

Selecting cash management models from a multiobjective perspective

Francisco Salas-Molina ·
Juan A. Rodríguez-Aguilar ·
Pablo Díaz-García

Received: 13 Oct 2016 / Accepted: 23 Aug 2017

Abstract This paper addresses the problem of selecting cash management models under different operating conditions from a multiobjective perspective considering not only cost but also risk. A number of models have been proposed to optimize corporate cash management policies. The impact on model performance of different operating conditions becomes an important issue. Here, we provide a range of visual and quantitative tools imported from Receiver Operating Characteristic (ROC) analysis. More precisely, we show the utility of ROC analysis from a triple perspective as a tool for: (i) showing model performance; (ii) choosing models; and (iii) assessing the impact of operating conditions on model performance. We illustrate the selection of cash management models by means of a numerical example.

Keywords Cash management models · ROC analysis · multiobjective · operating condition.

1 Introduction

A Receiver Operating Characteristic (ROC) curve is a useful graphical tool for classification models that illustrates its performance Metz (1978); Bradley (1997); Flach (2003); Fawcett (2006). One of the basic ideas behind ROC analysis is the study of the impact of operating conditions on model performance. Traditional ROC analysis was initially restricted to classification. However, a broader perspective was adopted in Drummond and Holte (2000, 2006); Elkan (2001); Hernández-

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

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

Pablo Díaz-García
Universidad Politécnica Valencia, Ferrándiz y Carbonell, s/n, 03801 Alcoy, Alicante, Spain
E-mail: pdiazga@txp.upv.es

Orallo et al. (2013), by following a cost-sensitive approach to analyze the economic implications of the operating conditions by means of ROC cost curves. In a particularly interesting attempt to cross borders, an extension of ROC analysis to regression models was proposed by Bi and Bennett (2003); Torgo (2005); Hernández-Orallo (2013).

On the other hand, cash flow management focuses on finding the balance between cash holdings and short-term investments. Cash managers make daily decisions about the firm optimal cash level for operational and precautionary purposes (Ross et al., 2002). In order to deal with this problem, 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); da Costa Moraes et al. (2015).

Multiple-criteria decision-making in economics and finance has rapidly evolved during the last decades (Ballesteros and Romero, 1998; Doumpos and Zopounidis, 2007; Steuer et al., 2007; Garcia-Bernabeu et al., 2016). Cost and risk are desired but conflicting objectives in cash management as suggested by Salas-Molina et al. (2016). Assuming that idle balances have an opportunity cost, a decrease in cash balances results in cost reductions but, at the same time, the risk of an overdraft penalty is increased. Holding and shortage costs are usually asymmetric, e.g., 0.01% versus 30%, in daily rate of the final cash balance. Then, without considering any other costs, having a zero balance would be optimal in terms of cost but suboptimal in terms of risk, since eventual negative cash balances may produce unexpected penalty costs.

The tradeoff between cost and risk for different models and operating conditions is at the core of decision support systems for cash managers. However, little work has been done in the comparison and selection of models from a multiobjective perspective. In this work, we argue that the evaluation of the total loss in a period of time, in terms of cost and risk, for any particular operating condition is a necessary management tool. Moreover, by adopting a double perspective, quantitative and graphical, the understanding of cash management models and their performance is enhanced. To this end, ROC analysis seems to be a promising tool also for cash management. In this paper, we rely on ROC analysis to achieve further insight on cash management through: (i) a graphical tool; (ii) a set of useful metrics; (iii) performance evaluation; and (iv) for a number of different operating conditions.

The main point of this paper is to show that graphical tools derived from ROC analysis can enhance the understanding and the decision-making process of cash managers in their purpose of selecting the best cash management model. Since there is no such a model that is the best for any context, this selection also implies the answer to the question of under what circumstances or operating conditions a model is better than another. From that, we focus on the following goals under a multiobjective viewpoint:

1. Showing models in the Cost-Risk space.
2. Choosing models according to risk preferences.
3. Deriving loss curves for different operating conditions.

This paper is organized around the objectives above. In Section 2, we provide some useful background about the model proposed by Stone (1972) that we later

use in the numerical examples and a recent formulation of the multiobjective cash management problem. Section 3 introduces a graphical representation to visualize the performance of cash management models in a bidimensional Cost-Risk space. Section 4 proposes techniques for choosing models in the Cost-Risk space. Section 5 presents a procedure to draw loss curves for different operating conditions. We illustrate the previous techniques and procedures with different examples. Finally, several concluding remarks will be given in Section 6.

2 Background

In this section, we provide useful background on the cash management problem (CMP) and a recent multiobjective formulation of the problem described in Salas-Molina et al. (2016). We first present the model proposed by Stone (1972) to control cash using forecasts as the main input that we later use in a numerical example. The reason to focus on the Stone model is double: firstly, by accepting forecasts as its main input, the Stone model performance is affected by an important operating condition such as the predictive accuracy; secondly, within a framework of selection, alternative Stone models can be considered by varying a single parameter which is the forecasting horizon.

2.1 The Stone cash management model

A number of cash management models have been proposed to control cash balances based on a set of levels or bounds da Costa Moraes et al. (2015); Gregory (1976); Srinivasan and Kim (1986). In these models, the cash balance is allowed to wander around between bounds and when one of them is reached, a transaction $x_t \in \mathbb{R}$ is made at day t to restore the balance to a target level. This transaction x_t is given by the particular characteristics of the chosen model and goes to/comes from an alternative short-term investment, such as treasury bills or other marketable securities.

To this end, Stone (1972) proposed the use of forecasts in cash management. To the best of our knowledge, it was the first time that the concept of cash flow forecasting appeared as a key input to any cash management model. The author focused on the generally attempted tasks performed by cash managers in practice. They usually look ahead to incorporate data from their cash forecasts. For instance, consider a firm with a current balance of 12, a target of 8 and an upper bound of 10. An automatic return model would require a transaction of -4 to restore the balance to 8, but if the cash manager expected additional outflows of 3 in the next 3 days, the balance would have been close to target without any transaction. From these tasks, Stone derived the idea of including forecasts in his model to reduce transactions by smoothing cash flows.

By considering cash flow forecasts for the next k days, this procedure aims to minimize cost through reducing unnecessary transactions and smoothing cash flows. It is assumed that the current cash balance, b_0 , before making any transaction decision, is known and that a forecast of the net cash flow, \hat{f}_t , that will occur on each day t over the next k days is available. Note that k is a sub-period within a time horizon of T days, i.e, $k \leq T$. Then, the expected level of cash balances k

days from now is the sum of the current level of cash balances and the sum of k daily net cash flows. This can be expressed as:

$$\hat{b}_{t+k} = b_0 + \sum_{t=1}^k \hat{f}_t = b_0 + \hat{S}_k \quad (1)$$

where the sum of net predicted cash flows over the next k days is denoted by \hat{S}_k .

Next, a number of simple rules are proposed to return to a desired target cash balance, z , based on two sets of control limits as shown in Figure 1. One set is defined by h_1 and h_0 as the upper and lower control limits for initially considering transactions. The other set is defined by $h_1 - \delta_1$ and $h_0 + \delta_0$ as the upper and lower limits that determine if a transaction will actually be made. The set of rules are summarized as follows:

1. If the current cash balance b_0 , before making any decision, is inside of the control limits defined by h_1 and h_0 , no action is taken.
2. If the control limits h_1 and h_0 are exceeded, the forecasts over the next k days is considered to decide whether a transaction should be made.
 - (a) If the expected cash balance in the next k days, \hat{b}_{t+k} exceed the control limits defined by $h_1 - \delta_1$ and $h_0 + \delta_0$, a transaction x_t is made to return the expected cash balance to the target level z in k days.
 - (b) No action taken otherwise.

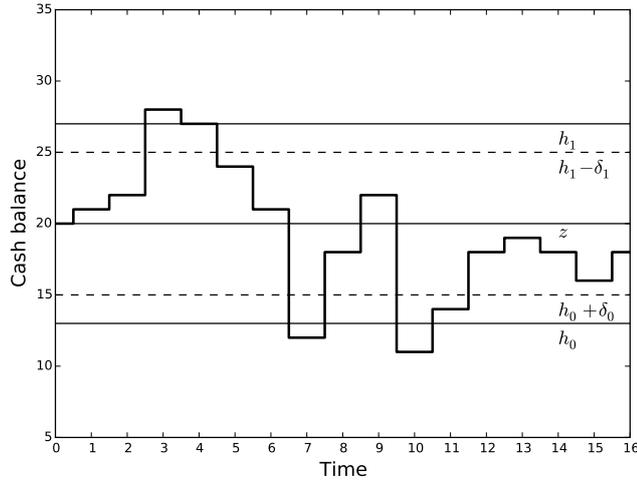


Fig. 1 Cash balance derived from the Stone model with $k = 3$ from the example in Table 1

According to these rules, when a transaction is made, the model returns the expected level of balances to the target level in k days rather than immediately returning the current balance to the target. Furthermore, the actual cash balance is the target plus the net cumulative forecast error. Summarizing, this procedure

to obtain the transaction x_t to perform on day t can be mathematically expressed as follows:

$$x_t = \begin{cases} z - \hat{b}_{t+k}, & \text{if } b_0 \geq h_1 \text{ and } \hat{b}_{t+k} > h_1 - \delta_1, \\ 0, & \text{if } h_0 < b_0 < h_1 \\ z - \hat{b}_{t+k}, & \text{if } b_0 \leq h_0 \text{ and } \hat{b}_{t+k} < h_0 + \delta_0. \end{cases} \quad (2)$$

An additional rule can be followed to make decisions at the beginning or at the end of the day, when the actual cash flow f_t is already known. From the following state transition law:

$$b_t = b_{t-1} + x_t + f_t \quad (3)$$

when the decision is made at the beginning of the day, then $b_0 = b_{t-1}$ holds and the current day forecast is included in the sum of forecast for the next k days. On the contrary, when the decision is made at the end of the day, $b_0 = b_{t-1} + f_t$ holds instead and the sum of forecasts cover the next k days. An example of the operation of this model presented in Stone (1972), when decisions are made at the end of the day, is summarized in Table 1.

Table 1 Comparison of the Stone 3-day look-ahead with no-day look-ahead procedure from Stone (1972) with \hat{f}_t = cash flow forecast; f_t = actual cash flow; e_t = forecasting error; x_t = control action; and b_t = cash balance. Initial cash balance $b_0 = 20$

Time	Forecast and actual flow			3-day look-ahead	No-day look-ahead		
				$k = 3$		$k = 0$	
t	\hat{f}_t	f_t	e_t	x_t	b_t	x_t	b_t
1	1	1	0	0	21	0	21
2	2	1	1	0	22	0	22
3	3	6	-3	0	28	-8	20
4	-1	-1	0	0	27	0	19
5	-2	-3	1	0	24	0	16
6	-3	-3	0	0	21	7	20
7	-8	-9	1	0	12	9	20
8	5	6	-1	0	18	0	26
9	6	4	2	0	22	-10	20
10	4	6	-2	0	11	0	26
11	5	3	2	-17	14	-9	20
12	4	4	0	0	18	0	24
13	0	1	-1	0	19	0	25
14	2	-1	3	0	18	0	24
15	-3	-2	-1	0	16	0	22
16	-1	2	-3	0	18	0	24

For illustrative purposes, let us consider this example in more detail. The target balance is $z = 20$, the control limits are $h_1 = 27$, and $h_0 = 13$; and $\delta_1 = \delta_0 = 2$, all figures in millions of euros. Since decisions are made at the end of the day, from a starting cash balance of 20, we add the actual cash flow, $f_1 = 1$, to obtain balance before decision $b_0 = 21$. Since the current balance is below $h_1 = 27$, no further action is taken. Consider now day $t = 10$, with balance before decision $b_0 =$

$22+6 = 28$, which is above $h_1 = 27$. Add the sum of predictions $\hat{S}_3 = 5+4-0 = 9$, to obtain the expected cash balance within 3 days, $\hat{b}_{13} = 28 + 9 = 37$, which is above $h_1 - \delta_1 = 25$. Then, we proceed to restore the balance to target $z = 20$ by transferring $x_{10} = 20 - 37 = -17$.

Stone claimed the utility of his model in comparison to a no-day look-ahead procedure by computing both the number and the total volume of transactions. In the example, it is clear that there are more transactions in the no-day look-ahead model than in the 3-day look-ahead model. Moreover, there is a greater total volume of transactions in the former than in the latter. However, it is not clear if this model would be useful for any value of k or if the predictive accuracy of forecasts would impact the performance of the procedure. This analysis can be done from a multiobjective perspective through the following formulation.

2.2 Formulation of the multiobjective cash management problem (MOCMP)

The MOCMP is characterized by its particular cost structure as follows. 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 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 loss function is defined as follows:

$$c(x_t) = \Gamma(x_t) + H(b_t) \quad (4)$$

where x_t is the transaction made at day t , b_t is the cash balance at the end of day t , $\Gamma(x_t)$ is a transfer cost function, and $H(b_t)$ stands for a holding/shortage cost function. The transfer cost function $\Gamma(x_t)$ is defined as:

$$\Gamma(x_t) = \begin{cases} \gamma_0^- - \gamma_1^- \cdot x_t & \text{if } x_t < 0, \\ 0 & \text{if } x_t = 0, \\ \gamma_0^+ + \gamma_1^+ \cdot x_t & \text{if } x_t > 0. \end{cases} \quad (5)$$

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

$$H(b_t) = \begin{cases} -u \cdot b_t & \text{if } b_t < 0; u > 0, \\ v \cdot b_t & \text{if } b_t > 0; v > 0. \end{cases} \quad (6)$$

For instance, consider an initial cash balance $b_{t-1} = 22$, a transaction given by a model $x_t = -8$, and an actual net cash flow $f_t = 6$, all amounts in millions of euro. Assume also transaction costs given by $\gamma_0^- = \gamma_0^+ = 50$ €, and $\gamma_1^- = \gamma_1^+ = 0.01\%$ per money unit. Consider a holding cost of $v = 0.05\%$ per money unit. Then, the final cash balance is given by $b_t = 22 - 8 + 6 = 20$ according to equation (3), the transaction cost is $\Gamma(x_t) = 850$ €, and the holding cost is $H(b_t) = 10000$ €. Consequently, the total cost for day t is $c(x_t) = 10850$ €.

Under this cost structure, the ultimate goal of the MOCMP defined in Salas-Molina et al. (2016) is to find the policy $X = \langle x_1, \dots, x_t, \dots, x_T \rangle$, with $x_t \in \mathbb{R}$, $t \in \mathbb{N}$, and $1 \leq t \leq T$ for a given planning horizon $T \in \mathbb{N}$. To this end, we can

measure the expected cost $C(m, T)$ for a model m and a planning horizon T by the average daily cost:

$$C(m, T) = \frac{1}{T} \sum_{t=1}^T c(x_t) = \frac{1}{T} \sum_{t=1}^T [F(x_t) + H(b_t)] \quad (7)$$

and the expected risk $R(m, T)$ either by the standard deviation or the upper semideviation of the daily cost. A common criticism made on the use of the standard deviation as a measure of risk is that it equally considers both positive and negative deviations. In this paper, we consider the upper semideviation to focus only on deviations above the mean:

$$R(m, T) = \left(E \left(\max\{c(x_t) - C(m, T), 0\}^2 \right) \right)^{1/2} \quad (8)$$

where E is the common expectation operator. By computing both the expected cost and risk of alternative models, we are in a position to locate models in the Cost-Risk space.

3 Showing models in the Cost-Risk space

Under the MOCMP framework, cash management models are evaluated in terms of cost and risk. The tradeoff between cost and risk for different models and operating conditions is a critical issue in the cash management decision-making process. In this paper, we define the Cost-Risk (CR) space as the graphical representation of cost in the horizontal axis and risk in the vertical axis. A point in the CR space shows the combined cost and risk performance of any particular model when applied to a cash flow data set $F = \langle f_1, \dots, f_N \rangle$ with $N \in \mathbb{N}$ real observations. The basic idea of the CR space is to show model asymmetry or bias (if any) towards cost and risk for any particular model. Thus, the particular location of a CR point provides useful information about a model for comparison and selection purposes.

Definition 1 Given a cash management model m and a particular set of model parameters p , a p -model, denoted by m_p , is an instance of model m with parameters determined by the elements of the vector of parameters $p \in \mathbb{R}^n$.

For instance, if we use m to refer to all Stone cash management models and define p as the tuple $(k, z, h_1, h_0, \delta_1, \delta_0) = (3, 20, 27, 13, 2, 2)$, m_p is an instance model of the Stone type defined by the set of parameters p . In what follows, we refer to Stone p -models as models unless stated otherwise.

Furthermore, we use the term operating condition to refer to any factor that may influence the performance of a model such as the cost scenario and the predictive accuracy of the forecasts used in the Stone model. Following Hernández-Orallo et al. (2014, 2013), we refer to the set of all relevant operating conditions as a context that can be expressed as a set of n parameters $\alpha = \langle \alpha_1, \dots, \alpha_n \rangle$. It is important to separate the notion of operating condition from the parameters of the model. The operating condition is somehow given to the decision-maker either internally (e.g., risk preferences), or externally, (e.g., cost scenarios). On the other hand, the parameters of the model, such as the look-ahead period or the control limits used, are adjustable by the decision-maker and represent the decision variables that need to be appropriately selected. Within this framework, we define:

Definition 2 A loss function is any function $\mathcal{L} : \mathbb{Y} \times \mathbb{Y} \rightarrow \mathbb{R}$ which compares elements in some output domain \mathbb{Y} .

Typical examples of this output domain \mathbb{Y} within the context of cash management are the cost and the risk of a cash management model computed, for instance, by means of equations (7) and (8).

Definition 3 Given a p -model m_p , a context α and a cash flow data set $F \in \mathbb{R}^N$, a p_α^* -model is the p -model which minimizes some loss function $\mathcal{L}(m_p, \alpha)$ when applied to F in context α .

We are then facing an optimization problem in which the goal of the cash manager is to find the p_α^* -model, i.e., the p -model with best performance for a given context. To this end, we can compare models by computing cost and risk over a given planning horizon. Since the ideal model is the one with zero cost-zero risk, one possible way to determine the best model is computing the distance to the ideal point $(0, 0)$. For normalization purposes, we use a baseline model m_0 , defined as a no-transaction policy with $x_t = 0$ for all $t \in T$. Thus, we finally compute the cost and risk for alternative models in the CR space using equations (7) and (8), and through the following cost and risk indexes:

$$\theta_1(m, T) = \frac{C(m, T)}{C(m_0, T)} \quad (9)$$

$$\theta_2(m, T) = \frac{R(m, T)}{R(m_0, T)}. \quad (10)$$

By representing points (θ_1, θ_2) for different models as shown in Figure 2 (a), we can derive three useful insights: (i) we can compare their relative performance by computing distances to the ideal point $(0, 0)$; (ii) we can determine if the model is non-trivial by observing if its CR coordinates are within the unit square; and (iii) we can conclude if the model is biased towards cost or risk by observing if the model is either over or under the increasing diagonal. For instance, from Figure 2(a) one can infer that model m_1 is better than m_2 because is closer to the ideal point. Note also that any model located out of the unit square performs worse than the baseline in terms of cost, risk or both.

Moreover, model m_1 is located in the upper left region meaning that it is able to reduce cost in a higher proportion than risk with respect to the baseline m_0 . We say then that model m_1 is biased towards risk from the ideal point perspective since it is more risky than costly, i.e., $\theta_1(m_1, T) < \theta_2(m_1, T)$. On the other hand, model m_2 is located in the lower right region of the CR space delimited by the increasing diagonal. This particular location shows that, in comparison to the baseline model m_0 , model m_2 reduces cost in a lower proportion than risk. We say then that model m_2 is biased towards cost from the ideal point perspective since it is more costly than risky, i.e., $\theta_1(m_2, T) > \theta_2(m_2, T)$. Models located in the increasing diagonal are equally costly than risky.

An important concept derived from the representation of particular models in the CR space is the notion of hybrid models (Hernández-Orallo et al., 2013). Given any two models, we can construct a hybrid model as a linear combination of the control actions derived from both models. This concept is graphically illustrated in Figure 2 (a) by means of the dashed line passing through m_1 and m_2 . Each point

in this line represents the expected CR performance of a hybrid model constructed as a linear combination of models with coefficients adding up to one.

In addition, when comparing alternative cash management models in the CR space, the concepts of Pareto optimality and convexity of the Pareto front provide further insights. By adding a number of alternative models to the Cost-Risk space from Figure 2 (a), we can build the Pareto front in which no model can reduce the level of cost (risk) without increasing the level of risk (cost) as shown in Figure 2 (b). More formally, any model m characterized by the point $(\theta_1(m, T), \theta_2(m, T))$ belonging to a non-empty feasible region $S \in \mathbb{R}^2$ is called Pareto optimal (Ballester and Romero, 1998; Miettinen, 2012) if there does not exist any other feasible point that can achieve the same or better performance for all the criteria while being strictly better for at least one criterion.

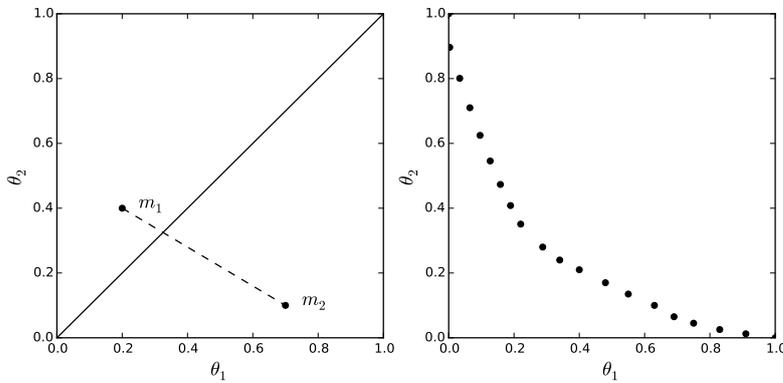


Fig. 2 (a) Two models in the normalized Cost-Risk space; (b) Pareto front of alternative models in the normalized Cost-Risk space

The Pareto front provides valuable information about alternative cash management models. Besides bias towards either cost or risk, the concept of convexity of the Pareto front is crucial to ensure that there is a trade-off between criteria in order to find an optimal solution among feasible alternatives. This trade-off measures the amount of achievement of one criterion that must be sacrificed in order to gain a unitary increase in the other criterion (Ballester and Romero, 1998). Furthermore, a number of distances to the ideal point $(0,0)$ can be computed to determine which model is best in terms of both cost and risk.

However, in the search for the best model, cash managers' preferences for cost and risk are not equal. A risky cash manager may accept a higher risks to achieve the same cost reduction than a conservative one. Next section considers risk preferences in the CR space to enhance the selection process.

4 Choosing models according to risk preferences

We mentioned in Section 3 that the normalized CR space allows comparisons of alternative cash management models in terms of cost and risk. The increasing diagonal divides the CR space into two regions: the upper-left half or risky region, and the bottom-right half or the conservative region. A neutral cash manager, i.e., with no bias for cost and risk, would choose models with similar values of cost and risk indexes or, equivalently, models close to the diagonal, where $\theta_1 = \theta_2$.

Recall that one model in the CR space dominates another either if it is no worse in terms of both cost and risk and it is strictly better either in cost or risk. All models are dominated by the model with the lowest risk from the cost perspective. Likewise, from the risk perspective, all models are dominated by the model with the lowest cost. From that, all non-dominated models form the Pareto front, i.e., the set of feasible alternatives for cash managers to choose according to their particular preferences.

Once models are presented in the CR space, cash managers have to choose the best compromise model that minimizes total loss in terms of cost and risk. Compromise programming (Zeleny, 1982; Yu, 1985; Ballestero and Romero, 1998) can then be used to solve the problem by minimizing the distance to the ideal point $(0, 0)$, usually unfeasible. Moreover, when considering particular risk preferences defined by w , the following family of weighted distance functions allows to determine the proximity of any particular point (θ_1, θ_2) to the ideal one:

$$\mathcal{L}_h(m, \alpha) = \left[w^h \cdot \theta_1^h + (1 - w)^h \cdot \theta_2^h \right]^{1/h} \quad (11)$$

where $h \in \mathbb{N}^+$ is the parameter defining the family of distance functions and weight $w \in [0, 1]$ denotes the particular risk preferences of the cash manager. Note that \mathcal{L}_1 is the Manhattan distance; \mathcal{L}_2 is the Euclidean distance, and \mathcal{L}_∞ is the Chebyshev distance. Only these cases are likely to be used in practice for interpretation and computational reasons (Ballestero, 2007; Ringuest, 2012). Computations with Manhattan distances are simpler and tend to produce solutions located further from the diagonal. On the other hand, Euclidean distances tend to produce more balanced solutions but with more computational burden due to non-linearity (Ballestero, 2007).

The concept of isometric curves in the CR space can be derived from the distance to the ideal, as combinations of θ_1 and θ_2 , that have equal performance and, thus, may represent models that cash managers are indifferent among. Risk preferences are then considered through the use of parameter w . Since we are dealing with a loss function to be minimized, the higher the value of w , the riskier the cash manager. From that, two families of CR isometrics can be considered by varying a parameter d for Manhattan ($h = 1$) and Euclidean distances ($h = 2$). For Manhattan distances, the total loss for a particular point can be graphically calculated as the sum of the distance to the y-axis and the distance to the x-axis. An interesting characterization of the solution for Manhattan distances can be found in Hernández-Orallo et al. (2013).

In what follows, we aim to characterize the solution to the problem when dealing with Euclidean distances. For $h = 2$, the total loss for a particular point can be graphically calculated as the sum of the squared distance to the y-axis, i.e., cost when risk is zero, and the squared distance to the x-axis, i.e., risk when cost

is zero. Then, a family of isometric curves can be defined by varying a parameter d so that the following expression holds:

$$w^2 \cdot \theta_1^2 + (1 - w)^2 \cdot \theta_2^2 = d^2. \quad (12)$$

This equation leads to consider isometric circles and ellipses as a tool to select among alternative cash management models as shown in Figure 3.

Definition 4 Given a loss function \mathcal{L}_2 and a context α , neutral cash management isometrics are represented by circles since $w = (1 - w)$, risky isometrics are represented by ellipses since $w > (1 - w)$ and conservative isometrics are represented by ellipses since $w < (1 - w)$.

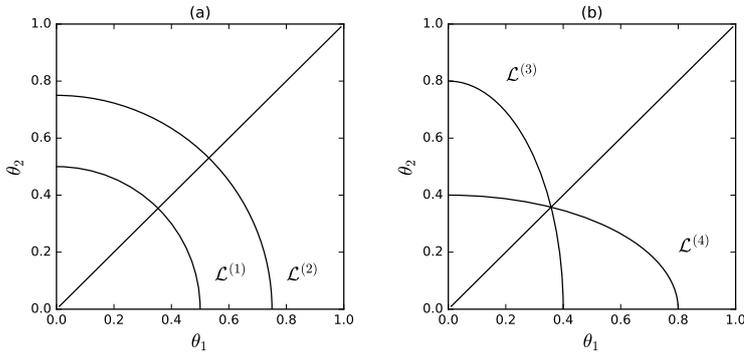


Fig. 3 (a) Isometrics for neutral cash managers; (b) Isometrics for risky/conservative cash managers

Neutral cash managers may consider a number of circles with increasing radius such as $\mathcal{L}^{(1)}$ and $\mathcal{L}^{(2)}$ in Figure 3(a). The model on the circle with minimum radius is the optimal model. Similarly, risky cash managers may consider a number of ellipses such as $\mathcal{L}^{(3)}$ in Figure 3(b) with vertical semiaxis larger than its horizontal semiaxis due to $w > (1 - w)$. Finally, conservative cash managers may consider a number of ellipses such as $\mathcal{L}^{(4)}$ in Figure 3(b) with vertical semiaxis smaller than its horizontal semiaxis due to $w < (1 - w)$. Then, the model on the ellipse with minimum horizontal/vertical semiaxis is the optimal model.

5 Deriving loss curves from operating conditions

In the previous sections, we showed different cash management models in the CR space and we also provided a tool to choose among them according to the particular risk preferences of cash managers. However, chances are that context changes may impact the performance of models. Therefore, cash managers may be interested in determining which model is best for a given range of operating conditions. In this section, we propose a method to derive loss curves that can be useful to assess the impact of operating conditions in the performance of cash management models.

Here, loss means the combined cost-risk performance computed using equation (11) according to the particular preferences of cash managers.

Recall from Section 3 that the term operating condition refers to any factor that may influence the performance of a model. A typical example of an operating condition when deploying models using forecasts as a key input is the prediction error. Let us assume that cash managers can produce forecasts with a given forecasting error ε with respect to a trivial forecast \bar{f} that is computed for a data set of N forecasts as follows:

$$\varepsilon = \frac{\sum_{t=1}^N (\hat{f}_t - f_t)^2}{\sum_{t=1}^N (\bar{f} - f_t)^2} \quad (13)$$

where \hat{f}_t is the prediction at time t , f_t is the real observation at the time t , and \bar{f} is the arithmetic mean of the real observations used to obtain predictions \hat{f}_t . Note that the closer ε is to zero, the better the predictive accuracy. On the contrary, if ε is close to one, the performance is similar to the mean as a naive forecast. Values greater than one show that the forecaster has no predictive ability.

Assume now that we want to obtain the best model for a range of prediction errors of particular interest for a hypothetical cash manager. To this end, we rely on a number of simulations. More precisely, our experimental design is based on multiple replicates of random cash balance paths under a Monte Carlo framework (Glasserman, 2003). These paths are obtained from a known initial cash balance and T independent draws from a real cash flow probability density function (pdf), whose cumulative sum determines a particular cash balance evolution. Whether this pdf is derived from a real data set or selected from one of the usual pdfs in financial problems, such as the Gaussian distribution, is at the discretion of the practitioner. Here, since we are provided with a real cash flow data set, we use an empirical pdf obtained from 1000 real cash flow observations from a Spanish industrial company.

For the sole purpose of this example, assume that our hypothetical cash manager knows, from the characteristics of the company cash flows, that the set of control limits that best worked in the past are $\{h_0, \delta_0, z, \delta_1, h_1\} = \{2.5, 0.4, 5.0, 0.4, 7.5\}$, all figures in millions of euros. In addition, our cash manager may produce cash flow predictions with an average prediction error in the range of $\varepsilon = [0, 1]$ in terms of normalized squared error from equation (13).

On the other hand, the current cost structure is determined by $\gamma_0^+ = 50$ €, $\gamma_0^- = 50$ €, $\gamma_1^+ = 0.1\%$, and $\gamma_1^- = 0.01\%$, of the transaction volume. Assume also a holding cost $v = 10\%$ per annum, and a daily penalty cost $u = 30\%$. Hence, in this example we set $\alpha = \langle \varepsilon, \gamma_0^+, \gamma_0^-, \gamma_1^+, v, u \rangle$ as the particular context of interest in which alternative models are going to be analyzed. In this context, there is only one variable operating condition, namely, the prediction error ε . Since this operating condition may present one or more different levels, we can express it as a tuple of d levels, $\alpha_i = \langle l_{i1}, \dots, l_{id} \rangle$, e.g., $\varepsilon = \langle 0, 0.25, 0.50, 0.75, 1 \rangle$, or as a continuous interval, $\alpha_i = [\alpha_{min}, \alpha_{max}]$, e.g., $\varepsilon = [0.5, 1]$.

Summarizing, the final goal is to derive a loss curve that shows the impact in equation (11) of the change of any variable operating condition in context α . For simplicity, we also set $h = 1$ and consider a neutral cash manager with $w = 0.5$, within a planning horizon of $T = 30$ days. Then, we proceed following the steps detailed in Algorithm 1.

Algorithm 1: Algorithm to build context loss curves

```

1 Inputs: Model  $m$ ; context  $\alpha$ ; cash flow data set  $F$ ; planning horizon  $T$ ;
   number of replicates  $r$ ;
2 Output: Context loss curve  $\mathcal{L}_h(m, \alpha)$ ;
3 Define a baseline model  $m_0$ ;
4 for each possible combination  $\alpha_i$  in  $\alpha$  do
5   for each replicate do
6     Generate a random cash balance path form  $F$ ;
7     Compute  $C(m_0, T)$  and  $R(m_0, T)$ ;
8     Compute  $C(m, T)$  and  $R(m, T)$ ;
9     Compute  $\mathcal{L}_h(m, \alpha_i)$ ;
10  end
11  Estimate and plot expected values for each  $\mathcal{L}_h(m, \alpha_i)$ ;
12 end

```

Since the only variable operating condition is the prediction error, by varying ε in the range $[0, 1]$, we are implicitly comparing different contexts. Then, we consider eleven possible prediction errors in steps of 0.1 from perfect prediction ($\varepsilon = 0$) to a naive prediction ($\varepsilon = 1$). For each value of ε , we generated 100 random cash balance paths of 30 days, equivalent to more than twelve working years and computed both the cost and risk using equation (7) and (8) for alternative Stone models and also for a trivial model m_0 consisting in taking no control action. These random experiments lead to an estimation of the expected loss of alternative Stone models for each possible prediction error. More precisely, we consider five alternative Stone cash management models for prediction horizons between $k = 0$ and $k = 4$ days. Note that no prediction is used in the case of $k = 0$. The comparative loss for models with $k = 0$ and $k = 1$ for the whole range of possible prediction errors is shown in Figure 4(a). The rest of models are not presented due to poorer performance.

Implicitly, by comparing models using predictions ($k \geq 1$) to another model using no prediction ($k = 0$), we are testing the ability of the Stone model to produce better results for a range of prediction errors. Our example results show that the 1-day look-ahead Stone model produced slightly better results for low prediction errors but rapidly increased the expected loss in comparison to the no-look ahead model. From that, unless our hypothetical cash manager is able to consistently produce forecasts with error below 0.2, the deployment of a 1-day Stone model is not recommended.

Since no model is the best model for the whole possible range of operating conditions, cash managers can rely on loss curves to select the best models for any particular context. Thus, the resulting loss curve depicted in 4(b) shows the minimum loss that can be achieved after selecting the best model for each prediction error. The first two point belong to the Stone model with $k = 1$ and the remaining to the no-look-ahead procedure. Note also that in the example we only compare Stone models varying the parameter k and for different prediction errors. In the search for more robust models, our method can be extended for all available cash management models and a wide range of operating conditions such as alternative cost scenarios.

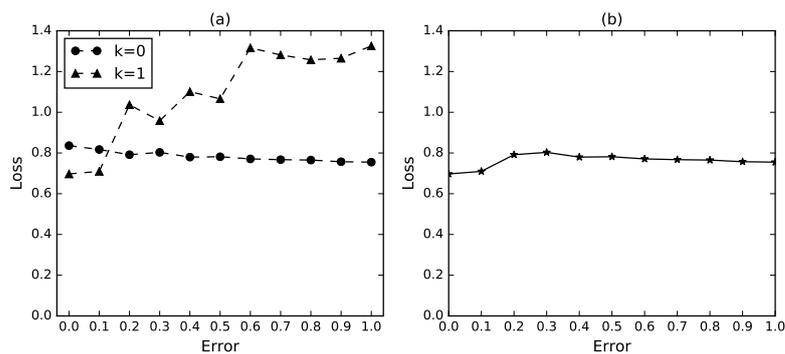


Fig. 4 (a) Loss curve for two alternative models and different prediction errors; (b) Expected loss after selecting the best model for each prediction error

6 Concluding remarks

This paper demonstrates that ROC analysis can be adapted to enhance the understanding of the multiobjective cash management problem. More precisely, we present three procedures for: (i) showing models in the Cost-Risk space; (ii) choosing models according to the risk preferences of cash managers; (iii) deriving loss curves for different operating conditions. These procedures allow cash managers to answer a number of important questions such as:

1. What is the performance of a model in terms of cost and risk?
2. What model is best according to cash managers risk preferences?
3. Under what circumstances a model is better than another?

All the previous questions are answered through the procedures presented here from a double perspective: quantitative and graphical. The essence of ROC analysis is then respected. Although the graphical analysis is limited to one, at most two operating conditions, the quantitative analysis can be generalized to any number of operating conditions.

A new interesting research line may be derived from loss curves for multiple operating conditions by considering multidimensional techniques such as factorial analysis. Further research would also be needed to assess under what circumstances an operating condition is more important than another.

Acknowledgments

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

References

- Ballester, 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.

- Ballesteros, E. and Romero, C. (1998). *Multiple criteria decision making and its applications to economic problems*. Springer Science & Business Media.
- Bi, J. and Bennett, K. P. (2003). Regression error characteristic curves. In *Proceedings of the 20th International Conference on Machine Learning (ICML-03)*, pages 43–50.
- Bradley, A. P. (1997). The use of the area under the roc curve in the evaluation of machine learning algorithms. *Pattern recognition*, 30(7):1145–1159.
- 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.
- Doumpos, M. and Zopounidis, C. (2007). Model combination for credit risk assessment: A stacked generalization approach. *Annals of Operations Research*, 151(1):289–306.
- Drummond, C. and Holte, R. C. (2000). Explicitly representing expected cost: An alternative to roc representation. In *Proceedings of the sixth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 198–207. ACM.
- Drummond, C. and Holte, R. C. (2006). Cost curves: An improved method for visualizing classifier performance. *Machine Learning*, 65(1):95–130.
- Elkan, C. (2001). The foundations of cost-sensitive learning. In *International joint conference on artificial intelligence*, volume 17, pages 973–978. Lawrence Erlbaum associates Ltd.
- Fawcett, T. (2006). An introduction to roc analysis. *Pattern recognition letters*, 27(8):861–874.
- Flach, P. A. (2003). The geometry of roc space: understanding machine learning metrics through roc isometrics. In *Proceedings of the 20th International Conference on Machine Learning (ICML-03)*, pages 194–201.
- García-Bernabeu, A., Benito, A., Bravo, M., and Pla-Santamaria, D. (2016). Photovoltaic power plants: a multicriteria approach to investment decisions and a case study in western spain. *Annals of Operations Research*, 245(1-2):163–175.
- Glasserman, P. (2003). *Monte Carlo methods in financial engineering*, volume 53. Springer Science & Business Media.
- Gregory, G. (1976). Cash flow models: a review. *Omega*, 4(6):643–656.
- Hernández-Orallo, J. (2013). Roc curves for regression. *Pattern Recognition*, 46(12):3395–3411.
- Hernández-Orallo, J., Flach, P., and Ferri, C. (2013). Roc curves in cost space. *Machine learning*, 93(1):71–91.
- Hernández-Orallo, J., Lachiche, N., and Martínez-Usó, A. (2014). Predictive models for multidimensional data when the resolution context changes. In *Workshop on Learning over Multiple Contexts at ECML*, volume 2014.
- Metz, C. E. (1978). Basic principles of roc analysis. In *Seminars in nuclear medicine*, volume 8, pages 283–298. Elsevier.
- Miettinen, K. (2012). *Nonlinear multiobjective optimization*, volume 12. Springer Science & Business Media.
- Ringuest, J. L. (2012). *Multiobjective optimization: behavioral and computational considerations*. Springer Science & Business Media.
- Ross, S. A., Westerfield, R., and Jordan, B. D. (2002). *Fundamentals of corporate finance*. McGraw-Hill, sixth edition.

-
- 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.
- Srinivasan, V. and Kim, Y. H. (1986). Deterministic cash flow management: state of the art and research directions. *Omega*, 14(2):145–166.
- Steuer, R. E., Qi, Y., and 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, B. K. (1972). The use of forecasts and smoothing in control limit models for cash management. *Financial Management*, 1(1):72.
- Torgo, L. (2005). Regression error characteristic surfaces. In *Proceedings of the eleventh ACM SIGKDD international conference on Knowledge discovery in data mining*, pages 697–702. ACM.
- Yu, P.-L. (1985). *Multiple criteria decision making: concepts, techniques and extensions*. Plenum Press.
- Zeleny, M. (1982). *Multiple criteria decision making*. McGraw-Hill.