Systemic Influences on Optimal Equity-Credit Investment

Agostino Capponi
Columbia University
New York NY 10027
USA
ac3827@columbia.edu

Christoph Frei
University of Alberta
Edmonton AB T6G 2G1
Canada
cfrei@ualberta.ca

This version: January 7, 2016
(Management Science 63 (2017), 2756–2771)

Abstract

We introduce an equity-credit portfolio framework taking into account the structural interaction of market and credit risk, along with their systemic dependencies. We derive an explicit expression for the optimal investment strategy in stocks and credit default swaps (CDSs). We exploit its representation structure and analyze the mechanisms driving the optimal investment decisions. The transmission of market risk premia is the key mechanism through which systemic influences affect the optimal investment strategy. We develop a novel calibration procedure and find that systemic dependencies are statistically significant when the model is fitted to historical time series of equity and CDS data. An empirical analysis with data of companies in the DJIA reveals the critical role of systemic risk in portfolio monitoring.

1 Introduction

The dynamic equity portfolio selection problem, originating from the seminal work of Merton (1969), has been subject of considerable investigation. Significant developments include the impact of trading constraints (see Cvitanić (2001) for a survey), the assessment of robustness against model misspecification as in Anderson et al. (2000) (see also Hansen et al. (2006) for a survey), and the effect of transaction costs as in the early work of Perold (1984) who consider a broad class of portfolio models, and the subsequent study of Liu (2007) for a stock portfolio.

In recent years, credit default swaps (CDSs) have been increasingly used by investors to trade on the credit quality of corporations or sovereign entities. This is because they are easier and quicker to trade than the risky bonds they are tied to.[1] For instance, an investor can implement a short credit position by buying protection through a credit swap. The investor pays a running spread premium and receives compensation for the loss at default. This is preferred to shorting a bond via a reverse repo transaction and sale. Despite the growing importance of portfolio credit derivatives strategies, to the best of our knowledge most of the studies have so far focused on risky bond portfolio problems. These include the works of of Wise and Bhansali (2002) who consider a structural default model, and of Kraft and Steffensen (2009) who account for default contagion and bankruptcy procedures. There have only been very few studies focusing on optimal credit swap allocation. These include Giesecke et al. (2014), who solve the static selection problem for a portfolio of credit swaps taking solvency, capital, and other trading constraints into account, and Bo and Capponi (2014), who consider a portfolio framework where a risk-averse investor can dynamically select his optimal CDS investment strategy.

The above surveyed literature has considered either stock or fixed-income securities, but never the two simultaneously. This stands in contrast with the behavior of leading market players, who often implement mixed equity-credit strategies.[2] Moreover, the recent financial crisis has shown that these two types of risk can reinforce each other and underlined the need of an integrated risk management framework.

Our study is the first to analyze the mechanisms through which the economic-wide interaction of market and credit risk affects the optimal mixed portfolio strategy of a rational investor. Previous studies have mainly focused on joint valuation of credit and equity derivatives. Mendoza et al. (2010) use time-changed processes to simultaneously price equity and credit derivatives of a given firm. Linetsky (2006) solves in closed form an extended Black-Scholes framework inclusive of bankruptcy. Mendoza and Linetsky (2014) develop a multi-name credit-equity model to compute no-arbitrage prices of contingent claims written on stocks using multivariate subordination.

Our model is rich enough to incorporate the most significant factors affecting default risk of a company identified by empirical research. These include (1) the performance of market, econometric and balance sheet variables, and (2) contagion effects manifested when other companies default or experience other severe financial distress. The dependence of the default intensity of a company on market information

[1] As reported by Bloomberg (2013), the world's biggest mutual fund PIMCO wagered over ten billion dollars in the CDS market in the year 2013.

[2] Chart 3 in the Citigroup perspective report released in July 2010 (see http://www.citibank.com/icg/global_markets/prime_finance/docs/CitiPerspectives_LiquiditySurvey_July2010.pdf for details) illustrates the breadth of its investment strategies which range from equity-focused to mixed to credit-focused.

such as the company's own stock price or a stock market index captures the interaction between market and credit risk. Macroeconomic indicators, such as interest rates or the gross domestic product, affect the credit quality of a portfolio; for instance, a rise in interest rates increases the costs of corporate financing, decreasing the market value of assets, and hence leading to a deterioration in the credit quality of the portfolio. The influence of other defaults on the credit risk of a specific company captures the *systemic* dependencies. This leads to two main channels of dependence between the price processes of the stocks in the portfolio. The first is via market risk, as in classical multidimensional default-free stock market models. The second is via systemic dependencies, as the default risk of each stock depends on the values and default states of other stocks in the portfolio.

We next list our two main contributions. The first contribution is to derive explicit representations for the optimal investment strategies. These depend on the underlying model parameters and on the sensitivity of CDS prices to price changes of the stocks. Despite the complex dependence structure of market and credit risk factors, we are able to recast the optimization problem as a linear system of equations, whose solution is the vector of optimal strategies. The dynamic interaction of market and credit risk leads to a structural form for the optimal mixed strategy which substantially differs from the case when equity and credit markets are treated separately. For instance, it is well known that logarithmic investors are myopic in the equity market; see, for instance, Section 2.3.3 in Brandt (2010). This is no longer the case in our model. Our optimal portfolio allocations depend in a nontrivial manner on the sensitivities of CDS prices to market risk (change in portfolio stock prices) and to credit risk (default of a portfolio name). This means that the optimal strategy has a crucial forward-looking component coming from these price sensitivities, so that it cannot be only estimated on the basis of historical data, such as expected returns, variances, and covariances of the available securities. As a result, portfolio managers interested in fixed-income investments need to consider a different estimation procedure than that used in the context of equity investment. In addition to current prices, they need to consider price sensitivities of fixed-income securities when computing the optimal investment strategies. The computation of these price sensitivities requires simulating the future paths of default intensities via efficient Monte-Carlo methods, taking into account systemic effects. We identify two key mechanisms driving the investment decisions. The equity and default risk premia are first subject to an amplification depending on the risk structure of the portfolio universe. The aggregate portfolio amplification is then transmitted via the sensitivity matrix to the investment strategy in the CDS of any individual company, and referred to as its systemic influence.

The second main contribution is, to the best of our knowledge, the first calibration algorithm of a hybrid equity-credit model accounting for systemic dependencies. The calibration is based on an iterative procedure which initially assumes the absence of systemic risk and then refines the model by incorporating the spreading consequences of defaults. Our iterative approach presents some similarities with the fictitious default algorithm by Eisenberg and Noe (2001), which recovers the unique clearing payment vector in an interbanking system starting from the assumption that all banks are sound and then progressively eliminating banks that become insolvent as a result of the failures of others. Using empirical data and the calibrated model, we show that systemic dependencies are significant for most companies in the Dow Jones Industrial Average 30 (DJIA). Tracking 30 large US companies, the DJIA reflects a broad overview of the economy and its different industries. Indeed, the DJIA is one of the oldest and most widely recognized stock market indices. Moreover, the investment grade companies in the DJIA are among the most liquid based on CDS volume traded. We find that the systemic influences of a small number of financial firms drive a large fraction of the cross-sectional demand for CDS securities, even if their credit spreads are not necessarily among the highest. Perhaps surprisingly, the demands which are most sensitive to systemic influences are for CDSs referencing industrial, as opposed to financial, firms. These industrial firms have large financial subsidiaries and their businesses are capital intensive, hence they are highly dependent on broad financial market conditions.

Our analysis indicates that systemic influences should be considered when monitoring mixed equity-credit portfolios. The investor needs to rebalance frequently his entire portfolio when there are changes in credit quality of systemically influential companies. Moreover, after identifying which positions are more sensitive to systemic influences, he would need to take suitable precautionary measures to deal with high loading/unloading of CDS positions.

The rest of the paper is organized as follows. Section 2 introduces the portfolio securities. Section 3 formulates the dynamic portfolio selection problem. Section 4 gives the explicit expression for the optimal strategy and analyze systemic influences. Section 5 performs an empirical analysis. Section 6 concludes the paper. Proofs of all technical results are delegated to an e-companion to this paper.

2 The Portfolio Securities

We consider a model of a financial market consisting of a risk-free money market account, M stocks and M credit default swaps (CDS) referencing the same companies as the stocks. The interest rate $r \geq 0$ is assumed to be constant. Throughout the paper, we fix a probability space where $\mathbb{P}$ denotes the subjective probability measure of the investor, reflecting his own views on market and credit risk of the economy. Section 2.1 describes the stock price process. Section 2.2 introduces the CDS price process.

2.1 Defaultable Stock Prices

The price at time t of the stock of company i is denoted by $S_i(t)$. In the classical Black-Scholes model, the dynamics of $S_i(t)$ is modeled using a geometric Brownian motion $\mathrm{d}S_i(t) = S_i(t)\big(\mu_i \,\mathrm{d}t + \sigma_i \,\mathrm{d}W_i(t)\big)$ with price $S_i(0) > 0$ at time 0. The coefficients $\mu_i \in \mathbb{R}$ and $\sigma_i > 0$ are the instantaneous drift and volatility, respectively, and $W_i(t)$ is a Brownian motion. By contrast, our model accounts for the possibility that the company defaults with an intensity $h_i(t, \mathbf{S}(t), \mathbf{F}(t))$, depending on the whole vector $\mathbf{S}(t) = (S_1(t), \ldots, S_M(t))$ of stock prices and on additional factors $\mathbf{F}(t)$, possibly correlated to $\mathbf{S}(t)$. For example, such factors can be balance sheet data of companies or macroeconomic indicators, chosen to guarantee a reasonably good fit of the default model to data. We refer to Section 5 for the empirical details. We use $h_i(t, \mathbf{S}(t), \mathbf{F}(t))$ to denote the investor's subjective default intensity of the i-th company, i.e., reflecting his own perception about the credit quality of the i-th company.

Denoting by $H_i(t)$ the default indicator process of the i-th company (it is zero before the default time τ_i, and jumps to 1 at the default time τ_i), the probability that company i defaults over an infinitesimal time interval $[t, t + \mathrm{d}t]$ equals

$$\mathbb{P}[H_i(t) = 0 \text{ and } H_i(t + \mathrm{d}t) = 1] = h_i(t, \mathbf{S}(t), \mathbf{F}(t))\,\mathrm{d}t.$$

For a d-dimensional Brownian motion $\mathbf{W}(t)$ and $\mathbf{\Sigma}_i \in \mathbb{R}^d$, the price dynamics of the i-th stock, adjusted for the possibility of default occurrence, is then given by

$$\mathrm{d}S_i(t) = S_i(t-)\big(\mu_i \,\mathrm{d}t + \mathbf{\Sigma}_i \,\mathrm{d}\mathbf{W}(t) - \big(\mathrm{d}H_i(t) - h_i(t, \mathbf{S}(t), \mathbf{F}(t))\,\mathrm{d}t\big)\big), \qquad S_i(0) > 0, \tag{2.1}$$

i.e., the instantaneous default probability $h_i(t, \mathbf{S}(t), \mathbf{F}(t))\,\mathrm{d}t$ is subtracted from $\mathrm{d}H_i(t)$ so that $\mathrm{d}H_i(t) - h_i(t, \mathbf{S}(t), \mathbf{F}(t))\,\mathrm{d}t$ is constant on average (its integral is a martingale), and hence the instantaneous drift of the stock price is still μ_i as in the Black-Scholes model. Conditional on a path of the stock prices, $H_i(t)$ is a Poisson process with intensity $h_i(t, \mathbf{S}(t), \mathbf{F}(t))$ up to the first jump time, independent of the other $H_j(t)$, $j \neq i$, and the Brownian motion $\mathbf{W}(t)$.

The default intensity of a firm may change because there is a change in (1) its own stock price, (2) the stock price of another company, or (3) the additional factors. Our model also has a sound mathematical foundation. We give further details in the e-companion to this paper, where we also formally prove the mathematical existence of the stock price dynamics given in (2.1) via Lemma A.1 under mild assumptions on the default intensities.

For future purposes, we define the volatility matrix $\mathbf{\Sigma}$ as the $(M \times d)$-dimensional matrix with rows $\mathbf{\Sigma}_i$, $i = 1, \ldots, M$, and assume that $\mathbf{\Sigma}\mathbf{\Sigma}^\top$ is invertible, where $\mathbf{\Sigma}^\top$ is the transpose of the matrix $\mathbf{\Sigma}$. This assumption is equivalent to the condition that $\mathbf{\Sigma}$ has rank M, and indicates that even in the absence of default risk, the price process of one stock cannot be replicated using the other stocks in the portfolio.

Remark 2.1. *Consider the case where the number of risk factors is $d = M + 1$, and*

$$\Sigma_{ij} = \begin{cases} \sigma_i \sqrt{1-\rho^2} & \text{if } j = i \\ \sigma_i \rho & \text{if } j = M+1 \\ 0 & \text{otherwise} \end{cases}$$

for constants $\rho \in (-1, 1), \sigma_1 > 0, \ldots, \sigma_M > 0$. We then have

$$\mathbf{\Sigma}_i\,\mathrm{d}\mathbf{W}(t) = \sigma_i\big(\sqrt{1-\rho^2}\,\mathrm{d}W_i(t) + \rho\,\mathrm{d}W_{M+1}(t)\big)$$

for all i. Under such a specification, the Brownian driver of each stock can be split into a systematic risk component $W_{M+1}(t)$ and an idiosyncratic risk component $W_i(t)$.

Hence, our multidimensional stock market model includes two underlying sources of risk dependencies. The first, via market risk, is captured by the volatility matrix $\mathbf{\Sigma}$. The second, via systemic risk, is captured by the default intensity process $h_i(t, \mathbf{S}(t), \mathbf{F}(t))$.

2.2 CDS Prices

We denote by $C_i(t)$ the price of the i-th stylized CDS at time t (with a continuously paid premium), seen from the point of view of the protection buyer. The maturity of the i-th CDS is denoted by T_i. We assume that all CDS contracts on the same name i have the same maturity T_i, possibly different from the maturity T_j of CDS contracts on name $j \neq i$. In reality, an investor can hold CDSs of different maturities on the same name at any point in time. The single maturity assumption per name allows us to reduce the complexity of the portfolio selection problem, but still gives a rich universe of securities for which we can analyze the impact of systemic dependencies on portfolio allocations. If the default of the i-th name happens at $\tau_i < T_i$, the protection buyer receives a payment equal to the product of the notional and the loss rate L_i from the protection seller. As is standard in the literature (see, for instance, Berndt et al. (2005) and Giesecke et al. (2014)), the loss rate is assumed to be the same, both in the views of the market and of the individual investor, hence there will be no premium for recovery risk. The protection buyer pays the spread premium ν_i to the protection seller for this coverage until the earlier of default time or maturity. Therefore, the cumulative dividend process of the unit notional i-th CDS received by the protection buyer is given by

$$D_i(t) = L_i H_i(t) - \nu_i \int_0^t (1 - H_i(s))\, \mathrm{d}s, \tag{2.2}$$

where $L_i H_i(t)$ is the payment at default time and $\nu_i \int_0^t (1-H_i(s))\, \mathrm{d}s$ is the cumulative payment before the occurrence of default. As for any traded derivative, the CDS price is equal to the expected discounted payoff under a risk-neutral probability measure $\mathbb{Q}$. Here, it is important to distinguish $\mathbb{Q}$ from the subjective probability measure $\mathbb{P}$ of the investor. Since CDS prices are determined by the market and not by a single investor, the dynamics of all other assets, namely money market account and stocks, are taken into account when pricing the CDS. This results in the choice of a risk-neutral probability measure $\mathbb{Q}$ determined by the market. We denote by $\lambda_i(t, \boldsymbol{S}(t))$ the default intensity of stock i under the risk-neutral probability measure $\mathbb{Q}$. Because $\mathbb{Q}$ is different from $\mathbb{P}$, the risk-neutral default intensity $\lambda_i(t, \boldsymbol{S}(t))$ differs from the subjective default intensity. In particular, the choice of the probability measure $\mathbb{P}$ is driven by a judicious selection of factors which have been found to be statistically significant by empirical research. As we will elaborate in Section 5, those include macroeconomic indicators (interest rates) as well as balance sheet variables (distance to default). We use $\boldsymbol{F}(t)$ to denote the factors whose choice is subjective to the individual investor. Correspondingly, $h_i(t, \boldsymbol{S}(t), \boldsymbol{F}(t))$ denotes the default intensity of the i-th stock under the investor's subjective probability measure. The model for the dynamics of the risk-neutral default intensities is an extension of that proposed by Linetsky (2006) to a systemic context, i.e., the default intensity of each company also depends on the market and default risk of others. We will exploit this dependence in the calibration algorithm developed in Section 5, where the subjective default intensity also depends on balance sheet data, while we achieve a good fit for the risk-neutral default intensity using only stock price data. Note that the factors driving the investor's subjective beliefs on default risk do not affect CDS prices, hence it is not necessary to propose a dynamical model for these factors in order to compute CDS prices. Technical conditions on $\lambda_i(t, \boldsymbol{S}(t))$ which guarantee that the measure $\mathbb{Q}$ and its related calculations are mathematically sound are given in the e-companion to this paper. As a consequence of the market's risk-neutral valuation, the price of the i-th CDS is given by

$$C_i(t) = \mathbb{E}^{\mathbb{Q}} \left[\int_t^{T_i} \mathrm{e}^{-r(u-t)} \mathrm{d}D_i(u) \middle| \boldsymbol{S}(t) \right], \tag{2.3}$$

where the $\mathbb{Q}$ expectation is conditional on the current state of information given by the vector of stock prices $\boldsymbol{S}(t)$. Notice that conditioning on $\boldsymbol{S}(t)$ is equivalent to considering all current and past information from the stock prices up to time t because all relevant information is reflected in the stock price processes, which are Markovian under the probability measure $\mathbb{Q}$. Because of this Markov property, the right-hand side of (2.3) is a function of t and $\boldsymbol{S}(t)$ only, and we can write $C_i(t) = \Phi_i(t, \boldsymbol{S}(t))$ for some function Φ_i. The expression $\Phi_i(t, s_1, \ldots, s_M)$ is the price of the i-th CDS at time t if the stock prices at time t are given by $S_1(t) = s_1, \ldots, S_M(t) = s_M$. We have that $\Phi_i(t, s_1, \ldots, s_M) = 0$ if $s_i = 0$, because at the default time the loss rate is paid and the CDS becomes worthless. Note that because of the default intensity specification, the market value of the i-th CDS depends on the prices of all stocks and not only on the i-th stock price.

We say that the i-th CDS is *untriggered* at t if the credit event of the i-th company has not yet occurred and $t \leq T_i$. We denote by $\mathbb{M}(t) \subseteq \{1, \ldots, M\}$ the set of untriggered CDSs at time t.

3 Dynamic Investment Problem

This section formulates the dynamic portfolio selection problem of the investor. We consider a rational investor who wants to maximize his expected utility from terminal wealth at time T by dynamically investing into the money market account, stocks and credit default swaps trading upfront. This reflects regulatory market practice established by the Big Bang protocol ISDA (2009), moving the CDS market towards fixed coupon CDS trading.[3]

The investor starts with a positive wealth $V(0)$ at time zero. As usual in optimal investment problems, we denote by $\pi_i(t)$ the proportion of total wealth invested in the i-th stock at time t. The investor can also purchase and sell a CDS written on the same company i. Next, we describe how CDS contracts initiated at different points of time can be treated in our portfolio problem. To fix ideas, we first consider discrete time instants $t_1, \ldots, t_N$ at which the investor rebalances his portfolio. We denote by $\alpha_i(t_j, T_i^j)$ the number of CDS contracts referencing entity i maturing at T_i^j that the investor buys or sells at time t_j. While the CDS entered at time t_{j+1} is a different contract, CDSs are standardized, and in particular single-name CDS contracts mature on standard quarterly end dates. Therefore, the investor can offset an existing CDS trade by taking an opposite position in the equivalent contract, agreed with a possibly different counterparty. Leaving aside counterparty risk considerations, the aggregate position in the credit default swaps is obtained by summing numbers of contracts entered at different instants of time as long as they all mature on the same date. In other words, suppose that the investor purchases CDS contracts referencing company i, first at time t_j and later at time t_ℓ. We would then have that the total position in the i-th CDS of the investor at time t_ℓ is $\alpha_i(t_j, T_i^j) + \alpha_i(t_\ell, T_i^\ell)$ as long as $T_i^j = T_i^\ell$.

We would like to remark here that the U.S. Securities and Exchange Commission (SEC) has taken steps to facilitate central clearing of over-the-counter derivatives so to eliminate counterparty risk.[4] Given the increasing number of centrally cleared single name CDS contracts, ignoring counterparty risk in our analysis would be very close to a realistic setting.[5]

[3]More specifically, the universe of names is split into investment grade names trading with fixed 1% coupon and the high yield names trading with 5% coupon.

[4]In the United States, the role of central counterparties is discussed in Title VII of the Dodd-Frank Wall Street Reform and Consumer Protection Act; see http://www.sec.gov/spotlight/dodd-frank/derivatives. Central counterparties (CCPs) are also an integral component of the European Market Infrastructure Regulation (EMIR) framework.

[5]The world's first clearing house for credit default swaps, ICE Clear Credit, clears more than 330 single name CDS contracts. These include North American, European and Emerging Markets indexes and North American, European and sovereign single names. Most of the credit default contracts written on companies listed in the DJIA are cleared by ICE Clear Credit. We also refer to https://www.theice.com/clear-credit for further details.

Notice that the maturities of the CDS contracts on the same name are indeed identical for all contracts entered in the same quarter because these mature on standard quarterly end dates. This means that it is enough to consider an optimization problem over a single quarter so that all contracts mature on the same date (see also Remark 4.3 later, which explains how to deal with different quarters). Consequently, we can denote by $\hat{\psi}_i(t)$ the net position in shares of the i-th CDS at time t for a given maturity T_i. For future purposes, we define $\psi_i(t) = \hat{\psi}_i(t)/V(t)$ to be the number of shares of the i-th CDS divided by the total wealth. Note, however, that we do not multiply by the CDS price, hence $\psi_i(t)$ is *not* the proportion of wealth invested in the i-th CDS. For CDS investment strategies, we use a slightly different parametrization than for stocks because credit swaps are zero-sum side bets valued at zero when the market price of the premium leg equals the corresponding price of the protection leg. Hence, differently from the stock, the CDS price can be zero even before the credit event has been triggered. Consequently, it is important to know the units and not only the dollar amount invested in the CDS.

The investor does not have intermediate consumption nor capital income to support his purchase of financial assets. Because of this self-financing condition, the position in the money market account is determined as residual from the chosen strategies in stocks and CDSs; see also Bielecki et al. (2008) for details. We use the notation $\boldsymbol{\pi} = (\pi_1, \dots, \pi_M)$ and $\boldsymbol{\psi} = (\psi_1, \dots, \psi_M)$ for the vector of stock and CDS investment strategies, respectively, and write $V^{\boldsymbol{\pi},\boldsymbol{\psi}}(t)$ to denote the investor's wealth at time t corresponding to a chosen strategy $(\boldsymbol{\pi}, \boldsymbol{\psi})$.

Our goal is to find the optimal strategy, which maximizes the expected utility from terminal wealth

$$\mathbb{E}\big[U\big(V^{\boldsymbol{\pi},\boldsymbol{\psi}}(T)\big)\big] \tag{3.1}$$

over all admissible strategies $(\boldsymbol{\pi}, \boldsymbol{\psi})$, where $T \in (0, \min(T_1, \dots, T_M))$. We define a strategy to be admissible if it has positive wealth and satisfies certain technical conditions, as given in Definition B.1 of the e-companion to this paper. In (3.1), $\mathbb{E}$ denotes the expectation under the subjective probability measure $\mathbb{P}$ of the investor. It is worth emphasizing here the key role played by the two probability measures. While the CDS price process is observed under the risk-neutral probability measure, the investor wishes to optimize his expected utility under the probability measure capturing his perception about risk factors. As we will see in the next section, the relation between the default intensities under the two probability measures plays a crucial role in determining the optimal investment decisions.

We consider logarithmic utility, i.e., $U(v) = \log(v)$. Our choice is driven by three main considerations. Firstly, logarithmic utility allows exemplifying typical systemic effects arising when the investor is active both on the equity and on the CDS market. Secondly, there is empirical evidence (see Gordon et al. (1972)) that wealthy investors, such as the highly specialized market players in the CDS market, maximize expected logarithmic utility. Thirdly, under a logarithmic utility function, the value function of the control problem decouples from the optimal investment strategy (this would not be the case if we use a power utility criterion) and leads to tractable representation formulas for systemic risk analysis.

4 Optimal Investment Strategies

This section analyzes the optimal investment strategies. Section 4.1 gives the explicit expressions. Section 4.2 discusses the mechanisms through which systemic influences affect the optimal allocation decisions. Section 4.3 gives an example of the contagion effects in the case of a cyclical dependence structure.

4.1 Explicit Formula for the Optimal Strategy

The optimal investment strategies have a crucial forward-looking component, given that they depend on a matrix which captures the sensitivities of CDS prices to the total risk of the investment universe. Such a matrix is given by $\boldsymbol{\Theta}(t, \boldsymbol{S}(t)) = (\Theta_{n,i}(t, \boldsymbol{S}(t)))_{n,i \in \mathbb{M}(t)}$ with entries

$$\Theta_{n,i}(t, \boldsymbol{S}(t)) = \begin{cases} L_i - C_i(t) + \frac{\partial \Phi_i}{\partial s_i}(t, \boldsymbol{S}(t)) S_i(t) & \text{if } i = n \\ \Phi_i\big(t, \boldsymbol{S}^{(n)}(t)\big) - C_i(t) + \frac{\partial \Phi_i}{\partial s_n}(t, \boldsymbol{S}(t)) S_n(t) & \text{if } i \neq n. \end{cases} \tag{4.1}$$

In the above expression, $\boldsymbol{S}^{(n)}(t)$ equals $\boldsymbol{S}(t)$ except for the n-th component of $\boldsymbol{S}^{(n)}(t)$, which equals zero. Then $\Phi_i\big(t, \boldsymbol{S}^{(n)}(t)\big)$ is the price of the i-th CDS when the n-th stock defaults. The partial derivatives terms $\frac{\partial \Phi_n}{\partial s_i}$ are used only to determine the strategies for positive s_i since $i \in \mathbb{M}(t)$. Indeed, no investment is made in a CDS or stock referencing a defaulted name. We show in the e-companion to this paper that the entry $\Theta_{n,i}(t, \boldsymbol{S}(t))$ may be interpreted as the contribution of the stock S_n to the error committed when hedging the compounded payment stream of the i-th CDS. For this reason, we refer to $\Theta_{n,i}(t, \boldsymbol{S}(t))$ as the (n, i)-th *hedging error*. This means that the matrix $\boldsymbol{\Theta}(t, \boldsymbol{S}(t))$ reflects only residual credit risk, i.e. remaining after eliminating the market risk component. We make the following mild assumption.

Assumption A. *The stochastic matrix* $\mathbf{\Theta}(t, \boldsymbol{S}(t))$ *has full rank.*

This can be shown to be equivalent to the uniqueness of the risk-neutral default intensity. In view of this equivalence, the CDS price dynamics are uniquely determined by the market.[6] It can be directly verified that if all default intensities λ_n's are constant, then the matrix $\mathbf{\Theta}(t)$ is diagonal with strictly positive entries $L_n - C_n(t)$, hence Assumption A trivially holds.

For the sake of readability, in the rest of the section we omit the argument $(t, \mathbf{S}(t), \mathbf{F}(t))$ in h_i and $(t, \mathbf{S}(t))$ in λ_i and $\mathbf{\Theta}$ so that h_i, λ_i and $\mathbf{\Theta}$ mean $h_i(t, \mathbf{S}(t), \mathbf{F}(t))$, $\lambda_i(t, \mathbf{S}(t))$ and $\mathbf{\Theta}(t, \mathbf{S}(t))$, respectively.

We are now ready to give our main result in the following theorem. The main insight behind its proof reported in the e-companion to this paper is to explore the structure of the investor's wealth process and use a suitable reparameterization of the strategies, so to recast the optimizer as the solution of a linear system of equations.

Theorem 4.1. *The optimal investment strategy in untriggered CDSs and stocks is given by*

$$\boldsymbol{\psi}(t) = \mathbf{\Theta}^{-1}\left[\left(\mathbf{\Sigma}\mathbf{\Sigma}^\top\right)^{-1}(\mu_n - r + h_n - \lambda_n)_{n\in\mathbb{M}(t)} + \left(\frac{h_n - \lambda_n}{\lambda_n}\right)_{n\in\mathbb{M}(t)}\right], \tag{4.2}$$

$$\boldsymbol{\pi}(t) = \left(\mathbf{\Sigma}\mathbf{\Sigma}^\top\right)^{-1}(\mu_n - r + h_n - \lambda_n)_{n\,\in\,\mathbb{M}(t)} - \left(\sum_{i\,\in\,\mathbb{M}(t)} \psi_i(t) S_n(t) \frac{\partial \Phi_i}{\partial s_n}(t, \mathbf{S}(t))\right)_{n\in\mathbb{M}(t)}. \tag{4.3}$$

The maximal expected utility is given by

$$\begin{aligned}\mathbb{E}\left[U\left(V^{\boldsymbol{\pi},\boldsymbol{\psi}}(T)\right)\right] &= \log(V(0)) + rT + \sum_{n=1}^{M} \mathbb{E}\left[\int_0^{\tau_n \wedge T} \left(-\log\left(\frac{\lambda_n}{h_n}\right) + \frac{\lambda_n}{h_n} - 1\right) h_n \,\mathrm{d}t\right] \\ &\quad + \frac{1}{2}\mathbb{E}\left[\int_0^T (\mu_n - r + h_n - \lambda_n)^\top_{n\,\in\,\mathbb{M}(t)} \left(\mathbf{\Sigma}\mathbf{\Sigma}^\top\right)^{-1} (\mu_n - r + h_n - \lambda_n)_{n\,\in\,\mathbb{M}(t)} \,\mathrm{d}t\right].\end{aligned} \tag{4.4}$$

By Assumption A, the hedging error matrix $\mathbf{\Theta}$ defined in (4.1) has full rank, hence its matrix inverse $\mathbf{\Theta}^{-1}$ in (4.2) exists. Moreover, we have used the notation $(\mu_n - r + h_n - \lambda_n)_{n\in\mathbb{M}(t)}$ to denote a vector with a number of entries equal to the number of untriggered CDSs. The volatility matrix $\mathbf{\Sigma}$ in (4.2)–(4.4) is understood as the submatrix of the original volatility matrix consisting of rows corresponding to untriggered CDSs. Therefore, $\mathbf{\Sigma}$ in (4.2)–(4.4) is effectively a matrix of dimension $|\mathbb{M}(t)| \times d$, where

[6] Bielecki et al. (2010) make a similar assumption (see Assumption 2.2 in their paper) in a special case of our market model consisting of one CDS and the stock of the reference entity.

$|\mathbb{M}(t)|$ is the number of untriggered CDSs at time t. Because the original matrix has rank M, the reduced matrix $\mathbf{\Sigma}$ in (4.2)–(4.4) has rank $|\mathbb{M}(t)|$ and hence the matrix $\mathbf{\Sigma}\mathbf{\Sigma}^\top$ is still invertible.

Eq. (4.4) indicates that the additional expected utility derived by the investor (over the safe investment in the money market account) consists of two components. The first is the contribution to the investor's expected utility coming from the CDS market and can be understood as follows. From the second-order Taylor approximation of the logarithmic function, we know that $-\log(x) + x - 1 \approx \frac{(x-1)^2}{2}$ when $x \approx 1$. Setting $x = \frac{\lambda_n}{h_n}$, we can see directly that the investor does not gain any utility from investing in the CDS referencing company n if $x = 1$, i.e., if his perception on the credit quality of company n coincides with the market views. On the contrary, as soon as risk-neutral and subjective default intensities differ, the investor always gains utility given that $\frac{(x-1)^2}{2} > 0$. If $x < 1$, the investor would purchase the CDS as he believes that the market is underpricing the credit risk of company n, while he would sell protection if he believes that the market is overpricing the credit risk of company n ($x > 1$). The second component quantifies the contribution to expected utility coming from stock investments and has a similar structure as in the classical Merton equity selection problem. The main difference of this second component is that in our model the price of risk includes both market and default risk.

The above discussion is also supported by the following corollary, which gives the sensitivity of the investor's wealth to market and credit risk.

Corollary 4.2. *The optimal strategies $\boldsymbol{\pi}(t)$ and $\boldsymbol{\psi}(t)$ are uniquely determined by*

$$\underbrace{\mathbf{\Sigma}^\top \left(\sum_{j \in \mathbb{M}(t)} \psi_j(t) \frac{\partial \Phi_j}{\partial s_n}(t, \boldsymbol{S}(t)) S_n(t) + \pi_n(t) \right)_{n \in \mathbb{M}(t)}}_{\text{relative wealth sensitivity in } \mathbf{W}} = \mathbf{\Sigma}^\top \left(\mathbf{\Sigma}\mathbf{\Sigma}^\top\right)^{-1} \left(\mu_n - r + h_n - \lambda_n\right)_{n \in \mathbb{M}(t)},$$

$$\underbrace{\sum_{i \in \mathbb{M}(t)} \psi_i(t) \big(\Phi_i\big(t, \boldsymbol{S}^{(n)}(t)\big) - C_i(t)\big) + \psi_n(t) L_n - \pi_n(t)}_{\text{relative wealth sensitivity in } H_n} = \frac{h_n - \lambda_n}{\lambda_n} \quad \text{for all } n \in \mathbb{M}(t).$$

The above expressions indicate that if the investor's subjective belief on the default risk of the n-th company is higher than the market, his optimal wealth would increase if the n-th company defaults. This is because the investor would choose to go long credit given that, according to his beliefs, the market is underpricing credit risk.

Next, we analyze the dual role played by the stock in the portfolio investment. We discuss this in the absence of systemic effects as this allows us to better explain the qualitative behavior without obscuring it with the complex dependence structure of risk factors. In other words, we consider the case where the default intensity of company n only depends on the price of the n-th stock, i.e., $h_n(t, \mathbf{S}(t)) = h_n(t, S_n(t))$ and the volatility matrix $\mathbf{\Sigma}$ is diagonal with entries σ_n for $n \in \mathbb{M}(t)$. Then $\mathbf{\Theta}$ becomes diagonal and

$$\psi_n(t) = \frac{1}{\Theta_{n,n}} \left(\frac{\mu_n - r + h_n - \lambda_n}{\sigma_n^2} + \frac{h_n - \lambda_n}{\lambda_n} \right).$$

Consider the case where the difference between investor's subjective and market risk perceptions on the n-th company is strictly positive, i.e., $\frac{\mu_n - r + h_n - \lambda_n}{\sigma_n^2} + \frac{h_n - \lambda_n}{\lambda_n} > 0$. Next assume that $\Theta_{n,n} > 0$, i.e., the investor purchases the credit default swap referencing company n. For the following discussion, let us consider the case where the risk-neutral default intensity λ_n is a decreasing function of the stock. This means that the default risk of a name becomes higher when the price of the corresponding stock declines, a common assumption made, for example, in Linetsky (2006). We analyze two distinct cases which can arise based on the sign of the default risk adjusted Merton ratio:

1. $\frac{\mu_n - r + h_n - \lambda_n}{\sigma_n^2} < 0$. In the classical equity portfolio selection problem, the investor would short-sell the stock since the Merton ratio is negative. However, here we must account for the dependence of the CDS price on the market risk of the stock, captured by the term $-\psi_n(t) S_n(t) \frac{\partial \Phi_n}{\partial s_n}(t, \mathbf{S}(t))$ in the stock investment strategy given by (4.3). Since an appreciation of the stock price decreases the default intensity and pushes down the CDS price, the partial derivative term is negative. Consequently, the previous term is positive. If it is higher than the absolute value of the Merton ratio, the investor would purchase the stock to *hedge* against the market risk incurred from the large CDS position. If it is smaller, the market risk accrued from the CDS position is low and the investor would *speculate* and short sell the stock.

2. $\frac{\mu_n - r + h_n - \lambda_n}{\sigma_n^2} > 0$. If stocks were the only instruments in the portfolio, the proportion of wealth invested in the n-th stock would be equal to the Merton ratio. As mentioned above, when CDSs are included we need to account for the dependence of the CDS price on the market risk of the stock. Since the term $-\psi_n(t) S_n(t) \frac{\partial \Phi_n}{\partial s_n}(t, \mathbf{S}(t))$ is positive, the investor would invest an even higher fraction of wealth in the n-th stock. This means that a portion of his stock shares is being used to speculate, and the remaining portion to hedge the market risk accumulated from the CDS position.

Remark 4.3. *Observe that the optimal investment strategy does not depend on the trading horizon T. This implies that we can split the optimization into subproblems with horizon of one quarter and clear the positions at the end of each quarter. In each of these quarters, the maturity of the CDS does not change so that we effectively consider strategies in CDSs with fixed maturity in the optimization problem. Realistically, when the position is cleared at the end of each quarter, the investor would need to enter into an offsetting position and thus incur transaction costs. We ignore these costs in the present study, which is primarily focused on the impact of systemic influences on optimal trading strategies.*

4.2 The Systemic Influences on Credit Swap Strategies

This section analyzes the structure of the explicit representation for the CDS investment strategy given in Theorem 4.1. We show that two mechanisms interact to shape the optimal CDS strategy: *amplification* and *transmission* of market risk premia. To explain it, we first rewrite the optimal CDS strategy as

$$\boldsymbol{\psi}(t) = \boldsymbol{\Theta}^{-1}\Big[\underbrace{\Big(\big(\boldsymbol{\Sigma}\boldsymbol{\Sigma}^{\top}\big)^{-1} + \boldsymbol{\Lambda}^{-1}\Big)}_{\text{amplifier 1}}\underbrace{(h_n - \lambda_n)_{n\in\mathbb{M}(t)}}_{\text{source 1}} + \underbrace{\big(\boldsymbol{\Sigma}\boldsymbol{\Sigma}^{\top}\big)^{-1}}_{\text{amplifier 2}}\underbrace{(\mu_n - r)_{n\in\mathbb{M}(t)}}_{\text{source 2}}\Big], \tag{4.5}$$

where $\boldsymbol{\Lambda} = \mathrm{diag}(\lambda_\mathrm{n})_{\mathrm{n}\in\mathbb{M}(\mathrm{t})}$ is the diagonal matrix whose entries are the risk-neutral default intensities λ_n for $n \in \mathbb{M}(t)$. We also recall that the hedging error matrix $\boldsymbol{\Theta}$ has been defined in (4.1). There are two sources of risk premia: the premium for default timing risk $h_n - \lambda_n$ and the equity risk premium $\mu_n - r$. Each of these sources is subject to an amplification. The amplification factor of the premium for default timing risk is $\big(\big(\boldsymbol{\Sigma}\boldsymbol{\Sigma}^{\top}\big)^{-1} + \boldsymbol{\Lambda}^{-1}\big)$, i.e., related to the sum of the inverse of market risk (captured by the covariance $\boldsymbol{\Sigma}\boldsymbol{\Sigma}^{\top}$ of stock returns) and of the inverse of default risk (captured by λ_n). For the equity risk premium, the size of the amplification is inversely related to the market risk. We can further decompose the first part of (4.5) as $\boldsymbol{\Theta}^{-1}\Big(\big(\boldsymbol{\Sigma}\boldsymbol{\Sigma}^{\top}\big)^{-1} + \boldsymbol{\Lambda}^{-1}\Big)(h_n - \lambda_n)_{n\in\mathbb{M}(t)} = \sum_{n\in\mathbb{M}(t)} \boldsymbol{\psi}^{(n)}$, where we define

$$\boldsymbol{\psi}^{(n)} = \boldsymbol{\Theta}^{-1}\Big(\big(\boldsymbol{\Sigma}\boldsymbol{\Sigma}^{\top}\big)^{-1} + \boldsymbol{\Lambda}^{-1}\Big)\boldsymbol{\kappa}^{(n)}, \tag{4.6}$$

with the vectors $\boldsymbol{\kappa}^{(n)}$ given by

$$\kappa_i^{(n)} = \begin{cases} h_n - \lambda_n & \text{if } i = n, \\ 0 & \text{otherwise.} \end{cases}$$

Expression (4.6) allows us to see the impact of systemic influences on the optimal CDS strategies. First, notice that the j-th component of $\boldsymbol{\psi}^{(n)}$ gives the impact of the premium for default timing risk of company n on the investment strategy in the j-th CDS, $j \neq n$. Hence, we call the j-th component of $\sum_{n \neq j} \boldsymbol{\psi}^{(n)}$ the *systemic influence* on the investment strategy in the j-th CDS because it gives the total impact, contributed by the premia for default timing risk of all portfolio names except for the j-th company, on the investment strategy in the j-th CDS. Next, we illustrate the transmission mechanism through which a change in premia for default timing risk of the portfolio entities affects the strategy in the j-th CDS. To this purpose, we first develop a decomposition of the matrix $\boldsymbol{\Theta}^{-1}$. Define the diagonal matrix $\boldsymbol{\Pi}$ whose n-th entry is given by $\Pi_{n,n}(t) = L_n - C_n(t) + \frac{\partial \Phi_n}{\partial s_n}(t, \boldsymbol{S}(t)) S_n(t)$. Then, the following asymptotic formula holds

$$\boldsymbol{\Theta}^{-1} = \boldsymbol{\Pi}^{-1} \left(\boldsymbol{\Theta} \boldsymbol{\Pi}^{-1} \right)^{-1} = \boldsymbol{\Pi}^{-1} \sum_{k=0}^{\infty} \left(\boldsymbol{I} - \boldsymbol{\Theta} \boldsymbol{\Pi}^{-1} \right)^{k}, \tag{4.7}$$

provided that all eigenvalues of $\boldsymbol{I} - \boldsymbol{\Theta}\boldsymbol{\Pi}^{-1}$ are strictly less than one in modulus. We refer to such a matrix as the *contagion matrix*. Combining (4.6) with (4.7), we have

$$\begin{aligned} \boldsymbol{\psi}^{(n)} &= \boldsymbol{\Theta}^{-1} \Big(\left(\boldsymbol{\Sigma} \boldsymbol{\Sigma}^{\top} \right)^{-1} + \boldsymbol{\Lambda}^{-1} \Big) \boldsymbol{\kappa}^{(n)} = \boldsymbol{\Pi}^{-1} \sum_{k=0}^{\infty} \underbrace{\left(\boldsymbol{I} - \boldsymbol{\Theta} \boldsymbol{\Pi}^{-1} \right)^{k}}_{\text{propagator}} \underbrace{\Big(\left(\boldsymbol{\Sigma} \boldsymbol{\Sigma}^{\top} \right)^{-1} + \boldsymbol{\Lambda}^{-1} \Big)}_{\text{amplifier}} \underbrace{\boldsymbol{\kappa}^{(n)}}_{\text{source}} \\ &= \boldsymbol{\Pi}^{-1} \Big(\underbrace{\Big(\left(\boldsymbol{\Sigma} \boldsymbol{\Sigma}^{\top} \right)^{-1} + \boldsymbol{\Lambda}^{-1} \Big) \boldsymbol{\kappa}^{(n)}}_{\text{round 0 of systemic influence}} + \underbrace{\left(\boldsymbol{I} - \boldsymbol{\Theta} \boldsymbol{\Pi}^{-1} \right) \Big(\left(\boldsymbol{\Sigma} \boldsymbol{\Sigma}^{\top} \right)^{-1} + \boldsymbol{\Lambda}^{-1} \Big) \boldsymbol{\kappa}^{(n)}}_{\text{round 1 of systemic influence}} \\ &\quad + \cdots + \underbrace{\left(\boldsymbol{I} - \boldsymbol{\Theta} \boldsymbol{\Pi}^{-1} \right)^{k} \Big(\left(\boldsymbol{\Sigma} \boldsymbol{\Sigma}^{\top} \right)^{-1} + \boldsymbol{\Lambda}^{-1} \Big) \boldsymbol{\kappa}^{(n)}}_{\text{round } k \text{ of systemic influence}} + \dots \end{aligned} \tag{4.8}$$

Assume that the investor has chosen the optimal strategy and there is now a change in the default premium of company n. This affects the investor's optimal level of credit risk exposure in the n-th company, which results primarily in a change in the n-th CDS position, and secondarily, in a further adjustment in the stock positions because of the dependence of the n-th CDS price on the corresponding stock and the correlation between the stocks. Because this mechanism is driven only by market risk correlation and not default dependencies, we refer to it as round 0 in (4.8). However, the adjusted position in the n-th CDS also changes the credit risk exposure in all other companies due to the systemic dependencies. As a consequence, the investor adjusts his positions in all other CDS to take the updated n-th CDS position into account (round 1 of systemic influences). A further consequence of these adjustments is an additional adjustment in all CDS positions, which results in round 2 of systemic influences, and goes on in rounds 3, 4, ..., resulting in smaller and smaller adjustments corresponding to consecutive powers of the contagion matrix $\boldsymbol{I} - \boldsymbol{\Theta}\boldsymbol{\Pi}^{-1}$. The aggregate outcomes of the procedure are the new optimal CDS strategies.

How fast this iterative procedure leads to the new optimal strategies depends on the contagion matrix $\boldsymbol{I} - \boldsymbol{\Theta}\boldsymbol{\Pi}^{-1}$. If $\boldsymbol{\Theta}$ is the diagonal matrix $\boldsymbol{\Pi}$, then $\boldsymbol{I} - \boldsymbol{\Theta}\boldsymbol{\Pi}^{-1}$ is a zero matrix and the adjustment is instantaneous (round 0 only). On the other hand, if $\boldsymbol{\Theta}$ has large off-diagonal entries so that the contagion matrix has eigenvalues close to one in modulus, the transmission effect is strong because the investment strategy in the j-th CDS is highly sensitive to market and default risk of other stocks in the portfolio. In this case, several rounds of adjustments are needed to closely approximate the new optimal strategy.

Using the decomposition formula (4.7), we may also rewrite the second term in (4.5) as

$$\boldsymbol{\Theta}^{-1}\big(\boldsymbol{\Sigma}\boldsymbol{\Sigma}^{\top}\big)^{-1}(\mu_n - r)_{n\in\mathbb{M}(t)} = \boldsymbol{\Pi}^{-1}\sum_{k=0}^{\infty}\underbrace{\big(\boldsymbol{I} - \boldsymbol{\Theta}\boldsymbol{\Pi}^{-1}\big)^{k}}_{\text{propagator}}\underbrace{\big(\boldsymbol{\Sigma}\boldsymbol{\Sigma}^{\top}\big)^{-1}}_{\text{amplifier}}\underbrace{(\mu_n - r)_{n\in\mathbb{M}(t)}}_{\text{source}},$$

which can be decomposed in an analogous way to the first term in (4.5) if we define the vectors $\tilde{\boldsymbol{\kappa}}^{(n)}$ by

$$\tilde{\kappa}_i^{(n)} = \begin{cases} \mu_n - r & \text{if } i = n, \\ 0 & \text{otherwise.} \end{cases}$$

The above expressions indicate that a change in the equity risk premium leads the investor to optimally rebalance the portfolio, which can be achieved by using an iterative procedure similar to that described above. However, the amplification factor for the equity risk premium only depends on the market risk of the investment universe. In contrast, either a low market risk or a low default risk leads to a large amplification of the premium for default timing risk.

4.3 Example

We have seen in the previous section that consecutive powers of the contagion matrix $\boldsymbol{I} - \boldsymbol{\Theta}\boldsymbol{\Pi}^{-1}$ are associated with subsequent rounds of transmission of systemic influences. This section illustrates the structure of these matrices in the case of a cyclical dependence structure. In networks with higher interconnectedness there would be stronger systemic influences through multiple channels of credit and market risk dependencies. Concretely, for $i = 1, \ldots, M$, we consider the default intensity specification

$$\lambda_n(t, \mathbf{S}(t)) = \begin{cases} \lambda_n(t, S_n(t), S_{n+1}(t)) & \text{if } n < M \\ \lambda_n(t, S_n(t), S_1(t)) & \text{if } n = M. \end{cases}$$

This means that the default of company n has a direct impact on the default intensity of its neighboring company $n-1$ if $n \neq 1$. The cycle is closed with the impact on the default intensity of company M caused by the default of company 1. This is a special case of the looping default model analyzed by Jarrow and Yu (2001) in the context of pricing, see also Yu (2005) for the construction of default times admitting interacting intensities. Under this specification, the matrix Θ defined in (4.1) takes the bi-diagonal form, where the non-zero entries are given by $\Theta_{n,n}$, $1 \leq n \leq M$, $\Theta_{n+1,n}$, $1 \leq n \leq M-1$, and $\Theta_{1,M}$. A direct computation shows that

$$\boldsymbol{I} - \boldsymbol{\Theta}\boldsymbol{\Pi}^{-1} = \begin{bmatrix} 0 & \cdots & & & & -\frac{\Theta_{1,M}}{\Theta_{1,1}} \\ -\frac{\Theta_{2,1}}{\Theta_{2,2}} & 0 & \cdots & & & 0 \\ 0 & -\frac{\Theta_{3,2}}{\Theta_{3,3}} & 0 & \cdots & & \vdots \\ \vdots & 0 & \ddots & 0 & \cdots & 0 \end{bmatrix}. \tag{4.9}$$

The off-diagonal entries of the above matrix are ratios of hedging errors. The spectral radius of this matrix is given by

$$\left| \frac{\Theta_{1,M}}{\Theta_{1,1}} \prod_{n=1}^{M-1} \frac{\Theta_{n+1,n}}{\Theta_{n+1,n+1}} \right|^{\frac{1}{M}}$$

For the spectral radius to be smaller than unity, we need at least one hedging error ratio to be sufficiently small. This means that the propagation effects along the chain must be attenuated by at least one company in the portfolio. When this is not the case, we expect the first round of transmission of risk premia to have a strong impact on the optimal CDS strategies.

When the expression in (4.7) has a nonzero term only for $k = 0$, we have that $\boldsymbol{\Theta}^{-1} = \boldsymbol{\Pi}^{-1}$, and hence the optimal strategies are only impacted by idiosyncratic effects, i.e., the optimal strategy in the n-th CDS is only influenced by the stock price of the n-th company.

The terms corresponding to indices $k \geq 1$ in the expansion (4.7) reflect the impact of subsequent rounds of systemic influence. More specifically, after one round we have the effects of direct contagion measured by the matrix (4.9). This indicates that, for each company, the optimal investment strategy in the CDS referencing it is adjusted to reflect influences of stock price changes of neighboring companies. We illustrate this behavior in the left panel of Figure 1.

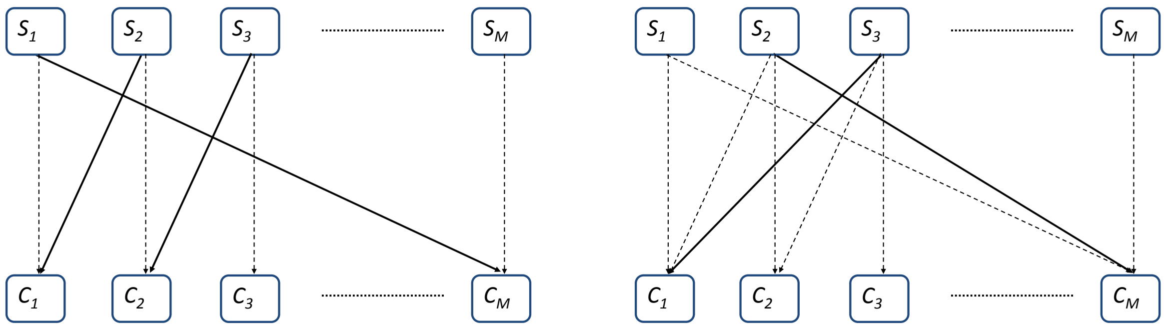

Figure 1: Left panel: The graph of dependencies after round 1 of systemic influence while the dash-dot lines refer to round 0. Right panel: The graph of dependencies after round 2 of systemic influence while the dash-dot lines refer to rounds 0 and 1 of systemic influence.

The impact of the second round of systemic influence is quantified by the matrix

$$\left(\boldsymbol{I}-\boldsymbol{\Theta}\boldsymbol{\Pi}^{-1}\right)^2=\begin{bmatrix} 0 & \cdots & 0 & \frac{\Theta_{M,M-1}}{\Theta_{M,M}}\frac{\Theta_{1,M}}{\Theta_{1,1}} & 0 \\ 0 & \cdots & 0 & 0 & \frac{\Theta_{1,M}}{\Theta_{1,1}}\frac{\Theta_{2,1}}{\Theta_{2,2}} \\ \frac{\Theta_{2,1}}{\Theta_{2,2}}\frac{\Theta_{3,2}}{\Theta_{3,3}} & 0 & \cdots & 0 & 0 \\ 0 & \frac{\Theta_{3,2}}{\Theta_{3,3}}\frac{\Theta_{4,3}}{\Theta_{4,4}} & 0 & \cdots & 0 \\ \vdots & \ddots & \vdots & & \vdots \end{bmatrix}$$

and captures second-order contagion effects, i.e., the influence on the investment strategy in the j-th CDS due to residual default risk transmitted to node j through its neighbor $j-1$, who has itself been influenced by its own neighbor $j-2$ in the previous round.

Iterating this procedure, we obtain that after $M-1$ rounds the CDS strategy in any company is affected by the risk premia of all other companies. As described in the previous section, this procedure results in smaller and smaller adjustments and the aggregate outcome of the corresponding rebalancing process gives the new optimal strategy. The above matrix structure also indicates that if the transmission power, as measured by the ratio of hedging errors, of a node along the chain of systemic propagations is small, then the overall transmission of systemic influences on the optimal strategies is attenuated.

5 Calibration and Empirical Analysis

We develop an empirical analysis based on stock and CDS data of companies listed in the Dow Jones Industrial Average 30 (DJIA) as of December 31, 2007. For each company, we extract the full term structure of CDS spreads and recovery rates from Markit. We take the historical time series of its stock price from the CRSP database and adjust for corporate actions happened during the period (splits, dividends, etc.). We obtain short-term, long-term debt, and number of common outstanding shares of each firm from quarterly Compustat files. We estimate stock drifts and volatilities from an out-of-sample two-year period 2006–2007. We use the correlation model for market risk factors given in Remark 2.1, and set the exposure of each company to the systematic risk factor to $\rho = 0.3372$. Driessen et al. (2013) find that the average realized equity return correlation for the DJIA, computed over the period 2002–2007 using an estimation window of 91 calendar days, is 0.3372 (see Table 6 therein).

Our investor assumes the subjective default intensity of company n at time t to depend on four main covariates: (1) $w_1^n(t)$: the firm's distance to default, (2) $w_2^n(t)$: the firm's trailing one-year stock return, (3) $w_3(t)$: the three-month Treasury bill rate, and (4) $w_4(t)$: the trailing one-year return on the DJIA. Duffie et al. (2009) find that such a model underestimates empirical default correlations. Moreover, they show that when an additional channel of correlation through frailties is added, the model provides a good fit to empirically observed sample default correlations. We have verified that the inclusion of frailties in our model only changes the size, but not the proportions, of systemic risk contributors. Since the latter is the key focus of our analysis, we opt for a simpler mode without dependence on different frailty paths.[7] We follow Duffie et al. (2007) and consider the following model for the default intensity of firm n:

$$h_n = \mathrm{e}^{K+\beta_1 w_1^n+\beta_2 w_2^n+\beta_3 w_3+\beta_4 w_4}.$$

Notice that balance sheet and interest rates variables can be incorporated into the factor $\mathbf{F}(t)$ so that the model is consistent with the theoretical specification given in Section 2.

Remark 5.1. *We use the methodology proposed by Vassalou and Xing (2002), see also Duffie et al. (2007) to estimate the distance to default. For each firm, we estimate short-term debt as the larger of the*

[7] The theoretical framework developed in the previous sections can accommodate the specification of a stochastic frailty process as a component of $\mathbf{F}(t)$. However, we would then need to analyze systemic risk separately for each frailty path, resulting in additional complexity without qualitatively changing the conclusions.

Compustat items "Debt in Current Liabilities" and "Current Liabilities", while the long term debt is taken from the Compustat item Long-Term Debt.[8] *The distance to default DF_t is defined as*

$$DF_t = \frac{\log\left(\frac{V_t}{L_t}\right) + \left(\mu_A - \frac{1}{2}\sigma_A^2\right) T}{\sigma_A \sqrt{T}},$$

where V_t is the market value of the firm's asset at t, and L_t is the default point estimated following the standard established by Moody's KMV, see Crosbie and Bohn (2002), and given by the short term debt plus one half of the long term debt at t. The time horizon T is chosen to be four quarters as in Duffie et al. (2007). We estimate $V_t = L_t + E_t$, where E_t is obtained by multiplying the price of the stock at time t by the Compustat item "Number of Common Shares Outstanding". Moreover, μ_A and σ_A are respectively the drift and the volatility of the asset value process. We estimate them using an out-of-sample two-year period 2006–2007 of asset values constructed as detailed above.

We use the maximum likelihood estimates reported in Table III of Duffie et al. (2009) given by: $K = -2.093$, $\beta_1 = -1.2$, $\beta_2 = -0.681$, $\beta_3 = -0.106$, $\beta_4 = 1.481$. These are the maximum likelihood estimates obtained by Duffie et al. (2009) from a data set containing 402,434 firm-months of data between January 1979 and March 2004. Since the model is used to decide on investment strategies in the year 2008 and given that we do not expect major differences if we included additional three years in the calibration window (i.e., if it went from 1979 to 2007), these estimates are satisfying for the systemic risk analysis conducted in this paper.

When an investment decision is made at t, the investor computes the quantities w_i^n's using stock prices at time t, balance sheet information released in the quarter immediately preceding time t (for example, if t is April 20, balance sheet data from March 31 is used), and interest rates at t.

The rest of the section is organized as follows. Section 5.1 gives the calibration algorithm for our default intensity model under the risk-neutral measure. We discuss experimental results obtained from the calibrated portfolio model in Section 5.2.

[8] For cases with missing debt data, if the missing value corresponds to year Y, quarter Q, then we take the closest observation preceding quarter Q in year Y to complete the data set.

5.1 A New Calibration Algorithm for the Systemic Risk Model

This section describes a novel algorithm for calibrating the risk-neutral default intensities taking systemic effects into consideration. We first specify the parametric model for the risk-neutral default intensity $\lambda_n(t, \mathbf{S}(t))$. We choose the following parsimonious specification, which is able to capture empirically observed features of defaults while allowing for an efficient calibration:

$$\lambda_n(t, \mathbf{S}(t)) = \alpha_n + \underbrace{\beta_n / S_n(t)}_{\text{idiosyncratic}} + \underbrace{\gamma_n \, \#\{j : S_j(t) = 0\}}_{\text{systemic}},$$

where α_n, β_n and γ_n are constants, and $\#\{j : S_j(t) = 0\}$ is the number of companies that have defaulted by time t. The idiosyncratic component $\beta_n / S_n(t)$ relates the default intensity of company n to the inverse of its stock price (the lower the stock price, the higher its default intensity). In the context of single name equity-credit joint valuation, a similar modeling choice has been adopted by Linetsky (2006) who chose a default intensity specification $\lambda_n(t, \mathbf{S}(t)) = \gamma (S_n(t))^{-p}$ for constants $\gamma, p > 0$. It also has been used for pricing convertible bonds, see for instance Duffie and Singleton (2001), Chapter 9.3 therein, where p is set to one as in our case. The systemic component $\gamma_n \, \#\{j : S_j(t) = 0\}$ captures feedback effects from defaults. The default intensity of company n increases by the factor γ_n if another company j defaults.

We estimate the parameters α_n, β_n and γ_n by calculating CDS prices implied by our default intensity specification and matching them to empirically observed data for 1-, 5-, 7- and 10-year maturities. The model implied CDS prices are computed by numerically simulating future paths of stock prices under the risk-neutral probability measure. We choose 1-, 5-, 7- and 10-year maturities for two main reasons. Firstly, CDSs with these maturities are the most liquid. Secondly, CDSs with longer maturities contain more information on the long-term credit risk, which is helpful for a good parameter estimation. In fact, Figure 2 indicates that our calibration procedure leads to low errors across all maturities. The optimal parameters α_n, β_n and γ_n are those minimizing the sum of the squared errors between model implied prices and observed data. Notice that a direct estimation is computationally intractable because one would need to solve a coupled minimization problem over ninety variables (α_n, β_n and γ_n for each of the 30 companies). Hence, we propose a more suitable approach where calibration is performed via a sequence of iterative steps, as described next:

- Step 1. We assume the absence of systemic risk (all γ_n's are zero) and estimate α_n and β_n. This leads to solving thirty separate optimization problems over two variables each, which are computationally tractable. We simulate stock prices on a discrete time grid, using an Euler scheme. While we will elaborate on how to make a suitable choice for the number of sample paths in Remark 5.2 below, a natural choice for the time step is one trading day so that we have a total of 2,520 steps over ten years, with 252 trading days per year. For every time step and stock, we compare independent pseudo-random variables uniformly distributed on $(0, 1)$ with the current value of the default intensity in order to determine the occurrence of default events.

- Step 2. For each $n = 1, \ldots, 30$, we solve an optimization problem over three variables α_n, β_n and γ_n as described next. For $j \neq n$, we use the values of α_j and β_j estimated in the first step, set $\gamma_j = 0$, and numerically simulate the stock price dynamics of stock j. The simulation of paths is done analogously to the first step. Notice that the stock price of company j, $j \neq n$, will affect the default intensity of company n, which in turn impacts the stock price dynamics of company n. We treat α_n, β_n and γ_n as free parameters and choose them so to minimize the squared error between model implied and empirically observed CDS prices.

- Step K, $K \geq 3$. For each $n = 1, \ldots, 30$, we solve an optimization problem over three variables α_n, β_n and γ_n as described next. For $j \neq n$, we use the values of α_j, β_j and γ_j estimated in step $K - 1$ and numerically simulate the stock price dynamics of stock j. The simulation of paths is done analogously to the first step. We then choose α_n, β_n and γ_n so to minimize the squared error between model implied and market prices of CDS.

Each of the above steps is repeated for each of the twenty trading days in the month of December 2007. We then take the average computed over these trading days as the final estimates for α_n, β_n and γ_n. We do this for each of the thirty companies in the index.

Remark 5.2. *The calibration procedure is very computationally intensive, but can be executed on modern personal computers. A procedure to determine a suitable number of paths is first needed to estimate the average time per iteration. Because the above described procedure is applied separately for each of the twenty trading days in December 2007, we can estimate the average time per iteration step based on a few trading days. We do not expect significant differences between close trading days. Using MATLAB* *on a personal computer with 2.2 GHz dual-core processors (Intel Core i7), we find estimated times of 22 min, 50 min and 69 min per iteration step and trading when considering 1,000, 2,000 and 3,000 sample paths per stock, respectively, while the computing time explodes for higher numbers of paths. Therefore, a suitable number of sample paths for a feasible implementation is 3,000. A further increase would reduce the number of needed iteration steps, but such a reduction would not outweigh the additional computational time. Moreover, the number of needed iteration steps cannot be easily estimated without running the full calibration and depends on the chosen calibration period.*

In order to reduce the number of steps needed to achieve a specified tolerance level for the approximation error, one might think of making the number of paths depending on the specific firm: more paths for firms with high default risk and fewer paths for others. However, note that for firms with low default risk, the algorithm converges faster automatically so that the additional steps require only little time. The reason is that in each step, we use as starting values the optimizers from the previous step. If the firm has low risk, the optimizers from the previous step will be close to the new optimizers so that each additional step will only take little time in the case of such a firm.

The above calibration procedure consists of a series of refinements to the systemic risk model, starting from the absence of systemic risk and stopping when the difference between the parameter values of two subsequent iteration steps is small enough. We choose 0.01 as the tolerance level for the sum of absolute differences in all parameters. Notice that a tolerance level for the sum is stronger than the same tolerance level for the maximum, given that we are considering here *absolute* differences. For completeness, we report both maxima and sums in Table 1, which shows the differences in parameter

	step 2 vs. step 1	step 3 vs. step 2	step 4 vs. step 3	step 5 vs. step 4	step 6 vs. step 5	step 7 vs. step 6	step 8 vs. step 7	step 9 vs. step 8	step 10 vs. step 9
maximum over different companies of absolute difference in parameters									
in α_n	0.0010	0.0008	0.0006	0.0002	0.0004	0.0001	0.0001	0.0002	0.0000
in β_n	0.1098	0.0668	0.0319	0.0121	0.0273	0.0045	0.0057	0.0073	0.0031
in γ_n	0.0063	0.0036	0.0044	0.0005	0.0003	0.0002	0.0004	0.0001	0.0002
sum over different companies of absolute difference in parameters									
in α_n	0.0116	0.0039	0.0030	0.0008	0.0012	0.0003	0.0003	0.0003	0.0001
in β_n	0.7048	0.2585	0.1838	0.0517	0.0641	0.0179	0.0228	0.0150	0.0060
in γ_n	0.0408	0.0087	0.0078	0.0012	0.0008	0.0008	0.0009	0.0003	0.0003
total	0.7573	0.2711	0.1946	0.0538	0.0661	0.0190	0.0241	0.0156	0.0063

Table 1: We first take the maximum or the sum, over all companies, of the absolute differences in parameter estimates from two consecutive steps of the calibration procedure, for each trading day in December 2007. We then take the average over these trading days. We report these values in the upper (lower) part of the table for the maximum (sum).

estimates from two consecutive steps of this procedure. These differences decrease fast, hence suggesting that the calibration procedure is converging quickly and providing a good fit to observed CDS and stock prices data. Since our tolerance level is achieved at step 10, we take the parameters estimates from the tenth step of the procedure. These are displayed in Table 2. For most of the companies, the idiosyncratic and systemic components of the default intensity are statistically significant, suggesting that model is able to capture systemic dependencies. Figure 2 reports the empirical distributions of pricing errors across

	α_n	β_n	γ_n		α_n	β_n	γ_n
AA	0.0002	0.0421**	0.0059***	HD	0.0046***	0.0517***	0.0009**
AXP	0.0005*	0.2456***	0.0019***	HON	0.0020***	0.0527***	0.0007**
AIG	0.0035***	0.1587***	−0.0007	INTC	−0.0002	0.0277***	0.0015***
BA	−0.0003	0.2868***	0.0006***	IBM	0.0042***	0.0075*	0.0001**
ATTINC	0.0045***	0.0562***	−0.0002	JNJ	0.0008***	0.0001***	0.0016***
MO	−0.0003	0.1847***	0.0027***	JPM	0.0012***	0.0496***	0.0002***
CAT	0.0030***	0.0755***	0.0013***	MCD	0.0013***	0.0873***	0.0010***
C	0.0009***	0.0943***	0.0026***	MRK	0.0001	0.0626***	0.0020***
KO	0.0006**	0.0853***	0.0006***	MSFT	0.0000	0.0002	0.0025***
DIS	0.0006*	0.0801***	0.0034***	PG	0.0018***	0.0442**	0.0005***
DD	0.0020***	0.0005	0.0022***	PFE	0.0016***	0.0362***	0.0004***
GE	0.0039***	0.0460***	−0.0008	MMM	0.0012***	0.0183*	0.0020***
GM	0.0469***	0.1636***	0.0155*	UTX	0.0037***	0.0325*	0.0015***
XOM	0.0003**	0.0246***	0.0004*	VZ	0.0037***	0.0054	−0.0002
HPQ	0.0002**	0.0203***	0.0050***	WMT	0.0004***	0.0475***	0.0004*

Table 2: Parameter estimates from the fifth step of the calibration procedure. The asterisks *, ** and *** indicate statistical significance at 5 %, 1 % and 0.1 % levels of the estimations compared to one-sided hypotheses of having zero values, using an appropriately scaled asymptotic normal distribution. The data is from different trading days over which the calibration is run.

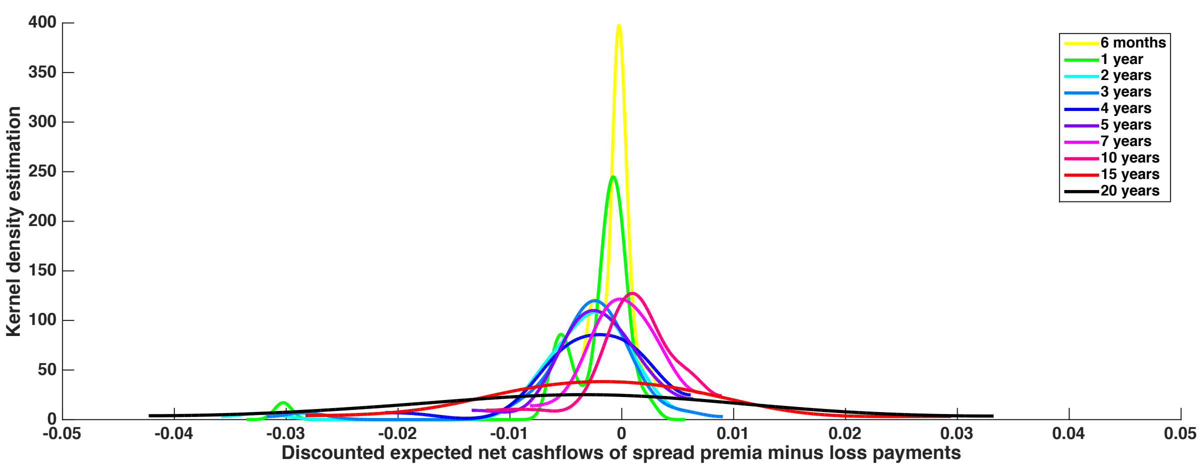

Figure 2: Estimated kernel density of the pricing errors across different maturities. It is based on CDS data of all 30 DJIA companies during December 2007. The symmetric distributions indicate the absence of any systematic pricing bias across different maturities.

different maturities. As expected, the errors are smaller for shorter maturities, but interestingly, the estimated distributions for different maturities are similar and only have different standard deviations. This indicates that the estimated model does not exhibit significant systematic pricing biases.

5.2 Empirical Results

We use the model calibrated using the algorithm described in the previous section, and analyze the optimal investment strategies on CDSs and equity. We consider CDSs with five-year maturity given that these are the most actively traded and liquid instruments.

Our main goal is to quantify the impact of systemic influences and their evolution over time based on the analysis done in Section 4. We find that companies having high systemic influences *on* others are not necessarily those which are highly sensitive to systemic influences *of* others. This can be understood by looking at the bars in the top panel of Figure 3. For each company j, the bar can be split into 29 terms contributing to the aggregate systemic influence $\sum_{n \neq j} \psi^{(n)}$ on the CDS strategy referencing the j-th company.[9] Here, we report the contribution (in absolute values) of the five most significant terms. These correspond to the companies C (Citigroup), AIG, GM (General Motors Company), MO (Altria Group), and AA (Alcoa). The largest contribution comes from C followed by AIG, whereas the influences of the other four are smaller. While the two companies with the biggest systemic influences belong to the financial sector, the most affected by these influences are GM, AIG, AA and GE (General Electric), whereof three of the four are non-financial companies. This can be explained as follows. Firstly, GM and AIG experienced financial distress in 2008. In the case of GM, it was largely caused by its financial services division GMAC which was bailed out by the U.S. Government. Secondly, the car manufacturing and aluminium businesses of GM and AA are very cyclical and capital intensive, hence depending on the broad financial markets conditions. Thirdly, GE has a big financial division called GE Capital, which focuses on loans and leases, making it more vulnerable to systemic influences. If we consider CDS premia as of December 30, 2008, we observe that GM has the highest 5-year CDS spread (0.955), followed by AIG (0.053) and Citigroup (0.019). Despite the relatively low credit spreads, Citigroup is systemically very influential. This is because in our model, the investor's subjective belief of default intensity of Citigroup is higher than what is perceived by the market. As a consequence, the investor purchases high volumes

[9] In order to get comparable values across different companies, we multiply this quantity with the 5-year CDS spread of the reference company j.

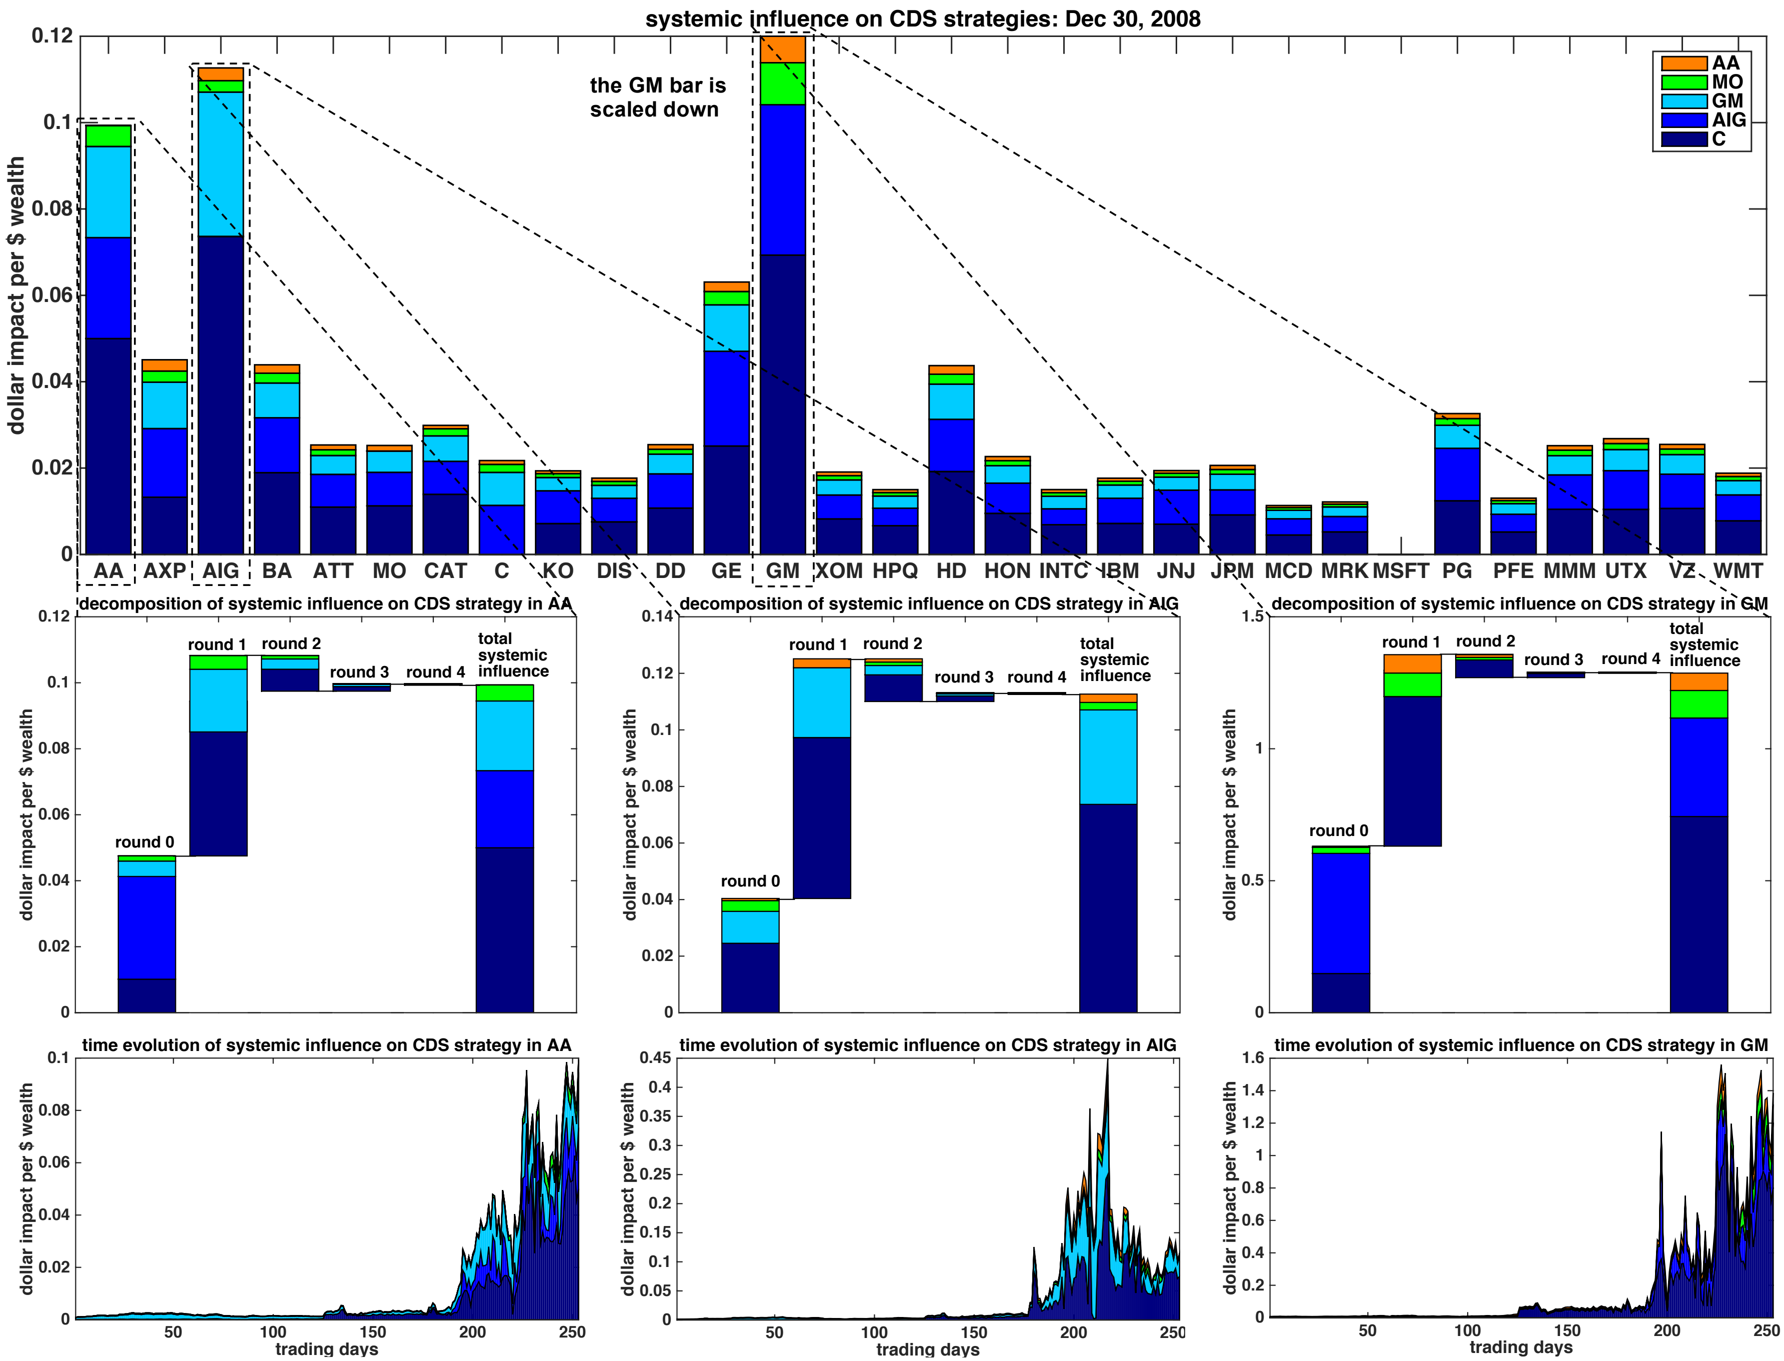

Figure 3: Systemic influence on CDS strategies: the main receivers (three highest bars in top panel) are GM, AIG and AA, while the main contributions (split of columns in top panel) come from C, AIG and GM. Middle: rounds of systemic propagation. Bottom: time evolution in 2008.

of credit protection on Citigroup, and adjusts his investment strategy in the CDSs referencing other companies because of the dependence structure of risk factors. To analyze systemic influences in more detail, we focus on the three systemically most affected companies, namely AA, AIG and GM and use the decomposition formula given in (4.8). While the sizes of their sensitivities to systemic influences are quite different, the middle panels of Figure 3 show that the qualitative behavior is similar for the three companies.[10] However, a change in the default risk premium of AIG leads to mainly direct adjustments in the CDS strategies in AA and GM, resulting from market risk correlation (round 0 in (4.8)). By contrast, several rounds of adjustments are needed if there is a change in the default risk premium of Citigroup.

Although portfolio strategies may change on a daily basis, the companies with the highest systemic influences (Citigroup, AIG and GM) are persistent over time, as shown in the bottom panels of Figure 3. Therein, as expected, we can also see that systemic influences were high in the fourth quarter of 2008, and negligible earlier in the year when the market was relatively more stable.

The investor is most active on the CDS market during the distressed last quarter of 2008. As it appears from Figure 4, he shifts from selling small amounts of protection on the five systemically most influential companies (in the first quarter) to purchasing high amounts of protection on these companies in the last quarter of 2008. Short CDS positions are especially high on AIG and C (Citigroup). This

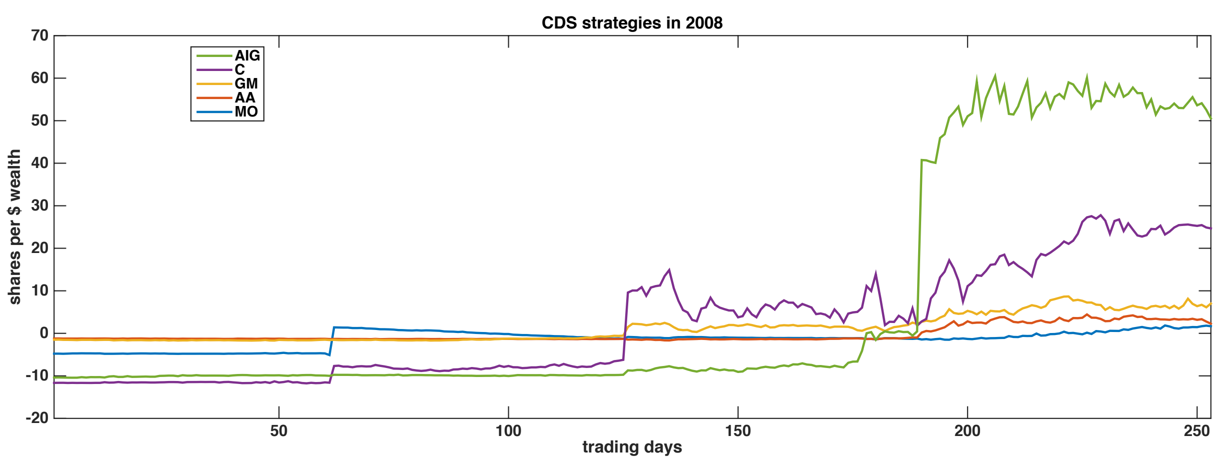

Figure 4: The optimal CDS investment strategies are computed using our calibrated model and the explicit expression given in Theorem 4.2. The jumps are due to changes in the subjective default intensities, as a result of quarterly updated balance sheet data. The investor's position in CDSs is highly subject to systemic influences which increase dramatically in the third and fourth quarter of 2008.

[10] We have verified that the eigenvalues of each contagion matrix used in the middle panel are strictly less than one in modulus, hence the decomposition formula is meaningful.

can be partly explained by the fact that AIG faced a liquidity crisis in late 2008 leading to the largest government bailout of a private company in U.S. history, while Citigroup suffered huge losses and was rescued in November 2008 via a massive stimulus package by the U.S. Government. Figure 4 also indicates how the portfolio changes from the end of a quarter to the beginning of the next quarter. These changes correspond to the jumps in the optimal investment strategies. We remark that when the investment strategy switches from a short to a long position, this is economically equivalent to clearing the position at quarter's end and taking a long position at the beginning of the new quarter.

Our analysis provides useful guidance for risk monitoring of mixed equity-credit portfolios. The investor can identify which companies should be subject to active monitoring using the systemic risk analysis developed in Section 4. This in turn has important risk management implications. The investor will need to rebalance frequently his entire portfolio when there are changes in credit quality of companies with high systemic influence. By exploiting our provided decomposition formulas, he will also be able to discriminate which positions are more sensitive to systemic influences and hence implement suitable precautionary measures. On the contrary, he will only need to adjust few portfolio positions when there are changes in credit quality of companies which are not systemically influent.

6 Conclusion

We have developed an equity-credit portfolio optimization framework taking into account the full interaction of market and credit risk, with special focus on the systemic dependencies. Despite the intricate interrelations of risk factors, we have provided explicit expressions for the optimal investment strategies in stocks and CDSs. We have analyzed the impact of systemic influences on the optimal strategies and presented a case study which illustrates precisely how subsequent rounds of systemic propagations contribute to optimal investment decisions. We have provided an iterative calibration procedure which progressively refines the systemic component of our model to incorporate information conveyed by historical time series of equity and CDS data. We have used the calibrated model to compute the optimal investment strategies during the year 2008. Our analysis reveals the presence of a small number of (mostly financial) companies with high systemic influences, as well as of a different set of companies which are highly sensitive to systemic influences. These findings indicate that systemic risk should be an integral component of portfolio monitoring, considering that significant loading/unloading of portfolio positions may be triggered by changes in the credit quality of systemically influential companies.

Acknowledgment

We would like to thank the Department Editor, Noah Gans, and two anonymous referees for valuable comments helping us improve the paper. Christoph Frei gratefully acknowledges financial support by the Natural Sciences and Engineering Council of Canada under grant RGPIN/402585-2011.

References

E. Anderson, L. Hansen, and T. Sargent (2000): Robustness, detection and the price of risk. *Working paper*, University of North Carolina, University of Chicago and Stanford University.

D. Becherer and M. Schweizer (2005): Classic solutions to reaction-diffusion systems for hedging problems with interacting Itô and point processes. *Annals of Applied Probability*, 15, 1111–1144.

A. Berndt, R. Douglas, D. Duffie, M. Ferguson, and D. Schranz (2005): Measuring default-risk premia from default swap rates and EDFs. *Working paper*, Graduate School of Business, Stanford University.

F. Biagini (2010): Second fundamental theorem of asset pricing. *Encyclopedia of Quantitative Finance*, editor: R. Cont, Wiley, 1623–1628.

T. Bielecki, S. Crepey, M. Jeanblanc, and M. Rutkowski (2010): Convertible bonds in a defaultable diffusion model. *Stochastic Analysis with Financial Applications*, editors: A. Kohatsu-Higa, N. Privault and S.J. Sheu, Birkhäuser Verlag, 255–298.

T. Bielecki, M. Jeanblanc, and M. Rutkowski (2008): Pricing and trading credit default swaps in a hazard process model. *Annals of Applied Probability*, 18, 2495–2529.

Bloomberg News, November 8, 2013: Report available at http://www.bloomberg.com/news/2013-11-08/pimco-said-to-wager-10-billion-in-default-swaps-credit-markets.html

L. Bo and A. Capponi (2014): Optimal investment in credit derivatives portfolio under contagion risk. *Mathematical Finance*, Forthcoming.

M. Brandt (2010): Portfolio choice problems. *Handbook of Financial Econometrics, Vol. 1: Tools and Techniques*, editors: Y. Aït-Sahalia and L. Hansen, North Holland, 269–336.

P. Crosbie and J. Bohn (2002): Modeling default risk. *Technical Report KMV*, LLC.

J. Cvitanić (2001): Theory of Portfolio Optimization in Markets with Frictions. *Handbook of Mathematical Finance*, Cambridge: Cambridge University Press.

J. Driessen, P. Maenhout and G. Vilkov (2013): Option-implied correlations and the price of correlation risk. *Working paper*, University of Amsterdam and INSEAD.

D. Duffie, A. Eckner, G. Horel, and L. Saita (2009): Frailty correlated default. *Journal of Finance*, 64, 2089–2123.

D. Duffie, L. Saita, and K. Wang (2007): Multi-period corporate default prediction with stochastic covariates. *Journal of Financial Economics*, 83, 635–665.

D. Duffie and K. Singleton (2001): Credit Risk: Pricing, Measurement, and Management. *Princeton Series in Finance.*

L. Eisenberg and T. Noe (2001): Systemic risk in financial systems. *Management Science*, 47, 236–249.

K. Giesecke, B. Kim, J. Kim, and G. Tsoukalas (2014): Optimal credit swap portfolios. *Management Science*, 60, 2291–2307

M. Gordon, G. Paradis, and C. Rorke (1972): Experimental evidence on alternative portfolio decision rules. *American Economic Review*, 62, 107–118.

L. Hansen, T. Sargent, G. Turmuhambetova, and N. Williams (2006): Robust control and model misspecification. *Journal of Economic Theory*, 128, 45–90.

ISDA News Release (2009): ISDA announces successful implementation of Big Bang? CDS protocol; determinations committees and auction settlement changes take effect. Available at http://www.isda.org/press/press040809.html.

H. Kraft and M. Steffensen (2009): Asset allocation with contagion and explicit bankruptcy procedures. *Journal of Mathematical Economics*, 20, 147–167.

H. Kunita (1984): Stochastic Differential Equations and Stochastic Flows of Diffeomorphism, *École d'Été de Probabilités de Saint-Flour XII–1982, Lecture Notes in Mathematics 1097*, Springer, 143–303.

R. Jarrow and F. Yu (2001): Counterparty risk and the pricing of defaultable securities. *Journal of Finance*, 56, 1765–1799.

D. Lando (1998): On Cox processes and credit risky securities. *Review of Derivatives Research*, 2, 99–120.

V. Linetsky (2006): Pricing equity derivatives subject to bankruptcy. *Mathematical Finance*, 16, 255–282.

J. Liu (2007): Portfolio selection in stochastic environments. *Review of Financial Studies*, 20, 1–39.

R. Mendoza-Arriaga, P. Carr, and V. Linetsky (2010): Time changed Markov processes in credit-equity modeling. *Mathematical Finance*, 20, 527–569.

R. Mendoza-Arriaga and V. Linetsky (2014): Multivariate subordination of Markov processes with financial applications. *Mathematical Finance*. Forthcoming.

R. Merton (1969): Lifetime portfolio selection under uncertainty: the continuous-time case. *Review of Economics and Statistics*, 51, 247–257.

A. Perold (1984): Large-scale portfolio optimization. *Management Science*, 30, 1143–1160.

P. Protter (2005): Stochastic Integration and Differential Equations. Second Edition, Volume 21 of *Stochastic Modelling and Applied Probability*, Springer, New York.

M. Vassalou and Y. Xing (2002): Default risk in equity returns. *Journal of Finance*, 59, 831–868.

M.B. Wise and V. Bhansali (2002): Portfolio allocation to corporate bonds with correlated defaults. *Journal of Risk*, 5, 39–58.

F. Yu (2005): Default correlation in reduced-form models. *Journal of Investment Management*, 3, 33–42.