapter as follows: Section 6.1 outlines the data and the estimation method, looking at different sets of parameters. Subsequently, Section 6.2 presents and interprets the main empirical results, focusing on the way of setting parameters that gave the most consistent results, while also highlighting sensitivities through comparisons. The chapter concludes in Section 6.3 with a summary of our key findings.

6.1 Estimation Method and Inputs

Data Sources and Key Variables

As previously stated in Chapter 4, the dataset for this study consists of quarterly public quantitative disclosures from nine selected CCPs, covering the studied period 2019-2024. For each CCP-quarter observation, the analysis uses three key variables, the Observed SITG (SR_{Observed}), the Cover 1 Stress Loss (E_1), and the Cover 2 Stress Loss ($D_{s,2}$).

Inverting the SITG Function and Parameter Sets Explored

The Cover 2 total SITG formula from the Cont and Ghamami (2025) framework, explained in Chapter 3, Section 3.4.5 , is central to this inversion. Taking Equation (3.11), in which we denote the expected Cover 1 stress loss as E_1, and $D_{s,2}$ represents the Cover 2 stress loss used as a baseline for total SITG requirements.

We set the total SITG $(S_2 + \tilde{S}_2)$ as $h_2(\alpha)$,

$$h_2(\alpha) = \left[\left(\frac{(q/\tilde{\pi}_2)^{1/\alpha} - 1}{(q/q_D)^{1/\alpha} - 1} \right) \cdot \left(\frac{E_1}{D_{s,2}} \right) - 1 \right] D_{s,2} \,. \tag{6.1}$$

Next, we want to find the value of α_{Implied} that makes the theoretical SITG function $h_2(\alpha)$ equal to SR_{Observed}, which is the actual SITG, by solving for α. We find the root of $f(\alpha) = h_2(\alpha) - SR_{\text{Observed}}$ numerically, searching for α values between 1.01 and 50. Equation (6.1) thus gives us the theoretical way to estimate α_{Implied}.

Since the model's outcomes are sensitive to its chosen probability parameters, we explored several different configurations to see how much our findings changed with different settings. Three parameter compositions were explored to check for sensitivity, primarily by varying the DF exhaustion probability q_D and the fixed target probability for the second-layer SITG, $\tilde{\pi}_2$. The probability of losses going over IM was held constant at $q = 0.01$.

The initial setup, Parameter Set 1 (Original Baseline), employed $q_D = 0.0050$ and $\tilde{\pi}_2 = 0.0040$. Parameter Set 2 (Alternative) utilised $q_D = 0.0050$ together with $\tilde{\pi}_2 = 0.0030$. Ultimately, Parameter Set 3 (with High-Convergence Focus) implemented $q_D = 0.0010$ and $\tilde{\pi}_2 = 0.0005$.

The rationale for primarily focusing the detailed presentation of results, in the following Section 6.2, on Parameter Set 3 with its conservative assumptions is its significantly higher success rate in giving valid α_{Implied} estimates, a point explained more in Table 6.1. Nonetheless, findings from Parameter Sets 1 and 2 also offer a valuable benchmark.

Related Risk Metrics

The Pareto-tail assumption, central to the EVT approach, Chapter 3, Section 3.3, lets us calculate key risk measures once α_{Implied} is estimated. Alternatively, if α_{Implied} cannot be derived for comparative purposes, a predetermined α value is used, $\alpha = 3.0$. For these calculations, the parameters E_1, q, and q_D are taken directly from the empirical data for each CCP quarter, corresponding to the specific set of parameters used to estimate α_{Implied}.

One of these measures is the Pareto scale parameter (κ), computed using the found α_{Implied} and the observed E_1, based on the relationship derived from Eq. (3.12) in Chapter 3,

$$\kappa = \frac{E_1}{\left(\frac{q}{q_D}\right)^{\frac{1}{\alpha_{\text{Implied}}}} - 1} \,. \tag{6.2}$$

Another significant statistic is the conditional value-at-risk (CVaR at q_D). This metric is computed following Equation (3.13) from Chapter 3,

$$\text{CVaR}_{q_D}((U - M)^+) = \left(\frac{\alpha_{\text{Implied}}}{\alpha_{\text{Implied}} - 1} \right) E_1 \,. \tag{6.3}$$

We determine the conditional expected loss beyond SITG using Equation (3.14) from Chapter 3, where the threshold u is set to SR_{Observed},

$$E[\text{Loss} \mid \text{Loss} > SR_{\text{Observed}}] = \max\left(0, SR_{\text{Observed}} + \frac{\kappa + SR_{\text{Observed}}}{\alpha_{\text{Implied}} - 1}\right). \tag{6.4}$$

6.2 Implied Alpha Results

All figures in this chapter reflect either Set 3, and Set 2 or Set 1 for robustness comparisons, as seen in Table 6.1, with specific parameter values noted in the captions. The graphs are specific for each CCP unless noted otherwise.

Comparative Success of Parameter Sets and Implied Alpha Statistics

The selection of the fixed probability parameters, q_D and $\tilde{\pi}_2$, significantly affects both whether we can estimate α_{Implied} and the statistical properties of its distribution. A summary of estimation success rates and key descriptive statistics for α_{Implied} under the three parameter sets is shown in Table 6.1.

Table 6.1: Implied Alpha Estimation: Success Rates and Descriptive Statistics

Parameter / Statistic	**Parameter Set 1**	**Parameter Set 2**	**Parameter Set 3**
q	0.0100	0.0100	**0.0100**
q_D	0.0050	0.0050	**0.0010**
$\tilde{\pi}_2$	0.0040	0.0030	**0.0005**
Success Rate	$\approx$ 16.7%	$\approx$ 36.7%	$\approx$ 75.8%
N (Successful Estimations)	36	79	**163**
α_{Implied} - Minimum	1.053	1.299	**1.066**
α_{Implied} - Mean	3.656	4.501	**2.675**
α_{Implied} - Maximum	47.871	31.204	**35.014**

Notes: Success rate based on 215 valid CCP-quarter observations. α_{Implied} *searched in the range* $[1.01, 50.0]$. *Parameter Set 3 (High-Convergence Focus) is used for the main detailed analysis.*

The results in Table 6.1 indicate that Parameter Set 3, which is characterised by $q = 0.01, q_D = 0.0010, \tilde{\pi}_2 = 0.0005$, works much better than the others, finding an α_{Implied} about 75.8% of the time. This high success rate likely indicates the model equation more frequently yields a numerically stable solution when using more conservative parameter settings, suggesting that using stricter assumptions for q_D and $\tilde{\pi}_2$ makes it possible to find a consistent α_{Implied} for more of the observed SITG levels. These more conservative assumptions suggest that many observed SITG values are easier to match with the Cont and Ghamami (2025) framework when the model assumes a regime with very low probability events, such as rare DF being used. With Set 3 in Table 6.1 , the mean α_{Implied} was approximately 2.7, much lower than the means from Set 1 (mean 3.656) and Set 2 (mean 4.501). This pattern suggests that parameter choices that help find stable solutions efficiently, also tend to link observed SITG levels with heavier tail risk, meaning a lower α_{Implied}. For the 52 instances, out of 215 valid observations, where no solution was found under Set 3, the numerical solver did not identify an intersection between the theoretical SITG function $h_2(\alpha)$ and the observed SITG when searching for α between 1.01 and 50.

Since Parameter Set 3 performed best out of all parameter setups, we focus our detailed analysis on it. Figure 6.1 provides a comparison of the average implied alphas for each CCP, showing results from all three parameter sets.

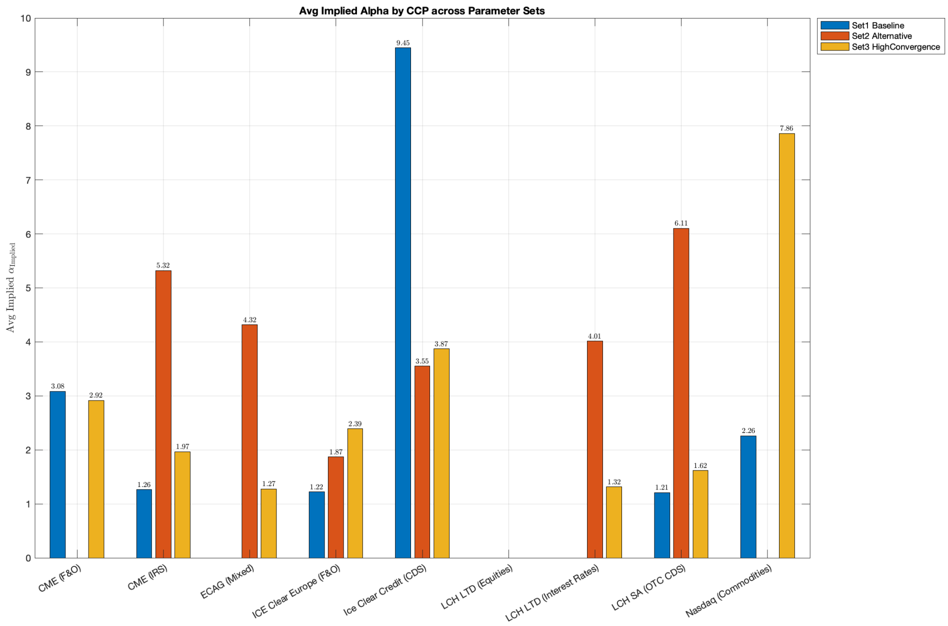

Figure 6.1: Average Implied Alpha by CCP Across Parameter Sets, highlighting how implied tail risks differ across CCPs and parameter choices

Notes: Bars illustrate how implied tail risks differ across CCPs, comparing average α_{Implied} *across the three parameter sets for quarters with valid solutions.*

Cross-Sectional Variation of Implied Alpha

In Figure 6.2, Panel (a) presents the mean α_{Implied} for each CCP as estimated using Parameter Set 3 from Table 6.1. To compare the sensitivity, Panel (b) is also used to show the corresponding results from Parameter Set 2, showing how these average implied alphas change when the model parameters are different. Big differences are observed across the CCPs' average α_{Implied}, for both parameter settings.

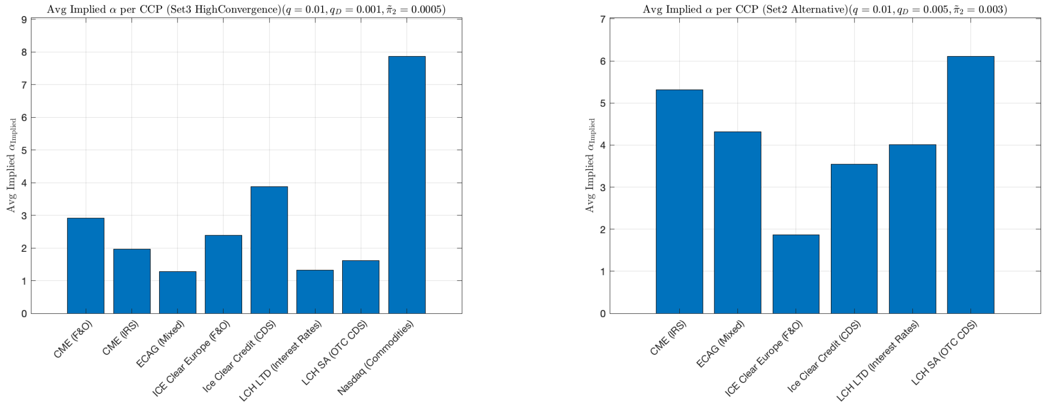

(a) Primary Analysis (Set 3: $\tilde{\pi}_2 = 0.0005, q_D = 0.0010$)

(b) Alternative (Set 2: $\tilde{\pi}_2 = 0.0030, q_D = 0.0050$)

Figure 6.2: Mean Implied Alpha by CCP: A Comparison of Parameter Sets, illustrating cross-CCP heterogeneity in implied risk. **Panel (a)** shows the Primary Analysis using Parameter Set 3. **Panel (b)** presents an Alternative Analysis using Parameter Set 2.

Notes: Bars in Figure 6.2, represent each CCP's average α_{Implied} for quarters with valid solutions under the respective parameter set, revealing differing implied risk profiles. Some CCPs do not appear in Panel (a) due to lack of valid implied alpha solutions under Set 3 (e.g., LCH LTD (Equities)).

Focusing on Panel (a) of Figure 6.2 that shows Parameter Set 3, the variation in average α_{Implied} is significant. For example, our calculations for Set 3 have a mean α_{Implied} of 1.275 for ECAG (Mixed) and 1.316 for LCH LTD (Interest Rates). Both values are quite low, which suggests that, with these parameter settings, the model implies very high chances of extreme events. Other CCPs, such as CME (IRS) with a mean α_{Implied} of 1.968 and LCH SA (OTC CDS) with 1.618, also show low average implied alphas. In contrast, Ice Clear Credit (CDS) shows a higher mean α_{Implied} of 3.873, and Nasdaq (Commodities) an even higher average of 7.862. These latter figures suggest their observed SITG values match up with a lower chance of extreme events when we use these specific model settings. In this context, the variation in α_{Implied} is not merely statistical but likely shows very different ways of setting SITG or different types of exposures. Nasdaq's high implied alpha might suggest a conservative SITG buffer relative to modelled risk, while CCPs like ECAG seem to imply that the model assumes they face frequent, severe losses. CCPs like CME (F&O), with a mean of 2.918, and ICE Clear Europe (F&O), having 2.389 as a mean, fall in between. It is noteworthy that for LCH LTD (Equities) we could not find a successful estimate for it in that calculation. The overall mean α_{Implied} of approximately 2.7 for Set 3 in Table 6.1suggests a general pattern, that for SITG levels to be consistent with the model using these parameters, the model often needs to assume a high level of extreme risk. As expected, Panel (b), displaying results for Set 2, generally shows higher average α_{Implied} values.

In summary, average α_{Implied} varies greatly by CCP, even under consistent assumptions. Some CCPs appear to operate with SITG levels that the model sees as needing very high extreme risk. This could mean they either face more risk, or there is a gap between the model's assumptions and practical implementation.

Visual Fit: $h_2(\alpha)$ vs. Observed SITG

Figure 6.3 offers a visual comparison of the model's fit to observed data under different parameters by plotting the theoretical $h_2(\alpha)$ function, given in Eq. (6.1), against average observed SITG levels. Panel (a) uses Parameter Set 3, illustrates $h_2(\alpha)$ curves often crossing the average observed SITG levels for different CCPs, which indicates the high success rate for thie parameter setup. Panel (b) in Figure 6.3, uses Parameter Set 1 for contrast, where fewer such intersections are found.

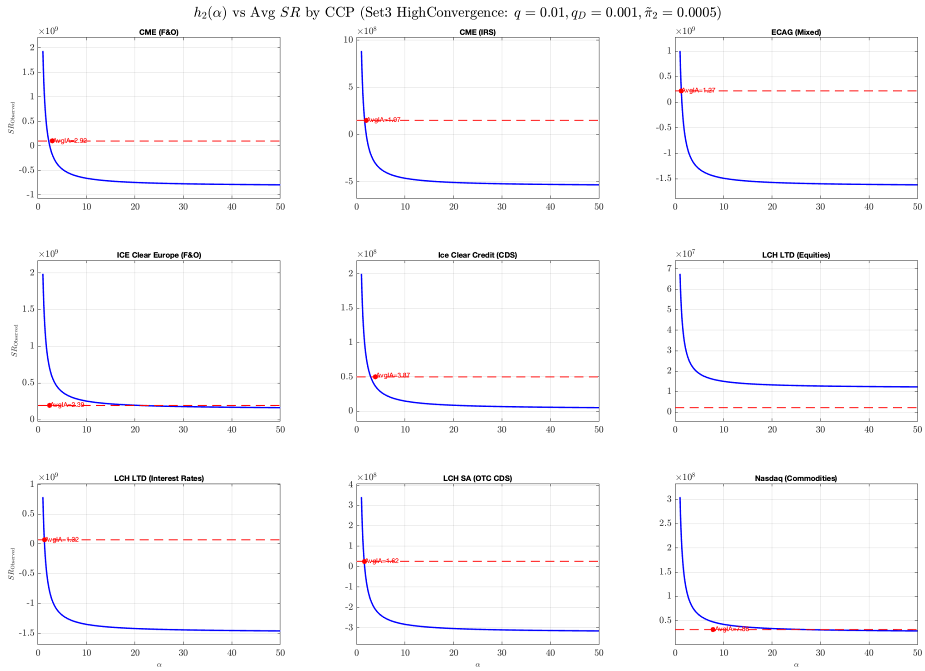

(a) Primary Analysis (Set 3: $\tilde{\pi}_2 = 0.0005, q_D = 0.0010$)

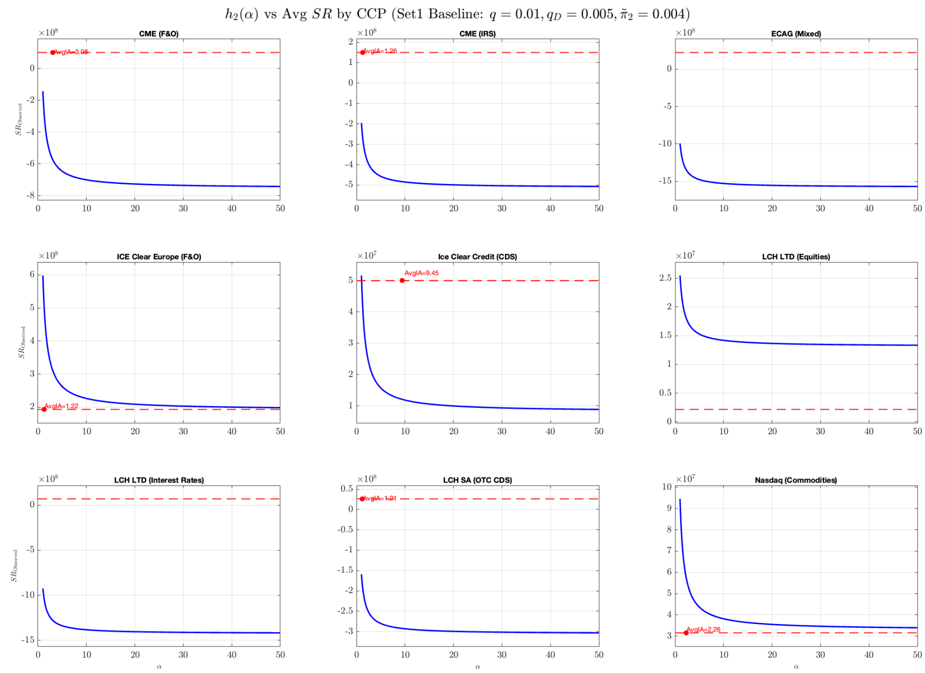

(b) Comparison (Set 1: $\tilde{\pi}_2 = 0.0040, q_D = 0.0050$)

Figure 6.3: Comparison of Theoretical $h_2(\alpha)$ (Blue) and Average Observed SITG (Red) by CCP, illustrating model fit and solution frequency. **Panel (a)** displays this comparison under the Primary Analysis parameters (Set 3). **Panel (b)** uses an alternative set of parameters for comparison (Set 1).

Notes: Each subplot in Figure 6.3 uses average values for $E_1, D_{s,2}$ specific to each CCP, along with their average observed SITG levels (red dashed line). The $h_2(\alpha)$ curve is drawn using the probability parameters from the set shown.

The $h_2(\alpha)$ curves under Parameter Set 3, as shown in Panel a, typically exhibit higher and more varied positioning. This characteristic seems to make it easier to find crossings with observed SITG than under Parameter Set 1, as shown in Panel b, where the curves are often lower and flatter, resulting in fewer solutions when attempting to match positive SITG values. The visual evidence in Panel b supports the higher convergence rate noted for Parameter Set 3. For some CCPs under Set 1, as shown in Panel b, the observed SITG lies entirely above or below the $h_2(\alpha)$ curve, meaning no implied alpha could be found, which shows how some parameter choices can make it mathematically impossible to match the real-world data.

Time-Series Variation and Distribution of Implied Alpha (Parameter Set 3)

To understand how the implied tail exponents change over time more clearly, Figure 6.4 presents the quarterly α_{Implied} estimations for each CCP using data from Parameter Set 3. This shows how they behaved over time when we could find solutions.

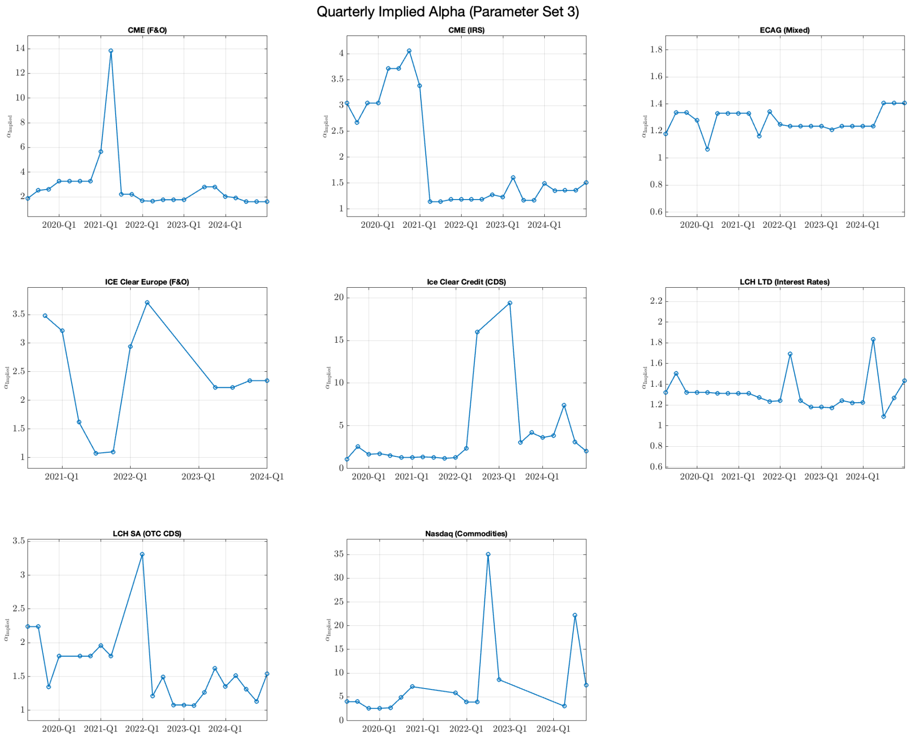

Figure 6.4: Quarterly Implied Alpha (2019 Q1 to 2024 Q4) with Parameter Set 3, showing temporal patterns in implied tail risk

Notes: Each subplot shows the quarterly α_{Implied} derived using Parameter Set 3 when a solution was found, revealing shifts in implied tail risk over time.

The time-series plots in Figure 6.4, which are based on Parameter Set 3, have more data points because this set was more successful at finding solutions to the equation $SR_{\text{observed}} = h(\alpha)$ stated in Eq. (6.1). For several CCPs, including ECAG (Mixed) and LCH LTD (Interest Rates), the α_{Implied} values are frequently low, and often below 2. This suggests the model interprets these CCPs as having high extreme risk when these parameters are used. Conversely, other CCPs, such as Ice Clear Credit (CDS) and Nasdaq (Commodities), generally keep higher α_{Implied} values for the whole period. These time series not only show the persistently low or high implied tail risk across CCPs, but also reflect sudden shifts during certain quarters, for example, sharp drops in α_{Implied} for CME (F&O) or spikes for Nasdaq around 2022. These shifts may indicate stress events in specific markets, or unusual SITG or stress loss reported in those quarters.

For a deeper understanding of the α_{Implied} distribution from Parameter Set 3 in Table 6.1, Figure 6.5 provides a histogram of all 163 successful α_{Implied} estimations. Complementing this, Figure 6.6 presents box plots showing the distribution of these quarterly estimations for each individual CCP.

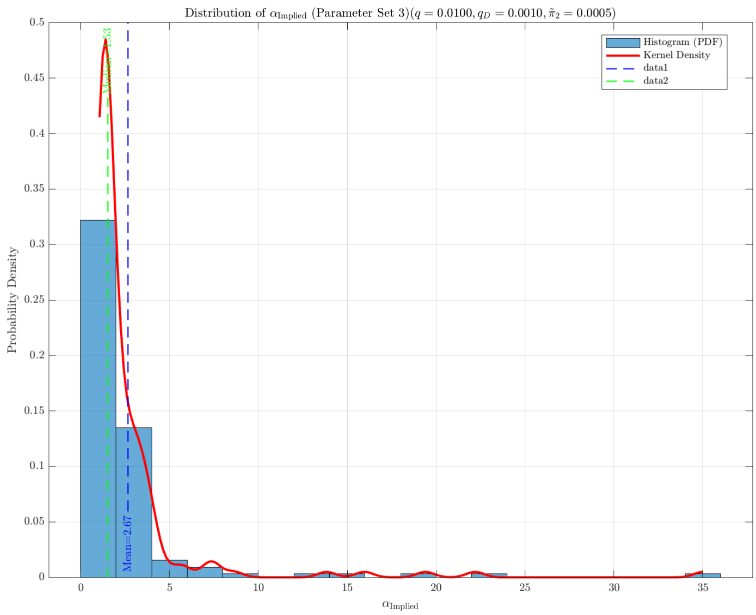

Figure 6.5: Overall Distribution of Implied Alpha Estimations. This histogram displays all 163 successful α_{Implied} values obtained using Parameter Set 3, revealing the overall shape and skewness of the estimations.

Notes: The histogram shows a right-skewed distribution, with a concentration of α_{Implied} *values below 5, consistent with the mean of approximately 2.7 for this parameter set.*

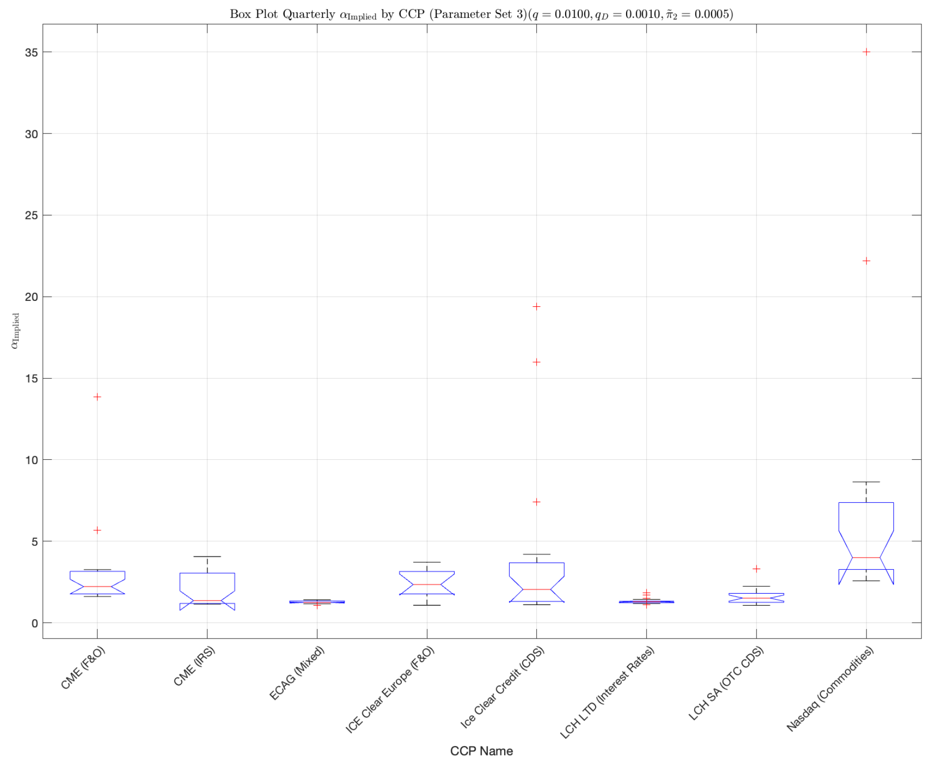

Figure 6.6: Distribution of Quarterly Implied Alpha Estimations by CCP, using Parameter Set 3. These box plots illustrate the distribution of quarterly α_{Implied} values for each CCP, highlighting inter-CCP differences and variations in central tendency and spread.

Notes: The box plots illustrate both the variation within each CCP in implied tail exponents over time and the notable differences in these distributions across the CCPs. For instance, CCPs like Nasdaq show a wide spread, while ECAG and LCH LTD (IR) exhibit more concentrated distributions.

The overall histogram in Figure 6.5 is right-skewed, with a clear concentration of α_{Implied} values below 5 and a long tail extending towards higher values. This observation matches the mean we mentioned earlier of approximately 2.7 for this parameter set. The box plots in Figure 6.6 illustrates both the variation within each CCP in implied tail exponents over time and the notable differences in these distributions across the CCPs in the sample. Boxplots show large spreads for Nasdaq and more tighter, less spread-out distributions for ECAG and LCH LTD (IR), which again shows the big differences in the tail risks suggested by the model. Notably, CCPs like Nasdaq show a wide spread of values, while ECAG and LCH LTD (IR) show consistently low and stable implied alphas, again supporting the idea that the model assumes ongoing high tail risk for them.

Derived Risk Metrics with Parameter Set 3: κ, CVaR, and Conditional Expected Loss

Having previously derived the implied alpha (α_{implied}) values for each CCP under Parameter Set 3, which represent the model-consistent tail exponents implied by their observed SITG levels (SR_{Observed}), as seen in Section 6.1, we now leverage these α_{implied} estimates to explore a set of related risk metrics.

The motivation for deriving these additional metrics is to further quantify and interpret the risk profiles suggested by each CCP's α_{implied}. Specifically, by using the α_{implied} obtained from Parameter Set 3, we can calculate the corresponding implied Pareto scale parameter (κ), conditional value-at-risk (CVaR_{q_D}), and conditional expected loss ($\mathbb{E}[L \mid L > SR_{\text{Observed}}]$). These derived values provide a comprehensive, nuanced perspective of the tail risk characteristics consistent with each CCP's observed SITG within the Cont and Ghamami (2025) framework.

Furthermore, analysing these metrics, particularly in comparison to values derived using a benchmark α of 3.0, allows for a more detailed understanding of the implications of the implied tail heaviness and the sensitivity of risk assessments to the tail exponent. This comparative analysis can highlight the quantitative impact of different tail assumptions on perceived risk.

We begin by examining the relationship between the implied Pareto scale parameter, κ, and the Cover 1 stress loss, E_1, in Figure 6.7. A weaker, though still positive, correlation, approximately 0.27, between the implied scale parameter κ and E_1 is evident from the scatter plot in Figure 6.7. This relationship mostly matches how they are linked by definition in Eq. (6.2). However, the scattering of the points also indicates that variations in α_{Implied} strongly affect κ.

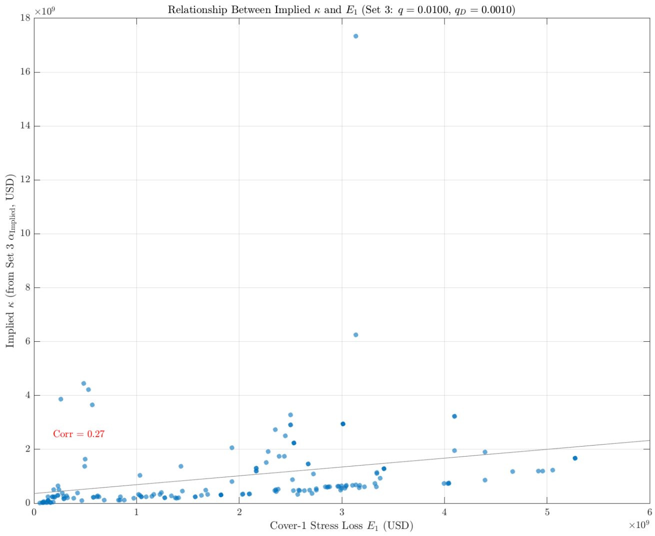

Figure 6.7: Relationship Between Implied κ and Cover 1 Stress Loss E_1

Notes: Each point represents a CCP-quarter where α_{Implied} was computed using Parameter Set 3. The correlation coefficient is approximately 0.27.

Table 6.2 illustrates the key results when we compare risk measures. These metrics are calculated using, the CCP's average α_{Implied}, where they are available and greater than 1, and a benchmark value of $\alpha = 3.0$. All underlying average values (SR_{Observed}, E_1) and the derived metrics come from our summarised calculations for the Parameter Set 3 analysis.

Table 6.2: Comparative Risk Metrics: Average α_{Implied} vs. Benchmark $\alpha = 3.0$

CCP Name	SR_{Obs}	E_1	α_{Impl}	CVaR_{q_D} **(Impl)**	CVaR_{q_D} $(\alpha = 3.0)$	$E[L \mid L > SR]$ **(Impl)**	$E[L \mid L > SR]$ $(\alpha = 3.0)$
CME (F&O)	0.100	3.463	2.918	5.270	5.195	1.655	1.650
CME (IRS)	0.150	1.800	1.968	3.661	2.700	1.142	1.005
ECAG (Mixed)	0.222	3.324	1.275	15.418	4.986	3.407	1.773
ICE Clear EU (F&O)	0.193	2.313	2.389	3.979	3.470	1.359	1.291
Ice Clear CR (CDS)	0.050	0.247	3.873	0.333	0.371	0.174	0.182
LCH LTD (Equities)	0.002	0.070	NaN	NaN	0.105	NaN	0.034
LCH LTD (IR)	0.069	2.850	1.316	11.880	4.275	2.186	1.338
LCH SA (OTC CDS)	0.026	0.833	1.618	2.181	1.249	0.497	0.400
Nasdaq (Comm)	0.032	0.351	7.862	0.401	0.526	0.186	0.199

Notes: All values expressed in USD billions. CVaR *and conditional expected loss are calculated using both the implied α and a fixed $\alpha = 3.0$. NaN denotes to no successful α_{Implied} estimation was achieved for that CCP.*

The data in Table 6.2 shows CCPs where Parameter Set 3 yields an average α_{Implied} that implies heavier tails, because $\alpha_{\text{Implied}} < 3.0$, such as ECAG (Mixed) with $\alpha_{\text{Implied}} \approx 1.3$ or LCH LTD (IR) with $\alpha_{\text{Implied}} \approx 1.3$. Their model consistent CVaR is much higher than the CVaR calculated using the benchmark $\alpha = 3.0$. For instance, ECAG's implied CVaR is 15.418 bn USD versus 4.986 bn USD under the benchmark. In contrast, for CCPs where α_{Implied} is close to or above 3.0, for example CME (F&O) with $\alpha_{\text{Implied}} \approx 2.9$ or Ice Clear CR with $\alpha_{\text{Implied}} \approx 3.9$, the implied

CVaR can be lower than or similar to the benchmark CVaR. This highlights that the choice of alpha, whether implied by the data or assumed as a benchmark, greatly affects the estimated risk of extreme events. The CVaR, as defined in Eq. (6.3), is $E_1\left(\frac{\alpha_{\text{Implied}}}{\alpha_{\text{Implied}}-1}\right)$, the multiplier term $\frac{\alpha_{\text{Implied}}}{\alpha_{\text{Implied}}-1}$ increases sharply as α_{Implied} approaches 1, which explains why the CVaR values are so large when implied alpha is low. These results underscore that SITG estimates and derived risk metrics are highly sensitive to tail-risk assumptions, which could be a challenge to setting risk management policies in practice.

Figures 6.8 and 6.9 provide a graphical representation of these comparisons, comparing SITG with risk metrics derived from Parameter Set 3 averages and the benchmark value of $\alpha = 3.0$.

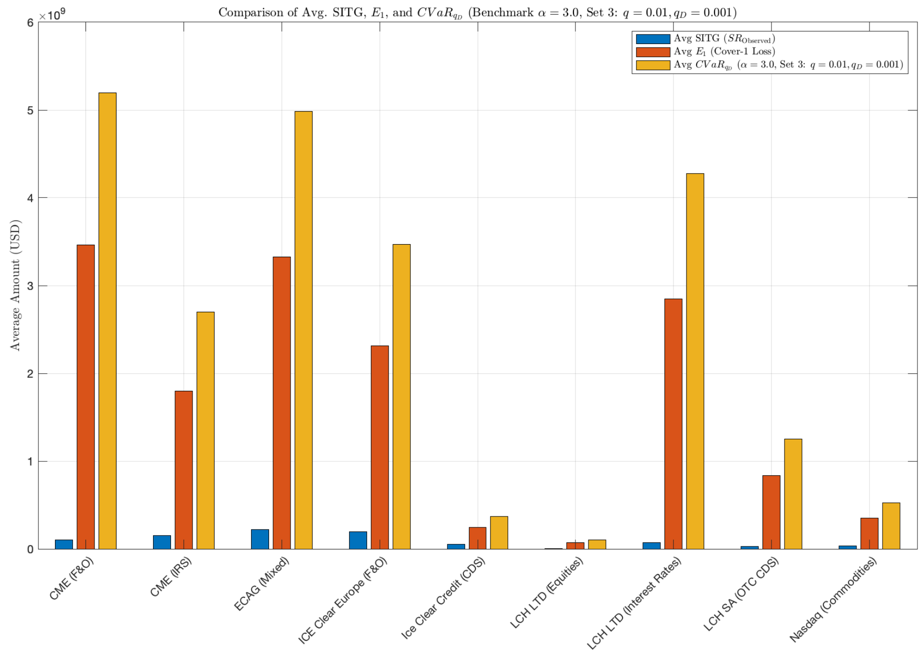

Figure 6.8: Comparison of Average SITG, E_1 (VaR), and CVaR ($\alpha = 3.0$), illustrating the relative magnitudes of different risk layers

Notes: Bars show mean SR_{Observed}*, mean* E_1*, and mean* CVaR*. Averages for* SR_{Observed} *and* E_1 *are based on Set 3 data. Values for* CVaR *are derived using the benchmark* $\alpha = 3.0$ *to standardize comparison.*

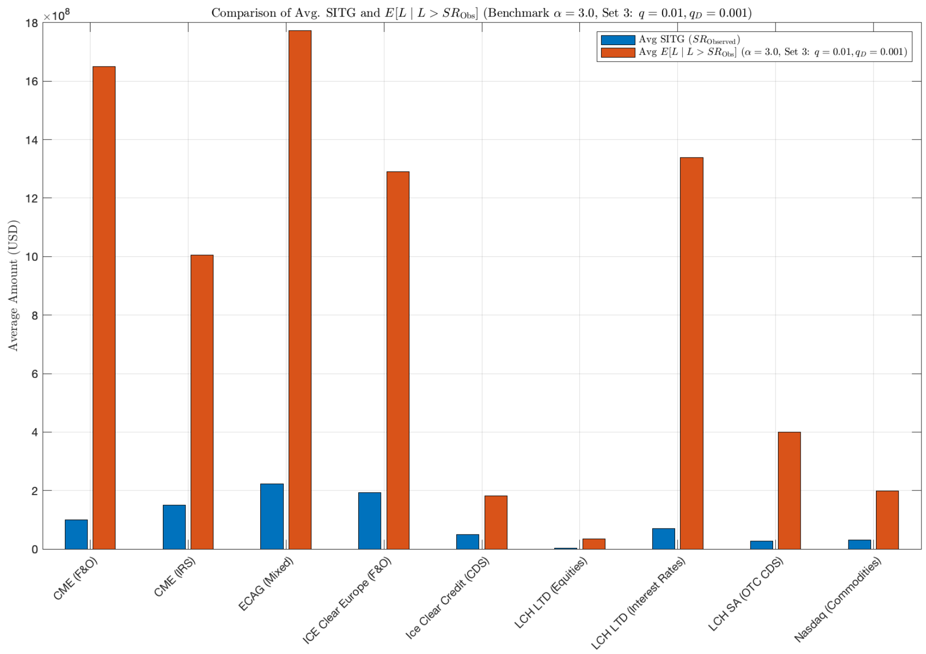

Figure 6.9: Average SITG vs. Conditional Expected Loss Beyond SITG ($\alpha = 3.0$), highlights potential uncovered tail exposures

Notes: Bars illustrate mean SR_{Observed} *and mean conditional expected loss* $E[Loss \mid Loss > SR_{\text{Observed}}]$. *Values for* $E[Loss \mid Loss > SR_{\text{Observed}}]$ *are derived using the benchmark* $\alpha = 3.0$ *to standardize comparison.*

The visualisations in Figures 6.8 and 6.9 clearly show that, even when applying a common benchmark value like $\alpha = 3.0$, the estimated CVaR and conditional expected losses are usually much higher than the average observed SITG levels reported by the sampled CCPs.

Limit Analysis

We explore the behaviour of $h_2(\alpha)$ at its limits from Eq. (6.1) to better understand why success rates differ across parameter sets. Using the sample averages for $E_1 \approx 1.702 \times 10^9$ USD and $D_{s,2} \approx 2.733 \times 10^9$ USD, our analysis gave varying limit values.

For Parameter Set 3 (with $\tilde{\pi}_2 = 0.0005$, $q_D = 0.0010$), the limits were $\lim_{\alpha \to 1^+} h_2(\alpha) \approx 8.600 \times 10^8$ and $\lim_{\alpha \to \infty} h_2(\alpha) \approx -5.189 \times 10^8$.

For Parameter Set 2 (where $\tilde{\pi}_2 = 0.0030$, $q_D = 0.0050$), we determined that $\lim_{\alpha \to 1^+} h_2(\alpha) \approx 1.238 \times 10^9$ and $\lim_{\alpha \to \infty} h_2(\alpha) \approx 2.231 \times 10^8$.

For Parameter Set 1 (using $\tilde{\pi}_2 = 0.0040$, $q_D = 0.0050$), the resulting limits were $\lim_{\alpha \to 1^+} h_2(\alpha) \approx -1.802 \times 10^8$ and $\lim_{\alpha \to \infty} h_2(\alpha) \approx -4.833 \times 10^8$.

The behaviour of $h_2(\alpha)$ under Parameter Set 3, which starts positive as $\alpha \to 1^+$ and turns negative as $\alpha \to \infty$, is the best for finding intersections with positive SR_{Observed} values when SITG is within a reasonable range. This matches its high success rate in deriving α_{Implied}. Parameter Set 2 also shows positive limits, which probably helps its moderate success. On the other side, Parameter Set 1, which shows negative limits at both ends of the alpha range for these average inputs, naturally has a lower success rate for finding solutions when SITG is positive.

6.3 Key SITG and Implied Tail Risk Insights

From the conducted implied alpha analysis in this chapter, where we worked backward, several key insights emerge, offering a broader understanding of CCP risk management. Chapter 6 introduced the α_{Implied} concept, derived by analytically inverting the SITG function ($h_2(\alpha)$) from the Cont and Ghamami (2025) model. We explored three parameter sets and identified one primary configuration, Parameter Set 3 the high-convergence set using $q = 0.01, q_D = 0.0010, \tilde{\pi}_2 = 0.0005$, that achieved a relatively high success rate in estimating α_{Implied} from observed CCP data. This implied alpha perspective provides a different viewpoint to the analysis in Chapter 5. While Chapter 5 often indicated that observed SITG levels were lower than theoretical requirements under fixed α assumptions, the current chapter shows how much extreme risk, meaning a low α_{Implied}, the model needs to assume to make many of these actual SITG amounts align with its framework.

Some of principal findings from this chapter are the parameter sensitivity. The sensitivity is key for successful estimations, the ability to find a valid α_{Implied} is strongly dependent on the choice of the probability parameters q_D and $\tilde{\pi}_2$. Parameter Set 3, with its more stringent probability values, did much better than Set 1 and Set 2 in producing reliable results (Table 6.1).

Secondly, successful model reconciliation often points to heavy tails, as the mean α_{Implied} of approximately 2.7 with Set 3 suggests that for the model to match up with observed SITG levels under these particular parameters, it often implicitly requires that the underlying loss distributions show quite heavy tails, meaning a higher chance of very large losses.

Thirdly, implied tail risk profiles show cross-CCP heterogeneity, with substantial differences in mean α_{Implied} across CCPs even under common parameters, as seen in Figure 6.2, Panel a. These values range from very low, at $\alpha_{\text{Implied}} \approx 1.3$ for ECAG implying very fat tails, to fairly high $\alpha_{\text{Implied}} \approx 7.9$ for Nasdaq Commodities, implying thinner tails.

Fourthly, a subset of observations remains inconsistent with the model, since a valid α_{Implied} solution could not be found for approximately one-quarter of observations even with the best-performing Parameter Set 3. This observation is particularly important because it might mean the theory itself or how we set its parameters has limits and cannot fully explain why all CCPs have the SITG they do in every quarter.

Finally, the derived risk metrics suggest potential understatement of true exposures, as the generally low average α_{Implied} values derived under Parameter Set 3 imply that model-consistent risk measures, such as CVaR and conditional expected loss beyond SITG, can be much larger than the actual SITG amounts, often by a lot as shown in Table 6.2. This point is even clearer when these metrics are compared against calculations using a benchmark value of $\alpha = 3.0$, in Figures 6.8 and 6.9.

Interpretation and Methodological Considerations

Our analysis in this chapter demonstrates that observed SITG levels often align with the Cont and Ghamami (2025) framework. However, to get this match, we usually need to choose parameters that help the model converge, such as Parameter Set 3. A central interpretation is that this consistency is usually achieved because the model assumes there is a lot of underlying extreme risk, which shows up as low α_{Implied} values. While this alignment might be seen as a form of model validation, it might just be a mathematical fit rather than something economically meaningful, particularly if the tail risks implied are unrealistically high for a given CCP. This process of deriving α_{Implied}, offers insights into the implicit assumptions that CCPs might be making regarding extreme tail risks, which could help adjust risk management plans or show where the model's assumptions do not match reality.

Several key interpretations and methodological considerations emerge from the analysis in this chapter. Regarding SITG adequacy and implied risk, if the Cont and Ghamami (2025) approach and the parameters from Set 3 are considered as a valid way for evaluation, the generally low α_{Implied} values offer a specific way to look at whether SITG is enough. These findings could be interpreted to suggest that current SITG levels at many CCPs may align with the model's adequacy criteria, but only if we make strong assumptions about how much extreme risk there is. An average α_{Implied} in this range points to distributions where extreme events are much more likely than would be assumed under thinner-tailed distributions. In practical terms, justifying current SITG levels often requires the model framework to assume a lot of extreme risk, based on its particular parameter settings.

The pivotal role of assumed probabilities, q_D and $\tilde{\pi}_2$, is also evident. The analysis clearly shows that both the success in estimating α_{Implied} and the size of the resultant values are highly sensitive to the exogenously chosen probability parameters, q_D and $\tilde{\pi}_2$. The achievement of an approximate 75.8% success rate using the specific values in Parameter Set 3 suggests that for the model to be matched with the actual data in most cases, it often requires assuming very conservative default fund coverage, for example $q_D = 0.0010$ implies a 99.9% confidence level for default fund adequacy, and a very tough target for how much protection SITG layers should give.

Furthermore, the model's significant sensitivity to unobservable probability parameters creates challenges for practical application and for deriving policy implications. The fact that the results are very sensitive to these unobservable probability parameters makes it hard to use this backward approach directly for regulators to set SITG rules or for CCPs to set their own SITG. Such use would demand strong justification and agreement on the specific values chosen for q, q_D, and particularly $\tilde{\pi}_2$. These choices are fundamental, as they affect what kind of risk profile we think there is and, consequently, how adequate SITG seems to be. Should regulators or CCPs consider the α_{Implied} values derived under a set like Parameter Set 3 to be too low to be realistic, it might raise questions if the actual SITG levels are adequate, or if the way the model is set up to understand them is suitable. Alternatively, it could reflect limitations of a "one-size-fits-all" approach, especially if CCP-specific risk management practices are different from the static assumptions built into the model.

Finally, the model's inherent limitations and scope must be acknowledged. The interpretations in this chapter are based on the specific theory from Cont and Ghamami (2025), which relies on assumptions such as Pareto tails and a static environment. The inability to determine an α_{Implied} for all observations, even with the best-performing parameter set, may highlight these inherent limitations of the model or point to unique characteristics of certain CCPs or specific quarters that are not completely explained by this specific model. Real-world practical constraints, such as liquidity dynamics during stress, the exact way regulations beyond SITG are applied, as discussed in Chapter 2, Section 2.2.3, or other operational frictions, are also outside the direct scope of this static, incentive-based model but could influence SITG rationale. Furthermore, the limit analysis demonstrated that for some parameter sets (like Set 1), the theoretical SITG function $h_2(\alpha)$ could give negative values for the whole α range when using average $\frac{E_1}{D_{s,2}}$ inputs, making it mathematically impossible to match actual SITG amounts, which are positive.

In summary, the implied alpha analysis conducted in this chapter offers a complex picture. While many observed SITG levels can be reconciled with the Cont and Ghamami (2025) model, this reconciliation frequently implies that current SITG levels are consistent with a lot of built-in extreme risk. This is particularly the case when specific model probability parameters are assumed to align the model's predictions with the actual data. We discuss the broader implications of these findings for SITG adequacy and systemic risk in Chapter 7.

7

Discussion and Conclusion

This chapter summarises the main findings from our investigation into central counterparty (CCP) skin in the game (SITG), based on the framework of Cont and Ghamami (2025), and connects the empirical analysis from public quantitative disclosure (PQD) data from 2019-2024 with the theoretical expectations. This summary of findings is presented in Section 7.1. Subsequently, the discussion covers what these findings mean for CCP practices, regulation, and financial stability in Section 7.2. We then address the study's limitations and suggest future research avenues in Section 7.3, before offering our concluding remarks in Section 7.4.

7.1 Summary of Findings

Examining CCP SITG through various analyses in Chapters 4, 5, and 6, reveals significant insights.

The primary finding from Chapter 5 and 6 is the significant gap observed between the reported SITG levels and those recommended by Cont and Ghamami (2025) in our sample of nine CCPs. When reasonable tail exponents, α between 2.0 and 4.0, and typical trigger probabilities are assumed, the model often suggests SITG should constitute a significant portion of the Cover 2 default fund ($D_{s,2}$). However, in practice, CCPs frequently hold less SITG, often much lower than what this theory suggests, as seen in Chapter 5, Figure 5.2. This gap was highlighted even more in Chapter 6, Section 6.2, where, even with the parameter set that worked best (Set 3), a valid implied alpha solution could not be found for approximately 24% of CCP quarterly observations, which suggests potential limits of the model's universal applicability. These findings indicate that SITG decisions are influenced by additional factors such as regulatory constraints, CCP risk tolerance, or capital budgeting frameworks.

The challenge of aligning the theoretical and the practical SITG level is directly connected to the limitations and sensitivity of the incentive model itself. The Cont and Ghamami (2025) model is highly sensitive to the chosen parameters. As shown in Chapter 5, theoretical SITG requirements change substantially with the tail exponent α and concentration ratios (c_1, r). Furthermore, Chapter 6 highlighted that changing just one probability parameter (e.g., $\tilde{\pi}_2$) could determine whether a valid implied alpha solution could be found, or greatly change its value, as seen in Table 6.1 and Figures 6.2 and 6.3. The model's reliance on a strict Pareto tail assumption for losses, shown in Chapter 3, and its static nature, which does not capture dynamic feedback or liquidity effects, discussed in Section 3.5 of the same chapter, are important limitations. While the model offers valuable theoretical insight, applying it for direct SITG calibration requires careful adjustment for each CCP's specific operational context.

Our findings on CCP capital structures, as seen in Chapter 4, support these modelling practices. The analysis shows that SITG levels remained mostly unchanged from 2019 to 2024, even as initial margin (IM) and default fund (DF) levels rose, illustrated in Figure 4.1. As a result, DF-to-SITG ratios increased, Figure 4.3, suggesting an increased reliance on member-funded resources compared to the CCP's own capital. No strong correlation was found between higher member concentration and increased SITG levels, seen in Figures 4.7, 4.8. The comparison in Chapter 5 showed that SITG suggested by theory can be much higher than what CCPs report. Moreover, the benchmark analysis in Chapter 6, assuming a moderately heavy-tailed scenario ($\alpha = 3.0$) under Parameter Set 3, indicated that potential losses (CVaR, or expected loss beyond SITG) could be much larger than the actual SITG levels (Figures 6.8 and 6.9). Taken together, these findings suggest that many CCPs may place a relatively small amount of their own capital at risk compared to the theoretical requirements for incentive alignment under extreme loss scenarios. This potentially increases reliance on member-funded resources and could reduce a CCP's motivation to invest properly in mitigating extreme tail exposures, raising questions about moral hazard if the CCP bears only a small fraction of such catastrophic losses.

7.2 Policy and Risk Management Implications

The findings of this thesis analysis have significant implications for CCP governance, regulatory frameworks, and the resilience of the financial system.

The persistent gap between theoretically derived SITG requirements and the actual SITG levels indicates that CCPs may benefit from reassessing the extent of their own capital contributions to their default waterfalls. It is the responsibility of CCPs to assess whether their current SITG contributions are adequate and address the extreme risks they deal with and offer sufficient incentives for effective risk management, especially as central clearing expands into more markets. The observation that SITG frequently remains stable during stress periods, while member contributions increase, should be looked at as an important factor when assessing risk capital internally.

From a regulatory perspective, our results emphasise the inherent challenge of developing a SITG rule based purely on risk that works universally and fits each CCP. Simpler, fixed percentage rules may prove insufficient or excessive compared to the more detailed needs suggested by models based on incentives. Given the sensitivity of the Cont and Ghamami (2025) model to parameter choices (like α and trigger probabilities), regulators might consider frameworks that involve stress-testing SITG across a range of plausible situations and tail assumptions, rather than relying on a single point estimate. The observation that even minor shifts in model parameters can greatly change the results (as seen with $\tilde{\pi}_2$ in the implied alpha analysis) suggests that any regulatory guidance based on such models should emphasise transparency and justification for the input chosen. Furthermore, as noted by Cont and Ghamami (2025), monolayer CCPs might require relatively more of their own capital to achieve the same level of incentive alignment and protection.

Inadequate SITG can pose systemic risk to financial stability if a CCP does not fully consider the potential costs of extreme events, which might lead to underinvestment in reducing risk. As central clearing continues to expand, for instance, into markets like US Treasuries, the consequences of a systemically important CCP being undercapitalised for extreme events could have severe consequences. Ensuring that CCPs have sufficient skin in the game is thus crucial for keeping confidence in the clearing system, especially during crises, and for protecting against systemic disruptions. For this reason, the calibration of SITG is a critical policy instrument in the preservation of financial stability.

7.3 Limitations and Future Research

As with any empirical study based on a particular theory, this thesis has several limitations that also suggest avenues for future research.

First, our empirical analysis primarily relies on quarterly public quantitative disclosure (PQD) data, seen in Chapter 4, which is aggregated and might not show how risks and resources change daily or even within the day. When employing the theoretical framework, estimations of specific member risk parameters were also required due to the granularity of PQD data. In addition, the manual extraction process for this PQD data was further complicated by inconsistencies in the manner in which CCPs report their information. Even if the dataset used was already standardised, the specific data fields often differed in pactice, which added a challenge when extracting the data.

Second, the Cont and Ghamami (2025) model, which forms the theoretical core of this thesis, is fundamentally static and depends on particular assumptions from extreme value theory (EVT), discussed in Chapter 3, Section 3.5. Most notably, it assumes a strict Pareto distribution for tail exposures, which, while common in financial modelling, might not fully capture the complexity of all CCP exposures. The model also does not capture dynamic feedback loops, liquidity effects during crises, or multiple periods risk management adjustments. Additionally, the theoretical framework, including the SITG formulas adapted for this thesis (Eqs. (3.8) and (3.9)), primarily models CCPs with a two layer SITG structure (S_2 and $\tilde{S}_2$). Our dataset included CCPs where a second SITG layer was not consistently reported or may not formally exist under their specific waterfall design, as only European CCPs were required under the new 2023 regulation to have the second SITG layer (European Parliament and Council 2021; European Commission 2023), which could affect the direct applicability of the full two layer model and the interpretation of α_{implied} for those entities.

Finally, the implied alpha analysis in Chapter 6 was successful in about 76% of cases with its best settings (Set 3), it still could not explain about 24% of observations and was very sensitive to parameter choices for all sets. This highlights potential areas where the model's setup or its specific settings do not quite match the real data for many CCPs in many quarters.

Acknowledging these points, future research could explore several major areas. The first area involves refining tail risk modelling. Future studies could explore the use of more flexible loss distributions, such as the generalised Pareto distribution (GPD) or hybrid approaches that combine empirical distributions with EVT for the extreme tail. If accessible, using more detailed data could greatly improve how accurately we estimate extreme risks and set up the model. Investigating alternative SITG triggers, and exploring different activation mechanisms beyond strict Pareto thresholds, could also be important, alongside modelling sequential defaults or liquidity pressures to provide further insights.

Another direction is exploring dynamics and systemic interactions. Developing multi-period models that reflect how margins, default fund contributions, and SITG can adjust across a market cycle, including potential pro-cyclical effects or liquidity pressures during default management, would be a useful next step.

Finally, research could also investigate drivers of SITG variation and model applicability. Further research could examine the reasons why different CCPs have different SITG levels, such as their governance structure, regulatory environment, operational model, or membership composition. An econometric approach might help explain why some CCPs hold more or less capital than others, following up on questions from Cont and Ghamami (2025). Examining how monolayer waterfall structures, as distinct from the multilayered ones primarily examined here, might further influence theoretical SITG requirements is another important direction, as suggested by Cont and Ghamami (2025) and noted as a potential area for future examination. This could also involve assessing the conditions under which the incentive based SITG framework of Cont and Ghamami (2025) is most applicable across different CCP designs and markets.

7.4 Conclusion

This thesis aimed to see how actual CCP SITG levels compare with the theory from Cont and Ghamami (2025). Our empirical comparison, using PQD data from 2019-2024, reveals a persistent and substantial gap: many CCPs report SITG levels much lower than what the model suggests might be needed to ensure incentives are adequate for extreme risks, even under conservative parameters.

Furthermore, the study highlights the model's significant sensitivity to the theoretical model's underlying assumptions, particularly concerning tail distributions and trigger probabilities for SITG layers. This sensitivity suggests that while the model provides a crucial benchmark and a strong basis for discussing moral hazard and incentive issues, it cannot be simply used as a fixed rule for setting SITG for all CCPs without careful consideration of each CCP's specific characteristics and justified input parameters. The challenges encountered in the implied alpha analysis also highlight how complex it is to make such theories match with observed market practices.

The findings raise important considerations for CCPs and regulators. If a CCP invests only a small amount of its own capital relative to potential extreme exposures, its incentives to manage tail risks might not be as strong. In an era where central clearing is expanding and CCPs are increasingly at the core of financial market infrastructure, ensuring these entities have sufficient and adequate 'skin in the game' that is set up correctly is vital to protect against systemic disruptions and keep market confidence. Overall, the findings show that both CCPs and regulators need to reflect on whether current SITG levels provide enough protection from extreme losses and, just as importantly, whether they provide CCPs with the right incentives to mitigate risks beforehand.

Bibliography

ABN-AMRO, Allianz, Barclays, and BlackRock (2020). *A Path Forward For CCP Resilience, Recovery, and Resolution*. Industry Report.

Avalos, F. and E. C. Mamatzakis (2018). Euro area unconventional monetary policy and bank resilience. *Monetary and Economic Department of the Bank for International Settlements*.

Biais, B., F. Heider and M. Hoerova (2016). Risk-Sharing or Risk-Taking? Counterparty Risk, Incentives, and Margins. *The Journal of Finance*. https://doi.org/10.1111/jofi.12396.

BIS (2008). *OTC derivatives market activity in the first half of 2008*.

– (2009). *Bank for International Settlements 79th Annual Report*.

– (2019). *Counterparty credit risk definitions and terminology*. CRE50.

– (2023). *Capital requirements for bank exposures to central counterparties*. CRE54.

CCP Global (2024). *PQD FAQ Guide*.

Clancy, L. (2025). *Academics call for tenfold jump in CCP capital cover*. Risk.net. URL: Risk.net.

CME Group (2025). *CPMI-IOSCO Quantitative Disclosures*. URL: https://www.cmegroup.com/solutions/clearing/cpmi-iosco-reporting.html (visited on 01/06/2025).

Commodity Futures Trading Commission (2011). *Derivatives Clearing Organizations, 17 CFR Part 39*

Cont, R. (2001). Empirical properties of asset returns: stylized facts and statistical issues.

– (2015). The end of the waterfall: default resources of central counterparties.

– (2017). Central Clearing and Risk Transformation. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.2955647.

Cont, R. and S. Ghamami (2025). Skin in the Game: Risk Analysis of Central Counterparties. *The Journal of Financial Market Infrastructures*. https://doi.org/10.21314/JFMI.2025.001.

Cont, R. and Y. H. Kan (2011). Statistical Modeling of Credit Default Swap Portfolios. *SSRN Electronic Journal*.

Cont, R. and T. Kokholm (2014). Central clearing of OTC derivatives: Bilateral vs multilateral netting. *Statistics & Risk Modeling* 31.1, 3–22. https://doi.org/10.1515/strm-2013-1161.

CPMI-IOSCO (2012). *Principles for Financial Market Infrastructures*. Basel: Bank for International Settlements.

CPMI-IOSCO (2015). *Public quantitative disclosure standards for central counterparties*. Basel: Bank for International Settlements.
– (2017). *Recovery of financial market infrastructures*. Basel: Bank for International Settlements.
– (2020). *Central counterparty default management auctions: issues for consideration*. Basel: Bank for International Settlements.
Duffie, D., M. Scheicher and G. Vuillemey (2015). Central clearing and collateral demand. *Journal of Financial Economics* 116.2, 237–256. https://doi.org/10.1016/j.jfineco.2014.12.006.
Duffie, D. and K. J. Singleton (2003). *Credit risk: pricing, measurement, and management*.
Duffie, D. and H. Zhu (2011). Does a Central Clearing Counterparty Reduce Counterparty Risk? *Review of Asset Pricing Studies* 1.1, 74–95. https://doi.org/10.1093/rapstu/rar001.
EACH (2021). *Carrots and sticks: How-the skin in the game incentivises CCPs to perform robust risk management*. EACH Paper.
Eurex Clearing AG (2025). *Regulatory standards and Disclosures*. Eurex. URL: https://www.eurex.com/ec-en/find/about-us/regulatory-standards (visited on 01/06/2025).
European Commission (2013). *Commission Delegated Regulation (EU) No 153/2013 of 19 December 2012 supplementing Regulation (EU) No 648/2012, Official Journal of the European Union, L 52, pp. 41-74*
European Commission (2023). *Commission Delegated Regulation (EU) 2023/840 specifying the methodology for calculating additional pre-funded dedicated own resources of central counterparties. Official Journal of the European Union, L 107, pp.29-38*
European Parliament and Council (2019). *Regulation (EU) 2019/2099 of the European Parliament and of the Council of 23 October 2019 amending Regulation (EU) No 648/2012, Official Journal of the European Union, L 322, pp.1-44*
European Parliament and Council (2021). *Regulation (EU) 2021/23 of the European Parliament and of the Council of 16 December 2020 on a framework for the recovery and resolution of central counterparties, Official Journal of the European Union, L 22, pp.1–102*
FSB (2022). *OTC Derivatives Market Reforms Implementation progress in 2022*.
– (2023). *Financial Resources and Tools for Central Counterparty Resolution Consultation report*. Consultation report.
Glasserman, P. and H. P. Young (2016). Contagion in Financial Networks. *Journal of Economic Literature* 54.3, 779–831. https://doi.org/10.1257/jel.20151228.
Gregory, J. (2014). *Central Counterparties*. Wiley.
Group of Twenty (G20) (2009). *Leaders Statement: The Pittsburgh Summit*. Communiqué. Group of Twenty (G20).
Haene, P. and A. Sturm (2009). *Optimal Central Counterparty Risk Management*.
Hansmann, H. (2014). All firms are cooperatives and so are governments. *Journal of Entrepreneurial and Organizational Diversity* 2.2, 1–10. https://doi.org/10.5947/jeod.2013.007.

Hart, O. (1996). The governance of exchanges: members' cooperatives versus outside ownership. *Oxford Review of Economic Policy* 12.4, 53–69. https://doi.org/10.1093/oxrep/12.4.53.

Heath, A. et al. (2016). CCPs and network stability in OTC derivatives markets. *Journal of Financial Stability* 27, 217–233. https://doi.org/10.1016/j.jfs.2015.12.004.

Heller, D. and N. Vause (2012). *Collateral requirements for mandatory central clearing of over-the-counter derivatives.*

HM Treasury (2021). *Expanded Resolution Regime: Central Counterparties.* United Kingdom: HM Treasury.

Hull, J. C. (2023). *Risk Management and Financial Institutions.* Wiley.

Intercontinental Exchange (2025). *Quarterly Clearing Disclosures.* URL: https://www.ice.com/clearing/quarterly-clearing-disclosures (visited on 01/06/2025).

Kuong, J. C.-F. and V. Maurin (2024). The Design of a Central Counterparty. *Journal of Financial and Quantitative Analysis* 59.3, 1257–1299. https://doi.org/10.1017/S0022109023000121.

London Stock Exchange Group (2025). *CCP Disclosures.* URL: https://www.lseg.com/en/post-trade/clearing/clearing-resources/ccp-disclosures (visited on 01/06/2025).

McNeil, A. J., R. Frey and P. Embrechts (2015). *Quantitative risk management: concepts, techniques and tools.* Revised edition. Princeton series in finance. Princeton Oxford: Princeton University Press.

Menkveld, A. J. (2017). Crowded Positions: An Overlooked Systemic Risk for Central Clearing Parties. *The Review of Asset Pricing Studies* 7.2, 209–242. https://doi.org/10.1093/rapstu/rax016.

Murphy, D., M. Vasios and N. Vause (2014). *An investigation into the procycliacality of risk-based initial margin models.* Financial Stability Paper No. 29. Bank of England.

Nasdaq Clearing (2025). *CPMI-IOSCO Disclosures.* URL: https://www.nasdaq.com/solutions/about-nasdaq-clearing (visited on 01/06/2025).

Pirrong, C. (2011). *The Economics of Central Clearing: Theory and Practice.* ISDA Discussion Paper Series. ISDA.

Rahman, A. (2015). *Over-the-counter (OTC) derivatives, central clearing and financial stability.*

SEC (2016). *Standards for Clearing Agencies, 17 CFR \S~40.17Ad-22*

Thiruchelvam, S. (2022). *BoE official signals tough stance on CCP skin in the game.* Risk.net.

Tsay, R. S. (2010). *Analysis of financial time series.* Third edition. Wiley series in probability and statistics. Hoboken, N.J: Wiley. https://doi.org/10.1002/9780470644560.

Tucker, P. (2014). Are Clearinghouses the New Central Banks? *Centre for International Governance Innovation.*

United Kingdom (2018). *European Union (Withdrawal) Act 2018, c. 16*

United States (2010). *Dodd-Frank Wall Street Reform and Consumer Protection Act*, Pub.L. 111203

Walker, J. (2023). *CCP skin in the game still dwarfed by member contributions*. Risk.net.

Wilkes, S. (2020). *EU hands CCP members a narrow win on skin in the game*. Risk.net.

A

Derivations of SITG Formulas for the Cover 2 Scenario

This appendix walks through the step-by-step derivations that adapt the general SITG (skin in the game) framework to the Cover 2 setup, the central case examined in this thesis. We use the notation and assumptions introduced in Chapter 3.

A.1 Derivation of First-Layer SITG (S_2)

In this appendix we detail how the first-layer SITG, S_2, is derived. The process starts with the Incentive Compatibility Constraint (ICC).

The first-layer SITG, S_2, is set using the ICC, requiring that the probability of excess losses from the largest defaulter (member (1)), which are $L_{(1)} = (U_{(1)} - M_{(1)})^+$, going above its own resources ($D_{(1)}$) plus S_2, must equal a target probability π_2. As noted in Chapter 3, π_2 must be less than or equal to q_D. The constraint is written as:

$$P(L_{(1)} > D_{(1)} + S_2) = \pi_2 \, .$$

To solve this, we use the unconditional Pareto probability for $L_{(1)}$ (from Eq. (3.4) in Chapter 3). We use its scale parameter $\kappa_{(1)}$ and set the threshold $x = D_{(1)} + S_2$. This gives us:

$$q \left(1 + \frac{D_{(1)} + S_2}{\kappa_{(1)}}\right)^{-\alpha} = \pi_2 \, .$$

To isolate the term $D_{(1)} + S_2$:

$$\begin{aligned}
\left(1 + \frac{D_{(1)} + S_2}{\kappa_{(1)}}\right)^{-\alpha} &= \frac{\pi_2}{q} \\
1 + \frac{D_{(1)} + S_2}{\kappa_{(1)}} &= \left(\frac{\pi_2}{q}\right)^{-1/\alpha} = \left(\frac{q}{\pi_2}\right)^{1/\alpha} \\
D_{(1)} + S_2 &= \kappa_{(1)} \left[\left(\frac{q}{\pi_2}\right)^{1/\alpha} - 1\right] . \qquad \text{(A.1)}
\end{aligned}$$

Next, plugging in the expression for $\kappa_{(1)}$ (using Eq. (3.5), which includes $E_{(1)}$, the tail exposure for member (1)):

$$D_{(1)} + S_2 = \frac{E_{(1)}}{\left(\frac{q}{q_D}\right)^{1/\alpha} - 1} \left[\left(\frac{q}{\pi_2}\right)^{1/\alpha} - 1\right] . \qquad \text{(A.2)}$$

To isolate S_2 itself, we use the relationships $D_{(1)} = c_1 D_{s,2}$ (from Eq. (3.7)) and $E_{(1)} = \frac{c_1}{c_1+c_2} D_{s,2}$ (from Eq. (3.6)):

$$S_2 = \frac{E_{(1)}}{\left(\frac{q}{q_D}\right)^{1/\alpha} - 1} \left[\left(\frac{q}{\pi_2}\right)^{1/\alpha} - 1\right] - D_{(1)}$$

$$S_2 = \frac{\frac{c_1}{c_1+c_2} D_{s,2}}{\left(\frac{q}{q_D}\right)^{1/\alpha} - 1} \left[\left(\frac{q}{\pi_2}\right)^{1/\alpha} - 1\right] - c_1 D_{s,2} \,. \tag{A.3}$$

Finally, factoring out $D_{s,2}$ gives us the formula for S_2 (as in Eq. (3.8)). This is the Cover 2 ($n = 2$) version of the general S_n formula:

$$S_2 = \left[\left(\frac{(q/\pi_2)^{1/\alpha} - 1}{(q/q_D)^{1/\alpha} - 1}\right)\left(\frac{c_1}{c_1 + c_2}\right) - c_1\right] D_{s,2} \,.$$

This formula for S_2 shows how the first-layer size connects to the largest member's characteristics (through concentration ratios c_1, c_2), the target default probabilities (π_2, q_D), and the tail risk parameter α.

A.2 Derivation of Second-Layer SITG ($\tilde{S}_2$)

To derive the second-layer SITG, $\tilde{S}_2$, we build on the first layer by introducing a second, more extreme loss threshold. This helps determine the additional resources needed to cover rarer, more severe losses. The loss thresholds are:

(a) $X_1 = D_{(1)} + S_2$, for which $P(L_{(1)} > X_1) = \pi_2$, so $X_1 = \kappa_{(1)} \left[\left(\frac{q}{\pi_2}\right)^{1/\alpha} - 1\right]$.

(b) $X_2 = D_{s,2} + S_2 + \tilde{S}_2$, for which $P(L_{(1)} > X_2) = \tilde{\pi}_2$, so $X_2 = \kappa_{(1)} \left[\left(\frac{q}{\tilde{\pi}_2}\right)^{1/\alpha} - 1\right]$.

To find $\tilde{S}_2$, we first write it as the difference between the total resources for the second threshold and other known resources:

$$\tilde{S}_2 = X_2 - (D_{s,2} + S_2) \,. \tag{A.4}$$

Then, substituting $S_2 = X_1 - D_{(1)}$ (from condition (a)) into this, we get:

$$\tilde{S}_2 = X_2 - D_{s,2} - (X_1 - D_{(1)}) = X_2 - X_1 - D_{s,2} + D_{(1)} \,. \tag{A.5}$$

Plugging in the formulas for X_1 and X_2 (which use $\kappa_{(1)}$) gives us:

$$\tilde{S}_2 = \kappa_{(1)} \left[\left(\frac{q}{\tilde{\pi}_2}\right)^{1/\alpha} - 1\right] - \kappa_{(1)} \left[\left(\frac{q}{\pi_2}\right)^{1/\alpha} - 1\right] - D_{s,2} + D_{(1)}$$

$$= \kappa_{(1)} \left[\left(\frac{q}{\tilde{\pi}_2}\right)^{1/\alpha} - \left(\frac{q}{\pi_2}\right)^{1/\alpha}\right] - D_{s,2} + D_{(1)} \,. \tag{A.6}$$

Finally, we substitute for $\kappa_{(1)}$ (using Eq. (3.5)), and the relationships for $E_{(1)}$ (Eq. (3.6)) and $D_{(1)}$ (Eq. (3.7)). Then, factoring out $D_{s,2}$ yields:

$$\tilde{S}_2 = \frac{E_{(1)}}{(q/q_D)^{1/\alpha} - 1} \left[\left(\frac{q}{\tilde{\pi}_2}\right)^{1/\alpha} - \left(\frac{q}{\pi_2}\right)^{1/\alpha}\right] - D_{s,2} + c_1 D_{s,2}$$

$$= \left[\left(\frac{(q/\tilde{\pi}_2)^{1/\alpha} - (q/\pi_2)^{1/\alpha}}{(q/q_D)^{1/\alpha} - 1}\right) \frac{c_1}{c_1 + c_2} + c_1 - 1\right] D_{s,2} \,. \tag{A.7}$$

This gives the formula for $\tilde{S}_2$ (as in Eq. (3.9)), the Cover 2 ($n = 2$) version of the general $\tilde{S}_n$ formula. This formula for $\tilde{S}_2$ shows how it depends on the two ICC thresholds $(\pi_2, \tilde{\pi}_2)$, the Pareto tail parameter α, and member concentration ratios. The use of two different thresholds, π_2 and $\tilde{\pi}_2$, is key here, as it allows us to size an additional layer for those rarer, more extreme losses.

A.3 Derivation of Boundary Probability ($\tilde{\pi}_{0,2}$)

The boundary probability ($\tilde{\pi}_{0,2}$) is useful for understanding the system's risk if the second layer of SITG ($\tilde{S}_2$), is not funded, especially when the first layer (S_2), is sized using its least conservative target probability. Specifically, $\tilde{\pi}_{0,2}$ is defined as the probability $P(\tilde{L}^{(1)}_{i,2} > D_{i,2})$. This represents the event that losses from the largest defaulter (member (1)) exhaust surviving member i's prefunded default fund contribution ($D_{i,2}$) in the Cover 2 scenario, thereby requiring unfunded resources. This probability is evaluated under the specific conditions where $\tilde{S}_2 = 0$ and the target probability for the first layer is $\pi_2 = q_D$.

The event $\tilde{L}^{(1)}_{i,2} > D_{i,2}$ (i.e., member i's prefunded DF contribution being exhausted by losses from member (1)) is equivalent to the total loss from member (1), denoted $L_{(1)} = (U_{(1)} - M_{(1)})^+$, exceeding the sum of the mutualised default fund ($D_{s,2}$), the first-layer SITG (S_2), and any second-layer SITG ($\tilde{S}_2$). The general formula for this probability, $P(L_{(1)} > D_{s,2} + S_2 + \tilde{S}_2)$, can be derived from the unconditional Pareto tail probability (Eq. (3.4) from Chapter 3). By substituting $\kappa_{(1)}$ using Eq. (3.5) and noting that for Cover 2, $D_{s,2}/E_{(1)} = (c_1 + c_2)/c_1$ (derived from Eq. (3.6)), this probability is given by

$$P(L_{(1)} > D_{s,2} + S_2 + \tilde{S}_2) = q\left[1 + \left(\left(\frac{q}{q_D}\right)^{1/\alpha} - 1\right)\left(\frac{c_1 + c_2}{c_1} + \frac{S_2}{E_{(1)}} + \frac{\tilde{S}_2}{E_{(1)}}\right)\right]^{-\alpha}. \tag{A.8}$$

To derive $\tilde{\pi}_{0,2}$, we start with Eq. (A.8) and apply the conditions $\tilde{S}_2 = 0$ and $\pi_2 = q_D$. Under the condition $\pi_2 = q_D$, S_2 simplifies to $S_2(\pi_2 = q_D) = \left(\frac{c_1}{c_1+c_2} - c_1\right) D_{s,2}$. Consequently, we have:

$$\frac{S_2(\pi_2 = q_D)}{E_{(1)}} = \frac{\left[\frac{c_1}{c_1+c_2} - c_1\right] D_{s,2}}{\frac{c_1}{c_1+c_2} D_{s,2}} = 1 - (c_1 + c_2)\,. \tag{A.9}$$

Plugging these conditions ($\tilde{S}_2 = 0$ and the simplified ratio from Eq. (A.9)) back into the general probability formula (Eq. (A.8)), we get:

$$\tilde{\pi}_{0,2} = q\left[1 + \left(\left(\frac{q}{q_D}\right)^{1/\alpha} - 1\right)\left(\frac{c_1 + c_2}{c_1} + (1 - (c_1 + c_2))\right)\right]^{-\alpha}. \tag{A.10}$$

Finally, simplifying the term $\left(\frac{c_1+c_2}{c_1} + 1 - (c_1 + c_2)\right)$ to $1 + (c_1 + c_2)\left(\frac{1-c_1}{c_1}\right)$ gives us the formula for $\tilde{\pi}_{0,2}$ (as in Eq. (3.10)):

$$\tilde{\pi}_{0,2} = q\Big[1 + \Big(\Big(\tfrac{q}{q_D}\Big)^{1/\alpha} - 1\Big)\Big(1 + (c_1 + c_2)\Big(\tfrac{1-c_1}{c_1}\Big)\Big)\Big]^{-\alpha}.$$

$\tilde{\pi}_{0,2}$ therefore shows the risk level if only the first layer exists (sized at its least conservative ICC, $\pi_2 = q_D$) and no second layer is funded. It is a key baseline for seeing if $\tilde{S}_2$ is really needed.

A.4 Derivation of Total SITG for Cover 2 ($S_2+\tilde{S}_2$)

The derivation of total SITG for the Cover 2 scenario simply involves summing the two previously derived layers, S_2 and $\tilde{S}_2$.

The total SITG is found by summing S_2 (as in Eq. (3.8)) and $\tilde{S}_2$ (as in Eq. (3.9)):

$$\begin{aligned} S_2 + \tilde{S}_2 &= D_{s,2}\left[\left(\frac{(q/\pi_2)^{1/\alpha} - 1}{(q/q_D)^{1/\alpha} - 1}\right)\frac{c_1}{c_1 + c_2} - c_1 \right. \\ &\quad \left. + \left(\frac{(q/\tilde{\pi}_2)^{1/\alpha} - (q/\pi_2)^{1/\alpha}}{(q/q_D)^{1/\alpha} - 1}\right)\frac{c_1}{c_1 + c_2} + c_1 - 1\right] \\ &= D_{s,2}\left[\frac{c_1}{(c_1 + c_2)((q/q_D)^{1/\alpha} - 1)}\left((q/\pi_2)^{1/\alpha} - 1 + (q/\tilde{\pi}_2)^{1/\alpha} - (q/\pi_2)^{1/\alpha}\right) - 1\right] \\ &= D_{s,2}\left[\frac{c_1}{c_1 + c_2}\frac{(q/\tilde{\pi}_2)^{1/\alpha} - 1}{(q/q_D)^{1/\alpha} - 1} - 1\right]. \end{aligned} \tag{A.11}$$

This gives us the formula for the total SITG (as in Eq. (3.11)). This is the Cover 2 ($n = 2$) version of the general $S_n + \tilde{S}_n$ formula. The structure of this total SITG formula shows that the overall requirement is driven by the more stringent target probability $\tilde{\pi}_2$, along with member concentration and tail risk. This reflects the idea that the total SITG should be strong enough to protect against even the most extreme losses.

B

Detailed Regression Results from Trends in CCP Financial Resources

This appendix summarizes the detailed regression outputs that support the exploratory analysis in Chapter 4, specifically Section 4.3. Each model regresses SITG on a single explanatory variable, with the coefficient tables using the generic label 'x1' for the respective predictor.

B.1 SITG vs. Required Initial Margin (IM)

Model: $\log(\text{SITG}) \sim \beta_0 + \beta_1 \log(\text{IM})$ (see Figure 4.5)

Estimated Coefficients

Predictor	Estimate	SE	tStat	pValue
(Intercept)	2.6763	0.48238	5.5481	8.4807e-08
x1	0.48931	0.046596	10.501	4.4988e-21

Model Statistics

Number of observations: 216, Error degrees of freedom: 214
Root Mean Squared Error: 0.481
$R^2 = 0.34$, Adjusted $R^2 = 0.337$
F-statistic vs. constant: 110 (p = 4.5e-21)

Summary: SITG $\propto$ IM$^{0.49}$

B.2 SITG vs. Required Default Fund (DF)

Model: $\log(\text{SITG}) \sim \beta_0 + \beta_1 \log(\text{DF})$ (see Figure 4.6)

Estimated Coefficients

Predictor	Estimate	SE	tStat	pValue
(Intercept)	-0.54605	0.44152	-1.2368	0.21754
x1	0.88596	0.04719	18.774	4.0023e-47

Model Statistics

Number of observations: 216, Error degrees of freedom: 214
Root Mean Squared Error: 0.364
$R^2 = 0.622$, Adjusted $R^2 = 0.620$
F-statistic vs. constant: 352 (p = 4e-47)

Summary: $\text{SITG} \propto \text{DF}^{0.89}$

B.3 SITG vs. Top 5 IM Concentration (%)

Model: $\text{SITG} \sim \beta_0 + \beta_1 \text{Top5IM\%}$ (see Figure 4.7)

Estimated Coefficients

Predictor	Estimate	SE	tStat	pValue
(Intercept)	1.9287e+08	1.637e+07	11.783	5.0195e-25
x1	-1.8675e+08	2.9575e+07	-6.3146	1.5418e-09

Model Statistics

Number of observations: 216, Error degrees of freedom: 214
Root Mean Squared Error: 6.81e+07
$R^2 = 0.157$, Adjusted $R^2 = 0.153$
F-statistic vs. constant: 39.9 (p = 1.54e-09)

Summary: $\text{SITG} = -186.75 \cdot \text{Top5IM\%} + 192.87$ (in millions)

B.4 SITG vs. Top 5 DF Concentration (%)

Model: $\text{SITG} \sim \beta_0 + \beta_1 \text{Top5DF\%}$ (see Figure 4.8)

Estimated Coefficients

Predictor	Estimate	SE	tStat	pValue
(Intercept)	1.7514e+08	1.3796e+07	12.695	6.7451e-28
x1	-1.7771e+08	2.8362e+07	-6.2658	2.0101e-09

Model Statistics

Number of observations: 216, Error degrees of freedom: 214
Root Mean Squared Error: 6.82e+07
$R^2 = 0.155$, Adjusted $R^2 = 0.151$
F-statistic vs. constant: 39.3 (p = 2.01e-09)

Summary: $\text{SITG} = -177.71 \cdot \text{Top5DF\%} + 175.14$ (in millions)