

Optimal Credit Swap Portfolios

Kay Giesecke* Baeho Kim[†] Jack Kim[‡] Gerry Tsoukalas[§]

June 29, 2009; this draft December 16, 2013[¶]

Management Science, forthcoming

Abstract

This paper formulates and solves the selection problem for a portfolio of credit swaps. The problem is cast as a goal program that entails a constrained optimization of preference-weighted moments of the portfolio value at the investment horizon. The portfolio value takes account of the exact timing of protection premium and default loss payments, as well as any mark-to-market profits and losses realized at the horizon. The constraints address collateral and solvency requirements, initial capital, position limits, and other trading constraints that credit swap investors often face in practice. The multi-moment formulation accommodates the complex distribution of the portfolio value, which is a nested expectation under risk-neutral and actual probabilities. It also generates computational tractability. Numerical results illustrate the features of optimal portfolios. In particular, we find that credit swap investment constraints can have a significant impact on optimal portfolios, even for simple investment objectives. Our problem formulation and solution approach extend to corporate bond portfolios and mixed portfolios of corporate bonds and credit derivatives.

*Department of Management Science & Engineering, Stanford University, Stanford, CA 94305-4026, email: giesecke@stanford.edu, web: www.stanford.edu/~giesecke.

[†]Korea University Business School, Seoul 136-701, Korea, email: baehokim@korea.ac.kr, web: <http://biz.korea.ac.kr/~baehokim/>.

[‡]JP Morgan, New York, email: jh.jackkim@gmail.com.

[§]The Wharton School, University of Pennsylvania, 3730 Walnut Street, 555 Jon M. Huntsman Hall, Philadelphia, PA 19104, email: gtsouk@wharton.upenn.edu, web: <https://opimweb.wharton.upenn.edu/profile/21137/>.

[¶]This paper was previously entitled “Fixed-Income Portfolio Selection.” We thank Anil Bangia, Vineer Bhansali, John Birge, Agostino Capponi, Andreas Eckner, Taosong He, George Lentzas, Pete Meindl, Hideyuki Takada and seminar participants at the 2009 American Mathematical Society Joint Mathematics Meeting, the 2008 INFORMS Annual Meeting, the Stanford Financial Mathematics Seminar, JP Morgan, and PIMCO for comments.

1 Introduction

The equity portfolio selection problem is treated in a substantial literature.¹ Despite its significance in practice, the problem of selecting a fixed-income portfolio of credit instruments such as corporate bonds and credit default swaps, is less well understood. This may be due to the unique cash flow and distributional features of credit instruments, which generate significant issues for the selection problem that are absent in the equity case. Importantly, the risk of default generates skewed return distributions. The clustering of defaults strongly influences the distribution of the portfolio return. Credit derivatives, which are routinely used by fixed-income investors to efficiently express investment views, involve additional issues. An investor can, for instance, implement a short credit position by buying protection through a credit swap. In this case, the investor pays a premium and receives compensation for losses if the underlying issuer defaults. A long credit position can be implemented by selling protection. Then, the investor receives a premium for promising to cover the loss due to default. There is yet another aspect. Suppose the credit quality of the issuer improves after the swap was entered into. The investor, who has sold protection for a fixed premium, can now realize a mark-to-market profit by buying protection on the issuer over the remaining term. Then, the swap is effectively closed out at a mark-to-market profit proportional to the difference between the initial premium and the prevailing market premium. This profit represents the market value of the swap to the protection seller. It is income in addition to the premiums already received.

This paper formulates and solves the selection problem for a portfolio of credit swaps. To address the cash flow implications of credit swap positions, we consider the portfolio's market value at the investment horizon. This value takes account of the exact timing of protection premium and default loss payments, as well as any mark-to-market profits and losses realized at the horizon. The selection problem is cast as a goal program (GP) that involves a two-stage constrained optimization of preference-weighted moments of the portfolio value. The decision variable is the vector of swap notional values. In the initial stage we perform the individual odd moment maximizations and even moment minimizations. In the second stage, we balance the multiple objectives by optimizing the sum of weighted distances from the individual moment optima. The constraints of the problem address collateral and solvency requirements, initial capital, position limits, and other trading constraints that credit swap investors often face in practice. Although we focus on credit swap portfolios, our problem formulation and solution approach extend to corporate bond portfolios and mixed portfolios of corporate bonds and credit derivatives.

¹The literature was pioneered by Markowitz (1952). Important developments include the treatment of parameter estimation errors as in Goldfarb & Iyengar (2003) and DeMiguel & Nogales (2009), transaction costs as in Perold (1984) and Lobo, Fazel & Boyd (2007), trading constraints as in Bonami & Lejeune (2009), higher moments as in Athayde & Flores (2004), and tail risk measures such as the lower partial moment (Jarrow & Zhao (2006)) or the value at risk (El Ghaoui, Oks & Oustry (2002)).

The GP formulation of the credit swap portfolio selection problem has important advantages over an alternative expected utility formulation. Not only can it address the strongly non-Gaussian features of the distribution of the portfolio value, but the multi-moment objective has also significant computational benefits. Although the full distribution of the portfolio value is generally intractable, the moments turn out to be less so. Analytical expressions for the moments can be obtained for a broad class of reduced-form, doubly-stochastic models of correlated default timing. Such models have been used for the pricing of credit derivatives (see Duffie & Garleanu (2001), Eckner (2009), Mortensen (2006), Papageorgiou & Sircar (2007) and others), as well as for predicting default events (see Chava & Jarrow (2004), Duffie, Saita & Wang (2006), and others).

The constraints that credit swap investors face have a significant impact on optimal portfolios. For example, collateral requirements and upfront protection premium payments impose a basic capital constraint. Position limits put restrictions on the swap notional that can be traded. Finally, the investor must be able to meet all payment obligations arising from the strategy after the swap positions were entered into. Our numerical results illustrate that these constraints result in non-trivial strategies even for simple single-moment optimization objectives. They also indicate that the optimal portfolios are relatively robust with respect to changes in the parameters of the default timing model.

This paper is the first to address the selection problem for a portfolio of credit derivatives. Prior research treats the corporate bond portfolio problem using expected utility or other objectives, including Akutsu, Kijima & Komoribayashi (2004), Kraft & Steffensen (2008), Meindl & Primbs (2006), and Wise & Bhansali (2002). These papers do not consider credit derivative positions and the mark-to-market valuation of portfolio positions. When accounting for mark-to-market profits and losses, the distribution of the portfolio value takes a complex form. It is a nested expectation under different probability measures. The inner expectation represents the market value of the portfolio at the horizon, and is taken under a risk-neutral pricing measure. The outer expectation is taken under the measure that describes the actual distribution of the risk factors governing the market value. A dynamic model of correlated default timing under actual and risk-neutral probabilities is required to address the portfolio problem.

The rest of this article is organized as follows. Section 2 discusses our basic assumptions and characterizes the credit swap portfolio value. Section 3 formulates investment constraints. Section 4 casts the selection problem as a GP over multiple moments of the portfolio value. Section 5 illustrates the implementation of the GP using a standard model of correlated default timing. Section 6 provides numerical results. Section 7 concludes. There are two technical appendices.

2 Strategy, capital, and portfolio value

2.1 Strategy and initial capital

We consider an investor selecting a portfolio of n credit swaps. A swap references a “name” i representing a firm, municipality, or sovereign country, on which default protection is being bought or sold. In a protection selling position, the investor covers the loss due to default, which is given by the product of the swap notional and the loss rate. In a protection buying position, the investor pays a premium for obtaining coverage against default. The swap expires at the time of default or the maturity date, whichever is earlier.

The protection premium is negotiated at contract inception (time 0). It has two parts: an upfront payment given by the product of the swap notional and the upfront rate $U_i \in [-1, 1]$, and a running payment given by the product of the swap notional and the swap spread S_i . A positive (negative) upfront rate represents a cash inflow (outflow) to the protection seller. The running payments occur every quarter, at dates (t_m) , where the final premium date is equal to the swap maturity date T .² We state the spread S_i on a per-quarter basis and ignore day counts and premium accruals. We assume that the loss rate ℓ_i is constant.

The investor intends to follow a “buy and hold” strategy for a fixed investment horizon $H < T$. At time 0, the investor designs the portfolio by selecting the vector $\delta = (\delta_1, \dots, \delta_n) \in \mathbb{R}^n$ of signed swap notional values. A positive δ_i represents a protection selling position with notional value δ_i , while a negative δ_i indicates a protection buying position with notational value $|\delta_i|$.

The investor starts with a fixed amount of capital C available for investment. The capital is initially invested in a risk-free money-market account that accrues interest at a constant, continuously compounding rate r . The funds in this account are used to cover payment obligations arising from the strategy. Any income generated by the strategy is invested in the account.

2.2 Portfolio value

The value at the horizon H of the investor’s portfolio is

$$\pi(\delta) = Ce^{rH} + \delta_1 P_1 + \dots + \delta_n P_n. \quad (1)$$

The first term represents the value of the risk-free account, and $\delta_i P_i$ represents the value of a swap position. Here, δ_i is the signed swap notional and P_i is the value at H of the cash flow generated by a swap with unit

²The assumption of a common maturity date for all swaps is made to maintain clarity in the exposition. It is straightforward to extend the analysis to include name-dependent maturity dates.

notional. Since a positive δ_i represents a protection selling position, we have

$$P_i = U_i e^{rH} + \sum_{t_m \leq H} e^{r(H-t_m)} S_i (1 - N_{t_m}^i) - \int_0^H \ell_i e^{r(H-s)} dN_s^i + V_i, \quad (2)$$

where N_t^i is the default indicator, which is 0 at any time t before default and 1 at any t after default.³ The first term in (2) represents the upfront premium payment at contract inception. The second term represents the running premium payments until time H . A premium payment S_i is contingent upon the survival to a payment date t_m , hence the “no-default” factor $(1 - N_{t_m}^i)$. The third term represents a one-time payment of the loss ℓ_i at the time of default, if default occurs before the horizon H . The fourth term represents the *mark-to-market value* of the swap at H , which is given by

$$V_i = \mathbb{E}^* \left(\sum_{t_m > H} e^{r(H-t_m)} S_i (1 - N_{t_m}^i) - \int_H^T \ell_i e^{r(H-s)} dN_s^i \mid \mathcal{F}_H \right), \quad (3)$$

where \mathbb{E}^* denotes the expectation under a risk-neutral measure \mathbb{P}^* .⁴ The mark-to-market value V_i is the difference between the value at H of any future premium payments and the value at H of a potential default payment of ℓ_i between H and T . It is 0 in case of default before H . The mark-to-market value is positive (corresponding to a profit) if the credit quality of the reference name has improved since contract inception. It is negative (corresponding to a loss) if the credit quality has deteriorated.

Due to the complex cash flow associated with a swap position, the distribution of the portfolio value (1) has significant non-Gaussian features. To illustrate, consider a protection selling position and a horizon H of 6 months. If the reference name defaults during the first quarter, then the investor pays the loss at default without receiving a spread payment. If the name defaults during the second quarter, the investor pays the loss after receiving a spread payment at the end of the first quarter. If the name survives to H , then the investor receives a spread payment at the end of the first quarter and at H . The investor also realizes a mark-to-market profit or loss at H . The distribution of the resulting position value $\delta_i P_i$ is asymmetric and multimodal. Each of the two default scenarios generates a mode in the left tail of the distribution. The no-default scenario generates a mode in the right tail. The distribution of the value (1) of a portfolio of selling and buying positions has complex features, which are not only driven by the characteristics of individual positions but also by the correlation between the positions.

³We fix a complete probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and an information filtration $(\mathcal{F}_t)_{t \geq 0}$ satisfying the usual conditions (see Protter (2004)). Here \mathbb{P} denotes the actual probability measure.

⁴The credit swap market is assumed to be free of arbitrage opportunities, so under mild technical conditions there exists an equivalent risk-neutral pricing measure. We fix a risk-neutral measure \mathbb{P}^* .

Our analysis ignores counterparty risk, the risk that a swap party fails to perform according to contract. This assumption appears reasonable because the credit swap market is increasingly moving towards central clearing, which alleviates counterparty risk for investors.⁵

Although we focus on a portfolio of credit swaps, the analysis of a portfolio of corporate bonds or other credit instruments is similar. A long bond position corresponds to a swap protection selling position (positive δ_i), and a short bond position corresponds to a protection buying position (negative δ_i). The market value V_i of a long bond position at H and the associated cash flow value P_i can be worked out using arguments similar to the ones applied above.

3 Investment constraints

Collateral requirements, position limits, and solvency requirements impose constraints on the swap investment strategy. This section formulates these constraints.

3.1 Initial capital constraint

At contract inception, a swap investor is required to post collateral (such as cash or other liquid securities) in a margin account, which we assume earns interest at the rate r . The collateral, which is given by the product of the collateral rate c_i and the swap notional $|\delta_i|$, reduces the exposure of the clearing house to the swap counterparties. The rate c_i is set by the clearing house, and may be a function of the notional, the nature of the position, and the expected position cash flows (such as scheduled premium payments).

The collateral requirement, along with the scheduled upfront payments, imposes a basic capital constraint on the investor. The collateral payments $|\delta_i|c_i$ and upfront premium payments $-\delta_i U_i$, summed over all positions, cannot exceed the capital C available to the investor at time 0:

$$\sum_{i=1}^n (|\delta_i|c_i - \delta_i U_i) \leq C. \quad (4)$$

The constraint (4) is analogous to a budget constraint that an equity investor faces. However, several unique issues distinguish the credit swap case. The first is that the constraint (4) does not guarantee the solvency of the investor during the investment period. To see this, consider a portfolio consisting of a single protection selling position. If the reference name defaults shortly after the swap was entered into, then even

⁵With central clearing, all trades are cleared centrally. A clearing house stands between a swap protection buyer and a seller, assuming counterparty risk on both sides of the trade. The original parties are thus insulated from each other and are only exposed to the clearing house. Risk reduction is achieved through the mutualization of losses: Should a trade party default, losses are shared among the clearing house members.

if (4) is met at time 0 the investor is not guaranteed to have sufficient funds at hand to cover the required loss payment. The constraint (4) does not address the possibility of insolvency.

The second issue is that the set of feasible notional values implied by (4) is not necessarily bounded for all values of c_i . To see this, consider again a portfolio consisting of a single protection selling position $\delta_1 > 0$, and suppose that $0 \leq c_1 < U_1$. In this case, (4) is equivalent to $-\frac{C}{U_1 - c_1} \leq \delta_1 \leq \infty$, so that δ_1 is unbounded from above. Allowing the investor to take arbitrarily large positions is not desirable or realistic. In practice, investors face additional constraints in the form of position limits.

3.2 Position limits

Position limits can either be imposed internally (by management) or externally (by the clearing house or the broader market). Internal position limits can serve multiple purposes. For example, investors may have diversification targets: they may want to have a minimum amount of exposure to a particular industry or they may want to limit their exposure to a particular name. At the portfolio level, investors may want to limit the total notional sizes of their positions. If an aggregate position becomes too large relative to the market, unwinding costs may dominate the investor's returns, possibly leading to catastrophic failures.⁶

In general, position limits can be asymmetric between protection selling and protection buying positions. To distinguish the two cases, we introduce two intervals for every name i . If the investor is selling protection (superscript s), then we require $0 \leq \delta_i \in [\underline{\delta}_i^s, \bar{\delta}_i^s]$. If the investor is buying protection (superscript b), then we require $0 \geq \delta_i \in [\underline{\delta}_i^b, \bar{\delta}_i^b]$. The interval bounds satisfy $\bar{\delta}_i^s \geq \underline{\delta}_i^s \geq 0 \geq \bar{\delta}_i^b \geq \underline{\delta}_i^b$. Similar constraints can be expressed at the portfolio level, on the aggregate protection selling and buying positions. Letting $\bar{\delta}^s \geq \underline{\delta}^s \geq 0 \geq \bar{\delta}^b \geq \underline{\delta}^b$ be the bounds for the selling and buying positions, the portfolio constraints are given by $\underline{\delta}^s \leq \sum_i \max(\delta_i, 0) \leq \bar{\delta}^s$ and $\underline{\delta}^b \leq \sum_i \min(\delta_i, 0) \leq \bar{\delta}^b$.

There may also be external position limits. In practice, investors may not be able to trade a continuum of notional values. Often, the notional values must be multiples of a given base value v_i , such as \$5 million. This can be addressed by requiring that $\delta_i \in \{x : x = v_i z, z \in \mathbb{Z}\}$. Moreover, minimum notional values may be required. In addition, feasible notional values are often bounded from above, because the market has finite size in any given name. Such bounds on δ take the form of internal limits.

⁶A recent example is the case of the “London Whale,” a credit trader at JP Morgan who accumulated excessively large swap positions in 2012. The large positions eventually had to be unwound, adversely impacting market prices and generating multi-billion dollar losses for the bank. These types of price impact effects have been extensively studied. For details, see Almgren & Chriss (2000), Bertsimas & Lo (1998), Perold (1984), Tsoukalas, Wang & Giesecke (2013), and references therein.

3.3 Solvency constraint

Although combining position limits with the initial capital constraint leads to more realistic feasible sets, it is not sufficient to ensure the investor's solvency, i.e., the ability to meet all payment obligations arising from the strategy after the positions were entered into. Assuming the investor cannot borrow additional capital, solvency is guaranteed if the present value at time 0 of the worst-case obligations associated with the strategy does not exceed $C + \sum_{i=1}^n \delta_i U_i$, the initial capital net of upfront premiums paid and received. The worst-case scenario is the situation where the names on which protection was sold (positive δ_i) default immediately after the positions were entered into, while the names on which protection was bought (negative δ_i) survive to H and carry virtually no default risk after that. That the names on which protection was bought carry no default risk after H means that the cost at H of default protection on these names is zero, and that the investor suffers the greatest possible loss in the mark-to-market value at H of these positions. Thus, in the worst-case scenario, the investor must make the default payments $\delta_i \ell_i$ on the protection selling positions, the premium payments $-\delta_i S_i$ at each $t_m \leq H$ on the protection buying positions, and cover the mark-to-market loss $-\sum_{t_m > H} \delta_i S_i e^{-r(t_m - H)}$ at H on the protection buying positions. Consequently, the present value at time 0 of the investors's worst-case obligations, denoted $g(\delta)$, takes the form

$$g(\delta) = \sum_{i=1}^n \max(\delta_i, 0) \ell_i - \sum_{i=1}^n \min(\delta_i, 0) \sum_{t_m \leq T} S_i e^{-rt_m}. \quad (5)$$

The first term above represents the present value of the default payments to be made by the investor on the protection selling positions. The second term represents the present value of the mark-to-market loss and the premium payments to be made by the investor on the protection buying positions. We require that the present value of the investor's worst-case obligations match the investor's initial capital C , adjusted by any upfront premium payments made and received:

$$C + \sum_{i=1}^n \delta_i U_i = g(\delta). \quad (6)$$

We adopt an equality constraint as opposed to an inequality constraint so that all the capital C committed to investment is fully exposed to risk.⁷ Note that the collateral rate c_i does not appear in (6) because any collateral pledged is returned (with interest) to the investor as long as the investor remains solvent.

The solvency constraint (6) can be used in conjunction with the initial capital and position limit con-

⁷On a more technical note, the equality constraint also precludes the second moment objective in Section 4 from attaining the trivial no risky investment, zero variance optimum – a case that is of no interest.

straints described above. The solvency constraint (6) implies the initial capital constraint (4) if the collateral rate c_i satisfies $c_i \leq \ell_i$ for all protection selling positions and $c_i \leq \sum_{t_m \leq T} S_i e^{-rt_m}$ for all protection buying positions (see Appendix A for a proof). The inequality $c_i \leq \ell_i$ states that the collateral rate imposed by the clearing house does not exceed the loss rate faced by the protection seller. The other inequality states that the collateral rate does not exceed the present value of all future spread rates faced by the buyer. While these bounds on the collateral rate are meaningful, they are not guaranteed to hold in practice. When setting c_i , the clearing house faces a trade-off: it must ensure that the rate is low enough to avoid impeding market liquidity, yet high enough to offer adequate protection against counterparty risk.

4 Optimizing the portfolio value

By selecting the vector of swap notional values δ , the investor seeks to optimize the value at the horizon H of the portfolio, which is given by $\pi(\delta) = Ce^{rH} + \delta^\top P$.⁸ Here, $P = (P_1, \dots, P_n)$. The set of admissible values of δ , denoted by Δ , is determined by the constraints discussed in Section 3.

The complex structure of $\pi(\delta)$ imposes restrictions on the choice of the optimization objective. The optimization of the expected utility of $\pi(\delta)$ appears difficult, because the full distribution of $\pi(\delta)$ generally required for this optimization is analytically intractable unless one imposes strong assumptions regarding default timing and position cash flows. To avoid an oversimplification, we propose to cast the selection problem as a goal program⁹ that involves a two-stage constrained optimization of preference-weighted moments of $\pi(\delta)$. Unlike the distribution, the moments of $\pi(\delta)$ turn out to be tractable for standard default timing models. The moments can also capture the non-Gaussian features of $\pi(\delta)$ highlighted in Section 2.2. This motivates a moment-based formulation of the portfolio problem.

The credit swap investor wishes to maximize the expected value of $\pi(\delta)$ while minimizing the variance of $\pi(\delta)$. At the same time, as noted in Kraus & Litzenberger (1983), Simkowitz & Beedles (1978), and Briec, Kerstens & Jokung (2007), the investor prefers positive skewness. Moreover, the investor prefers low kurtosis in order to avoid tail risk, i.e., a large contribution of extreme events to the variance. Generalizing, the investor seeks to maximize odd moments and minimize even moments of $\pi(\delta)$; see Scott & Horvath

⁸The optimization of the portfolio value is more natural than the optimization of the portfolio return $\frac{\pi(\delta) - C}{C}$, because the decision variable is the vector of swap notional values (rather than a vector of portfolio weights).

⁹Goal programming is a method by which one can address problems that involve multiple goals. The often conflicting nature of simultaneous objectives results in solution methods that offer a compromise between the optimality of individual goals, depending on the decision maker's preference for each objective. See Charnes, Cooper & Ferguson (1955) for an early application of goal programming and Jones & Tamiz (2010) for a recent overview of the method.

(1980) for a formal justification. This leads to a multi-objective optimization problem:

$$\min_{\delta \in \Delta} (-1)^k Z_k(\delta), \quad k \geq 1, \quad (7)$$

where $Z_k(\delta)$ is the mean of $\delta^\top P$ for $k = 1$, the variance for $k = 2$, and the k th standardized central moment for $k \geq 3$. Formally,

$$\begin{aligned} Z_1(\delta) &= \delta^\top M_1 \\ Z_2(\delta) &= \delta^\top M_2 \delta \\ Z_k(\delta) &= (\delta^\top M_k \bigotimes_{i=1}^{k-1} \delta) Z_2(\delta)^{-k/2}, \quad k \geq 3, \end{aligned}$$

where M_k is the mean vector of P for $k = 1$, the covariance matrix for $k = 2$, and the k -dimensional k th central moment tensor for $k \geq 3$. Thus, $Z_3(\delta)$ is the skewness and $Z_4(\delta)$ the kurtosis of $\delta^\top P$. Here and below, we consider the moments of $\delta^\top P = \pi(\delta) - Ce^{rH}$ rather than those of the portfolio value $\pi(\delta)$, because the constant Ce^{rH} does not affect the central moments. The constant only shifts the mean, and that does not affect the solution of the problem (7).

In general, the multi-objective problem (7) does not have a solution that satisfies all optimizations simultaneously.¹⁰ Therefore, we balance the objectives through a goal program (GP). The optimization proceeds in two steps. We first find the *moment goal* Z_k^* , defined as the stand-alone optimal k th moment of $\delta^\top P$ among all admissible strategies δ . Formally,

$$Z_k^* = (-1)^k \min\{(-1)^k Z_k(\delta) : \delta \in \Delta\}.$$

We then seek to minimize the preference-weighted deviations $d_k(\delta) = (-1)^k (Z_k(\delta) - Z_k^*) \geq 0$.¹¹ Given nonnegative preference parameters w_k, γ_k and d_0 , we consider the problem

$$\min_{\delta \in \Delta} \Theta(\delta) = \sum_{k=1}^K w_k (d_0 + d_k(\delta))^{\gamma_k}. \quad (8)$$

The objective function $\Theta(\delta)$ is a sum of linearly and exponentially weighted deviations from the moment goals. The parameters γ_k and w_k govern the weights, and reflect the investor's priorities over different moment goals. The w_k can include normalizing factors that scale each moment appropriately so that the

¹⁰The mean optimal portfolio, for instance, is likely to include large notional values on names that have high expected default payments or large spread movements. This composition is unlikely to be optimal for the portfolio variance.

¹¹A variable $d_k(\delta)$ is nonnegative, because $Z_k(\delta) \leq Z_k^*$ for k odd and $Z_k(\delta) \geq Z_k^*$ for k even.

comparison is meaningful between different moments and different exponents. The parameter d_0 controls the monotonicity of preferences.¹² The preference parameters determine the investor's marginal rate of substitution of moment i for moment j , which is given by $MRS_{ij} = \frac{\partial \Theta / \partial d_i}{\partial \Theta / \partial d_j} = \frac{w_i \gamma_i (d_0 + d_i)^{\gamma_i - 1}}{w_j \gamma_j (d_0 + d_j)^{\gamma_j - 1}}$. Setting $\gamma_k = 1$ for all k , we obtain a GP with linear weights and constant MRS , while setting $w_k = 1$ for all k , we obtain a polynomial goal program.¹³ The maximum order of the problem is specified by K , so for instance $K = 4$ represents the 4-moment problem. Individual moment objectives can be deactivated by setting the corresponding preference parameters to zero.

The GP formulation (8) decomposes the swap portfolio problem into simpler solvable problems and then iteratively attempts to find solutions that preserve, as closely as possible, the individual moment goals while also reflecting the investor's priority over the different moments.

5 Implementing the GP

This section illustrates the implementation of the GP (8) using a standard intensity-based model of correlated default timing, which is formulated in Section 5.1. The required moments are calculated in Section 5.2. Numerical results will be presented in Section 6.

5.1 Default timing model

5.1.1 Actual intensity

We suppose a name i defaults at an intensity λ_i representing the conditional default rate relative to the actual measure \mathbb{P} . We adopt a doubly-stochastic formulation, see Chava & Jarrow (2004), Duffie & Garleanu (2001), Duffie et al. (2006), Eckner (2009), Mortensen (2006), Papageorgiou & Sircar (2007) and others. The intensity λ_i is a specified function Λ_i of an idiosyncratic risk factor X_i and a systematic risk factor X_0 . The risk factors are independent of one another. Conditional on a realization of (X_i, X_0) , name i defaults at the first jump time of an inhomogeneous Poisson process with rate $\Lambda_i(X_i, X_0)$. Given a realization of the

¹²Since $d_k(\delta) \geq 0$, the function $(d_0 + d_k(\delta))^{\gamma_k}$ is monotonically increasing in γ_k for any $d_0 \geq 1$. In other words, the higher γ_k , the more important is the role that the associated moment plays in the overall objective. Conversely, for $d_0 < 1$, the function may become nonmonotonic.

¹³Lai (1991), Chunchinda, Dandapani, Hamid & Prakash (1997), Sun & Yan (2003) and others formulate polynomial goal programs to select a portfolio of stocks. There are, however, key technical distinctions of the GP (8), which make it a different mathematical problem. We can no longer subsume the second moment objective as a constraint in the individual optimizations. In the aforementioned papers, the decision variables are the relative portfolio weights. Hence the unit variance constraint can be subsumed in the individual optimizations at the first stage and the weights re-scaled after the second stage to satisfy the constraints. However, in our problem, the decision variables are the swap notional values restricted by nonlinear constraints. Therefore the second moment is explicitly incorporated in the two stages (as an objective in the first and a constraint in the second). Both stages are governed by the constraints.

systematic factor X_0 , which drives the intensities of all reference names, the default times are independent of one another. This property facilitates the calculation of the moments M_k required for the GP, and renders the parameter estimation problem computationally tractable. However, it also rules out possible contagion effects from defaults (see Azizpour, Giesecke & Schwenkler (2010)).

As shown in Section 5.2, the only requirement on the risk factor dynamics is that they facilitate a computationally tractable characterization of certain transforms of cumulative intensities. Examples of models that meet this requirement include affine jump-diffusions (Duffie, Pan & Singleton (2000)) and quadratic diffusions (Leippold & Wu (2002)). Here, for purposes of illustration, we assume that the factors follow Cox-Ingersoll-Ross (CIR) processes, which are members of the affine family:

$$dX_i(t) = \kappa_{X_i}[\theta_{X_i} - X_i(t)]dt + \sigma_{X_i}\sqrt{X_i(t)}dW_i(t), \quad i = 0, 1, \dots, n \quad (9)$$

where the W_i s are independent Brownian motions under \mathbb{P} . We assume that $2\kappa_{X_i}\theta_{X_i} > \sigma_{X_i}^2$ so that $X_i(t)$ is strictly positive.¹⁴ The intensities are given by

$$\lambda_i = X_i + \omega_i X_0 \quad (10)$$

where the ω_i s are the factor loadings of the common risk factor.¹⁵ These loadings control and differentiate the level of exposure that a firm has to the common risk factor, and therefore they modulate the correlation structure of the default times. As shown by Duffie & Garleanu (2001), certain restrictions apply to the parameters for the intensity λ_i to remain a CIR process. Denoting the parameter triples of the idiosyncratic and common factors by $(\kappa_{X_i}, \theta_{X_i}, \sigma_{X_i})$ and (κ, θ, σ) , respectively, we require that

$$\kappa_{X_i} = \kappa \quad (11)$$

$$\sigma_{X_i} = \sqrt{\omega_i}\sigma. \quad (12)$$

There are no restrictions on the mean level parameters θ_{X_i} . The resulting CIR parameter triple for the intensity λ_i is given by

$$(\kappa_i, \theta_i, \sigma_i) = (\kappa, \theta_{X_i} + \omega_i\theta, \sqrt{\omega_i}\sigma). \quad (13)$$

¹⁴Below, we always assume that this condition is met when a new CIR process is introduced.

¹⁵It should be noted that the ω_i s are relative in that if they are all scaled by a common constant, equivalent intensity dynamics can be obtained by scaling the common risk factor parameters.

The parameter triple (13) can be estimated from historical default experience and risk factor observations by the method of maximum likelihood, as in Chava & Jarrow (2004) and Duffie et al. (2006). The doubly-stochastic property implies that the likelihood function is the product of the risk factor-conditional likelihood functions of the firms' survival events.

5.1.2 Risk-neutral intensity

The mark-to-market value (3) of a swap position at the horizon H is governed by an intensity λ_i^* under the risk-neutral measure \mathbb{P}^* . This intensity represents the risk-neutral conditional default rate.¹⁶ We need to specify the dynamics of λ_i^* under actual and risk neutral probabilities because the moments M_k are given in terms of nested expectations under actual and risk-neutral measures.

Our measure change from actual to risk-neutral probabilities incorporates a premium for default timing risk as an adjustment to the actual intensity λ_i , and accounts for the premium for diffusive mark-to-market fluctuations through adjustments to the Brownian motions driving the λ_i . Berndt, Douglas, Duffie, Ferguson & Schranz (2005) provide regression results and an empirical time-series fitting in which the risk-neutral intensity λ_i^* is modeled as an affine function of the actual intensity plus a stochastic noise term. Motivated by this formulation, we take

$$\lambda_i^*(t) = \hat{\alpha}_i + \beta_i \lambda_i(t) + \hat{u}_i(t), \quad (14)$$

where $\hat{\alpha}_i$ and β_i are constant parameters and the \hat{u}_i s are stochastic noise terms. With (10),

$$\lambda_i^*(t) = \hat{\alpha}_i + \beta_i X_i(t) + \omega_i \beta_i X_0(t) + \hat{u}_i(t). \quad (15)$$

Since the factor X_i and X_0 are CIR processes, λ_i^* will also be a CIR process as long as the noise term \hat{u}_i is a CIR process, and parameter restrictions are met.

We extend the formulation of Berndt et al. (2005) in that we incorporate correlation between intensities via (10). In the model of λ_i^* specified by (15), the intensities remain correlated through the common risk factor X_0 . Given the structure of the transformation (15), we will decompose (14) into two independent transformations – a component attributed to the transformation of the idiosyncratic factor process $X_i \rightarrow X_i^*$, and another one describing $X_0 \rightarrow X_0^*$. We achieve (14) by an explicit specification of these two transformations. This is consistent with the empirical results of Eckner (2008), who finds that the default timing risk premium is composed of an idiosyncratic component and a component ascribed to the correlation

¹⁶Artzner & Delbaen (1995) show that if there is an intensity under \mathbb{P} , then there is also one under \mathbb{P}^* .

risk. As we will see, this decomposition leads to a natural formulation for the diffusive premium as well, and is convenient for the moment computations that follow in Section 5.2.

To this end, we work with $2n + 2$ state variables – n idiosyncratic risk factors, one common factor, and $n + 1$ factors that modulate the relationship between the actual and risk-neutral factors. The modulator terms, denoted by u_i , are CIR processes under \mathbb{P} that track the stochastic evolution of the default timing risk premia. We posit the following relationship between the actual factor processes X_i and their risk-neutral counterparts X_i^* :

$$X_i^*(t) = \alpha_i + \beta_i X_i(t) + u_i(t), \quad i = 0, 1, \dots, n \quad (16)$$

with $\beta_0 = 1$. Here the α_i s and β_i s are constant parameters and the u_i s are again independent CIR processes with dynamics under \mathbb{P} given by

$$du_i(t) = \kappa_{u_i}[\theta_{u_i} - u_i(t)]dt + \sigma_{u_i}\sqrt{u_i(t)}dB_i(t), \quad i = 0, 1, \dots, n. \quad (17)$$

The B_i s are independent Brownian motions under \mathbb{P} , also independent of the W_i s of (9). From (15),

$$\lambda_i^*(t) = \alpha_i + \beta_i X_i(t) + u_i(t) + \omega_i \beta_i (\alpha_0 + X_0 + u_0(t)) \quad (18)$$

$$= X_i^*(t) + \omega_i \beta_i X_0^*(t) \quad (19)$$

which gives us a factor decomposition of the risk-neutral intensity λ_i^* , analogous to (10) for the actual intensity λ_i . Just as λ_i is driven by (X_i, X_0) , the transformed factors (X_i^*, X_0^*) now drive λ_i^* . Equation (18) gives us our intended formulation of (14) where $\hat{\alpha}_i = \alpha_i + \omega_i \beta_i \alpha_0$ and $\hat{u}_i = u_i + \omega_i \beta_i u_0$. There is now a correlation structure that is also embedded in the stochastic noise term \hat{u}_i , which captures fluctuations in the default timing risk premium.¹⁷ The noise \hat{u}_i follows a CIR process if

$$\kappa_{\hat{u}_i} = \kappa \quad (20)$$

$$\sigma_{\hat{u}_i} = \sqrt{\omega_i \beta_i} \sigma, \quad i = 0, 1, \dots, n, \quad (21)$$

with $\omega_0 = \beta_0 = 1$. These restrictions ensure that the risk-neutral factors and intensities remain CIR pro-

¹⁷While Eckner (2008) also decomposes the risk premia in terms of idiosyncratic and common influences, he chooses to model the measure change as a scaling of the factors and a modification of the factor loading. Hence, unlike our formulation, fluctuations in the default timing risk premia are driven solely by the scaled common factor. In our formulation the fluctuations inherit the correlation structure of the intensities.

cesses. The resulting CIR parameter triple of the risk-neutral intensity under \mathbb{P} is

$$(\bar{\kappa}_i^*, \bar{\theta}_i^*, \bar{\sigma}_i^*) = (\kappa, \frac{\alpha_i + \omega_i \beta_i \alpha_0}{\kappa} + \beta_i \theta_{X_i} + \omega_i \beta_i \theta + \theta_{u_i} + \omega_i \beta_i \theta_{u_0}, \sqrt{\omega_i \beta_i} \sigma). \quad (22)$$

For the individual risk factors, $(\bar{\kappa}_{X_i}^*, \bar{\theta}_{X_i}^*, \bar{\sigma}_{X_i}^*) = (\kappa, \frac{\alpha_i}{\kappa} + \beta_i \theta_{X_i} + \theta_{u_i}, \sqrt{\omega_i \beta_i} \sigma)$, $i = 0, 1, \dots, n$. While our model specification allows the flexibility of having firm-specific scale (α) and shift (β) parameters, in practice these parameters may be estimated per sector or industry for model parsimony.

Equation (22) specifies the actual measure dynamics of λ_i^* . To obtain the risk-neutral dynamics, we only need to specify how the Brownian motions driving $(X_0, X_1, \dots, X_n, u_0, u_1, \dots, u_n)$ are adjusted when the measure is changed. We take the risk-neutral measure \mathbb{P}^* to be the measure under which, along with the previous intensity transformations for the default point processes, $(W_0^*, W_1^*, \dots, W_n^*, B_0^*, B_1^*, \dots, B_n^*)$ is an $(2n + 2)$ -dimensional standard Brownian motion where

$$W_i(t) = W_i^*(t) - \frac{\eta}{\sqrt{\omega_i}} \int_0^t \sqrt{X_i(s)} ds \quad (23)$$

$$B_i(t) = B_i^*(t) - \frac{\eta}{\sqrt{\omega_i \beta_i}} \int_0^t \sqrt{u_i(s)} ds, \quad (24)$$

for $\omega_i, \beta_i > 0$ and $i = 0, 1, \dots, n$. The parameter η governs the risk premium for the diffusive volatility of mark-to-market values. It is appropriately scaled to adjust the Brownian motions driving the individual factors and modulators. This synchronization preserves model parsimony and enforces the κ restrictions that ensure that the overall risk-neutral intensity processes remain CIR processes under \mathbb{P}^* .

Now we substitute (23) and (24) into the \mathbb{P} -dynamics of X_i^* and λ_i^* , to rewrite them in terms of the \mathbb{P}^* -Brownian motions. This results in CIR dynamics for X_i^* and λ_i^* under \mathbb{P}^* specified by

$$(\kappa_{X_i}^*, \theta_{X_i}^*, \sigma_{X_i}^*) = (\kappa + \eta\sigma, \frac{\alpha_i + \kappa(\beta_i \theta_{X_i} + \theta_{u_i})}{\kappa + \eta\sigma}, \sqrt{\omega_i \beta_i} \sigma) \quad (25)$$

$$(\kappa_i^*, \theta_i^*, \sigma_i^*) = (\kappa + \eta\sigma, \frac{\alpha_i + \omega_i \beta_i \alpha_0 + \kappa(\beta_i \theta_{X_i} + \omega_i \beta_i \theta + \theta_{u_i} + \omega_i \beta_i \theta_{u_0})}{\kappa + \eta\sigma}, \sqrt{\omega_i \beta_i} \sigma). \quad (26)$$

The overall measure change that governs both the point process and diffusion transformations can be made precise by the corresponding Radon-Nikodym derivative, which is a product of the exponential martingale terms for the independent transforms. The default process, by construction, remains doubly-stochastic in the risk-neutral measure. Technical details are provided in Appendix B.

Given the parameters (13) of the actual intensity, the actual parameters (22) and risk-neutral parameters (26) of the risk-neutral intensity can be estimated from time series of single-name, index and tranche swap

market rates (see Eckner (2008) and Azizpour, Giesecke & Kim (2011)). The doubly-stochastic property generates significant computational advantages for the estimation problem.

5.2 Moment calculations

Given the Markov doubly-stochastic property of the default times under actual and risk-neutral probabilities, and the affine relationship between the risk-neutral and actual intensities, we can express the moments M_k of $P = (P_1, \dots, P_n)$ in closed form. The computations are explicit but tedious; we only outline them here.¹⁸ The arguments do not depend on the specific CIR dynamics of the risk factors; they also apply to other risk factor dynamics (such as quadratic or other affine models).

The mean vector $M_1 = (\mu_i)$ of P is given by $\mu_i = \mathbb{E}(P_i)$. From (2), we have

$$\mu_i = U_i e^{rH} + \sum_{t_m \leq H} e^{r(H-t_m)} S_i p_i(t_m) - \mathbb{E} \left(\int_0^H \ell_i e^{r(H-s)} dN_s^i \right) + \mathbb{E}(V_i) \quad (27)$$

where $p_i(t)$ is the actual probability of name i to survive to time t . By iterated expectations and the doubly-stochastic property of the default time (see Section 5.1.1), we get

$$p_i(t) = \mathbb{E} \left(\exp \left(- \int_0^t \lambda_i(s) ds \right) \right). \quad (28)$$

Because the actual intensity λ_i follows an affine model, the survival probability (28) is an exponentially affine function of $\lambda_i(0)$ that is specified by the parameter triple (13). The coefficients of that function satisfy ordinary differential equations (see Duffie et al. (2000)). The CIR dynamics of λ_i ensure that the differential equations are solvable in closed form. The resulting analytical expression for (28) leads to an analytical expression also for the third term in (27). By integration by parts, we get

$$\mathbb{E} \left(\int_0^H e^{r(H-s)} dN_s^i \right) = e^{-rH} (1 - p_i(H)) + r \int_0^H e^{-rs} (1 - p_i(s)) ds. \quad (29)$$

In order to compute the fourth term $\mathbb{E}(V_i)$ in (27), a nested expectation under different probability measures, we first compute the mark-to-market value V_i . By the Markov property of the default timing model, the risk-neutral conditional expectation $V_i = (1 - N_H^i) f(\lambda_i^*(H))$ for a function f that is governed by the risk-neutral dynamics of the risk-neutral intensity λ_i^* specified in Section 5.1.2. Exploiting this affine specification along with the doubly-stochastic property of the default time under risk-neutral probabilities, the results of Duffie et al. (2000) imply that f is a sum of exponentially affine functions whose coefficients

¹⁸The detailed computations are available upon request.

satisfy ordinary differential equations. The CIR dynamics ensure that these equations are solvable in closed form. Moreover, by iterated expectations and the doubly-stochastic property of the default time under actual probabilities, we get

$$\mathbb{E}(V_i) = \mathbb{E} \left(f(\lambda_i^*(H)) \exp \left(- \int_0^H \lambda_i(s) ds \right) \right). \quad (30)$$

To compute the expectation (30), we first note that the risk-neutral intensity λ_i^* is an affine function of the intensity λ_i plus an independent affine noise term (see (14)). This, along with the fact that f has an exponentially affine representation, allows us to use the results of Duffie et al. (2000) to express $\mathbb{E}(V_i)$ as a sum of exponentially affine functions of the initial risk factor values. Once again, the CIR dynamics of the risk factors and noise terms ensure that these functions take a closed form.

Proceeding to the computation of higher-order moments, the covariance matrix $M_2 = (\mu_{ij})$ of P takes the form $\mu_{ij} = \mathbb{E}(P_i P_j) - \mu_i \mu_j$. The skewness tensor $M_3 = (\mu_{ijk})$ of P is given by

$$\mu_{ijk} = \mathbb{E}(P_i P_j P_k) - \mu_i \mathbb{E}(P_j P_k) - \mu_j \mathbb{E}(P_i P_k) - \mu_k \mathbb{E}(P_i P_j) + 2\mu_i \mu_j \mu_k. \quad (31)$$

All higher-order moment tensors can be expanded similarly. To compute these tensors, we must compute terms of the form $\mathbb{E}(\prod_{j=1, \dots, k} P_{i_j})$ for $k \geq 2$. The computation of these terms uses the same arguments as the computation of the first moment, along with the property that the default times of the reference names are independent conditional on a realization of the risk factors. The computation simplifies under the assumption that $\ell_i = \ell$ for all names i . We omit the details.

6 Numerical results

This section provides numerical results for the default timing model of Section 5.1. In particular, we examine the implications for optimal swap portfolios of the investment constraints discussed in Section 3. We also show that optimal portfolios are relatively robust to model parameter estimation errors.

6.1 Assumptions

We consider a portfolio of $n = 15$ reference names. The names are ordered according to their credit quality, as measured by the protection premiums: name 1 is the least risky name, while name 15 is the most risky one. Table 1 reports the model parameters. The parameters represent a realistic range of low, medium, and high-risk names based on the empirical estimates of Duffie et al. (2006) and Eckner (2008). The parameters

Name	S_i	U_i	c_i	ω_i	θ_{X_i}	σ_{X_i}	$\theta_{X_i}^*$	$\sigma_{X_i}^*$	α_i	β_i	θ_{u_i}	Mean	SD	Skew
0	N/A	N/A	N/A	1	0.023	0.114	0.159	0.114	0.030	1.00	0.042	N/A	N/A	N/A
1	25	-2.3	0.5	0.09	0.002	0.034	0.003	0.021	0.000	0.38	0.001	0.006	0.041	-15.5
2	25	0.4	0.5	0.16	0.003	0.046	0.010	0.025	0.003	0.31	0.001	0.017	0.050	-12.3
3	25	3.9	0.5	0.27	0.005	0.059	0.012	0.039	0.003	0.44	0.002	0.026	0.063	-9.2
4	25	7.1	0.5	0.24	0.009	0.056	0.027	0.037	0.007	0.45	0.002	0.039	0.069	-8.1
5	125	-3.7	2.4	0.48	0.034	0.079	0.047	0.041	0.012	0.27	0.002	0.051	0.145	-4.5
6	125	-0.8	2.4	0.56	0.009	0.085	0.032	0.062	0.008	0.52	0.005	0.074	0.100	-6.5
7	125	0.4	2.4	1.32	0.055	0.131	0.035	0.063	0.005	0.23	0.006	0.033	0.183	-3.1
8	125	2.3	2.4	0.54	0.011	0.084	0.031	0.071	0.005	0.71	0.007	0.080	0.099	-6.2
9	125	8.0	2.4	0.09	0.062	0.034	0.110	0.025	0.025	0.55	0.001	0.080	0.151	-3.6
10	125	19.7	2.4	1.11	0.095	0.120	0.101	0.084	0.013	0.49	0.009	0.075	0.166	-2.4
11	125	19.8	2.4	0.03	0.069	0.020	0.183	0.017	0.044	0.73	0.000	0.118	0.135	-3.5
12	125	23.1	2.4	0.74	0.098	0.098	0.114	0.089	0.003	0.82	0.010	0.075	0.151	-2.4
13	125	34.9	2.4	1.06	0.087	0.117	0.164	0.115	0.018	0.95	0.016	0.108	0.119	-2.4
14	125	38.9	2.4	1.91	0.053	0.158	0.101	0.152	0.002	0.93	0.032	0.116	0.101	-2.5
15	125	44.1	2.4	2.55	0.159	0.182	0.225	0.144	0.029	0.63	0.025	0.089	0.114	-1.5

Table 1: Model parameters for the $n = 15$ reference names in the numerical study. The names are sorted according to their riskiness (name 1 represents the least risky name while name 15 represents the riskiest name). The quarterly spreads S_i are expressed in basis points, and are for a $T = 5$ year maturity swap. The upfront rates U_i and collateral rates c_i are expressed in % (of notional). Name 0 represents the systematic risk factor X_0 ; the factor loading is specified by ω_i , see (10). The CIR parameter triple (13) of the actual intensity is $(\kappa_i, \theta_i, \sigma_i)$ where $\kappa_i = \kappa = 0.414$ and κ is the actual measure mean reversion speed of X_0 . The CIR parameter triple (26) of the risk-neutral intensity is $(\kappa_i^*, \theta_i^*, \sigma_i^*)$ where $\kappa_i^* = \kappa^* = 0.357$ and κ^* is the risk-neutral mean reversion speed of X_0 . The parameter (α_i, β_i) specifies the relation between the actual and risk-neutral risk factors, see (16). The parameter θ_{u_i} specifies the mean reversion level of the noise term (17). Other parameters include the loss rate $\ell_i = 0.6$, risk-free rate $r = 0.02$, and risk premium parameter $\eta = -0.5$. The last three columns report the mean, standard deviation (SD), and skewness of the distribution of the value P_i of a unit-notional protection selling position in an individual swap.

regarding the change of measure are consistent with the fitted sector data of Berndt et al. (2005).¹⁹ The investment horizon is $H = 1$ year and the swap maturity is $T = 5$ years. The loss rate is $\ell_i = 0.6$ and the risk-free rate is $r = 0.02$.

We compute the moments of the portfolio mark-to-market using the formulas developed in Section 5.2. Time integrals in formulas such as (29) are discretized on a quarterly grid, consistent with credit derivatives market convention. We substitute the computed moments into the constrained GP (8), which can be solved by standard nonlinear optimization routines. We perform the higher-order optimizations via nonlinear interior-point methods over randomized initial values. To deal with the existence of local optima we use randomization and perturbations of the boundaries. In conjunction with the constraints (see Section 3), the premium and loss payments define effective ranges for each name over which we generate 1000 random points used as initial values. We keep track of the running minimum of the objective value over the iterations.

¹⁹The parameter values were randomly sampled from a uniform distribution centered at the empirical estimates.

$(\gamma_1, \gamma_2, \gamma_3)$	(1,0,0)	(0,1,0)	(0,0,1)	(1,1,0)	(0,1,1)	(1,1,1)	(3,3,1)	(3,2,1)	(1,1,3)	(1,3,1)
Mean	158.2	102.7	82.6	109.4	102.5	109.4	115.7	128.1	101.6	104.0
Variance	5,109.9	13.2	20,338.2	16.5	14.1	17.0	33.2	129.9	136.4	13.3
Skewness	-1.5	-4.2	15.5	-3.7	-2.0	-2.5	-1.2	-1.0	12.6	-3.6
δ_1	0.0	49.6	-3,506.7	49.7	36.9	39.4	21.8	0.0	-269.0	44.7
δ_2	0.0	31.6	2.6	33.7	27.1	29.8	23.2	0.0	40.2	30.9
δ_3	0.0	18.7	3.6	20.7	18.6	20.8	18.1	0.0	25.2	19.6
δ_4	0.0	14.8	3.1	17.9	15.5	18.6	20.9	14.5	20.6	16.2
δ_5	0.0	4.0	0.9	4.9	4.9	5.9	6.0	5.0	5.8	4.6
δ_6	0.0	8.0	3.5	10.7	9.2	11.9	16.6	23.7	11.8	9.3
δ_7	0.0	2.3	1.2	2.4	3.0	3.2	2.2	0.0	3.5	2.6
δ_8	0.0	7.8	4.3	10.6	9.0	11.9	17.8	28.3	11.6	9.2
δ_9	0.0	3.0	0.3	4.6	3.9	5.6	8.3	14.4	4.5	3.7
δ_{10}	0.0	-2.1	1.4	2.8	-3.0	3.7	5.6	10.9	3.0	-2.0
δ_{11}	0.0	-3.2	0.2	6.1	3.8	7.3	14.7	33.2	-4.1	4.1
δ_{12}	0.0	-2.7	1.6	3.3	-3.8	4.3	6.9	13.8	-3.6	-2.6
δ_{13}	0.0	-5.6	3.2	5.3	-8.3	6.8	15.6	39.1	-7.4	-5.3
δ_{14}	0.0	-8.6	6.9	6.8	-13.1	8.9	22.6	59.6	-11.3	-8.2
δ_{15}	628.2	-7.0	3.6	4.1	-10.1	-5.8	13.5	35.5	-9.4	-7.0

Table 2: Optimization results with solvency constraint (6): optimal mean, variance, and skewness of the portfolio value, and the corresponding optimal swap notional values δ_i , for each of several preference triples $(\gamma_1, \gamma_2, \gamma_3)$. The parameter γ_1 specifies the weight of the mean, γ_2 the weight of the variance, and γ_3 the weight of the skewness in the GP (8). The other parameters $d_0 = w_k = 1$. A value δ_i indicates the optimal notional value of protection to be bought or sold. A positive δ_i indicates a protection selling position while a negative δ_i indicates a protection buying position. A value δ_i of 0 indicates the exclusion of a name from the portfolio. The investment capital is given by $C = 100$. All model parameters are specified in Table 1.

After each randomized session we also perturb the available investment capital C by small increments.²⁰

6.2 Base case with alternative moment preferences

We consider the mean-variance-skewness problem with capital $C = 100$, preference monotonicity parameter $d_0 = 1$ and linear weights $w_k = 1$ for all k . Table 2 reports the optimization results for various exponential weights $(\gamma_1, \gamma_2, \gamma_3)$ under the solvency constraint (6) (position limits will be introduced later).²¹ Figure 1 visualizes the mean-variance-skewness tradeoff of different optimal portfolios.

Columns 2 – 6 of Table 2 provide an interesting comparison between standalone and paired alternative investment strategies. The (1, 0, 0) portfolio represents an investor maximizing the expected value of the portfolio. The optimal strategy is to sell protection on the riskiest name 15, which involves the greatest protection premium among all 15 names. The optimal notional value $\delta_{15} = 628.2$ is as large as permitted by the solvency constraint.²² This strategy corresponds to a “high-risk, high-reward” strategy that some

²⁰It is a standard neighborhood search technique in nonlinear programming to go through $C \pm \epsilon$ for small values of ϵ to check for significant turning points.

²¹To ensure commensurability and to make the comparison meaningful between moment deviations with different exponential weights γ_k , the linear weights w_k have different units.

²²Curiously, the value at $H = 1$ year of a unit-notional protection selling position on name 15 does not have the highest mean




Figure 1: Mean-variance-skewness tradeoff of different optimal portfolios. A dot corresponds to a specific preference triple $(\gamma_1, \gamma_2, \gamma_3)$ and represents an optimal portfolio in the mean-variance plane. The shade of a dot represents the skewness. The base parameters are given in Table 1. The investment capital $C = 100$.

relative value funds employ in practice. The expected portfolio value in this case is 158.2, corresponding to an expected return of 58.2% over the one-year investment period. While this value is much higher than that resulting from any other strategy, the overall risk profile is unattractive if we consider the other moments. The portfolio has an extremely large variance of 5109.9 and a negative skewness of -1.5 .

The $(0, 1, 0)$ portfolio corresponds to a minimum variance strategy. This strategy has recently been emphasized in the equity portfolio selection literature, see DeMiguel & Nogales (2009) and others. It is motivated by the fact that mean equity returns are hard to estimate accurately. In view of these estimation errors and the resulting instability of the portfolio policies, the investor may be better off ignoring the mean return as a performance metric altogether. In our credit swap setting, the optimal $(0, 1, 0)$ strategy consists of buying protection on some of the riskiest names, while selling protection on the safest names. As such, it resembles a classic hedging strategy. The minimum variance is 13.2. However this is accompanied by a low expected portfolio value of 102.7, which is just slightly higher than the portfolio value obtained by leaving the capital C in the risk-free account, and a portfolio skewness of -4.2 that is hard to ignore.

On the other hand, we may also seek to maximize the skewness of the portfolio. This corresponds to the $(0, 0, 1)$ portfolio. The maximum portfolio skewness attainable is 15.5, while the variance and the expected

(see Table 1). To appreciate this observation, note that selling protection on name 15 rewards the investor with the highest possible upfront premium payment. Following the constraint (6), this allows for a position of the highest possible notional amount among all the names in the portfolio. Name 15 thus represents the best tradeoff between mean position value and notional size. This feature distinguishes credit swap portfolios from equity portfolios.

value of the portfolio are the worst among all test cases. The optimal strategy consists of purchasing a very large amount of protection on the safest name 1, which has the highest absolute skewness (see Table 1). To help fund the purchase, protection is sold on the remaining names in the portfolio. In light of their full 3-moment profiles, none of the single-moment strategies offer attractive portfolios.

Columns 5 and 6 of Table 2 represent mean-variance optimal ($\gamma_3 = 0$) and skewness-variance optimal ($\gamma_1 = 0$) portfolios, respectively. Mean-variance optimality enjoys a relatively high mean of 109.4, and a relatively low variance of 16.5. However, the portfolio has a significant negative skewness, indicating the significance of including the third moment into the optimization. Skewness-variance optimality improves both the variance and skewness but sacrifices a significant portion of the mean.

The mixed preference portfolios (columns 7 – 11) strike a better balance between the different moments. We see that raising one parameter while keeping the others fixed, generally increases the corresponding moment at the expense of the others. This is especially noticeable in the case of $(3, 2, 1) \rightarrow (3, 3, 1)$ in which the portfolio variance improves from 129.9 to 33.2, but the mean decreases from 128.1 to 115.7 and the skewness drops from -1.0 to -1.2 . Whether this is a worthy tradeoff depends on the investor’s risk profile.

Note that the optimal δ_i s vary drastically across the different moment preferences. For instance, in the $(1, 1, 1)$ portfolio we buy protection on the riskiest reference name 15 for a notional value of 5.8, while if we reduce the weight of the skewness goal in the $(3, 3, 1)$ portfolio, we sell protection on that name for a notional of 13.5. Mean-variance optimality involves selling protection on name 10, whereas variance-skewness optimality involves buying protection on the same name. In general, the optimal allocations are non-trivial, counterbalancing multiple trade-offs resulting from both the non-Gaussian features of the individual swap value distributions and from the non-linear nature of the solvency constraint. These non-linear effects are further illustrated in the next section.

6.3 The effect of position limits

The solvency constraint guarantees the investor’s solvency and bounds the swap notional values. Nonetheless, some portfolios in Table 2 involve excessive positions in certain names. For instance, the $(1, 0, 0)$ portfolio involves selling protection on name 15 with a notional amount that is over six times larger than the investor’s capital C . Similarly, the $(0, 0, 1)$ portfolio involves purchasing protection on name 1 with a notional amount that is over 35 times larger than C . To understand the implications of position limits on optimal portfolios, we re-run the solvency constrained optimization, restricting each δ_i to a range of $[-25, 25]$ (or $1/4$ of the investor’s initial wealth). While the range is arbitrary, it is chosen with the following trade-off

$(\gamma_1, \gamma_2, \gamma_3)$	(1,0,0)	(0,1,0)	(0,0,1)	(1,1,0)	(0,1,1)	(1,1,1)	(3,3,1)	(3,2,1)	(1,1,3)	(1,3,1)
Mean	123.4	102.4	94.6	111.8	102.5	110.0	112.5	116.3	100.2	104.3
Variance	106.4	15.2	48.2	21.9	15.6	19.5	23.0	35.8	25.5	15.4
Skewness	-1.0	-1.9	1.5	-2.0	-1.3	-1.9	-1.8	-1.2	0.6	-2.0
δ_1	0.0	25.0	25.0	25.0	25.0	25.0	25.0	19.2	25.0	25.0
δ_2	0.0	25.0	20.2	25.0	25.0	25.0	25.0	22.3	23.8	25.0
δ_3	0.0	25.0	14.8	25.0	21.6	25.0	25.0	17.9	17.5	25.0
δ_4	0.0	20.7	12.6	25.0	18.0	22.8	25.0	21.2	15.6	20.7
δ_5	19.6	5.6	-25.0	8.0	5.7	7.1	7.0	6.1	-1.3	5.9
δ_6	25.0	11.3	-25.0	16.9	10.6	14.2	15.8	17.1	10.2	11.9
δ_7	0.0	3.3	3.0	4.2	3.5	3.9	3.4	2.2	3.6	3.4
δ_8	25.0	11.0	8.1	16.5	10.5	14.0	15.9	18.4	10.3	11.6
δ_9	25.0	4.2	-25.0	7.0	4.4	6.5	6.9	8.6	5.1	4.7
δ_{10}	25.0	-2.9	2.7	4.2	-3.4	4.2	4.3	5.9	-2.9	3.0
δ_{11}	25.0	-4.4	3.5	8.4	4.3	8.0	9.7	15.5	-25.0	5.0
δ_{12}	25.0	-3.8	-3.8	4.9	-4.3	4.9	5.2	7.2	4.3	-3.5
δ_{13}	25.0	-7.8	-9.1	6.9	-9.2	7.2	8.9	16.6	5.9	-7.2
δ_{14}	25.0	-11.9	-25.0	8.4	-14.5	9.2	11.8	24.1	-19.3	-11.1
δ_{15}	25.0	-9.7	-10.3	5.1	-11.3	-8.9	7.1	14.4	-11.3	-9.3

Table 3: Optimization results with solvency constraint (6) and position limits $\delta_i \in [-25, 25]$: optimal mean, variance, and skewness of the portfolio value, and the corresponding optimal swap notional values δ_i , for each of several preference triples $(\gamma_1, \gamma_2, \gamma_3)$. The parameter γ_1 specifies the weight of the mean, γ_2 the weight of the variance, and γ_3 the weight of the skewness in the GP (8). The other parameters $d_0 = w_k = 1$. A value δ_i indicates the optimal notional value of protection to be bought or sold. A positive δ_i indicates a protection selling position while a negative δ_i indicates a protection buying position. A value δ_i of 0 indicates the exclusion of a name from the portfolio. The values δ_i hitting the limit are highlighted in bold. The investment capital is given by $C = 100$. All model parameters are specified in Table 1.

in mind: It must be narrow enough to have a tangible effect on the optimal allocations but wide enough so that it does not prevent the investor from exposing all capital to risk (a necessary condition imposed by the solvency constraint). Table 3 reports the optimization results.

Imposing position limits has the desired effect of limiting exposures to the predefined range. However, these benefits come at a cost. This is particularly visible in the single-moment optimization cases (columns 2-4). Because we are shrinking the feasible set of admissible notional values, the resulting optimal portfolios significantly underperform the ones attainable in the unconstrained case. In particular, the highest attainable mean is now 123.4, the lowest variance is 15.2 and the highest skewness is limited to 1.5. Furthermore, the optimal allocations in each case are more complex and spread out than the ones reported previously. The mean-optimal portfolio involves selling protection on the higher-risk names, however the different trade-offs at play here (between mean and solvency-governed notional capacity) lead to a non-trivial split across the risky names. The complex nature of the proposed allocations is further highlighted in the optimal variance and skewness portfolios, where it is optimal to buy and sell protection on a variety of low risk and high risk

names. In general, position limits will force the investor to trade protection on a larger number of names.²³ While this reduces single-objective optimality, it does have the benefit of making these single-objective portfolios more diversified, thus reducing their variance. This is particularly visible in the mean-optimal case, where the variance is reduced by a factor of 50 and in the skewness-optimal case where the variance is reduced by a factor of 400.

In the mixed cases (columns 5-11), one cannot necessarily make similar statements in terms of added diversification benefits. As the portfolios already trade off multiple moments, the allocations tend to be diverse by design, with less extreme exposures in any particular name, and if position limits improve one moment, it generally comes at the expense of another. For example, adding position limits to the (1, 1, 1) portfolio improves the mean and skewness of the portfolio over the unconstrained case. However, this comes at the expense of increased variance.

6.4 Robustness of optimal portfolios

Next, we consider the sensitivity of optimal portfolios to small changes in the model parameters. Understanding this sensitivity is important, for several reasons. First, if the selection problem is solved sequentially over shorter horizons, then changes in allocations incur costs that impact overall returns. The costs increase with the discrepancies in adjacent allocations, which are governed by the changes in the fitted model parameters. Second, in any single one-period setting, parameter estimation errors may play a role. While the parameter inference problem for models of correlated default timing such as the one considered here is relatively well-understood (see Berndt et al. (2005), Duffie et al. (2006), Eckner (2009) and Eckner (2008)), not all parameters can be estimated with high accuracy. It would be undesirable if estimation errors were to drastically change the optimal investment policies. Hence, we examine the effect of perturbing some of the parameters that are difficult to estimate accurately. These include the equilibrium mean levels θ_i of the intensities (see equation (13)) and the parameters η , α_i and β_i specifying the change from actual to risk-neutral probabilities (see Section 5.1.2).

We start from the base parameters given in Table 1, and focus on the (1, 1, 1)-portfolio of Table 2. We perturb the global measure change parameter η , the common risk factor equilibrium level θ , and the common risk factor measure change adjustment α_0 , each over a range of $\pm 25\%$. Figure 2 shows the optimal notional values δ_i for all 15 names. For the first name, the optimal notional ranges are [39.42, 40.16], [38.96, 40.24] and [39.17, 40.21], for the η , θ and α_0 perturbations, respectively. Given that the total notional committed for the base (1, 1, 1)-portfolio is 183.79, a turnover range of less than 1.28 only corresponds to 0.7% of

²³This generates higher transaction costs in practice.

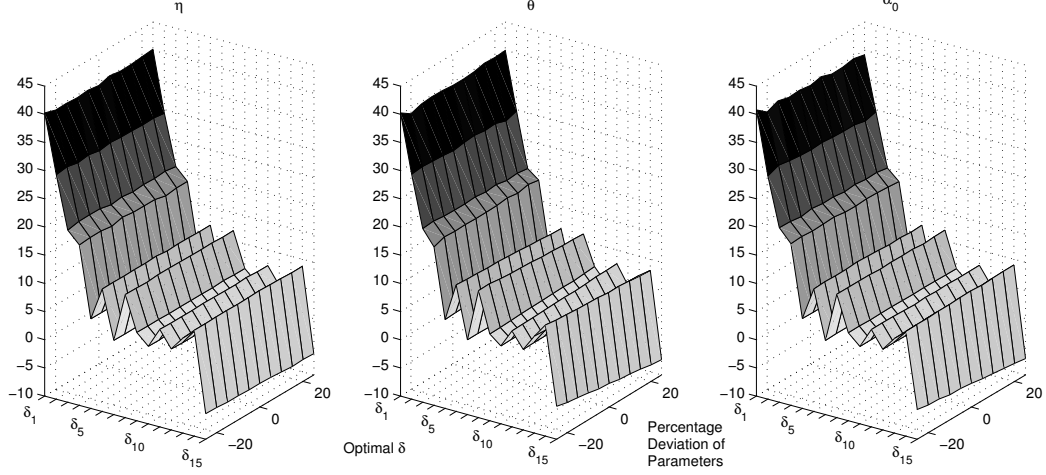


Figure 2: Optimal notional values δ_i for the $(1, 1, 1)$ -portfolio of Table 2 when the common risk factor equilibrium level θ and measure change parameters η and α_0 are perturbed over a range of $\pm 25\%$, with step size of 5%. The base parameters are given in Table 1. The investment capital $C = 100$.

the committed notional. For the second name, the respective ranges are even narrower at $[29.37, 30.11]$, $[29.13, 30.38]$ and $[29.23, 30.27]$. For the subsequent names the deviations are further reduced, and even for the most volatile θ perturbations, the maximum total deviation of the 15 names amounts to less than 7.63% of the committed notional. Thus, the optimal portfolios are remarkably robust to univariate parameter perturbations. A definite trend across the names is that θ perturbations have a more pronounced effect on the optimal allocations while the solutions are more stable under perturbations to the measure change parameters.

This trend is also indicated in Figure 3, which shows percentage changes in the optimal mean, variance, and skewness of the $(1, 1, 1)$ -portfolio with η , θ and α_0 perturbations of $\pm 25\%$. Since the moments are highly nonlinear functions of the parameters and the moments are components of a nonlinear optimization problem, not much can be said from an analytical point of view about monotonicity trends. However, what is apparent is that θ perturbations again exhibit the most salient effects. The optimal variance and skewness are constricted to relatively narrow bands for the η and α_0 perturbations, while the band for the mean is somewhat larger. The θ perturbations prompt relatively greater but not excessive changes to all optimal moments.




Figure 3: Percentage changes to the optimal mean, variance, and skewness of the $(1, 1, 1)$ -portfolio of Table 2 when the common risk factor equilibrium level θ and measure change parameters η and α_0 are perturbed over a range of $\pm 25\%$, with step size of 5% . The base parameters are given in Table 1. The investment capital $C = 100$.

6.5 Alternative GP specifications

Above we have focused on a goal program objective function (8) with polynomial weights ($w_k = 0$). An alternative parametrization of (8) uses linear rather than polynomial weights. Here, we set $\gamma_k = 1$, $d_0 = 0$ and control the different moment goals using the linear weight triple (w_1, w_2, w_3) .²⁴ Table 4 reports the optimization results with the solvency constraint and Table 5 those with solvency constraint and position limits $[-25, 25]$. The results in columns 2-7, where the weights are either 1's or 0's, are identical in linear and polynomial formulations of the goal weights. This is expected given the form of (8). In contrast, the mixed cases in the last four columns lead to different allocations and objective values.

The formulation with linear weights is arguably simpler than the polynomial formulation, requiring the investor to specify less parameters. However, the linear formulation implies a constant marginal rate of substitution, treating moments as perfect substitutes of one another. This can have counterintuitive consequences. In particular, the linear formulation seems to put less weight on the skewness goal. None of the mixed weight portfolios can achieve positive skewness, even in the extreme case where the skewness weight is chosen to be three times larger than the other weights.

Another potential issue of the linear model has to do with the effects of adding position limits (see Table 5). These limits can sometimes improve the highest-weighted moments in mixed portfolios. For instance, the $(3, 2, 1)$ portfolio has a mean of 110.9 without position limits, and a higher mean of 112.3 after impos-

²⁴The w_k need to be expressed in different units to ensure the commensurability of the moments in the objective.

(w_1, w_2, w_3)	(1,0,0)	(0,1,0)	(0,0,1)	(1,1,0)	(0,1,1)	(1,1,1)	(3,3,1)	(3,2,1)	(1,1,3)	(1,3,1)
Mean	158.2	102.7	82.6	109.4	102.5	109.4	109.9	110.9	111.2	104.9
Variance	5109.9	13.2	20338.2	16.5	14.1	17.0	17.1	18.7	20.3	13.6
Skewness	-1.5	-4.2	15.5	-3.7	-2.0	-2.5	-3.2	-2.5	-1.9	-3.4
δ_1	0.0	49.6	-3506.7	49.7	36.9	39.4	45.7	41.4	35.8	43.4
δ_2	0.0	31.6	2.6	33.7	27.1	29.8	32.8	30.9	28.4	30.5
δ_3	0.0	18.7	3.6	20.7	18.6	20.8	21.2	20.8	21.1	19.7
δ_4	0.0	14.8	3.1	17.9	15.5	18.6	18.6	19.0	19.2	16.5
δ_5	0.0	4.0	0.9	4.9	4.9	5.9	5.4	5.6	7.3	4.8
δ_6	0.0	8.0	3.5	10.7	9.2	11.9	11.3	12.2	12.8	9.7
δ_7	0.0	2.3	1.2	2.4	3.0	3.2	2.7	2.7	4.4	2.7
δ_8	0.0	7.8	4.3	10.6	9.0	11.9	11.3	12.4	12.9	9.5
δ_9	0.0	3.0	0.3	4.6	3.9	5.6	5.1	5.7	7.0	4.0
δ_{10}	0.0	-2.1	1.4	2.8	-3.0	3.7	3.1	3.6	-2.9	2.5
δ_{11}	0.0	-3.2	0.2	6.1	3.8	7.3	6.6	8.0	8.7	4.5
δ_{12}	0.0	-2.7	1.6	3.3	-3.8	4.3	3.7	4.3	5.8	3.0
δ_{13}	0.0	-5.6	3.2	5.3	-8.3	6.8	5.8	7.6	8.7	-5.2
δ_{14}	0.0	-8.6	6.9	6.8	-13.1	8.9	7.5	10.1	11.1	-8.2
δ_{15}	628.2	-7.0	3.6	4.1	-10.1	-5.8	4.5	6.2	7.5	-7.0

Table 4: Optimization results with solvency constraint (6): optimal mean, variance, and skewness of the portfolio value, and the corresponding optimal swap notional values δ_i , for each of several preference triples (w_1, w_2, w_3) . The parameter w_1 specifies the weight of the mean, w_2 the weight of the variance, and w_3 the weight of the skewness in the GP (8). The other parameters $d_0 = 0$ and $\gamma_k = 1$. A value δ_i indicates the optimal notional value of protection to be bought or sold. A positive δ_i indicates a protection selling position while a negative δ_i indicates a protection buying position. A value δ_i of 0 indicates the exclusion of a name from the portfolio. The investment capital is given by $C = 100$. All model parameters are specified in Table 1.

ing position limits. The same effect is observed in the $(3, 3, 1)$ portfolio. This may be counterintuitive to investors because adding constraints is associated with moving away from optimality in the single-objective setting. In contrast, the polynomial formulation does not exhibit this behavior.

There are other formulations of the objective function that may be of interest in practice. For example, the moment deviations $d_k(\delta)$ in the GP objective function (8) could be normalized by the moment goals $|Z_k^*|$, provided these are strictly positive. This would render the summands in the objective function dimensionless and ensure commensurability. On the other hand, this type of normalization would ignore absolute deviations which may also be important to investors. The optimization may lead to counterintuitive allocations if the moment goals differ by orders of magnitude.²⁵

7 Conclusion

We have analyzed the selection problem for a portfolio of credit swaps. The problem was cast as a goal program involving a constrained optimization of preference-weighted moments of the portfolio value at

²⁵The optimization results for this formulation and other extensions are available upon request.

(w_1, w_2, w_3)	(1,0,0)	(0,1,0)	(0,0,1)	(1,1,0)	(0,1,1)	(1,1,1)	(3,3,1)	(3,2,1)	(1,1,3)	(1,3,1)
Mean	123.4	102.4	94.6	111.8	102.5	110.0	111.8	112.3	104.8	104.5
Variance	106.4	15.2	48.2	21.9	15.6	19.5	22.0	22.6	16.4	15.4
Skewness	-1.0	-1.9	1.5	-2.0	-1.3	-1.9	-1.9	-1.8	-1.5	-2.0
δ_1	0.0	25.0	25.0	25.0	25.0	25.0	25.0	25.0	25.0	25.0
δ_2	0.0	25.0	20.2	25.0	25.0	25.0	25.0	25.0	25.0	25.0
δ_3	0.0	25.0	14.8	25.0	21.6	25.0	25.0	25.0	19.1	24.9
δ_4	0.0	20.7	12.6	25.0	18.0	22.8	25.0	25.0	17.4	20.7
δ_5	19.6	5.6	-25.0	8.0	5.7	7.1	8.0	7.4	6.7	6.0
δ_6	25.0	11.3	-25.0	16.9	10.6	14.2	16.5	15.7	12.0	12.1
δ_7	0.0	3.3	3.0	4.2	3.5	3.9	4.3	3.8	4.3	3.5
δ_8	25.0	11.0	8.1	16.5	10.5	14.0	16.2	15.7	12.1	11.9
δ_9	25.0	4.2	-25.0	7.0	4.4	6.5	7.1	7.1	6.5	4.9
δ_{10}	25.0	-2.9	2.7	4.2	-3.4	4.2	4.3	4.4	-2.3	3.1
δ_{11}	25.0	-4.4	3.5	8.4	4.3	8.0	8.5	9.3	8.0	5.3
δ_{12}	25.0	-3.8	-3.8	4.9	-4.3	4.9	5.0	5.2	5.8	-3.4
δ_{13}	25.0	-7.8	-9.1	6.9	-9.2	7.2	7.1	8.4	-6.3	-6.9
δ_{14}	25.0	-11.9	-25.0	8.4	-14.5	9.2	8.7	10.8	-12.2	-10.8
δ_{15}	25.0	-9.7	-10.3	5.1	-11.3	-8.9	5.3	6.6	-10.8	-9.1

Table 5: Optimization results with solvency constraint (6) and position limits $\delta_i \in [-25, 25]$: optimal mean, variance, and skewness of the portfolio value, and the corresponding optimal swap notional values δ_i , for each of several preference triples (w_1, w_2, w_3) . The parameter w_1 specifies the weight of the mean, w_2 the weight of the variance, and w_3 the weight of the skewness in the GP (8). The other parameters $d_0 = 0$ and $\gamma_k = 1$. A value δ_i indicates the optimal notional value of protection to be bought or sold. A positive δ_i indicates a protection selling position while a negative δ_i indicates a protection buying position. A value δ_i of 0 indicates the exclusion of a name from the portfolio. The values δ_i hitting the limit are highlighted in bold. The investment capital is given by $C = 100$. All model parameters are specified in Table 1.

the investment horizon. The portfolio value takes account of the exact timing of the protection premium and default loss payments associated with the swap positions, as well as any mark-to-market profits and losses realized at the horizon. The constraints address collateral and solvency requirements, position limits, and other trading restrictions that credit swap investors face in practice. The multi-moment formulation accommodates the complex distribution of the portfolio value, which is a nested expectation under risk-neutral and actual probabilities. It also generates computational tractability. Numerical results illustrate the features of optimal portfolios. In particular, they highlight the significance of the credit swap investment constraints. They also indicate that the optimal portfolios are relatively robust with respect to changes in the parameters of the default timing model.

A Initial capital and solvency constraints

Proposition A.1. *Suppose that the collateral rate c_i satisfies $c_i \leq \ell_i$ for all protection selling positions and $c_i \leq \sum_{t_m \leq T} S_i e^{-rt_m}$ for all protection buying positions. Then the solvency constraint (6) implies the initial capital constraint (4).*

Proof. We need to show that for all δ_i , the following inequality holds:

$$\sum_{i=1}^n \left(-\delta_i U_i + |\delta_i| c_i \right) \leq \sum_{i=1}^n \left(-\delta_i U_i + \delta_i^+ \ell_i - \delta_i^- \sum_m S_i e^{-rt_m} \right),$$

where $\delta_i^+ = \max(\delta_i, 0)$ and $\delta_i^- = \min(\delta_i, 0)$. The inequality holds for the sum if it holds for each element i of the sum. A sufficient condition is therefore that for all i ,

$$-\delta_i U_i + |\delta_i| c_i \leq -\delta_i U_i + \delta_i^+ \ell_i - \delta_i^- \sum_m S_i e^{-rt_m}.$$

Simplifying both sides and using the identity $|\delta_i| = \delta_i^+ - \delta_i^-$, we have

$$\begin{aligned} (\delta_i^+ - \delta_i^-) c_i &\leq \delta_i^+ \ell_i - \delta_i^- \sum_m S_i e^{-rt_m} \\ 0 &\leq \delta_i^+ (\ell_i - c_i) - \delta_i^- \left(\sum_m S_i e^{-rt_m} - c_i \right). \end{aligned}$$

The above inequality needs to hold for all δ_i , whether buying or selling protection. For all protection selling positions, $\delta_i > 0 \Rightarrow \delta_i^+ = \delta_i$ and $\delta_i^- = 0$. The inequality reduces to $c_i \leq \ell_i$. For all protection buying positions, $\delta_i < 0 \Rightarrow \delta_i^+ = 0$ and $\delta_i^- = \delta_i$. The inequality reduces to $c_i \leq \sum_m S_i e^{-rt_m}$. \square

B Measure change

This appendix details the measure change in Section 5.1. Let $\{\lambda_i, \lambda_i^*, X_j, u_j\}$ be the (strictly positive) \mathbb{P} -CIR processes specified in Sections 5.1.1 and 5.1.2. The default indicator process (N^1, \dots, N^n) is a doubly-stochastic Poisson process driven by $X = (X_0, X_1, \dots, X_n, u_0, u_1, \dots, u_n)$ with \mathbb{P} -intensity $(\lambda_1, \dots, \lambda_n)$. Given these dynamics, for any $\eta \in \mathbb{R}$ and finite horizon $T > 0$, $\frac{\eta}{\sqrt{\omega_i}} \sqrt{X_i}$ and $\frac{\eta}{\sqrt{\omega_i \beta_i}} \sqrt{u_i}$ satisfy Novikov's condition for T , and $\int_0^T \lambda_i(s) ds$ and $\int_0^T \lambda_i^*(s) ds$ are both finite almost surely for each i . Let

$$\begin{aligned} Z_t^N &= \prod_{i=1}^n \exp \left(\int_0^t \log \left(\frac{\lambda_i^*(s)}{\lambda_i(s)} \right) dN_s^i + \int_0^t (\lambda_i(s) - \lambda_i^*(s)) ds \right) \\ Z_t^W &= \prod_{i=0}^n \exp \left(-\frac{\eta}{\sqrt{\omega_i}} \int_0^t \sqrt{X_i(s)} dW_s - \frac{\eta^2}{2\omega_i} \int_0^t X_i(s) ds \right) \\ Z_t^B &= \prod_{i=0}^n \exp \left(-\frac{\eta}{\sqrt{\omega_i \beta_i}} \int_0^t \sqrt{u_i(s)} dB_s - \frac{\eta^2}{2\omega_i \beta_i} \int_0^t u_i(s) ds \right) \end{aligned}$$

for $\omega_0 = \beta_0 = 1$. Then, a strictly positive \mathbb{P} -martingale $Z = (Z_t)_{t \leq T}$ is given by $Z_t = Z_t^N \times Z_t^W \times Z_t^B$. To show that Z is a martingale we argue as follows. The Novikov conditions for $\frac{\eta}{\sqrt{\omega_i}} \sqrt{X_i}$ and $\frac{\eta}{\sqrt{\omega_i \beta_i}} \sqrt{u_i}$ are sufficient for Z_t^W and Z_t^B to define martingales. A sufficient condition for Z_t^N to form a martingale is that $\int_0^T (\frac{\lambda_s^*}{\lambda_s} - 1)^2 \lambda_s ds < \infty$ almost surely for any finite horizon T (Liptser & Shiriyayev (1989)). We have that both $\lambda_t^* < \infty$ and $\lambda_t < \infty$ almost surely for every t , therefore, upon expansion of the integrand, we only need to verify that the quotient $\frac{\lambda_t^{*2}}{\lambda_t}$ is finite almost surely. However CIR processes are finite variance processes so $\lambda_t^{*2} < \infty$ almost surely, and the Feller condition bounds the denominator λ_t away from 0 almost surely. Hence the given quotient is indeed almost surely finite.

The process Z defines an equivalent probability measure \mathbb{P}^* on \mathcal{F}_T via $d\mathbb{P}^* = Z_T d\mathbb{P}$. On $[0, T]$, the default indicator process (N^1, \dots, N^n) is a doubly-stochastic Poisson process driven by $\{X(t) : t \in [0, T]\}$, with intensity $(\lambda_1^*, \dots, \lambda_n^*)$ relative to \mathbb{P}^* and $\{\mathcal{F}_t : 0 \leq t \leq T\}$. The fact that λ_i^* is adapted to the filtration generated by X , along with the fact that $\int_0^T \lambda_i^*(s) ds$ is finite almost surely, guarantees that the doubly-stochastic property is preserved. Furthermore, the processes W_i^* and B_i^* , $i = 0, 1, \dots, n$, defined by (23)–(24) are \mathbb{P}^* -standard Brownian motions on $[0, T]$ relative to $\{\mathcal{F}_t : 0 \leq t \leq T\}$.

References

- Akutsu, N., M. Kijima & K. Komoribayashi (2004), ‘A portfolio optimization model for corporate bonds subject to credit risk’, *Journal of Risk* **6**(2), 31–48.
- Almgren, Robert & Neil Chriss (2000), ‘Optimal execution of portfolio transactions’, *Journal of Risk* **3**(2), 5–29.
- Artzner, P. & F. Delbaen (1995), ‘Default risk insurance and incomplete markets’, *Mathematical Finance* **5**, 187–195.
- Athayde, G. & R. Flores (2004), ‘Finding a maximum skewness portfolio: A general solution to three-moments portfolio choice’, *Journal of Economic Dynamics and Control* **28**, 1335–1352.
- Azizpour, S., K. Giesecke & B. Kim (2011), ‘Premia for correlated default risk’, *Journal of Economic Dynamics and Control* **35**(8), 1340 – 1357.
- Azizpour, S., K. Giesecke & G. Schwenkler (2010), Exploring the sources of default clustering. Working Paper, Stanford University.

- Berndt, A., R. Douglas, D. Duffie, M. Ferguson & D. Schranz (2005), Measuring default risk premia from default swap rates and edfs. Working Paper, Stanford University.
- Bertsimas, Dimitris & Andrew Lo (1998), 'Optimal control of execution costs', *Journal of Financial Markets* **1**(1), 1–50.
- Bonami, P. & M. Lejeune (2009), 'An exact solution approach for portfolio optimization problems under stochastic and integer constraints', *Operations Research* **57**(3), 650–670.
- Briec, W., K. Kerstens & O. Jokung (2007), 'Mean-variance-skewness portfolio performance gauging: A general shortage function and dual approach', *Management Science* **53**(1), 135–149.
- Charnes, A., W. W. Cooper & R.O. Ferguson (1955), 'Optimal estimation of executive compensation by linear programming', *Management Science* **1**(2), 138–151.
- Chava, S. & R. Jarrow (2004), 'Bankruptcy prediction with industry effects', *Review of Finance* **8**, 537–569.
- Chunhachinda, P., K. Dandapani, S. Hamid & A.J. Prakash (1997), 'Portfolio selection and skewness: Evidence from international stock markets', *Journal of Banking and Finance* **21**, 143–167.
- DeMiguel, V. & F. Nogales (2009), 'Portfolio selection with robust estimation', *Operations Research* **57**(3), 560–577.
- Duffie, D., J. Pan & K. Singleton (2000), 'Transform analysis and asset pricing for affine jump-diffusions', *Econometrica* **68**(6), 1343–1376.
- Duffie, D., L. Saita & K. Wang (2006), 'Multi-period corporate default prediction with stochastic covariates', *Journal of Financial Economics* **83**(3), 635–665.
- Duffie, D. & N. Garleanu (2001), 'Risk and evaluation of collateralized debt obligations', *Financial Analysts Journal* **57**(1), 41–59.
- Eckner, A. (2008), Risk premia in structured credit derivatives. Working Paper, Stanford University.
- Eckner, A. (2009), 'Computational techniques for basic affine models of portfolio credit risk', *Journal of Computational Finance* **15**, 63–97.
- El Ghaoui, L., M. Oks & F. Oustry (2002), 'Worst-case value-at-risk and robust portfolio optimization: a conic programming approach', *Operations Research* **51**(4), 543–556.

- Goldfarb, D. & G. Iyengar (2003), 'Robust portfolio selection problems', *Mathematics of Operations Research* **28**(1), 1–38.
- Jarrow, R. & F. Zhao (2006), 'Downside loss aversion and portfolio management', *Management Science* **52**, 558–566.
- Jones, D. & M. Tamiz (2010), *Practical Goal Programming*, Springer, New York.
- Kraft, H. & M. Steffensen (2008), 'How to invest optimally in corporate bonds: a reduced-form approach', *Journal of Economic Dynamics and Control* **32**, 348–385.
- Kraus, A. & R. Litzenberger (1983), 'Skewness preference and the valuation of risk assets', *Journal of Finance* **31**, 1085–1100.
- Lai, T. Y. (1991), 'Portfolio selection with skewness: a multiple-objective approach', *Review of Quantitative Finance and Accounting* **1**, 293–305.
- Leippold, M. & L. Wu (2002), 'Asset pricing under the quadratic class', *Journal of Financial and Quantitative Analysis* **37**, 271–295.
- Liptser, R. & A. Shiriyayev (1989), *Statistics of Random Processes*, Springer-Verlag, Berlin.
- Lobo, M., M. Fazel & S. Boyd (2007), 'Portfolio optimization with linear and fixed transaction costs', *Annals of Operations Research* **151**(2), 341–365.
- Markowitz, H. (1952), 'Portfolio selection', *Journal of Finance* **7**(1), 77–91.
- Meindl, P. & J. Primbs (2006), Corporate bond portfolio optimization with transaction costs. Working Paper, Stanford University.
- Mortensen, A. (2006), 'Semi-analytical valuation of basket credit derivatives in intensity-based models', *Journal of Derivatives* **13**, 8–26.
- Papageorgiou, Evan & Ronnie Sircar (2007), 'Multiscale intensity models and name grouping for valuation of multi-name credit derivatives', *Applied Mathematical Finance* **15**(1), 73–105.
- Perold, A. (1984), 'Large-scale portfolio optimization', *Management Science* **30**(10), 1143–1160.
- Protter, P. (2004), *Stochastic Integration and Differential Equations*, Springer-Verlag, New York.

- Scott, R. & P. Horvath (1980), 'On the direction of preference for moments of higher order than variance', *Journal of Finance* **35**, 915–919.
- Simkowitz, M. & W. Beedles (1978), 'Diversification in a three-moment world', *Journal of Financial and Quantitative Analysis* **13**, 927–941.
- Sun, Q. & Y. Yan (2003), 'Skewness persistence with optimal portfolio selection', *Journal of Banking and Finance* **27**(6), 1111–1121.
- Tsoukalas, G., J. Wang & K. Giesecke (2013), Dynamic portfolio execution. Working Paper.
- Wise, M. B. & V. Bhansali (2002), 'Portfolio allocation to corporate bonds with correlated defaults', *Journal of Risk* **5**(1), 39–58.