

Efficient Monte Carlo Valuation of Corporate Bonds in Financial Networks

Dohyun Ahn*

The Chinese University of Hong Kong

Agostino Capponi†

Columbia University

February 13, 2026

Abstract

Valuing corporate bonds in systemic economies is challenging due to intricate webs of inter-institutional exposures. When a bank defaults, cascading losses propagate through the network, with payments determined by a system of fixed-point equations lacking closed-form solutions. Standard Monte Carlo methods cannot capture rare yet critical default events, while existing rare-event simulation techniques fail to account for higher-order network effects and scale poorly with network size. To overcome these challenges, we propose a novel approach—*Bi-Level Importance Sampling with Splitting*—and characterize individual bank defaults by decoupling them from the network’s complex fixed-point dynamics. This separation enables a two-stage estimation process that directly generates samples from the banks’ default events. We demonstrate theoretically that the method is both scalable and asymptotically optimal, and validate its effectiveness through numerical studies on empirically observed networks.

KEYWORDS: FINANCIAL NETWORKS, IMPORTANCE SAMPLING, MONTE CARLO METHOD, EISENBERG-NOE FRAMEWORK, BOND VALUATION

1 Introduction

In modern financial systems, institutions are intricately linked through common balance sheet holdings (Capponi and Larsson, 2015; Greenwood et al., 2015; Duarte and Eisenbach, 2021) and interbank liabilities. This interconnectedness can be destabilizing and lead to the propagation of

*Department of Systems Engineering and Engineering Management, E-mail: dohyun.ahn@cuhk.edu.hk

†Department of Industrial Engineering and Operations Research and Columbia Business School, E-mail: ac3827@columbia.edu

contagion and defaults during distressed periods. As demonstrated during financial crises, such as the 2008 global financial crisis, a financial institution can withstand a moderate shock, yet be forced into default by contagion due to the propagation of losses originating from other distressed firms within the network (Gai et al., 2011).

The risk management literature has extensively analyzed financial contagion mechanisms and their impact on institutional and systemic stability, accounting for network exposures and default-triggered frictions. However, a critical gap remains: the pricing of corporate bonds subject to network-induced credit risk has been largely overlooked, despite its practical importance. This stands in contrast to the vast literature on corporate bond pricing, a central topic in financial engineering for decades. Our paper addresses this gap by proposing an efficient valuation framework for bonds of firms exposed to systemic credit risk.

The fundamental challenge in this valuation stems from the network effects themselves. Their complexity precludes any closed-form expression for the bond price, which is a key distinction from classical non-networked models, where analytical characterizations are typical. Arguably, the most tractable computational remedy would be the use of standard Monte Carlo simulation, where random shocks to each firm’s assets are generated from the underlying distributions to simulate the resulting contagion, and the bond price is computed at an equilibrium where the cascade of losses and defaults stops. However, this method is notoriously inefficient in capturing rare but critical default events that are the key drivers of credit risk.

A judiciously designed Monte Carlo method can dramatically improve efficiency through variance reduction. The theoretically optimal (zero-variance) estimator, however, requires prior knowledge of the very quantity being estimated and is thus impractical to implement. Consequently, this paper focuses on designing a tractable Monte Carlo method that achieves near-optimal variance reduction by efficiently sampling extreme events that cause defaults. While Monte Carlo methods for rare-event simulation have been studied extensively, the literature on their application to financial networks remains sparse, and the few existing methods face three significant challenges:

- (i) *Higher-order network effects.* Classical rare-event simulation techniques often rely on an explicit characterization of the event of interest. As previously noted, this is not possible in financial networks where contagion extends beyond direct, first-order exposures. The failure of one institution can trigger losses at its counterparties, which in turn propagate distress to their own counterparties, initiating indirect, higher-order cascades. These dynamics can cause shocks to circulate and amplify until the system settles into a new equilibrium, often described by a fixed-point equation. Consequently, the critical events we focus on are defined only implicitly, making traditional rare-event simulation methods inapplicable.
- (ii) *The curse of dimensionality.* Despite the implicit nature of these events, a few Monte Carlo methods have been proposed for estimating default probabilities, particularly in the classical model of Eisenberg and Noe (2001). While these methods lead to accurate estimates

for small networks (e.g., ten or fewer institutions), their computational time increases exponentially with network size. Given that real-world financial systems involve hundreds or thousands of interconnected institutions, a scalable method whose performance is robust to high dimensionality is imperative.

- (iii) *Network model sensitivity.* The existing methods mentioned above suffer from a further limitation: they are tailored to a specific model. They are built upon a linear programming reformulation of the equilibrium in the model of Eisenberg and Noe (2001), which only accounts for contagion through interbank liabilities. This rigid structure prevents the incorporation of other critical, non-linear contagion channels, such as fire-sale externalities and bankruptcy costs. As these non-linear effects have received significant recent attention for their realism, this motivates the design of a flexible Monte Carlo method applicable to a wide range of modern financial network models.

To the best of our knowledge, our paper is the first to introduce an efficient Monte Carlo method for bond valuation that simultaneously addresses all three of these challenges. Our model incorporates common asset holdings and interbank liabilities, which are widely recognized as the two most prominent channels of financial contagion (Clerc et al., 2016; Glasserman and Young, 2016). Specifically, distressed institutions may be forced to sell off some of their illiquid asset holdings to meet their debt obligations. This fire sale depresses asset prices and subsequently creates a negative externality, reducing the value of the same assets held on the balance sheets of all other banks. Furthermore, if a distressed bank defaults on its obligations even after the asset liquidation, it distributes all its residual assets to pay its creditors. Due to the shortfall in payments, the creditors recover only a fraction of their claims, and these credit losses can then trigger further defaults, creating loss amplification among insolvent institutions.¹

In this model, we first explicitly characterize an institution’s default condition conditional on the shock realizations of all other institutions. As discussed in (i), default events cannot be defined explicitly in the full-dimensional space of external shocks due to loss amplification and default cascade dynamics. This complexity makes it difficult to develop efficient simulation methods that sample directly from the default region. We overcome this by identifying a critical threshold for a target institution’s external assets, below which it defaults given the shocks to the rest of the system. This is achieved by strategically decoupling the target institution from the network and analyzing the impact of amplified losses from the separated system back onto the target. This impact is what defines the critical threshold. Crucially, the calculation of this threshold is computationally efficient, with a complexity that is linear in the network size. Furthermore, this decoupling strategy is not limited to our specific financial model but is broadly applicable. Accordingly, our characterization provides a unified solution to the challenges outlined in (i) to (iii).

¹Although we exclude bankruptcy costs and other inefficiencies that might arise at default from our model, our methodology is robust enough to accommodate more general frameworks that incorporate them alongside the two contagion channels.

We then leverage this default set characterization to design a novel and efficient Monte Carlo method for pricing institutional bonds: Bi-Level Importance Sampling with Splitting (BLISS). The proposed method employs a splitting technique that decomposes the default event into a sequence of two nested, more tractable sub-events. This bi-level framework aligns perfectly with the decoupling approach mentioned above: the outer-level event corresponds to shock realizations in non-target institutions, while the inner-level event represents the conditional default of the target institution itself. To achieve significant variance reduction, BLISS integrates this splitting framework with two distinct importance sampling schemes, one for each level. Specifically, we apply exponential tilting to the outer-level sampling and optimal importance sampling—a provably ideal technique for this purpose (Asmussen and Glynn, 2007)—to the inner-level sampling.

Our key technical contribution lies in establishing theoretical performance guarantees for the BLISS estimator by proving its asymptotic optimality across two distinct asymptotic regimes that capture rare and non-rare default events: large-asset and low-volatility. The first regime is characterized by increasing external assets, while the second regime focuses on situations where individual asset volatility decreases. In these two regimes, default becomes increasingly rare. For each regime, we propose an effective choice of the exponential tilting parameter for the outer-level importance sampling in the BLISS method, leading to substantial variance reduction. Although the optimal parameter—defined as the minimizer of the estimator’s second moment—would be ideal, this stochastic optimization is analytically intractable, a well-established challenge in the field (Glasserman, 2003). To circumvent this issue, we construct a deterministic surrogate for the second moment in each regime and identify its minimizer instead, which serves as a proxy for the optimal parameter.

The remainder of the paper is organized as follows. Section 2 provides a review of the related literature. Section 3 introduces the modeling framework and formulates the bond pricing problem. In Section 4, we present the decoupling approach and the BLISS methodology noted above. Section 5 analyzes theoretical performance guarantees for the BLISS algorithm, demonstrating in particular its asymptotic optimality. Section 6 validates our theoretical results through numerical experiments based on a model calibrated to data from the European Banking Authority (EBA). Section 7 concludes. Proofs of all theoretical results are deferred to the appendix.

2 Literature Review

The valuation of interbank claims in the presence of network contagion poses significant computational challenges. While theoretical frameworks for clearing payments and default cascades are well-developed, translating these models into tractable valuation tools requires addressing the curse of dimensionality that arises from network interactions and non-linear contagion channels. This section reviews three related strands of literature: network-based financial contagion models, existing approaches to debt valuation in networks that highlight computational limitations, and simulation methodologies for related problems that inform our proposed approach.

Network-based financial contagion has been extensively studied, beginning with the seminal framework of Eisenberg and Noe (2001), which characterizes equilibrium clearing payments through a fixed-point formulation that accounts for cascading defaults arising from failures in fulfilling interbank liabilities. Subsequent research has extended this framework along several dimensions. Contagion channels have been enriched to include asset fire sales (Cifuentes et al., 2005), bankruptcy costs (Rogers and Veraart, 2013; Glasserman and Young, 2015), and cross-holdings (Weber and Weske, 2017). Other studies have incorporated institutional features such as central clearing counterparties (Amini et al., 2016a; Ahn, 2020), multiple maturities (Kusnetsov and Veraart, 2019), collateral requirements with contract termination (Ghamami et al., 2022), and contagious bank runs (Veraart, 2025). The model we employ builds upon the framework of Cifuentes et al. (2005) to capture both liability and fire-sale contagion channels, which we extend to include external liabilities and stochastic shocks to external assets.

Building upon these models, other studies have investigated computational and policy questions that motivate our methodological contribution. Elsinger et al. (2006) develop an empirical approach to assess systemic risk, Capponi et al. (2016) study the impact of the concentration of interbank liabilities on systemic losses, and Chen et al. (2016) analyze the interplay between the two contagion channels of asset liquidation and interbank liabilities. Other studies have characterized the sensitivity of clearing payments (Liu and Staum, 2010; Feinstein et al., 2018) to non-interbank assets and network structure, intervention policies (Capponi and Chen, 2015; Ahn and Kim, 2019; Bernard et al., 2022), systemic stress scenario selection (Ahn et al., 2023), computation of systemic risk measures (Ararat and Meimanjan, 2023), and worst-case network effects under partial information (Ahn et al., 2024; Hu et al., 2024). We refer to Capponi (2016) and Glasserman and Young (2016) for comprehensive surveys.

While the literature on clearing payments focuses on the ex-post consequences of shocks, computing the ex-ante value of interbank claims requires integrating over all possible realizations of external shocks. This problem is fundamentally constrained by the curse of dimensionality. The methodology of Gouriéroux et al. (2012) decomposes the space of external shocks into 2^n default scenarios for a network of n institutions and then evaluates expectations conditional on each scenario. However, this process is computationally impractical for large networks, meaning explicit analytical solutions are often restricted to highly simplified cases, such as those with few firms or no interbank liabilities. Barucca et al. (2020) develop a general network valuation framework that synthesizes both clearing and ex-ante valuation models, while also accounting for uncertainty in external assets. However, analytical solutions remain elusive for complex, realistic systems, specifically when we take into account non-linear contagion channels such as asset liquidation and bankruptcy costs.

The study of Cossin and Schellhorn (2007) provides an analytical valuation framework for corporate debt in a network economy. There are significant differences between their study and ours. First, their queueing-network-based model features strategic default, and banks default when

equity value turns negative and shareholders no longer find it profitable to operate the firm. By contrast, we model liquidity-induced defaults that occur when a bank’s asset value falls below its liabilities. Second, Cossin and Schellhorn (2007) derive asymptotic formulas for equity and debt prices in the steady state under the assumption that firms hold sufficiently large cash reserves. We provide an exact simulation framework that does not rely on steady-state or large-cash assumptions, enabling precise valuation across a broader range of financial conditions.

Despite its practical significance, the development of efficient simulation methods for financial networks remains an underexplored area. The few existing studies have focused on estimating firm-specific default probabilities via conditional Monte Carlo (Ahn and Kim, 2018), mixed importance sampling (Ahn and Zheng, 2021), and conditional importance sampling (Ahn and Zheng, 2025). A similar problem has been explored in the context of distribution networks, where Blanchet et al. (2019) introduce importance sampling and conditional Monte Carlo estimators to compute systemic failure probabilities.² However, as mentioned earlier, these approaches are not well-suited for large-scale systems. The first three methods rely on the exponential partitioning of the external asset space, in line with the approach of Gouriéroux et al. (2012), whereas the latter requires solving a linear program in each iteration. Hence, these methods suffer from a significant computational burden that scales poorly with network size. Furthermore, these algorithms are designed for estimating expectations or probabilities defined over a set characterized by linear constraints, which restricts their use to models that exclude the non-linear effects of fire sales (e.g., the original Eisenberg-Noe framework). Lastly, the above schemes require that random shocks to external assets be elliptically distributed. This condition is not met under the common modeling practice of using normally distributed log-returns, and thus, these existing techniques are inapplicable in this standard scenario. To address these complexities, we introduce a novel approach leveraging bi-level sampling coupled with conditioning on the external assets of non-target institutions.

Our bi-level sampling approach is conceptually analogous to rare-event simulation techniques developed in the portfolio credit risk literature. To be more specific, Glasserman and Li (2005), Glasserman et al. (2008), Kang and Shahabuddin (2005), and Bassamboo et al. (2008) propose two-stage importance sampling methods that separate the sampling of common risk factors in a credit portfolio from that of idiosyncratic risk factors. This mirrors our approach in that we also split random variables into two groups and apply different importance sampling schemes to each. The key distinction, however, lies in how the target rare event and the corresponding sampling distributions are characterized. In our setting, these characterizations are driven by network contagion effects—a factor absent from the credit portfolio models discussed in the aforementioned papers. In those models, an obligor’s default does not directly affect another obligor’s financial status; instead, their dependence is captured through elliptical copulas. Hence, while the underlying two-stage concept is similar, the methods from the portfolio credit risk literature are not applicable to the problem of bond valuation in financial networks.

²We refer the reader to Gandy and Veraart (2017) and Glasserman and de Larrea (2023) for different approaches to sampling constrained networks and their application in financial networks.

Assets	Liabilities and Net Worth
External Liquid Assets S_i^0	External Liabilities \bar{p}_{i0}
Interbank Assets $\bar{p}_{ji}, j \neq i$	Interbank Liabilities $\bar{p}_{ij}, j \neq i$
Illiquid Assets e_i with price \bar{q}	
Net Worth w_i	

Table 1: Bank i 's initial balance sheet

3 Model and Problem Formulation

3.1 Basic Framework

We build on the framework of Cifuentes et al. (2005), extending it to incorporate liabilities outside the financial network and stochastic external liquid asset values. We consider a network of n banks indexed by $1, \dots, n$. Each bank i 's initial balance sheet consists of the following items:

- $\bar{p}_{ij} \geq 0$: the nominal liabilities (or payment obligation) to bank j , with $\bar{p}_{ii} = 0$;
- $\bar{p}_{i0} > 0$: the nominal liabilities (or payment obligation) to external entities;
- $S_i^0 > 0$: the initial value of external liquid assets;
- $e_i > 0$: the units of illiquid assets with a nominal unit price of \bar{q} .

Table 1 provides a decomposition of the balance sheet into asset and liabilities. Given this setup, bank i 's total liabilities are denoted by $\bar{p}_i = \bar{p}_{i0} + \sum_{j=1}^n \bar{p}_{ij}$, and its initial net worth is given by $w_i = S_i^0 + \sum_{j=1}^n \bar{p}_{ji} + e_i \bar{q} - \bar{p}_i$. We denote the total illiquid asset holdings across all banks in the system by $\bar{e} := \sum_{i=1}^n e_i$. In line with previous studies (e.g., Cifuentes et al., 2005; Chen et al., 2016; Bernard et al., 2022), all banks are assumed to hold the same type of illiquid assets. This single representative asset can be viewed as an approximation of multiple correlated assets (Weber and Weske, 2017).

Assuming that all liabilities share the same maturity, we model the value of bank i 's external liquid assets at maturity as

$$S_i = S_i^0 \exp \left(-\frac{\sigma_i^2}{2} + \sum_{k=1}^i \Lambda_{ik} Z_k \right), \quad (1)$$

where $(Z_1, \dots, Z_n) \sim \mathcal{N}(\mathbf{0}_n, \mathbf{I}_n)$, $\mathbf{\Lambda} = (\Lambda_{ij})$ is a lower triangular matrix, $\mathbf{\Lambda} \mathbf{\Lambda}^\top = \mathbf{\Sigma}$ is a covariance matrix, $\sigma_i^2 := \sum_{k=1}^i \Lambda_{ik}^2$, and $\mathbf{0}_n$ and \mathbf{I}_n are the n -dimensional vector of zeros and the n -by- n identity matrix, respectively. In other words, S_1, \dots, S_n follow correlated lognormal distributions with $S_i^0 = \mathbb{E}[S_i]$. The risk-free rate is assumed to be zero without loss of generality. We next assume that all liabilities are of equal seniority; that is, when a bank's asset value at maturity falls short of the total liabilities, its debts are repaid to the creditors proportionally to the payment obligations. To illustrate this, we introduce the relative liability matrix $\mathbf{\Pi} = (\pi_{ij}) \in \mathbb{R}_+^{n \times n}$, where $\pi_{ij} := (\bar{p}_{ij}/\bar{p}_i) \mathbf{1}_{\{\bar{p}_i > 0\}}$ represents the proportion of bank i 's liabilities to bank j . Then, if p_j represents

bank j 's total debt payment for each $j = 1, \dots, n$, bank i 's interbank assets, i.e., the total payment received by bank i , become $\sum_{j=1}^n \pi_{ji} p_j$.

For each realization \mathbf{s} of the external liquid asset vector \mathbf{S} , banks with insufficient liquid assets need to sell off illiquid assets to fulfill their payment obligations. Specifically, banks first use their external liquid assets s_i and interbank assets ($\sum_j \pi_{ji} p_j$) to meet the payment obligation. When these assets are not sufficient (i.e., $s_i + \sum_j \pi_{ji} p_j < \bar{p}_i$) and the market price of illiquid assets at the maturity is given by q , bank i either sells $(\bar{p}_i - s_i - \sum_j \pi_{ji} p_j)/q$ units of illiquid assets if the residual asset value is enough to fulfill the obligation (i.e., $s_i + qe_i + \sum_j \pi_{ji} p_j \geq \bar{p}_i$), or liquidates all its illiquid assets to distribute the residual assets to its creditors via the pro rata rule otherwise.

The price q of illiquid assets is determined by the *inverse demand function* $Q : [0, \bar{e}] \rightarrow \mathbb{R}_+$, which maps the total amount of illiquid assets in the market to the price. Consequently, the equilibrium price $q(\mathbf{s})$ and the clearing payments $\{p_i(\mathbf{s})\}_{i=1}^n$ are characterized by a solution to the following fixed-point equations:

$$\begin{cases} q = Q \left(\sum_{i=1}^n \left\{ \frac{(\bar{p}_i - s_i - \sum_j \pi_{ji} p_j)^+}{q} \wedge e_i \right\} \right), \\ p_i = \bar{p}_i \wedge \left(s_i + qe_i + \sum_{j=1}^n \pi_{ji} p_j \right), \quad i = 1, \dots, n, \end{cases} \quad (2)$$

where $x^+ := \max\{x, 0\}$ and $x \wedge y := \min\{x, y\}$ for any $x, y \in \mathbb{R}$. It follows from Theorem 2 of Amini et al. (2016b) that the solution to the above equations is uniquely determined under some mild assumptions on the inverse demand function Q listed below.

Assumption 1 We make the following assumptions:

- (a) $\bar{q} = Q(0) > Q(\bar{e}) > 0$;
- (b) $Q(x)$ is continuous and strictly decreasing in $x \in [0, \bar{e}]$;
- (c) $xQ(x)$ is strictly increasing in $x \in [0, \bar{e}]$.

Given the fixed-point characterization in (2), we say that bank i is solvent if $p_i(\mathbf{s}) = \bar{p}_i$ and that it defaults otherwise. To rule out trivial cases where some banks are initially insolvent, we additionally assume that $p_i(\mathbf{S}^0) = \bar{p}_i$ for all $i = 1, \dots, n$.

3.2 Problem Specification

Under the modeling framework in Section 3.1, we aim to calculate the price of bank n 's bond with a face value of one, i.e.,

$$\hat{p}_n := \mathbb{E} \left[\frac{p_n(\mathbf{S})}{\bar{p}_n} \right],$$

and this quantity can be decomposed as follows:

$$\hat{p}_n = 1 - \mathbb{P}(p_n(\mathbf{S}) < \bar{p}_n) + \mathbb{E}\left[\frac{p_n(\mathbf{S})}{\bar{p}_n} \mathbb{1}_{\{p_n(\mathbf{S}) < \bar{p}_n\}}\right], \quad (3)$$

where $\{p_n(\mathbf{S}) < \bar{p}_n\}$ denotes the event of bank n 's default and $\mathbb{1}_A$ equals one if A occurs and zero otherwise. Based on the above decomposition of the bond price, we develop an efficient Monte Carlo method to estimate the expectation

$$\gamma := \mathbb{E}\left[h(\mathbf{S}) \mathbb{1}_{\{p_n(\mathbf{S}) < \bar{p}_n\}}\right] \quad (4)$$

for any function $h(\cdot)$ uniformly positive and bounded on \mathbb{R}_+^n . This is because the second and last terms in (3) are equal to the expectation (4) when $h(\cdot) \equiv 1$ and $h(\cdot) = p_n(\cdot)/\bar{p}_n$, respectively. Furthermore, as the random variables associated with these two terms are strongly correlated, an efficient Monte Carlo method for (4) can be paired with common random numbers (Glasserman, 2003). This approach effectively estimates both terms while reducing the overall variance in the calculation of the bond price \hat{p}_n .

Estimating expectations of the form $\mathbb{E}[g(\mathbf{X}) \mathbb{1}_{\{\mathbf{X} \in A\}}]$, where g is a function and \mathbf{X} is a random vector, has long been a central problem in the Monte Carlo simulation literature. Our objective (4) falls within this class of problems. Designing an efficient Monte Carlo method for such an objective typically requires a careful geometric understanding of the target set A , which helps us generate more samples from this set to reduce variance. In our setting, however, this poses a major challenge: the explicit representation of the target set is notoriously difficult due to the implicit fixed-point characterization of the equilibrium price and clearing payments in (2). To address this difficulty, prior research has focused on models without illiquid assets and developed Monte Carlo methods that exploit a linear programming reformulation of (2) (Ahn and Kim, 2018; Ahn and Zheng, 2021, 2025). As noted earlier, however, this approach not only incurs a substantial computational burden as the network size increases but cannot be extended to our general framework with illiquid assets.

4 Network Decoupling and Bi-Level Importance Sampling

We address the above-mentioned issue in our general setup by identifying a novel representation of the default event $\{p_n(\mathbf{s}) < \bar{p}_n\}$. This finding enables us to develop our proposed algorithm for estimating γ in (4), referred to as Bi-Level Importance Sampling with Splitting (BLISS).

4.1 Reformulation of Default Events

In this subsection, we focus on analyzing bank n 's default event $\{p_n(\mathbf{s}) < \bar{p}_n\}$ for any fixed $\mathbf{s} \in \mathbb{R}_+^n$. We first note that a bank's solvency is fundamentally dependent on its external liquid assets. Therefore, to exclude the trivial possibility that a bank could remain solvent even with a complete

loss of these assets, we impose the following condition, which is typically satisfied in practice. As will be shown in Section 6.2, it also holds in the real-world example based on actual data.

Assumption 2 For each bank $i = 1, \dots, n$, the value of its total nominal liabilities exceeds the combined nominal value of its illiquid and interbank assets by at least one monetary unit, i.e., $\bar{p}_i \geq 1 + \bar{q}e_i + \sum_{j \neq i} \pi_{ji} \bar{p}_j$.

The main challenge in simulating default events in networks is that a bank's solvency depends implicitly on the entire system's fixed-point equations in (2). This implicit characterization prevents direct application of standard rare-event simulation techniques. We overcome this by decoupling the target bank from the network through a fictitious system that allows us to explicitly characterize the threshold external asset value below which the bank defaults. Using the above assumption, we establish a tractable representation by considering a *fictitious* system in which bank n sells all its illiquid assets and satisfies its payment obligation regardless of its financial status. In this system, for each realization \mathbf{s} of \mathbf{S} , the equilibrium asset price $\tilde{q}(\mathbf{s}_{-n})$ and the clearing payment of banks 1 through $n - 1$, denoted by $\tilde{p}_1(\mathbf{s}_{-n}), \dots, \tilde{p}_{n-1}(\mathbf{s}_{-n})$, are the solution to the following fixed-point equations:

$$\begin{cases} q = Q \left(e_n + \sum_{i=1}^{n-1} \left\{ \frac{(\bar{p}_i - s_i - \pi_{ni} \bar{p}_n - \sum_{j=1}^{n-1} \pi_{ji} p_j)^+}{q} \wedge e_i \right\} \right), \\ p_i = \bar{p}_i \wedge \left(s_i + qe_i + \pi_{ni} \bar{p}_n + \sum_{j=1}^{n-1} \pi_{ji} p_j \right), \quad i = 1, \dots, n-1. \end{cases}$$

It is easy to check that the solution to these equations is also unique based on Assumption 1. Furthermore, we define

$$v_n(\mathbf{s}_{-n}) := \bar{p}_n - \tilde{q}(\mathbf{s}_{-n})e_n - \sum_{j=1}^{n-1} \pi_{jn} \tilde{p}_j(\mathbf{s}_{-n}), \quad (5)$$

where the second and third terms represent the values of bank n 's illiquid assets and interbank assets, respectively, in this fictitious system. The quantity $v_n(\mathbf{s}_{-n})$ in equation (5) represents the minimum external liquid asset value that bank n requires to remain solvent, conditional on the external asset realizations \mathbf{s}_{-n} of all other banks. In the fictitious system, bank n 's net worth is therefore $s_n - v_n(\mathbf{s}_{-n})$. The following proposition establishes that this threshold fully characterizes bank n 's default condition in the original system. Specifically, it shows that the n -dimensional implicit default event, defined through the fixed-point equations in (2), can be equivalently represented as a simple one-dimensional threshold condition on bank n 's external assets.

Proposition 1 (Default Condition Reformulation) *Suppose that Assumptions 1 and 2 hold. Then, bank n defaults in the original system if and only if it has a negative net worth in the above-mentioned fictitious system, i.e., $\{\mathbf{s} \in \mathbb{R}_+^n : p_n(\mathbf{s}) < \bar{p}_n\} = \{\mathbf{s} \in \mathbb{R}_+^n : s_n < v_n(\mathbf{s}_{-n})\}$.*

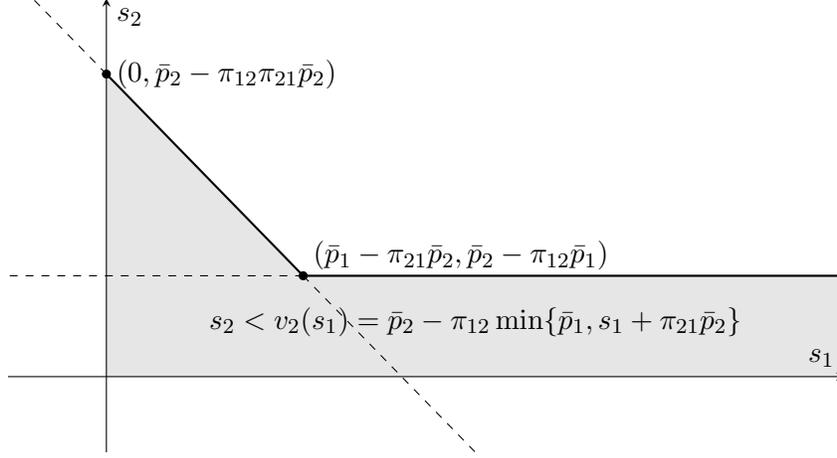


Figure 1: A graphical illustration of the default event $\{\mathbf{s} \in \mathbb{R}_+^2 : p_2(\mathbf{s}) < \bar{p}_2\}$ in a 2-bank system without fire sales. The shaded area represents the set of liquid asset values that cause bank 2 to default. The thick solid line depicts the value of $v_2(s_1)$ for different s_1 .

The essence of Proposition 1 lies in characterizing the minimum external asset value $v_n(\mathbf{s}_{-n})$ required for bank n to remain solvent, given the realized external asset values \mathbf{s}_{-n} of other banks. This result is a nontrivial extension of Lemma 1 in Ahn et al. (2024) under the Eisenberg-Noe framework without fire sales to our generalized modeling framework in Section 3.1. While their result is derived by using a linear program reformulation of the clearing payment characterization and its duality, such a tool is not applicable in our setup due to the nonlinearity introduced by asset liquidation in the clearing payment characterization; see the first equation in (2). Instead, we achieve this generalization by decoupling bank n 's external assets from the fixed-point equations in (2) through the aforementioned fictitious system approach. Without this approach, deriving the critical threshold $v_n(\mathbf{s}_{-n})$ would be challenging because s_n is inherently coupled with the fixed-point equations in (2); consequently, one must resort to numerically solving the system of equations to determine the threshold. See Figure 1 for a description of a default event in a 2-bank system based on the result in Proposition 1.

4.2 BLISS: Introduction

Using the explicit default characterization from Proposition 1, we now construct our simulation methodology. We leverage the fact that the threshold $v_n(\mathbf{S}_{-n})$ suggests a natural splitting. First, at the outer level, we sample the external assets of banks 1 to $n - 1$. Then, at the inner level, we conditionally sample bank n 's external assets. In the special case of $h \equiv 1$, it can be easily checked that such a splitting (or conditional Monte Carlo) approach results in smaller estimation errors than the standard Monte Carlo method:

$$\text{Var}(\mathbb{1}_{\{p_n(\mathbf{s}) < \bar{p}_n\}}) = \text{Var}(\mathbb{1}_{\{S_n < v_n(\mathbf{s}_{-n})\}})$$

$$\begin{aligned}
&= \text{Var}\left(\mathbb{E}\left[\mathbb{1}_{\{S_n < v_n(\mathbf{s}_{-n})\}} \mid \mathbf{S}_{-n}\right]\right) + \mathbb{E}\left[\text{Var}\left(\mathbb{1}_{\{S_n < v_n(\mathbf{s}_{-n})\}} \mid \mathbf{S}_{-n}\right)\right] \\
&> \text{Var}\left(\mathbb{E}\left[\mathbb{1}_{\{S_n < v_n(\mathbf{s}_{-n})\}} \mid \mathbf{S}_{-n}\right]\right) \\
&= \text{Var}\left(\mathbb{P}(S_n < v_n(\mathbf{S}_{-n}) \mid \mathbf{S}_{-n})\right),
\end{aligned}$$

where the strict inequality holds because $\text{Var}\left(\mathbb{1}_{\{S_n < v_n(\mathbf{s}_{-n})\}} \mid \mathbf{S}_{-n} = \mathbf{x}\right) > 0$ for all $\mathbf{x} \in \mathbb{R}_+^{n-1}$.

To generalize this splitting approach for more flexible choices of h and achieve greater variance reduction, we propose a new algorithm, which we call *Bi-Level Importance Sampling with Splitting (BLISS)*, for estimating an equivalent (split) version of (4):

$$\gamma = \mathbb{E}_{\mathbf{S}_{-n}} \left[\mathbb{E}_{S_n} \left[h(\mathbf{S}) \mathbb{1}_{\{S_n < v_n(\mathbf{s}_{-n})\}} \mid \mathbf{S}_{-n} \right] \right]. \quad (6)$$

This algorithm applies two types of importance sampling methods to the outer-layer sampling of \mathbf{S}_{-n} and the inner-layer sampling of S_n in the splitting framework, from which its name is derived.

To be more specific, based on the definition of $\{S_i\}_{i=1}^n$ in (1), we first generate a sample of $\mathbf{Z}_{-n} = (Z_1, \dots, Z_{n-1})^\top$ by *exponential tilting* to form a sample of \mathbf{S}_{-n} . As is well documented in the literature, exponential tilting—also known as exponential twisting or exponential change of measure—is a standard importance sampling (or measure change) technique widely used for rare-event simulation problems involving light-tailed random variables (see, e.g., Glasserman, 2003). A well-known property of this method is that when applied to a normal distribution, the tilted distribution remains normal with the same variance, but its mean is shifted to increase the probability of the rare event of interest (Bucklew, 2004). Accordingly, we take a sample of \mathbf{Z}_{-n} not from its original distribution $\mathcal{N}(\mathbf{0}_{n-1}, \mathbf{I}_{n-1})$ but from its mean-shifted version $\mathcal{N}(\boldsymbol{\mu}_{-n}, \mathbf{I}_{n-1})$, where the choice of the mean vector $\boldsymbol{\mu}_{-n} = (\mu_1, \dots, \mu_{n-1})^\top$ will be discussed in Section 5. Then, the corresponding likelihood ratio for debiasing is given by

$$\mathcal{L}_{\text{out}}(\mathbf{Z}_{-n}) = \exp\left(\frac{1}{2}\|\boldsymbol{\mu}_{-n}\|^2 - \boldsymbol{\mu}_{-n}^\top \mathbf{Z}_{-n}\right), \quad (7)$$

and a sample of \mathbf{S}_{-n} can be obtained by setting $i = n$ and replacing Z_1, \dots, Z_{n-1} in (1) with their samples.

Next, conditional on the samples \mathbf{z}_{-n} and \mathbf{s}_{-n} of \mathbf{Z}_{-n} and \mathbf{S}_{-n} , respectively, we sample Z_n using *near-optimal importance sampling* to construct a sample of S_n . To illustrate this, we define

$$\ell_n(\mathbf{z}_{-n}) := \frac{1}{\Lambda_{nn}} \left(\log S_n^0 - \log v_n(\mathbf{s}_{-n}) - \frac{\sigma_n^2}{2} + \sum_{k=1}^{n-1} \Lambda_{nk} z_k \right). \quad (8)$$

Then, by (1), it is straightforward to see that the default event $\{S_n < v_n(\mathbf{s}_{-n})\}$ can be recast as $\{Z_n < -\ell_n(\mathbf{z}_{-n})\}$. Furthermore, for any random variable X and any given set A , the optimal importance sampling distribution for estimating $\mathbb{P}(X \in A)$ is known as the conditional distribution

Algorithm 1 BLISS(μ_1, \dots, μ_{n-1})

- 1: **procedure** OUTER-LAYER SAMPLING(μ_1, \dots, μ_{n-1})
 - 2: Sample $Z_k = z_k$ from the normal distribution $\mathcal{N}(\mu_k, 1)$ for $k = 1, \dots, n-1$
 - 3: Set $s_i = S_i^0 \exp(-\sigma_i^2/2 + \sum_{k=1}^i \Lambda_{ik} z_k)$ for $i = 1, \dots, n-1$
 - 4: Compute $v_n(\mathbf{s}_{-n})$ and $\mathcal{L}_{\text{out}}(\mathbf{z}_{-n})$ using (5) and (7), respectively
 - 5: **end procedure**
 - 6: **procedure** INNER-LAYER SAMPLING(z_1, \dots, z_{n-1})
 - 7: Calculate $\ell_n(\mathbf{z}_{-n})$ in (8)
 - 8: Sample $Z_n = z_n$ from the truncated normal distribution with density $\psi(\cdot|\mathbf{z}_{-n})$ in (9)
 - 9: Set $s_n = S_n^0 \exp(-\sigma_n^2/2 + \sum_{k=1}^n \Lambda_{nk} z_k)$
 - 10: Compute $\mathcal{L}_{\text{in}}(\mathbf{z}_{-n})$ using (10)
 - 11: **end procedure**
 - 12: Set $T = h(s_1, \dots, s_n) \mathcal{L}_{\text{in}}(\mathbf{z}_{-n}) \mathcal{L}_{\text{out}}(\mathbf{z}_{-n})$
 - 13: **return** T
-

of X given A (Asmussen and Glynn, 2007). Hence, especially for the case of $h \equiv 1$ and when $\mathbf{Z}_{-n} = \mathbf{z}_{-n}$ and $\mathbf{S}_{-n} = \mathbf{s}_{-n}$, the optimal importance sampling distribution of Z_n for estimating the inner conditional expectation in (6) is characterized by the truncated standard normal density

$$\psi(x|\mathbf{z}_{-n}) := \frac{\phi(x) \mathbb{1}_{\{x < -\ell_n(\mathbf{z}_{-n})\}}}{\Phi(-\ell_n(\mathbf{z}_{-n}))} \quad \forall x \in \mathbb{R}, \quad (9)$$

where $\phi(\cdot)$ and $\Phi(\cdot)$ are the density function and cumulative distribution function, respectively, of the univariate standard normal distribution. We use this sampling distribution for general cases of h , leading to the term *near-optimal* importance sampling. The associated likelihood ratio is given by

$$\mathcal{L}_{\text{in}}(\mathbf{z}_{-n}) = \frac{\phi(\cdot)}{\psi(\cdot|\mathbf{z}_{-n})} = \Phi(-\ell_n(\mathbf{z}_{-n})). \quad (10)$$

Note that this likelihood ratio does not depend on the samples of Z_n because the sampling distribution of Z_n is defined on the interval $(-\infty, -\ell_n(\mathbf{z}_{-n}))$, i.e., $\mathbb{1}_{\{Z_n < -\ell_n(\mathbf{z}_{-n})\}} = 1$.

Consequently, our proposed BLISS estimator is finally characterized as

$$h(\tilde{S}_1, \dots, \tilde{S}_n) \mathcal{L}_{\text{in}}(\tilde{\mathbf{Z}}_{-n}) \mathcal{L}_{\text{out}}(\tilde{\mathbf{Z}}_{-n}),$$

where $\tilde{Z}_k \sim \mathcal{N}(\mu_k, 1)$ for $k < n$, $\tilde{Z}_n = (Z_n | Z_n < -\ell_n(\tilde{\mathbf{Z}}_{-n}))$, and $\tilde{S}_i = S_i^0 \exp(-\sigma_i^2/2 + \sum_{k=1}^i \Lambda_{ik} \tilde{Z}_k)$ for $i = 1, \dots, n$. Note that this estimator is unbiased. The complete description of our BLISS procedure is provided in Algorithm 1.

It is worth highlighting that this algorithm's computational complexity is at most linear in the number of banks (n) because the sampling cost increases linearly in n and the calculation of $v_n(\mathbf{s}_{-n})$ requires at most n iterations of the so-called fictitious default algorithm (Weber and Weske, 2017). This is a significant improvement over current state-of-the-art algorithms for rare-event simulation

in financial networks, which exhibit exponential computational complexity relative to the number of banks (see Section 1).

5 Analysis of the BLISS Estimator

We provide theoretical performance guarantees for the BLISS algorithm introduced in the previous section. We mainly focus on situations where the default event is rare. In this case, the naive Monte Carlo estimator is significantly inefficient. For example, when $h \equiv 1$, the standard error of the naive Monte Carlo estimator of γ in (4), quantified by its L^2 norm, is $\sqrt{\gamma}$. This implies that the error is not only greater than the mean γ for sufficiently small γ but also decays much slower than γ as $\gamma \rightarrow 0$. To mathematically evaluate the efficiency of the proposed method, we analyze its performance in two asymptotic regimes governed by a parameter m , in which the default probability $\gamma = \gamma_m$ decreases to 0 as m tends to ∞ . These regimes, detailed in Sections 5.1 and 5.2, are:

- Large asset regime: The initial values of the external liquid assets, S_1^0, \dots, S_n^0 , scale with the parameter m , while their return volatility, $\mathbf{\Lambda}$, remains unchanged;
- Small volatility regime: The volatility of the asset return, $\mathbf{\Lambda}$, decreases with the parameter m , while the initial asset values, S_1^0, \dots, S_n^0 , are fixed.

In the following subsections, we show that the BLISS estimator is *asymptotically optimal* in these two regimes.³ This property—a standard notion of efficiency in rare-event simulation—is formally defined as follows:

Definition 1 Let Γ_m denote an unbiased estimator of γ_m . Then, we say that the estimator Γ_m is asymptotically optimal if

$$\lim_{m \rightarrow \infty} \frac{\log \mathbf{E}[\Gamma_m^2]}{2 \log \mathbf{E}[\Gamma_m]} = 1.$$

The asymptotic optimality of an estimator indicates that its L^2 norm and first moment share the same asymptotic rate of change, thereby achieving significant variance reduction for large m compared to the naive Monte Carlo method. See Asmussen and Glynn (2007) for more details on this efficiency notion.

5.1 Large Asset Regime

In practice, the value of external assets is significantly larger than the values of interbank assets and equities. Motivated by this, we consider an asymptotic regime where S_1^0, \dots, S_n^0 in (1) are

³In Appendix C, we assess the performance of the BLISS estimator in an asymptotic regime where increasing volatility drives the default probability toward one. The analysis, however, is restricted to the case of $h \equiv 1$ in (4).

replaced by $S_{1,m}^0, \dots, S_{n,m}^0$, which grow large as m increases. Based on this regime, the function $\ell_n(\cdot)$ in (8) is converted into

$$\ell_{n,m}^A(\mathbf{x}) := \frac{1}{\Lambda_{nn}} \left(\log S_{n,m}^0 - \frac{\sigma_n^2}{2} - \log(v_n(s_m^A(\mathbf{x}))) + \sum_{k=1}^{n-1} \Lambda_{nk} x_k \right) \text{ for } \mathbf{x} \in \mathbb{R}^{n-1},$$

where $s_m^A(\mathbf{x})$ is an $(n-1)$ -dimensional vector whose i -th element is $S_{i,m}^0 \exp(-\sigma_i^2/2 + \sum_{k=1}^i \Lambda_{ik} x_k)$. Then, our BLISS estimator in the large asset regime becomes

$$h(S_{1,m}^A, \dots, S_{n,m}^A) \mathcal{L}_{\text{in}}(\mathbf{Z}_{-n,m}^A) \mathcal{L}_{\text{out}}(\mathbf{Z}_{-n,m}^A), \quad (11)$$

where $\mathbf{Z}_{-n,m}^A := (Z_{1,m}^A, \dots, Z_{n-1,m}^A) \sim \mathcal{N}(\boldsymbol{\mu}_{-n,m}^A, \mathbf{I}_{n-1})$, $Z_{n,m}^A = (Z_n | Z_n < -\ell_{n,m}^A(\mathbf{Z}_{-n,m}^A))$, and $S_{i,m}^A = S_{i,m}^0 \exp(-\sigma_i^2/2 + \sum_{k=1}^i \Lambda_{ik} Z_{k,m}^A)$ for $i = 1, \dots, n$.

The efficiency of our estimator (11) depends on the choice of the shifted mean vector $\boldsymbol{\mu}_{-n,m}^A$ in the outer-layer importance sampling. We specifically choose this value to reduce the variance of the estimator. Let $\bar{\Phi}(\cdot) := 1 - \Phi(\cdot)$, $\boldsymbol{\lambda}_n = (\Lambda_{n1}, \dots, \Lambda_{n(n-1)})^\top$, and $\kappa_n := \sigma_n^2/2 + \log(v_n(\mathbf{0}))$. Then, as m tends to ∞ , an asymptotic upper bound of the second moment is given by

$$\begin{aligned} & \tilde{\mathbb{E}} \left[h(S_{1,m}, \dots, S_{n,m})^2 \mathcal{L}_{\text{in}}(\mathbf{Z}_{-n,m}^A)^2 \mathcal{L}_{\text{out}}(\mathbf{Z}_{-n,m}^A)^2 \right] \\ & \leq \text{constant} \times \mathbb{E} \left[\bar{\Phi}(\ell_{n,m}^A(\mathbf{Z}_{-n}))^2 \exp \left(\frac{1}{2} \|\boldsymbol{\mu}_{-n,m}^A\|^2 - (\boldsymbol{\mu}_{-n,m}^A)^\top \mathbf{Z}_{-n} \right) \right] \\ & \approx \text{constant} \times \exp \left(\max_{\mathbf{x}} \left\{ -\ell_{n,m}^A(\mathbf{x})^2 + \frac{1}{2} \|\boldsymbol{\mu}_{-n,m}^A\|^2 - (\boldsymbol{\mu}_{-n,m}^A)^\top \mathbf{x} - \frac{1}{2} \|\mathbf{x}\|^2 \right\} \right) \\ & \leq \text{constant} \times \exp \left(\max_{\mathbf{x}} \left\{ -\frac{(\log S_{n,m}^0 - \kappa_n + \boldsymbol{\lambda}_n^\top \mathbf{x})^2}{\Lambda_{nn}^2} + \frac{1}{2} \|\boldsymbol{\mu}_{-n,m}^A\|^2 - (\boldsymbol{\mu}_{-n,m}^A)^\top \mathbf{x} - \frac{1}{2} \|\mathbf{x}\|^2 \right\} \right), \end{aligned}$$

where $\tilde{\mathbb{E}}$ and \mathbb{E} are expectations under the new measure and the original measure, respectively, the first inequality follows from the boundedness of the function $h(\cdot)$, the approximation stems from the Laplace method (Wong, 2001) and the fact that $\log \bar{\Phi}(x)^2 \sim -x^2$ as $x \rightarrow \infty$, and the last inequality holds since $\ell_{n,m}^A(\mathbf{x}) \geq (\log S_{n,m}^0 - \kappa_n + \boldsymbol{\lambda}_n^\top \mathbf{x})/\Lambda_{nn}$ for all \mathbf{x} .

We find $\boldsymbol{\mu}_{-n,m}^A$ that minimizes the above deterministic upper bound of the second moment, meaning that we solve

$$\min_{\boldsymbol{\mu}_{-n,m}^A} \left\{ \max_{\mathbf{x}} \left\{ -\frac{(\log S_{n,m}^0 - \kappa_n + \boldsymbol{\lambda}_n^\top \mathbf{x})^2}{\Lambda_{nn}^2} + \frac{1}{2} \|\boldsymbol{\mu}_{-n,m}^A\|^2 - (\boldsymbol{\mu}_{-n,m}^A)^\top \mathbf{x} - \frac{1}{2} \|\mathbf{x}\|^2 \right\} \right\}.$$

It is easy to see that the objective function is convex in $\boldsymbol{\mu}_{-n,m}^A$ and concave in \mathbf{x} , and thus, the minimax theorem (Sion, 1958) holds. Accordingly, it is equivalent to solving

$$\max_{\mathbf{x}} \left\{ \min_{\boldsymbol{\mu}_{-n,m}^A} \left\{ -\frac{(\log S_{n,m}^0 - \kappa_n + \boldsymbol{\lambda}_n^\top \mathbf{x})^2}{\Lambda_{nn}^2} + \frac{1}{2} \|\boldsymbol{\mu}_{-n,m}^A\|^2 - (\boldsymbol{\mu}_{-n,m}^A)^\top \mathbf{x} - \frac{1}{2} \|\mathbf{x}\|^2 \right\} \right\},$$

where the inner optimum is achieved at $\boldsymbol{\mu}_{-n,m}^A = \mathbf{x}$. Then, the outer optimization can be recast as $\max_{\mathbf{x}} \{ -(\log S_{n,m}^0 - \kappa_n + \boldsymbol{\lambda}_n^\top \mathbf{x})^2 / \Lambda_{nn}^2 - \|\mathbf{x}\|^2 \}$, and one can easily check that its maximum is achieved at $\mathbf{x} = -(\log S_{n,m}^0 - \kappa_n)(\boldsymbol{\lambda}_n \boldsymbol{\lambda}_n^\top + \Lambda_{nn}^2 \mathbf{I})^{-1} \boldsymbol{\lambda}_n$. Consequently, we use

$$\boldsymbol{\mu}_{-n,m}^A = -(\log S_{n,m}^0 - \kappa_n)(\boldsymbol{\lambda}_n \boldsymbol{\lambda}_n^\top + \Lambda_{nn}^2 \mathbf{I})^{-1} \boldsymbol{\lambda}_n \quad (12)$$

as the shifted mean vector for the outer-layer importance sampling, and we denote the corresponding BLISS estimator (11) by Γ_m^A .

Given this setup, the asymptotic optimality of the BLISS estimator is demonstrated in the following theorem.

Theorem 1 (Asymptotic Optimality of Γ_m^A) *Suppose that Assumptions 1 and 2 hold. Then, the BLISS estimator Γ_m^A is asymptotically optimal in the large asset regime.*

Moreover, the following proposition quantifies the variance reduction achieved by using our proposed mean in (12) for outer-layer importance sampling, relative to the approach in which the outer-layer importance sampling is not applied (i.e., the standard case of setting $\boldsymbol{\mu}_{-n,m}^A = \mathbf{0}_{n-1}$).

Proposition 2 (Effectiveness of Optimizing $\boldsymbol{\mu}_{-n,m}^A$) *Suppose that Assumptions 1 and 2 hold. Let Γ_m^0 denote the BLISS estimator (11) with $\boldsymbol{\mu}_{-n,m}^A = \mathbf{0}_{n-1}$. Then, in the large asset regime, we have*

$$\lim_{m \rightarrow \infty} \frac{\log \tilde{\mathbb{E}}[(\Gamma_m^A)^2]}{\log \mathbb{E}[(\Gamma_m^0)^2]} = 2 - \frac{\Lambda_{nn}^2}{\sigma_n^2}.$$

Proposition 2 indicates that correlation among external assets magnifies the performance advantage of our proposed BLISS estimator, Γ_m^A , over its counterpart without outer-layer importance sampling. Specifically, when assets are uncorrelated, the equality $\Lambda_{nn}^2 = \sigma_n^2$ holds, and the second moments of both estimators decay at the same rate. In contrast, higher correlation decreases the ratio $\Lambda_{nn}^2 / \sigma_n^2$, which causes the second moment of Γ_m^A to decay exponentially faster.

5.2 Small Volatility Regime

We now consider another regime where the volatility of external asset returns vanishes. In particular, we replace $\boldsymbol{\Lambda}$ by $\boldsymbol{\Lambda}/m$, which diminishes to zero as m increases. Based on this regime, we replace the function $\ell_n(\cdot)$ in (8) with

$$\ell_{n,m}^B(\mathbf{x}) := \frac{m}{\Lambda_{nn}} \left(\log S_n^0 - \frac{\sigma_n^2}{2m^2} - \log \left(v_n(s_m^B(\mathbf{x})) \right) + \frac{1}{m} \sum_{k=1}^{n-1} \Lambda_{nk} x_k \right) \text{ for } \mathbf{x} \in \mathbb{R}^{n-1},$$

where $s_m^{\mathbf{B}}(\mathbf{x}) \in \mathbb{R}^{n-1}$ is a vector whose i -th element is $S_i^0 \exp(-\sigma_i^2/(2m^2) + \sum_{k=1}^i \Lambda_{ik}x_k/m)$. Then, our BLISS estimator in the small volatility regime becomes

$$h(S_{1,m}^{\mathbf{B}}, \dots, S_{n,m}^{\mathbf{B}}) \mathcal{L}_{\text{in}}(\mathbf{Z}_{-n,m}^{\mathbf{B}}) \mathcal{L}_{\text{out}}(\mathbf{Z}_{-n,m}^{\mathbf{B}}), \quad (13)$$

where $\mathbf{Z}_{-n,m}^{\mathbf{B}} := (Z_{1,m}^{\mathbf{B}}, \dots, Z_{n-1,m}^{\mathbf{B}}) \sim \mathcal{N}(\boldsymbol{\mu}_{-n,m}^{\mathbf{B}}, \mathbf{I}_{n-1})$, $Z_n^{\mathbf{B}} = (Z_n | Z_n < -\ell_{n,m}^{\mathbf{B}}(\mathbf{Z}_{-n,m}^{\mathbf{B}}))$, and $S_{i,m}^{\mathbf{B}} = S_i^0 \exp(-\sigma_i^2/(2m^2) + \sum_{k=1}^i \Lambda_{ik}Z_k/m)$ for $i = 1, \dots, n$.

Analogously to Section 5.1, we choose $\boldsymbol{\mu}_{-n,m}^{\mathbf{B}}$ to minimize an asymptotic upper bound of the second moment of the BLISS estimator. Using a similar argument, it is easy to check that the bound is minimized at

$$\boldsymbol{\mu}_{-n,m}^{\mathbf{B}} = m \times \underset{\mathbf{x} \in \mathbb{R}^{n-1}}{\text{argmin}} \left\{ \ell_n^{\mathbf{B}}(\mathbf{x})^2 + \|\mathbf{x}\|^2 \right\}, \quad (14)$$

where $\ell_n^{\mathbf{B}}(\mathbf{x}) := \lim_{m \rightarrow \infty} \ell_{n,m}^{\mathbf{B}}(m\mathbf{x})/m$ for all \mathbf{x} . Accordingly, we denote by $\Gamma_m^{\mathbf{B}}$ the BLISS estimator (13) that uses (14) as the shifted mean for the outer-layer importance sampling. In the following two results, we show the asymptotic optimality of our proposed estimator $\Gamma_m^{\mathbf{B}}$ and its superiority over the case without outer-layer importance sampling.

Theorem 2 (Asymptotic Optimality of $\Gamma_m^{\mathbf{B}}$) *Suppose that Assumptions 1 and 2 hold. Then, the BLISS estimator $\Gamma_m^{\mathbf{B}}$ is asymptotically optimal in the small volatility regime.*

Proposition 3 (Effectiveness of Optimizing $\boldsymbol{\mu}_{-n,m}^{\mathbf{B}}$) *Suppose that Assumptions 1 and 2 hold. Let $\Gamma_m^{\mathbf{O}}$ denote the BLISS estimator (13) with $\boldsymbol{\mu}_{-n,m}^{\mathbf{B}} = \mathbf{0}_{n-1}$. Then, in the small volatility regime, we have*

$$\lim_{m \rightarrow \infty} \frac{\log \mathbb{E}[(\Gamma_m^{\mathbf{B}})^2]}{\log \mathbb{E}[(\Gamma_m^{\mathbf{O}})^2]} = \frac{\min_{\mathbf{x}} \{ \ell_n^{\mathbf{B}}(\mathbf{x})^2 + \|\mathbf{x}\|^2 \}}{\min_{\mathbf{x}} \{ \ell_n^{\mathbf{B}}(\mathbf{x})^2 + \|\mathbf{x}\|^2/2 \}} \geq 1.$$

Moreover, the above limit is strictly greater than one if and only if $\boldsymbol{\lambda}_n \neq \mathbf{0}$.

Proposition 3 characterizes the difference in the decay rates of the second moments of $\Gamma_m^{\mathbf{B}}$ and $\Gamma_m^{\mathbf{O}}$ for the small volatility regime. Although this characterization differs from that for the large asset regime (Proposition 2), the primary implication remains the same: the two estimators' performances are comparable for uncorrelated assets, but the superiority of $\Gamma_m^{\mathbf{B}}$ over $\Gamma_m^{\mathbf{O}}$ in estimation efficiency grows with correlation.

6 Computational Performance and Empirical Validation

To provide numerical validation of the preceding theoretical analyses and discussions, we begin with a synthetic dataset that illustrates the computational efficiency of the BLISS estimator against the state-of-the-art as network size scales. We subsequently analyze a real-world dataset to evaluate its simulation efficiency relative to the naive Monte Carlo method as the rarity of default events varies. In all implementations, we use MATLAB R2023a on a Linux workstation equipped with Dual Intel Xeon E5-2697 2.6GHz CPU and 128 GB RAM.

6.1 Computational Complexity: A Toy Example

We consider n -bank financial networks with $\bar{p}_i = 5$, $\bar{p}_{i0} = 4$, $e_i = 0$, $S_i^0 = S^0$, $\Lambda_{ii} = \sigma$, and $\Lambda_{ij} = 0$ for all $i = 1, \dots, n$ and $j \neq i$, where the values of S^0 and σ will be specified later. We construct two relative liability matrices $\mathbf{\Pi}^c$ and $\mathbf{\Pi}^r$, representing complete and ring networks, respectively:

$$\mathbf{\Pi}^c = \begin{bmatrix} 0 & \frac{.2}{n-1} & \cdots & \frac{.2}{n-1} \\ \frac{.2}{n-1} & 0 & \ddots & \vdots \\ \vdots & \ddots & \ddots & \frac{.2}{n-1} \\ \frac{.2}{n-1} & \cdots & \frac{.2}{n-1} & 0 \end{bmatrix} \quad \text{and} \quad \mathbf{\Pi}^r = \begin{bmatrix} 0 & .2 & 0 & \cdots & 0 \\ 0 & 0 & .2 & \ddots & \vdots \\ \vdots & \vdots & \ddots & \ddots & 0 \\ 0 & 0 & \cdots & 0 & .2 \\ .2 & 0 & \cdots & 0 & 0 \end{bmatrix}.$$

Both matrices satisfy $\sum_{j=1}^n \pi_{ij} = (\bar{p}_i - \bar{p}_{i0})/\bar{p}_i = 0.2$ for all $i = 1, \dots, n$. Under these configurations, we evaluate the relative error and running time of our BLISS method against the conditional importance sampling (CIS) method (Ahn and Zheng, 2023, 2025) and the naive Monte Carlo (MC), when network size (n) increases from 4 to 12. Note that the relative error is defined as the standard error divided by the target estimate. We exclude illiquid assets in this example because, as alluded to earlier, the CIS method cannot be used once fire sales are incorporated. For simplicity, we set $h \equiv 1$ and aim to estimate bank n 's default probability in this example.

Figure 2 presents the associated numerical results for varying values of S^0 and σ , where we use 10,000 simulation trials to implement each of the three methods. The left column corresponds to the results with the complete network case, while the right column depicts those with the ring network configuration. The top and bottom rows compare the relative errors and running times, respectively, of the three methods across three different (S^0, σ) scenarios: $(5, 0.1)$, $(5, 0.08)$, and $(5.5, 0.1)$. In these three scenarios, the probability estimates are of order 10^{-2} , 10^{-3} , and 10^{-4} , respectively, consistently across all values of n .⁴

The figure demonstrates that the BLISS method completely dominates the state-of-the-art CIS method in terms of both relative error and running time. Specifically, BLISS achieves relative errors an order of magnitude smaller than those of CIS, and this advantage becomes more pronounced as target events become rarer (i.e., as probability estimates decrease). Furthermore, BLISS maintains constant running times regardless of network size, whereas CIS experiences exponential growth in running time as the network expands. Although the MC method is comparable to BLISS in speed, it suffers from substantially higher relative errors. Finally, the above observations turn out to be consistent across different network structures.

⁴In this example, for fixed values of (S^0, σ) , the probability estimates remain largely unaffected by network size. This is because bank n 's interbank assets do not change with n , and the parameter setups are homogeneous across all banks.

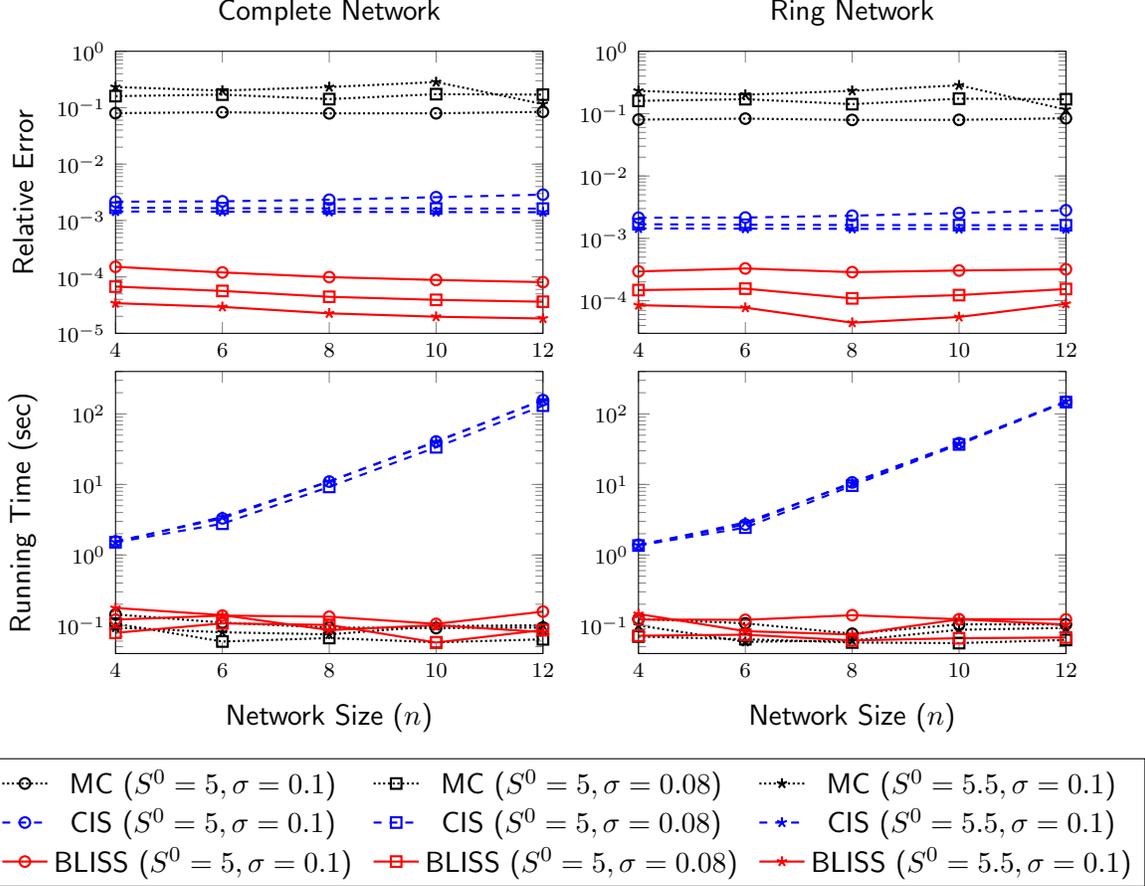


Figure 2: Performance comparison between MC, CIS, and BLISS for the toy example in Section 6.1 in terms of relative error (top panels) and running time (bottom panels) when network size changes. The first and second columns correspond to the cases of the complete and ring networks, respectively.

6.2 Simulation Efficiency: Evidence from European Banking Network

In this subsection, we draw on the European Banking Authority (EBA) stress test dataset (European Banking Authority, 2018), which covers major European banks in 2017, to assess the simulation efficiency of the BLISS method. We first describe our data collection and model calibration process and then illustrate the numerical experiments and their outcomes.

6.2.1 Data Selection and Calibration

In the EBA stress test dataset, we use its Internal Ratings-Based (IRB) total exposure value and Tier 1 capital, both sourced from the dataset, as proxies for initial total assets and initial net worth (w_i), respectively. Each bank i 's total liabilities (\bar{p}_i) are defined as the difference between the above two values. We further define its total interbank assets, ($\sum_{j \neq i} \bar{p}_{ji}$), as the IRB exposures to other institutions reported in the dataset. While the dataset covers 48 banks, we exclude 12

institutions with no interbank exposures, resulting in a final sample of 36 banks ($n = 36$). The raw data on total assets, net worth, and interbank assets for these 36 banks (in millions of euros) are provided in Table 2; the same dataset is used in Bernard et al. (2022). We use the numbers in the first column of the table for the subscripts of variables; e.g., HSBC’s net worth is denoted by $w_1 = 125,976$. Additionally, the final column of Table 2 presents the return volatilities of the corresponding equities in 2017, computed from daily stock price data provided by MarketWatch. For the ten banks not listed on stock exchanges, equity volatilities are set equal to the 2017 volatility of the EURO STOXX Banks Index and marked with *. We note that throughout 2017, European risk-free rates remained negligible, consistent with our framework. We therefore maintain the assumption of a zero risk-free rate.

Using the raw data in Table 2, we construct external liabilities (\bar{p}_{i0}) by imposing a common ratio (α) of external liabilities to total liabilities across all banks; that is, $\bar{p}_{i0} = \alpha \bar{p}_i$ for all i . This ratio is calibrated such that the aggregate interbank liabilities match the aggregate interbank assets within the network, i.e., $\sum_{i=1}^n (\bar{p}_i - \bar{p}_{i0}) = \sum_{i=1}^n \sum_{j=1}^n \bar{p}_{ji}$, or equivalently, $\alpha = 1 - (\sum_{i=1}^n \sum_{j=1}^n \bar{p}_{ji}) / (\sum_{i=1}^n \bar{p}_i)$. Moreover, the initial value of external liquid assets (S_i^0) is set to a fixed portion (β) of its nominal non-interbank assets: $S_i^0 = \beta(w_i + \bar{p}_i - \sum_{j=1}^n \bar{p}_{ji})$. The remainder of the non-interbank assets defines the nominal value of illiquid assets: $e_i = (1 - \beta)(w_i + \bar{p}_i - \sum_{j=1}^n \bar{p}_{ji})$. We use the inverse demand function given by $Q(x) = \exp(-\nu x)$ as in Cifuentes et al. (2005), Chen et al. (2016), and Bernard et al. (2022), with $\nu = 2.5 \times 10^{-8}$ arbitrarily chosen. We also set $\beta = 40\%$ in our experiments. This choice is informed by T02.03.1 in ECB Banking Supervision (2018), which shows that liquid assets accounted for approximately 40% of the total assets of major European banks consistently throughout 2017. One can easily check that Assumptions 1 and 2 are satisfied under these specifications.

Moreover, we calculate the volatility σ_i of each bank i based on the Merton model (Merton, 1974); specifically, we solve a fixed-point equation: $\sigma_i = (w_i/A_i)\sigma_i^E/\Delta(\sigma_i)$, where w_i , A_i , and σ_i^E denote bank i ’s net worth, total assets, and equity volatility, respectively, reported in Table 2. The term $\Delta(\sigma_i)$ represents the Black–Scholes delta of a call option on the bank’s total asset value, with the strike price set to its total liabilities, a zero risk-free rate, and a one-year maturity, and depends on the asset volatility σ_i . While the solution to the above fixed-point equation represents the volatility of the total assets, we use it as a proxy for the volatility of the external liquid assets. With this, we define the matrix $\mathbf{\Lambda}$ as $\mathbf{\Lambda} = \text{diag}(\sigma_1, \dots, \sigma_{36}) \cdot \mathbf{R}$, where \mathbf{R} is the matrix in Table 4 of Appendix B for the correlated case and the identity matrix for the uncorrelated case. In both cases, the diagonal elements of the matrix $\mathbf{\Lambda}\mathbf{\Lambda}^\top$ correspond to $\sigma_1^2, \dots, \sigma_{36}^2$.

Since bilateral interbank exposure data are unavailable, we reconstruct three network topologies—core-periphery, complete, and ring—used in Chen et al. (2016); see Figure 3 for their graphical illustration. This reconstruction is based on an entropy-minimization method in Upper and Worms (2004), defining the core as the ten largest banks by total assets (Banks 1 to 10). The core-periphery network is a relatively sparse, non-extreme structure widely observed in real-world financial sys-

No.	Bank name	Total asset	Net worth	Interbank asset	Equity vol.
1	HSBC Holdings Plc	1,322,909	125,976	117,004	16.46%
2	Groupe Cr�dit Agricole	1,047,925	84,292	97,114	23.34%
3	BNP Paribas	1,012,707	84,417	63,604	22.17%
4	ING Groep N.V.	780,776	50,325	76,469	17.21%
5	Deutsche Bank AG	758,140	57,631	58,015	31.54%
6	Groupe BPCE	668,255	59,490	32,956	*19.41%
7	Soci�t� G�n�rale S.A.	642,940	49,514	53,400	25.53%
8	Lloyds Banking Group Plc	590,827	40,948	8,817	17.49%
9	Banco Santander S.A.	565,109	77,283	36,878	22.70%
10	Barclays Plc	562,002	60,765	49,797	21.87%
11	Co�peratieve Rabobank U.A.	547,353	37,204	14,461	*19.41%
12	The Royal Bank of Scotland Group Plc	490,122	44,577	23,685	23.29%
13	Nordea Bank	437,347	28,008	40,127	18.00%
14	Group Cr�dit Mutuel	430,308	45,578	44,606	*19.41%
15	UniCredit S.p.A.	395,077	54,703	39,890	34.86%
16	ABN AMRO Group N.V.	367,487	19,618	14,942	19.82%
17	Commerzbank AG	314,214	25,985	42,564	28.94%
18	Intesa Sanpaolo S.p.A.	309,144	43,466	36,125	21.10%
19	Banco Bilbao Vizcaya Argentaria S.A.	276,960	46,980	75,226	22.20%
20	Danske Bank	273,199	20,302	14,926	16.13%
21	Svenska Handelsbanken	253,639	12,954	7,339	16.53%
22	KBC Group NV	226,455	16,552	17,128	18.27%
23	Skandinaviska Enskilda Banken	209,082	13,452	14,944	14.99%
24	Landesbank Baden-W�rttemberg	206,824	12,795	57,434	*19.41%
25	Swedbank	202,830	11,356	6,522	15.90%
26	DZ BANK AG	202,301	19,923	35,800	*19.41%
27	Bayerische Landesbank	193,192	9,393	22,731	*19.41%
28	Erste Group Bank AG	147,479	15,368	10,756	21.08%
29	Belfius Banque SA	127,440	8,141	35,597	*19.41%
30	Landesbank Hessen-Th�ringen	122,115	8,180	15,767	*19.41%
31	Banco de Sabadell S.A.	108,282	11,111	1,559	28.60%
32	OP Financial Group	91,467	9,973	7,299	*19.41%
33	Norddeutsche Landesbank	71,764	6,229	16,037	*19.41%
34	Bank of Ireland Group plc	68,264	7,617	4,537	29.86%
35	Raiffeisen Bank International AG	61,499	9,839	5,416	31.87%
36	Allied Irish Banks Group plc	48,157	11,028	10,064	56.28%

Table 2: Balance sheet characteristics of 36 major European banks from the 2018 EBA stress test. Total assets, net worth, and interbank assets are in millions of euros. Equity volatility is annualized. Banks marked with * are not publicly traded. We proxy their equity volatility using the volatility of the EURO STOXX Banks index. We use this data to calibrate the model as described in Section 6.2.1.

tems (Craig and von Peter, 2014). By contrast, the complete and ring structures represent two extremes of connectivity: the complete network is maximally dense, with all 36 banks interconnected (1,260 edges), while the ring network is minimally connected and sparse, containing only 74 edges.

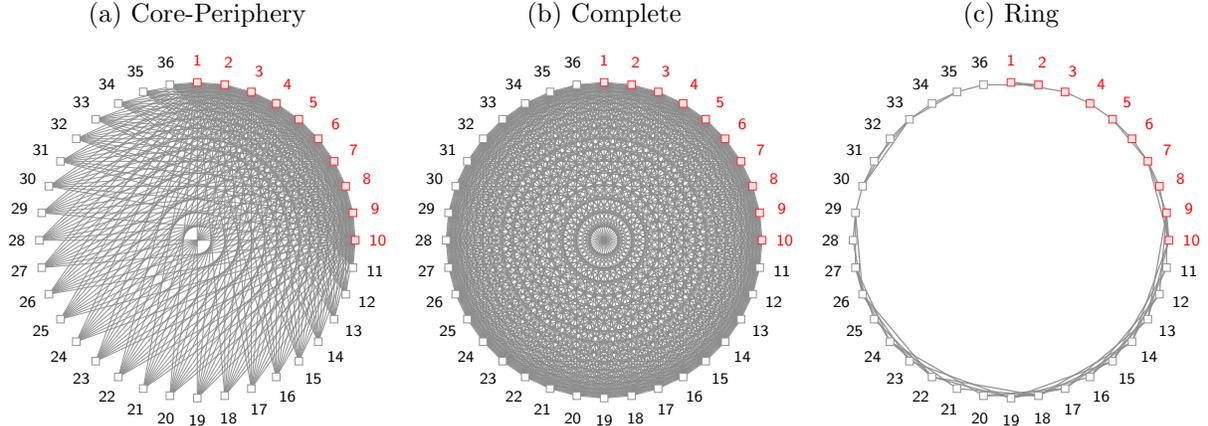


Figure 3: The recovered interbank network structures from the 2018 EBA stress test data. For each structure, nodes represent individual banks, and edges stand for their interbank exposures. The red-colored nodes correspond to the core banks (banks 1 to 10).

6.2.2 Numerical Experiments and Results

In this subsection, we compute the bond price defined in (3) in the setup described above. We select Allied Irish Banks Group plc as the target institution. As one of the largest banks in the Republic of Ireland, it ranks 36th in the EBA dataset. This choice is motivated by the fact that it exhibits the highest volatility among the banks under consideration, making the assessment of its bond price particularly meaningful. We consider scenarios for both correlated and uncorrelated assets. We use standard Monte Carlo (MC) as a benchmark to evaluate the performance of our BLISS methodology. Furthermore, to highlight the importance of outer-layer importance sampling, we compare our approach against a version of BLISS that excludes this step, which we refer to as Inner-Layer Importance Sampling (ILIS). For all numerical experiments in this subsection, we use 10^6 simulation trials to implement MC, ILIS, and BLISS.

Figures 4 and 5 illustrate the performance of the above-mentioned three methods (MC, ILIS, and BLISS) for the core-periphery network under varying asset values and volatilities, respectively.⁵ The asset multiplier and volatility multiplier represent scaling factors applied to the initial price S^0 and matrix $\mathbf{\Lambda}$; these figures correspond to large asset and small volatility regimes, respectively. We report standard errors and efficiency ratios. The efficiency ratio is defined as the ratio of the product of variance and runtime for the crude Monte Carlo method to that of the target method. It is widely regarded as a comprehensive measure of efficiency, accounting for differences in both variance and runtime (Glasserman et al., 2008; Ahn and Zheng, 2025). In Table 3, we additionally provide yield estimates for each of the asset and volatility multipliers we consider, from which we confirm that higher asset multipliers and lower volatility multipliers both result in rarer default events and, consequently, smaller spreads.

⁵The figures for the complete and ring networks are nearly identical to those for core-periphery networks and are thus deferred to Appendix B.

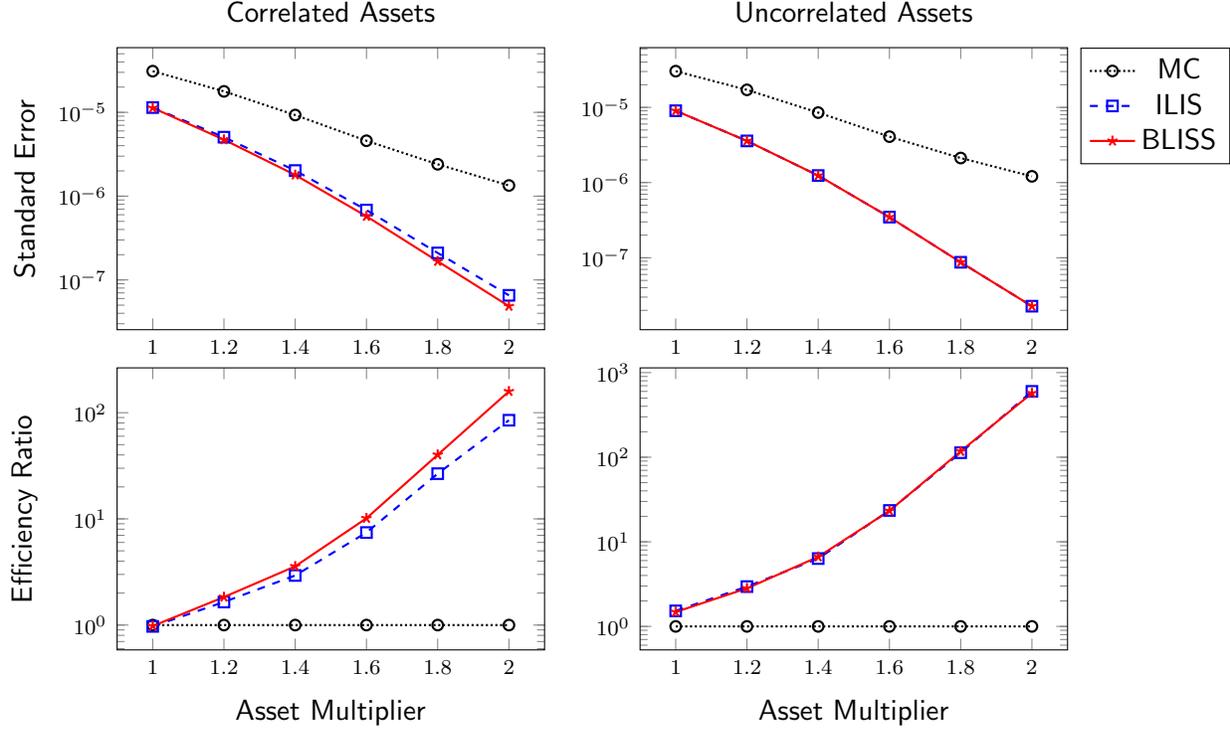


Figure 4: Performance comparison between MC, ILIS, and BLISS for the real-world example in Section 6.2 with the core-periphery network when asset value changes. The top panels compare standard errors of the three methods, and the bottom panels exhibit their efficiency ratios. The left panels correspond to the case where liquid assets are correlated across different banks, while the right panels describe the case of uncorrelated assets.

In both figures, BLISS demonstrates superior performance in terms of both standard errors and efficiency ratios, which remains robust across all cases we consider. As predicted by our theoretical results, the performance gap widens significantly in rare-event regimes, i.e., where the price approaches one. Conversely, crude Monte Carlo performs poorly in all cases. A comparison between ILIS and BLISS reveals that when assets are uncorrelated, the two methods perform almost identically. This is because the outer-layer importance sampling in BLISS is effective only under correlation (specifically, $\mu_{-n,m}^A$ and $\mu_{-n,m}^B$ become zero vectors in the absence of correlation). However, when assets are correlated, ILIS provides negligible variance reduction compared to MC but incurs a higher computational cost, resulting in poorer efficiency ratios. In contrast, BLISS significantly outperforms MC in these scenarios, underscoring the critical importance of the outer-layer sampling strategy. Furthermore, it is commonly observed that the performance gap between ILIS and BLISS gets larger as default events become rarer, which corresponds to our theoretical findings in Propositions 2 and 3.

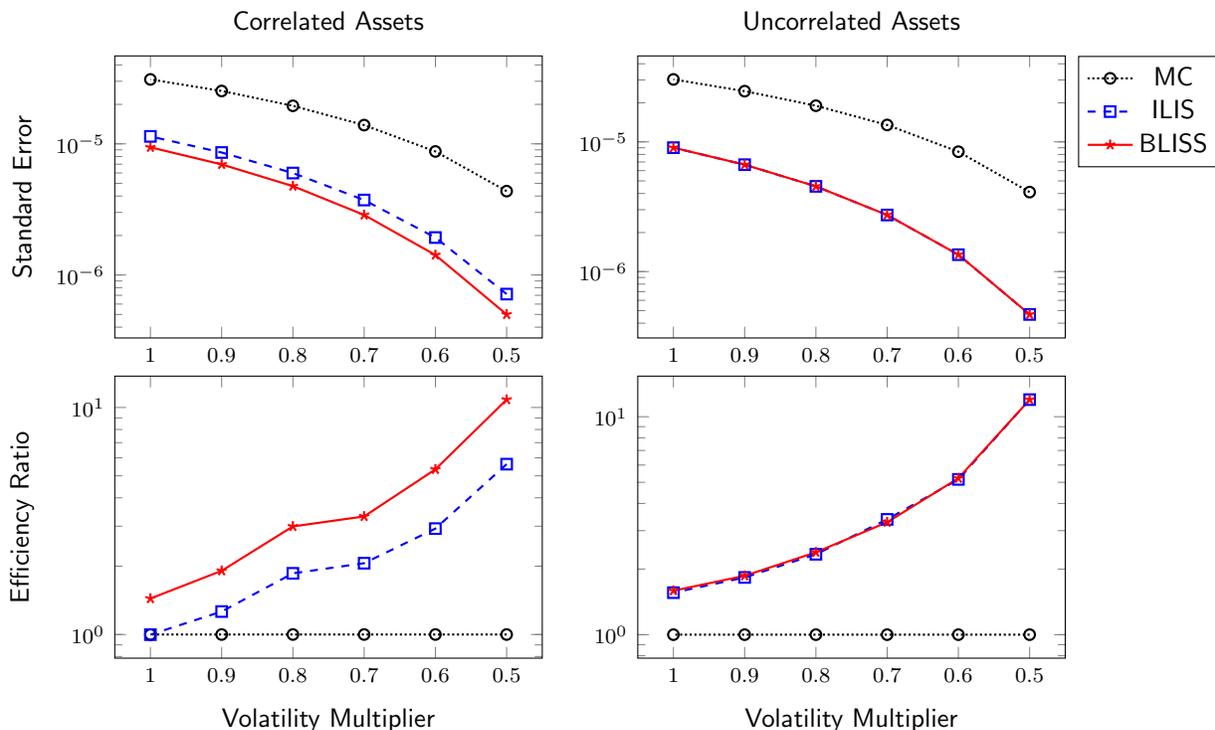


Figure 5: Performance comparison between MC, ILIS, and BLISS for the real-world example in Section 6.2 with the core-periphery network when volatility changes. The top panels compare standard errors of the three methods, and the bottom panels exhibit their efficiency ratios. The left panels correspond to the case where liquid assets are correlated across different banks, while the right panels describe the case of uncorrelated assets.

Asset Mult.	Yield (correlated)	Yield (uncorrelated)	Vol. Mult.	Yield (correlated)	Yield (uncorrelated)
1	132	128	1	132	128
1.2	48	45	0.9	94	91
1.4	14	13	0.8	62	59
1.6	4	3	0.7	35	34
1.8	1.1	0.9	0.6	16	15
2	0.4	0.3	0.5	5	5

Table 3: Yield estimates (in basis points) for different asset and volatility multipliers. The left and right tables show the estimates under the experimental settings of Figures 4 and 5, respectively.

7 Conclusion

In this paper, we have developed an efficient simulation framework for pricing corporate bonds of institutions, which account for network-induced credit risk. The valuation of corporate securities within these complex, interconnected systems is challenging because of the lack of closed-form solutions and the inefficiency of standard simulation methods in capturing rare default events.

To overcome these challenges, we proposed a novel computational framework—Bi-Level Importance Sampling with Splitting (BLISS)—to estimate bond prices in a network setting with simultaneous clearing of interbank liabilities which additionally accounts for fire sales. Specifically, we introduced a judiciously designed fictitious system that helps detach the target institution from the intertwined system, enabling us to reformulate its default condition in a tractable way. This allowed us to decompose the valuation problem into a tractable two-layer sampling process and apply exponential tilting for the outer-layer sampling and optimal importance sampling for the inner-layer sampling to achieve significant variance reduction. Our theoretical analysis demonstrated that (1) BLISS’s computational complexity grows linearly with the network size and (2) it is asymptotically optimal in large-asset and low-volatility regimes, underscoring its robustness in handling rare events. Numerical experiments clearly confirmed the effectiveness of our methodology in high-dimensional rare-event settings.

Several promising directions for future research emerge from this work. First, the BLISS framework could be extended to value more complex securities, including credit derivatives, contingent convertible bonds, or collateralized debt obligations, where network effects play an increasingly important role. Second, incorporating dynamic network structures, where interbank exposures evolve over time, would enhance the realism of the model and enable term-structure pricing of corporate bonds. Third, the methodology could be adapted to stress testing and regulatory capital assessment, where financial institutions must evaluate tail-risk scenarios across large portfolios efficiently.

A Proofs of the Main Theoretical Results

Proof of Proposition 1. For any $\mathbf{s}, \mathbf{s}' \in \mathbb{R}_+^n$ satisfying $\mathbf{s} \leq \mathbf{s}'$, let

$$\begin{cases} q^1 = Q \left(\sum_{i=1}^n \left\{ \frac{(\bar{p}_i - s'_i - \sum_j \pi_{ji} p_j(\mathbf{s}))^+}{q(\mathbf{s})} \wedge e_i \right\} \right), \\ \mathbf{p}^1 = \bar{\mathbf{p}} \wedge \left(\mathbf{s}' + q(\mathbf{s}) \mathbf{e} + \mathbf{\Pi}^\top \mathbf{p}(\mathbf{s}) \right). \end{cases}$$

and

$$\begin{cases} q^{k+1} = Q \left(\sum_{i=1}^n \left\{ \frac{(\bar{p}_i - s'_i - \sum_j \pi_{ji} p_j^k)^+}{q^k} \wedge e_i \right\} \right), \\ \mathbf{p}^{k+1} = \bar{\mathbf{p}} \wedge \left(\mathbf{s}' + q^k \mathbf{e} + \mathbf{\Pi}^\top \mathbf{p}^k \right) \end{cases}$$

for all $k \geq 1$. Then, by Assumption 1(b), it is easy to check that $q(\mathbf{s}) \leq q^1 \leq q^2 \leq \dots$ and $\mathbf{p}(\mathbf{s}) \leq \mathbf{p}^1 \leq \mathbf{p}^2 \leq \dots$. Furthermore, we have $q^k \rightarrow q(\mathbf{s}')$ and $\mathbf{p}^k \rightarrow \mathbf{p}(\mathbf{s}')$ as $k \rightarrow \infty$. Thus, we have

$$q(\mathbf{s}) \leq q(\mathbf{s}') \quad \text{and} \quad \mathbf{p}(\mathbf{s}) \leq \mathbf{p}(\mathbf{s}') \quad \text{for any } \mathbf{s}, \mathbf{s}' \in \mathbb{R}_+^n \text{ satisfying } \mathbf{s} \leq \mathbf{s}'. \quad (15)$$

Fix $\tilde{s}_1, \dots, \tilde{s}_{n-1} \in \mathbb{R}_+$. Define a real-valued function $\zeta(\cdot)$ by

$$\zeta(s_n) = s_n + q(\tilde{s}_1, \dots, \tilde{s}_{n-1}, s_n)e_n + \sum_{j=1}^{n-1} \pi_{jn} p_j(\tilde{s}_1, \dots, \tilde{s}_{n-1}, s_n) - \bar{p}_n,$$

which is strictly increasing by (15). Observe that $\zeta(s_n) < 0$ for all s_n small enough since we have $\bar{p}_n > Q(0)e_n + \sum_{j=1}^{n-1} \pi_{jn} \bar{p}_j$, whereas $\zeta(s_n) > 0$ for all s_n large enough. Hence, it is straightforward that there exists a unique value s_n^* satisfying $\zeta(s_n^*) = 0$. Then, we have

$$\frac{\bar{p}_n - s_n^* - \sum_j \pi_{jn} p_j(\tilde{s}_1, \dots, \tilde{s}_{n-1}, s_n^*)}{q(\tilde{s}_1, \dots, \tilde{s}_{n-1}, s_n^*)} = e_n \quad \text{and} \quad p_n(\tilde{s}_1, \dots, \tilde{s}_{n-1}, s_n^*) = \bar{p}_n.$$

This suggests that $q(\tilde{s}_1, \dots, \tilde{s}_{n-1}, s_n^*) = \tilde{q}(\tilde{s}_1, \dots, \tilde{s}_{n-1})$ and $p_i(\tilde{s}_1, \dots, \tilde{s}_{n-1}, s_n^*) = \tilde{p}_i(\tilde{s}_1, \dots, \tilde{s}_{n-1})$ for $i = 1, 2, \dots, n-1$ by the uniqueness of $q(\cdot)$, $\tilde{q}(\cdot)$, $\mathbf{p}(\cdot)$, and $\tilde{\mathbf{p}}(\cdot)$. Therefore, $s_n^* = v_n(\tilde{s}_1, \dots, \tilde{s}_{n-1})$. Accordingly, we have

$$\begin{aligned} p_n(\tilde{s}_1, \dots, \tilde{s}_{n-1}, s_n) < \bar{p}_n &\Leftrightarrow \zeta(s_n) < 0 \\ &\Leftrightarrow s_n < s_n^* = v_n(\tilde{s}_1, \dots, \tilde{s}_{n-1}). \end{aligned}$$

This completes the proof since $\tilde{s}_1, \dots, \tilde{s}_{n-1}$ are chosen arbitrarily. \square

Proof of Theorem 1. We prove this result in three steps.

Step 1. Let $I(Y) = \int_{-\infty}^{\infty} \bar{\Phi}(Y + Ax)^2 \phi(x) dx$, where $A \in \mathbb{R}$ is a constant. We claim that

$$I(Y) \sim \frac{(1 + 2A^2)^{3/2}}{2\pi Y^2} \exp\left(-\frac{Y^2}{1 + 2A^2}\right) \quad \text{as } Y \rightarrow \infty. \quad (16)$$

By Mill's ratio, we observe that

$$\bar{\Phi}(u)^2 = \frac{1}{2\pi u^2} e^{-u^2} (1 + \delta(u)),$$

where $\delta(u) = O(u^{-2})$ as $u \rightarrow \infty$. By substituting the leading order term of Mill's ratio, we define $\iota(\cdot)$ that approximates the integrand of I :

$$\iota(x) = \frac{\exp(\Psi(x))}{(2\pi)^{3/2} (Y + Ax)^2}, \quad (17)$$

where $\Psi(x) := -(Y + Ax)^2 - x^2/2 = -Y^2 - 2AYx - (A^2 + 1/2)x^2$. Then $x^* := -2AY/(1 + 2A^2)$ is the unique solution to the first-order equation $\Psi'(x) = 0$.

By the Laplace method (Wong, 2001), the prefactor term $g(x) := 1/(Y + Ax)^2$ in the above integral is approximated by $g(x^*) = (1 + 2A^2)^2/Y^2$ as $Y \rightarrow \infty$. Also, a simple calculation shows that $\Psi(x) = -Y/(1 + 2A^2) - (1 + 2A^2)(x - x^*)^2/2$. Hence, by combining these results with (17),

we have

$$I(Y) \sim \frac{(1+2A^2)^2}{2\pi Y^2} \exp\left(-\frac{Y^2}{1+2A^2}\right) \int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}} e^{-\frac{(1+2A^2)}{2}(x-x^*)^2} dx$$

as $Y \rightarrow \infty$. Finally, (16) holds since the Gaussian integral evaluates to:

$$\int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}} e^{-\frac{(1+2A^2)}{2}(x-x^*)^2} dx = \sqrt{\frac{1}{1+2A^2}}.$$

One can similarly show that

$$\int_{-\infty}^{\infty} \bar{\Phi}(Y+Ax)\phi(x) dx \sim \frac{\sqrt{1+A^2}}{\sqrt{2\pi Y}} \exp\left(-\frac{Y^2}{2+2A^2}\right) \text{ as } Y \rightarrow \infty.$$

Step 2. We define $\tau_n := \sigma_n^2/2 + \log(\bar{p}_n - Q(e_n)e_n - \sum_{j=1}^{n-1} \pi_{jn}\bar{p}_j) > 0$, where the inequality holds by Assumptions 1 and 2. Then, recalling that $\kappa_n := \sigma_n^2/2 + \log(v_n(\mathbf{0})) < \infty$, it is easy to see that $\tau_n \leq \sigma_n^2/2 + \log(v_n(s_m^A(\mathbf{x}))) \leq \kappa_n$ for all \mathbf{x} and m . This implies that for any fixed \mathbf{x} , $\ell_{n,m}^A(\mathbf{x})$ increases to infinity as m grows. We let $\boldsymbol{\mu}_m = \boldsymbol{\mu}_{-n,m}^A$ for notational convenience. Then, based on the above observation, we find the following relationship for the second moment of the BLISS estimator:

$$\begin{aligned} \tilde{\mathbb{E}}[(\Gamma_m^A)^2] &\leq h_u^2 \mathbb{E} \left[\bar{\Phi}(\ell_{n,m}^A(\mathbf{Z}_{-n}))^2 \exp\left(-\boldsymbol{\mu}_m^\top \mathbf{Z}_{-n} + \frac{1}{2} \|\boldsymbol{\mu}_m\|^2\right) \right] \\ &= h_u^2 \int_{\mathbb{R}^{n-1}} \bar{\Phi}(\ell_{n,m}^A(\mathbf{x}))^2 e^{-\boldsymbol{\mu}_m^\top \mathbf{x} + (1/2)\|\boldsymbol{\mu}_m\|^2} \phi_{n-1}(\mathbf{x}) d\mathbf{x} \\ &= h_u^2 e^{\|\boldsymbol{\mu}_m\|^2} \int_{\mathbb{R}^{n-1}} \bar{\Phi}(\ell_{n,m}^A(\mathbf{x}))^2 \phi_{n-1}(\mathbf{x} + \boldsymbol{\mu}_m) d\mathbf{x} \\ &= h_u^2 e^{\|\boldsymbol{\mu}_m\|^2} \int_{\mathbb{R}^{n-1}} \bar{\Phi}(\ell_{n,m}^A(\mathbf{z} - \boldsymbol{\mu}_m))^2 \phi_{n-1}(\mathbf{z}) d\mathbf{z} \\ &\leq h_u^2 e^{\|\boldsymbol{\mu}_m\|^2} \int_{\mathbb{R}} \bar{\Phi}(\Lambda_{nn}^{-1}(\log S_{n,m}^0 - \kappa_n - \boldsymbol{\lambda}_n^\top \boldsymbol{\mu}_m + \|\boldsymbol{\lambda}_n\|y))^2 \phi(y) dy, \end{aligned} \quad (18)$$

where h_u denotes the uniform upper bound of $h(\cdot)$.

Next, we have the following identity:

$$\begin{aligned} \boldsymbol{\lambda}_n^\top \boldsymbol{\mu}_m &= -(\log S_{n,m}^0 - \kappa_n) \boldsymbol{\lambda}_n^\top (\boldsymbol{\lambda}_n \boldsymbol{\lambda}_n^\top + \Lambda_{nn}^2 \mathbf{I})^{-1} \boldsymbol{\lambda}_n \\ &= -\frac{\log S_{n,m}^0 - \kappa_n}{\Lambda_{nn}^2} \boldsymbol{\lambda}_n^\top \left(\mathbf{I} - \frac{1}{\sigma_n^2} \boldsymbol{\lambda}_n \boldsymbol{\lambda}_n^\top \right) \boldsymbol{\lambda}_n \\ &= -\frac{\log S_{n,m}^0 - \kappa_n}{\Lambda_{nn}^2} \left(\boldsymbol{\lambda}_n^\top \boldsymbol{\lambda}_n - \frac{(\boldsymbol{\lambda}_n^\top \boldsymbol{\lambda}_n)^2}{\sigma_n^2} \right) \\ &= -(\log S_{n,m}^0 - \kappa_n) \frac{\boldsymbol{\lambda}_n^\top \boldsymbol{\lambda}_n}{\sigma_n^2} \\ &= -(\log S_{n,m}^0 - \kappa_n) (1 - \Lambda_{nn}^2 / \sigma_n^2), \end{aligned}$$

where the second equality holds by the Sherman-Morrison formula and $\sigma_n^2 = \boldsymbol{\lambda}_n^\top \boldsymbol{\lambda}_n + \Lambda_{nn}^2$, and similarly, it is easy to check that $\|\boldsymbol{\mu}_m\|^2 = (\log S_{n,m}^0 - \kappa_n)^2 (1 - \Lambda_{nn}^2 / \sigma_n^2) / \sigma_n^2$. Therefore, we have

Right-hand side of (18)

$$\begin{aligned} &= h_u^2 \exp \left\{ \frac{(\log S_{n,m}^0 - \kappa_n)^2}{\sigma_n^2} \left(1 - \frac{\Lambda_{nn}^2}{\sigma_n^2} \right) \right\} \int_{\mathbb{R}} \bar{\Phi} \left(\frac{\log S_{n,m}^0 - \kappa_n}{\Lambda_{nn}} \left(2 - \frac{\Lambda_{nn}^2}{\sigma_n^2} \right) + \frac{\|\boldsymbol{\lambda}_n\|}{\Lambda_{nn}} y \right)^2 \phi(y) dy \\ &\sim \frac{h_u^2 (1 + A^2)^2}{2\pi Y_m^2 \sqrt{1 + 2A^2}} \exp \left(-\frac{Y_m^2}{1 + A^2} \right) \quad \text{as } m \rightarrow \infty, \end{aligned}$$

where $Y_m = \Lambda_{nn}^{-1} (\log S_{n,m}^0 - \kappa_n)$, $A = \|\boldsymbol{\lambda}_n\| / \Lambda_{nn}$, and the asymptotic equivalence holds by (16).

Step 3. Analogous to Step 2, we obtain the following asymptotic relationship for the first moment of the BLISS estimator as m grows large:

$$\begin{aligned} \tilde{\mathbb{E}}[\Gamma_m^{\mathbf{A}}] &\geq h_l \mathbb{E} \left[\bar{\Phi}(\ell_{n,m}^{\mathbf{A}}(\mathbf{Z}_{-n,m}^{\mathbf{A}})) \right] \\ &= h_l \int_{\mathbb{R}^{n-1}} \bar{\Phi}(\ell_{n,m}^{\mathbf{A}}(\mathbf{x})) \phi_{n-1}(\mathbf{x}) d\mathbf{x} \\ &\geq h_l \int_{\mathbb{R}} \bar{\Phi}(\Lambda_{nn}^{-1} (\log S_{n,m}^0 - \tau_n + \|\boldsymbol{\lambda}_n\| y)) \phi(y) dy \\ &\sim h_l \frac{\sqrt{1 + A^2}}{\sqrt{2\pi} \tilde{Y}_m} \exp \left(-\frac{\tilde{Y}_m^2}{2 + 2A^2} \right), \end{aligned}$$

where $\tilde{Y}_m = \Lambda_{nn}^{-1} (\log S_{n,m}^0 - \tau_n)$, $A = \|\boldsymbol{\lambda}_n\| / \Lambda_{nn}$, $h_l > 0$ denotes the uniform lower bound of $h(\cdot)$, and the asymptotic equivalence holds by (16).

Accordingly, we conclude that

$$\liminf_{m \rightarrow \infty} \frac{\log \tilde{\mathbb{E}}[(\Gamma_m^{\mathbf{A}})^2]}{2 \log \tilde{\mathbb{E}}[\Gamma_m^{\mathbf{A}}]} \geq \lim_{m \rightarrow \infty} \frac{Y_m^2}{\tilde{Y}_m^2} = 1.$$

Since $\tilde{\mathbb{E}}[(\Gamma_m^{\mathbf{A}})^2] \geq \tilde{\mathbb{E}}[\Gamma_m^{\mathbf{A}}]^2$, the desired result follows. \square

Proof of Proposition 2. Define Y_m , \tilde{Y}_m and A as in the proof of Theorem 1. According to the proof of Theorem 1, we have

$$\tilde{\mathbb{E}}[(\Gamma_m^{\mathbf{A}})^2] \lesssim \frac{h_u^2 (1 + A^2)^2}{2\pi Y_m^2 \sqrt{1 + 2A^2}} \exp \left(-\frac{Y_m^2}{1 + A^2} \right) \quad \text{as } m \rightarrow \infty.$$

Using a similar argument, we also observe that as $m \rightarrow \infty$,

$$\tilde{\mathbb{E}}[(\Gamma_m^{\mathbf{A}})^2] \gtrsim \frac{h_l^2 (1 + A^2)^2}{2\pi \tilde{Y}_m^2 \sqrt{1 + 2A^2}} \exp \left(-\frac{\tilde{Y}_m^2}{1 + A^2} \right),$$

which results in $\log \tilde{\mathbb{E}}[(\Gamma_m^A)^2] \sim -Y_m^2/(1+A^2)$ since $Y_m \sim \tilde{Y}_m$.

Similarly, we obtain

$$\frac{h_l^2(1+2A^2)^{3/2}}{2\pi\tilde{Y}_m^2} \exp\left(-\frac{\tilde{Y}_m^2}{1+2A^2}\right) \lesssim \mathbb{E}[(\Gamma_m^0)^2] \lesssim \frac{h_u^2(1+2A^2)^{3/2}}{2\pi Y_m^2} \exp\left(-\frac{Y_m^2}{1+2A^2}\right),$$

and thus, $\mathbb{E}[(\Gamma_m^0)^2] \sim -Y_m^2/(1+2A^2)$ as $m \rightarrow \infty$. Consequently, we get

$$\lim_{m \rightarrow \infty} \frac{\log \tilde{\mathbb{E}}[(\Gamma_m^A)^2]}{\log \mathbb{E}[(\Gamma_m^0)^2]} = \frac{1+2A^2}{1+A^2} = 2 - \frac{\Lambda_{nn}^2}{\sigma_n^2}.$$

Hence, the proof is complete. \square

Proof of Theorem 2. We first obtain the following relationship of the first moment of the BLISS estimator for all m sufficiently large:

$$\begin{aligned} \tilde{\mathbb{E}}[\Gamma_m^B] &\geq h_l \mathbb{E}\left[\bar{\Phi}(\ell_{n,m}^B(\mathbf{Z}_{-n,m}))\right] \\ &= h_l \int_{\mathbb{R}^{n-1}} \bar{\Phi}(\ell_{n,m}^B(\mathbf{z})) \phi_{n-1}(\mathbf{z}) d\mathbf{z} \\ &= h_l m^{n-1} \int_{\mathbb{R}^{n-1}} \bar{\Phi}(\ell_{n,m}^B(m\mathbf{x})) \phi_{n-1}(m\mathbf{x}) d\mathbf{x} \\ &\geq h_l m^{n-1} \int_{\mathbb{R}^{n-1}} \frac{\ell_{n,m}^B(m\mathbf{x})}{\ell_{n,m}^B(m\mathbf{x})^2 + 1} \phi(\ell_{n,m}^B(m\mathbf{x})) \phi_{n-1}(m\mathbf{x}) d\mathbf{x} \end{aligned}$$

where $h_l > 0$ denotes the uniform lower bound of $h(\cdot)$, and the last inequality stems from the fact that $\bar{\Phi}(t) \geq t\phi(t)/(t^2+1)$ for all t . By a straightforward application of the Laplace method (Wong, 2001), the following holds for some constants $B > 0$:

$$\int_{\mathbb{R}^{n-1}} \frac{\ell_{n,m}^B(m\mathbf{x})}{\ell_{n,m}^B(m\mathbf{x})^2 + 1} \phi(\ell_{n,m}^B(m\mathbf{x})) \phi_{n-1}(m\mathbf{x}) d\mathbf{x} \sim \frac{B \ell_{n,m}^B(m\mathbf{x}^*) e^{-\ell_{n,m}^B(m\mathbf{x}^*)^2/2 - \|\mathbf{m}\mathbf{x}^*\|^2/2}}{(\ell_{n,m}^B(m\mathbf{x}^*)^2 + 1) m^{(n-1)/2}} \text{ as } m \rightarrow \infty,$$

where $\mathbf{x}^* = \operatorname{argmin}_{\mathbf{x}} \{\ell_n^B(\mathbf{x})^2 + \|\mathbf{x}\|^2\}$ and $\ell_n^B(\mathbf{x}) = \lim_{m \rightarrow \infty} (1/m) \ell_{n,m}^B(m\mathbf{x})$.

By writing $\boldsymbol{\mu}_m = \boldsymbol{\mu}_{-n,m}^B$ for brevity, we next find an upper bound of the second moment of the estimator for all m large enough:

$$\begin{aligned} \tilde{\mathbb{E}}[(\Gamma_m^B)^2] &\leq h_u^2 \mathbb{E}\left[\bar{\Phi}(\ell_{n,m}^B(\mathbf{Z}_{-n}))^2 \exp\left(-\boldsymbol{\mu}_m^\top \mathbf{Z}_{-n} + \frac{1}{2} \|\boldsymbol{\mu}_m\|^2\right)\right] \\ &= h_u^2 \int_{\mathbb{R}^{n-1}} \bar{\Phi}(\ell_{n,m}^B(\mathbf{z}))^2 e^{-\boldsymbol{\mu}_m^\top \mathbf{z} + (1/2) \|\boldsymbol{\mu}_m\|^2} \phi_{n-1}(\mathbf{z}) d\mathbf{z} \\ &= h_u^2 e^{\|\boldsymbol{\mu}_m\|^2} \int_{\mathbb{R}^{n-1}} \bar{\Phi}(\ell_{n,m}^B(\mathbf{z}))^2 \phi_{n-1}(\mathbf{z} + \boldsymbol{\mu}_m) d\mathbf{z} \\ &= h_u^2 e^{\|\boldsymbol{\mu}_m\|^2} m^{n-1} \int_{\mathbb{R}^{n-1}} \bar{\Phi}(\ell_{n,m}^B(m\mathbf{x}))^2 \phi_{n-1}(m\mathbf{x} + \boldsymbol{\mu}_m) d\mathbf{x} \end{aligned}$$

$$\leq h_u^2 e^{\|\boldsymbol{\mu}_m\|^2} m^{n-1} \int_{\mathbb{R}^{n-1}} \frac{\phi(\ell_{n,m}^{\mathbb{B}}(m\mathbf{x}))^2}{\ell_{n,m}^{\mathbb{B}}(m\mathbf{x})^2} \phi_{n-1}(m\mathbf{x} + \boldsymbol{\mu}_m) d\mathbf{x}$$

where h_u represents the uniform upper bound of $h(\cdot)$ and the second inequality holds because $\bar{\Phi}(t) \leq \phi(t)/t$ for all $t > 0$. Since $\boldsymbol{\mu}_m = m\mathbf{x}^*$, the Laplace method (Wong, 2001) implies that

$$e^{\|\boldsymbol{\mu}_m\|^2} \int_{\mathbb{R}^{n-1}} \frac{\phi(\ell_{n,m}^{\mathbb{B}}(m\mathbf{x}))^2}{\ell_{n,m}^{\mathbb{B}}(m\mathbf{x})^2} \phi_{n-1}(m\mathbf{x} + \boldsymbol{\mu}_m) d\mathbf{x} \sim \frac{B e^{-\ell_{n,m}^{\mathbb{B}}(m\mathbf{x}^*)^2 - \|m\mathbf{x}^*\|^2}}{\ell_{n,m}^{\mathbb{B}}(m\mathbf{x}^*)^2 m^{(n-1)/2}} \text{ as } m \rightarrow \infty.$$

Consequently, we have

$$\liminf_{m \rightarrow \infty} \frac{\log \tilde{\mathbb{E}}[(\Gamma_m^{\mathbb{B}})^2]}{2 \log \tilde{\mathbb{E}}[\Gamma_m^{\mathbb{B}}]} \geq \lim_{m \rightarrow \infty} \frac{\log(m^{(n-1)/2} e^{-\ell_{n,m}^{\mathbb{B}}(m\mathbf{x}^*)^2 - \|m\mathbf{x}^*\|^2})}{\log(m^{n-1} \ell_{n,m}^{\mathbb{B}}(m\mathbf{x}^*)^2 e^{-\ell_{n,m}^{\mathbb{B}}(m\mathbf{x}^*)^2 - \|m\mathbf{x}^*\|^2})} = 1.$$

Since $\tilde{\mathbb{E}}[(\Gamma_m^{\mathbb{B}})^2] \geq \tilde{\mathbb{E}}[\Gamma_m^{\mathbb{B}}]^2$, the result follows. \square

Proof of Proposition 3. Following the same approach as in the proof of Theorem 2, we observe that

$$\begin{aligned} \log \tilde{\mathbb{E}}[(\Gamma_m^0)^2] &\leq \log \left(h_u^2 m^{n-1} \int_{\mathbb{R}^{n-1}} \frac{\phi(\ell_{n,m}^{\mathbb{B}}(m\mathbf{x}))^2}{\ell_{n,m}^{\mathbb{B}}(m\mathbf{x})^2} \phi_{n-1}(m\mathbf{x}) d\mathbf{x} \right) \sim -m^2 \left(\ell_n^{\mathbb{B}}(\bar{\mathbf{x}})^2 + \frac{\|\bar{\mathbf{x}}\|^2}{2} \right); \\ \log \tilde{\mathbb{E}}[(\Gamma_m^0)^2] &\geq \log \left(h_l^2 m^{n-1} \int_{\mathbb{R}^{n-1}} \frac{\ell_{n,m}^{\mathbb{B}}(m\mathbf{x})^2 \phi(\ell_{n,m}^{\mathbb{B}}(m\mathbf{x}))^2}{(\ell_{n,m}^{\mathbb{B}}(m\mathbf{x})^2 + 1)^2} \phi_{n-1}(m\mathbf{x}) d\mathbf{x} \right) \sim -m^2 \left(\ell_n^{\mathbb{B}}(\bar{\mathbf{x}})^2 + \frac{\|\bar{\mathbf{x}}\|^2}{2} \right) \end{aligned}$$

as $m \rightarrow \infty$, where $\bar{\mathbf{x}} = \operatorname{argmin}_{\mathbf{x}} \{ \ell_n^{\mathbb{B}}(\mathbf{x})^2 + \|\mathbf{x}\|^2/2 \}$. Similarly, we obtain

$$\begin{aligned} \log \tilde{\mathbb{E}}[(\Gamma_m^{\mathbb{B}})^2] &\geq \log \left(h_l^2 e^{\|m\mathbf{x}^*\|^2} m^{n-1} \int_{\mathbb{R}^{n-1}} \frac{\ell_{n,m}^{\mathbb{B}}(m\mathbf{x})^2 \phi(\ell_{n,m}^{\mathbb{B}}(m\mathbf{x}))^2}{(\ell_{n,m}^{\mathbb{B}}(m\mathbf{x})^2 + 1)^2} \phi_{n-1}(m\mathbf{x} + m\mathbf{x}^*) d\mathbf{x} \right) \\ &\sim -\ell_{n,m}^{\mathbb{B}}(m\mathbf{x}^*)^2 - \|m\mathbf{x}^*\|^2 \\ &\sim -m^2 \left(\ell_n^{\mathbb{B}}(\mathbf{x}^*)^2 + \|\mathbf{x}^*\|^2 \right) \end{aligned}$$

as $m \rightarrow \infty$. Therefore, by letting $\rho_n = \log(S_n^0/v_n(\mathbf{S}_{-n}^0)) > 0$, we get

$$\begin{aligned} \lim_{m \rightarrow \infty} \frac{\log \mathbb{E}[(\Gamma_m^{\mathbb{B}})^2]}{\log \mathbb{E}[(\Gamma_m^0)^2]} &= \frac{\ell_n^{\mathbb{B}}(\mathbf{x}^*)^2 + \|\mathbf{x}^*\|^2}{\ell_n^{\mathbb{B}}(\bar{\mathbf{x}})^2 + \|\bar{\mathbf{x}}\|^2/2} \\ &\geq 1 + \frac{\|\mathbf{x}^*\|^2}{2\ell_n^{\mathbb{B}}(\mathbf{x}^*)^2 + \|\mathbf{x}^*\|^2} \\ &\geq 1 + \frac{\Lambda_{nn}^2 \|\mathbf{x}^*\|^2}{2(\rho_n + \boldsymbol{\lambda}_n^\top \mathbf{x}^*)^2 + \Lambda_{nn}^2 \|\mathbf{x}^*\|^2}. \end{aligned} \tag{19}$$

We further see that the second term in (19) is strictly positive if and only if $\boldsymbol{\lambda}_n = \mathbf{0}$. This completes the proof. \square

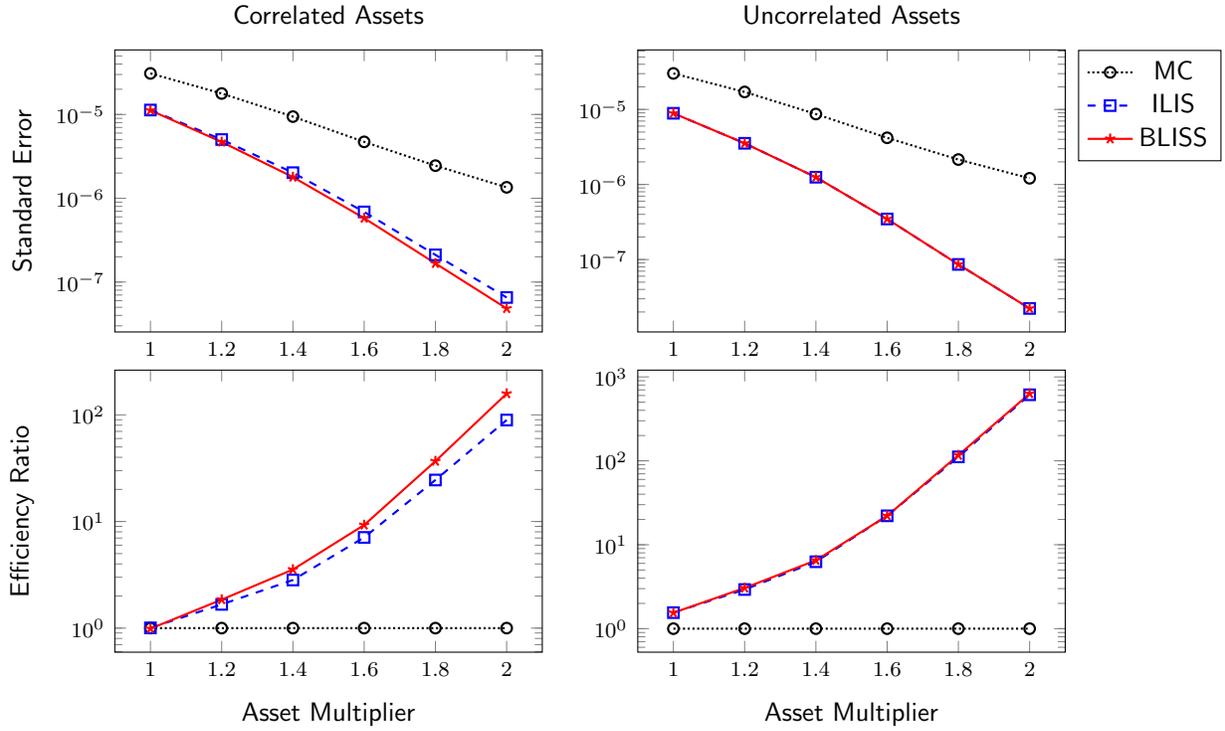


Figure 6: Performance comparison between MC, ILIS, and BLISS for the real-world example in Section 6.2 with the complete network when asset value changes

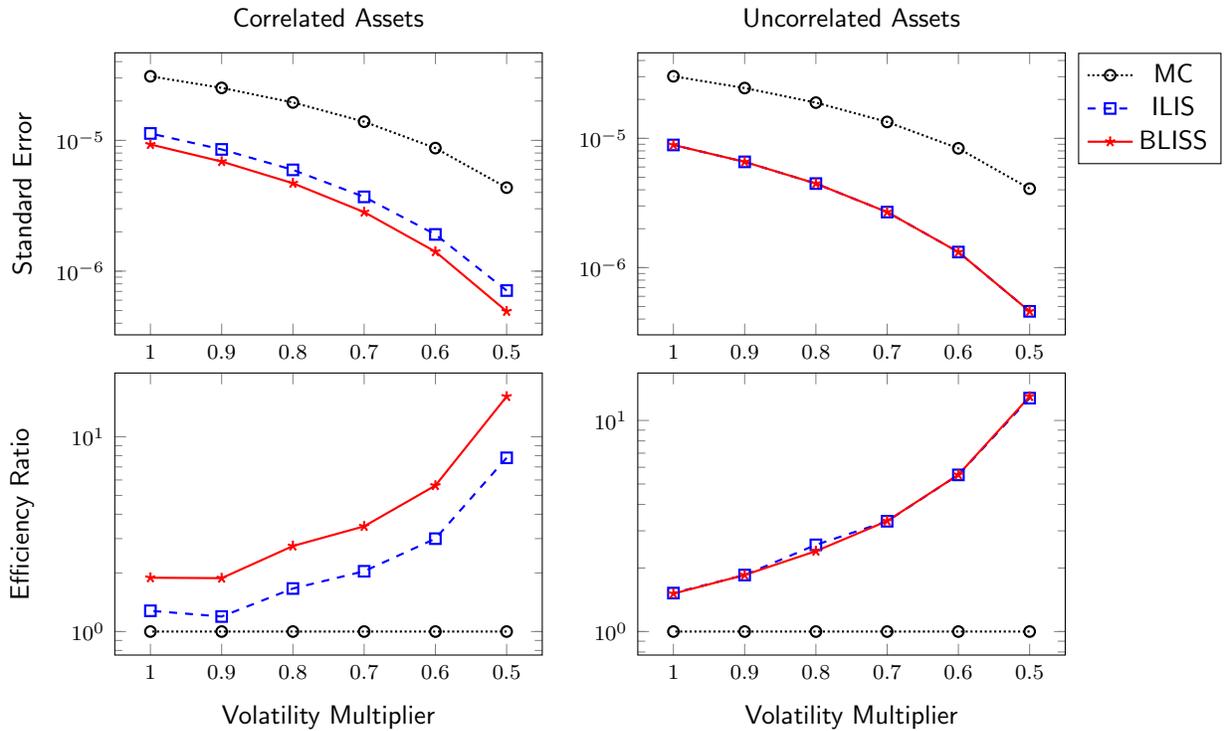


Figure 7: Performance comparison between MC, ILIS, and BLISS for the real-world example in Section 6.2 with the complete network when volatility changes

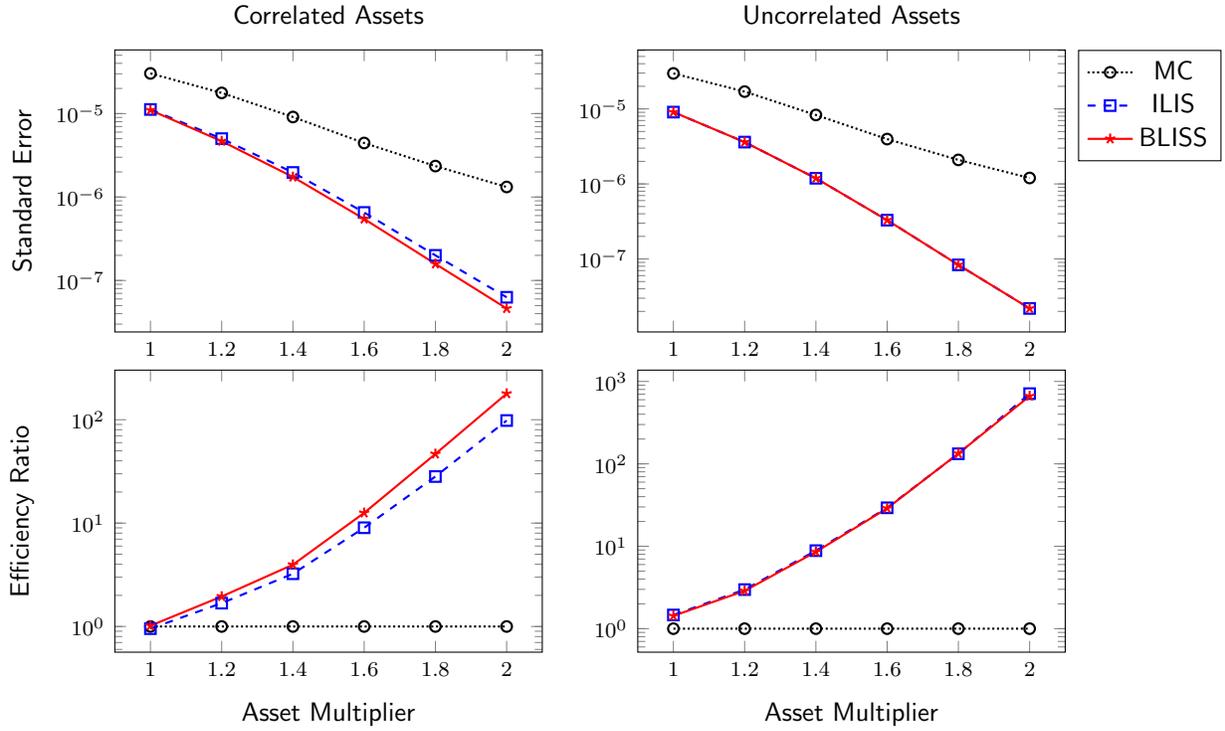


Figure 8: Performance comparison between MC, ILIS, and BLISS for the real-world example in Section 6.2 with the ring network when asset value changes

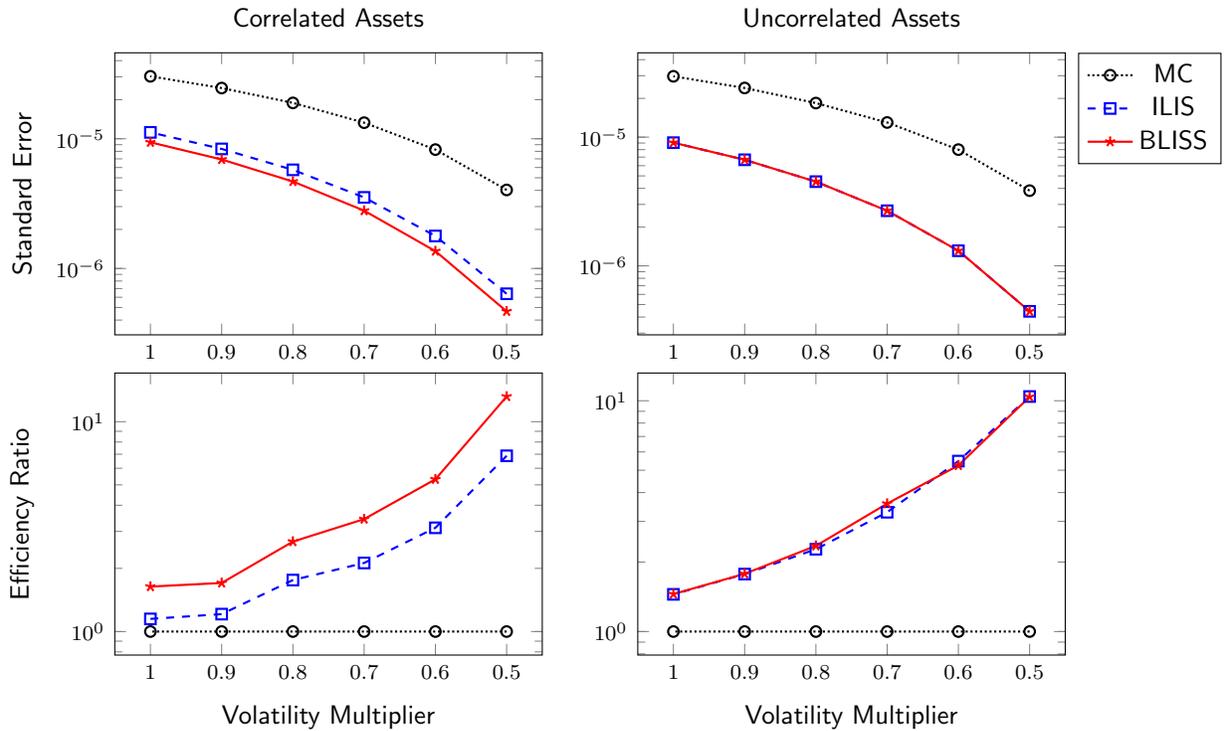


Figure 9: Performance comparison between MC, ILIS, and BLISS for the real-world example in Section 6.2 with the ring network when volatility changes

C Additional Theoretical Result

To reflect crisis scenarios characterized by soaring asset volatility, we further evaluate the performance of the BLISS estimator under a large volatility regime. In particular, we scale the external asset volatility by replacing $\mathbf{\Lambda}$ with $m\mathbf{\Lambda}$, where larger values of m represent periods of greater market stress. Consequently, we replace the function $\ell_n(\cdot)$ in (8) with

$$\ell_{n,m}^{\mathbf{c}}(\mathbf{x}) := \frac{1}{m\Lambda_{nn}} \left(\log S_n^0 - \frac{m^2\sigma_n^2}{2} - \log(v_n(s_m^{\mathbf{c}}(\mathbf{x}))) + m \sum_{k=1}^{n-1} \Lambda_{nk} x_k \right)$$

where $s_m^{\mathbf{c}}(\mathbf{x})$ is an $(n-1)$ -dimensional vector whose i -th element is $S_i^0 \exp(-m^2\sigma_i^2/2 + m \sum_{k=1}^i \Lambda_{ik} x_k)$.

In this regime, external asset values decrease with m , and thus, this regime can also be viewed as a small asset regime. Consequently, default is no longer a rare event, making the last term in (3) easy to estimate even without the BLISS estimator. Conversely, the complementary probability $\mathbb{P}(p_n(\mathbf{S}) = \bar{p}_n)$, represented by the first two terms in (3), now becomes a rare-event probability. We therefore focus on estimating this new rare-event probability using a variant of the BLISS estimator, constructed as follows:

$$\Gamma_m^{\mathbf{c}} := \Phi(\ell_{n,m}^{\mathbf{c}}(\mathbf{Z}_{-n,m}^{\mathbf{c}})) \mathcal{L}_{\text{out}}(\mathbf{Z}_{-n,m}^{\mathbf{c}}), \quad (20)$$

where the samples of $\mathbf{Z}_{-n,m}^{\mathbf{c}} := (Z_{1,m}^{\mathbf{c}}, \dots, Z_{n-1,m}^{\mathbf{c}}) \sim \mathcal{N}(\boldsymbol{\mu}_{-n,m}^{\mathbf{c}}, \mathbf{I}_{n-1})$ are drawn via outer-layer importance sampling. This construction is based on Proposition 1, which shows that, conditional on $\mathbf{Z}_{-n,m}^{\mathbf{c}}$, bank n 's solvency event $\{p_n(\mathbf{S}) = \bar{p}_n\}$ is equivalent to the event $\{Z_n > -\ell_{n,m}^{\mathbf{c}}(\mathbf{Z}_{-n,m}^{\mathbf{c}})\}$. The probability of this equivalent event is given by $\Phi(\ell_{n,m}^{\mathbf{c}}(\mathbf{Z}_{-n,m}^{\mathbf{c}}))$ and forms the first term in the estimator (20). Here, the shifted mean vector $\boldsymbol{\mu}_{-n,m}^{\mathbf{c}}$ for the outer-layer importance sampling is chosen in a way similar to the argument in Section 5.1 as

$$\boldsymbol{\mu}_{-n,m}^{\mathbf{c}} = -\frac{m\sigma_n^2}{2} (\boldsymbol{\lambda}_n \boldsymbol{\lambda}_n^{\top} + \Lambda_{nn}^2 \mathbf{I})^{-1} \boldsymbol{\lambda}_n.$$

The asymptotic optimality of this new BLISS estimator is established in the following result:

Corollary 1 (Asymptotic Optimality of $\Gamma_m^{\mathbf{c}}$) *Suppose that Assumptions 1 and 2 hold. Then, the BLISS estimator $\Gamma_m^{\mathbf{c}}$ for estimating the solvency probability $\mathbb{P}(p_n(\mathbf{S}) = \bar{p}_n)$ is asymptotically optimal in the large volatility regime.*

Proof. The proof follows the same approach as that of Theorem 1 and is therefore omitted. \square

References

Ahn, D. (2020). Shock amplification in financial networks with applications to the CCP feasibility. *Quantitative Finance*, 20(7):1045–1056.

- Ahn, D., Chen, N., and Kim, K.-K. (2024). Robust risk quantification via shock propagation in financial networks. *Operations Research*, 72(1):1–18.
- Ahn, D. and Kim, K.-K. (2018). Efficient simulation for expectations over the union of half-spaces. *ACM Transactions on Modeling and Computer Simulation*, 28(3):Article 23.
- Ahn, D. and Kim, K.-K. (2019). Optimal intervention under stress scenarios: A case of the Korean financial system. *Operations Research Letters*, 47(4):257–263.
- Ahn, D., Kim, K.-K., and Kwon, E. (2023). Multivariate stress scenario selection in interbank networks. *Journal of Economic Dynamics and Control*, 154:104712.
- Ahn, D. and Zheng, L. (2021). Efficient simulation for linear programming under uncertainty. In Kim, S., Feng, B., Smith, K., Masoud, S., Zheng, Z., Szabo, C., and Loper, M., editors, *2021 Winter Simulation Conference (WSC)*, pages 1–12. IEEE.
- Ahn, D. and Zheng, L. (2023). Conditional importance sampling for convex rare-event sets. In Corlu, C. G., Hunter, S. R., Lam, H., Onggo, B. S., Shortle, J., and Biller, B., editors, *Proceedings of the 2023 Winter Simulation Conference*. IEEE.
- Ahn, D. and Zheng, L. (2025). Efficient simulation of polyhedral expectations with applications to finance. *Mathematics of Operations Research*. Articles in Advance.
- Amini, H., Filipović, D., and Minca, A. (2016a). To fully net or not to net: Adverse effects of partial multilateral netting. *Operations Research*, 64(5):1135–1142.
- Amini, H., Filipović, D., and Minca, A. (2016b). Uniqueness of equilibrium in a payment system with liquidation costs. *Operations Research Letters*, 44(1):1–5.
- Ararat, Ç. and Meimanjan, N. (2023). Computation of systemic risk measures: A mixed-integer programming approach. *Operations Research*, 71(6):2130–2145.
- Asmussen, S. and Glynn, P. W. (2007). *Stochastic Simulation: Algorithms and Analysis*. Springer.
- Barucca, P., Bardoscia, M., Caccioli, F., D’Errico, M., Visentin, G., Caldarelli, G., and Battiston, S. (2020). Network valuation in financial systems. *Mathematical Finance*, 30(4):1181–1204.
- Bassamboo, A., Juneja, S., and Zeevi, A. (2008). Portfolio credit risk with extremal dependence: Asymptotic analysis and efficient simulation. *Operations Research*, 56(3):593–606.
- Bernard, B., Capponi, A., and Stiglitz, J. E. (2022). Bail-ins and bailouts: Incentives, connectivity, and systemic stability. *Journal of Political Economy*, 130(7):1805–1859.
- Blanchet, J., Li, J., and Nakayama, M. K. (2019). Rare-event simulation for distribution networks. *Operations Research*, 67(5):1383–1396.
- Bucklew, J. A. (2004). *Introduction to Rare Event Simulation*. Springer, New York, NY.

- Capponi, A. (2016). Systemic risk, policies, and data needs. *INFORMS TutORials in Operations Research*, pages 185–206.
- Capponi, A. and Chen, P.-C. (2015). Systemic risk mitigation in financial networks. *Journal of Economic Dynamics and Control*, 58:152–166.
- Capponi, A., Chen, P.-C., and Yao, D. D. (2016). Liability concentration and systemic losses in financial networks. *Operations Research*, 64(5):1121–1134.
- Capponi, A. and Larsson, M. (2015). Price contagion through balance sheet linkages. *The Review of Asset Pricing Studies*, 5(2):227–253.
- Chen, N., Liu, X., and Yao, D. D. (2016). An optimization view of financial systemic risk modeling: The network effect and the market liquidity effect. *Operations Research*, 64(5):1089–1108.
- Cifuentes, R., Ferrucci, G., and Shin, H. S. (2005). Liquidity risk and contagion. *Journal of the European Economic Association*, 3(2/3):556–566.
- Clerc, L., Giovannini, A., Langfield, S., Peltonen, T., Portes, R., and Scheicher, M. (2016). Indirect contagion: The policy problem. *ESRB: Occasional Paper Series*, (2016/09).
- Cossin, D. and Schellhorn, H. (2007). Credit risk in a network economy. *Management Science*, 53(10):1604–1617.
- Craig, B. and von Peter, G. (2014). Interbank tiering and money center banks. *Journal of Financial Intermediation*, 23(3):322–347.
- Duarte, F. and Eisenbach, T. M. (2021). Fire-sale spillovers and systemic risk. *The Journal of Finance*, 76(3):1251–1294.
- ECB Banking Supervision (2018). Supervisory banking statistics, fourth quarter 2017. https://www.bankingsupervision.europa.eu/ecb/pub/pdf/ssm.supervisorybankingstatistics_fourth_quarter_2017_201804.en.pdf.
- Eisenberg, L. and Noe, T. H. (2001). Systemic risk in financial systems. *Management Science*, 47(2):236–249.
- Elsinger, H., Lehar, A., and Summer, M. (2006). Risk assessment for banking systems. *Management Science*, 52(9):1301–1314.
- European Banking Authority (2018). EU-wide stress testing 2018. <https://www.eba.europa.eu/risk-and-data-analysis/risk-analysis/eu-wide-stress-testing/eu-wide-stress-testing-2018>.
- Feinstein, Z., Pang, W., Rudloff, B., Schaanning, E., Sturm, S., and Wildman, M. (2018). Sensitivity of the Eisenberg-Noe clearing vector to individual interbank liabilities. *SIAM Journal on Financial Mathematics*, 9(4):1286–1325.

- Gai, P., Haldane, A. G., and Kapadia, S. (2011). Complexity, concentration and contagion. *Journal of Monetary Economics*, 58(5):453–470.
- Gandy, A. and Veraart, L. A. M. (2017). A Bayesian methodology for systemic risk assessment in financial networks. *Management Science*, 63(12):4428–4446.
- Ghamami, S., Glasserman, P., and Young, H. P. (2022). Collateralized networks. *Management Science*, 68(3):2202–2225.
- Glasserman, P. (2003). *Monte Carlo Methods in Financial Engineering*. Springer, New York.
- Glasserman, P. and de Larrea, E. L. (2023). Maximum entropy distributions with applications to graph simulation. *Operations Research*, 71(5):1908–1924.
- Glasserman, P., Kang, W., and Shahabuddin, P. (2008). Fast simulation of multifactor portfolio credit risk. *Operations Research*, 56(5):1200–1217.
- Glasserman, P. and Li, J. (2005). Importance sampling for portfolio credit risk. *Management Science*, 51(11):1643–1656.
- Glasserman, P. and Young, H. P. (2015). How likely is contagion in financial networks? *Journal of Banking & Finance*, 50:383–399.
- Glasserman, P. and Young, H. P. (2016). Contagion in financial networks. *Journal of Economic Literature*, 54(3):779–831.
- Gouriéroux, C., Héam, J.-C., and Monfort, A. (2012). Bilateral exposures and systemic solvency risk. *Canadian Journal of Economics/Revue canadienne d'économie*, 45(4):1273–1309.
- Greenwood, R., Landier, A., and Thesmar, D. (2015). Vulnerable banks. *Journal of Financial Economics*, 115(3):471–485.
- Hu, F., Mitchell, D., and Tompaidis, S. (2024). Robust financial networks. *Operations Research*, 72(5):1827–1842.
- Kang, W. and Shahabuddin, P. (2005). Fast simulation for multifactor portfolio credit risk in the t -copula model. In *Proceedings of the 2005 Winter Simulation Conference*, pages 1859–1868.
- Kusnetsov, M. and Veraart, L. A. M. (2019). Interbank clearing in financial networks with multiple maturities. *SIAM Journal on Financial Mathematics*, 10(1):37–67.
- Liu, M. and Staum, J. (2010). Sensitivity analysis of the Eisenberg-Noe model of contagion. *Operations Research Letters*, 38(5):489–491.
- Merton, R. C. (1974). On the pricing of corporate debt: The risk structure of interest rates. *The Journal of Finance*, 29(2):449–470.

- Rogers, L. C. G. and Veraart, L. A. M. (2013). Failure and rescue in an interbank network. *Management Science*, 59(4):882–898.
- Sion, M. (1958). On general minimax theorems. *Pacific Journal of Mathematics*, 8(1):171–176.
- Upper, C. and Worms, A. (2004). Estimating bilateral exposures in the German interbank market: Is there a danger of contagion? *European Economic Review*, 48(4):827–849.
- Veraart, L. A. M. (2025). Modelling contagious bank runs. *Available at SSRN*. <https://ssrn.com/abstract=5176815>.
- Weber, S. and Weske, K. (2017). The joint impact of bankruptcy costs, fire sales and cross-holdings on systemic risk in financial networks. *Probability, Uncertainty and Quantitative Risk*, 2(9):1–38.
- Wong, R. (2001). *Asymptotic Approximations of Integrals*. Society for Industrial and Applied Mathematics, Philadelphia, PA.