

Portfolio Optimization with Conditional Value-at-Risk Budgets

Kris Boudt^a

Peter Carl^b

Brian G. Peterson^c

^a*K.U.Leuven and Lessius, Naamsestraat 69, 3000 Leuven, Belgium.*

Email: kris.boudt@econ.kuleuven.be

^b*Guidance Capital Management, Chicago, IL. Email: pcarl@gsb.uchicago.edu*

^c*Cheiron Trading, Chicago, IL. Email: brian@braverock.com*

The art of successful portfolio management is not only to be able to identify opportunities, but also to balance them against the risks that they create in the context of the overall portfolio.

- Robert Litterman [1996, p. 73]

Risk budgets are frequently used to estimate and allocate the risk of a portfolio by decomposing the total portfolio risk into the risk contribution of each component position. Many approaches to portfolio allocation use *ex post* methods for constructing risk budgets and take the variance as a risk measure. In this paper, however, we use *ex ante* methods to evaluate the component contribution to Conditional Value at Risk (CVaR) and to allocate risk. The proposed minimum CVaR concentration portfolio draws a balance between the investor's return objectives and the diversification of risk across the portfolio. For a portfolio invested in bonds, equities, and commodities, we show that the minimum CVaR concentration portfolio offers an attractive compromise between the good risk-adjusted return properties of the minimum CVaR portfolio and the positive return potential and low portfolio turnover of an equal-weighted portfolio.

The outline of the paper is as follows. First, we review the definition and estimation of CVaR portfolio budgets. We then describe several portfolio allocation strategies that use the portfolio component CVaR risk

budget as an objective or constraint in the portfolio optimization problem. The paper concludes with an application of the risk budget allocation methodology to optimize portfolios allocating across asset classes.

PORTFOLIO CVaR BUDGETS

DEFINITIONS

The first step in the construction of a risk budget is to define how portfolio risk and its risk contributions should be measured. A naïve approach is to set the risk contribution equal to the stand-alone risk of each portfolio component. This approach is overly simplistic and neglects important diversification or multiplication effects of the component units being exposed differently to the underlying risk factors. Using game theory, Denault [2001] has shown that the only satisfactory risk allocation principle is to measure the risk contribution as the weight of the position in the portfolio times the partial derivative of the portfolio risk R_w with respect to that weight:

$$C_{(i)}R_w = w_{(i)} \frac{\partial R_w}{\partial w_{(i)}}. \quad (1)$$

The standard deviation, VaR and CVaR of a portfolio are all linear in position size. By Euler's theorem we have that for such risk measures the total portfolio risk equals the sum of the risk contributions in (1).

Previous work by Chow and Kritzman [2001], Litterman [1996], Maillard, Roncalli and Teiletche [2010], Peterson and Boudt [2009], and Scherer [2002] study the use of portfolio standard deviation and value-at-risk (VaR) budgets. In his book "Risk budgeting", Pearson [2002, p.7] notes that "value-at-risk has some well known limitations, and it may be that some other risk measure eventually supplants value-at-risk in the risk budgeting process". Unlike value-at-risk, conditional value-at-risk (CVaR) has all the properties a risk measure should have to be coherent and is a convex function of the portfolio weights (see Artzner, Delbaen, Eber and Heath [1999] and Pflug [2000]). Moreover, CVaR provides less incentive to load on to tail risk above the VaR level.

We develop a risk budgeting framework for portfolio Conditional Value at Risk (CVaR). Portfolio CVaR can be expressed in monetary value or percentage returns. Our goal is to apply the CVaR budget in an investment strategy based on quantitative analysis of the assets returns. We therefore choose to define CVaR in percentage returns.

Denote by $r_{w,t}$ the return at time t on the portfolio with weight vector w . To simplify notation, we omit the time index whenever no confusion is possible and assume that the density function of r_w is continuous. At a preset probability level denoted α , which is typically set between 1 and 5 percent, the portfolio VaR is the negative value of the α -quantile of the portfolio returns. The portfolio CVaR is the negative value of the expected portfolio return when that return is less than its α -quantile:

$$CVaR_w(\alpha) = -E[r_w \mid r_w \leq -VaR_w(\alpha)], \quad (2)$$

with E the expectation operator. The CVaR contribution is the weight of the position in the portfolio times the partial derivative of the portfolio CVaR with respect to that weight:

$$C_{(i)} CVaR_w = w_{(i)} \frac{\partial CVaR_w}{\partial w_{(i)}}, \quad (3)$$

where $w_{(i)}$ is the portfolio weight of position i and there are N assets in the investment universe ($i = 1, \dots, N$).

For ease of interpretation, the CVaR contributions are standardized by the total CVaR. This yields the percentage CVaR contributions:

$$\%C_{(i)} CVaR_w = \frac{C_{(i)} CVaR_w}{CVaR_w}. \quad (4)$$

An interesting summary statistic of the portfolio's CVaR allocation is what we call the portfolio CVaR Concentration, defined as the largest Component CVaR of all positions:

$$C_w(\alpha) = \max_i C_{(i)} CVaR_w(\alpha). \quad (5)$$

As we will show later, minimizing the portfolio CVaR concentration leads to portfolios with a relatively low CVaR and a balanced CVaR allocation.

ESTIMATION

The actual risk contributions can be estimated in two ways. A first approach is to estimate the risk contributions by replacing the expectation in (2) with the sample counterpart evaluated at historical or simulated data. In a portfolio optimization setting the risk contributions needs to be evaluated for a large number of possible weights and therefore fast and explicit estimators are needed. A more elegant approach for optimization problems is therefore to derive the analytical formulae of the risk contributions. If the returns at time t are conditionally normally distributed with mean μ_t and covariance matrix Σ_t , then CVaR at time t is given by:

$$CVaR_w(\alpha) = -w' \mu_t + \sqrt{w' \Sigma_t w} \frac{\phi(z_\alpha)}{\alpha}, \quad (6)$$

with z_α the α -quantile of the standard normal distribution and ϕ the standard normal density function. The contribution to CVaR is then:

$$C_{(i)} CVaR_w(\alpha) = w_{(i)} \left[-\mu_{(i)t} + \frac{(\Sigma_t w)_{(i)}}{\sqrt{w' \Sigma_t w}} \frac{\phi(z_\alpha)}{\alpha} \right]. \quad (7)$$

Financial returns are usually non-normally distributed. In the empirical application, we use the modified CVaR (contribution) estimator proposed by Boudt, Peterson and Croux [2008]. Based on Cornish-Fisher expansions, the modified CVaR estimate is an explicit function of the comoments of the underlying asset returns. It has been shown to deliver accurate estimates of CVaR (contributions) for portfolios with non-normal returns. To save space, we refer the reader to Boudt, Peterson and Croux [2008] for the exact definition of this estimator and to the Appendix for details on the implementation of this estimation method as used in this paper. Throughout the paper we set the loss probability α to 5%.

CVaR RISK BUDGETS IN PORTFOLIO OPTIMIZATION

Previously, risk budgets based on portfolio standard deviation and value-at-risk have been used either as an ex post or ex ante tool for tuning the portfolio allocation.

In the ex post approach, the portfolio is first optimized without taking the risk allocation into account. Next the risk budget of the optimal portfolio is estimated and risk budget violations are adjusted on a marginal basis. The rationale for this is that the risk contributions in (1) can be interpreted as the marginal risk impact of the corresponding position. Because of transaction costs, traders and portfolio managers often update their portfolios incrementally, which makes the marginal interpretation of risk contribution useful in practice (Litterman [1996], Stoyanov, Rachev and Fabozzi [2009]). When the risk contribution of a position is zero, Litterman [1996] calls this the “best hedge” position for that portfolio component. The positions with the largest risk contributions are called “hot spots”. If a risk contribution is negative, a small increase in the corresponding portfolio weight leads to a decrease in the portfolio risk. Keel and Ardia [2010] show however that reallocation of the portfolio based on these risk contributions is limited in two ways. First of all, as a sensitivity measure, they are only precise for infinitesimal changes, but for realistic reallocations, these approximations can be poor.

Second, they assume changing a single position keeping fixed all other positions. In the presence of a full investment constraint, this is unrealistic.

While various volatility-weighted portfolio allocation methods have existed for many years, the ex ante use of risk budgets in portfolio allocation is more recent. Qian's [2005] "Risk Parity Portfolio" allocates portfolio variance equally across the portfolio components. Maillard, Roncalli and Teileche [2010] call this the "Equally-Weighted Risk Contribution Portfolio" or, simply, the Equal-Risk Contribution (ERC) portfolio. They derive the theoretical properties of the ERC portfolio and show that its volatility is located between those of the minimum variance and equal-weight portfolio. Zhu, Li and Sun [2010] study optimal mean-variance portfolio selection under a direct constraint on the contributions to portfolio variance.

Our first contribution to this recent literature is to use CVaR contributions rather than variance contributions as an objective or constraint in portfolio optimization. By integrating the CVaR budget into their optimal portfolio policy, investors can directly optimize downside risk diversification. The rationale for this is founded on Scaillet [2002] establishing a direct link between CVaR contributions and downside risk concentration. More precisely, let $r_{(i)}$ be the return on position i . Scaillet [2002] shows that the contributions to CVaR correspond to the conditional expectation of the return of the portfolio component when the portfolio loss is larger than its VaR loss:

$$C_{(i)}CVaR_w(\alpha) = -E[w_{(i)}r_{(i)} \mid r_w \leq -VaR_w(\alpha)]. \quad (8)$$

From (1) and (8) it follows that the percentage CVaR contribution can be rewritten as the ratio between the expected return on the position at the time the portfolio experiences a beyond VaR loss and the expected value of these beyond VaR portfolio losses:

$$\%C_{(i)}CVaR_w(\alpha) = \frac{E[w_{(i)}r_{(i)} \mid r_w \leq -VaR_w(\alpha)]}{E[r_w \mid r_w \leq -VaR_w(\alpha)]}. \quad (9)$$

In almost all practical cases, the denominator in (7) is negative such that a high positive percentage CVaR contribution indicates that the position has a large loss when the portfolio also has a large loss. The higher the percentage CVaR, the more the portfolio downside risk is concentrated on that asset and vice versa.

Our second contribution is that we propose two strategies for using the CVaR budgets in portfolio optimization in order to balance the maximum return, minimum downside risk, and maximum downside risk diversification objectives of an investor. The first strategy is the Minimum CVaR Concentration portfolio

(MCC), which uses the downside risk diversification criterion as an objective rather than a constraint. More formally, the MCC portfolio allocation is given by:

$$w^{MCC} = \arg \min_w C_w(\alpha), \quad (10)$$

with the portfolio's CVaR concentration $C_w(\alpha)$ as defined in (5).

The second strategy consists of imposing bound constraints on the percentage CVaR contributions. This may be viewed as a direct substitute for a risk diversification approach based on position limits. It has the ERC constraint as a special case:

$$\% C_{(1)} CVaR_w(\alpha) = \dots = \% C_{(N)} CVaR_w(\alpha) = 1/N. \quad (11)$$

Note that for a portfolio that has the ERC property, the relative weights are inversely proportional to the marginal impact of the position on the portfolio CVaR:

$$\frac{w_{(i)}}{w_{(j)}} = \frac{\partial CVaR_w(\alpha) / \partial w_{(j)}}{\partial CVaR_w(\alpha) / \partial w_{(i)}}. \quad (12)$$

It follows that the ERC allocation strategy yields portfolios that give higher weights to assets with a small marginal risk impact and down-weights the investments with a high marginal risk (the so-called “hot spots” in Litterman [1996]).

The next paragraphs study the properties of these two approaches in more detail. In the empirical section, we will compare these CVaR budget based portfolio allocation rules with the more standard Minimum CVaR (MC) and Equal-Weight (EW) portfolios:

$$w^{MC} = \arg \min_w CVaR_w(\alpha) \text{ and } w^{EW} = (1/N, \dots, 1/N)'. \quad (13)$$

PROPERTIES OF THE MINIMUM CVaR CONCENTRATION PORTFOLIO

For the derivation of the properties of the minimum CVaR concentration portfolio, it is useful to rewrite the portfolio CVaR concentration as the portfolio CVaR times the largest percentage CVaR contribution:

$$C_w(\alpha) = CVaR_w(\alpha) \max\{\% C_{(1)} CVaR_w(\alpha), \dots, \% C_{(N)} CVaR_w(\alpha)\}. \quad (14)$$

The first factor in (14) is minimized by the minimum CVaR portfolio. The second factor attains its lowest value when the portfolio has the ERC property, since $\max\{\% C_{(1)} CVaR_w(\alpha), \dots, \% C_{(N)} CVaR_w(\alpha)\} \geq 1/N$.

By minimizing the product of these two factors, the MCC portfolio strikes a balance between the objectives of

portfolio risk diversification and total risk minimization. The first order conditions of the MCC portfolio give some further interesting insights:

$$\begin{aligned} \frac{\partial C_w(\alpha)}{\partial w} = CVaR_w(\alpha) \frac{\partial \max\{\% C_{(1)} CVaR_w(\alpha), \dots, \% C_{(N)} CVaR_w(\alpha)\}}{\partial w} \\ + \max\{\% C_{(1)} CVaR_w(\alpha), \dots, \% C_{(N)} CVaR_w(\alpha)\} \frac{\partial CVaR_w(\alpha)}{\partial w} = 0. \end{aligned} \quad (15)$$

We see that a necessary condition for the MCC portfolio to have the exact ERC property is also that the derivative of the portfolio CVaR is zero. Since CVaR is a convex function, this only happens if the (unconstrained) minimum CVaR portfolio has the ERC characteristic, in which case it coincides with the MCC portfolio. Compared with the unconstrained minimum CVaR portfolio, we have that the CVaR of the MCC portfolio is higher, but the risk is less concentrated. In fact, we show in Appendix that the percentage CVaR of the minimum CVaR portfolio coincides with the component's portfolio weight:

$$\% C_{(i)} CVaR_{w^{MC}} = w_{(i)}^{MC}. \quad (16)$$

It is well known that the minimum CVaR portfolio generally suffers from the drawback of portfolio concentration. By (16) this carries directly over to the CVaR allocation.

Note also in (15) that there can be multiple portfolio weights for which the first order conditions are satisfied. For this reason, a global optimizer is needed to find the MCC portfolio. We used the differential evolution algorithm developed by Price, Storn and Lampinen [2008].

The MCC objective can easily be combined with a return target. This serves the general purpose of maximizing return subject to some level of risk, while also minimizing risk concentration at that risk level. We define the mean-CVaR concentration efficient frontier as the collection of all portfolios that achieve the lowest degree of CVaR concentration for a return objective. For a given return target \bar{r} , the mean-CVaR concentration efficient portfolio solves:

$$\min_w C_w(\alpha) \quad \text{s.t.} \quad w' \mu \geq \bar{r}. \quad (17)$$

The minimum CVaR portfolio under the ERC constraint in (11) is an alternative to the MCC portfolio for attaining a balance between the objectives of portfolio risk diversification and total risk minimization. On our data examples, the two portfolios were always very similar.

The advantage of the MCC portfolio over the ERC constrained minimum CVaR portfolio is that it is computationally simpler and also will yield a solution if the ERC constraint is not feasible or conflicts with

other constraints. Since most real-world portfolios are constructed with an explicit or implicit return objective and other constraints, being able to combine with other objectives and constraints is an important consideration for asset managers that is often incompatible with the published literature on utilizing risk metrics in portfolio construction. Note also that the properties of the MCC portfolio generalize to any minimum Risk Concentration portfolio, as long as the portfolio risk measure is a one-homogeneous function of the portfolio weights.

PORTFOLIO ALLOCATION USING PERCENTAGE CVaR CONSTRAINTS

The risk allocation can also be controlled by imposing explicit constraints on the percentage CVaR allocations. This process operates in much the same way that portfolio managers impose weight constraints on portfolios.

Such percentage CVaR contribution constraints reduce the feasible space in a way that depends on the return characteristics. Stoyanov, Rachev and Fabozzi [2009] study in detail the effect of the component return characteristics on the total portfolio CVaR. To build further intuition via a stylized example we plot in Exhibit 1 the percentage CVaR contributions for a two-asset portfolio with asset returns that have a bivariate normal distribution with means μ_1 and μ_2 , standard deviations σ_1 and σ_2 and a correlation of 0.5. Of course, the percentage CVaR contribution is zero and one if the weight is zero and one, respectively. In between these values, the percentage CVaR displays an S-shape. The dotted lines in this figure illustrate the effect on the feasible space for portfolio weight 1 of imposing an upper 60% bound on the percentage CVaR contributions of the two assets. This implies that the percentage CVaR contribution of asset 1 has to be between 40% and 60%. In the top figure, the two assets are identical. In this case, the feasible space is centered around the equal-weight portfolio. In the middle and bottom figure, asset 1 is more attractive than asset 2 since it has either a lower volatility or a higher expected return. We see that this leads to a shift of the feasible space to the right, with allowed portfolio weights around 60%. The set of possible weights satisfying the box constraints on the percentage CVaR contributions changes in an intuitively appealing way when differences in return and volatility are allowed.

[Insert Exhibit 1 about here]

For general portfolios with non-normal returns, there is no explicit representation of the percentage CVaR constraint as weight constraint available for investment. A general purpose portfolio solver that can handle such

percentage CVaR contribution constraints is available in the R package PortfolioAnalytics of Boudt, Carl and Peterson [2010].

EMPIRICAL RESULTS

In this section we apply the CVaR decomposition methodology to optimize portfolios that allocate across asset classes. The analysis is based on the January 1976 – June 2010 monthly total USD returns of the Merrill Lynch Domestic Master index (bonds), the S&P 500 index (US stocks), the MSCI EAFE index (Europe, Asia and Far East stocks) and the S&P Goldman Sachs commodity index. The data are obtained from Datastream. We will start with a static two asset bond-equity portfolio, and expand to a larger portfolio for studying the effects of rebalancing under various constraints and objectives. We impose in all portfolio allocations the full investment constraint and exclude short sales.

STATIC BOND-EQUITY PORTFOLIO

The simple bond-equity portfolio application in Exhibit 2 illustrates the impact of the portfolio policy on the risk allocation. Portfolio managers frequently rely on the heuristic approach of applying position limits to ensure diversification. Such a simple approach may ignore the individual risks of the portfolio assets and their risk dependence. A first example is the equal-weight portfolio, which is popular in practice because it does not require any information on the risk and return and supposedly provides a diversified portfolio. A second popular position constrained bond-equity portfolio is the 60/40 portfolio, investing 60% in bonds and 40% in equity. The first two lines in Exhibit 2 show the estimated risk allocation of these portfolios.¹ We see that position limits clearly fail to produce portfolios with an ex ante risk diversification: respectively 97% and 86% of the portfolio CVaR is caused by the equity investment in the equal-weight and 60/40 portfolios.

[Insert Exhibit 2 about here]

Note also in Exhibit 2 that the equal-weight and 60/40 portfolios have a relatively high level of total portfolio CVaR. Rockafellar and Uryasev [2002], among others, recommend the minimum CVaR portfolio to investors wanting to avoid extreme losses. For our sample, the minimum CVaR portfolio has a monthly 95% CVaR of 2.44%, which is less than half the CVaR of the equal-weight and 60/40 portfolios. However, the portfolio risk is

¹ All moments are estimated by their historical sample counterpart. Over the January 1976 – June 2010 period, the annualized average monthly return of the bond and US equity is 7.55% and 10.25%, respectively. This difference in average return is traded off with the higher CVaR of the equity (9.97%) compared to the one of the bond (2.46%).

still heavily concentrated in one asset: the bond allocation is responsible for 97% of portfolio CVaR in the minimum CVaR portfolio.

This paper proposes the Minimum CVaR Concentration (MCC) portfolio for investors interested in having both a high ex ante downside risk diversification and a low total portfolio CVaR. We see in Exhibit 2 that for this sample the MCC portfolio has the highest CVaR diversification possible: it is an equal risk contribution portfolio with a 23% part in equity. It has only a slightly higher CVaR than the minimum CVaR portfolio, but also a higher average return.

Finally, we also consider substituting the 60/40 weight allocation with a 60/40 risk allocation. 81% of this percentage risk constrained portfolio is invested in bonds. Like for the MCC portfolio, the price for risk diversification is a slight increase in the portfolio CVaR compared to the minimum CVaR portfolio, but this is also compensated by a higher average return.

MEAN-CVaR CONCENTRATION EFFICIENT FRONTIER

In comparison with the ERC portfolio of Qian [2005], the MCC portfolio has the advantage that it may be easily combined with many other investor objectives and constraints (such as return targets or drawdown constraints). In Exhibits 3 and 4 we implement the mean-CVaR concentration efficient frontier in (16) by adding a return target to the minimum CVaR concentration objective. The investment universe is expanded by including the GSCI and EAFE index.² For benchmarking, we also consider the minimum standard deviation (StdDev) and minimum CVaR portfolios. The mean-CVaR approach is potentially more appealing than the standard mean-StdDev approach because it trades off return with risk of extreme losses rather than volatility. In our application, the estimated CVaR accounts for the non-normality of the financial return series (see Appendix for more details).

[Insert Exhibits 3 and 4 about here]

The upper panel of Exhibit 3 plots the mean-CVaR frontiers, while the lower panel shows the annualized mean return of the portfolios against the largest percentage CVaR contribution. A joint reading of these plots is needed to understand the trade-off between the maximum return, minimum risk, and minimum risk concentration objectives. Exhibit 4 presents the corresponding weight and CVaR allocation plots.

² The GSCI index has a relatively low annualized monthly return (5.41%) and high risk (monthly CVaR of 12.78%). With an annualized return of 8.68% and monthly CVaR of 10.87%, the EAFE index offers a higher return than the bond, but a lower return and higher risk than the S&P 500.

Consider first the portfolios without return constraint. The minimum StdDev and CVaR portfolio offer an annualised return of 7.5% for a monthly 95% CVaR of 2.5% and 2.4%, respectively. The minimum CVaR concentration portfolio is more risky: it has a CVaR of 3.4%, but it has a higher return (7.7%) and its investments are more diversified across all four assets: the weight of the US bond, S&P 500, EAFE and GSCI indices are 63.3%, 13.4%, 10.4% and 11.9%, respectively. In contrast, the minimum CVaR (and StdDev) portfolios are concentrated in the bond, which has a 94.7% (86.2%) weight in these portfolios.

Exhibit 4 shows that imposing a return constraint on the minimum StdDev and CVaR portfolios leads to a higher allocation to the S&P 500 and a reduction in the bond investment. Of course this leads to portfolios with a higher return and risk, but interestingly as long as the return target is below 8.2% it reduces the risk concentration of the portfolio as can be seen from the lower figure in Exhibit 3. From that point onwards, the S&P 500 index becomes the largest risk contributor and higher returns are traded off with both a higher total portfolio risk and risk concentration. In this application, the difference between the mean-StdDev and mean-CVaR frontier is relatively small. For portfolios with a return above the 8.2% threshold, the mean-StdDev and mean-CVaR efficient portfolios are similar. For lower returns, the mean-StdDev portfolios tend to invest more in the commodities index, because of its negative correlation (-6%) with the bond returns. A bigger difference in variation of the mean-StdDev and mean-CVaR frontiers can be expected for larger portfolios with more non-normality in the underlying asset return distributions.

The mean-CVaR concentration efficient portfolio is very different from the mean-CVaR and mean-StdDev efficient portfolios. On this data set, the mean-CVaR concentration efficient frontier has three distinct segments. Unconstrained, the mean-CVaR concentration efficient frontier is an equal risk contribution portfolio with an annualized return of 7.77%. For a target return between 7.77% and 8.13%, the portfolio CVaR concentration increases from 0.87% to 1.12%, but the portfolio CVaR decreases from 3.47% to 3.35%. This is due to a reallocation from the more risky commodity investment into bonds, as can be seen in Exhibit 4. At the end of this segment, the portfolio is no longer invested in commodities. Bonds dominate the portfolio budget allocation with a 71% share. On the second segment, the bond allocation shrinks to zero, while the shares of the S&P 500 and the EAFE index rise from 16% to 51% and from 13% to 49%, respectively. On this angle portfolio, the S&P 500 and EAFE index contribute each for 50% to the portfolio CVaR, which is now 9.8% compensated by a target return of 9.48%. The portfolios on the final segment of the frontier replace gradually the EAFE investment with the S&P 500. Since this asset offers the highest return, it is also the endpoint of the long-only constrained mean-CVaR concentration efficient frontier.

DYNAMIC INVESTMENT STRATEGIES

Let us now consider a dynamic portfolio invested in bonds, US equity, Europe, Asia and Far East equity, and commodities. The portfolio is rebalanced quarterly to satisfy either an equal-weight, minimum CVaR or minimum CVaR concentration (MCC) objective. The risk budgets that are optimized are all conditional on the information available at the time of rebalancing. We give more details on the estimation in the Appendix. Since part of the estimation is based on rolling samples of eight years and the data span is January 1976 – June 2010, the optimized weights are only available for the quarters 1984Q1 – 2010Q3. In all aspects, the MCC portfolio and ERC constrained minimum CVaR portfolios are very similar. We therefore discuss in the text only the results for the MCC portfolio, but for completeness the exhibits show the results for both portfolios.

We discuss first the results for the equal-weight, minimum CVaR and MCC portfolios. We then analyze the sensitivity of the minimum CVaR and MCC portfolios to the inclusion of a weight or risk allocation constraint.

RESULTS UNCONSTRAINED PORTFOLIOS

The left and right panels of Exhibit 5 plot the weight and CVaR allocations of the equal-weight, minimum CVaR, and MCC portfolios. We find that for almost all periods the minimum CVaR portfolio is highly invested in the bond index, while the MCC portfolio is more balanced across all asset classes. As predicted by theory, the risk allocation of the minimum CVaR portfolio coincides with its weight allocation and the risk allocation of the MCC portfolio is close to the equal risk contribution state. The CVaR of the equal-weight portfolio is dominated by the S&P 500 and EAFE stocks. The diversification potential of the bond is not fully exploited by the equal-weight portfolio, since for many quarters it has a negative risk contribution. This indicates that increasing the weight of the bond would marginally decrease portfolio risk. The reason for the bad performance of weight constraints in ensuring ex ante risk diversification is the non-linear dependence of portfolio CVaR contributions on the weights. Reaching the portfolio manager's goal of ensuring risk diversification is therefore more efficiently achieved via direct constraints on the risk budget contributions rather than on the weights.

Exhibit 6 plots the ex ante portfolio risk estimates. As expected, the CVaR of the MCC portfolio is for all quarters in between the CVaR of the minimum CVaR portfolio and the CVaR of the equal-weight portfolio.

[Insert Exhibits 5 and 6 about here]

The solid black lines in the lower and upper panels of Exhibit 7 plot the ratio of the monthly cumulative out-of-sample returns of the minimum CVaR and MCC portfolios versus the cumulative returns of the equal-

weight portfolio over the period January 1984-June 2010. The value of the chart is less important than the slope of the line. If the slope is positive, the strategy in the numerator is outperforming the equal-weight strategy, and vice versa. The vertical grey bars denote bear markets defined by Ellis [2005] as periods with a decline in the S&P 500 index of 12 per cent or more. The left side of the bar corresponds to the market peaks and the right side to the stock market trough.³ We see in Exhibit 6 that the minimum CVaR portfolio, having a large allocation to the bond, outperforms the equal-weight and MCC portfolios at times of serious stock market downturn. The performance of the MCC portfolio seems to be a middle ground between the performance of the equal-weight and minimum CVaR portfolios. It offers an attractive compromise between the good performance of the minimum CVaR portfolio in adverse markets and the upward potential of the equal-weight portfolio. A final observation is that periods where one strategy is outperforming the other are relatively long and indicate the possibility of applying market timing strategies on top of these allocations.

Exhibit 8 reports the annualized out-of-sample average return on the portfolios. When computed over the whole period, the minimum CVaR and MCC portfolios performed within 16 bps of one another. This is a small margin, given the long period of time presented. The risk statistics computed from the out-of-sample returns confirm the ex ante risk estimates from Exhibit 6. The value of the annualized standard deviation and monthly historical CVaR of the MCC portfolio is in between those of the minimum CVaR portfolio and the equal-weight portfolio.⁴ The equal-weight portfolio has extremely large drawdowns. Over the sample it has four drawdowns higher than 10%, while the minimum CVaR and MCC only have one. In the credit crisis the equal-weight portfolio suffered a drawdown of 48%, which is triple the drawdown of the MCC portfolio.

Splitting the sample into bull/bear periods, we see a much bigger variation in relative performance. The return for the minimum CVaR portfolio trailed the MCC portfolio by more than 200 bps during equity bull markets, yet outperformed during bear markets by more than 1000 bps. The minimum CVaR and MCC portfolio have thus each their appeal depending on the market environment. This might lead to risk timing the portfolio allocation, whereby the investor selects his risk appetite based on broad market conditions. In a secular bull market, he might choose the MCC portfolio because of its relative outperformance in exchange for the risk of slightly larger losses. In a secular bear market, the minimum CVaR portfolio might be more appealing because of its conservatism.

³ For our sample, the bear market periods are September-November 1987, June-October 1990, July-August 1998 and November 2007-February 2009.

⁴ The historical CVaR is the average out-of-sample portfolio return when the return is below its 5% empirical quantile.

[Insert Exhibits 7 and 8 about here]

The last panel of Exhibit 8 summarizes the out-of-sample trade-off between minimum downside risk and maximum downside risk diversification of the portfolios. It reports for each strategy the median and maximum of all losses exceeding 10% as well as the median and maximum value of the largest component percentage contribution to those losses. Recall from (14) that the MCC portfolio is designed to have both a low downside risk and high downside risk diversification. This is confirmed by the data. Compared to the minimum CVaR portfolio, the value of the extreme losses on the MCC portfolio are similar, but in the MCC portfolio the contribution to these losses are less concentrated. The equal-weight portfolio is most effective in diversifying its downside risk exposure, but this comes at the price of having also the highest level of downside risk. Its median and maximum loss exceeding 10% is 16% and 22%, respectively, while for the MCC portfolio, these are only 12% and 14%.

Finally, we consider the portfolio turnover of the strategies, defined by DeMiguel, Garlappi and Uppal [2009] as the average sum of the absolute value of the trades across the N available assets:

$$\text{Turnover} = \frac{1}{NT_*} \sum_{t=1}^{T_*-1} \sum_{i=1}^N |w_{(i)t+1} - w_{(i)t}|, \quad (18)$$

where $w_{(i)t+1}$ is the weight of asset i at the start of rebalancing period $t+1$, $w_{(i)t}$ is the weight of that asset before rebalancing at t and T_* is the total number of rebalancing periods. This turnover quantity can be interpreted as the average percentage of wealth traded in each period. The portfolio turnover is the lowest for the equal-weight portfolio (1.26%). The MCC portfolio has a significantly lower turnover (1.74%) than the minimum CVaR portfolio (2.14%).

In conclusion, the minimum CVaR portfolio has the lowest out-of-sample risk but a high risk concentration and turnover. The equal-weight strategy has the lowest turnover and risk concentration, but highest total risk. The proposed MCC portfolio is on all these dimensions the second best. It achieves an attractive compromise between low overall risk, good upside return, high diversification, and low turnover.

SENSITIVITY TO WEIGHT AND RISK ALLOCATION CONSTRAINTS

Portfolio managers might wish to impose their diversification objective through a position limit or risk allocation constraint on the minimum CVaR or MCC portfolios. We investigate in Exhibits 6-9 the sensitivity of

the portfolios to an upper 40% position limit or an upper 40% CVaR allocation limit. The choice of 40% is arbitrary, but it is consistent with the 40% allocation to equity in the stylized 60/40 bond-equity portfolio.

The upper two plots in Exhibit 9 present the weight allocations of the constrained minimum CVaR portfolios. We see that the 40% upper bound on the portfolio weights and risk allocations is stringent for almost all periods. Under these constraints, the component CVaR contribution of the minimum CVaR portfolio no longer coincides with the weight allocation. The investment in the bond typically contributes less to CVaR risk than its portfolio weight. Its contribution is for some months even negative under the position limit. The bottom figures in Exhibit 9 show the weight and risk allocation of the MCC portfolio under a 40% upper bound on the portfolio weights. We see that in spite of the weight constraint, the risk of the MCC portfolio is still more equally spread out than for the minimum CVaR portfolio where for some periods the S&P 500 and EAFE investments cause more than half of portfolio risk.

[Insert Exhibit 9 about here]

From the weight and CVaR allocation plots, it is clear that adding position or risk allocation limits pushes the minimum CVaR and MCC portfolio towards an allocation that is closer to the equal-weight portfolio. Consequently, the return, risk, and turnover properties of these constrained portfolios are closer to the equal-weight portfolio, as can be seen in Exhibits 6-9. Note also that the effect on returns of the risk contribution constraint is smaller than for the corresponding weight constraint.

CONCLUSION

An extensive empirical application of ex ante risk budget methods to dynamic allocation across bonds, commodities, domestic and international equity illustrated the out of sample effectiveness of risk budgets in generating portfolios that have low portfolio risk and risk concentration, high diversification, and low portfolio turnover. A first strategy is to impose bound constraints on the percentage CVaR contributions. This provides a direct substitute and improvement to the commonly practiced risk diversification approach based on position limits. A second strategy consists of minimizing the largest component CVaR contribution, which directly addresses risk diversification, even in portfolios with non-normally distributed assets. The properties of these approaches as described in this paper compare favorably relative to the equal-weight and minimum risk portfolios. Unconstrained, the Minimum CVaR Concentration (MCC) portfolio is typically very similar to the

equal-risk-contribution portfolio of Qian [2005]. Furthermore, it may be easily combined with many other investor objectives and constraints (such as return targets or drawdown constraints). Investors can thus optimally balance their maximum return, minimum downside risk, and maximum downside risk diversification objectives through an ex ante use of conditional value-at-risk (CVaR) budgets in portfolio optimization.

ENDNOTES

The authors thank David Ardia, Bernhard Pfaff, and Dale Rosenthal for helpful comments and the National Bank of Belgium for financial support. The code to replicate the analysis is available in the R packages PerformanceAnalytics of Carl and Peterson [2010] and PortfolioAnalytics of Boudt, Carl and Peterson [2010]. More information on the utilization of these packages for the estimation and optimization of portfolio CVaR budgets can be found at www.econ.kuleuven.be/kris.boudt/public/riskbudgets.htm.

REFERENCES

- Artzner, Philippe, Delbaen, Freddy, Eber, Jean-Marc and David Heath. "Coherent measures of risk." *Mathematical Finance*, Vol. 6 (1998), pp. 203-228.
- Boudt, Kris, Peterson, Brian G. and Christophe Croux. "Estimation and decomposition of downside risk for portfolios with non-normal returns." *The Journal of Risk*, Winter 2008, pp. 79-103.
- Boudt, Kris, Carl, Peter and Peterson, Brian G. PortfolioAnalytics. R package version 0.0.6, 2010.
- Carl, Peter and Peterson, Brian G. PerformanceAnalytics: Econometric tools for performance and risk analysis. R package version 1.0.2.1, 2010.
- Chow, George, and Mark Kritzman. "Risk budgets." *The Journal of Portfolio Management*, Winter 2001, pp. 56-60.
- DeMiguel, Victor, Garlappi, Lorenzo and Raman Uppal. "Optimal versus naïve diversification: How inefficient in the 1/N portfolio strategy?" *Review of Financial Studies*, Vol. 22 (2009), pp. 1915-1953.
- Denault, Michel. "Coherent allocation of risk capital." *The Journal of Risk*, Fall 2001, pp. 1-33.
- Ellis, J.H. Ahead of the curve: a commonsense guide to forecasting business and market cycles. Harvard Business Press, 2005.
- Keel, Simon and David Ardia. "Generalized marginal risk." AERIS Capital working paper, 2010.
- Litterman, Robert B. "Hot SpotsTM and hedges." *The Journal of Portfolio Management*, 1996 Special Issue, pp. 52-75.

Maillard, Sébastien, Roncalli, Thierry and Jérôme Teiletche. “On the properties of equally-weighted risk contributions portfolios.” *The Journal of Portfolio Management*, Summer 2010, pp. 60-70.

Martellini, Lionel and Volker Ziemann. “Improved estimated of higher-order comoments and implications for portfolio selection.” *Review of Financial Studies*, Vol. 23 (2010), pp. 1467-1502.

Pearson, N.D. Risk Budgeting. New York: John Wiley and Sons, 2002.

Peterson, Brian, and Kris Boudt. “Component VaR for a non-normal world.” *RISK*, November 2008, pp. 78-81.

Pflug, G. Ch. “Some remarks on the value-at-risk and the conditional value-at-risk.” In S. Uryasev, ed., *Probabilistic Constrained Optimization: Methodology and Applications*, Dordrecht: Kluwer, 2000, pp. 272-281.

Price, K.V., Storn, R.M. and Lampinen, J.A. *Differential Evolution - A Practical Approach to Global Optimization*. Springer-Verlag, 2006.

Qian, Edward E. “Risk parity portfolios: Efficient portfolios through true diversification of risk.” *Panagora Asset Management*, September 2005.

Rockafellar, Ralph T. and Stanislav Uryasev. “Optimization of Conditional Value-At-Risk.” *The Journal of Risk*, Spring 2000, pp. 21-41.

Scaillet, Olivier. “Nonparametric estimation and sensitivity analysis of expected shortfall.” *Mathematical Finance*, Vol. 14 (2002), pp. 74-86.

Scherer, B. *Portfolio Construction & Risk Budgeting*. London: Risk Books, 2002.

Stoyanov, Stoyan, Rachev, Svetlozar T. and Frank J. Fabozzi. “Sensitivity of portfolio VaR and CVaR to portfolio return characteristics.” *Universität Karlsruhe working paper*, 2009.

Zhu, Shushang, Li, Duan and Xiaoling Sun. “Portfolio selection with marginal risk control.” *The Journal of Computational Finance*, Fall 2010, pp. 1-26.

APPENDIX

CVaR BUDGET OF MINIMUM CVaR PORTFOLIO

The first order conditions of the minimum CVaR portfolio are that

$$\frac{\partial CVaR_w}{\partial w_{(1)}} = \dots = \frac{\partial CVaR_w}{\partial w_{(N)}} = 0. \quad (19)$$

By l'Hôpital's rule we have that

$$\%C_{(i)}CVaR_w = \frac{w_{(i)} \frac{\partial CVaR_w}{\partial w_{(i)}}}{\sum_{i=1}^N w_{(i)} \frac{\partial CVaR_w}{\partial w_{(i)}}} \rightarrow \frac{w_{(i)}}{\sum_{i=1}^N w_{(i)}} \text{ when } \max\left\{\frac{\partial CVaR_w}{\partial w_{(1)}}, \dots, \frac{\partial CVaR_w}{\partial w_{(N)}}\right\} \rightarrow 0. \quad (20)$$

Hence the weight and CVaR allocation coincide for the minimum CVaR portfolio.

DETAILS ON ESTIMATION METHOD

Because of the non-normality in the data, we use the modified CVaR estimator of Boudt, Peterson and Croux [2008]. Its implementation requires an estimate of the first four moments of the portfolio returns. For the static portfolio, the moment estimates are their samples counterparts. For the dynamic portfolio allocation, time-varying conditional moment estimates are obtained as follows. We first center the returns around an exponentially weighted average of the returns over the past eight years. The centered returns are modeled as a GARCH(1,1) process whose parameters are estimated by Gaussian quasi-maximum likelihood using all data available from inception up to the time for which the CVaR estimate is needed. We then compute the innovations as the centered returns divided by their volatility estimate. The correlation, coskewness and cokurtosis matrices of these innovations are then estimated as the rolling eight year sample correlation, coskewness and cokurtosis matrix of a winsorized version of these innovations. The winsorization ensures the outlier-robustness of the estimates and is described in Boudt, Peterson and Croux [2008]. In addition or as an alternative to the robust parameter estimation we employ here, the modified CVaR estimator may also use other improved estimators of higher-order comoments such as those put forward by Martellini and Ziemann [2010], providing additional options for asset managers to fine-tune the methodology for their portfolio.

EXHIBITS

Exhibit 1: Percentage CVaR contribution of asset 1 in function of its portfolio weight for a two-asset portfolio with asset returns that have a bivariate normal distribution with means μ_1 and μ_2 , correlation 0.5 and standard deviations σ_1 and σ_2 , respectively.




Exhibit 2: Weight and CVaR allocation of bond-equity portfolios, together with the in-sample annualized mean and monthly 95% CVaR over the period January 1976-June 2010.

	Weight allocation		CVaR allocation		Ann. mean	95% CVaR
	Bond	Equity	Bond	Equity		
Equal-weight	50%	50%	3.47%	96.53%	8.90%	4.87%
60/40 weight	60%	40%	13.79%	86.21%	8.63%	4.03%
Min CVaR	96.86%	3.14%	96.86%	3.14%	7.63%	2.44%
Min CVaR concentration	77.01%	22.91%	50%	50%	8.17%	3.00%
60/40 risk allocation	81.23%	18.77%	60%	40%	8.05%	2.80%

Exhibit 3: Annualized mean returns versus the monthly portfolio CVaR and largest percentage CVaR contribution for the mean-StdDev, mean-CVaR and mean-CVaR concentration efficient portfolios. The frontier is estimated using all January 1976-June 2010 monthly returns.




Exhibit 4: Weight and CVaR allocation of mean-StdDev, mean-CVaR and mean-CVaR concentration efficient portfolios for various levels of annualized portfolio returns. The frontier is estimated using all January 1976-June 2010 monthly returns.




Exhibit 5: Stacked bar weight and CVaR contribution plots for the equal-weight, minimum CVaR and minimum CVaR concentration portfolios invested in the Merrill Lynch US bond, S&P 500, MSCI EAFE and S&P GSCI indices. The portfolios are rebalanced quarterly.

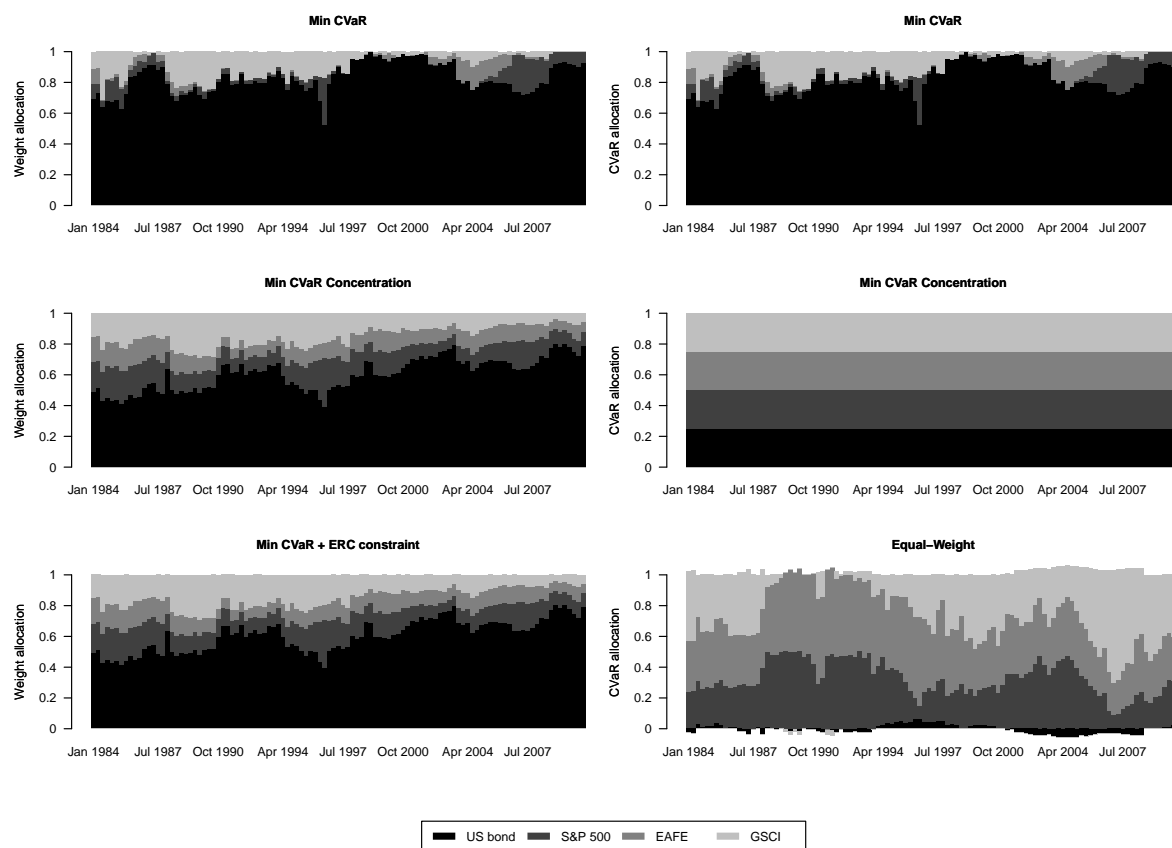


Exhibit 6: Monthly CVaR of the equal-weight and risk budget optimized portfolios invested in the Merrill Lynch US bond, S&P 500, MSCI EAFE and S&P GSCI indices. The portfolios are rebalanced quarterly. The shaded regions indicate a bear market regime.




Exhibit 7: Relative performance of the risk budget optimized portfolios versus the equal-weight portfolio invested in the Merrill Lynch US bond, S&P 500, MSCI EAFE and S&P GSCI indices. The portfolios are rebalanced quarterly. The shaded regions indicate a bear market regime.




Exhibit 8: Summary statistics of monthly out-of-sample returns on investment strategies over the period January 1984 - June 2010.

		Equal Weight	Min CVaR			Min CVaR Concentration		
			40% Position		40% CVaR ERC	40% Position		
			Limit	Alloc Limit			Limit	
<i>Full period (in %)</i>								
Ann. Mean		7.32	8.07	7.72	7.99	8.22	8.23	7.63
Ann. StdDev		10.42	4.53	8.77	5.35	5.78	5.78	8.46
Monthly Hist CVaR		7.42	2.34	6.15	2.95	3.35	3.35	5.91
Portfolio turnover		1.26	2.14	3.55	2.64	1.74	1.74	1.51
<i>Bear stock market (in %)</i>								
Ann. Mean		-24.36	6.31	-17.29	-0.61	-3.81	-3.79	-16.51
Ann. StdDev		15.46	6.00	13.03	7.10	7.44	7.44	12.58
Monthly Hist CVaR		13.71	3.30	11.05	5.37	6.28	6.27	10.80
<i>Normal/Bull stock market (in %)</i>								
Ann. Mean		13.37	8.40	12.50	9.63	10.52	10.52	12.24
Ann. StdDev		8.08	4.20	6.92	4.81	5.16	5.16	6.67
Monthly Hist CVaR		4.09	2.04	3.50	2.27	2.38	2.38	3.38
<i>Drawdowns higher than 10%</i>								
Credit crisis*		0.48	0.09	0.37	0.13	0.15	0.15	0.38
Dot-com bubble burst**		0.25		0.19				0.17
Asian-Russian crisis***		0.12						
Black Monday****		0.11		0.12				
<i>Summary statistics on level and concentration of portfolio losses exceeding 10%</i>								
$-w_t' r_t$	median	0.16	0.11	0.13	0.12	0.12	0.12	0.12
	max	0.22	0.17	0.21	0.15	0.14	0.14	0.19
$\max_i \frac{(w_{(i)t} r_{(i)t})}{w_t' r_t}$	median	0.43	0.97	0.51	0.85	0.83	0.83	0.58
	max	0.45	1.02	0.74	0.97	0.92	0.92	0.66

* May-Oct 2008 for the Min CVaR strategy, June 2008-Feb 2009 for all other styles. ** Sep 2000-Sep 2002.

*** April-Aug 1998. **** Sep-Nov 1987.

Exhibit 9: Stacked bar weight and CVaR contribution plots for the constrained minimum CVaR portfolios invested in the Merrill Lynch US bond, S&P500, MSCI EAFE and S&P GSCI indices. The portfolios are rebalanced quarterly.

