Journal of Banking & Finance 67 (2016) 119–134

Contents lists available at ScienceDirect

Journal of Banking & Finance

journal homepage: www.elsevier.com/locate/jbf

Systematic multi-period stress scenarios with an application to CCP risk management

Alan De Genaro *

Dept of Economics, University of Sao Paulo, R: Luciano Gualberto, 908 Sao Paulo, Brazil
Cetip S.A. Mercados Organizados, Av. Brigadeiro Faria Lima, 1663 Sao Paulo, Brazil

ARTICLE INFO

Article history:
Received 10 September 2013
Accepted 23 December 2015
Available online 8 January 2016

JEL classification:
G23
G17
G18

Keywords:
Stress scenarios
Central Counterparty
Monte Carlo simulation
Margin requirements
Maximum block-entropy

ABSTRACT

In the aftermath of the financial crisis of 2007–2008 regulators in multiple jurisdictions have laid the foundation of new regulatory standards aiming at strengthening systemic resilience. Among different initiatives, mandatory central clearing of standardized OTC derivatives has been one of the most prominent. Because OTC derivatives entail default management procedures that are far more complex than listed derivatives, risk management procedures have to follow suit. The recent paper by Vicente et al. (2015) propose an innovative way to calculate margin requirements by using multi-period robust optimization (RO) methods that accounts for important differences between OTC and listed derivatives default procedures. Motivated by this methodology, this paper proposes a hybrid framework to construct discrete uncertainty sets, in which each element of this set can be seen as multi-period stress scenarios, which are necessary to solve the RO problem faced by the Central Counterparty (CCP). When applied to determine the margin requirements, the present method provides both qualitative and quantitative results that outperform other robust optimization models such as Ben-Tal and Nemirovski (2000) and Bertsimas and Pachamanova (2008).

1. Introduction

The economic and financial crisis that began in 2007 demonstrated significant weaknesses in the resiliency of banks and other market participants to financial and economic shocks. Losses were significantly higher with OTC transaction which, being bilaterally cleared, triggered systemic contagion and spillover risks given a default of a big player or a small one with a complex network of counterparties.

In response, the Group of Twenty (G20) initiated a reform program in 2009 to reduce the systemic risk from OTC derivatives. In many jurisdictions important changes in market regulation are in place, among the most notable are the Dodd-Frank Wall Street Reform and Consumer Protection Act in the USA and the new European Market Infrastructure Regulation (EMIR). Both initiatives have imposed that certain OTC transactions must be cleared through a Central Counterparty (CCP).

In very broad terms, a CCP can reduce systemic risk by interposing itself as a counterparty to every trade, performing multilateral netting, and providing various safeguards and risk management practices to ensure that the failure of a clearing member to the CCP does not affect other members. Thereafter, through a process called "novation", the CCP becomes the counterparty to all trades and is responsible for crediting and debiting participants accounts and monitoring proper levels of margin requirements. Therefore CCPs are entities whose main function is to novate contracts between the trading parties, becoming the seller to every buyer, and buyer to every seller. By doing so, CCPs relieve their clients of counterparty risk, which they themselves manage by, e.g. calling margins and collecting default funds to mutualize possible losses.

Unlike banks and other financial institutions, CCPs do not bring additional market risk to the clearing system once all its exposure is originally offset. On the other hand, CCPs face counterparty credit risk under two distinct dimensions. The first relates to the potential financial loss associated with the default by one or more of the CCPs participants. Because the CCP is responsible for discharging all the obligations of any defaulting participant it may have to cover losses that arise during this process. Generally speaking, such losses are associated with the market and liquidity risks inherent in the settlement of the defaulting participants positions. The second dimension is the credit risk associated with the collateral that is part of the CCPs safeguard system. If this collateral

* Address: Dept of Economics, University of Sao Paulo, R: Luciano Gualberto, 908 Sao Paulo, Brazil.
E-mail address: adg@usp.br

http://dx.doi.org/10.1016/j.jbankfin.2015.12.011
0378-4266/

needs to be used in connection with the failure of its issuer, the result will be a loss for the CCP, "wrong way risk". As noted above, a CCP is exposed to market risk only in the event of default by one or more participants. The CCP is responsible for unwinding all of a defaulting participants original obligations and can settle each obligation on the original due date or, when possible, close out these positions against other participants under prevailing market conditions. Therefore the CCP may suffer a financial loss if market prices and/or rates move in an adverse direction to a position that is being closed out.

In general, to control its own risk, CCPs have implement mechanisms known as safeguards, designed to assure the fulfillment of their obligations to non-defaulting participants even if one or more debtors default. Generally speaking, safeguards not only enable the CCP to discharge its obligations to compliant participants on a timely basis but also cover any costs associated with the default.

Guidelines for CCP risk management and best practices have been put forth by the BIS and IOSCO in a series of joint papers known as the BIS/IOSCO papers where the most notorious article of this series was the so-called "*Principles for Financial Market Infrastructure*", CPSS-IOSCO (2012). The European Securities Markets Association (ESMA) also recently issued the EMIR initiative for Financial Market Infrastructures. Being principles-based, such regulations leave considerable room for CCPs to propose and implement risk-management structures that are better suited to their particular asset-classes and local market characteristics.

While CCPs have been widely employed in exchange-traded futures and options for decades they had limited applications on OTC markets prior to the recent financial crisis. However, the Dodd-Frank and EMIR, and similar measures under consideration elsewhere, will dramatically expand the volume of cleared transactions. We also observe that, as a consequence, academic literature on CCPs has proliferated, among them Duffie and Zhu (2011), Pirrong (2011) and Heller and Vause (2012) are the most notable.

Effective clearing through CCPs mitigates systemic risk by lowering the risk that defaults propagate from counterparty to counterparty. However, as pointed out by Pirrong (2011), although CCPs are intended to reduce systemic risk in the financial system, it must also be recognized that CCPs can create, or contribute to, systemic risk. In other words, CCP might be simply another flawed too-big-fail entity that leads market participants to increase their exposure under the incorrect assumption that their derivatives positions have zero counterparty risk. Therefore it is crucial to understand the various mechanisms by which CCPs affect the financial system in order to assess their contribution to financial stability. Furthermore, in an environment where CCPs are playing a growing role, ensure that margin requirements are enough to cover the resulting losses from the default of a counterparty should be the first concern for all market participants.

Historically CCPs have relied on mainstream risk metrics[1] such as Value at Risk and stress testing in order to define the adequate amount of collateral necessary for each participants portfolio (i.e. margin requirement). While these metrics, borrowed mainly from the banking and asset management industry, are useful for measuring potential losses for a given portfolio over a fixed time horizon (mark-to-market risk), it does not handle the dynamic nature of unwinding a defaulter's portfolio (default management process). In this sense the paper of Vicente et al. (2015) that introduced the CORE approach represents a big step towards the more accurate depiction the risk management problem faced by a CCP. According to the CORE approach, margining requirements are determined based on a risk methodology which estimates potential losses relative to the close-out process of a defaulters portfolio in multiple asset-class, multi-market Central Counterparty.

[1] The Standard Portfolio Analysis of Risk, or SPAN, developed by the Chicago Mercantile Exchange in 1988 is the most used system for calculating margin requirements for futures and options on futures.

1.1. Structure and contribution of this paper

The recent papers of Vicente et al. (2015) and Avellaneda and Cont (2013) introduced the idea of using robust optimization to formulate the multi-period closeout problem faced by a CCP and have shown excellent results as a computationally efficient alternative to existing methods for margining requirements. However, neither of these studies addressed the question of how uncertainty sets should be constructed, as a matter of fact, these papers highlight the importance of considering a broad set of scenarios for calculating closeout risk figures, including but not limited to "zig–zag" scenarios. The example presented in Vicente et al. (2015) illustrates a closeout strategy where, given only two "zig–zag" scenarios, the algorithm deliberately postpone the liquidation of the call option in order to get the most of its value.

This paper intends to shed some light on two key questions that are fundamental to advancing our knowledge on designing margin models for CCPs: (1) procyclicality and (2) the tradeoff between severity and plausibility.

Margin models that prescribe higher margins during periods of increased price volatility, or as the prices of securities decline, exacerbate procyclicality. This can be an undesirable property, as a rise in margin requirements during a period of market stress could cause market participants to face funding strains. These considerations are supported by empirical and theoretical evidence. For example, Brunnermeier et al. (2009) model an economy in which investors face funding liquidity risk due to the possibility of future margin calls or losses on existing positions.

Kamhi (2009) and more recently Murphy et al. (2014) suggest that one way to mitigate the procyclicality of margin rules is to make them less dependent on near-term market conditions. This involves determining price volatility using long historical data sets, making sure that past extreme events are captured in the data. If there are no extreme events in the data, then stress scenarios can be used to simulate such outcomes. Here we point out that there exist substantial similarities, although not necessarily explored, between designing uncertainty sets and the problem of defining trust region for stress testing.

Studer (1999) was one of the first to develop a systematic, probabilistic approach to stress testing. He considered trust regions for market risk factors which were connected higher-dimensional sets with a prescribed probability and introduced the maximum loss risk measure, defined to be the maximum loss of a market portfolio over all scenarios in the trust region.

Thus, instead of recurring to the assumption of elliptical distribution for describing uncertainty sets (trust region) as usually adopted in robust optimization (stress testing) problems, in the current work we have developed a more realistic multivariate (empirically motivated) model that captures the three stylized facts of financial asset returns literature: fat tails, volatility clustering, and tail dependence. So this paper contributes to this literature by providing a methodology that allows simulating dynamic stress scenarios which combined with an optimal closeout strategy provides margin requirements.

Regulation also has recognized that, subject to being adequately severe, stress scenarios should not be implausible, CPSS-IOSCO (2012). Federal Reserve institutionalized the use of supervisory stress tests for establishing minimum capital standards in 2010 through its now annual Comprehensive Capital Assessment and Review (CCAR) for all banking organizations with more than $50 billion in total assets, as well as systemically important financial institutions (SIFI). As pointed out by Christensen et al. (2015) financial stress tests, including those that have examined the Fed's

own financial position, usually only consider a small number of hand-picked scenarios. A benefit of this choice relies on avoiding the model risk involved in the choice of a risk factor distribution. However, this advantage comes at a price. Not assuming any risk factor distribution stress testers cannot judge whether the stress scenarios are really dangerous or sufficiently plausible.

Therefore, this paper presents a hybrid two-step approach to construct the uncertainty set (trust region) handling, effectively, the severity and plausibility concept. During the first step we define the boundaries of the uncertainty set using Extreme Value Theory (EVT). The next step involves generating multi-period trajectories via Monte Carlo simulation (MCS) that reproduce stylized facts present in financial time series. Additionally, in order to avoid the curse of dimensionality usually present in Monte Carlo problems we have resorted to a data reduction scheme that relies on the concept of maximum block-entropy.

This modeling strategy has a number of important benefits both for robust optimization and stress scenarios exercises. First, by relying on Monte Carlo simulation it is possible to generate uncertainty sets that are very flexible and are not limited to a specific parametric formulation, such as ellipsoids. Furthermore, it is also possible to incorporate almost any model developed by the financial econometric literature to describe asset prices dynamics. Second, the size of the uncertainty set, which represents the desired level of robustness, can be determined based on a statistical model designed for handling extreme occurrences. It is worth to mention one advantage of this method over Breuer et al. (2009) which consists in not imposing any trade-off between plausibility and severity of scenarios. Third, we keep the robust optimization problem, even for large-scale problems, solvable with standard optimization packages.

Finally, the robust optimization formulation presented here addresses the stress scenario plausibility problem in a much more general way than Breuer and Csiszár (2013) because the loss function is solved exactly (i.e. without any approximations) and the worst-case scenario is obtained *simultaneously* with the closeout strategy and consequently can motivate further studies.

This paper contains, besides this introduction, 5 more sections. Section 2 presents the main elements of CCP risk management and how the intertemporal closeout problem for the CCP can be written as a robust optimization problem. Next, in Section 3, we describe, in comparative terms, how to construct uncertainty sets and how to generate plausible stress scenarios. Section 4 presents a two-step methodology that uses Monte Carlo simulation and the concept of block-entropy for generating maximum entropy trajectories. Practical applications of the proposed hybrid framework are described in Section 5. Finally, in Section 6 we assess the quality of our framework for generating stress scenarios when applied to the problem of determining margin requirements. Appendix A presents the results of a comparative study between different ways of handling dependence using Copulae.

2. Foundations for risk management in CCPs

2.1. Motivation

Central clearing alters the allocation of performance risk that is inherent in derivatives trades. In a traditional OTC transaction, the original counterparties remain at risk to the failure of each other to perform on their obligations for the life of the contract. In contrast to such bilateral trades, CCP becomes the buyer to the original seller and the seller to the original buyer. If either buyer or seller defaults, the CCP is committed to pay all that is owed to the non-defaulting party. To meet its obligations, the CCP has recourse to a safeguard structure formed by collateral posted by those who clear through it and financial contributions made by its members and owners.

Yet as regulators have sought to implement the central clearing mandate globally, the clearing structure[2] have only marginally been considered in the policy debate. Their attention seems to relegate to second place any discussion on clearing arrangements because a system cleared by a CCP can be organized in very different architectures. In the simplest case, all market participants directly connect to the CCP. In other, by far more frequent, cases the CCP clears for a restricted number of institutions the General Clearing Members, or GCMs which in turn clear for other participants, and so on in a hierarchy of tiers. Galbiati and Soramäki (2012) presents a general model of clearing, designed to study the mechanics of any clearing network and offers a numerical application of this model, to study the statistical properties of exposures under a broad spectrum of clearing network. Without any lack of generality, this paper is based on a three-layer structure formed by Clearing Member, Broker and final beneficiary owner, the latter is the level where margin requirements are calculated.

In general, CCPs set collateral requirements with the intent that, given participant default, the likelihood that this participant will suffer a loss that exceeds the amount of margin held is very small. To do this, CCPs typically set initial margin to reflect their estimate of the riskiness of the underlying transaction. For instance, they typically charge higher margins on instruments with more volatile prices, and on less liquid instruments that take a CCP longer to cover in the event of a default. Crucially, CCPs typically do not establish initial margin based on the creditworthiness of the party to a contract.

As mentioned, CCPs, more often than not, determine margin requirements based on VaR or stress test metrics that measure potential credit exposure of a given portfolio with just a single number that represents the potential financial loss in the context of static portfolios and fixed holding period. The problem with this approach as it stands is that it does not take into account explicitly the dynamic nature of unwinding a portfolio. In other words, it quantifies potential mark-to-market (MTM) losses, but does not anticipate the cost of unwinding a portfolio. Due to liquidity constraints, the latter is expected to be more severe than MTM losses prior to liquidation. As a way to overcame this limitation Vicente et al. (2015) present the CORE approach which explicitly recognize that closeout processes are dynamic, so the portfolio's risk profile changes as positions are liquidated and/or settled through time. According to this view, the management problem of a CCP given a participant default consists in fulfilling all obligations with minimal market impact, so the whole process is undisruptive to other (non-defaulting) market participants. This process entails having the adequate resources for closing out the defaulters portfolio at prevailing market conditions – which can be extremely adverse. In particular, for either multi-asset or multi-market portfolios the closeout process has to take into account important practical, real life constraints (or frictions), such as different settlement cycles, cash flow mismatches and heterogeneous liquidity profiles.

Thus, from the CCP's perspective, the question that arises is how to organize and schedule the liquidation of portfolio to mitigate risk. Some portfolios have natural offsets between instruments of different liquidity for instance, interest rate futures and interest rate swaps may have the same exposure to the term-structure of interest rates, but very different market liquidity. Just closing-out all positions in the portfolio as soon as possible may in fact increase CCP risk, which is not the case when using an orderly (structured) close-out process. An appropriate choice of the liquidation strategy for a portfolio is crucial. A poor choice may give rise

[2] As known as CCP tiering.

to unwarranted losses over the liquidation period due to market fluctuations, resulting in high risk estimates; a good choice could take advantage of risk offsets between instruments that can be liquidated simultaneously, thus minimizing the overall closeout risk.

2.2. The multi-period risk management CCP problem

Typically, a portfolio is defined as a collection of N financial instruments, represented by the vector $\boldsymbol{\theta}(0) = [\theta_i(0), \ldots, \theta_N(0)]$, where each $\theta_i(0)$ corresponds to the total absolute quantity of an instrument i at an initial time $t = 0$. The uncertainty about futures states of the world is represented by a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and a set of financial risks $\mathcal{M}$ defined on this space, where these risks are interpreted as portfolio or position losses over some fixed time horizon. Additionally, a risk measure is defined as a mapping $\psi : \mathcal{M} \to \mathbb{R}$ with the interpretation that $\psi(L)$ gives the amount of cash that is needed to back a position with loss L.

In its turn, asset values are defined as a function of time and a I-dimensional random vector of t-measurable risk factors $Z_{t,i}$ by $f(t, Z_{t,1}, Z_{t,2}, \ldots Z_{t,I})$ where $f : \mathbb{R}_+ \times \mathbb{R}^I \to \mathbb{R}$ is a measurable function. For any portfolio comprised of N instruments and initial quantities $\boldsymbol{\theta}(0) = [\theta_i(0), \ldots, \theta_N(0)]$ the realized portfolio variation over the period $[t-1, t]$ is given by:

$$L_{[t,t-1]} = \sum_i^N \theta_i(0)[f_i(t, \mathbf{Z}_t) - f_i(t-1, \mathbf{Z}_{t-1})] \tag{1}$$

Here $f_i(t, \mathbf{Z}_t)$ describes the exact risk mapping function which is not only limited to linear functions.

In a typical stress tests exercise we are interested in assessing what happens to the portfolio losses given an adverse shock on the risk factors. In other words, it pictures the same portfolio at two different moments (today and future) and then calculates the maximum loss over all possible future scenarios $\mathbf{Z}_{t+m}$:

$$\sup_{\mathbf{Z}_{t+m} \in \mathcal{S}} L_{[t+m,t]}(\mathbf{Z}_t) = \sum_i^N \theta_i(0)[f_i(t+m, \mathbf{Z}_{t+m}) - f_i(t, \mathbf{Z}_t)] \tag{2}$$

In the CCP context, the equation above implicitly assumes that all defaulter's position would have to be unwind at same time $(t+m)$ and in the same market scenario, $\mathbf{Z}_{t+m}$. While this assumption is a fair proxy for the potential losses relative to the unwinding process of a defaulter's portfolio comprising just one asset class, e.g. equities or futures, this is not necessarily true for highly heterogeneous portfolio (e.g. listed and OTC derivatives), as pointed by Vicente et al. (2015).

Therefore a more realistic formulation arises when the closeout strategy $(\boldsymbol{\theta}(t))_{t \geqslant 0}$ is explicitly incorporated into the loss function. In this context, a closeout strategy $(\boldsymbol{\theta}(t))_{t \geqslant 0}$ is understood as a rule for liquidating all portfolio positions given some finite time horizon. A closeout strategy can be as simple as liquidating all positions as soon as possible, or it may have a more sophisticated structure where a set of rules and/or restrictions have to be followed. Therefore, the resulting loss function after including the closeout strategy is given by:

$$L_{[t,t+\tau]}(\mathbf{Z}_t, \boldsymbol{\theta}(t)) = \sum_t^{t+\tau} \sum_i^N \theta_i(t)[f_i(t+1, \mathbf{Z}_{t+1}) - \phi f_i(t, \mathbf{Z}_t)] \tag{3}$$

Here ϕ takes value equal to 1 for positions subjected to periodical marking to market, such as futures and futures-style options, and zero otherwise. Additionally, τ represents the liquidation horizon, meaning that all positions must be liquidated until $t + \tau$.

Once characterized the loss function for a CCP it is possible to state mathematically its problem of finding the optimal closeout strategy, $\theta_i^\star(t)$, as:

$$\sup_{\boldsymbol{\theta}(t) \in \mathcal{A}} L_{[t,t+\tau]}(\mathbf{Z}_t, \boldsymbol{\theta}(t)) = \sum_t^{t+\tau} \sum_i^N \theta_i(t)[f_i(t+1, \mathbf{Z}_{t+1}) - \phi f_i(t, \mathbf{Z}_t)] \tag{4}$$

where the set $\mathcal{A}$ describes all conditions imposed to a feasible solution.

In reality, of course, future price paths are not known at time t and uncertainty should be incorporated into this problem. Furthermore, depending on how uncertainty is introduced, which is postpone to the next section, into this problem there exist different techniques for finding the optimal solution, e.g, dynamic programing or robust optimization.

Finally, once the optimal closeout strategy has been determined, the CCP will require $L(\mathbf{Z}_\tau, \boldsymbol{\theta}^\star(\tau))$ as collateral in the form of cash or high liquid securities to support the losses incurred during the unwinding of a defaulter portfolio.

2.3. Solving the multi-period risk management CCP problem

The optimization problem posed above resemble the one found in studies of long-term investment planning in the context of a utility function based on consumption, and placed the problem in the realm of dynamic programming. However, closed-form solutions of this kind can be derived only under strong assumptions on the investors behavior and the structure of the asset price process, and cannot be easily generalized when market frictions, e.g., transaction costs, are included.

Bertsimas and Pachamanova (2008) have pointed out that recent advances in computer technology have reduced the significance of the ability to solve the multi-period portfolio problem in closed form, and have made discrete time portfolio optimization models more tractable. Additionally, a significant amount of research has been performed in the field of stochastic programming applied to portfolio optimization, for instance Steinbach (2001). The main idea of the stochastic programming approach is to represent future realizations of returns by a scenario tree, and find portfolio weights at each point in time that maximize the expected value of the returns over the whole time horizon minus some measure of risk. Although this analysis is very flexible from a modeling standpoint, its computational complexity increases exponentially with the number of time periods.

As originally posed, the solution of (4) requires some form of intertemporal optimization technique, for instance, dynamic programing. In its turn, dynamic programing characterization of uncertainty relies on the average or expected performance of the system in the presence of uncertain effects. Although expected-value optimization is often the most adequate approach for many applications in finance, risk management emphasis on worst case scenarios and tail events requires a different set of tools. Indeed, in the CCP context it is important to notice that the closeout strategy must be robust enough to minimize the losses considering all possible states and, more importantly, in adverse extreme situations such as the one represented by stress scenarios.

Robust optimization has emerged as a leading methodology for addressing uncertainty in optimization problems. While stochastic optimization's pursuit is to immunize the solution in some probabilistic sense, robust optimization constructs a solution that is optimal for any realization of uncertainty in a given set called uncertainty set, $\mathcal{S}$. Typically, the original uncertain optimization problem is converted into a equivalent deterministic form by duality arguments and then solved using standard optimization algorithms. Ben-Tal et al. (2000) were the first to suggest using robust optimization to deal with the curse of dimensionality in multi-period portfolio optimization problems. As pointed by Bertsimas and Pachamanova (2008), the formulation by Ben-Tal et al. (2000) can be viewed as an extension of the Certainty

Equivalent Controller (CEC) procedure from dynamic programming. The CEC represents a deterministic approach to uncertainty at each stage, it applies the policy that is optimal when all uncertain quantities are fixed at their expected values. An important disadvantage of the CEC approach is that risk is not factored in Ben-Tal et al. (2000) incorporate risk by allowing future asset returns to vary in ellipsoidal sets whose size is determined by the user and depends on his aversion to uncertainty. The robust counterpart of the multi-period portfolio optimization problem can then be formulated as a second order cone problem (SOCP). Although this method appears simplified, the computational results obtained by the authors indicate that this robust optimization formulation for ellipsoidal uncertainty sets outperforms stochastic programming algorithms both in terms of efficiency and in terms of optimal strategy selection.

Fortunately, the objective function (4) has important mathematical properties that allow the problem to be tackled with robust optimization methods. Avellaneda and Cont (2013) have shown that the objective function is convex; furthermore they show that the original problem can be written, by means of duality arguments, as a robust optimization problem. Hence, determining the optimal liquidation strategy consists in maximizing a convex function under the linear equality/inequality constraints. Therefore, the CCP problem that corresponds to the robust policy is given by the solution of the minimax problem:

$$\sup_{\theta(t)\in\mathcal{A}} \inf_{\mathbf{Z}_t\in\mathcal{S}} L(\mathbf{Z}_t, \theta(t)) \tag{5}$$

where the set $\mathcal{A}$ describes all conditions imposed to a feasible solution, such as zero inventory in $t+\tau$, monotonically decreasing inventory and the maximum number of contracts liquidated per day. Further details can be found in Vicente et al. (2015) and Avellaneda and Cont (2013).

Here the worst-case scenario is computed *simultaneously* with the minimization over $\theta(t)$. If the closeout strategy is efficient in this scenario, then the results associated with less severe scenarios must result in smaller losses. In a similar way Polak et al. (2010) have shown that, the worst-case optimization has provided better results than other techniques for portfolio optimization and risk mitigation such as expected value maximization or variance minimization.

After we have stated how the risk management problem faced by the CCP can be solved it is possible to move forward to define the uncertainty set $\mathcal{S}$.

3. Uncertainty sets and plausibility

In robust optimization, the description of the uncertainty of the parameters is formalized via uncertainty sets which can represent or may be formed by differences of opinions on future values of certain parameters. Bertsimas et al. (2013) have pointed out that the computational experience has suggested that with well-chosen uncertainty sets, robust models yields tractable optimization problems whose solutions performs as well or better than other approaches. On the other hand, with poorly chosen uncertainty sets, robust models may be overly-conservative or even computationally intractable.

At this point the definition of uncertainty sets resembles the definition of plausible stress scenarios present on systematic stress tests literature. Then, ignoring for a moment the concept of closeout strategy, the stress test of a particular portfolio is carried out by computing:

$$\sup_{\mathbf{Z}_t\in\mathcal{S}} L(\mathbf{Z}_{t+m},) \tag{6}$$

When comparing (6) with (5) we realize that the role played by $\mathcal{S}$ in both problems is vital. For instance, stress scenarios ($\mathbf{Z}\in\mathcal{S}$) which are not too severe will, according to (6), result in reduced losses and therefore are not so informative for risk management policy purposes, such as motivating the adoption of risk-reducing strategies. Likewise, an uncertainty set that is too narrow could not represent all uncertainty present in the robust optimization problem. On the other hand, stress scenarios ($\mathbf{Z}\in\mathcal{S}$) which produce high values for (6) are also liable to some level of questioning because if the value of $\mathbf{Z}$ is unrealistic this could undermine the credibility of stress test results.

Therefore the quality of this kind of problem depends crucially on the definition of $\mathcal{S}$. In stress testing the bias toward historical experience can lead to the risk of ignoring plausible but harmful scenarios which have not yet happened in history.

As a way to ensure that the quality of stress testing depends less on individual skills and more on models/methods, the Basel Committee on Banking Supervision (2005) issued a recommendation to construct stress scenarios observing two main dimensions: plausibility and severity. Severity tends to be the easiest component to take into account, because risk managers can ultimately look at historical data and define scenarios based on the largest movement observed for each risk factor. On the other hand, plausibility requires that after setting individual stress scenarios S_i and S_j for risk factors i and j the resulting joint scenario (S_i, S_j) makes economic sense.

The current literature in stress test has provided different approaches for generating plausible stress scenarios. A first attempt in this direction was made by Studer (1999) and Breuer and Krenn (1999), who developed what is called "traditional systematic stress tests". In particular, Studer considered elliptical multivariate risk factor distributions and proposed to quantify the plausibility of a realization by its Mahalanobis distance:

$$\sup_{\mathbf{Z}:Maha(\mathbf{Z})\leqslant k} L(\mathbf{Z}) \tag{7}$$

This approach was extended by Breuer et al. (2012) to a dynamic context where multi-period plausible stress scenarios can be generated. While this approach introduced the systematic treatment of plausibility for generating stress scenarios it has problems of its own. First, the maximum loss over a Mahalanobis ellipsoid depends on the choice of coordinates. Second, the Mahalanobis distance as a plausibility measure reflects only the first two moments of the risk factor distribution. It is important to notice that the argument for working with elliptical approximations to the risk factor is the relative tractability of the elliptical case. However, this ability to handle the problem has a cost of ignoring the fact that an given extreme scenario should be more plausible if the risk factor distribution has fatter tails.

Recently, Breuer and Csiszár (2013) proposed a new method to overcomes the shortcomings of Studer's method. The authors introduce a systematic stress testing for general distributions where the plausibility of a mix scenario is determined by its relative entropy $D(\mathcal{Q}\|\nu)$ with respect to some reference distribution ν, which could be interpreted as a prior distribution. A second innovation of their approach was the usage of mixed scenarios instead pure scenarios. Putting all elements together Breuer and Csiszár (2013) propose the following problem:

$$\sup_{Q:D(\mathcal{Q}\|\nu)\leqslant k} \mathbb{E}_Q(L) := MaxLoss(L,k) \tag{8}$$

This method represents an important contribution to the literature as it moves the analysis of plausible scenarios beyond elliptical distributions. The authors show that under some regularity conditions the solution to (8) is obtained by means of the Maximum Loss Theorem and some concrete results are presented for linear and quadratic approximations to the loss function. Even though the solution of (8) can be shown theoretically to exist and to be unique by the Maximum Loss Theorem it is not that

simple in practice for complex portfolios, in particular with non-linearities in the loss function, since it involves the evaluation of a integral over the possibly high-dimensional sample space.

In terms of how to construct uncertainty sets, much of the Robust Optimization literature assumes an underlying parametric structure a priori for describing disturbances. For instance, Ben-Tal et al. (2000) suggest an ellipsoidal uncertainty sets for describing the uncertainty set:

$$\mathcal{S} = \{\mathbf{Z}_t : \|\Sigma^{-1/2}(\mathbf{Z}_t - \bar{\mathbf{Z}}_t)\|_2 \leqslant \rho, \quad \forall\, t \in [1, T_{max}]\} \tag{9}$$

where Σ represents the covariance matrix among risk factors and $\bar{\mathbf{Z}}_t$ is a vector comprised of the mean values for each risk factor. The parameter ρ, as know as robustness budget, is a subjetive value chosen by the decision maker to reflect his attitude towards risk; the larger is ρ, the more risk averse he is. Therefore risk is incorporated by allowing future asset returns to vary in ellipsoidal sets whose size ρ is determined by the user and depends on his aversion to uncertainty.

A more general approach was developed by Bertsimas et al. (2004) where uncertainty on future returns is described by polyhedral rather than ellipsoidal uncertainty sets. They studied the robust counterparts of linear optimization problems when the total distance (according to a pre-specified norm) between the realized uncertain coefficients and their nominal values is restricted to be less than a robustness budget ρ. Based on this structure, a flexible construction for the uncertainty set has been proposed in the multi-period asset allocation context by Bertsimas and Pachamanova (2008). The authors propose to describe the uncertainty by means of polyhedral sets, which in its simplest form, assumes that i-th asset returns vary in intervals, such as:

$$\mathcal{S} = \{\mathbf{Z}_t : S(t)_i^l \leqslant Z(t)_i \leqslant S(t)_i^u \quad \forall\, t \in [1, T_{max}],\ i = 1, \ldots, I\} \tag{10}$$

where T_{max} is the maximum holding period and I is the number of risk factors.

The length of these intervals can be determined, for example, as a percentage of their standard deviations or by means of bootstrapping as pioneered by Tütüncü and Koeing (2004). The solution produced with this assumption can be viewed as a worst-case nominal policy in the sense that the program will protect against uncertainty by setting all returns to their lowest possible values the end points of the intervals. However, this approach may be overly conservative or implausible. In practice, there is usually some form of correlation amongst future returns, and it rarely happens that all uncertain returns take their worst-case values simultaneously. It may therefore be desirable to incorporate some kind of variability and correlation of asset returns. So the uncertainty set takes the form:

$$\mathcal{S} = \{\mathbf{Z}_t : \|\Sigma_t^{-1/2}(\mathbf{Z}_t - \bar{\mathbf{Z}}_t)\|_d \leqslant \rho_t, \quad \forall\, t \in [1, T_{max}]\} \tag{11}$$

As in (9) Σ_t represents the covariance matrix among risk factors and $\bar{\mathbf{Z}}_t$ is a vector comprised of the mean values for each risk factor. However, the Bertsimas and Pachamanova (2008) approach given by (11) generalizes (9) in two directions. First, by defining a norm called the d-norm that includes the polyhedral norms L_1 and the L_∞ as special cases, and they demonstrate that for $d = \sqrt{n}$, where n is the dimension of the vector of uncertain returns, the d-norm approaches the L_2 norm. Thus, if we use the L_2 norm we obtain Ben-Tal et al. (2000) formulation. Second, by making time-varying the robustness budget parameter ρ_t and the covariance matriz Σ_t.

It is worth to mention that embedded in all uncertainty sets $\mathcal{S}$ presented above is the assumption that uncertainty is modeled by a continuous variable, so in this cases the corresponding robust optimization problem is called continuous minimax problem. On the other hand, if the uncertainty set $\mathcal{S}$ is a finite set, e.g. $\mathcal{S} = \{Z_1, Z_2, \ldots, Z_k\}$, the corresponding optimization problem is called discrete minimax and no reformulation is necessary to preserve the structural properties (linearity, convexity, etc.) of the original optimization problem. Consequently, when the uncertainty set is a finite set the resulting robust optimization problem is larger but theoretically no more difficult than the non-robust version of the problem.

Finally, we finalize this section by noting that the definition of the set $\mathcal{S}$ is equally important for robust optimization and systematic stress testing, which, despite its terminology, depend on the definition of $\mathcal{S}$ to obtain their results. As already mentioned a poorly chosen uncertainty sets, can produce results which are overly-conservative (implausible) or even computationally intractable.

4. A hybrid framework for generating $\mathcal{S}$

A key element either in robust optimization problem or stress testing is to define an uncertainty set of possible realizations of the uncertain parameters and then optimize against worst-case realizations within this set. Fortunately, as pointed by Bertsimas et al. (2013), there are several theoretically motivated and experimentally validated proposals for constructing good uncertainty sets in robust optimization problems. According to the authors, most of these proposals share a common paradigm: they combine a priori reasoning with mild assumptions on the uncertainty to motivate the construction of the set.

Traditional approaches, such as Bertsimas and Sim (2004) and Chen et al. (2010), typically assume that $\mathbf{Z}$ is a random variable whose distribution $\mathbb{P}$ is not known exactly. In general, these approaches make only mild a priori assumptions about certain structural features of $\mathbb{P}$. For example, they may assume that $\mathbb{P}$ has an elliptical distribution. These approaches then look for a set $\mathcal{S}$ that would satisfy two key properties: (i) being computationally tractable (ii) involving either deterministic or probabilistic guarantees of constraints against violation, e.g $\mathbf{Z}^T x \leqslant b\ \forall\, \mathbf{Z} \in \mathcal{S}$. We can observe that while elliptical distributions have been an usual assumption for describing uncertainty on robust optimization problems empirical research in financial time-series indicates the presence of a number of facts which are not properly captured by this class of distributions. Concretely, empirical literature indicates that asset returns display a number of so-called stylized facts: fat tails, volatility clustering and tail dependence.

In the current work we propose a framework for designing uncertainty sets which is flexible enough to capture all the three stylized facts present in financial asset returns: fat tails, volatility clustering, and tail dependence. In our approach, rather than imposing several assumptions about certain structural features of $\mathbb{P}$, we generate several scenarios for the uncertain variables $\mathbf{Z}$ based on Monte Carlo simulation and then optimize against the elements within this set. So this sampling procedure allows to incorporate with a relative low cost different state-of-art statistical models for describing the uncertainty on $\mathcal{S}$ and produces a robust optimization formulation with a finite uncertainty set:

$$\mathcal{S} = \{\mathbf{Z}_{1,t}(\omega), \mathbf{Z}_{2,t}(\omega), \ldots, \mathbf{Z}_{N,t}(\omega)\} \quad \forall\, t \in [1, T_{max}] \tag{12}$$

where $\mathbf{Z}_i$ represents a realization (ω) for the source of uncertainty under consideration.

Therefore, we propose a hybrid two-step approach to define discrete scenarios in $\mathcal{S}$. Thus, as a *first step* we define envelope scenarios, S^i_{env}, for each risk factor in $\mathcal{M}$ using Extreme Value Theory (EVT). Technically, EVT is a limit law for extremes just like the central limit theory is a limit law for the mean. Using these laws, it is possible to find limiting distributions to model only the tails of the sample instead of imposing a distribution to the entire sample.

Also, since EVT methods are implemented separately to the two tails, the skewness is implicitly accounted for Embrechts et al. (2001), McNeil and Frey (2000), among others, showed that EVT methods fit the tails of heavy-tailed financial time series better than more conventional distributional approaches. Note that by recognizing that worst-case scenarios are indeed associated with extraordinary events it is very straightforward to employ elements from EVT literature to provide the ground for constructing $\mathcal{S}$.

Basically, envelope scenarios would represent the worst-case scenario (WCS) for each risk factor and have the highest level of severity embedded:

$$S^i_{env} = [S(t)^i_l, S(t)^i_u], \quad \forall\, t \in [1, T_{max}] \quad \text{and} \quad i = 1, \ldots, I \tag{13}$$

where T_{max} is the maximum holding period and I is the number of risk factors.

In this context it is assumed that tails of asset returns have a cumulative Generalized Pareto Distribution given by:

$$F_{(\xi,\sigma)}(x) = 1 - \left(1 + \frac{\xi x}{\sigma}\right)^{-\frac{1}{\xi}} \tag{14}$$

The support is $z \geqslant 0$ for $\xi \geqslant 0$ and $0 \leqslant z \leqslant -1/\xi$ for $\xi < 0$.

This particular distributional choice is motivated by a limit result in EVT which establish convergence results for a broad class of distribution functions G and its distribution function of the excesses over certain (high) threshold k denoted by $G_k(x)$. The Pick ands–Balkema–de Haan theorem states that as the threshold k becomes large, the distribution of the excesses over the threshold tends to the Generalised Pareto distribution, provided the underlying distribution G satisfies some regularity conditions. In the class of distributions for which this result holds are essentially all the common continuous distributions of statistics, and these may be further subdivided into three groups according to the value of the parameter ξ in the limiting GPD approximation to the excess distribution[3].

Once the parameters of the GPD have been estimated the upper and lower bounds of (13) can be obtained using the percentile of (14):

$$S(t)^i_l = \inf\{x \in \Re : F_{(\xi,\sigma)}(x) \geqslant \alpha_l\}, \quad \forall\, t \in [1, T_{max}] \quad \text{and} \quad i = 1, \ldots, I \tag{15}$$

$$S(t)^i_u = \inf\{x \in \Re : F_{(\xi,\sigma)}(x) \geqslant \alpha_u\}, \forall\, t \in [1, T_{max}] \quad \text{and} \quad i = 1, \ldots, I \tag{16}$$

A compelling argument for this approach is the capacity of defining the size of the uncertainty set based on probabilistic terms instead of risk aversion parameters which can be either hard to estimate or that are attached to a specific family of utility functions. Thus, as it is standard in the risk management literature, for daily data, and assuming a confidence level of 99.96%, it is expected that actual variation would exceed $S(t)_l$ one out of 10 years.

Note that the envelopes proposed above are similar to polyhedral uncertainty set proposed by Bertsimas and Pachamanova (2008), but now they are determined using a recognized statistical technique for dealing with extremes. While these envelopes can be viewed as a worst-case scenario so it will protect against uncertainty by setting all returns to their lowest (highest) possible values the end points of the intervals, this approach can produce implausible scenarios. In practice, there is usually some form of dependence among future returns, and it rarely happens that all uncertain returns take their worst-case values simultaneously. It may therefore be desirable to incorporate some kind of variability and correlation of asset returns. A second limitation with polyhedral uncertainty set comes from the fact that for some no-linear instruments, such as options, the maximum loss may not occurs when the underlying asset hits its envelope scenario. For instance, a stress scenario for a portfolio with an ATM long straddle position is a scenario where the underlying remains unchanged. However for an outright position this scenario generates almost no losses. The literature calls this the dimensional dependence of maximum loss and suggests that effective stress scenarios are not only made up of those exhibiting extreme variation.

Therefore a *second step* is taken in order to overcome these two limitations. So, let $\Gamma(\boldsymbol{\alpha}, \Phi)$ denotes a particular dynamic data generating process (DGP) which depends on a parameter vector $\boldsymbol{\alpha}$ and an arbitrary distribution Φ. The second step consists in generating trajectories for each risk factor along the holding period which are expected to fill as uniformly as possible the state space comprised by the upper, $S(t)_u$, and lower bounds, $S(t)_l$:

$$\mathcal{S} = \{S(t) : S(t)^i_l \leqslant S(t) \leqslant S(t)^i_u, \quad \forall\, t \in [1, T_{max}] \quad \text{and} \quad i = 1, \ldots, I\} \tag{17}$$

where T_{max} is the maximum holding period and I is the number of risk factors.

Note that given a particular realization ω there is no sufficient condition imposed to assure that $S(t, \omega) \in \mathcal{S}$. Therefore an additional structure is required to meet this condition. One way of assuring that $S(t, \omega) \in \mathcal{S}$ is by recurring to the concept of *exit time*:

$$\tau(\omega) := \inf\{t \geqslant 0 | S(t, \omega) \notin \mathcal{S}\} \tag{18}$$

Therefore, it is possible for every trajectory generated by Monte Carlo to construct the *simulated stopped process* as:

$$\bar{S}(\omega, t) = \begin{cases} S(\omega, t), & \text{if} \quad \tau > t \\ S(\omega, \tau), & \text{if} \quad \tau \leqslant t \end{cases} \tag{19}$$

In this situation the use of stopped process is an artifice that assures for any ω the *simulated stopped process* is by construction contained in $\mathcal{S}$.

It is worth to mention that the dependence structure among risk factors will be part of the DGP specification. So the plausibility among risk factors will be embedded into the dependence structure established, in other words, by assuming a dependence structure where tail dependence is present it is expected that simulated trajectories will exhibit average dependence as well as extreme co-movements. To illustrate how the *simulated stopped process* behaves we present two pictures: as one can see in Fig. 1 there are three simulated paths which are not necessarily contained in $\mathcal{S}$ while one can be seen in Fig. 2 how the simulated stopped process would look like.

It's worth mentioning that modifying the properties of a set of simulated trajectory is a method which has already been adopted in other studies. Avellaneda et al. (2001) in the context of option pricing proposed an algorithm for correcting price-misspecification and finite sample effects arising during the Monte Carlo simulation by assigning probability weights to the simulated paths. Differently from Avellaneda et al. (2001) who changed the simulated price distribution, our approach need to assure $S(t, \omega) \in \mathcal{S}$ pathwise rather than in distribution.

Finally, even though the definition of the worst-case scenarios is based on GPD other fat tail technique could be selected, such as asymmetric t-distribution proposed by Hansen (1994). On the other hand, if the statistical method seems too rigid for describing the worst-case scenario and some degree of subjectiveness is required, there is no limitation in replacing these statistical methods by hand-picked scenarios if necessary. The usage of hand-picked stress scenarios, although questionable among academics,

[3] The case $\xi > 0$ corresponds to the heavy-tailed distributions whose tails decay like power functions such as the Pareto, Student's t, Cauchy, Burr, loggamma and Frechet distributions. The case $\xi = 0$ corresponds to distributions like the normal, exponential, gamma and lognormal, whose tails decay exponentially; we call such distributions thin-tailed. The final group of distributions are short-tailed distributions $\xi < 0$ with a finite right endpoint like the uniform and beta distributions.

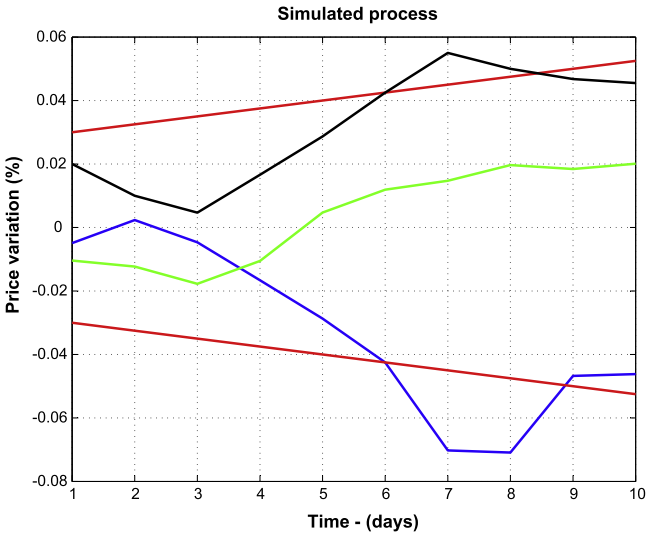

Fig. 1. Sample trajectories drawn using Monte Carlo, $S(\omega,t)$. Red lines are the envelope scenarios. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

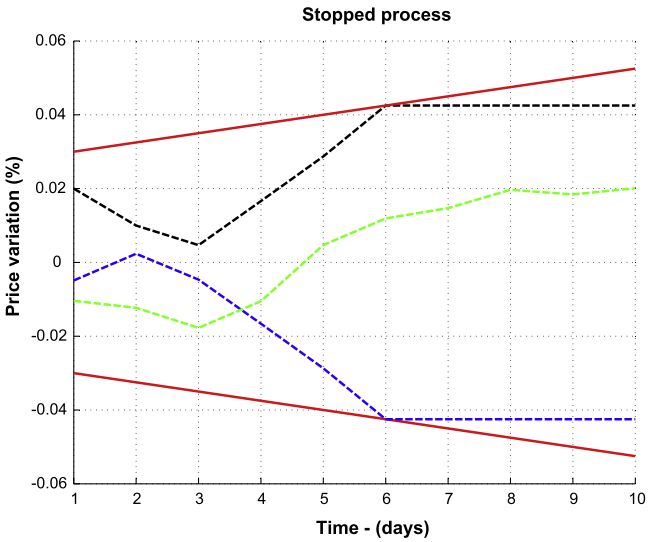

Fig. 2. Sample trajectories drawn using simulated stopped process, $\bar{S}(\omega,t)$. Red lines are the envelope scenarios. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

is commonly implemented by practitioner because it produces outcomes which are intuitive and easy to interpret[4]. In 2005 the Committee on the Global Financial System (2005) set up a group to review what banks and securities firms perceived to be material risks for them at that time, based on the stress tests they were running. They point out that hand-picked scenarios are potentially more relevant to determining the risk profile of the firm, but they are labor-intensive and involve considerably more judgment. With this in mind, some firms invite senior managers, front desks and economists to these discussions in order to secure objectivity and support for the scenario-setting process. Due to its underlying properties and historical usage it is reasonable to expect that this method will continue to be used, but from now on combined with quantitative tools of a probabilistic nature.

[4] Senior management tends to prefer simple and intuitive models rather black-box tools, particularly after the credit crunch where decisions were made based on models, which later were proved as very imprecise.

4.1. The curse of dimensionality and maximum entropy trajectories

A typical Monte Carlo experiment for risk measurement usually involves thousands of samples, which, combined with a 10 days horizon can yield millions of scenarios for just one portfolio. On the other hand, as pointed by Avellaneda and Cont (2013), the optimal closeout strategy for unwinding a portfolio can be solved using Linear Programing (LP) techniques and the problem's dimensionality grows as $\mathcal{O}(N \times T_{max})$, where N is the number of scenarios and T_{max} denotes the maximum horizon for unwind a given portfolio. As one can easily observe, the computational power required to solve the problem increases as the number of scenarios grow, and therefore the number of scenarios should be carefully defined to maintain the optimization problem tractable.

Our approach for maintaining the problem of finding the solution of the closeout strategy tractable consists in choosing a subset of paths from the original Monte Carlo experiment which present the highest variation (or volatility) along the holding period. Therefore, our approach consists in properly choosing a second plausible set $\mathcal{S}^{\star}, \mathcal{S}^{\star} \subset \mathcal{S}$, formed by those trajectories with highest entropy. The entropy quantifies the average uncertainty, disorder or irregularity generated by a process or system per time unit and, it is the primary subject of fundamental results in information and coding theory (Shannon's noiseless coding theorem) and statistical mechanics (second law of thermodynamics). It is not surprising, therefore, that this notion, appropriately generalized and transformed, is ubiquitous in many fields of mathematics and science when randomness or random-like behavior is at the heart of the theory or model being studied. The application of entropy in finance can be regarded as the extension of the information entropy and the probability entropy. It can be an important tool in portfolio selection and asset pricing. Philippatos and Wilson (1972) were the first two researchers who applied the concept of entropy to portfolio selection while Zhou et al. (2013) review the concepts and principles of entropy and provide an updated list of papers that have applied entropy in finance over the last 30 years.

With respect to estimating the entropy $\mathcal{E}(t)$ of a given time series, a good starting point might be the Shannon n-gram (block) entropy suggested in Ou (2005) and Philippatos and Wilson (1972). In a very general case, a given sequence of N observations $x_1, x_2, \ldots, x_N$ is first partitioned into subvectors of length L with an overlap of one time step, which are further divided into subtrajectories (delay vectors) of length $n < L$. Real-valued observations $x_i \in \mathbb{R}$ are discretised by mapping them onto λ non-overlapping intervals $A^{\lambda}(x_i)$. The precise choice of those intervals (also called states) denoted by A^{λ} would depend on the range of values taken by x_i. Hence a certain subtrajectory $x_1, x_2, \ldots, x_n$ of length n can be represented by a sequence of states $A_1^{\lambda}, A_2^{\lambda}, \ldots, A_n^{\lambda}$. The authors then define the n-gram entropy (entropy per block of length n) to be:

$$\mathcal{E}_n = -\sum_{\chi} p\left(A_1^{\lambda}, A_2^{\lambda}, \ldots, A_n^{\lambda}\right) \log_{\lambda} p\left(A_1^{\lambda}, A_2^{\lambda}, \ldots, A_n^{\lambda}\right) \tag{20}$$

In the above equation the summation is done over all possible state sequences $\chi \in A_1^{\lambda}, A_2^{\lambda}, \ldots, A_n^{\lambda}$. The probabilities $p(A_1^{\lambda}, A_2^{\lambda}, \ldots, A_n^{\lambda})$ are calculated based on all subtrajectories $x_1, x_2, \ldots, x_n$ contained within a given subvector of length L. In general, processes or variables with entropy equal to 0 are deterministic, in our context, trajectories with block-entropy close to zero should present low price variation.

In the same way as in the historical VaR approach, where returns are sorted in ascending order and the VaR number at a k per cent level of confidence is determined as the $(100-k)$ percentile of these sorted returns, trajectories will be chosen based on its sorted entropy. That is, entropies, along with their assigned trajectory, are sorted in ascending order: $\mathcal{E}_1(Z_1) \leqslant \ldots \leqslant \mathcal{E}_m(Z_j)$ where m is used to denote an ordering based on each entropy trajectory and $\mathcal{E}_i(Z_j)$ denotes the n-gram entropy for the $j-th$ trajectory Z_j calculated as defined by (20). Therefore, we sort the simulated stopped trajectories based on their entropy, starting from the smallest value (probably close to zero) to the highest entropies.

Therefore after raking every simulated trajectory based on its block entropy our strategy consists in choosing the $k-th$ largest values for every risk factor to form the final stress scenarios, $\mathcal{S}^{\star}$. For those risk factors which are jointly modeled we repeat this calculation for every factor in the group choosing those realization ω_i for each factor l. The final set will be formed by the union of all realizations of every factor. So, if we have two risk factors and we choose only the two largest values for each factor, we will end up with 4 scenarios for each risk factors, in other words, if realization ω_1 is the largest for factor one and ω_3 is the largest for factor two we keep both realization in our final set of stress scenarios $\{\omega_1, \omega_3\}$ for each factor. In this way it is possible to preserve the dependence structure presented in the data.

To illustrate this concept, the pictures below present two sets of simulated trajectories grouped according to their block-entropy over a 10-days horizon:

As one can see in Fig. 3, 1000 trajectories were simulated and later classified as low entropy according to the block-entropy estimator. It can be observed that the vast majority of the trajectories are contained in the [−0.05, 0.05] interval, with just a few exceptions hitting the envelope scenarios. This lack of coverage means that a number of price paths will not be taken into consideration for determining the optimal closeout strategy, which might lead to risk underestimation. On the other hand, as can be seen in Fig. 4, 1000 simulated trajectories were drawn and ranked as high entropy. In this second set one can observe, as expected, that trajectories with highest entropy indeed presented wider variations. Additionally, these trajectories performed better in making the coverage of state space formed by the envelopes scenarios more uniform.

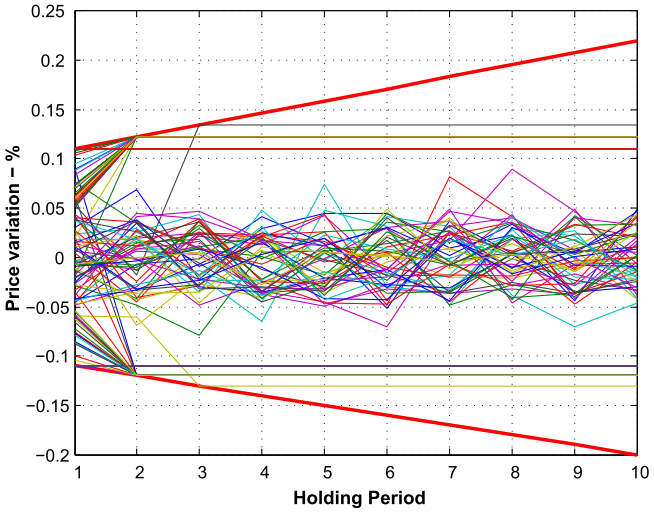

Fig. 3. Low entropy trajectories.

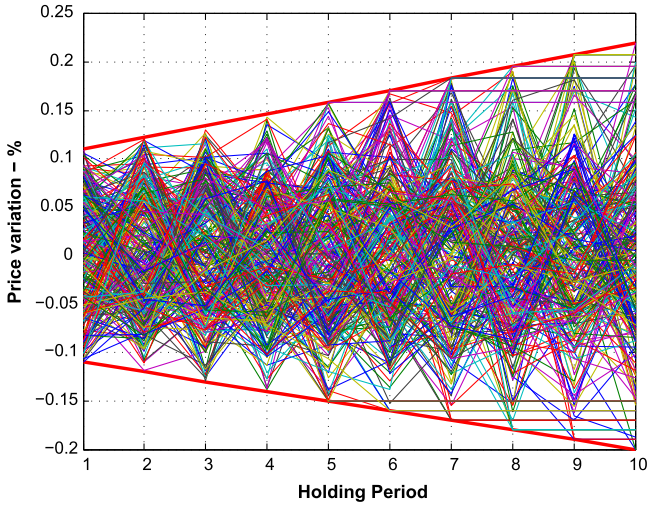

Fig. 4. High entropy trajectories.

5. Application: generating scenarios for $\mathcal{S}$

Once the framework for generating $\mathcal{S}^{\star}$ has been presented in the previous section, a more detailed walk through description is provided here for methodological reasons. The first step involves defining the envelope scenarios:

$$S(t)_l^i = \inf\{x \in \Re : F_{(\xi,\sigma,l)}(x) \geqslant \alpha_l\}, \quad \forall\, t \in [1, T_{max}], \quad i = 1, \ldots, I \tag{21}$$

$$S(t)_u^i = \inf\{x \in \Re : F_{(\xi,\sigma,u)}(x) \geqslant \alpha_u\}, \quad \forall\, t \in [1, T_{max}], \quad i = 1, \ldots, I \tag{22}$$

Here $F_{(\xi,\sigma,\cdot)}(x)$ is the Generalized Pareto Distribution:

$$F_{(\xi,\sigma,\cdot)}(x) = 1 - \left(1 + \frac{\xi x}{\sigma}\right)^{-\frac{1}{\xi}} \tag{23}$$

where the last function argument denotes that this function is independently defined to each tail. The support is $z \geqslant 0$ for $\xi \geqslant 0$ and $0 \leqslant z \leqslant -1/\xi$ for $\xi < 0$.

Once the parameters of the GPD have been estimated, the upper and lower bounds of (13) are obtained using the percentiles of (23) for the level of severity of $\alpha_l = 0.0004$ and $\alpha_l = 0.9996$.

Next we proceed by specifying the DGP, $\Gamma(\alpha, \Phi)$ that will be used for generating paths of the simulated stopped process, $\bar{S}(\omega, t)$:

$$\mathcal{S} = \{\bar{S}(\omega,t) : S(t)_l^i \leqslant \bar{S}(\omega,t) \leqslant S(t)_u^i, \quad \forall\, t \in [1, T_{max}], \quad i = 1, \ldots, I\} \tag{24}$$

In order to model dynamic volatility and fat tails together, the method implemented for generating asset returns follows a generalization of the two-step procedure of McNeil and Frey (2000) where assets volatility are modeled by GARCH-type models and tails distributions of GARCH innovations are modeled by EVT.

Therefore, following the same spirit of the model employed by McNeil and Frey (2000), assume that the Data Generating Process describing asset returns is given by:

$$r_t = \omega + \rho r_{t-1} + \sigma_t \epsilon_t \tag{25}$$

where $\{\epsilon_t\}$ is a sequence of independent Student's t distribution with v degrees of freedom to incorporate fat tails often presented in asset returns.

Furthermore, assume that:

$$\sigma_t^2 = \alpha_0 + \sum_{j=1}^{q} \alpha_j \epsilon_{t-j}^2 + \sum_{j=1}^{p} \beta_j \sigma_{t-j}^2 \tag{26}$$

Eqs. (25) and (26) define the standard AR-GARCH(p,q) model of Bollerslev (1987). Parameter restrictions are required to ensure positiveness of the conditional variance σ_t in (26). Assuming

$\alpha_j \geqslant 0, j = 1, \ldots, q$, and $\beta_j \geqslant 0, j = 1, \ldots, p$ is sufficient for this. Both necessary and sufficient conditions were derived by Nelson and Cao (1992). A more sophisticated formulation where $p > 1$ and $q > 1$ or asymmetric models such as Glosten et al. (1993) or tempered stable GARCH models as Menn and Rashev (2005) could also be adopted here. In this paper we shall concentrate on (26) assuming $p = q = 1$. This is done for two reasons. First, the GARCH(1,1) model is by far the most frequently applied GARCH specification. Second, we want to keep a more parsimonious specification once we are handling a large scale problem.

The second step according to McNeil and Frey (2000) requires modeling the marginal distributions for standardized innovations $z_t := r_t/\sigma_t$ of each risk factor. To accomplish that, a non-parametric smooth kernel density estimator is implemented for describing the center of the data while the Generalized Pareto Distribution for the upper and lower tails above the threshold is adopted.

To incorporate the dependence structure among different risk factors we adopted the "*t*-copula" which has received much attention in the context of modeling multivariate financial return data. A number of papers such as Mashal and Zeevi (2002) and Breymann et al. (2003) have shown that the empirical fit of the *t* copula is generally superior to that of the so-called Gaussian copula. One reason for this is the ability of the *t* copula to better capture the phenomenon of dependent extreme values, which is often observed in financial return data.

A *d*-dimensional copula is a d-dimensional distribution function on $[0,1]^d$ with standard uniform marginal distributions. Sklars Theorem (see for example Nelsen (1999), Theorem 2.10.9) states that every density function *F* with margins $F_1, \ldots, F_d$ can be written as

$$F(x_1, \ldots, x_d) = C(F_1(x_1), \ldots, F_d(x_d)) \tag{27}$$

for a *t*-copula, which is uniquely determined on $[0,1]^d$ for distributions *F* with absolutely continuous margins. Conversely any copula *C* may be used to join any collection of univariate dfs $F_1, \ldots, F_d$ using (27) to create a multivariate df F with margins $F_1, \ldots, F_d$.

The *t* copula is defined as:

$$C(\mathbf{u}) = \int_{-\infty}^{\mathbf{t}_\nu^{-1}(\mathbf{u_1})} \cdots \int_{-\infty}^{\mathbf{t}_\nu^{-1}(\mathbf{u_d})} \frac{\Gamma(\frac{\nu+\mathbf{d}}{2})}{\Gamma(\frac{\nu}{2})\sqrt{(\nu\pi)^{\mathbf{d}}|\mathbf{P}|}} \left(\mathbf{1} + \frac{\mathbf{x}'\mathbf{P}^{-\mathbf{1}}\mathbf{x}}{\nu}\right)^{-\frac{\nu+\mathbf{d}}{2}} \mathbf{dx} \tag{28}$$

where t_ν^{-1} denote the quantile function of a standard univariate t_ν distribution and *P* is the correlation matrix implied by the dispersion matrix Σ.

The *t* copula has a number of desirables features, such as tail dependence different from zero. The coefficients of tail dependence provide asymptotic measures of the dependence in the tails of the bivariate distribution of (X_1, X_2). The coefficient of upper tail dependence of X_1 and X_2 is

$$\lim_{q \to 1} \mathbb{P}(X_2 > F_2^{-1}(q) | X_1 > F_1^{-1}(q)) = \lambda_u \tag{29}$$

provided a limit $\lambda \in [0,1]$ exists. *Mutatis Mutantis* the lower tail dependence can also be obtained. For the copula of an elliptically symmetric distribution like the *t* the two measures λ_u and λ_d coincide, and a simple formula was found by Embrechts et al. (2001). Additionally, Embrechts et al. (2001) found that for continuously distributed random variables with the t-copula the coefficient of tail dependence is given by:

$$\lambda = 2t_{\nu+1}\left(-\sqrt{\nu+1}\sqrt{1-\rho}/\sqrt{1+\rho}\right) \tag{30}$$

where ρ is the off-diagonal element of *P*.

Although it is quite common to use the word correlation when referring to relations between different risk factors, there are other forms of dependence that are not strictly speaking linear in nature. Hence they are not measured properly using a correlation statistic, which pertains strictly to the realm of linear relations. From Eq. (30) we can see that even if correlation between any two risk factor is null there still exist a positive probability of extreme joint movement. This is a desirable feature in risk management because under extreme market conditions the co-movements of asset returns do not typically preserve the linear relationship observed under ordinary conditions.

5.1. Results

In this subsection we have summarized some aspects of statistical inference, as well as, the outcomes for the models that we have proposed and estimated. Our dataset contains daily closing spot prices of the two most important Brazilian stock indexes, Ibovespa and IBrX-50 (indexes compiled by BM&FBOVESPA made up of securities that are the most actively traded and liquid in Brazil) and the S&P 500 over the period from January 2, 2002 to October 10, 2014. Although the dimension of the problem is small, the example illustrates the qualitative properties of the proposed methods.

The first step requires estimating the parameters for the AR-GARCH processes as defined by Eqs. (25) and (26). In general non-Gaussian GARCH models parameters are estimated by quasi-maximum likelihood (QMLE). Bollerslev and Wooldridge (1992) showed that the QMLE still delivers consistent and asymptotically normal parameter estimates even if the true distribution is non-normal. In our approach the efficiency of the filtering process, i.e., the construction of z_t, is of paramount importance. This is so because the filtered residuals serve as an input to both the EVT tail estimation and the copula estimation. This suggests that we should search for an estimator which is efficient under conditions of non-normality. Therefore, as pioneered by Bollerslev (1987) and adopted here, model's parameters can be estimated by maximizing the exact conditional t-distributed density with ν degrees of freedom rather than an approximate density.

Having completed the calibration of GARCH parameters and hence obtained sequences of filtered residuals we now consider estimation of the tail behavior by using a Pareto distribution for the tails and the Gaussian kernel for the interior of the distribution. A crucial issue for applying EVT is the estimation of the beginning of the tail. Unfortunately, the theory does not say where the tail should begin. We know that we must be sufficiently far out in the tail for the limiting argument to hold, but we also need enough observations to reliably estimate the parameters of the GPD. There is no correct choice of the threshold level. While McNeil and Frey (2000) use the "mean-excess-plot" as a tool for choosing the optimal threshold level, some authors, such as Mendes (2005), use an arbitrary threshold level of 90% confidence level (i.e. the largest 10% of the positive and negative returns are considered as the extreme observations). In this paper we define upper and lower thresholds such that 10% of the residuals are reserved for each tail.

Estimating the parameters of a copula or the spectral measure is an important part of our framework. The literature on this topic provides two ways of estimating t-copula parameters: a fully parametric method or a semi-parametric method. The first method, that has been termed as the inference functions for margins (IFM) method Joe (1996), relies on the assumption of parametric univariate marginal distributions. The parameters of the margins are first estimated, and then each parametric margin is plugged into the copula likelihood, and this full likelihood is maximized with respect to the copula parameters. The success of this method obviously depends upon finding appropriate parametric models for the margins, which may not always be so straightforward if they show evidence of fat tails and/or skewness. Alternatively, without

making any substantial assumptions for the margins, the univariate semi-parametric cumulative distribution functions can by plugged into the likelihood to yield a canonical maximum likelihood (CML) method, (see Romano (2002)). Here the semi-parametric approach is employed for estimating parameters of the t-copula. To do so, firstly the marginal distributions are obtained by transforming the standardized residuals to uniform variates by the semi-parametric empirical CDF derived above, then estimate the t-copula parameters using a canonical maximum likelihood (CML).

Maximum likelihood estimates of the parameters for each estimated model are presented in Table 1 along with asymptotic standard error in parentheses:

For each of our three stock indexes, Table 1 reports the results of the maximum likelihood estimation of each model's parameters. As the outcomes are very similar across indexes, the following comments apply for all of them. First, the mean equations present weak, although statistically significance, evidence on the first daily autocorrelation of the index returns[5]. Second, for the three indexes, GARCH coefficients α_1 and β_1 are significant at the 1% and their sum is less than one implying that the GARCH model is stationary, thought the volatility is fairly persistent since $(\alpha_1 + \beta_1)$ is close to one. Third, the estimated number of degrees of freedom of the conditional t-distribution is smaller than 13 which suggests that the returns on the selected indexes are conditionally non-normally distributed.

In the middle of the Table 1 the ML estimation of the GPD parameters fitted to the observations in excess over the thresholds are presented. For all indexes, the estimated ξ is found to be positive for both lower tail and upper tail of the filtered residuals distribution. This indicate that the tails on both sides of the distribution are heavy and all moments up to the forth are finite. Finally, the parameters describing the dependence are presented at bottom of the table. The number of degrees of freedom estimated for the t-copula is equal to 5.27 which is a very small value given the fact that for $\nu = 4$ the kurtosis of a Student-t random variable with ν degrees of freedom is infinite.

Once the parameters in each model have been estimated the next step consists of generating maximum entropy trajectories for up to 10 days by Monte Carlo simulation. Thus, based on parameters of Table 1, 100,000 samples for the simulated stopped process were drawn for each risk factor over a 10-days horizon forming the set $\mathcal{S}$. To illustrate this process, 10,000 trajectories for each risk factor are grouped according to its entropy and displayed on figure below:

The uncertainty set $\mathcal{S}^\star$ formed by 10,000 trajectories with highest block-entropy can be seen in Figs. 5–7. The red dotted lines represent the envelope scenario which were estimated using Extreme Value Theory and a confidence level of $\alpha_l = 0.04\%$ and $\alpha_u = 99.96\%$. We observe that paths in this set cover almost all possible variation along the holding period which is desirable from a risk viewpoint because in this case the optimal closeout strategy will be determined considering a broad range of future outcomes for each risk factor and no potentially harmful price variation is left aside.

For comparative purposes, other 10,000 trajectories ranked with low entropy are also presented as one can see in Figs. 8–10 and here some price variations are not accounted for determining the optimal closeout strategy and therefore will potentially underestimate the margin requirements. In addition, we notice that from the 100,000 trajectories that were originally generated only 244 of them hit the envelopes scenarios (boundaries of our problem) and therefore were stopped. This value represents a very high acceptance rate for our sampling algorithm.

Finally, it is worth to mention one advantage of this method over Breuer et al. (2009) which consists in not imposing any trade-off between plausibility and severity of scenarios. If one needs to work with more severe scenarios, it is only necessary to set higher WCS and the Monte Carlo simulation will fill the state space comprised by the new envelope of more severe plausible stress scenarios.

Table 1
Estimates of GARCH(1,1), tail distribution and t-copula.

	Ibovespa	IBrX 50	S&P 500
	AR-GARCH		
ω	8.04E-4	9.2384E-04	7.5074E-04
	(2.7378E-04)	(2.5220E-04)	(1.4626E-04)
ρ	−0.0029	0.0252	−0.0643
	(0.0196)	(0.0196)	(0.0199)
α_0	4.90E-06	4.5790E-06	1.23E-06
	(1.4098E-06)	(1.2544E-06)	(3.5279E-07)
α_1	0.0617	0.0634	0.0898
	(0.0109)	(0.0088)	(0.01)
β_1	0.9222	0.9193	0.9035
	(0.0086)	(0.0111)	(0.011)
t(ν)	12	10	7
	(2.3)	(1.95)	(0.97)
	Lower Tail		
ξ	0.083	0.07466	0.057
	(1.55E-2)	(5.47E-8)	(8.36E-6)
σ	0.7001	0.6901	0.7288
	(1.21E-3)	(5.17E-9)	(4.26E-5)
	Upper Tail		
ξ	0.1387	0.1396	0.1072
	(1.55E-2)	(5.47E-8)	(8.36E-6)
σ	0.6175	0.6412	0.6206
	(1.21E-5)	(5.17E-9)	(4.26E-5)
	t-copula		
Ibovespa	1	0.9686	0.6106
IBrX 50	0.9686	1	0.6082
S&P 500	0.6106	0.6082	1
DoF		5.27	

5.2. Robustness check

Even though t-copula presents desirable properties for handling high dimension distributions of financial variables, the dependence structure among pairs of variables might vary substantially, ranging from independence to complex forms of non-linear dependence, and, in the case of the t-copula, all dependence is captured by only two parameters, the correlation coefficients and the number of degrees of freedom. Due to this potential limitation, we assess the t-copula results comparing its outcomes with a more flexible structure given by a fast-growing technique known as pair-copula originally proposed by Joe (1996).

Pair-copulas, being a collection of potentially different bivariate copulas, is a flexible and very appealing concept. The method for construction is hierarchical, where variables are sequentially incorporated into the conditioning sets, as one moves from level 1 (tree 1) to tree $d-1$. The composing bivariate copulas may vary freely, from the parametric family to the parameters values.

On the other hand, for high-dimensional distributions, there is a significant number of possible pair-copulae constructions. For example, there are 240 different constructions for a five-dimensional density. To help organizing them, Bedford and Cooke (2001, 2002) have introduced a graphical model denoted the *regular vine*. The class of regular vines is still very general and embraces a large number of possible pair-copula decompositions. Consider again the joint distribution F with density f and

[5] Here a word of caution about the interpretation of this result is required. We are not suggesting the returns are predictable or exhibit some form of time series momentum as investigated by Moskowitz et al. (2012) but that the AR(1) filter is required to remove memory in the data.

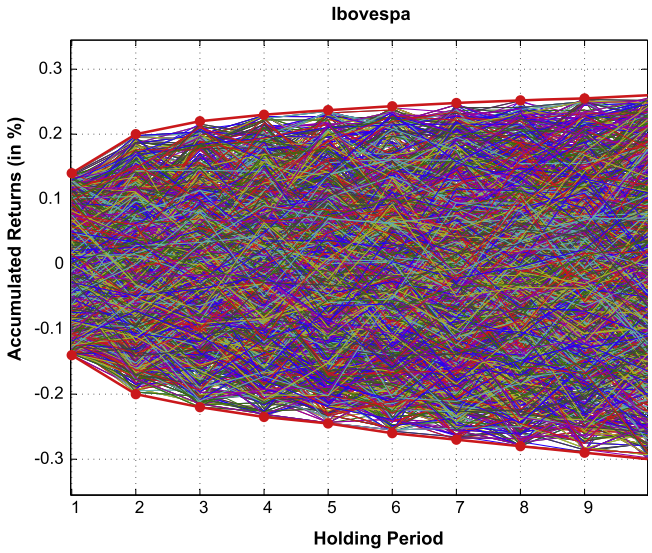

Fig. 5. High entropy trajectories for Ibovespa generated using Monte Carlo, $\bar{S}(\omega, t)$.

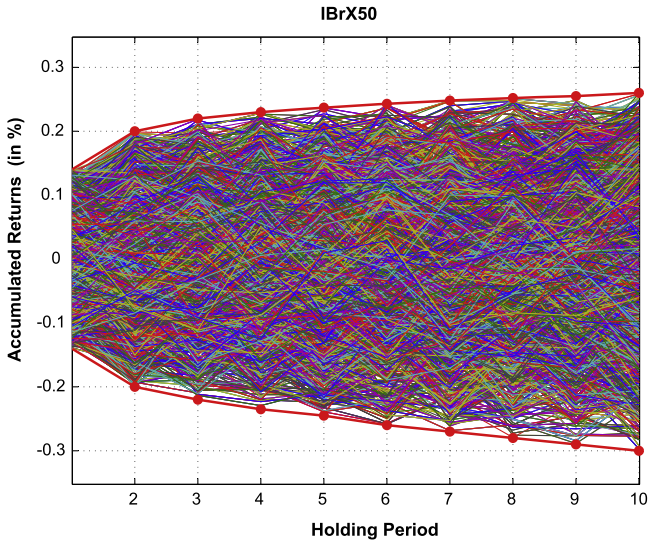

Fig. 6. High entropy trajectories for IBrX50 generated using Monte Carlo, $\bar{S}(\omega, t)$.

with strictly continuous marginal c.d.f.s $F_1, \ldots, F_d$ with densities f_i. First note that any multivariate density function may be uniquely (up to relabel of variables) decomposed as:

$$f(x_1, \ldots, f_d) = f_d(x_d) \cdot f(x_{d-1}|x_d) \cdot f(x_{d-2}|x_{d-1}, x_d) \ldots f(x_1|x_2, \ldots, x_d) \tag{31}$$

The conditional densities in (31) may be written as functions of the corresponding copula densities. That is, for every j:

$$f(x|v_1, v_2, \ldots, v_d) = C_{xv_j|\mathbf{v}_{-\mathbf{j}}}\left(F(x|\mathbf{v}_{-\mathbf{j}}), (F(x|\mathbf{v}_{-\mathbf{j}})\right) \cdot f(x|\mathbf{v}_{-\mathbf{j}}) \tag{32}$$

where $\mathbf{v}_{-\mathbf{j}}$ denotes the d-dimensional vector $\mathbf{v}$ excluding the jth component and $C_{xv_j|\mathbf{v}_{-\mathbf{j}}}(\cdot, \cdot)$ represents a bivariate marginal copula density.

In general, under appropriate regularity conditions, a multivariate density can be expressed as a product of pair-copulae, acting on several different conditional probability distributions. It is also clear that the construction is iterative in its nature, and that given a specific factorization, there are still many different

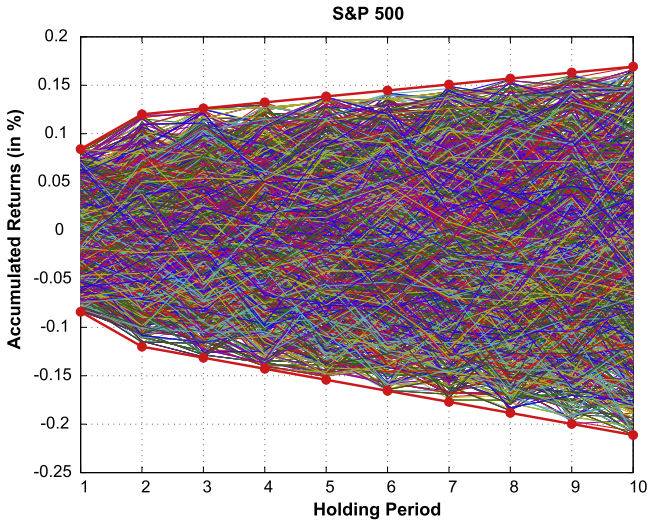

Fig. 7. High entropy trajectories for S&P 500 generated using Monte Carlo, $\bar{S}(\omega, t)$.

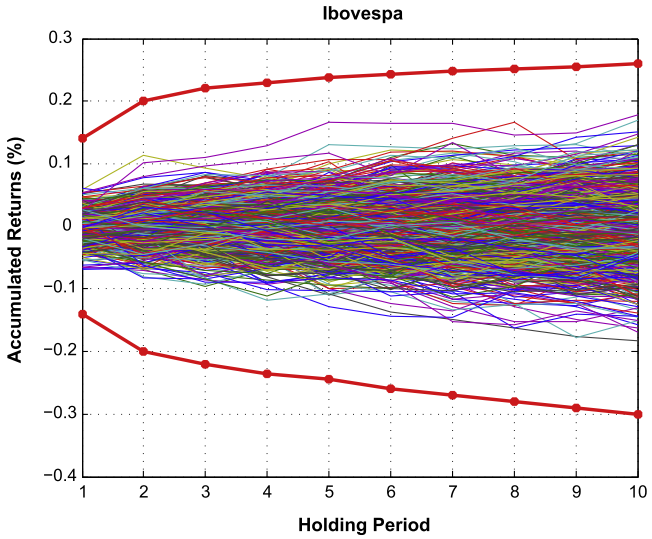

Fig. 8. Low entropy trajectories for Ibovespa generated using Monte Carlo, $\bar{S}(\omega, t)$.

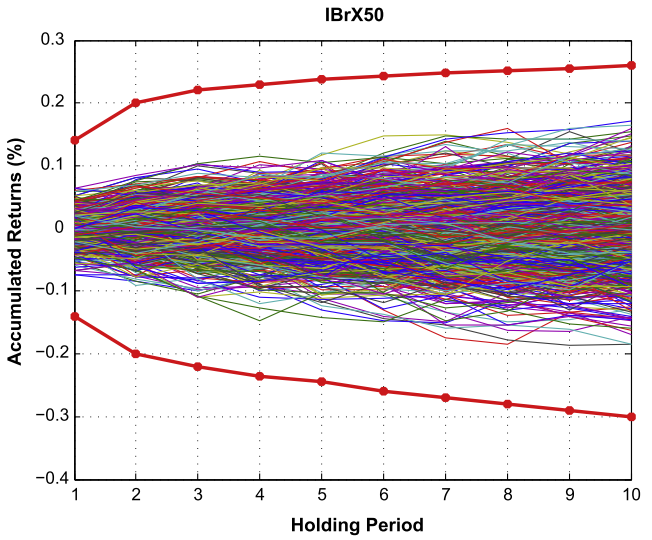

Fig. 9. Low entropy trajectories for IBrX50 generated using Monte Carlo, $\bar{S}(\omega, t)$.

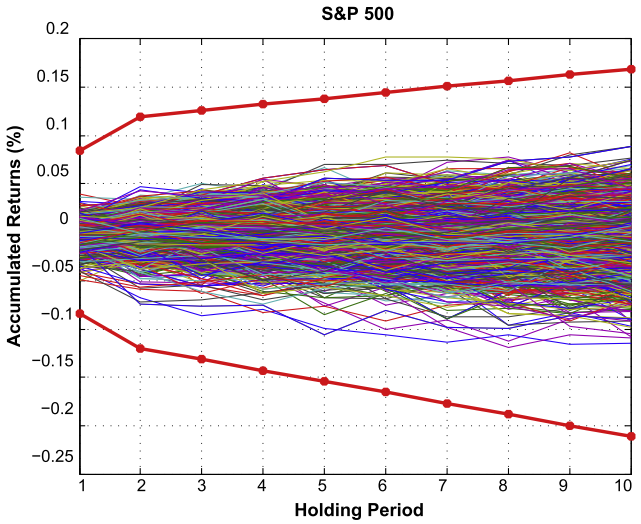

Fig. 10. Low entropy trajectories for S&P 500 generated using Monte Carlo, $\bar{S}(\omega, t)$.

re-parameterizations. Here, we concentrate on D-vine, which is a special case of regular vines (Kurowicka, 2004), where the specification is given in form of a nested set of trees. In this case, it is not essential that all bivariate copulas should belong to the same family. This is exactly what we are searching for, since, recall, our objective is to construct (or estimate) a multivariate distribution which best represents the data at hand, which might be composed by completely different margins (symmetric, asymmetric, with different dynamic structures, and so on) and, more importantly, could be pair-wise joined by more complex dependence structures possessing linear and/or non-linear forms of dependence, including tail dependence, or could be joined independently. For example, one may combine the following types[6] of (bivariate) copulas: Gaussian (no tail dependence, elliptical); t-student (equal lower and upper tail dependence, elliptical); Clayton (lower tail dependence, Archimedean); Gumbel (upper tail dependence, Archimedean); BB8 (different lower and upper tail dependence, Archimedean).

As summarized at appendix A, we performed a comparative study among *t*-copula and 18 possible pair copulae and we found that even though a pair copula presents a better fit in our example its results are not substantially different from a t-copula when applied to measure tail risk in a portfolio. Additionally, despite its appealing flexibility for describing dependence among risk factors, pair copula requires a additional step where one should specify the dependence structure itself, which in a real world application, might be a high dimension problem. Therefore, we decided to move forward in our study with a more parsimonious construction given by the *t*-copula.

6. Comparison among methods of determining margin requirements

In this last session we want to show the performance of methodology to generate trajectories with maximum entropy applied to the problem that motivated its development. Therefore maximum entropy trajectories were generated according to the framework established in the previous sections and shall be used as part of the robust optimization problem that a Central Counterparty solves when employing the concept of optimal closeout strategy for determining margining requirements. To illustrate our results, 7 synthetic portfolios were formed as outlined below (See Table 2):

Table 2
Synthetic portfolios.

	Strategy	Instrument type	Description
Portfolio 1	Outright	Futures	Long 1000 Ibovespa Futures
Portfolio 2	Outright	Futures	Short 1000 IBrX-50 Futures
Portfolio 3	Outright	Futures	Long 1000 S&P 500 Futures
Portfolio 4	pairs trade	Futures & ETF	Long 1000 Ibovespa ETF & Short 1000 IBrX-50 Futures
Portfolio 5	pairs trade	Futures	Long 1000 Ibovespa Futures & Short 1000 S&P 500 Futures
Portfolio 6	pairs trade	Futures & NDF	Long 1000 Ibovespa NDF & Short 1000 IBrX-50 Futures
Portfolio 7	pairs trade	Futures & NDF	Long 1000 Ibovespa NDF & Short 1000 S&P 500 Futures

These portfolios were created to reproduce different asset classes and strategies that are common in multi-asset or multi-market CCPs. The first three portfolios are outright positions using stock indexes futures. Portfolios 4 and 5 illustrate two pairs of trading strategies using exchange-traded instruments. While portfolio 4 is a market-neutral intra-sector pairs trade using two Brazilian indexes, portfolio 5 is an inter-sector pairs trade where the investor would be long the sector that has the best outlook and short the sector that could be vulnerable to a downturn. Finally, portfolios 6 and 7 share the same economic fundamentals of the two previous ones but differ in the type of asset invested, where an OTC contract is combined with an Exchange-Trade Derivative.

In addition, to reproduce the liquidity constraints that a CCP can face during the closeout process, we establish that only 500 contracts can be traded daily for each contract. A second constraint that can arise in real life problems that was included here is a minimum time interval between the moment a default is detected and the beginning of the execution of the default procedures. For futures contracts and ETFs, which are exchange-traded contracts, we set that the execution starts in T + 1 of the default, on the other hand, for OTC instruments, NDFs, the CCP will typically organize an auction among clearing participants and these instruments will not be able to be liquidated before T + 5 of the default. A simplification adopted in order to make the interpretation of the results easier was to normalize all prices to $ 100.

For these 7 synthetic portfolios we calculated, using the CORE approach proposed by Vicente et al. (2015), the amount of collateral required when the uncertainty set is constructed as described in the previous sections. The results obtained, called baseline model, are compared in various dimensions. First we maintain the original robust optimization problem (optimal closeout and discrete minimax problem) but the uncertainty set is modified in two ways: (i) maximum entropy trajectories are replaced by a set of low entropy trajectories; (ii) historical scenarios[7] – price variation up to 10 days from Jan 2, 2002 to Dec 31 2014. The second dimension analyzed was to replace the discrete minimax problem, where the uncertainty set is given by a finite set of scenarios, by its continuous counterparty (optimal closeout and continuous minimax problem). Here we have implemented three different approaches to describe the uncertainty set: (i) independent intervals as Tütüncü and Koeing (2004); (ii) ellipsoids as proposed by Ben-Tal and Nemirovski (2000); (iii) polyhedral uncertainty as presented in Bertsimas and Pachamanova (2008). Finally, we ignore the optimal

[6] See Joe (1996) for a copula catalog.

[7] Historical Scenarios are the most common approach used in practice for generating scenarios. Clearing House such as ICE, LCH and CME have adopted ten years of historical data for calculating margin requirements

Table 3
Margin requirements across different models.

	Optimal closeout						Naive closeout	
	Discrete minimax			Continuous minimax				
	SMC high entropy	SMC low entropy	Historical	Tütüncü & Koenig	Ben-Tal & Nemirovski	Bertsimas & Pachamanova	Studer - Normal	Studer - t
Portfolio 1	15,979	6816	12,630	13,900	9924	14,057	7599	14,066
Portfolio 2	15,674	7906	14,446	13,900	9534	13,505	7278	13,485
Portfolio 3	9742	5603	8513	8300	7009	9,333	5469	9108
Portfolio 4	5725	4461	3610	21,017	4792	6,813	1912	3518
Portfolio 5	9300	7399	8935	19,113	3592	9,324	5707	10,560
Portfolio 6	6989	5001	5358	27,017	8791	6,904	3825	7036
Portfolio 7	17,100	12,477	15,591	20,149	6927	11,217	11,414	21,120

closeout strategy and pursue the Studer approach (Naive closeout) where the maximum loss is determined considering a fixed time horizon. In this case we have followed the implementation suggested by Flood and Korenko (2013) for two particular elliptical distributions: (i) The multivariate Normal distribution; (ii) the multivariate t-distribution.

The results found on Table 3 are very encouraging and indicates a positive gain of our baseline model over different models and methodologies for determining margin requirements. Portfolios 1–3 comprise outright positions so it is expected that the worst-case scenario should be as close as possible to the envelope. This is true for our baseline model (for portfolio 1: $15,979 = 500 \times 100 \times (-0.1368) + 500 \times 100 \times (-0.1827)$), on the other hand, the value calculated using the low entropy scenarios (6,816) is far away of this value and half of the historical scenario. Portfolios 4–5 comprise liquid instruments and represent a relative trade between two stock indexes and therefore we expect some level of risk offsetting. Our baseline model provides this risk offsetting. In this case the low entropy scenario recognize this dependence among risk factors while the methodology of Tütüncü and Koeing (2004) treat these two risk factors independently and overestimate the risk of these strategies. Portfolios 6–7 are also relative trades just like portfolio 4–5, but this time using instruments of different asset classes, futures and NDFs. It is expected that due to the lag for beginning of the execution of the default procedures these portfolios shall be settle in more than 2 days and consequently this position is subject to a higher price variation risk. Our baseline model recognizes this potential scenario and results in a higher margin requirement. The Studer method with the t distribution provided in some cases an amount of collateral higher than our baseline model, however this arises from the myopic treatment of this approach where all positions are unwind in the end – carrying more risk; the main pitfall of this method relies on the fact that this approach ignores completely the liquidity constraint by assuming that there is sufficient liquidity to buy/sell the defaulter portfolio without impacting the prices. The method which provided results more in line with our baseline models is Bertsimas and Pachamanova (2008). This is not surprising once Bertsimas and Pachamanova (2008) also recognize dependence among risk factors and the shape of the uncertainty set is not elliptical. However, the risk requirement for portfolio 2 using the Bertsimas and Pachamanova (2008) approach is smaller than the historical and this is not desirable once the historical scenarios can be seen as a minimum amount because the margin requirements should be enough to cover at least all prices variation seen in the past.

7. Final remarks

Central Counterparties (CCPs) play an important role in the new regulatory arena. Even though CCPs are intended to reduce systemic risks in the financial system, it must also be recognized that CCPs can create, or contribute to, systemic risks. Default by one or more participants may expose the CCP to both market risk and liquidity risk, since the Clearinghouse must ensure the timely settlement of all the trades it guarantees as the CCP. Market risk following a default event corresponds to the potential difference in value between the rights and obligations of the defaulting participant, given the market conditions prevailing at the time of failure treatment. To mitigate market risk arising from a default, most of CCPs have implemented risk sharing mechanisms throughout the chain of responsibilities. These mechanisms are based on three classic components: defaulter pays (DP), survivor pays (SP) and third party pays (TP). Prudent risk management practices recommends that the first element of the safeguard structure is the collateral posted by clients to cover the risk of their positions and obligations. Therefore in an environment where CCPs are playing a growing role, ensure that margin requirements are enough to cover the resulting losses from the default of a counterparty should be the first concern for all market participants.

The CORE[8] methodology, presented in Vicente et al. (2015), was the first attempt to answer how to organize and schedule the liquidation of heterogeneous portfolio to mitigate risk. Even though Vicente et al. (2015) and Avellaneda and Cont (2013) have formulated the CORE model as a robust optimization problem under convex constraints, neither of them described how uncertainty sets are constructed. Therefore, the main objective of this paper is to fill this gap by proposing a flexible method to construct discrete uncertainty sets, in which each element of this set can be seen as multi-period stress scenarios, that shall be used in the optimization step.

Our approach consists in generating multi-period trajectories via Monte Carlo simulation (MCS) that reproduce the stylized facts present in financial times series. Additionally, in order to avoid the curse of dimensionality usually present in Monte Carlo problems we have resorted to a data reduction scheme, which relies on the concept of maximum block-entropy. Finally, we align ourselves with the recommendation from the Basel Committee on Banking Supervision (2005) for stress tests regarding the plausibility and severity of stress scenarios, as well as avoiding procyclicality of margin requirements. When applied to determine the margin requirements, the present method provides both qualitative and quantitative results that outperform other robust optimization models such as Ben-Tal and Nemirovski (2000) and Bertsimas and Pachamanova (2008).

Finally, it is also worth mentioning that this framework for generating uncertainty sets could be used outside the CCP world for any financial institution that need to conduct systematic stress testing in a dynamic context.

[8] BM&FBOVESPA (Sao Paulo based Clearing House) implemented the CORE model and since August 2014 it has been the model for determining margin requirements for every participant trading derivatives (exchange-traded and OTC – except single-name equity derivatives) in Brazil. As of March, 3 2015 the total of margin requirements calculated using both CORE and the framework described in this paper is approximately USD 70 Billions.

Acknowledgements

The views and opinions expressed herein are those of the author and do not necessarily reflect the views of CETIP. The author is indebted to Julio Thomaz, Luiz Chiyoda and many former colleagues at BM&FBOVESPA. A special thank is due to Luis Vicente who provided valuable suggestions. I am also grateful to Marco Avellaneda, Rama Cont, Claudio Flores and conference participants at the 2015 Financial Engineering and Banking Society (FEBS) and Brazilian Finance Association meeting 2015 for their comments on an earlier version of this paper. The errors that are found herein are mine alone.

Appendix A. Pair copulae

This appendix presents the results of a comparative study performed between pair copulae and *t*-copula. We chose three risk factors to perform our analysis: Ibovespa, IBrX-50 and S&P 500. Even though pair-copulae is a very flexible approach for handling complex dependence among variable it requires the imposition of significant number of pairs combinations. To overcome this issue, we implemented an automated way to select pairs using Akaike criteria. For every pair of copula we estimate its parameters using maximum likelihood and choose the one with smaller AIC. We tested a total of 18 different pair-copulae: independent; Gaussian; t-copula; Clayton; Gumbel; Frank; Joe; BB1; BB6; BB7; BB8; rotated Clayton; rotated Gumbel; rotated Joe; rotated BB1; rotated BB6; rotated BB7 and rotated BB8. We also estimated the standard *t*-copula using the same dataset. We depicts the results on table below:

The configuration which presented the smaller AIC was given by:

In fact, from Table 4 we can observe the superior performance of pair-copulae when considered the statistical criteria. As a reality check of our models we simulated 10,000 samples using each configuration and computed percentiles (0.01–0.99) for each risk factor and compared them with the empirical distribution (see Table 5).

We conclude from Tables 7 and 8 that both methodologies are capable to produce results in line with those observed in practice (Table 6). Finally, we observe that pair-copula provides results very close to those obtained with *t*-copula, however, it embeds an extra and time consuming step formed by the process of obtaining the right configuration of each pair of copula, so we did not find enough elements to replace the parsimonious *t* -copula. It is important to mention, we could find different results for a particular risk factors but for a considerable number of them the results found here prevailed.

Table 4
Fitting results.

	Pair Copula with AIC	*t*-copula
Log-Likelihood	565.8	521.1
AIC	−1121.6	−1038

Table 5
Best pair-copula configuration.

Pair copula	Family	Parameters
IBrX-Ibovespa	BB1	(0.15, 1.29)
S&P -Ibovespa	BB8	(2.45, 0.76)
IBrX-S&P given Ibovespa	Rotated Clayton	−0.073

Table 6
Empirical percentiles for each risk factor.

	0.01	0.99
Ibovespa	−3.87	3.78
IBrX	−3.57	3.76
S&P	−1.48	1.73

Table 7
Losses percentiles for each risk factor using pair-copulae.

	0.01	0.99
Ibovespa	−4.000	4.98
IBrX	−3.282	3.721
S&P	−1.425	1.650

Table 8
Losses percentiles for each risk factor using *t*-copula.

	0.01	0.99
Ibovespa	−4.020	4.988
IBrX	−3.282	3.716
S&P	−1.426	1.652

References

Avellaneda, M. Cont, R. 2013. Close-Out Risk Evaluation (CORE): a new risk management approach for central counterparties. Manuscript Available at SSRN: http://ssrn.com/abstract=2247493.
Avellaneda, M., Buff, R., Friedman, C.A., Grandechamp, N., Kruk, L., Newman, J., 2001. Weighted Monte Carlo: a new technique for calibrating asset-pricing models (2001). Int. J. Theor. Appl. Finance 4, 1–29.
Ben-Tal, A., Margelit, T., Nemirovski, A., 2000. Robust modeling of multi-stage portfolio problems. In: Frenk, H., Roos, K., Terlaky, T., Zhang, S. (Eds.), High-performance Optimization. Kluwer Academic Publishers, Dordrecht, pp. 303–328.
Ben-Tal, A., Nemirovski, A., 2000. Robust solutions of linear programming problems contaminated with uncertain data. Math. Program 88, 411–424.
Bertsimas, D., Gupta, V., Kallus, N., 2013. Data-driven robust optimization. Manuscript.
Bertsimas, D., Pachamanova, D., Sim, M., 2004. Robust linear optimization under general norms. Oper. Res. Lett. 32, 510–516.
Bertsimas, D., Pachamanova, D., 2008. Robust multiperiod portfolio management in the presence of transaction costs. Comput. Oper. Res. 35, 3–17.
Bertsimas, D., Sim, M., 2004. The price of robustness. Oper. Res. 52 (1), 35–53.
Bollerslev, T., 1987. A conditionally Heteroskedastic time series model for speculative prices and rates of return. Rev. Econ. Stat. 69, 542–547.
Basel Committee on Banking Supervision, 2005. International Convergence of Capital Measurement and Capital Standards: A Revised Framework. Technical Report, Bank for International Settlements.
Bedford, T., Cooke, R.M., 2001. Probability density decomposition for conditionally dependent random variables modeled by vines. Ann. Math. Artif. Intell. 32, 245–268.
Bedford, T., Cooke, R.M., 2002. Vines – a new graphical model for dependent random variables. Ann. Stat. 30 (4), 1031–1068.
Bollerslev, T., Wooldridge, J.M., 1992. Quasi-maximum likelihood estimation and inference in dynamic models with time varying covariances. Econ. Rev. 11, 143–172.
Breuer, T., Csiszár, I., 2013. Systematic stress tests with entropic plausibility constraints. J. Bank. Finance 37, 1552–1559.
Breuer, T., Krenn, G., 1999. Stress testing. Guidelines on Market Risk 5, Oesterreichische Nationalbank, Vienna.
Breuer, T., Jandačka, M., Rheinberger, K., Summer, M., 2009. How to find plausible, severe and useful stress scenarios. Int. J. Cent. Banking 3, 205–224.
Breuer, T., Jandacka, M., Menća, J., Summer, M., 2012. A systematic approach to multi-period stress testing of portfolio credit risk. J. Bank. Finance 36 (2), 332–340.
Breymann, W., Dias, A., Embrechts, P., 2003. Dependence structures for multivariate high-frequency data in finance. Quant. Finance 3, 1–14.
Brunnermeier, M., Pedersen, L., 2009. Market liquidity and funding liquidity. Rev. Financial Stud. 22 (6), 2201–2238.
Chen, W., Sim, M., Sun, J., Teo, C., 2010. From CVaR to uncertainty set: implications in joint chance constrained optimization. Oper. Res. 58 (2), 470–485.
Christensen, J.H.E., Lopez, J.A., Rudebusch, G.D., 2015. A probability-based stress test of federal reserve assets and income. J. Monetary Econ. 73, 26–43.
Committee on Payment and Settlement Systems and Technical Committee of the International Organization of Securities Commissions, 2012. Principles for financial markets infrastructures.

Duffie, D., Zhu, H., 2011. Does a Central Clearing Counterparty reduce counterparty Risk? Working Paper n. 46, Rock Center for Corporate Governance at Stanford University.
Embrechts, P., McNeil, A., Straumann, D., 2001. Correlation and dependency in risk management: properties and pitfalls. In: Dempster, M., Moffatt, H. (Eds.), Risk Management: Value at Risk and Beyond. Cambridge University Press.
Flood, M., Korenko, G., 2013. Systematic scenario selection. working paper. Office of Financial Research – U.S Dept of the Treasury.
Galbiati, M., Soramäki, K., 2012. Clearing networks. J. Econ. Behav. Organiz. 83, 609–626.
Glosten, L.R., Jagannathan, R., Runkle, D., 1993. On the relation between the expected value and the volatility of the nominal excess return on stocks. J. Finance 48, 1779–1801.
Hansen, B.E., 1994. Autoregressive conditional density estimation. Int. Econ. Rev. 35 (3), 705–730.
Heller, D., Vause, N., 2012. Collateral requirements for mandatory central clearing of over-the-counter derivatives. BIS Working Papers No 373.
Joe, H., 1996. Families of m-variate distributions with given margins and m(m-1)/2 bi- variate dependence parameters. In: Ruschendorf, L., Schweizer, B., Taylor, M. D. (Ed.), Distributions with Fixed Marginals and Related Topics.
Kamhi, N., 2009. Procyclicality and margin requirements. Bank of Canada Financial System Review.
Kurowicka, D., Cooke, R.M., 2004. Distribution – free continuous bayesian belief nets. In: Fourth International Conference on Mathematical Methods in Reliability Methodology and Practice, Santa Fe, New Mexico.
Mashal, R., Zeevi, A., 2002. Beyond correlation: extreme comovements between financial assets. Unpublished, Columbia University.
McNeil, A.J., Frey, R., 2000. Estimation of tail-related risk measures for heteroscedastic financial time series: an extreme value approach. J. Empirical Finance 7, 271–300.
Mendes, B.V., 2005. Asymmetric extreme interdependence in emerging equity markets. Appl. Stochastic Models Bus. Ind. 21, 483–498.
Menn, C., Rashev, S., 2005. A GARCH option pricing model with α-stable innovations. Eur. J. Oper. Res. 163, 201–209.
Moskowitz, T., Ooi, Y., Pedersen, L., 2012. Time series momentum. J. Financial Econ. 104 (2), 228–250.
Murphy, D., Vasios, M., Vause, N., 2014. An investigation into the procyclicality of risk-based initial margin models. Bank of England Financial Stability Paper 29.
Nelsen, R.B., 1999. An Introduction to Copulas. Springer, New York.
Nelson, D.B., Cao, C., 1992. Inequality constraints in the univariate GARCH model. J. Bus. Econ. Stat. 10, 229–235.
Ou, J., 2005. Theory of portfolio and risk based on incremental entropy. J. Risk Finance 6, 31–39.
Philippatos, G., Wilson, C., 1972. Entropy, market risk, and the selection of efficient portfolios. Appl. Econ. 4, 209–220.
Pirrong, C., 2011 The economics of central clearing: theory and practice. ISDA – Discussion Paper Series.
Polak, G., Rogers, D., Sweeney, D., 2010. Risk management strategies via minimax portfolio optimization. Eur. J. Oper. Res. 207, 409–419.
Romano, C., 2002. Calibrating and simulating copula functions: an application to the Italian stock market, Working paper n. 12, CIDEM.
Steinbach, M., 2001. Markowitz revisited single-period and multi-period mean-variance models. SIAM Rev. 43, 31–85.
Studer, G., 1999. Market risk computation for nonlinear portfolios. J. Risk 1 (4), 33–53.
Tütüncü, R., Koeing, M., 2004. Robust asset allocation. Ann. Oper. Res. 132, 157–187.
Vicente, L., Cerezetti, F., De Faria, S., Iwashita, T., Pereira, O., 2015. Managing risk in multi-asset class, multimarket central counterparties: the CORE approach. J. Bank. Finance 51, 119–130.
Zhou, R., Ru Cai, R., Tong, G., 2013. Applications of entropy in finance: a review. Entropy 15, 4909–4931.