Empowering cash managers through compromise programming

Francisco Salas-Molina

Hilaturas Ferre, S.A., Les Molines, 2, 03450 Banyeres de Mariola, Spain,
E-mail:francisco.salas.molina@gmail.com

David Pla-Santamaria

Universitat Politècnica de València, Ferrándiz y Carbonell, 03801 Alcoy, Spain,
E-mail:dplasan@upv.es

Juan A. Rodríguez-Aguilar

IIIA-CSIC, Campus UAB, 08913 Cerdanyola, Spain,
E-mail:jar@iiia.csic.es

July 19, 2017

Abstract

Typically, the cash management literature focuses on optimizing cost, hence neglecting risk analysis. In this chapter, we address the cash management problem from a multiobjective perspective by considering not only cost but also the risk of cash policies. We propose novel measures to incorporate risk analysis as an additional goal in cash management. Next, we rely on compromise programming as a method to minimize the sum of weighted distances to an ideal point where both cost and risk are minimum. These weights reflect the particular preferences of cash managers when selecting the best policies that solve the multiobjective cash management problem. As a result, we suggest three alternative solvers to cover a wide range of possible situations: Monte Carlo methods, linear programming, and quadratic programming. We also provide a Python software library with an implementation of the proposed solvers ready to be embedded in cash management decision support systems. We finally describe a framework to assess the utility of cash management models when considering multiple objectives.

Keywords: Cash management; multiobjective; risk; mathematical programming; Python.

1 Introduction

Cash management is concerned with optimizing costs of short-term cash policies of a company. Different cash management models have been proposed in which the focus has been placed on a single objective, namely, minimizing cost (Gregory, 1976; Srinivasan and Kim, 1986; da Costa Moraes et al., 2015). However, risk analysis must also be incorporated as an additional goal to be minimized in cash management. Hence, cost and risk are usually desired but conflicting objectives. On the one hand, there is a trade-off between holding and transaction costs. On the other hand, cost reductions are achieved by reducing cash balances but, at the same time, the risk of an overdraft increases. In this chapter, we address the multiobjective cash management problem (MOCMP), which can be defined as a decision-making problem in which, given a set of past cash flow observations determining an initial cash balance, the goal is to find the best sequence of control actions, which is called a policy, in terms of cost and risk.

In order to solve the MOCMP, we rely on compromise programming (CP) (Zeleny, 1982; Yu, 1985; Ballestero and Romero, 1998; Ballestero and Pla-Santamaria, 2004) as a suitable technique to derive the best policies by minimizing weighted distances to an ideal point, where both cost and risk independently take minimum values subject to the restrictions of the problem. Under the CP framework, these weights reflect the particular preferences of cash managers. An important advantage of CP in practical applications is the possibility to specify these preferences in a deployment phase. Then, we follow a two-step decision-making process in which we present alternatives before selection. In the context of cash management, a set of alternative policies is obtained in a learning phase and presented to cash managers. Later, a policy is selected according to their particular preferences in a deployment phase.

We first describe a general formulation of the MOCMP as suggested by Salas-Molina et al. (2016). More precisely, we consider a cash balance that starting at an initial value fluctuates according to a particular cash flow process in absence of control actions. At any time, cash managers can take a control action by increasing/decreasing the cash balance but paying a transaction cost. The resulting cash balance at the end of a particular time period is finally determined by the control action and the net cash flow occurred and it is charged with some holding cost. Since risk analysis is incorporated in the MOCMP, we pay particular attention to the pros and cons of different risk measures, such as the variance or the semivariance of daily costs. Moreover, we focus on the problem of estimating large losses as an issue of special concern for cash managers by defining novel risk measures that are able to capture the effect of large losses.

Once cash managers have defined the set of decision criteria, usually cost and risk but may be others, and the particular objective functions that best fit their requirements, they are in a position to use CP to find and select the best policies that solve the MOCMP. To this end, we propose different solvers to cover a wide range of possible situations: (i) Monte Carlo methods; (ii) linear programming (LP); and (iii) quadratic programming (QP). Monte Carlo methods allow for a simulation strategy presenting policies before selection according to the particular risk preferences of cash managers. On the other hand, the linear and quadratic programming counterparts of compromise programming models result in an more automated decision-making technique when preferences and the extreme values of both cost and risk objectives can be reasonably estimated by cash managers.

In this chapter, we also consider an additional question as posed by Daellenbach (1974): *Are cash management models worthwhile*? We aim to answer this question from a multiobjective perspective, when less is better, by comparing the loss derived from a given policy to the loss derived from a trivial policy that takes no control action and hence lets cash balance freely wander around. As a result, we provide a formal definition of the cash management utility problem (CMUP) within a multiobjective framework.

Despite the recent advances in cash management and multiobjective decision-making, there is a lack of supporting technology to aid the transition from theory to practice. In order to fill this gap, we provide a free software library in Python for practitioners interested in either building software applications based on CP to solve the MOCMP or performing their own experiments.

In this chapter, we contribute to empower cash managers through compromise programming by: (i) defining novel risk measures to incorporate risk analysis in cash management; (ii) suggesting three alternative solvers of the MOCMP to cover a wide range of possible real-world situations; (iii) providing a Python software library with the proposed solvers ready to be embedded in cash management decision support systems; and (iv) providing a framework to assess the ability of cash management models when dealing with multiple objectives.

In what follows, we first review previous works on quantitative cash management. Then, we formulate the MOCMP in Section 3. Since risk analysis is incorporated in the definition of the MOCMP, we explore alternative risk measures in Section 4. Next, we propose three different solvers of the MOCMP in Section 5. We formulate the cash management utility problem in Section 6. Finally, we provide some concluding remarks in Section 7.

2 Literature review

In this section, we review the most recent cash management literature that is relevant to this chapter. The quantitative approach to the cash management problem dates back to the works by Baumol (1952), in a deterministic context, and by Miller and Orr (1966), for stochastic cash flows. We refer the interested reader to the surveys by Gregory (1976) and Srinivasan and Kim (1986) for works dated up to 1986 and to da Costa Moraes et al. (2015) for subsequent cash management research. We next focus on the most recent contributions to the field of quantitative cash management.

A common assumption in recent cash management works is the use of diffusion processes to represent cash flows. Premachandra (2004) followed such an assumption to propose a generalized version of the Miller and Orr (1966) model. Baccarin (2002, 2009) also used a diffusion process to control cash management systems with generalized cost functions. Bar-Ilan et al. (2004) represented cash flows as a superposition of a Brownian motion and a compound Poisson process. Similarly, da Costa Moraes and Nagano (2014) assumed Gaussian cash flows in their experiments to compare two approximate techniques to solve the cash management problem. A closely related topic is whether future cash flows can be predicted and be ultimately used to reduce cost. Gormley and Meade (2007) claimed the utility of forecasting in cash management and Salas-Molina et al. (2017) proved that predictive accuracy is highly correlated with cost savings.

After defining the main characteristics of the cash flow process under study, researchers have proposed a number of alternative models to control cash balances. Bensoussan et al. (2009) proposed a model with dividends and uncertain capital gains of idle cash balances invested in stock. Melo and Bilich (2013) proposed an expectancy balance model to deal with uncertainty of both deterministic and stochastic cash flows grouped into intervals of occurrence. Herrera-Cáceres and Ibeas (2016) proposed a model predictive control approach in which a given cash balance function is used as a reference trajectory to be followed by means of the appropriate control actions.

Cash management models require a method to solve the problem, i.e., to derive the policy that will ultimately be deployed by cash managers. Dynamic programming (Chen and Simchi-Levi, 2009), finite element methods (Baccarin, 2009), or approximate techniques such as genetic algorithms (Gormley and Meade, 2007) or particle swarm optimization (da Costa Moraes and Nagano, 2014) are some examples. An important question regarding alternative solvers is the optimality of solutions, which is a desired objective, but that has to be balanced with computational and deployment cost.

Finally, it is important to highlight that cash management models presented in the literature have focused only on cost as a minimization goal with the exception of Salas-Molina et al. (2016). The authors proposed a multiobjective approach in which both the cost and the risk of alternative policies are minimized. As a measure of risk, they proposed the use of the standard deviation and the upper semideviation of daily costs. In this chapter, we follow this approach to propose different risk measures and three different methods to solve the MOCMP.

3 Formulation of the MOCMP

In this section, we first formulate the MOCMP as proposed by Salas-Molina et al. (2016). Within a single objective framework, consider a firm with a given cash flow process expressed either as a probability distribution or a as set of past cash flows observations. The cash management problem (CMP) is defined as an optimization problem that aims to find the best policy $X = \{x_t : t = 1, 2, \ldots, n\}$ with $x_t \in \mathbb{R}$ that minimizes some objective function over a time horizon of n days. Positive (negative) control actions are charged with a fixed cost $\gamma_0^+(\gamma_0^-)$ and a variable cost $\gamma_1^+(\gamma_1^-)$ per money unit. In addition, cash balances at the end of each time step are charged with either a holding cost v for positive cash balances or a penalty cost u for a negative cash balances. As a result, a general daily cost function $c(x_t)$ is expressed as:

$$c(x_t) = \Gamma(x_t) + H(b_t) \tag{1}$$

where x_t is the transfer made at time step t, $\Gamma(x_t)$ is a transfer cost function, $H(b_t)$ is a holding/shortage cost function, and b_t is the cash balance at the end of time step t, determined by the next cash balance state equation:

$$b_t = b_{t-1} + x_t + f_t \tag{2}$$

being f_t the net cash flow occurred at time step t. However, since decisions are made in advance to real cash flow, both predicted cash flows $(\hat{f}_t)$ and predicted cash balances $(\hat{b}_t)$ are used instead. The transfer cost function $\Gamma(x_t)$ is defined as:

$$\Gamma(x_t) = \begin{cases} \gamma_0^- - \gamma_1^- \cdot x_t & \text{if} \quad x_t < 0, \\ 0 & \text{if} \quad x_t = 0, \\ \gamma_0^+ + \gamma_1^+ \cdot x_t & \text{if} \quad x_t > 0. \end{cases} \tag{3}$$

Additionally, the holding/shortage cost function is expressed as:

$$H(\hat{b}_t) = \begin{cases} -u \cdot \hat{b}_t & \text{if} \quad \hat{b}_t < 0; u > 0, \\ v \cdot \hat{b}_t & \text{if} \quad \hat{b}_t > 0; v > 0. \end{cases} \tag{4}$$

Given an initial cash balance b_0, the solution to the CMP, namely, the policy X, that minimizes the sum of transaction and holding costs, up to time step n, can be obtained by solving the following optimization problem:

$$\min C(X) = \min \sum_{t=1}^{n} \Big(\Gamma(x_t) + H(\hat{b}_t) \Big) \tag{5}$$

subject to:

$$\hat{b}_t = \hat{b}_{t-1} + \hat{f}_t + x_t \tag{6}$$

$$\hat{b}_t \geq 0 \tag{7}$$

$$x_t \in S \tag{8}$$

$$t = 1, 2, \ldots, n. \tag{9}$$

Since cash managers usually discard policies including overdrafts, we restrict the feasibility space to non-negative cash balances which is equivalent to set $u = \infty$ in equation (4). Set S contains all possible transactions determined by the cash management model, e.g., the cash management model proposed by Miller and Orr (1966), which is based on two bounds and a target level. Nevertheless, cash manager may be also interested in the risk of alternative policies. As a result, given a cost structure and an initial cash balance, we aim to solve the MOCMP by finding the best policy X, that delivers the best combination in terms of cost and risk over a planning horizon of n time steps:

$$\min_{X} \quad [C(X), R(X)] \tag{10}$$

subject to $X \in S$, where $C(X)$ and $R(X)$ denote general cost and risk functions, respectively. In order to include risk in the analysis of cash policies, we next consider alternative measures of risk.

4 Risk analysis in cash management

In this section, we aim to answer the question: how can we measure risk in cash management? To this end, we first provide a basic framework for risk analysis in cash management. Next, we define a number of risk measures, and we finally summarize the pros and cons for each of the suggested risk measures.

4.1 Measuring risk in cash management

One may hypothesize that risk is incorporated in the decision-making process of cash management by considering high penalty costs on negative cash balances. This view implies that high cost policies are also high risk policies. However, within the range of low cost policies, decision-makers may prefer, for instance, policies with the lowest variability in cost due to the less uncertainty involved. Intuitively, risk is associated to uncertainty, danger, chance of loss or damage. It is not the damage itself but the chance of it, the possibility of occurrence. A general definition of risk in a financial context can be found in McNeil et al. (2005), who consider risk as any event or action that may adversely affect an organization's ability to achieve its objectives and execute its strategies. To some extent, managers can choose the risks that a business takes (Brealey and Myers, 2003). Quantitatively, risk is also linked to unexpected losses. For example, risk management is an important task in investment because different assets offer different degrees of risk. In the well-known mean-variance model for portfolio selection proposed by Markowitz (1952), profitability is measured by the mean of returns, and risk by the variance of returns over a given period of time in the past.

The notion of risk is closely related to the concept of randomness. To some extent, the particular variability of future cash flows provokes risk. For example, suppose that two different cash managers operate under the same cost scenario given by current bank conditions. Suppose also that, at the end of the year, total cash management costs are exactly the same for both of them. Who did better? Apparently, the answer is that both performed equally well. However, if we are told that one of the cash managers deals with very stable and foreseeable cash flows and the other one faces highly variable and unpredictable cash flows, the answer would be different. In practice, there are different approaches to measure risk in a financial context (McNeil et al., 2005):

1. Notional-amount approach. For instance, the risk of a portfolio of assets is defined as the sum of the notional values of the individual assets of a portfolio. In this case, the higher the values the higher the risk.

2. Factor-sensitivity measures. These measures provide the change in value associated to a given change in one of the underlying risk factors. For instance, the greeks in portfolios of derivatives.

3. Scenario-based measures. In this approach, a number of future scenarios are considered, e.g., a 10% increase in the USD/EUR exchange rate. Risk is then measured by the the maximum loss produced under all scenarios considered.

4. Risk measures based on loss distributions. These measures are based on statistical quantities that describe the distribution of a random variable over a given period of time. Examples include the Value-at-Risk, the Conditional Value-at-Risk and the variance, which we here accommodate to a cash management context. All of them summarize in a single value, the risk contained in a distribution modeling loss.

4.2 Alternative measures of risk

Since most modern risk measures are based on loss distributions (McNeil et al., 2005; Glasserman, 2003), we next consider risk measures for cash management based on loss distributions. To model risk from a probabilistic approach, let c be a cost random variable on the probability space defined by $(\Omega, \mathcal{C}, P)$. An element c in Ω is a realization of an experiment, $\mathcal{C}$ is the set of all possible events and P is the probability of an event. Consider that $c(x_t)$ is a general cost function $c : X \times \mathcal{T} \rightarrow \mathbb{R}$, that associates a cost to each control action x_t in policy X deployed at time $t \in \mathcal{T}$. The probability that random variable $c(x_t)$ is below some value c_0 is given by the cumulative distribution function:

$$F_c(c_0) = P(c(x_t) \leq c_0). \tag{11}$$

Thus, we first propose to measure the risk of policy X as the probability that $c(x_t)$ is above c_0, given by:

$$P(c(x_t) \geq c_0) = 1 - F_c(c_0). \tag{12}$$

Similarly to the definition of Value-at-Risk (McNeil et al., 2005), we here suggest to synthetically describe this cumulative distribution function by its moments such as the mean and variance, or by a quantile such as the Cost-at-Risk.

Definition 1. The Cost-at-Risk (CaR) of a cash policy X at a confidence level $\alpha \in [0, 1]$ is given by the smallest number c_0 such that the probability that the cost $c(x_t)$ exceeds c_0 is no larger than $1 - \alpha$, formally:

$$CaR(X, c, \alpha) = \inf\{c_0 \in \mathbb{R} | P(c(x_t) \geq c_0) \leq 1 - \alpha\}, \ \ \forall x_t \in X \tag{13}$$

or alternatively:

$$CaR(X, c, \alpha) = \inf\{c_0 \in \mathbb{R} | F_c(c_0) \geq \alpha\}, \ \ \forall x_t \in X. \tag{14}$$

Notice that the CaR of policy X depends on the definition of cost function c and threshold α. Typical values for α are 0.95 or 0.99. Figure 1 illustrates the notion of CaR. Say that from a number of experiments, the empirical average daily cost is distributed as shown in the figure. As an example, assume also that a Weibull distribution (Weibull et al., 1951) is the function that best fits the empirical data. If α is 0.95, we can then expect that the cost exceeds 2200 with probability 0.05.

Probably the major drawback of CaR is that it does not provide information about the severity of losses beyond c_0. In the usual case of heavy-tailed distributions, the estimation of large losses is an important question to be considered. In the example of Figure 1, the last two bars could be located at points 5000 an 5500 and the CaR would remain unaltered. Significant advantages over CaR are provided by the Conditional Cost-at-Risk ($CCaR$) measure, which we define as the conditional excess expectation, similarly to the definition of Conditional Value-at-Risk in Rockafellar and Uryasev (2002).

Definition 2. We define the Conditional Cost-at-Risk ($CCaR$) value of a cash policy X at a confidence level $\alpha \in [0, 1]$ as:

$$CCaR(X, c, \alpha) = E[c(x_t) | c(x_t) > c_0], \ \ \forall x_t \in X. \tag{15}$$

where c_0 is the cost such that the probability that $c(x_t)$ exceeds c_0 is no larger than $1 - \alpha$.

An additional advantage of $CCaR$ is that it is a coherent measure of risk in the sense of Artzner et al. (1999). In practice, when r out of n realizations $\{c(x_1), \ldots, c(x_r)\}$ of a given policy X are above c_0, the $CCaR$ value can be obtained as:

$$CCaR(X, c, \alpha) = E\left(\{c(x_1), \ldots, c(x_r)\}\right) = c_0 + \frac{1}{r}\sum_{t=1}^{n} \max(c(x_t) - c_0, 0) \quad (16)$$

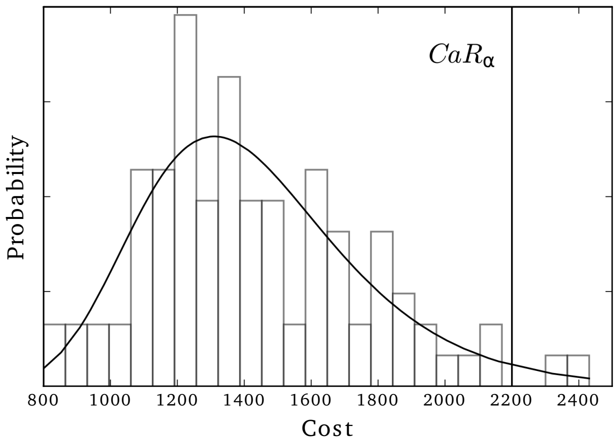

Figure 1: Probability density function fitted to the empirical histogram of a cost random variable and its CaR_α quantile.

Alternatively, the problem of large losses is also taken into account by using variance as a risk measure. Since the mean-variance model by Markowitz (1952), variance has been extensively used in finance. Moreover, its simplicity and ease of computation in experimental environments makes variance a good risk measure. Variance makes no distinction between positive and negative deviations and cash managers are usually more interested in positive deviation of cost. Semivariance or upside/downside deviation solves this problem (Ballestero, 2005; Pla-Santamaria and Bravo, 2013). In the context of cash management, we propose to calculate the risk of a policy X by computing the variance of daily costs as follows:

Definition 3. The variance (V) of a cash policy X deployed over n time steps is obtained as:

$$V(X, c) = \frac{1}{n}\sum_{t=1}^{n}(c(x_t) - E(c))^2, \quad \forall x_t \in X. \quad (17)$$

where $E(c)$ is the expected cost of policy X over n.

The underlying assumption on the use of variance is that the more disperse the costs within a policy around the expected cost, the higher the risk of the policy. However, since cash managers are probably more interested in upside deviations of cost rather than downside deviations, an upper partial moment such as the semivariance may be considered as an alternative measure of risk.

Following McNeil et al. (2005), given an exponent k and a reference point c_0, we here propose an additional measure of risk for a policy X as follows:

Definition 4. The k-Upper Partial Moment (UPM) with respect to c_0 of a cash policy X is obtained as:

$$UPM(X, c, k, c_0) = \int_{c_0}^{\infty} (c - c_0)^k dF_c \tag{18}$$

where F_c is the cumulative distribution of the density function of cost c.

Note that if $k = 0$, then $UPM(X, c, 0, c_0) = P(c \geq c_0)$, is the probability that the cost exceeds the reference c_0. Additionally, if $k = 1$, then $UPM(X, c, 1, c_0)$ is the expected upper deviation of cost from the reference c_0. Finally, when $k = 2$ and c_0 is set to the expected cost, then $UPM(X, c, 2, E(c))$ is the upper semivariance of cost. However, since common planning and control practices in most organizations are typically performed in discrete intervals indexed by time step t, the $UPM(X, c, k, c_0)$ can be computed in discrete time as:

$$UPM(X, c, k, c_0) = E\left((\max\{c(x_t) - c_0, 0\})^k\right) \tag{19}$$

4.3 Summary of risk measures

As a summary, the pros and cons for the aforementioned risk measures are presented in Table 1. When dealing with risk, cash managers are usually concerned not only with average variation but also with abnormal or extreme values (McNeil et al., 2005; Glasserman, 2003). The risk of large losses must then be considered and, although CaR considers heavy tails, it does not provide information about the severity of large losses. Thus $CCaR$, variance or UPM can be used instead. The use of standard deviation is preferred to variance because it presents the same units as cost, i.e., money units, and numerical comparisons are then possible. However, a drawback must be pointed out against variance or standard deviation since there is no distinction between positive and negative deviations. This problem is easily solved by considering UPM, such as the upper semivariance. Non-linearity is another important aspect to be considered, specially when using this measure as part of an objective function to be minimized. Linear objective functions and linear constrains are usually preferred in mathematical programming. In this sense, the $CCaR$ value should be considered as a good risk measure since it can be easily expressed as a linear function in an optimization problem.

Table 1: Advantages and disadvantages of alternative risk measures.

Measure	Advantages	Disadvantages
Cost-at-Risk	Considers heavy tails	No large losses
Conditional Cost-at-Risk	Large losses and linear	Selection of cost c_0
Variance	Large losses	Symmetric, quadratic
Upper partial moments	Large losses	Non-linear for $k \geq 2$

In what follows, we focus on risk measures that allow to formulate the MOCMP as a linear program such as the $CCaR$, or as quadratic program such as variance or standard deviation. Notice that by using an empirical statistic, we make no assumption on the underlying probability distribution.

5 Compromise models to solve the MOCMP

Recall from the introduction that we aim to derive cash policies that minimize a weighted loss function in which both cost an risk are desired objectives. To this end, we rely on compromise programming models and three different solvers: (i) Monte Carlo methods; (ii) linear programming; and (iii) quadratic programming. While Monte Carlo methods provide approximate solutions, both linear and quadratic programming guarantee the optimality of solutions.

Compromise programming is based on the concept of ideal point and the Zeleny's axiom of choice (Zeleny, 1974), which states that alternatives that are closer to the ideal are better than those that are further. The concept of ideal point is at the core of compromise programming. When less is better, the minimum values for each objective subject to the constraints of the problem determine the ideal point. In the context of the MOCMP, the ideal point in a bidimensional cost-risk space is the point with zero cost and zero risk that simultaneously minimizes $C(X)$ and $R(X)$ in objective function (10). Since this ideal point is usually unfeasible, it is necessary to look for compromise solutions by minimizing the distance to this ideal point. A general distance function between two bidimensional points $P_1 = (x_1, y_1)$ and $P_2 = (x_2, y_2)$ is the Minkowski distance of order h, defined as:

$$\left(|x_1 - x_2|^h + |y_1 - y_2|^h\right)^{1/h}. \tag{20}$$

By computing the distance between the ideal point $(0, 0)$ and any particular point, we are in a position to determine whether a given solution is better than another one. However, when the scale used to measure goals is different, in order to avoid a meaningless comparison, each goal has to be normalized. In the MOCMP, we can define a cost index (θ_1) and a risk index (θ_2) as follows:

$$\theta_1(X) = \frac{C(X) - C_{min}}{C_{max} - C_{min}} \tag{21}$$

$$\theta_2(X) = \frac{R(X) - R_{min}}{R_{max} - R_{min}} \tag{22}$$

where C_{max} (R_{max}) and C_{min} (R_{min}) are, respectively, the maximum and minimum values of cost function C (risk function R) subject to the constraints of the problem. Note that due to normalization $\theta_1, \theta_2 \in [0, 1]$. Consequently, the closer to the ideal point $(0, 0)$, the better the solution. Moreover, when considering particular goal preferences, weighted distances must be computed instead. From that, CP proposes a family of normalized distance functions including weights that determine the decision-maker's risk preferences as:

$$\mathcal{L}_h = \left[w_1^h \cdot \theta_1^h + w_2^h \cdot \theta_2^h\right]^{1/h}. \tag{23}$$

Note that $\mathcal{L}_1$ is the Manhattan distance; $\mathcal{L}_2$ is the Euclidean distance, and $\mathcal{L}_\infty$ is the Chebyshev distance. They are the most used distances in practice for interpretation and computational reasons (Ringuest, 1992; Ballestero, 2007).

On the other hand, weights w_1 and w_2 in equation (23) reflect the particular preferences of cash managers. As a result, considering Manhattan distances to avoid non-linearity, we can formulate the MOCMP as the following CP model:

$$\min \ [w_1 \cdot \theta_1(X) + w_2 \cdot \theta_2(X)] \tag{24}$$

$$X \in S. \tag{25}$$

Next, we consider three alternative methods to solve the MOCMP: (i) Monte Carlo methods; (ii) linear programming; and (iii) quadratic programming. We prefer Monte Carlo methods when exploring alternatives within a bounded set, when the specific cost/risk preferences are not known in the learning phase, or when there are reasonable doubts about the minimum/maximum values in (21) and (22). However, if we can express the objectives and the constraints in (24)-(25) as linear functions, and we know both weights w_1 and w_2, and the extreme values in (21) and (22), we can automate the cash management decision-making process by solving the MOCMP without cash managers' intervention using linear programming. In addition, if any of the objective functions is quadratic (e.g., when using variance as a measure of risk), we can use quadratic programming. As a result, we can solve both linear and quadratic programs using state-of-the-art solvers such as CPLEX or Gurobi.

5.1 Solving the MOCMP by Monte Carlo methods

Assume that we want to deploy a policy of the Miller and Orr (1966) type based on three control bounds: a lower bound l_1, a target level l_2 and an upper bound l_3. Cash balances are allowed to wander around between bounds l_1 and l_3, and when any of these bounds is reached, a control action is taken to restore the cash balance to the target level as described in the following expression:

$$x_t = \begin{cases} l_2 - \hat{b}_{t-1}, & \text{if} \quad \hat{b}_{t-1} > l_3 \\ 0, & \text{if} \quad l_1 \leq \hat{b}_{t-1} \leq l_3 \\ l_2 - \hat{b}_{t-1}, & \text{if} \quad \hat{b}_{t-1} < l_1 \end{cases} \tag{26}$$

where $\hat{b}_{t-1}$ is the cash balance previous to control action x_t.

As a result, under the framework of compromise programming for $\mathcal{L}_1$ in equation (23), solving the MOCMP for a given planning horizon of n time steps is equivalent to finding the set $\{l_1, l_2, l_3\}$, which minimizes the weighted Manhattan distance to the ideal point $(0, 0)$:

$$\min \ \left[w_1 \cdot \frac{C(X) - C_{min}}{C_{max} - C_{min}} + w_2 \cdot \frac{R(X) - R_{min}}{R_{max} - R_{min}} \right] \tag{27}$$

subject to:

$$\hat{b}_t = \hat{b}_{t-1} + \hat{f}_t + x_t \tag{28}$$

where $X = \{x_t : t = 1, 2, \ldots, n\}$ with x_t according to equation (26) and bounds satisfying $0 \geq l_1 \geq l_2 \geq l_3$. We here measure cost by the average daily cost and risk by the standard deviation of daily cost as follows:

$$C(X) = E(C) = \frac{1}{n} \sum_{t=1}^{n} c(x_t) \tag{29}$$

$$R(X) = \left(\frac{1}{n} \sum_{t=1}^{n} (c(x_t) - E(C))^2 \right)^{1/2}. \tag{30}$$

An advantage of CP in practical applications is the possibility to specify these preferences in a deployment phase. Then, we follow a two-step decision-making process in which alternatives are presented before selection. Thus, since weights w_1 and w_2 are unknown at this point, we aim to obtain a Pareto efficient set of solutions (Yu, 1985). In other words, we want to derive an efficient frontier with policies X not dominated by any other policy in terms of cost and risk. A suitable and simple method to obtain this efficient set is Monte Carlo simulation (Glasserman, 2003). Monte Carlo methods are based on performing a high number of random experiments that are later evaluated in some outcome domain. The law of large numbers ensures that the estimations derived from this analysis converge to real values as the number of experiments increases. An example of a Monte Carlo method is summarized in the steps detailed in Algorithm 1.

Algorithm 1: Montecarlo method to solve the multiobjective cash management problem

Inputs: Model m; set of goals indexes θ_i; cost context β; cash flow data set F; time period n; replicates r;
Output: Estimation of the context efficient set of solutions;
for *each replicate* **do**
    Generate a random solution X;
    for *each goal* **do**
        Compute $g_i(X)$;
    end
end
for *each goal* **do**
    Compute θ_i ;
end
Compute the efficient set;

As an illustrative example, consider a cost context β defined by the following cost scenario: $\gamma_0^+ = \gamma_0^- = 200$ €, $\gamma_1^+ = 0.1\%$, $\gamma_1^- = 0$, $v = 0.1\%$, $u = 30\%$. Assume also the following sequence of expected cash flows ($\hat{F}$) starting at an initial cash balance of 20, all figures in millions of euros, is:

$$\hat{F} = [1, 1, 6, -1, -3, -3, -9, 6, 4, 6, 3, 4, 1, -1, -2, 2]. \tag{31}$$

After applying Algorithm 1 with 10,000 replicates to our example, we obtain the efficient set summarized in Table 2. We depict in Figure 2 the efficient frontier derived from the set of (θ_1, θ_2) values in Table 2. As expected, there is a cost-risk tradeoff and lower costs can only be achieved by accepting higher risks. Cash managers can obtain a compromise solution by selecting the policy with the minimum Manhattan distance to the ideal point $(0, 0)$ according to their risk/cost preferences. In the case of unbiased preference for cost or risk, i.e., $w_1 = w_2$, the best solution to our example is policy 2 with control bounds 1, 11 and 20. However, a conservative cash manager may choose policy 4 in order to reduce risk but accepting a higher cost.

Table 2: Example efficient set for a Miller-Orr model with three levels.

Id	l_3	l_2	l_1	Cost	Risk	θ_1(cost)	θ_2(risk)	$\theta_1 + \theta_2$
1	20	10	1	12963	5687	0,00	1,00	1,00
2	20	11	1	14275	4626	0,14	0,33	0,47
3	20	12	0	15213	4541	0,25	0,27	0,52
4	20	13	4	16150	4467	0,35	0,23	0,58
5	28	13	13	21400	4162	0,93	0,03	0,96
6	28	14	13	22025	4108	1,00	0,00	1,00

The closer to the ideal point $(0, 0)$, the better the policy. However, not all policies are relevant to the decision-maker according to their risk preferences. As suggested in in Ballestero (1998), we can express risk preferences as a parameter $r_0 \in \mathbb{R}_+$, equivalent to the number of marginal units of risk (θ_2) that the cash manager is willing to accept in order to achieve a decrease of one marginal unit of cost (θ_1). Linking r_0 and (w_1, w_2), if $r_0 = 0.5$, a conservative cash manager is willing to accept only 0.5 units of risk for each unit of decreased cost, then $w_1 = 0.33$ and $w_2 = 0.67$.

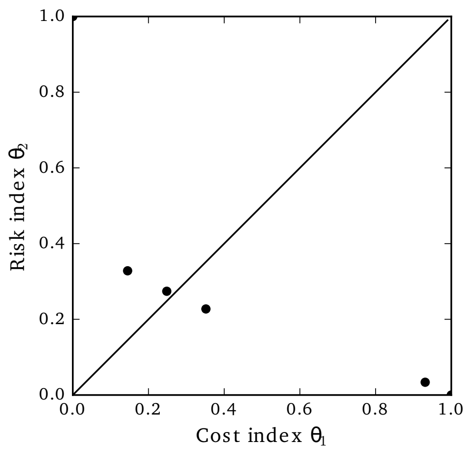

Figure 2: Example of efficient set.

Moreover, it is proven that the utility optimum for a decision-maker lies on the normalized efficient frontier between points L and L_∞. On the one hand, bound L is the point minimizing the linear loss function $r_0\theta_1 + \theta_2$ on the normalized efficient frontier. On the other hand, bound L_∞ is the intersection of $\theta_1 = \theta_2$ with the efficient frontier. For instance, consider a conservative cash manager with $r_0 = 0.5$. From Table 2, bound L, with minimum $r_0\theta_1 + \theta_2$, is policy 3, and bound L_∞, when cost index θ_1 approximately equals risk index θ_2, coincides with policy 3. However, when the specific risk preferences of cash managers are known (or known to lie in a given interval), and there is no reasonable estimation doubt about the minimum/maximum values for both cost and risk, a more straightforward procedure can be followed by relying on linear or quadratic programming.

5.2 Solving the MOCMP by linear programming

From the set of alternative risk measures considered in Section 4, consider now $CCaR$ as a measure of risk. Recall that $CCaR$ is defined as the expected cost above a given reference c_0. Thus, minimizing $CCaR$ is equivalent to minimizing the sum of positive cost deviations from cost reference c_0, which can be cast as a Goal Programming (GP) model (Abdelaziz et al., 2007; Aouni et al., 2014). GP aggregates multiple objectives to obtain the solution that minimizes the sum of deviations between achievement and the aspiration levels of the goals. Then, we proceed as follows:

1. We define the goals that are relevant for the cash manager, e.g., cost and risk.
2. We set the aspiration level or target (τ_i), for each goal g_i, with $1 \leq i \leq q$.
3. We introduce both positive (δ_i^+) and negative (δ_i^-) deviation auxiliary variables to connect individual goal achievement and targets.

In the GP setting, the particular risk preferences of cash managers can be incorporated to determine the relative importance of each goal by means of a set of positive (w_i^+) and negative weights (w_i^-). Then, a general Weighted Goal Programming (WGP) model is expressed as follows:

$$\min \sum_{i=1}^{q} (w_i^+ \delta_i^+ + w_i^- \delta_i^-) \tag{32}$$

subject to:

$$g_i + \delta_i^- - \delta_i^+ = \tau_i \tag{33}$$

$$\delta_i^-, \delta_i^+ \geq 0, \quad i = 1, 2, \ldots, q. \tag{34}$$

It is important to highlight the close link between CP and GP models. Indeed, a CP minimization problem for Manhattan distances ($h = 1$) is analytically equivalent to a GP problem when both target values (τ_i) and negative deviations (δ_i^-) are set to zero, and positive deviations are set to:

$$\delta_i^+ = \theta_i = \frac{g_i - g_{min,i}}{g_{max,i} - g_{min,i}}. \tag{35}$$

Let us consider again our MOCMP with two goals, namely, cost and risk, aggregated through the CP model encoded in equations (24)-(25). Note that cost function $c(x_t)$ in equation (1) is clearly non-linear. In order to linearize it, we rewrite control action x_t as the difference of two non-negative variables:

$$x_t = x_t^+ - x_t^-. \tag{36}$$

Let us introduce two binary variables $z_t^+, z_t^- \in \{0, 1\}$ linked to controls x_t^+, x_t^- by means of the following constraints:

$$k \cdot z_t^+ \leq x_t^+ \leq K \cdot z_t^+ \tag{37}$$

$$k \cdot z_t^- \leq x_t^- \leq K \cdot z_t^- \tag{38}$$

where $K(k)$ is a very large (small) number. Note that constraint (37) ensures that $z_t^+ = 1$ when x_t^+ occurs, $z_t^+ = 0$, otherwise. Similarly, constraint (38) ensures that $z_t^- = 1$ when x_t^- occurs, $z_t^- = 0$, otherwise. Furthermore, we impose $z_t^+ + z_t^- \leq 1$ to avoid the simultaneous occurrence of x_t^+ and x_t^-. As a result, we can rewrite equation equation (1) as follows:

$$c(x_t) = \gamma_0^+ \cdot z_t^+ + \gamma_1^+ \cdot x_t^+ + \gamma_0^- \cdot z_t^- + \gamma_1^- \cdot x_t^- + v \cdot \hat{b}_t. \tag{39}$$

After reasonably setting C_{min} and R_{min} to zero, due to the fact that both zero cost and zero risk policies can be independently achieved, we define the following cost and risk indexes for policy X:

$$\theta_1(X) = \frac{1}{C_{max}} \sum_{t=1}^{n} c(x_t) \tag{40}$$

$$\theta_2(X) = \frac{1}{R_{max}} \sum_{t=1}^{n} \delta_t^+, \tag{41}$$

where δ_t^+ is the positive deviation from a given cost reference c_0, equivalent to $CCaR$. Then, we formulate the following LP model using the total cost as a measure of cost, and $CCaR$ as a measure of risk:

$$\min \quad \left[\frac{w_1}{C_{max}} \sum_{t=1}^{n} c(x_t) + \frac{w_2}{R_{max}} \sum_{t=1}^{n} \delta_t^+ \right] \tag{42}$$

subject to:

$$\hat{b}_t = \hat{b}_{t-1} + \hat{f}_t + x_t^+ - x_t^- \tag{43}$$

$$c(x_t) = \gamma_0^+ \cdot z_t^+ + \gamma_1^+ \cdot x_t^+ + \gamma_0^- \cdot z_t^- + \gamma_1^- \cdot x_t^- + v \cdot \hat{b}_t \tag{44}$$

$$z_t^+ + z_t^- \leq 1 \tag{45}$$

$$k \cdot z_t^+ \leq x_t^+ \leq K \cdot z_t^+ \tag{46}$$

$$k \cdot z_t^- \leq x_t^- \leq K \cdot z_t^- \tag{47}$$

$$c(x_t) - \delta_t^+ \leq c_0 \tag{48}$$

$$\hat{b}_t \geq b_{min} \tag{49}$$

$$w_1 + w_2 = 1 \tag{50}$$

$$\sum_{t=1}^{n} c(x_t) \leq C_{max} \tag{51}$$

$$\sum_{t=1}^{n} \delta_t^+ \leq R_{max} \tag{52}$$

$$z_t^+, z_t^- \in \{0, 1\} \tag{53}$$

$$x_t^+, x_t^-, \hat{b}_t, \delta_t^+ \geq 0 \tag{54}$$

where the main decision variables are control actions x_t^+ and x_t^-. In practice, C_{max} and R_{max} can be regarded as budget limitations for both cost and risk, leading to unfeasible policies when these constraints are no satisfied.

Since we use cash flow forecasts, cash managers may be interested to protect themselves against forecasting errors. This protection can be achieved through a minimum cash balance b_{min}. For instance, by setting a minimum cash balance equivalent to the maximum forecasting error, we transform an optimization problem affected by uncertainty into its robust counterpart as proposed by Soyster (1973) and Ben-Tal et al. (2009). It is also important to highlight that we do not impose any additional constraint on the form of the policy, apart from non-negativity. We refer to that kind of policies as being produced by an unconstrained cash management model.

As a numerical example consider again the cost context β, and the set of expected cash flows $(\hat{F})$ for the next $n = 16$ days detailed in Section 5.1. Using the total cost as a measure of cost and the $CCaR$ as a measure of risk, we can solve the MOCMP by minimizing objective function (42), with 96 decision variables detailed as follows:

- 16 ordering transactions x_t^+;
- 16 returning transactions x_t^-;
- 16 auxiliary binary variables z_t^+ for fixed costs of ordering transactions;
- 16 auxiliary binary variables z_t^- for fixed costs of returning transactions;
- 16 expected cash balance variables $\hat{b}_t$;
- 16 positive deviation variables δ^+.

Assume that a cash manager is biased for cost such that $w_1 = 0.67$ and $w_2 = 0.33$. For precautionary purposes, she sets a minimum cash balance of two standard deviations of the expected cash flow ($b_{min} = 7$). The solution of this MOCMP results in the optimal cash policy and balance shown in Figure 3. This policy produces a total cost of 133,600 €, equivalent to an average daily cost of 8,350 €, and a total risk of of 10,800 €, in terms of $CCaR$ with respect to cost reference $c_0 = 8,000$ €, representing a combined 62% of the total maximum budget constraints determined by $C_{max} = 0.15$ and $R_{max} = 0.15$, both figures in millions of euros.

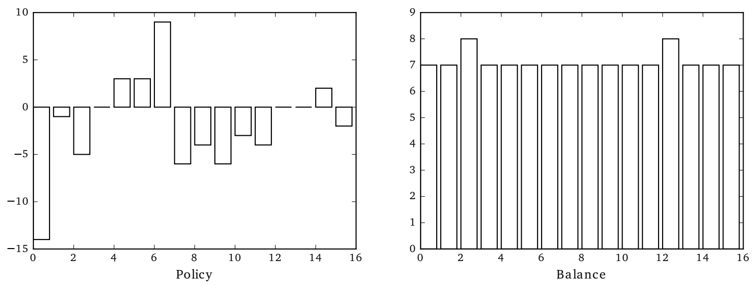

Figure 3: Policy and balance for the example using LP.

5.3 Solving the MOCMP by quadratic programming

Consider now daily cost variance as a measure of risk. We denote $\boldsymbol{c}$ as an $n \times 1$ vector of daily costs, and $\boldsymbol{d}$ as an $n \times 1$ vector of cost deviations around the average which can be computed as:

$$\boldsymbol{d} = \boldsymbol{c} - \frac{\mathbf{1} \cdot \mathbf{1}^T \cdot \boldsymbol{c}}{n} \tag{55}$$

where $\mathbf{1}$ is an $n \times 1$ vector of ones. Cash managers aiming to minimize only cost variance, or tantamount standard deviation, can derive optimal policies through the following quadratic objective function:

$$\min \quad \frac{\boldsymbol{d}^T \cdot \boldsymbol{d}}{n}. \tag{56}$$

Similarly to Section 5.2, let us consider an $n \times 1$ vector of positive (negative) transactions $\boldsymbol{x}^+ (\boldsymbol{x}^-)$ and an $n \times 1$ vector of expected balances $\hat{\boldsymbol{b}}$. We can then rewrite the state transition law in equation (43) in matrix notation as follows:

$$\hat{\boldsymbol{b}} = \hat{\boldsymbol{b}}_0 + L \cdot \left(\hat{\boldsymbol{f}} + \boldsymbol{x}^+ - \boldsymbol{x}^- \right) \tag{57}$$

where $\hat{\boldsymbol{b}}_0$ is an $n \times 1$ vector with all entries set to the initial cash balance, and L is an $n \times n$ lower triangular matrix with elements $l_{ij} = 1$ for all $i \geq j$. Furthermore, vector $\boldsymbol{c}$ can be computed by means of the following expression:

$$\boldsymbol{c} = \gamma_0^+ \cdot \boldsymbol{z}^+ + \gamma_1^+ \cdot \boldsymbol{x}^+ + \gamma_0^+ \cdot \boldsymbol{z}^- + \gamma_1^- \cdot \boldsymbol{x}^- + v \cdot \hat{\boldsymbol{b}} \tag{58}$$

where $\boldsymbol{z}^+, \boldsymbol{z}^- \in \mathcal{B}^n$ are, respectively, $n \times 1$ vectors of positive and negative binary variables, and $\mathcal{B}^n$ is an n-dimensional binary space. As a result, we can aggregate average cost and variance as a measure of risk to formulate the MOCMP as the following quadratic program:

$$\min \quad \left[\frac{w_1}{C_{max}} \frac{\mathbf{1}^T \cdot \boldsymbol{c}}{n} + \frac{w_2}{R_{max}} \frac{\boldsymbol{d}^T \cdot \boldsymbol{d}}{n} \right] \tag{59}$$

subject to:

$$\hat{\boldsymbol{b}} = \hat{\boldsymbol{b}}_0 + L \cdot \left(\hat{\boldsymbol{f}} + \boldsymbol{x}^+ - \boldsymbol{x}^- \right) \tag{60}$$

$$\boldsymbol{c} = \gamma_0^+ \cdot \boldsymbol{z}^+ + \gamma_1^+ \cdot \boldsymbol{x}^+ + \gamma_0^+ \cdot \boldsymbol{z}^- + \gamma_1^- \cdot \boldsymbol{x}^- + v \cdot \hat{\boldsymbol{b}} \tag{61}$$

$$\boldsymbol{d} = \boldsymbol{c} - \frac{\mathbf{1} \cdot \mathbf{1}^T \cdot \boldsymbol{c}}{n} \tag{62}$$

$$\boldsymbol{z}^+ + \boldsymbol{z}^- \leq \mathbf{1} \tag{63}$$

$$k \cdot \boldsymbol{z}^+ \leq \boldsymbol{x}^+ \leq K \cdot \boldsymbol{z}^+ \tag{64}$$

$$k \cdot \boldsymbol{z}^- \leq \boldsymbol{x}^- \leq K \cdot \boldsymbol{z}^- \tag{65}$$

$$\hat{\boldsymbol{b}} \geq \boldsymbol{b}_{min} \tag{66}$$

$$\frac{\mathbf{1}^T \cdot \boldsymbol{c}}{n} \leq C_{max} \tag{67}$$

$$\frac{\boldsymbol{d}^T \cdot \boldsymbol{d}}{n} \leq R_{max} \tag{68}$$

$$\boldsymbol{x}^+, \boldsymbol{x}^-, \hat{\boldsymbol{b}} \in \mathbb{R}^n_+ \tag{69}$$

$$\boldsymbol{z}^+, \boldsymbol{z}^+ \in \mathcal{B}^n \tag{70}$$

$$w_1 + w_2 = 1 \tag{71}$$

where vectors $\boldsymbol{x}^+(\boldsymbol{x}^-)$ are the main decision variables; $\boldsymbol{b}_{min}$ is an $n \times 1$ vector with all elements set to a given minimum balance.

Following with our example with $w_1 = 0.67$, $w_2 = 0.33$ and $b_{min} = 7$, if we set $C_{max} = 0.15$ and $R_{max} = 10$ millions of euros, we obtain the optimal cash policy and balance shown in Figure 4. This policy produces a total cost of 140,150 €, equivalent to an average daily cost of 8,787 €, and a total risk in terms of variance of 3,970,029, equivalent to a standard deviation of 1,992 €.

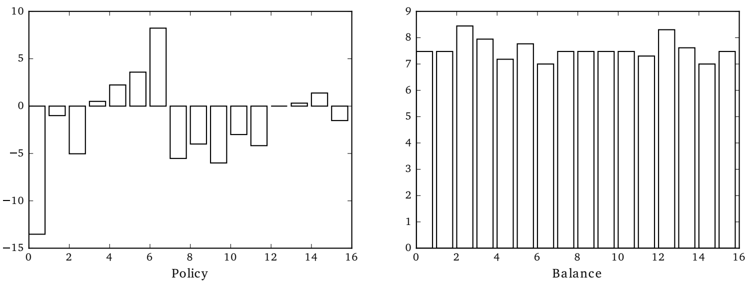

Figure 4: Policy and resulting cash balance for the example using QP.

5.4 Discussion

As a summary of the MOCMP solvers, some additional comments must be done on the pros and cons of the previous techniques. First, the selection of R_{max} could be tricky in the LP and QP approaches due to the difference in scale of concepts such as $CCaR$ or variance. When using variance as a measure of risk, standard deviation is a more known concept that can be used as a proxy to set the maximum accepted risk in terms of variance by squaring the maximum standard deviation value. In addition, when cash managers require to analyze either quantitatively or graphically the impact of h in equation (23) as well as of risk preferences, the Monte Carlo approach is a more suitable alternative. Otherwise, a more automated decision-making procedure can be followed by solving the MOCMP using linear or quadratic programming.

6 On the utility of cash management models

At this point, we formulate a fundamental question along the lines of Daellenbach (1974), who posed the following general question: *Are cash management models worthwhile*? Under a general CP framework, the answer to Daellenbach's question is equivalent to comparing the loss derived from policy X to the loss derived from a baseline policy X_0. As a result, we here introduce the Cash Management Utility Problem (CMUP) as follows:

Definition 5. The Cash Management Utility Problem is defined in a multiobjective framework, when less is better, as the problem of determining if policy X is preferred to a baseline policy X_0, formally expressed as:

$$[g_1(X), \ldots, g_q(X)] \leq [g_1(X_0), \ldots, g_q(X_0)] \tag{72}$$

where operator $\leq$ means that $g_i(X) \leq g_i(X_0)$ holds for all i in the range $[1, q]$, and at least, there is one i such that $g_i(X) < g_i(X_0)$.

Then, setting $X_0 = 0$, as a baseline policy consisting in taking no control action, any policy X is worthwhile if it is able to reduce the value of at least one of the considered objective functions in comparison to X_0. For instance, in the case of considering only cost, the previous comparison is equivalent to: $C(X) < C(X_0)$. Considering both general cost and risk measures as in equation (42), we here provide further insight by extending the question posed by Daellenbach (1974) to a cost-risk framework:

Definition 6. A policy X is preferred to a No-Trans policy $X_0 = 0$, in terms of cost and risk indexes θ_1 and θ_2, when less is better, if:

$$\theta_1(X) + \theta_2(X) \leq \theta_1(X_0) + \theta_2(X_0) \tag{73}$$

subject to:

$$X \in S. \tag{74}$$

The implications of the CMUP are twofold. First, practitioners may be interested in finding the external conditions that must hold to ensure the utility of a non-trivial policy. An example of this issue was pointed out by Constantinides and Richard (1978), showing that a No-Trans policy is the best alternative in terms of cost when $\gamma_1^+ > u$ and $\gamma_1^- > v$. Second, researchers may be interested in establishing the particular characteristics that both cash management models and alternative cost and risk measures must present in order to avoid non-triviality. As an example, consider the average daily cost as a measure of cost and the daily cost variance as a measure of risk as in Section 5.3. This setting reduces the CMUP to:

$$\mathbf{1}^T \cdot \boldsymbol{c} + \boldsymbol{d}^T \cdot \boldsymbol{d} \leq \mathbf{1}^T \cdot v \cdot \hat{\boldsymbol{b}}_{t,0} + \boldsymbol{d}_0^T \cdot \boldsymbol{d}_0 \tag{75}$$

subject to:

$$\boldsymbol{d} = \boldsymbol{c} - \frac{\mathbf{1} \cdot \mathbf{1}^T \cdot \boldsymbol{c}}{n} \tag{76}$$

$$\boldsymbol{d}_0 = v \cdot \hat{\boldsymbol{b}}_{t,0} - \frac{\mathbf{1} \cdot \mathbf{1}^T v \cdot \hat{\boldsymbol{b}}_{t,0}}{n}. \tag{77}$$

In other words, the utility of a particular cash management model in the previous multiobjective framework is given by the combined ability of the model to reduce both the cost and risk impact by introducing some control actions summarized in policy X. Furthermore, the CMUP can also be viewed as a precautionary tool to avoid unnecessary efforts in forecasting and mathematical programming tools when some inputs of the problem reduce the utility of the policy.

7 Concluding remarks

Within a dynamic context characterized by increasing uncertainty, cash managers can be empowered by following an integrated approach in which not only cost but also risk are optimized. To this end, we propose alternative measures to incorporate risk analysis into a multiobjective formulation of the cash management problem. We pay particular attention to the problem of estimating large losses as an issue of special concern for cash managers. As a result, apart from usual measures of risk such as variance or standard deviation, we introduce CaR, $CCaR$, and UPM as suitable measures to capture the effect of large losses.

To solve the MOCMP, we rely on a general compromise programming framework to find policies that minimize weighted distances to an ideal (but usually unfeasible) point of zero cost and zero risk. Once the cost and risk objective functions are defined, we propose three different solvers within the framework of compromise programming: (i) Monte Carlo methods; (ii) linear programming; and (iii) quadratic programming. Summarizing, two-stage Monte Carlo methods require intervention of cash managers to choose policies. On the other hand, the linear and quadratic programming counterparts of compromise programming models result in a more automated decision-making process when risk preferences and both cost and risk maximum budgets can be reasonably estimated by cash managers. We also make publicly-available the Python code for the three solvers used in the numerical examples. This represents a good starting point for practitioners interested in either designing cash management decision support systems or performing their own experiments.

Finally, we further elaborate on the utility of cash management models by formalizing the problem from a multiobjective perspective in which we compare the loss derived from a given policy to the loss derived from a trivial policy. The analysis of the impact of any cash management model in cost-risk reductions is useful for either avoiding unnecessary efforts when estimated benefits are low, or realizing the potential when estimated benefits are high. This problem formulation depends on the particular measures for the set of goals considered. This fact opens a number of interesting future research lines aiming at establishing the particular conditions that must hold to ensure the utility of cash management models.

Appendix: Python for cash management

In an attempt to fill the gap between theory and practice in cash management and multiobjective decision-making, we next provide the link to a Python software library containing the three proposed MOCMP solvers. We used this library to perform the examples in Sections 5.1, 5.2 and 5.3:

`https://github.com/PacoSalas/Empowering-cash-managers-CP.git`

Acknowledgements

Work partially funded by projects 2014 SGR 118 and Collectiveware TIN2015-66863-C2-1-R (MINECO/ FEDER).

References

Abdelaziz, F. B., Aouni, B., and El Fayedh, R. (2007). Multi-objective stochastic programming for portfolio selection. *European Journal of Operational Research*, 177(3):1811–1823.

Aouni, B., Colapinto, C., and La Torre, D. (2014). Financial portfolio management through the goal programming model: Current state-of-the-art. *European Journal of Operational Research*, 234(2):536–545.

Artzner, P., Delbaen, F., Eber, J.-M., and Heath, D. (1999). Coherent measures of risk. *Mathematical finance*, 9(3):203–228.

Baccarin, S. (2002). Optimal impulse control for cash management with quadratic holding-penalty costs. *Decisions in Economics and Finance*, 25(1):19–32.

Baccarin, S. (2009). Optimal impulse control for a multidimensional cash management system with generalized cost functions. *European Journal of Operational Research*, 196(1):198–206.

Ballestero, E. (1998). Approximating the optimum portfolio for an investor with particular preferences. *Journal of the Operational Research Society*, 49:998–1000.

Ballestero, E. (2005). Mean-semivariance efficient frontier: A downside risk model for portfolio selection. *Applied Mathematical Finance*, 12(1):1–15.

Ballestero, E. (2007). Compromise programming: A utility-based linear-quadratic composite metric from the trade-off between achievement and balanced (non-corner) solutions. *European journal of operational research*, 182(3):1369–1382.

Ballestero, E. and Pla-Santamaria, D. (2004). Selecting portfolios for mutual funds. *Omega*, 32(5):385–394.

Ballestero, E. and Romero, C. (1998). *Multiple criteria decision making and its applications to economic problems*. Springer Science & Business Media.

Bar-Ilan, A., Perry, D., and Stadje, W. (2004). A generalized impulse control model of cash management. *Journal of Economic Dynamics and Control*, 28(6):1013–1033.

Baumol, W. J. (1952). The transactions demand for cash: An inventory theoretic approach. *The Quarterly Journal of Economics*, 66(4):545–556.

Ben-Tal, A., El Ghaoui, L., and Nemirovski, A. (2009). *Robust optimization*. Princeton University Press.

Bensoussan, A., Chutani, A., and Sethi, S. P. (2009). Optimal cash management under uncertainty. *Operations Research Letters*, 37(6):425–429.

Brealey, R. A. and Myers (2003). *Principles of corporate finance*. McGraw-Hill.

Chen, X. and Simchi-Levi, D. (2009). A new approach for the stochastic cash balance problem with fixed costs. *Probability in the Engineering and Informational Sciences*, 23(04):545–562.

Constantinides, G. M. and Richard, S. F. (1978). Existence of optimal simple policies for discounted-cost inventory and cash management in continuous time. *Operations Research*, 26(4):620–636.

da Costa Moraes, M. B. and Nagano, M. S. (2014). Evolutionary models in cash management policies with multiple assets. *Economic Modelling*, 39:1–7.

da Costa Moraes, M. B., Nagano, M. S., and Sobreiro, V. A. (2015). Stochastic cash flow management models: A literature review since the 1980s. In *Decision Models in Engineering and Management*, pages 11–28. Springer International Publishing.

Daellenbach, H. G. (1974). Are cash management optimization models worthwhile? *Journal of Financial and Quantitative Analysis*, 9(04):607–626.

Glasserman, P. (2003). Monte carlo methods in financial engineering.

Gormley, F. M. and Meade, N. (2007). The utility of cash flow forecasts in the management of corporate cash balances. *European Journal of Operational Research*, 182(2):923–935.

Gregory, G. (1976). Cash flow models: a review. *Omega*, 4(6):643–656.

Herrera-Cáceres, C. A. and Ibeas, A. (2016). Model predictive control of cash balance in a cash concentration and disbursements system. *Journal of the Franklin Institute*, 353(18):4885–4923.

Markowitz, H. (1952). Portfolio selection. *The journal of finance*, 7(1):77–91.

McNeil, A. J., Frey, R., and Embrechts, P. (2005). *Quantitative risk management: Concepts, techniques and tools*. Princeton university press.

Melo, M. A. and Bilich, F. (2013). Expectancy balance model for cash flow. *Journal of Economics and Finance*, 37(2):240–252.

Miller, M. H. and Orr, D. (1966). A model of the demand for money by firms. *The Quarterly journal of economics*, 80(3):413–435.

Pla-Santamaria, D. and Bravo, M. (2013). Portfolio optimization based on downside risk: a mean-semivariance efficient frontier from dow jones blue chips. *Annals of Operations Research*, 205(1):189–201.

Premachandra, I. (2004). A diffusion approximation model for managing cash in firms: An alternative approach to the miller–orr model. *European Journal of Operational Research*, 157(1):218–226.

Ringuest, J. L. (1992). *Multiobjective optimization: behavioral and computational considerations*. Springer Science & Business Media.

Rockafellar, R. T. and Uryasev, S. (2002). Conditional value-at-risk for general loss distributions. *Journal of Banking and Finance*, 26(7):1443–1471.

Salas-Molina, F., Martin, F. J., Rodríguez-Aguilar, J. A., Serrà, J., and Arcos, J. L. (2017). Empowering cash managers to achieve cost savings by improving predictive accuracy. *International Journal of Forecasting*, 33(2):403–415.

Salas-Molina, F., Pla-Santamaria, D., and Rodriguez-Aguilar, J. A. (2016). A multi-objective approach to the cash management problem. *Annals of Operations Research*, pages 1–15.

Soyster, A. L. (1973). Technical note—convex programming with set-inclusive constraints and applications to inexact linear programming. *Operations research*, 21(5):1154–1157.

Srinivasan, V. and Kim, Y. H. (1986). Deterministic cash flow management: state of the art and research directions. *Omega*, 14(2):145–166.

Weibull, W. et al. (1951). A statistical distribution function of wide applicability. *Journal of applied mechanics*, 18(3):293–297.

Yu, P.-L. (1985). *Multiple criteria decision making: concepts, techniques and extensions*. Plenum Press.

Zeleny, M. (1974). A concept of compromise solutions and the method of the displaced ideal. *Computers & Operations Research*, 1(3):479–496.

Zeleny, M. (1982). *Multiple criteria decision making*. McGraw-Hill.