

Document downloaded from:

<http://hdl.handle.net/10251/82639>

This paper must be cited as:

Salas-Molina, F.; Pla Santamaría, D.; Rodríguez Aguilar, JA. (2016). A multi-objective approach to the cash management problem. *Annals of Operations Research*. 1-15.
doi:10.1007/s10479-016-2359-1



The final publication is available at

<http://dx.doi.org/10.1007/s10479-016-2359-1>

Copyright Springer Verlag (Germany)

Additional Information

A multi-objective approach to the cash management problem

Francisco Salas-Molina ·
David Pla-Santamaria ·
Juan A. Rodriguez-Aguilar

Received: date / Accepted: date

Abstract Cash management is concerned with optimizing costs of short-term cash policies of a company. Different optimization models have been proposed in the literature whose focus has been only placed on a single objective, namely, on minimizing costs. However, cash managers may also be interested in risk associated to cash policies. In this paper, we propose a multi-objective cash management model based on compromise programming that allows cash managers to select the best policies, in terms of cost and risk, according to their risk preferences. The model is illustrated through several examples using real data from an industrial company, alternative cost scenarios and two different measures of risk. As a result, we provide cash managers with a new tool to allow them deciding on the level of risk to take in daily decision-making.

Keywords Cash management · multi-objective decision-making · risk preferences

1 Introduction

Cash flow management is concerned with the efficient use of a company's cash and short-term investments. Decision making in cash flow management focuses on keeping the balance between what the company holds in cash and what has been placed in short-term investments, such as deposit accounts or treasury bills. In other words, cash managers have to make daily decisions about the amount of

Francisco Salas-Molina
Hilaturas Ferre, S.A., Les Molines, 2, 03450 Banyeres de Mariola, Alicante, Spain
E-mail: francisco.salas@hifesa.com

David Pla-Santamaria
Escuela Politécnica Superior de Alcoy, Ferrándiz y Carbonell, s/n, 03801 Alcoy, Alicante, Spain
E-mail: dplasan@upv.es

Juan A. Rodriguez-Aguilar
IIIA-CSIC, Campus UAB, 08913 Cerdanyola, Catalonia, Spain
E-mail: jar@iiiia.csic.es

transactions between cash and any other kind of available liquid assets, resulting in an increase or decrease of the cash level.

The corporate cash management problem (CMP) usually handles two types of assets, namely, cash balance and short-term assets, such as marketable securities or interest bearing accounts. Controlling cash balances usually implies the use of a set of bounds or control limits. For example, a high bound, and a low bound that, when reached, trigger some transaction, i.e., a movement of cash, to restore the balance to a target level. The set of transactions scheduled over a future time horizon is called a policy. Finally, policies are determined by cash management by following a number of simple rules.

A number of cash management models have been proposed to control cash balances based on a set of levels or bounds. A comprehensive review of models, from the first proposals to the most recent contributions, can be found in Gregory (1976); Srinivasan and Kim (1986) and da Costa Moraes et al (2015). The CMP was first proposed from an inventory control point of view by Baumol (1952) in a deterministic way. Later on, Miller and Orr (1966) followed a stochastic approach assuming that cash balance changes are totally random. From these two seminal works many other models were developed, each one focusing on a particular dimension of the problem. While Girgis (1968) considered continuous net cash flows with both fixed and linear transaction costs, Eppen and Fama (1969) focused on discrete net cash flows with only variable transaction costs. The use of forecasts in the corporate cash management problem was first introduced by Stone (1972) as a way of smoothing cash flows. More recently, Gormley and Meade (2007) claimed the utility of cash flow forecasts in the cash management problem. However, the focus has been only placed on minimizing costs.

Since the CMP can also be viewed as a sequential decision-making problem with a set of cash balances, a probability density function of cash flows, a set of transactions and an objective function. This leads to a dynamic programming approach initially proposed by Eppen and Fama (1969) and Neave (1970) and more recently followed by Penttinen (1991); Chen and Simchi-Levi (2009); Melo and Bilich (2013). Finally, different strategies such as impulse control techniques in Baccarin (2009) and evolutionary models in da Costa Moraes and Nagano (2014) have been recently used to solve the CMP.

Notice that the particular characteristics of cash flows (or cash balances) are an important dimension of the CMP. In this sense, cash flows representations used in the literature range from deterministic (Baumol, 1952) to complete uncertainty (Miller and Orr, 1966). The most usual cash flow probability distribution in the literature is the Gaussian in the form of a Wiener process (Baccarin, 2009; Constantinides and Richard, 1978; Premachandra, 2004). In Penttinen (1991), a double exponential distribution was considered. However, empirical distributions obtained from real data sets are hardly used with the exception of Gormley and Meade (2007).

From the practitioner's point of view, several issues arise when dealing with real cash management problems. First, real cash flow data may not fit a theoretical probability distribution. This situation forces cash managers to find alternative ways of eliciting the best cash management policies, for instance, simulation. Second, the focus is typically placed on a single objective, namely, cost. With the exception of Zopounidis (1999), cash management and multi-criteria decision-making are not usually linked concepts in the financial literature. However, the

analysis of risk of the policies proposed by a cash management model is necessary if an overdraft is meant to be avoided. This kind of analysis is widely used in portfolio selection when high expected returns and low variances are desired objectives (Markowitz, 1952; Ballestero and Romero, 1998; Ballestero and Pla-Santamaria, 2004; Steuer et al, 2007). Furthermore, risk preferences are also an important issue for decision-makers. In this sense, an example of techniques for approximating the utility optimum when considering risk preferences can be found in Ballestero (1998). However, the design of cash management models that analyze both cost and risk and that consider the particular risk preferences of cash managers remains a rather unexplored problem.

Against this background, our main contribution is a novel multiobjective cash management model. More precisely, we argue that cost-risk analysis is a suitable management tool for obtaining efficient cash policies for the CMP. Our model is conceived to obtain policies that minimize both cost and risk incorporating the risk preferences of cash managers. Furthermore, it is able to handle empirical cash flow probability distributions. In addition to our model, we provide metrics and graphical tools to characterize cash policies in terms of cost and risk.

To illustrate our model, a real cash flow data set is used to select the best compromise policy in terms of cost and risk by extending the Miller and Orr's cash management model (Miller and Orr, 1966). We empirically evaluate our cash management model on three alternative scenarios with increasing costs and employing two different risk measures, namely, the standard deviation and the upper semi-deviation of daily cost. Our empirical results show that our model is robust to cost variations and that no difference is observed in the selection of the best compromise policies between the two measures of risk.

Then, in Section 2, the multi-objective cash management problem is defined and details of the proposed approach are given. This approach is later illustrated in Section 3 with a real example using the transaction rules proposed by the Miller and Orr's model, but relaxing most of their assumptions. In addition, a benchmark for alternative cost scenarios is performed and a graphical characterization of cash policies is provided. Finally, we discuss the results derived from a numerical case study.

2 Formalizing the multi-objective cash management problem

The main objective in managing cash is to keep the amount of available cash as low as possible while still keeping the company operating efficiently. In addition, companies may place idle cash in short-term investments (Ross et al, 2002). Then, the cash management problem can be viewed as a trade-off between holding and transaction costs. On the one hand, holding costs are usually opportunity costs due to idle cash that could be allocated in alternative investments. Holding too much cash is then inefficient but holding too little may produce high shortage costs. On the other hand, transaction costs are associated to the movement of cash from/into a cash account into/from any other short-term asset available, for example, treasury bills and other marketable securities. Summarizing, if a company tries to keep balances too low, holding cost will be reduced but undesirable situations of shortage will force to sell available marketable securities, hence increasing transaction costs. In contrast, if the balance is too high, low trading costs

will be produced due to unexpected cash flow but the company will carry high holding costs because no interest is earned on cash. Therefore, there is a target cash balance which the company must determine.

Consequently, the CMP can be defined as an optimization problem whose goal is to find the best sequence of transactions $X = \langle x_1, x_2, \dots, x_T \rangle \in \mathbb{R}$, what is called a policy, that minimizes a cost function $C(X, T)$ over a time horizon of T days. The CMP is characterized by its particular cost structure. Any ordering transaction into a cash account may have a cost, which may include a fixed part (γ_0^+) and a variable part (γ_1^+). On the other hand, a return transaction from a cash account may also have a cost with a fixed part (γ_0^-) and a variable part (γ_1^-). Furthermore, at the end of the day, a holding cost v per money unit is charged if a positive cash balance occurs, or a penalty cost u per money unit is charged if a negative cash balance occurs. According to this cost structure, a general daily cost, C_t , can be defined as:

$$C_t(x_t, b_t) = \Gamma(x_t) + L(b_t) \quad (1)$$

where $\Gamma(x_t)$ is a transfer cost function, x_t is the transfer made on day t , $L(b_t)$ is a holding/shortage cost function, and b_t is the cash balance at the end of day t , determined by the following continuity function:

$$b_t = b_{t-1} + x_t + f_t \quad (2)$$

being f_t the net cash flow at day t . We define the transfer cost function $\Gamma(x_t)$ as:

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

Additionally, we express the holding/shortage cost function as:

$$L(b_t) = \begin{cases} -u \cdot b_t & \text{if } b_t < 0; u > 0, \\ v \cdot b_t & \text{if } b_t > 0; v > 0. \end{cases} \quad (4)$$

Under the framework of the aforementioned cost functions, the CMP can be formally defined as a sequential decision-making problem where, given a data set of past cash flows $F = \langle f_{t-n}, \dots, f_{t-1} \rangle \in \mathbb{R}$ that determines an initial cash balance b_{t-1} , the goal is to find the policy $X^* = \langle x_t, x_{t+1}, \dots, x_{t+T} \rangle$ that minimizes the expected daily cost over the time horizon T :

$$X^* = \underset{X}{\operatorname{argmin}} E(C(X, T)) \quad (5)$$

where:

$$E(C(X, T)) = \frac{1}{T} \sum_{t=1}^T C_t(x_t, b_t) = \frac{1}{T} \sum_{t=1}^T [\Gamma(x_t) + L(b_t)] \quad (6)$$

subject to $X \in S$, where each element of X is the decision variable or the transaction made at day t over a time horizon of T days, and S is a set of feasible strategies defined by the constraints of the problem. Note that the net cash flow f_t over the time horizon T is unknown, but not its empirical probability distribution given by F .

In the CMP, not only cost but also risk deserves the attention of cash managers. On the one hand, managers can choose the risk that the business takes (Myers and Brealey, 2003). From that, risk measures are often used as management tool (McNeil et al, 2005). On the other hand, risk management allows to determine the cash buffer that companies need to hold for precautionary purposes (Ross et al, 2002). Intuitively, risk is associated to uncertainty, danger, chance of loss or damage. Quantitatively, risk is also linked to the randomness of costs. In the CMP, randomness is introduced by the particular variability of future cash flows that ultimately produces random costs.

The main performance measure in cash management is cost and the goal is to minimize it. But cost is a random variable characterized by its mean and its standard deviation. Consequently, cash managers are interested in minimizing both, the average daily cost as measure of cost and the standard deviation of cost as a measure of risk. The utility of introducing risk in the decision-making process is better understood through the following example. Consider two cash policies, X_1 and X_2 , producing the same average cost $\mu_1 = \mu_2 = 20$ money units but with different standard deviations, say, $\sigma_1 = 2$ and $\sigma_2 = 10$. Assuming normally distributed costs, a rational cash manager must select policy X_1 because of the lower volatility, because of the lower probability that the cost is above a given cost reference. We say, then, that X_2 is a riskier policy than X_1 .

Different approaches to measuring risk have been proposed in finance. Historically, variance has been the dominating risk measure in finance (McNeil et al, 2005) due to Mean-Variance portfolio selection model by Markowitz (1952). Moreover, its simplicity and ease of computation in experimental environments makes variance (and standard deviation) a good risk measure. By using an empirical statistic, no assumption on the underlying probability distribution is required. A common criticism to the use of variances points out that it makes no distinction between positive and negative deviations. Although this fact is only a problem when the cost distribution is asymmetric, it can be solved by using semi-variances or upside/downside deviations as in Ballesteros (2005); Pla-Santamaria and Bravo (2013).

In this work, we use the standard deviation (and the upper semi-deviation) to keep unaltered the units of cost and risk. Most of the applied cash management models in the literature consider risk incorporated into the opportunity cost and/or the shortage cost (Penttinen, 1991; Gormley and Meade, 2007; da Costa Moraes and Nagano, 2014). In this work, we consider these traditional measures of risk but we also add additional criteria such as the standard deviation of the cost associated with cash or the upper semi-deviation. In this way, we reflect other risk definitions that could be relevant to a broader population of decision makers. These supplemental risk measures are cornerstone in finance, but surprisingly, its use is not widespread in cash management. Some relevant works that consider risk as volatility on the topic are Whalen (1966) where the demand of cash is analyzed, Gao et al (2013) when defining the determinants of corporate cash policy and Pinkowitz et al (2016) or Bates et al (2009) that take into account the cash flow volatility when working with industry cash holdings.

In cash management, cost and risk are often desired but conflicting objectives. Cost reductions are achieved by reducing cash balances but, at the same time, the risk of an overdraft is increased. From that, we define the multi-objective cash management problem (MOCMP) as a sequential decision-making problem where,

given a set of past cash flows that determines an initial cash balance, the goal is to find the best compromise policy in terms of cost and risk over the considered time horizon:

$$X^* = \underset{X}{\operatorname{argmin}} \quad [E(C(X, T)), R(C(X, T))] \quad (7)$$

subject to $X \in S$, where $R(C(X, T))$ denotes a general risk function that can be evaluated, for example, as the standard deviation of $C(X, T)$:

$$R(C(X, T)) = \left(\frac{1}{T} \sum_{t=1}^T (C_t - E(C))^2 \right)^{1/2}, \quad (8)$$

where the arguments of functions $C(X, T)$ and $C_t(x_t, b_t)$ are omitted on the right-hand side of the equation for ease of notation, or as the upper semi-deviation of cost:

$$R(C(X, T)) = \left(E \left[\max\{C_t - E(C), 0\}^2 \right] \right)^{1/2}. \quad (9)$$

Summarizing, the MOCMP presents the following attributes:

1. It is a sequential decision-making problem with a set of cash balances, a probability density function of cash flows (theoretical or empirically derived from past cash flows), a set of transactions and a multi-objective function of cost and risk.
2. Cost and risk are desired but conflicting objectives that need to be minimized according to the risk preferences of the cash manager.
3. The constraints of the problem are defined by the particular characteristics of the cash management model used.
4. It is a stochastic problem where the user is able to learn from experience.

Based on the previous characteristics, we propose a general procedure to solve the MOCMP by following the next steps:

1. Define the objective functions for cost and risk.
2. Obtain a set of feasible policies evaluated in terms of cost and a risk.
3. Derive an efficient frontier in the cost-risk (C-R) space, i.e., a Pareto efficient set of points where minimum risk can be obtained for any given cost value.
4. Approximate the best compromise policy X^* according to the risk preferences of the cash manager.

Firstly, a definition of the objective function is required. A common and intuitive way to optimize a multi-objective function when rewards can be easily translated into monetary value is the use of a linear scalarization function (Roijers et al, ???). This function may be the weighted sum of the values for each objective (in our case, cost and risk). However, we consider these weights to be unknown in the learning phase. Later, when the selection phase occurs, a policy is chosen among those computed in the learning phase according to the risk preferences of the cash manager. Motivation for this assumption is double: first, we want to avoid the computational burden of minimizing cost and risk for all possible combinations of risk preferences or weights; second, we consider easier and natural a decision-making process where alternative policies are presented before selection, allowing what-if analysis.

Therefore, a compromise programming (CP) approach (Zeleny, 1982; Ballester and Romero, 1998; Yu, 2013) is applied to optimize the general minimization multi-objective function defined in equation (7). This CP approach is based on the concept of ideal point and on the Zeleny's axiom of choice (Zeleny, 1982). The ideal point in the MOCMP is the point where both minimum cost and risk occur simultaneously. This point is clearly unfeasible and it is necessary to look for compromise solutions that can be evaluated, for example, computing the distance to the ideal point.

Once the objective function is defined, we obtain a set of available policies evaluated in terms of cost and the standard deviation of the daily cost as a measure of risk. At this point, we are in a position to derive an efficient frontier combining cost and risk with the set of daily policies X , that are not dominated by any other policy. Indeed, in this context, all policies are dominated by the policy with the lowest risk from the cost perspective. Likewise, from the risk perspective, all policies are dominated by the policy with the lowest cost. Graphically, the efficient frontier in the C-R space, is built from a set of two-dimensional points, $(C(X, T), R(X, T))$. Here, $C(X, T)$ is the expected daily cost and $R(X, T)$ is the risk obtained using policy X over T days.

In the next step, cash managers have to choose a policy from the efficient frontier according to their risk/cost preferences. This is attained through the use of two normalized indexes (Ballester and Romero, 1998). First, a cost index θ_1 , defined as:

$$\theta_1(X, T) = \frac{C(X, T) - C_{min}}{C_{max} - C_{min}} \quad (10)$$

where $C(X, T)$ is the expected daily cost over T computed for a particular policy X using equation (6), and C_{max} and C_{min} are, respectively, the maximum and minimum daily cost obtained from the set of values that forms the efficient frontier. Second, a risk index, θ_2 , defined as:

$$\theta_2(X, T) = \frac{R(X, T) - R_{min}}{R_{max} - R_{min}} \quad (11)$$

where $R(X, T)$ is a risk measure computed for a particular policy X over T using one of the previously suggested risk functions, and R_{max} and R_{min} are, respectively, the maximum and minimum values of R obtained from the set of points under consideration. Note that θ_1 and θ_2 range in the interval $[0,1]$, and that the ideal point $(\theta_1, \theta_2) = (0, 0)$, is clearly unfeasible. However, the closer to this ideal point, the better the solution.

From the efficient frontier, a number of useful metrics can be derived to select the best policies according to the particular risk preferences of cash managers. In what follows, a real case study illustrating these concepts is developed.

3 An illustrative example using the Miller and Orr's model

In this section, a real case of our multi-objective approach to the CMP using the model proposed by Miller and Orr (1966) is presented using three alternative cost scenarios and two different risk measures, namely, the standard deviation and the upper semi-deviation of daily cost as defined above. This real case uses cash flow data from an industrial Spanish company. The section is divided in five

parts: (i) we introduce some background about the Miller-Orr's model; (ii) we derive a normalized cost-risk efficient frontier of cash policies; (iii) risk preferences of are considered to select the best compromise policy; (iv) this best policy will be compared to other baseline policies; (v) two useful metrics are proposed to characterize the cost-risk efficient frontier.

3.1 The Miller-Orr's cash management model

The first approach to the CMP considering stochastic cash flows, was made by Miller and Orr (1966) where cash flows are characterized as a sequence of independent and symmetric Bernoulli trials. They supposed that cash balance will either increase or decrease by m dollars with probability $p = 1/2$. The main assumptions of this approach are independence, stationarity, zero-drift, and the absence of regular swings in daily cash flow. Moreover they ignored shortage and variable transaction costs.

This model sets its policy based on a low bound¹, l , and a high bound, h , and the cash balance is allowed to wander around between these bounds. When h is reached a withdrawal transfer will be made to restore the balance to a target level of z . In the same way, when the cash balance reaches l , a positive transfer will be made to again restore the balance to z . Accordingly, the transfer occurring at time t , x_t , is elicited by comparing the cash balance at the beginning of the day, b_{t-1} , to the low and high bounds:

$$x_t = \begin{cases} z - b_{t-1}, & \text{if } b_{t-1} > h \\ 0, & \text{if } l < b_{t-1} < h \\ z - b_{t-1}, & \text{if } b_{t-1} < l \end{cases} \quad (12)$$

By relaxing most of the Miller and Orr's assumptions, below we use their model in order to allow: first, the use of real world cash flow probability distributions; second, the introduction of variable cost of transactions and shortage costs; and third, the selection of the best policy determined by the set of control limits $\{h^*, z^*, l^*\}$ that minimizes cost and risk, according to the particular risk preferences of cash managers. Note that this problem represents a particular case of MOCMP formalized in Section 2 where the policy X is determined by the control limits $\{h \geq z \geq l\}$ according to equation (12).

The objective is determining the best cash policies based on the minimization of cost and risk. Here, instead of directly solving the general MOCMP as defined in Section 2 by obtaining the set X of daily transfers, we introduce the Miller and Orr's cash management model in the problem. Consequently, we help cash managers to learn the set of control limits $\{h, z, l\}$ that will ultimately determine the sequence X of daily transfers according to equation (12).

In our case study, we use a data set $F \in \mathcal{R}$ from a Spanish industrial company composed by 1000 observations of cash flows at non-bank-holidays for a period of about 4 years. Note that cash balances are readily derived from an initial cash balance adding the sum of cash flows up to the day we are interested in. To illustrate

¹ Miller and Orr initially obtained optimal values setting the low bound to zero. However, setting a low bound distinct from zero for precautionary purposes was given by Ross et al (2002).

our model, we make computations using three different cost scenarios as shown in Table 1. Firstly, we consider a Low, Medium and High cost scenario with increasing costs selecting values along the lines of those suggested in da Costa Moraes and Nagano (2014). Secondly, we consider that the into-account transaction costs are higher than the from-account ones since costs of obtaining funds are usually higher. And finally, it is important to note that the from-account transaction cost is lower than the holding cost ($\gamma_1^- < v$), and that into-account transaction cost is lower than the shortage cost ($\gamma_1^+ < u$). Otherwise, no transaction would be made since the unitary costs of transferring money are higher than those of holding the same amount of money as shown in Constantinides and Richard (1978).

Table 1 Three alternative scenarios with increasing costs including holding/shortage and transaction costs.

Cost	Alternative cost scenarios		
	Low	Medium	High
Holding cost v	0.02 %	0.05 %	0.1 %
Shortage cost u	10 %	20 %	30 %
Fixed into account γ_0^+	5 €	50 €	200 €
Fixed from account γ_0^-	2 €	20 €	80 €
Variable into account γ_1^+	0.2 %	0.5 %	1 %
Variable from account γ_1^-	0.002 %	0.005 %	0.01 %

3.2 Deriving the cost-risk efficient frontier

In order to derive the efficient frontier in the cost-risk space, we develop a complete grid search over different policies derived from the Miller and Orr's model. These policies are determined by the control limits $\{h, z, l\}$. Since our cost structure presents a high shortage cost, we expect better policies to be given by high values of control limits. More precisely, we iterate over 57400 feasible combinations of $\{h, z, l\}$ where $h \geq z \geq l$, and the respective ranges are $[h - \sigma, h + 3\sigma]$, $[z - \sigma, z + 3\sigma]$, $[l - \sigma, l + 3\sigma]$, in steps of 10000 €, where σ is the standard deviation of cash flows in the data set. Since the number of possible combinations of policies is not high enough to produce unaffordable computational times, we follow a complete search process. Moreover, grid search allows a random simulation of policies and the selection of the best set without setting any constraint. As long as the search is deep enough, meaning a thorough exploration of alternatives, this selection will be closer to the optimal policies.

A visual representation of the efficient frontier derived from the results obtained for the Medium cost scenario using the standard deviation as a measure of risk is shown in Figure 1. Under the CP framework, the closer to the ideal point $(0, 0)$, the better the policy. A sample of the best policies of the efficient frontier from Figure 1 is summarized in Table 2, where Manhattan distances to the ideal point are computed for each policy. From that, an average cash manager (without defined preferences for risk or cost) would select the policy denoted with id 10 because the distance to the ideal point is minimum.

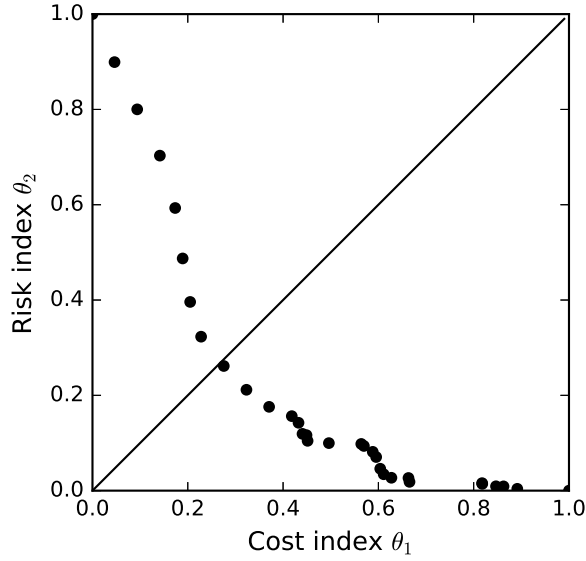


Fig. 1 Efficient frontier in the cost-risk (θ_1, θ_2) for the Medium cost scenario using the standard deviation as a measure of risk

3.3 Selecting policies considering risk preferences

In the context of the CMP, risk preferences of the cash managers are worth considering. Recall from Section 2, that the goal is minimizing the multi-objective function in equation (7). In our illustrative example, maximizing utility amounts to minimizing a weighted loss function $\mathcal{L}(\theta_1, \theta_2)$:

$$\mathcal{L}(\theta_1, \theta_2) = w_1\theta_1 + w_2\theta_2 \quad (13)$$

where weights, w_1 and w_2 , are used to introduce the cash manager's risk preferences in the problem. For instance, when $2w_1 = w_2$, cost is considered half as important as risk, i.e., risk is twice more important than cost. This is an example of a conservative cash manager. On the other hand, when $w_1 = 2w_2$, cost is considered twice more important than risk, i.e., risk is half as important as cost, as it may be the case for a risky cash manager.

According to Ballesteros and Romero (1998), risk preferences can be considered by introducing a parameter, $r_0 \in \mathbb{R}^+$. Translated into the cash management problem, r_0 indicates 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 $r_0 = 0.5$, our 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 $w_2 = 0.67$. Moreover, in an attempt to bound the location of the best policy in the normalized cost-risk space when considering risk preferences, it is proven in Ballesteros (1998) that the utility optimum lies on the normalized efficient frontier between points L and L_∞ , which are defined as follows:

Table 2 Central sample of the efficient frontier from Figure 1 (Id=policy identifier; h =high control limit; z =target balance; l =low control limit; Cost=average daily cost in €; Risk=standard deviation of daily cost in €; θ_1 =cost index; θ_2 =risk index; $\theta_1 + \theta_2$ = Manhattan distance to the ideal point $(\theta_1, \theta_2) = (0, 0)$; and we use $r_0 = 0.5$).

Id	h	z	l	Cost	Risk	θ_1	θ_2	$\theta_1 + \theta_2$	$r_0\theta_1 + \theta_2$
7	720000	570000	270000	358	528	0,205	0,396	0,601	0,498
8	730000	580000	280000	360	486	0,228	0,323	0,551	0,437
9	740000	590000	290000	363	451	0,276	0,262	0,537	0,399
10	750000	600000	300000	366	423	0,323	0,212	0,535	0,373
11	760000	610000	310000	369	402	0,371	0,176	0,547	0,361
12	770000	620000	320000	372	391	0,418	0,156	0,575	0,365
13	680000	490000	360000	373	383	0,432	0,142	0,575	0,359
14	720000	490000	360000	373	370	0,441	0,119	0,560	0,340
15	730000	490000	360000	374	368	0,449	0,116	0,566	0,341
16	750000	490000	360000	374	362	0,451	0,105	0,556	0,330
17	760000	500000	370000	377	359	0,496	0,100	0,596	0,348
18	800000	490000	370000	381	358	0,564	0,098	0,662	0,380

- L is the point minimizing the linear loss function $r_0\theta_1 + \theta_2$ on the normalized efficient frontier.
- L_∞ is the intersection of $\theta_1 = \theta_2$ with the efficient frontier.

Bounds L and L_∞ , referring here to cash policy performance when less is better, delimit the compromise set that can be taken as a surrogate for approximating the best policies for cash managers with different risk preferences defined by the parameter r_0 . According to this procedure, for a conservative cash manager with $r_0 = 0.5$, from Table 2 we obtain that bound $L = 16$ when $r_0\theta_1 + \theta_2$ is minimum, and that bound $L_\infty = 9$ when $\theta_1 \approx \theta_2$.

3.4 Benchmarking selected policies and risk measures

In this section, we compare in Table 3 the best compromise policy ($Id = 10$) with three policies in terms of cost and risk for a cash manager without particular preferences for cost and risk. Four different policies were considered: (i) a No-Trans policy with no cash movement which is also used as a baseline for comparison purposes; (ii) a policy derived from a set of bounds obtained from the values proposed by Miller and Orr (1966) with low bound, $l = 0$; and (iii) a policy derived from a Miller-Orr set of bounds but setting a low bound, $l = 2 \cdot \sigma_F$, for precautionary purposes as recommended by Ross et al (2002), with σ_F set to the standard deviation of our real cash flow data set F ; (iv) the best compromise policy derived from our multi-objective approach. The comparison yields that the Miller-Orr policy with $l = 0$ produced the highest cost-risk due to the absence of a low bound. Setting a low bound allowed cash managers to improve the performance of the baseline: for Medium and High costs when using the standard deviation and only for High costs when using the upper semi-deviation. Clearly, the highest cost-risk reductions were achieved using the best compromise policies from our multi-objective approach.

Moreover, no difference was found in the selection of best policies when using the standard deviation or the upper semi-deviation as a measure of risk. Since the

Table 3 Relative cost-risk performance of policies for alternative cost scenarios using two measures of risk. The lower the percentage the better.

Policy	Standard deviation			Upper semi-deviation		
	Low	Medium	High	Low	Medium	High
No-Trans policy	100	100	100	100	100	100
Miller-Orr policy with $l = 0$	541	359	229	851	572	367
Miller-Orr policy with $l = 2 \cdot \sigma_F$	161	85	51	264	138	80
Best compromise policy	21	22	23	29	30	31

measure definition is different, the final cost-risk value is also different. However, the best compromise policies selected were the same in both cases due to the small differences observed between the standard deviation and the upper semi-deviation of daily cost. From that, we hypothesize a certain degree of symmetry in the cost random variable, at least enough to produce the same selection of the best compromise policy.

3.5 Characterizing the cost-risk efficient frontier

From the analysis of the efficient frontier, a number of insights might be derived. Firstly, in Figure 1 we observe significant discontinuities. This fact can be quantified by computing the elasticity of the efficient frontier proposed by Ballestero and Pla-Santamaria (2004). Risk elasticity measures the percentage change in risk following a 1% change in cost for two consecutive policies from the efficient frontier. Negative values indicate that risk decreases as cost increases. This amount indicates the magnitude of the risk sensitivity to changes in cost. For instance, in our example, policies 15 and 16 are equivalent in terms of cost but not as much in terms of risk as shown in Table 2. For the same cost, policy 15 is much riskier than policy 16. Thus, the latter should be chosen. A visualization of the elasticity of the efficient frontier for the Medium cost scenario and using the standard deviation as a measure of risk is shown in Figure 2. In this figure, risk discontinuities from the efficient frontier can be easily identified. From the cost perspective, changes in cost per unit of risk can be similarly computed.

Moreover, we propose the use of a Sharpe-like ratio (SLR) that can be easily computed for each policy as synthetic and quick performance measure. The Sharpe ratio (Sharpe, 1966, 1994) was introduced to measure the performance of mutual funds as a reward-to-variability ratio. Since high reward and low variability are desired objectives, the higher the Sharpe ratio the better. Its utility is based on the combination of two goals in a single figure. Likewise, here we import this concept and propose a Sharpe-like ratio that synthetically measures cost and risk in one single figure. We define the SLR as the geometric average of both cost and risk of each policy:

$$SLR(X) = \sqrt{\theta_1 \cdot \theta_2}. \quad (14)$$

where θ_1 and θ_2 are normalized indexes. Normalization is necessary to avoid bias towards attributes. In this case, normalization is done by dividing cost and risk by their respective averages to avoid the problem of close-to-zero values in the extremes of the efficient frontier. For comparison purposes, the lower the SLR the better the policy.

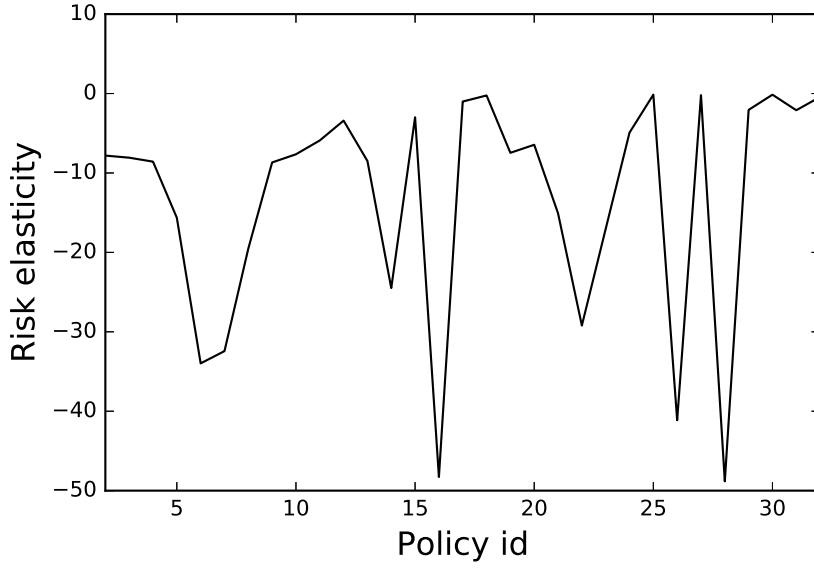


Fig. 2 Elasticity of the efficient frontier

By combining cost, measured by the average, and risk, measured by the standard deviation, in the SLR, cash managers are able to estimate the global performance of the policy in just one single value. For instance, consider two policies X_1 and X_2 , characterized by mean values 20 and 25, and standard deviations 4 and 2 respectively. Even though policy X_1 presents a lower cost, a quick and simple computation shows that policy X_2 is a better choice because $SLR(X_1) > SLR(X_2)$. From that, when comparing a reduced number of policies, the SLR provides a rapid performance estimator, likewise the Sharpe ratio does when comparing, for example, two mutual funds.

Furthermore, we propose a more refined version of the SLR to incorporate risk preferences of cash managers by means of a set of weights. In general, given the cash manager's cost and risk preferences expressed as weights w_1 and w_2 , such that $0 \leq w_1, w_2 \leq 1$, with $w_1 + w_2 = 1$, we define the Weighted Sharpe Like Ratio (WSLR) of policy X as:

$$WSLR(X) = [\theta_1^{w_1} \cdot \theta_2^{w_2}]^{1/(w_1+w_2)}. \quad (15)$$

Note that when $w_1 = w_2$, the cash manager has no preference and expressions (14) and (15) are equivalent. Figure 3 shows the WSLR for the Medium cost scenario and using the standard deviation for different pair of weights. Recall that the higher the value for w_1 the higher the preference for risk of the cash manager. Our case study is characterized by a clear decreasing trend up to policy 10, from which a horizontal evolution, in terms of the WSLR value, is observed regarding the rest of policies. This fact indicates that the policies ranging from Id 10 to 32 exhibit a stable behavior, in terms of the WSLR value, when varying a cash manager's preferences. A closer look at the particular policies shows that this

reduced average risk is associated to policies with higher control limits, producing higher cash balances but reducing the probability of an overdraft.

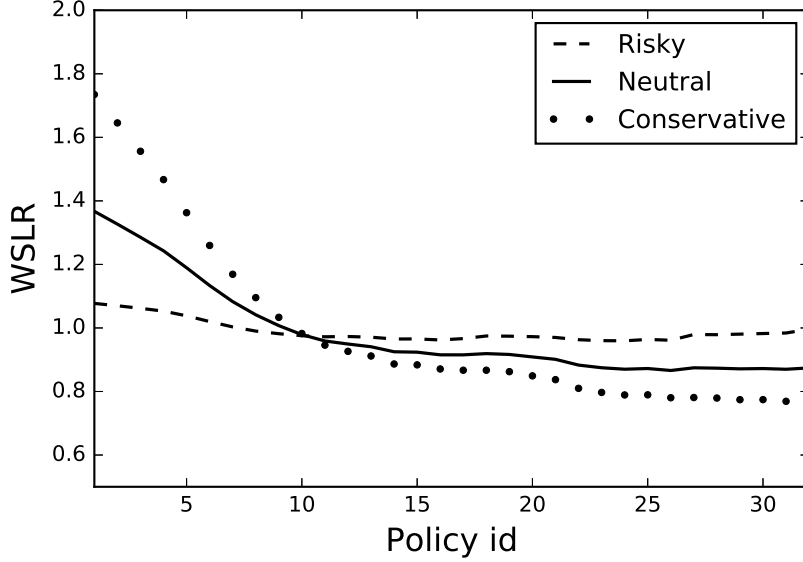


Fig. 3 Weighted Sharpe-like ratio for different risk preferences and each of the policies. In the Risky case, $w_1 = 0.8, w_2 = 0.2$, in Neutral, $w_1 = 0.5, w_2 = 0.5$, and in Conservative, $w_1 = 0.2, w_2 = 0.8$

An interesting property can be derived from the analysis of the WSLR plot for different weights as shown in Figure 3. The existence of policies, around policy 10, where the WSLR value is almost equal for any particular risk preference of the cash manager. Such a preference-neutral policy is characterized by similar values of θ_1 and θ_2 and is directly linked to the efficient frontier and the increasing diagonal in the normalized cost-risk space presented in Figure 1. There exists a line-to-point correspondence between the diagonal in the normalized cost-risk space and the preference-neutral policy. Policies in this diagonal are insensitive to the risk preferences of the cash manager. Hence, the following proposition characterizes and links preference-neutral policies to the normalized cost-risk space.

Definition 1 A policy X characterized by normalized cost and risk indexes θ_1 and θ_2 , is preference-neutral to risk preferences w_1 and w_2 with respect to the utility function \mathcal{L} , iff for every $w_1, w_2 \in [0, 1]$ such that $w_1 + w_2 = 1$, we have that $\mathcal{L}(\theta_1, \theta_2, w_1, w_2) = k$ for some $k \in \mathbb{R}$, i.e., the utility function is constant.

Proposition 1 A policy X whose normalized cost and risk indexes θ_1 and θ_2 satisfy that $\theta_1 = \theta_2$ is preference-neutral to the arithmetic, geometric and harmonic weighted average utility functions and its utility is θ_1 .

Proof Since $\theta_1 = \theta_2$ and $w_1 + w_2 = 1$, in the case of the arithmetic mean:

$$w_1 \cdot \theta_1 + w_2 \cdot \theta_2 = w_1 \cdot \theta_1 + (1 - w_1) \cdot \theta_1 = w_1 \cdot \theta_1 + \theta_1 - w_1 \cdot \theta_1 = \theta_1. \quad (16)$$

For the geometric mean, such as the WSLR presented above:

$$[\theta_1^{w_1} \cdot \theta_1^{w_2}]^{1/(w_1+w_2)} = [\theta_1^{w_1} \cdot \theta_1^{1-w_1}] = \theta_1^{w_1+1-w_1} = \theta_1 \quad (17)$$

and for the harmonic mean:

$$\frac{w_1 + w_2}{\frac{w_1}{\theta_1} + \frac{w_2}{\theta_2}} = \frac{1}{\frac{w_1}{\theta_1} + \frac{1-w_1}{\theta_1}} = \frac{\theta_1^2}{w_1 \cdot \theta_1 + \theta_1 - w_1 \cdot \theta_1} = \frac{\theta_1^2}{\theta_1} = \theta_1 \quad (18)$$

which is constant in all three cases and independent of the risk preferences w_1 and w_2 .

A final useful remark on the elasticity and Sharpe-like ratio measures is appropriate. Not only they provide a useful characterization of an efficient frontier for different risk preferences, but they also allow comparing different cash management models and different cost structures. To this end, cash managers can plot the elasticity and WSLR curves for each model under comparison.

4 Conclusions

In the cash management problem, cash managers must control the amount of risk that their company take. Here, we propose a new method that differs from previous cash management approaches by considering risk and cash managers' risk preferences. Although different approaches to measuring risk have been proposed in the literature, standard deviation has been a dominating risk measure in finance. Along this direction, we propose a multi-objective model to minimize cost and risk in cash management that can employ either the standard deviation or the upper semi-deviation of cost as a measure of risk.

By relying on compromise programming, the best set of non-dominated policies in terms of cost and risk is presented to cash managers to be selected according to their particular risk preferences. Moreover, several examples using a real cash flow data set are given for three alternative cost scenarios and two different measures of risk, namely, the standard deviation and the upper semi-deviation of daily cost. Our empirical results show that the policies derived from our cash management model outperformed three different baseline policies: a No-trans policy, a Miller-Orr policy with low level set to zero, and a Miller-Orr policy with low level set to two times the standard deviation of historical cash flows. It is important to highlight that our cash management model produced the very same best compromise policies for the two risk measures that we considered, the standard deviation and the upper semi-deviation of daily cost.

Finally, we also provide cash managers with further insights in the evaluation of cash policies through elasticity plots and a weighted Sharpe-like ratio. More precisely, elasticity plots point out the magnitude of risk changes, which are not captured in the normalized cost-risk space. On the other hand, the Sharpe-like offers a rapid performance estimator to compare a reduced number of policies, similarly to the way the performance of mutual funds is compared. An interesting property of this ratio is the preference-neutrality of policies with the same cost and risk indexes.

Cash managers have now new management tools to control the amount of risk they take in their decision-making processes. Both less uncertain and more stable policies can be selected by using the proposed expected performance analysis. In this sense, further research is required to consider more complex cost functions and to explore the utility of forecasts to further reduce the uncertainty associated to cash management.

References

- 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
- Ballesteros E (1998) Approximating the optimum portfolio for an investor with particular preferences. *Journal of the Operational Research Society* 49(9):998–1000
- Ballesteros E (2005) Mean-semivariance efficient frontier: A downside risk model for portfolio selection. *Applied Mathematical Finance* 12(1):1–15
- Ballesteros E, Pla-Santamaria D (2004) Selecting portfolios for mutual funds. *Omega* 32(5):385–394
- Ballesteros E, Romero C (1998) Multiple criteria decision making and its applications to economic problems. Springer Science & Business Media
- Bates TW, Kahle KM, Stulz RM (2009) Why do US firms hold so much more cash than they used to? *The journal of finance* 64(5):1985–2021
- Baumol WJ (1952) The transactions demand for cash: An inventory theoretic approach. *The Quarterly Journal of Economics* 66(4):545–556
- Chen X, 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 GM, Richard SF (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 MB, Nagano MS (2014) Evolutionary models in cash management policies with multiple assets. *Economic Modelling* 39:1–7
- da Costa Moraes MB, Nagano MS, Sobreiro VA (2015) Stochastic cash flow management models: A literature review since the 1980s. In: *Decision Models in Engineering and Management*, Springer International Publishing, pp 11–28
- Eppen GD, Fama EF (1969) Cash balance and simple dynamic portfolio problems with proportional costs. *International Economic Review* 10(2):119–133
- Gao H, Harford J, Li K (2013) Determinants of corporate cash policy: Insights from private firms. *Journal of Financial Economics* 109(3):623–639
- Girgis NM (1968) Optimal cash balance levels. *Management Science* 15(3):130–140
- Gormley FM, 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
- Markowitz H (1952) Portfolio selection. *The journal of finance* 7(1):77–91
- McNeil AJ, Frey R, Embrechts P (2005) *Quantitative risk management: Concepts, techniques and tools*. Princeton University Press

- Melo MA, Bilich F (2013) Expectancy balance model for cash flow. *Journal of Economics and Finance* 37(2):240–252
- Miller MH, Orr D (1966) A model of the demand for money by firms. *The Quarterly Journal of Economics* 80(3):413–435
- Myers SC, Brealey RA (2003) *Principles of corporate finance*, seventh edn. McGraw-Hill
- Neave EH (1970) The stochastic cash balance problem with fixed costs for increases and decreases. *Management Science* 16(7):472–490
- Penttinen MJ (1991) Myopic and stationary solutions for stochastic cash balance problems. *European Journal of Operational Research* 52(2):155–166
- Pinkowitz L, Stulz RM, Williamson R (2016) Do US firms hold more cash than foreign firms do? *Review of Financial Studies* 29(2):309–348
- Pla-Santamaria D, 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
- Roijers DM, Vamplew P, Whiteson S, Dazeley R (????) A survey of multi-objective sequential decision-making. *Journal of Artificial Intelligence Research* 48(1):67–113
- Ross SA, Westerfield R, Jordan BD (2002) *Fundamentals of corporate finance*, sixth edn. McGraw-Hill
- Sharpe WF (1966) Mutual fund performance. *Journal of business* 39(1):119–138
- Sharpe WF (1994) The sharpe ratio. *The Journal of portfolio management* 21(1):49–58
- Srinivasan V, Kim YH (1986) Deterministic cash flow management: state of the art and research directions. *Omega* 14(2):145–166
- Steuer RE, Qi Y, Hirschberger M (2007) Suitable-portfolio investors, nondominated frontier sensitivity, and the effect of multiple objectives on standard portfolio selection. *Annals of Operations Research* 152(1):297–317
- Stone BK (1972) The use of forecasts and smoothing in control-limit models for cash management. *Financial Management* 1(1):72–84
- Whalen EL (1966) A rationalization of the precautionary demand for cash. *The Quarterly Journal of Economics* 80(2):314–324
- Yu PL (2013) *Multiple-criteria decision making: concepts, techniques, and extensions*. Springer Science & Business Media
- Zeleny M (1982) *Multiple criteria decision making*. McGraw-Hill
- Zopounidis C (1999) Multicriteria decision aid in financial management. *European Journal of Operational Research* 119(2):404–415