



UNIVERSITAT
POLITÈCNICA
DE VALÈNCIA



Data-driven decision-making and its application to the corporate cash management problem

Francisco Salas Molina

September 2017

Advisors: Juan A. Rodríguez-Aguilar
Francisco J. Martín
Pablo Díaz García

Contents

Abstract	iii
Contents	iii
1 Introduction	13
1.1 Motivation	13
1.2 Challenges	22
1.3 Contributions	25
1.4 Structure	34
2 Background and related research	39
2.1 Definition of the cash management problem	39
2.2 A review of cash management models	43
2.3 Analysis	78
I Data-driven cash flow analysis	85
3 The savings hypothesis	87
3.1 Motivation	87

3.2	Description and data preprocessing	88
3.3	Building forecasting models	91
3.4	Forecasting models' comparison	96
3.5	Does a better forecast produce better policies?	99
3.6	Summary	107
4	Empirical analysis of daily cash flow	111
4.1	Motivation	111
4.2	Data summary	114
4.3	A cross-validated test for non-linearity	121
4.4	The impact of data transformations	127
4.5	Summary	129
II	Multiobjective cash management	133
5	A multiobjective approach to the cash management problem	135
5.1	Motivation	135
5.2	Formalizing the multiobjective cash management problem	137
5.3	An illustrative example using the Miller and Orr's model	143
5.4	On the utility of cash management models	152
5.5	Summary	154
6	On the use of distance indexes to find robust policies	157
6.1	Motivation	157
6.2	On the use of distances as a measure of robustness	160
6.3	An illustrative example using the Miller and Orr's model	166
6.4	Discussion	169
6.5	Summary	170

7	Selecting cash management models from a multiobjective perspective	173
7.1	Motivation	174
7.2	Showing models in the cost-risk space	176
7.3	Choosing models according to risk preferences	180
7.4	Deriving loss curves from operating conditions	183
7.5	Summary	186
8	Boundless multiobjective models for cash management	189
8.1	Motivation	189
8.2	Linearizing the cash management problem	191
8.3	Reformulating the Gormley-Meade's model as a linear-quadratic program	194
8.4	Boundless multiobjective models for cash management	198
8.5	Sensitivity analysis to forecasting errors	201
8.6	Summary	205
9	Compromise programming models for cash management	207
9.1	Motivation	207
9.2	Risk analysis in cash management.	209
9.3	Compromise models to solve the MOCMP	215
9.4	Summary	226
III	Managing multiple bank accounts	229
10	Robust cash management with multiple bank accounts	231
10.1	Motivation	231
10.2	Formulation of the MBACMP.	233
10.3	Theoretical results	236
10.4	Robust formulations of the MBACMP	238
10.5	A new two-stage cost-sensitive robust formulation of the MBACMP	241
10.6	Empirical case study	245

10.7 Summary	250
11 PyCaMa: Python for cash management	253
11.1 Motivation	253
11.2 Problem formulation	255
11.3 Software description	257
11.4 An illustrative example	259
11.5 Impact	261
11.6 Summary	262
12 Conclusions and future work	265
12.1 Summary of challenges and contributions	265
12.2 Data-driven multiobjective decision-making	271
12.3 Future work	273
A Terminology and acronyms	277
Bibliography	281

List of Figures

1.1	Cash management as a control problem.	15
1.2	The cash management problem.	16
1.3	Cost of holding and transferring cash.	19
1.4	Contribution 1. The savings hypothesis.	27
1.5	Contributions 2 and 3. Empirical evidence on daily cash flow. . .	29
1.6	Contributions 4, 5, and 6. Multiobjective cash management. . .	31
1.7	Contributions 7 and 8. Boundless models	32
1.8	Contribution 9 and 10. Managing multiple bank accounts . . .	33
1.9	Integrating the contributions of this thesis.	35
2.1	The Baumol model.	44
2.2	The Tobin model.	46
2.3	The Miler-Orr model.	49
2.4	Monthly net cash flow example with two fixed days of payment. .	53
2.5	Eppen-Fama model representation with two return points. . . .	55
2.6	Structure of the Stone model with two sets of limits	62

3.1	Histogram for data sets 1 and 2 compared to a normal distribution.	89
3.2	Mean square error comparison for data set 1	99
3.3	Mean square error comparison for data set 2	99
3.4	Savings for data set 1 and random forest	106
3.5	Savings for data set 2 and radial basis functions	106
3.6	Potential cost savings analysis methodology	108
4.1	Poincaré map with lag 1 for time-series 1 and 2.	118
4.2	Simplified flow chart for our non-linearity test	123
4.3	A basic decision tree	125
5.1	Efficient frontier in the cost-risk space	146
5.2	Elasticity plot	149
5.3	Weighted Sharpe-like ratio	151
6.1	Distances between cash policies evaluated over two data sets. .	164
6.2	Efficient frontier in the cost-robustness space	167
7.1	Alternative models in the normalized cost-risk space	178
7.2	Pareto front of alternative models in the normalized cost-risk space	180
7.3	(a) Isometrics for neutral cash managers; (b) Isometrics for risky/conservative cash managers	182
7.4	(a) Loss curve for two alternative models and different prediction errors; (b) Expected loss after selecting the best model for each prediction error	185
8.1	The common two-assets setting in the cash management problem.	192
8.2	The dynamic simple policy of Gormley-Meade.	194

8.3	Run time for different planning horizons.	201
8.4	Median cost-risk loss of the BM for different forecasting errors.	205
9.1	Probability density function	213
9.2	Example of efficient set.	219
9.3	Policy and balance using LP.	223
9.4	Policy and resulting cash balance using QP	225
10.1	A cash management system with three bank accounts.	234
10.2	An example of a loop between accounts.	237
10.3	A cash management system with two bank accounts.	244
10.4	Total loss due to estimation errors.	246
10.5	A cash management system with two bank accounts and an investment account.	247
11.1	A cash management system with three accounts.	255
11.2	Optimal balances for the cost minimization example.	261

List of Tables

2.1	Assumptions of the Baumol and Miller-Orr models.	61
2.2	Comparison Stone with no-look-ahead procedure	64
2.3	Open research questions in cash management	82
2.4	Characteristics of the main models	83
3.1	Data set summary.	90
3.2	Model selection according to average error ratio	98
3.3	Average predictive error	100
3.4	Cost scenarios	102
3.5	Average daily savings	104
4.1	Data sets statistical summary	116
4.2	Correlation and seasonality test results.	120
4.3	Results of the test for non-linearity	126
4.4	Results of the test for non-linearity after transformations	128
4.5	Summary after transformations	129

5.1	Three alternative scenarios	144
5.2	Central sample of the efficient frontier	147
5.3	Relative cost-risk performance	148
6.1	Central sample of the efficient frontier	168
6.2	Evaluation of policies	169
8.1	Data sets description. Figures in thousands of €.	203
9.1	Advantages and disadvantages of alternative risk measures.	214
9.2	Example efficient set for a Miller-Orr model with three levels.	219
10.1	Example error data set in millions of €.	245
10.2	Cost structure data (β) for our case study.	247
10.3	Minimum balances in euros used the case study.	248
10.4	Results of the empirical case study. Cost figures in euros.	249
11.1	Python inputs and outputs of PyCaMa	259
11.2	Cost structure data for the example.	259
A.1	Acronyms and abbreviations	280

"A man of knowledge lives by acting,
not by thinking about acting."

Carlos Castaneda

"The specialist,
who knows increasingly more about less,
in the limit, all about nothing,
is as regressive as the generalist,
who knows increasingly less about more,
in the limit, nothing about all."

Jorge Wagensberg

"The tool that serves as intermediary
between theory and practice,
between thought and observation,
is mathematics; it is mathematics
which builds the linking bridges,
and gives them ever more reliable forms."

David Hilbert

Dedication

To those who taught me,
no matter what,
no matter when.

Agradecimientos

Mi agradecimiento más grande a mi familia, a mis hermanas María Inmaculada y Lucía, a mis sobrinas Mar y Neus, y sobre todo a mis padres Francisco y Francisca, por toda una vida dedicada a los suyos. A todos mis profesores, a todos, por despertar en mí el deseo por aprender. A mis amigos y compañeros de trabajo de Hilaturas Ferre, por estar siempre pendientes.

Mi agradecimiento más especial a mis directores Juan Antonio, Francisco y Pablo, y por extensión, al Instituto de Investigación en Inteligencia Artificial de Barcelona, a BigML y a la Universidad Politécnica de Valencia por permitir que esta tesis se convierta en realidad. Todos ellos, el científico, el emprendedor y el profesor han enriquecido enormemente este viaje con sus conocimientos y con su visión personal.

Mi agradecimiento más profundo a la mejor de mis historias, Cristina.

Banyeres de Mariola
Septiembre 2017
Francisco Salas Molina.

Abstract

This thesis investigates the cash management problem from a multidimensional perspective. Cash management focuses on finding the balance between cash holdings and short-term investments. Typically, cash managers make decisions based usually on a firm's optimal cash balance for operational and precautionary purposes. We here explore the opportunities for improved decision-making derived from modeling cash flow uncertainty with the help of data-driven procedures within a multiobjective context. On the one hand, cash managers may achieve cost savings by forecasting future cash flows. To this end, we perform an empirical analysis of daily cash flow time-series to take advantage of modern machine learning techniques as a key step to connect data analysis and optimization methods in cash management. On the other hand, cash managers may be interested not only in the cost but also in the risk associated to decision-making. Thus, we address the cash management problem from a multiobjective perspective focusing on both cost and risk. In addition, under the current situation of time-varying financial circumstances, the selection of cash management models according to operating conditions and its robustness are worth considering questions. We also show the utility of forecasts through a new cash management model which outperforms the state-of-the-art by guaranteeing optimal solutions. Since most firms usually deal with cash management systems with multiple accounts, we develop a framework to formulate and solve the multiple bank accounts cash management problem. Finally, in an attempt to fill the gap between theory and practice, we also provide a software library in Python for practitioners interested in building decision support systems for cash management.

Resumen

Esta tesis investiga el problema de gestión de tesorería desde un punto de vista multidimensional. La gestión de tesorería trata de equilibrar la cantidad que se mantiene en efectivo y la que se dedica a inversiones a corto plazo. Normalmente, los tesoreros toman decisiones basándose en el nivel óptimo de tesorería por motivos operativos y de precaución. En esta tesis exploramos las oportunidades para mejorar la toma de decisiones derivadas de modelar la incertidumbre presente en los flujos de caja con la ayuda de procedimientos basados en datos en un entorno multiobjetivo. Por un lado, los tesoreros pueden conseguir ahorros a través de la previsión de tesorería. Para ello, realizamos un estudio empírico con el objetivo de aprovechar las más recientes técnicas de aprendizaje automático como paso clave para conectar el análisis de los datos disponibles con los procesos de optimización en la gestión de tesorería. Por otro lado, los tesoreros pueden estar interesados no solo en el coste sino también en el riesgo asociado a sus decisiones. Por esta razón, tratamos el problema de gestión de tesorería desde una perspectiva multiobjetivo, considerando tanto el coste como el riesgo. Además, debido a la cambiante situación financiera actual, exploramos la selección de modelos de gestión de tesorería en función de diferentes condiciones operativas y de su robustez. También demostramos la utilidad de las previsiones a través de un nuevo modelo de gestión de tesorería que mejora el estado del arte al garantizar soluciones óptimas. Como la mayoría de las empresas trabaja con sistemas de tesorería con múltiples cuentas bancarias, desarrollamos un marco para la formulación y solución del problema de gestión de tesorería con múltiples cuentas bancarias. Finalmente, en un intento de acercar teoría y práctica, también ofrecemos una librería de software en Python para usuarios interesados en la construcción de sistemas de ayuda a la toma de decisiones en gestión de tesorería.

Resum

Esta tesi investiga el problema de gestió de tresoreria des d'un punt de vista multidimensional. La gestió de tresoreria tracta d'equilibrar la quantitat que es manté en efectiu i la que es dedica a inversions a curt termini. Normalment, els tresorers prenen decisions basant-se en el nivell òptim de tresoreria per motius operatius i de precaució. En aquesta tesi explorem les oportunitats per millorar la presa de decisions derivades de modelitzar la incertesa present en els fluxos de caixa amb l'ajuda de procediments basats en dades. Per un costat, els tresorers poden aconseguir estalvis de costos mitjançant la previsió de tresoreria. Per tal d'aconseguir-ho, realitzem d'un estudi empíric amb l'objectiu d'aprofitar les més recents tècniques d'aprenentatge automàtic per connectar l'anàlisi de les dades disponibles amb els processos d'optimització en la gestió de tresoreria. Per altra banda, els tresorers poden estar interessats no sols en el cost sinó també en el risc associat a les seues decisions. Per tant, tractem el problema de gestió de tresoreria des d'un punt de vista multiobjectiu, fixant-se tant en el cost com en el risc. A més a més, degut a la canviant situació financera actual, explorem la selecció de models de gestió de tresoreria en funció de diferents condicions operatives i de la seua robustesa. També demostrem la utilitat de les previsions mitjançant un nou model de tresoreria que millora l'estat de l'art al garantir solucions òptimes. Com que la majoria d'empreses treballa amb sistemes de tresoreria amb múltiples comptes bancaris, desenvolupem un marc per a la formulació i solució del problema de gestió de tresoreria amb múltiples comptes bancaris. Finalment, en un intent d'apropar teoria i pràctica, també oferim un llibreria en Python per a usuaris interessats en la construcció de sistemes d'ajuda a la presa de decisions en la gestió de tresoreria.

Chapter 1

Introduction

This thesis investigates the cash management problem from a multiobjective perspective. Cash management focuses on finding the balance between cash holdings and short-term investments. Typically, cash managers make daily decisions based usually on a firm's optimal cash balance for operational and precautionary purposes. In this thesis, we argue that opportunities for improved decision-making by both modeling cash flow uncertainty and by considering additional objectives are significant in cash management. On the one hand, cash managers may achieve cost savings by forecasting future cash flows. On the other hand, they may be interested not only in the cost but also in the risk associated to cash management decision-making. In what follows, we first motivate the key research topics presented in this thesis. Next, we describe the challenges and how our research contributes to tackle such topics and, finally, we summarize the structure of this thesis.

1.1 Motivation

In addition to collections from customers and payments to vendors, cash managers handle bank account balances, short-term loans and investments on a daily basis. Although necessarily influenced by long-term decisions, the corporate cash management problem is mainly a short-term financing problem. The focus is placed on how companies manage their cash and other liquid assets like interest bearing accounts or marketable securities. However, uncertainty about the near future is an important issue for cash managers. Cash flow fore-

casts play then a key role in cash management as a way to reduce uncertainty. As a result, it is common practice to predict future cash flows in an attempt to maintain average cash balances sufficient to face payments. However, we observe that there is always some degree of inaccuracy in cash flow forecasts and, consequently, cash management involves not only costs but also the risk associated to short-term financial operations as the target goals to optimize.

It was John Maynard Keynes (1936) who, in *The General Theory of Employment*, first identified three motives for holding cash:

1. The **transaction motive**, which is the need for cash for the current transaction of personal and business exchanges.
2. The **precautionary motive**, which is the desire for security as to the future cash equivalent of a certain proportion of total resources to act as a financial reserve.
3. The **speculative motive** or the object of securing profit from knowing better than the market what the future will bring forth. The goal is to take advantage of future investment opportunities.

From the above motives, one can infer what kind of problems cash managers face and what kind of solutions they should seek to do their job. On the one hand, they have to pay for the bills but, at the same time, they have to collect payments from customers. The result from this initial task is the cash balance obtained as previous balances plus the net cash flow which, in turn, is determined by the difference between collections and disbursements in a given period. The higher this difference the better. However, the amount of cash that companies hold in cash is an important variable. Cash is the main resource of a company and cash managers must use it efficiently due to the cost associated to obtain cash either from banks or from shareholders.

Cash is the life blood of a company and, from an operational point of view, has different origins and ends. Incoming cash comes from customers, banks, shareholders or even from public institutions in the form of grants or subsidies. Outgoing cash goes to vendors, employees and, again, banks, public institutions and shareholders in the form of taxes and dividends respectively. Collections, disbursements, investments, control, efficiency and costs involved are some of areas where many decisions have to be made by cash managers on a daily basis. In this sense, information technology and all the related tools and techniques can be of great help.

Intuitively, the cash management problem (CMP) can be viewed as a control problem. Think of a water tank like the one in Figure 1.1. The level of the tank needs to be monitored to keep the water between two bounds, for instance, a low bound and a high bound. To this end, some control actions can be taken to increase or decrease the level of water. Replace water with money and you will be dealing with the typical vision of the cash management problem.




Figure 1.1: Cash management as a control problem.

More precisely, the cash management problem aims to keep the balance between what a company holds in cash and what is placed in alternative investments. For example, consider a company with two bank accounts depicted as numbered circles in Figure 1.2. Account 1 receives payments from customers and it is also used to send payments to suppliers. Both inflows and outflows are summarized through the net cash flow f_1 . Account 2 represents the amount of alternative investments available to be converted into cash through transaction x_1 when needed. In addition, idle cash balances from account 1 can be allocated in account 2 for a profit through transaction x_2 . The sequence of control actions deployed over a period of time, for instance, the next five working days,

is called a policy and it is charged with a given cost. As a result, the challenge for cash managers is to find a policy that optimizes some objective function.




Figure 1.2: The cash management problem.

In what follows, we focus on the main dimensions of the cash management problem as a framework to identify open research questions.

1.1.1 Cash management models

A first attempt to face the cash management problem can be derived from its main traditional tasks (Myers and Brealey, 2003; Ross, Westerfield, and Jordan, 2002):

1. How is cash to be collected and paid out?
2. How much cash should a firm hold?

The first of these tasks deals with two commonsense goals: speeding up collections and control disbursements, or more plainly, "collect early and pay late" (Ross, Westerfield, and Jordan, 2002). From a cash manager point of view, the primary concern is speeding up timing of receipts into available funds and controlling the release and timing of disbursements. In the second one, cash managers have a choice between holding cash and investing it in short-term liquid assets. From a cost perspective, they should minimize borrowing costs of working capital needs and, at the same time, maximize yields from short term investments of temporary cash excess. They should keep minimum cash balances in order to meet future cash requirements without incurring in unexpected costs. If the timing and amount of future cash flow was known, no cash buffer would be necessary. However, this situation is far from reality.

In an attempt to solve the CMP, 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) and Srinivasan and Kim (1986)

and Costa Moraes, Nagano, and Sobreiro (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. For instance, 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). More recently, Gormley and Meade (2007) claimed the utility of cash flow forecasts in the cash management problem.

All previous models are based on assuming a set of bounds to control cash balances. However, notice that the ultimate goal of the cash management problem is not to find a set of bounds that defines a policy, but the policy itself. Consequently, we argue that a new class of cash management models can be explored by imposing no restriction on the form of control policy.

1.1.2 Cash flow process

Cash flow statistical characterization is also a key issue in understanding corporate cash management. Separation between inflows and outflows, or receipts and disbursements, is the basic breaking down, but a more detailed separation can be of help when trying to extract patterns from data. In this sense, Stone and Miller (1981) and Stone and Miller (1987) suggest the utility of problem structuring, or breaking down a problem in different subproblems, to appropriately handle cash flow forecasting as a key task in cash management. They proposed the separation of cash flows in two streams: major cash flows and non-major cash flows. Major cash flows are defined as flows that are not generalized from past history but are easy to handle in daily forecasting since most are known in both timing and amount. Transfers between bank accounts, payroll, taxes and loan payments are examples of major flows. Non-major cash flows are all other transactions usually related to customers and vendors. Note that the separation between customers and vendors produces an implicit division between inflows and outflows. Summarizing, a detailed representation of the cash flow process is a mandatory step in cash management.

In addition, common assumptions on the statistical properties of cash flows include: (i) normality, meaning that its values are centered around the average following a Gaussian distribution; (ii) absence of correlation, meaning that its values are not correlated with each other; (iii) and stationarity, meaning

that its mean and variance are constant with time. However, little empirical evidence on the statistical properties of cash flow have been provided with the exception of Mullins and Homonoff (1976), Emery (1981), and Pindado and Vico (1996). Negative normality tests were reported in Mullins and Homonoff (1976) for a manufacturing company. Later on, Emery (1981) reported normally distributed cash flow, after data transformation, and Pindado and Vico (1996) provided negative normality and independence results on 36 companies but considering daily cash flow for one month. Both Emery (1981) and Pindado and Vico (1996) reported the influence of day-of-week effect on cash flows in line with the works of Stone and Wood (1977), Stone and Miller (1981), Miller and Stone (1985), and Stone and Miller (1987).

1.1.3 Costs in cash management

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, Westerfield, and Jordan, 2002). Then, the cash management problem can be viewed as a trade-off between holding and transaction costs as shown in Figure 1.3. 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 optimize.

1.1.4 Desired objectives

In the cash management literature, 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, analyzing additional objectives derived from a cash management model is necessary if a full understanding of the problem is meant to be achieved. An interesting additional goal in cash management is

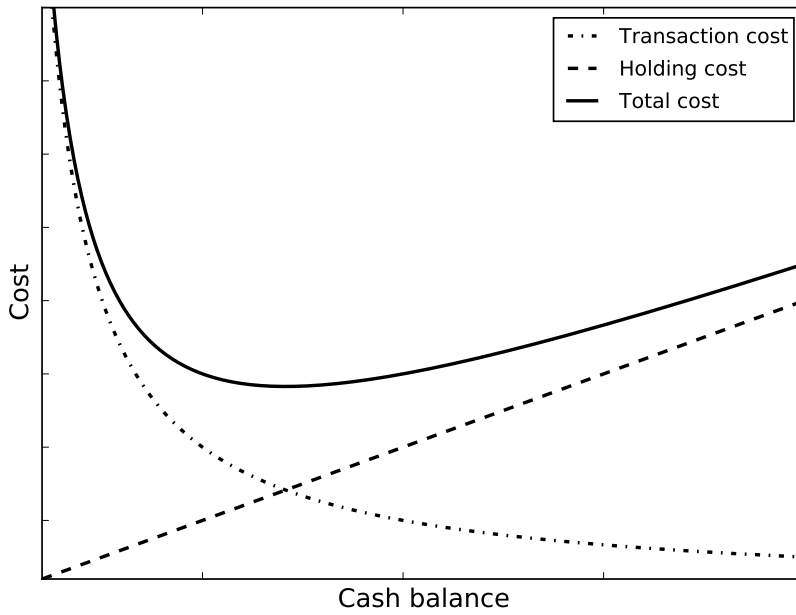


Figure 1.3: Cost of holding and transferring cash.

the risk of policies. Risk 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, Qi, and Hirschberger, 2007). Furthermore, risk preferences are also an important issue for decision-makers. 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 consider both cost and risk (and possibly other goals) by incorporating the particular preferences of cash managers remains a rather unexplored problem.

Intuitively, risk is associated to uncertainty and possible loss and, to some extent, managers can choose the risk that a business takes (Brealey and Myers, 2003). Consequently, the reduction of risk makes financial planning easier, but one of the main problems cash managers face everyday is that they do not exactly know future cash flows. This fact introduces much uncertainty in their daily work. High cash buffers could be a solution, but companies would be

incurring in high opportunity costs by leaving idle an amount of money that otherwise could be generating a return by investing it. In the particular context of cash management, uncertainty is introduced by the particular variability of future cash flows. Risk management is then an important task in decision-making, and since different cash strategies entail different degrees of risk, a quantitative approach to measure risk is required.

1.1.5 Solving the cash management problem

Cash management poses a general optimization problem, namely, determining the policy that optimizes some objective function. Within a single objective framework, one suitable method to solve this problem is dynamic programming, which was initially proposed by Eppen and Fama (1969) and Neave (1970), and more recently followed by Penttinen (1991), Chen and Simchi-Levi (2009), and Melo and Bilich (2013). On the other hand, compromise programming (CP) and goal programming (GP) (Zeleny, 1982; Yu, 2013; Ballestero and Romero, 1998; Ballestero and Pla-Santamaria, 2004; Bravo, Ballestero, and Pla-Santamaria, 2012; Pla-Santamaria and Bravo, 2013) are possible approaches to deal with multiple objectives. When the CMP is formulated as a linear/quadratic program, we can benefit from state-of-the-art mathematical programming solvers to obtain optimal solutions. Alternative approaches such as Monte Carlo methods or meta-heuristics may provide a sufficiently good solution when non-linear objectives or constraints are considered. In this case, only approximate solutions can be achieved. Monte Carlo methods (Glasserman, 2003) are based on the law of large numbers which ensures that estimates derived from a random sample converge to the real value as the size of the sample increases. On the other hand, meta-heuristics such as evolutionary algorithms (Gormley and Meade, 2007; Costa Moraes and Nagano, 2014) deploy a given search strategy over a large set of feasible solutions.

1.1.6 Managing multiple bank accounts

In the cash management literature, cash management systems with multiple bank accounts have received little attention from the research community with the exception of Baccarin (2009). Indeed, there is a lack of cash management models that are able to handle multiple bank accounts. However, cash management systems with multiple bank accounts are the rule rather than the exception in most firms. Then, a formal framework to both design and analyze such systems is required.

1.1.7 *Integrating machine learning and optimization*

In an attempt to reduce risk and uncertainty, forecasting deserves the attention of cash managers. Experience indicates that the success of cash management is closely related to reliable forecasts of future cash flows (Stone and Miller, 1981). Planning and forecasting are required when there is a need to know when an event will occur. Because of that, predicting cash flow is one of the most useful tools for cash management. Any effort done in the line of reducing uncertainty should be rewarded.

From an operational perspective, cash flow forecasting is usually performed on a daily basis. Weekly and monthly predictions can be of help as well, but cash management is typically associated to short term decision-making, mainly decisions to be made daily. Nonetheless, the great majority of corporate cash managers report dissatisfaction with their efforts at daily cash forecasting (Miller and Stone, 1985). Several possible explanations can be argued to be the reason for these failures (Stone and Miller, 1987): major flow separation, component identification, information system support and pattern resolution. The uncertain nature of industrial markets and the lack of materials requirements planning systems can be additional reasons.

There is no doubt that the main resource available to cash flow forecasting is data. The size of business data bases may continue to increase on a daily basis as a result of decision-making or transaction recording. Data-driven decision-making refers to the practice of basing decisions on the analysis of data, rather than purely on intuition. In the context of business and finance, decision-making is performed on a daily basis and it is essentially based on data. Consequently, there is a need to benefit from data and, at the same time, there is also a need for a sound strategy to incorporate data into the decision-making process.

To this end, machine learning or, in a broader sense, artificial intelligence represents a suitable alternative to integrate forecasts in the CMP. Synthetically, machine learning is concerned with the question of how to construct computer programs that automatically improve with experience (Mitchell et al., 1997). More formally, machine learning is the field of artificial intelligence that deals with making computers modify or adapt their actions so that these actions get more accurate, where accuracy is measured by how well the chosen actions reflect the correct ones (Marsland, 2009). When machine learning techniques are utilized in the development of alternatives for the decision-maker, the resulting systems are referred to as intelligent decision support systems, and several examples can be found in Doumpos and Grigoroudis (2013). However,

there is a lack of research on the link between machine learning and the cash management problem.

1.2 Challenges

Next, we focus on a number of challenges derived from the open research questions identified above that we tackle in this thesis. To this end, consider a typical situation in which one or more bank **accounts** are characterized by a cash balance which starts in an initial value and fluctuates according to a particular **cash flow process** in absence of control actions. At any time, cash managers can take control actions by increasing/decreasing the cash balance, paying a cost defined by some **cost function**. Given an initial condition and a cash flow process, a cash management **model** defines the sequence of control actions, namely, a policy, according to some predefined rules. The ultimate goal of the cash management problem is to find the policy that optimizes some objective function with one or more **objectives** over a time horizon. Any procedure or algorithm used to find the best policy is called a cash management **solver**. These are the main dimensions of the CMP. We next consider a number of challenges worth tackling.

Decision-making in cash flow management has been supported by different models designed to establish a policy, i.e., a number control actions by increasing/decreasing the cash balance. The main implication for cash management derived from the assumption of a distribution modeling a cash flow process is its predictability. As mentioned in Section 1.1.2, we must explore empirical cash flow data sets from which some patterns may be extracted. Seasonality and other alternative features from cash flow can be used to increase predictive accuracy of any forecasting technique used in cash management. Thus, in this thesis we consider the next question:

***Question 1: The savings hypothesis.** Can cash flow predictive accuracy achieve cost savings in the cash management problem?*

Furthermore, common statistical assumptions on daily cash flow include normality, stationarity, absence of correlation and also linearity, meaning that cash flows are proportional to a set of explanatory variables. However, some results questioning these assumptions detailed in Section 1.1.2 lead us to face in this thesis the following question:

Question 2: Cash flow assumptions. *Are common statistical assumptions of daily cash flow supported by recent empirical data?*

In addition, data transformation is usually considered as a necessary previous step by time-series analysis techniques to achieve normality and linearity as required as common assumptions by the cash management literature. In this thesis, we face the underlying question:

Question 3: Data transformations. *Is it always possible to achieve a Gaussian, noise-free and linear time-series through data transformations?*

If the answer to this question was negative, alternative approaches such as non-linear time-series forecasting could be considered as a key tool in cash management. On the other hand, it is easy to understand that most decision-making problems must take into account multiple objectives. This fact is particularly true in cash management where several, but often conflicting, goals are pursued by decision-makers with possibly different preferences for each goal. Usually, different goals are often translated into economic terms or introduced in the optimization problem as additional constraints. However, since risk is inherent to any decision-making process that involves a certain degree of uncertainty, as it is the case of the CMP, in this thesis we formulate the next question:

Question 4: Risk as an additional goal. *Can we incorporate risk as an additional goal to the cash management problem?*

Under the current situation of time-varying financial circumstances, cash managers, may be interested in identifying the best compromise policies in terms of cost and risk that are also robust to cash flow regime changes. As a result we pose the following question:

Question 5: Robust models. *Can we provide a robust counterpart for any cash management model?*

A closely related topic to the previous research question is the concept of operating condition, which we use in this context to refer to any factor that may influence the performance of a model. Being able to analyze the tradeoff between cost and risk for different models and operating conditions can be of help for cash managers. In this thesis, we also consider the following open research question:

Question 6: Operating conditions. *Under what circumstances or operating conditions a model is better than another?*

The cash management models mentioned in Section 1.1.1 are based on a set of control limits. They aim to obtain the best set of bounds that minimizes holding and transaction costs. However, the ultimate goal of the cash management problem is not to find the best set of bounds, but the best sequence of control actions. We argue that the constraints imposed by these models in the form of bounds reduce the space of possible control actions. In this thesis, we aim to answer the following question:

Question 7: Boundless models. *Are control bounds really necessary in cash management?*

An additional question derives from the methods used to derive policies as mentioned in Section 1.1.5. How are these policies determined in practice? Alternative approaches include methods providing optimal solutions such as dynamic programming and approximate methods such as genetic algorithms. The inclusion of additional goals, possibly expressed as non-linear functions, lead us to consider in this thesis the question:

Question 8: Optimal solutions. *Can we obtain optimal solutions for the multiobjective cash management problem?*

Cash management systems with multiple bank accounts present particular characteristics that require a different approach. Cash managers aim to obtain a feasible policy that minimizes some objective function. When dealing with multiple bank accounts, a feasible policy requires the allocation of available

funds in a set of accounts. Furthermore, transactions between any pair of bank accounts are permitted at a certain cost provided that there is enough balance in the source account. As pointed out in Section 1.1.6, the lack of a formalization of the CMP for multiple bank accounts forces us to consider in this thesis the question:

Question 9: Multiple bank accounts. *Can we derive optimal policies for cash management systems with multiple bank accounts?*

Selecting the best policy to keep the balance between what a company holds in cash and what is placed in alternative investments in cash management systems with multiple bank accounts is by no means straightforward. Automating decision-making in cash management is also a challenge that is worth tackling. This leads to the final research question that we pose in this thesis:

Question 10: Software for cash management. *Can we automate decision-making in cash management through the use of software?*

Summarizing, in Section 1.1 we established the framework for the study of the CMP based on six dimensions: models, cash flow process, cost functions, objectives, solvers and number of bank accounts considered. This framework allows us to formulate a number of important research questions. The answer to these questions becomes the main contributions of this thesis that we next describe.

1.3 Contributions

In this section, we enumerate the main contributions of this thesis and how they are linked to the aforementioned research questions within the framework of the six dimensions of the CMP discussed in Section 1.1. By means of this six-dimensional framework, we map each contribution to a family of cash management problems defined by the particular values of each dimension.

1.3.1 *The savings hypothesis*

Our first contribution relies on bound-based cash management models that accept empirical cash flow process, considering linear cost functions, a single objective, namely, cost, and that are solved by Monte Carlo methods (Glasserman, 2003) for only one bank account as shown in Figure 1.4.

The utility of cash flow forecasts has received little attention in the literature with the exception of Stone (1972) and Gormley and Meade (2007). A measure of quality of any forecasting technique is its predictive accuracy (Makridakis, Wheelwright, and Hyndman, 2008) and, under an economic perspective, predictive accuracy must be mapped to estimated cost savings. Surprisingly, it is unknown whether even small improvements in cash flow predictive accuracy may lead to savings that could perhaps amount to millions of euros in total. We claim this step as a mandatory one, specially when improving forecasting accuracy may be correlated with cost savings, that we here call the savings hypothesis. As a result, our first contribution is a test to accept or reject the savings hypothesis as an answer to Question 1. More precisely, we provide a procedure to estimate how much companies can save by improving predictive models and, consequently, the cost of not predicting, i.e., the missed savings minus the cost of implementing the model.

Moreover, we present and compare different forecasting methods including linear and non-linear models. Using two real data sets from companies in the textile industry in Spain, as a proof of concept: (i) we show empirically that forecasting accuracy is highly correlated with savings in cash management and, thus, a comparison in terms of accuracy and savings between different forecasting models is performed; (ii) we argue that the effect of forecasting accuracy on cash management can be estimated in advance and, thus, we propose a new methodology for estimating this effect. As a result, cash managers are then able to know if extra effort in improving accuracy is worthwhile.

1.3.2 *Cash flow assumptions*

In our second and third contribution, we focus only on the empirical properties of the cash flow process as shown in Figure 1.5. First, in an attempt to answer Question 2, we provide recent evidence on the statistical properties of 54 real cash flow data sets from small and medium companies in Spain, with annual revenue up to 10 million euro each. To the best of our knowledge, this is the most comprehensive empirical study on daily cash flow so far. We base this statement on the range of statistical properties considered, and on both the

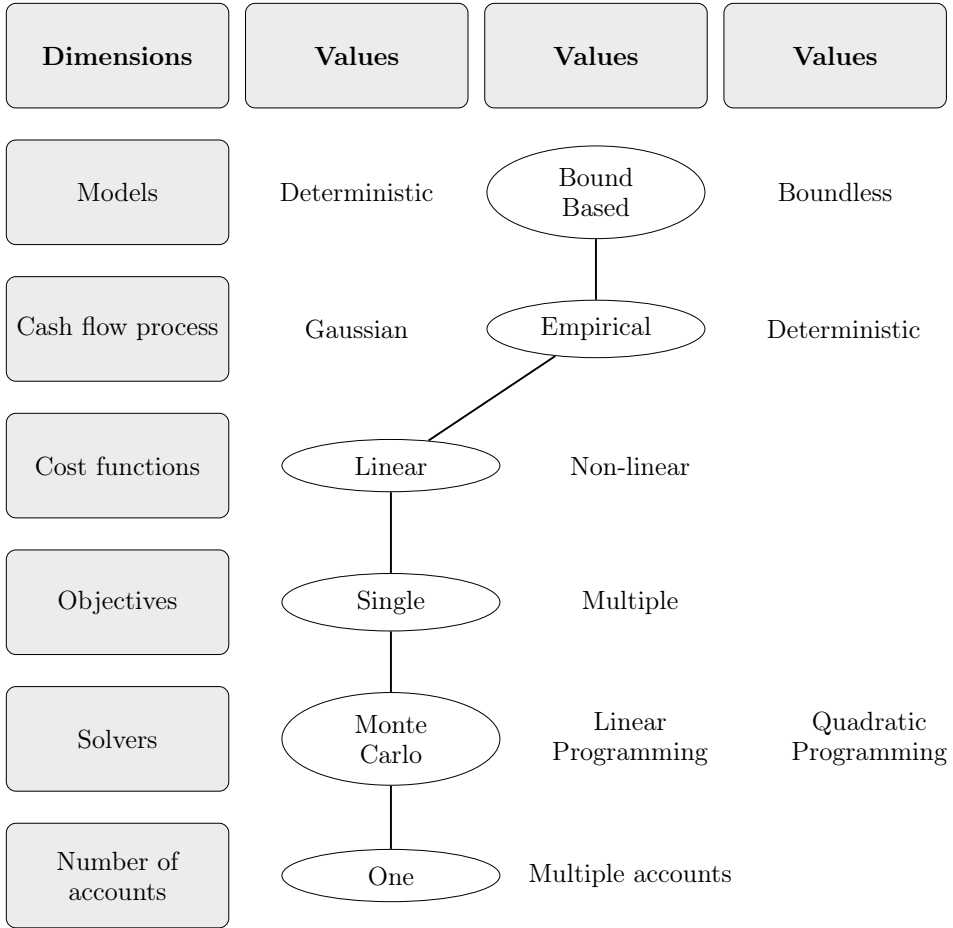


Figure 1.4: Contribution 1. The savings hypothesis.

number and the length of the data sets, which amounts to 58,005 observations in total.

Contrarily to what it is assumed in the literature, our results show the unlikely occurrence of normality, stationarity and absence of correlation, and this fact leads to consider alternative forecasting linear and non-linear forecasting models. Consequently, we also consider linearity, meaning that cash flow is proportional either to another variable or to a combination of explanatory variables. In order to determine whether a non-linear model is preferable to a linear one, we propose a new cross-validated test for non-linearity that is able

to consider any non-linear functional form. Our test is based on time-series cross validation (Hyndman and Athanasopoulos, 2013) to avoid the potential problem of overfitting the data.

1.3.3 Data transformations

As a solution to the previous problems of lack of normality and linearity, data transformation is usually considered as a general technique by state-of-the-art time-series forecasting techniques. Therefore, we face the underlying additional Question 3, namely, whether it is always possible to achieve a Gaussian, noise-free and linear time-series through data transformations. We rely both on common statistical tests and on our novel non-linearity test to answer this question and we find that: (i) outlier treatment and Box-Cox transformation is not enough to achieve normality; (ii) outlier treatment produce mixed results in terms of noise reduction and information loss; (iii) linear models obtain no benefit from outlier treatment; (iv) a Box-Cox transformation produces similar results in our cross-validated non-linearity test. These results suggest that non-linear models represent a more realistic alternative for time-series forecasting.

1.3.4 Risk as an additional goal

Our contributions 4, 5, and 6 focus also on bound-based models using empirical cash flow data and linear cost functions, but extending the number of objectives to consider not only cost but also risk in the search for the best policies and a single bank account as shown in Figure 1.6.

One of the most common features of most decision-making problems is the optimization of multiple objectives. However, current research on cash management does not approach the CMP from a multiobjective perspective. In this thesis, we propose a multiobjective approach to the CMP. Since risk is inherent to any decision-making process that involves a certain degree of uncertainty, we incorporate risk analysis in the CMP providing an answer to Question 4. We further provide theoretical results on the characterization of compromise solutions, and also a procedure to deal with discrete efficient frontiers and uncertain risk preferences.

Dimensions	Values	Values	Values
Models	Deterministic	Bound Based	Boundless
Cash flow process	Gaussian	Empirical	Deterministic
Cost functions	Linear	Non-linear	
Objectives	Single	Multiple	
Solvers	Monte Carlo	Linear Programming	Quadratic Programming
Number of accounts	One	Multiple accounts	

Figure 1.5: Contributions 2 and 3. Empirical evidence on daily cash flow.

1.3.5 Robust models

In addition, there is a need for providing support to cash managers to test policies under a changing financial context. Then, we introduce a robustness index based on a data-driven procedure using distances in the cost-risk space. This index helps identify the best compromise policies in terms of cost and risk, which are also robust to cash flow regime changes, hence answering Question 5.

1.3.6 *Operating conditions*

The concept of operating condition is closely related to changing contexts. Therefore, the analysis of cost and risk for different models and operating conditions is at the core of decision support systems for cash managers. Consequently, we also introduce the use of ROC (Receiver Operating Characteristic) analysis and loss curves in cash management. We allow model comparison through the use of graphical tools to: (i) show models in the cost-risk space; (ii) choose models according to risk preferences; (iii) derive cost-risk curves for different operating conditions.

1.3.7 *Boundless models*

Our seventh and eighth contribution are new cash management models that use forecasts derived from empirical cash flow processes to determine optimal multiobjective policies with linear cost functions by relying on linear and quadratic programming as shown in Figure 1.7.

In addition to the lack of risk analysis in state-of-the-art cash management models, most of them have a common feature: they are based on setting control limits or bounds. Within the framework of Bound-Based Models, cash balance is allowed to wander around between some bounds, usually a high bound and a low bound. When any of these bounds is reached, a control action is made to restore the balance to some target level. At this point, we face Question 7 by testing if these bounds are really necessary to derive optimal policies. In all previous models, cash managers have to determine the bounds which minimize transaction and holding costs. However, the ultimate goal of the cash management problem is not to find the best set of bounds, but the best sequence of control actions.

In this thesis, we propose a new class of cash management models without bounds in order to enlarge the search space for alternative policies. Furthermore, we show that usual constraints imposed by these models, in the form of bounds, are not necessary to derive optimal policies. As a result, we first linearize the CMP in order to guarantee the optimality of solutions. Then, we use this approach to reformulate the Gormley and Meade (2007) model as a linear program ensuring optimality. Finally, we propose an equivalent but simpler model that we call the Boundless Model (BM) in which no constraint or restriction is imposed on the form of the policy.

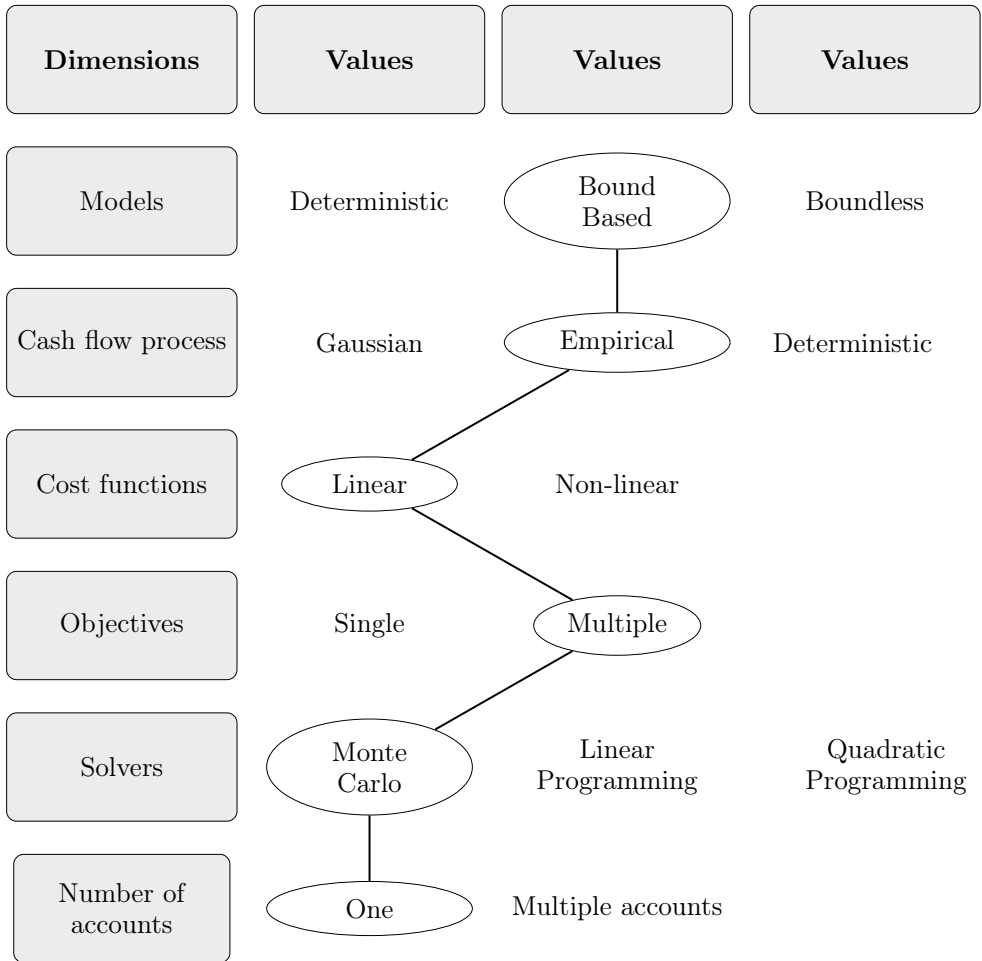


Figure 1.6: Contributions 4, 5, and 6. Multiobjective cash management.

1.3.8 Optimal solutions

In an attempt to answer Question 8, we propose two different solvers for cash management problems based on linear programming (LP) and quadratic programming (QP). Linear and quadratic programming counterparts of compromise programming models result in an automated and optimal decision making technique when preferences and the extreme values of both cost and risk objectives can be reasonably estimated by cash managers.

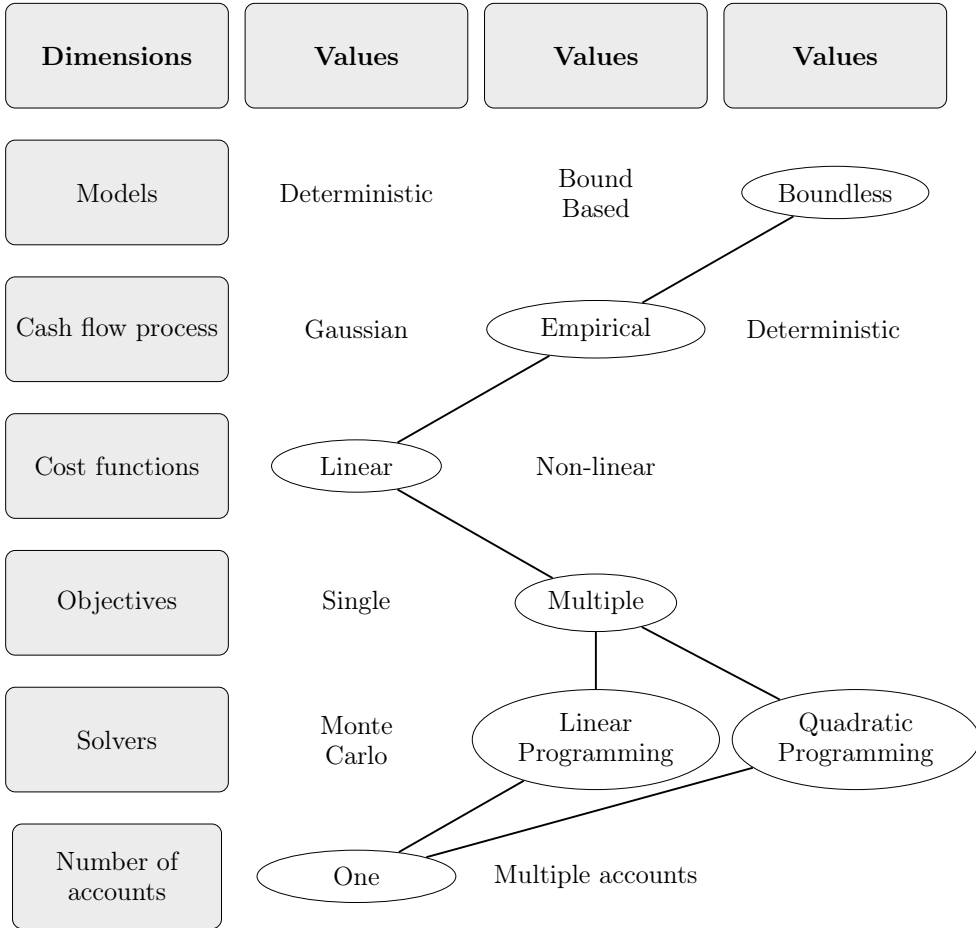


Figure 1.7: Contributions 7 and 8. Boundless models for multiobjective cash management.

1.3.9 Multiple bank accounts

Figure 1.8 shows graphically our ninth and tenth contribution as a method to automate cash management systems with multiple bank accounts relying on linear programming for a single objective with linear cost functions. Nonetheless, it is important to highlight that our model can be readily extended to consider multiple objectives.

Cash managers usually deal with multiple banks to receive payments from customers and to send payments to suppliers, employees and other creditors.

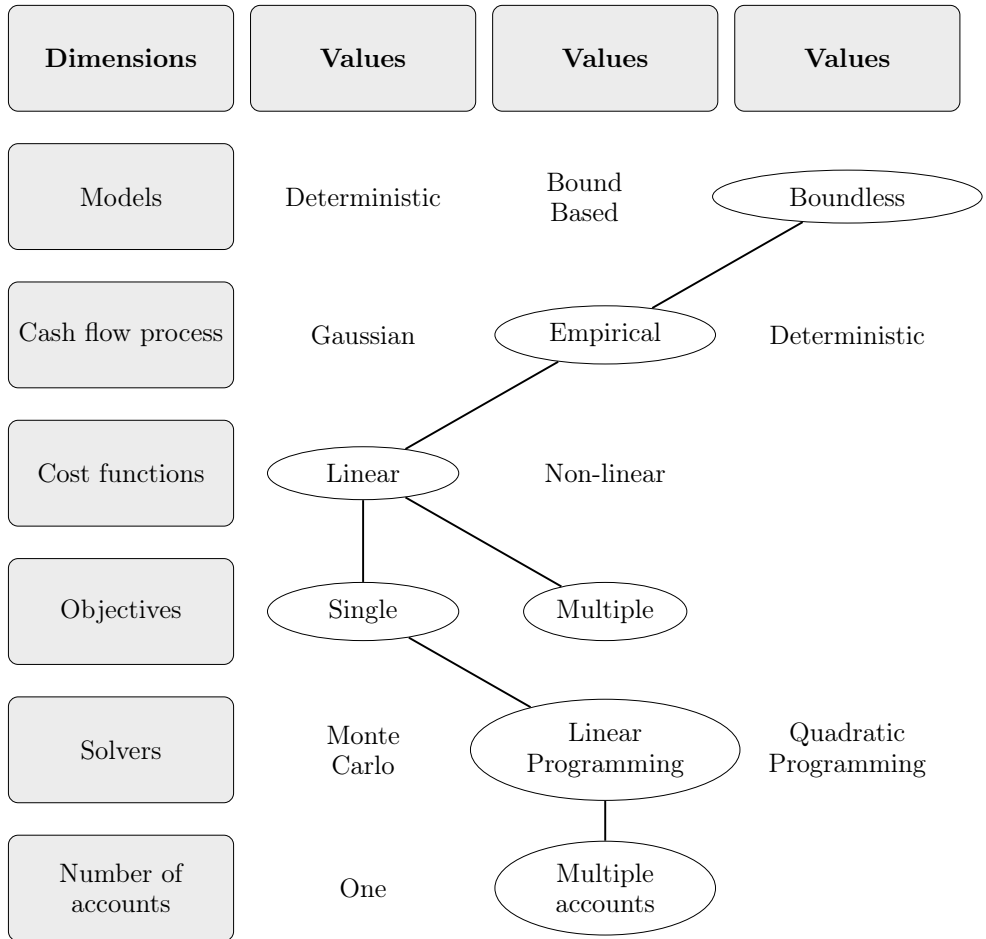


Figure 1.8: Contribution 9 and 10. Automating cash management systems with multiple bank accounts.

Hence, a cash management system can be viewed as a set of bank accounts and their relationship. Operating a cash management system implies a number of transactions between accounts to maintain the system in a state of equilibrium, meaning that there exist enough cash balance to face payments and avoid an overdraft. In the cash management literature, systems with multiple bank accounts have received little attention of the research community with the exception of Baccarin (2009). However, there is a lack of cash management models that are able to handle multiple bank accounts. Hence, our ninth

contribution answers Question 9 by means of a formalization of the CMP with multiple bank accounts. In order to handle the uncertainty introduced by cash flow forecasts, we provide robust solutions through linear programming.

1.3.10 Software for cash management

In an attempt to reduce the gap between theory and practice, we also provide a software tool called PyCaMa to define, analyze and solve the CMP with multiple bank accounts as an answer to research Question 10. More precisely, we provide a software library in Python that can be used for cash management systems with multiple bank accounts and also for a single bank account. Furthermore, PyCaMa can also solve the CMP considering both cost and risk of the policies.

1.3.11 Integrating contributions

As a final contribution of this thesis, Figure 1.9 represents a logical procedure to be followed by cash managers to improve daily decision-making. Cash managers provided with both a cash flow data set and a cash management model are now able to perform a savings test to know if predictive accuracy may lead to enough savings. If these savings are worthwhile, then a detailed empirical study on cash flow data will necessary improve the understanding of the cash management problem. As a result, a number of suitable forecasters can be used as a key input to: (i) follow a multiobjective approach in which not only cost but also risk is considered; (ii) use a new family of boundless models that are able to consider a wider input space than state-of-the-art cash management models; (iii) solve the cash management problem within a more realistic framework with multiple bank accounts.

1.4 Structure

This thesis is organized in three parts, twelve chapters, including this introductory chapter, and one appendix with useful terminology and acronyms:

Chapter 2 reviews cash management related work focusing on models proposed to determine optimal policies. We discuss the main contributions of each model highlighting common assumptions and those characteristics closely related to our six dimensions framework established in Section 1.1 for cash management analysis. We finally provide a comparative study of cash management models.

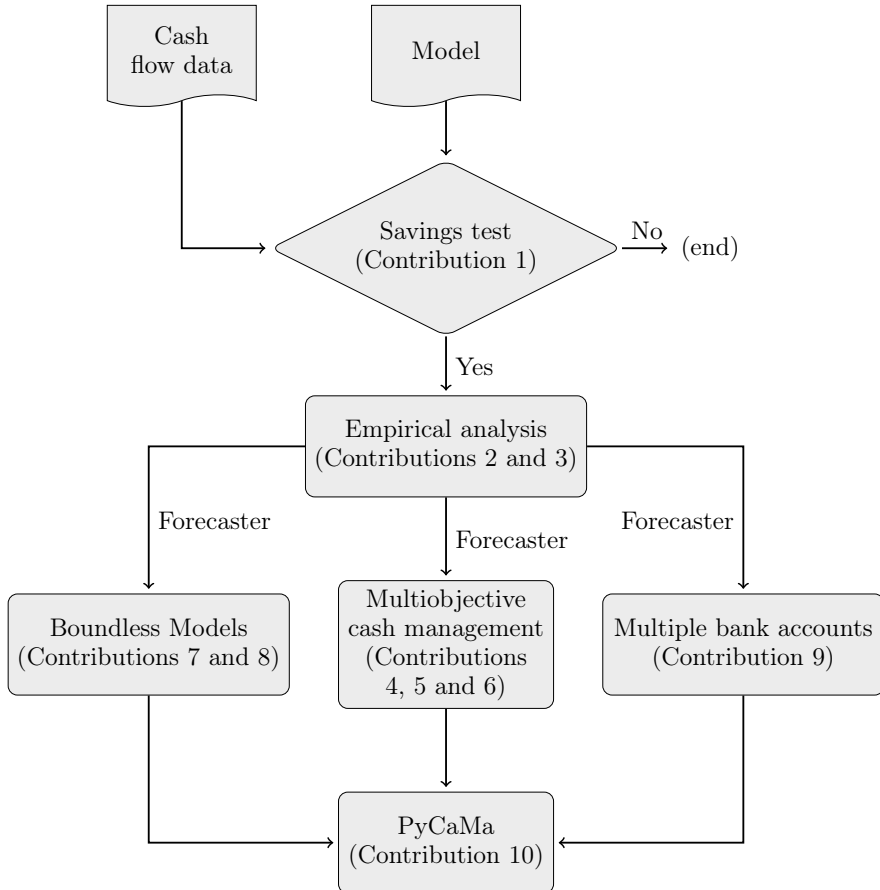


Figure 1.9: Integrating the contributions of this thesis.

Chapter 3 presents the savings hypothesis contribution as a proof of concept to assess the utility of forecasting in cash management. The underlying idea behind the savings hypothesis is that cash managers can estimate in advance how much savings can be achieved using any particular cash management model and any cash flow process within the framework of the current cost structure. Thus, the notion of simulation is of central importance. We first show that alternative forecasting techniques are able to produce better predictions than a trivial forecaster. Then, we demonstrate that better forecasts lead to cost savings, and we finally provide a simulation methodology to estimate savings derived from predictive accuracy. If the result of this proof of concept is positive, i.e., the estimated savings are worthy, then cash managers may consider

better forecasting techniques and alternative cash management models. The material contained in this chapter has been published in:

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

Chapter 4 provides a comprehensive empirical study on daily cash flow based on 54 real cash flow data sets as a contribution regarding the common cash flow assumptions in the cash management literature. We propose a novel cross-validated test for time-series non-linearity and we also analyze the impact of data transformations. The material contained in this chapter has been published in:

- Salas-Molina, F., Rodríguez-Aguilar, J. A., Serrà, J. Empirical analysis of daily cash flow time-series and its implications for forecasting. (2016). [arXiv:1611.04941v3](https://arxiv.org/abs/1611.04941v3).

Chapter 5 approaches the cash management problem from a multiobjective perspective. We motivate the use of risk as an additional goal in cash management and we further describe our multiobjective approach including a method to solve the multiobjective CMP. The material contained in this chapter has been published in:

- Salas-Molina, F., Pla-Santamaria, D., Rodríguez-Aguilar, J. A. (2016). A multiobjective approach to the cash management problem. *Annals of Operations Research*.

Chapter 6 aids cash managers to select the best policies under a changing context by incorporating robustness as an additional goal to cost and risk within a multiobjective framework. We propose to calculate robustness as a multiple criteria distance index that helps identify the robust policies that are also efficient in terms of cost and risk. Using this multiobjective index, cash managers can transform a cash management model in a more robust one. The material contained in this chapter has been published in:

- Salas-Molina, F., Pla-Santamaria, D., Rodríguez-Aguilar, J. A. (2017). On the use of multiple criteria distance indexes to find robust cash management policies. *Information Systems and Operational Research*.

Chapter 7 adapts Receiver Operating Characteristic (ROC) analysis to cash management. We provide a range of visual and quantitative tools for showing

model performance; choosing models; and assessing the impact of operating conditions on model performance. The material contained in this chapter has been published in:

- Salas-Molina, F., Rodríguez-Aguilar, J. A., Díaz-García, P. (2017). Selecting cash management models from a multiojective perspective. *Annals of Operations Research*.

Chapter 8 proposes a new boundless cash management model based on forecasts that guarantees the optimality of solutions in a multiobjective framework in which cost and risk are goals to be minimized. We also offer an extensive evaluation of the impact of forecasting errors using real data sets from 54 different companies.

Chapter 9 formulates alternative compromise programming models for cash management as a method to minimize the sum of weighted distances to an ideal point where both cost and risk are minimum. We discuss the pros and cons of different measures of risk. Furthermore, we propose and compare three different solvers to cover a wide range of possible situations such as Monte Carlo methods, linear programming and quadratic programming.

Chapter 10 develops as a sound formal framework to automate cash management systems with multiple bank accounts. In this chapter, we approach the cash management problem from a multidimensional perspective to provide a formal specification of cash management systems with multiple bank accounts. We provide a methodological framework that allows to define multidimensional cash management systems and to obtain a solution by using state-of-the-art mathematical programming solvers.

Chapter 11 introduces PyCaMa, a Python module for multiobjective cash management based on linear programming that allows to derive optimal policies for cash management with multiple bank accounts in terms of both cost and risk. The material contained in this chapter has been published in:

- Salas-Molina, F., Rodríguez-Aguilar, J. A., Díaz-García, P. (2017). PyCaMa: Python for cash management. [arXiv:1702.05005](https://arxiv.org/abs/1702.05005)

Chapter 12 concludes summarizing the main contributions of this thesis and highlighting future lines of action.

Chapter 2

Background and related research

In this chapter, we summarize and analyze the research appearing in the literature which is relevant to this thesis. We mainly focus on the alternative solutions, in the form of cash management models, that have been proposed to the cash management problem. Remaining related works on other topics will be covered in their respective chapter. In order to put these works into perspective, we dedicate the first part of this chapter to define the cash management problem. Next, we introduce models and we finally present a comparative analysis highlighting the main characteristics of these models according to the framework provided in Chapter 1.

2.1 Definition of the cash management problem

The Cash Management Problem (CMP) was first addressed from an inventory point of view by Baumol (1952). Cash was then viewed as a stock, as a buffer, like an inventory of wheat or bolts. Holding cash has a cost because of being idle but, at the same time, transferring idle money to alternative investments is also costly. How much money should companies keep in cash to operate efficiently? This question defines the CMP. However, additional information about the cash flows' properties and both holding and transaction costs are required. Consequently, we can synthetically define the CMP by means of

a statement containing all the required inputs to determine the amount of cash balance that optimizes a given objective function, usually holding and transaction costs. Probably, the most detailed definition of the CMP was given by Eppen and Fama (1969), as follows:

"To meet its day-to-day transactions requirements a firm keeps a cash balance, either in the form of cash on hand or as a bank deposit. In controlling the level of the cash balance, the goal of the firm is to choose an operating policy or decision rule which minimizes discounted costs over some horizon period. An incentive to keep the cash balance low arises from the fact that each unit of positive cash leads to holding cost since cash has alternative uses (i.e., as dividends or as investment in earning assets). On the other hand, if the cash balance goes below zero (or some other finite level), penalty costs will be incurred as a result of delay in meeting demands for cash. Finally, the firm can increase its cash balance either by raising new capital or by selling some earning asset, and the cash balance can be reduced by paying dividends or investing in earning assets. Adjusting the cash balance will in either case involve transactions or transfer costs. To complicate the problem, day-to-day inflows and outflows of cash are, at least to some extent, unpredictable."

The first approach by Baumol (1952) was deterministic, meaning that the basic premise was that the demand for money was assumed to be perfectly known. Later on, Miller and Orr (1966) approached the problem from a stochastic perspective by considering a symmetric Bernoulli process in which both inflows and outflows were exactly of the same size and had probability 1/2. These two approaches were considered extreme cases by Stone (1972), who postulated that real-world cash flows were neither completely certain as in Baumol (1952), nor completely unpredictable as in Miller and Orr (1966). On the contrary, Stone considered that most companies can forecast their cash flows although there is always some uncertainty in this task.

The main goal of cash management was briefly stated in Gregory (1976). Cash flow management is concerned with the efficient use of the company's cash and short-term investments. In this review the author pointed out that *under stable economic conditions, this is a matter of deciding when to transfer assets and*

how much. As a result, one may consider that the development of models to answer these questions is a task worth tackling.

Note that from the seminal work by Baumol (1952), the CMP was viewed as an inventory control problem which maps a number of inputs to an output solution. However, cash management covers a wider range of administrative tasks which deserve to be mentioned, though they are beyond the scope of this thesis. A first definition of the corporate cash management can be found in Fabozzi and Masonson (1985): *the actual management of incoming and outgoing cash on a daily basis is managed effectively.* However, this definition has a remarkable problem: the use of the word management, not only once but twice, to define management itself. Common definitions of management break down the term in four tasks or processes: plan, organize, execute and control. Therefore, we next reformulate the previous definition to increase clarity:

cash management tries to effectively plan, organize, execute and control incoming and outgoing cash.

Moreover, Fabozzi and Masonson (1985) extended their definition by enumerating five basic elements of corporate cash management:

1. Collections: speeding up timing of receipts into available funds.
2. Disbursements: controlling the release and timing of payments.
3. Concentration: inexpensively mobilizing funds from outlying banks to one location for their efficient use.
4. Investments: maximizing yield within acceptable levels of risk and maturity.
5. Information and control: externally obtaining accurate, timely data on bank balances, bank deposits, and internally obtaining data for accurate short-term cash forecasts.

In the most general case, cash managers have to deal with collections, disbursements, concentration, investments and information and control. Interestingly, information and control are placed in the last position but, if we carefully look at the problem, we will see that all of them would fail to succeed if they were not based on timely data and if they were not appropriately controlled. In this work, we will particularly focus on information and control.

Another comprehensive definition can be found in Srinivasan and Kim (1986). In this cash management review, the authors first focus on *seeking answers to several questions: What does the cash manager do? What kinds of decisions does the literature suggest the cash manager is involved in?* The answers to these questions contain the concepts of planning and control: the cash management function has the responsibility to mobilize, control and plan the firm's cash resources. Finally they consider four different decision types that fall under these decision processes:

1. Cash position management.
2. Short-term borrowing.
3. Short-term investing.
4. Cash forecasting.

As expected, the definitions by Fabozzi and Masonson (1985) and Srinivasan and Kim (1986) have many common points. It seems clear that cash management is a problem with two directions: one into the company and another one out from the company, inflows and outflows, collections and disbursements. On the other hand, there is a clear interest on controlling cash balances or cash positions in the different bank accounts needed for operational purposes. Therefore, cash management also focuses on allocating balances among bank accounts to avoid overdrafts. Moreover, this cash mobilization can be done by short-term borrowing or by selling short-term assets, such as interest-bearing securities. Excess cash positions can be used for short-term investing. Finally, an important point is also mentioned: forecasting. This thesis focuses on the use of data-driven procedures such as forecasting in cash management with a goal in mind: improving daily decision-making processes.

Finally, a recent review on stochastic cash management models since the 1980s can be found in Costa Moraes, Nagano, and Sobreiro (2015). Although clearly referring back to the seminal works by Baumol (1952) and Miller and Orr (1966), and to the previous reviews by Gregory (1976) and Srinivasan and Kim (1986), new perspectives on cash management are given through the concepts of profitability, liquidity and risk such as any other investment decision.

From the basic characteristics of cash management and its several perspectives given above, a number of cash management models have been proposed in the literature. These models represent alternative ways to determine the ideal amount of available cash to be held by firms that we next review.

2.2 A review of cash management models

In the following sections, we present a chronological review of the most relevant previous works on the cash management problem. We respect the original author's notation and clarify issues on notation when necessary for comparison purposes. After a detailed description, we discuss how each cash management model contributed to the research questions stated in Section 1.2.

2.2.1 The Baumol model

It was 1952 when Baumol first introduced the inventory control approach to the cash management problem. Baumol (1952) expected that inventory theory and monetary theory can learn from one another. However, a number of important assumptions were made to, using the exact Baumol's words, abstract from precautionary and speculative demands. The most important one was that transactions are perfectly foreseen and occur in a steady stream. Consequently, we are dealing with a deterministic model. Baumol assumed that an outflow of T dollars occurs for a given period of time in a steady stream. On the other hand, in order to offset these outflows, inflows can be obtained by borrowing or by withdrawing from an investment at the cost of i dollars per dollar per period. One more assumption is made by considering that these withdrawals are done in lots of C dollars evenly spaced in time with a fixed cost of b dollars (see Figure 2.1). Under these constraints, cash managers make T/C withdrawals for a given period and total cost is given by:

$$\frac{bT}{C} + \frac{iC}{2} \quad (2.1)$$

where the first part of the equation is the number of transactions multiplied by the unitary fixed cost of each of transaction and, the second one is the average cash balance multiplied by the cost of holding this balance.

Then, the goal for cash managers is choosing C so that equation (2.1) is minimized. Setting the derivative of the total cost with respect to C to zero, we obtain that the value of C that minimizes (2.1) is given by the following expression:

$$C = \sqrt{\frac{2bT}{i}}. \quad (2.2)$$

This initial simple model was extended to consider the possibility of withholding some or all of the receipts from investment and keeping the cash until is needed. From the total outflow T , the cash manager can decide to invest




Figure 2.1: The Baumol model.

I dollars and withhold the remainder R dollars. In this case, Baumol also considered fixed (b_w) and variable (k_w) withdrawing costs and fixed (b_d) and variable (k_d) depositing or investing costs. Then the total cost of cash operations is given by the sum of the total cost of withholding the R dollars and investing the I dollars, and the cost of obtaining cash for the remainder of the period as follows:

$$\frac{T - I}{2} i \frac{T - I}{T} + b_d + k_d I + \frac{C}{2} i \frac{I}{T} + (b_w + k_w C) \frac{I}{C}. \quad (2.3)$$

Again the optimal value of C is given by the equation (2.2) with $b = b_w$. Now partially differentiating equation (2.3) with respect to I we obtain the optimum average cash balance before drawing on invested receipts begins:

$$R = T - I = \frac{C}{2} + \frac{b_w T}{C i} + \frac{T(k_d + k_w)}{i}. \quad (2.4)$$

Substituting equation (2.2) in (2.4) we obtain a simpler expression:

$$R = C + T \left(\frac{k_d + k_w}{i} \right). \quad (2.5)$$

Note that the analysis assumes a continuous steady stream of outflows and discontinuous inflows, but it also applies to the case of continuous receipts and discontinuous payments.

On the application of this model to the real corporate cash management problem, it is interesting to say that, in his work, Baumol pointed out that his model was "only a suggestive oversimplification". The emphasis was on the description of the demand for money by firms or household economies. The steady stream of payments and the absence of receipts during the relevant period make this model impractical in many real applications. However, a first step in the inventory control approach to the cash management problem was done. Interestingly, Baumol also envisioned the inherent task of forecasting the cash flow by stating that, with sufficient foresight, if receipts can meet payments, savings in the use of cash can be achieved.

*Baumol (1952) initiated the inventory approach to the cash management problem and highlighted the utility of forecasts as stated in **Question 1**.*

2.2.2 The Tobin model

Tobin (1956) argued that cash requirements depend inversely on the rate of interest for a given volume of transactions governed by the lack of synchronization of receipts and disbursements. The higher the lack of synchronization the higher the need for transaction balances. However, there is no need for holding these balances in cash. Instead, cash managers have the opportunity to keep these balances in assets with higher yields, such as bonds or marketable securities. At the time that cash is needed these assets would be shifted into cash again to face payments. Consequently, it is likely that the amount of cash held for transaction purposes is related inversely to the interest rate on such alternative assets.

Given an interest rate r , the problem is to find the relationship between what is held in cash and what is held in alternative assets so as to maximize interest

earnings, net of transaction costs. At the beginning of each period $t = 0$, an amount Y is held by the cash manager that is uniformly disbursed until the end of the period $t = 1$ when no cash is available, as shown in Figure 2.2. Thus, the total transactions balance is $T(t) = Y(1 - t)$, with $0 \leq t \leq 1$. However this total $T(t)$ can be divided between cash $C(t)$ and bonds $B(t)$ such that $T(t) = C(t) + B(t)$, with $B(t)$ yielding an interest r per time period. Three different questions are then faced by Tobin: (i) given r and a fixed number n of transactions, determine the optimal timing and amounts to be held in cash and bonds; (ii) given r but a variable number n of transactions, determine the optimal n^* ; and (iii) how does n^* depend on r ?

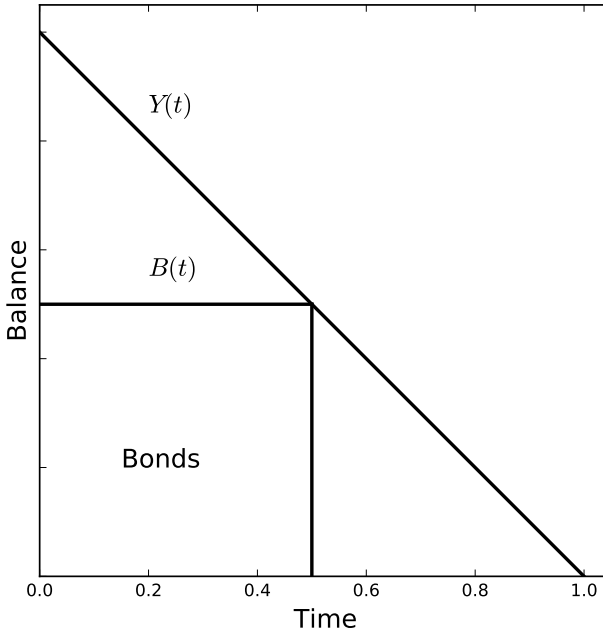


Figure 2.2: The Tobin model.

Considering a transaction x between bonds and cash, the transaction cost is given by $a + b \cdot x$, with $a, b > 0$. Then, for the general case, Tobin proves that the average amount of bonds is given by:

$$\bar{B} = \frac{n-1}{2n} Y \left(1 - \frac{4b^2}{r^2} \right) \tag{2.6}$$

with $n \geq 2$ and $r \geq 2b$. In order to determine the optimal number of transactions, the next profit function is maximized:

$$\pi_n = \frac{n-1}{2n} Yr \left(1 - \frac{2b}{r}\right)^2 - na \quad (2.7)$$

which is a decreasing function of n . Then, the optimal number of transactions n^* is greater than 2, when $1/12 Yr(1 - 2b/r)^2 \geq a$ holds. Finally, the relation of the optimal number of transactions n^* , and the interest rate is given by equation (2.6). Since B_n is an increasing function of n , and n^* varies directly with r , the optimal proportion of bonds also varies directly with r , and consequently, the proportion of cash varies inversely with r , for high enough rates.

Later on, Smith (1986) proposed a *Dynamic Baumol-Tobin Model of Money Demand* based on a stochastic model described by Frenkel and Jovanovic (1980). However, this Baumol-Tobin model has much more to do with the Constantinides and Richard (1978) model than with the first proposals by Baumol (1952) and Tobin (1956). More recently, Mierzejewski (2011) followed Tobin's approach, according to which companies hold cash as a behavior towards risk, to propose a theoretical model of equilibrium in markets of cash balances which is beyond the scope of this thesis.

Summarizing, the Tobin (1956) model is also a deterministic model dealing with a uniform cash flow such as the Baumol (1952) but incorporating the interest rate as a key parameter. In addition, Tobin considered not only fixed costs but also variable transaction costs between two alternative assets, namely, bonds and cash.

*Tobin (1956) considered an operating condition such as the interested rate as key factor to include in cash management as we identified in **Question 6**.*

2.2.3 The Miller-Orr model

The Baumol (1952) model was later considered by Miller and Orr (1966) to apply reasonably well to much of the household sector, particularly to salary-earning households. Baumol focused on the cost of putting idle cash to work by means of the wide variety of interest-bearing securities of very low risk and very quickly convertible to cash. Although this is a more common practice in

business firms than in household economies, Miller and Orr argued that cash balance does not fluctuate steadily but irregularly for many companies, hence resulting in an impractical application of the Baumol model.

In "A Model of Demand for Money by Firms", Miller and Orr developed a simple model following an opposite approach to Baumol by considering stochastic cash flows. From a predictability point of view, Miller and Orr leaped from the perfect knowledge of cash flows in Baumol model to cash flows generated by a stationary random walk, from a deterministic approach to completely stochastic cash flows. Particularly, Miller and Orr considered cash flows to be 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 features of this approach are independence, stationarity, zero-drift, and the absence of regular swings in the cash flow. Moreover they ignored shortage and variable transaction costs.

In a first attempt to deal with the corporate cash management problem, they assumed that companies seek to minimize long-run average cost of managing cash balances under some simple policy. This policy sets a lower bound, zero, and an upper bound, h , where the cash balance is allowed to wander around between the lower and upper levels. We say then that the Miller and Orr (1966) model is a Bound-Based Model (BBM). Apart from the cash balance, the model also assumes the existence of a second asset of any kind like interest bearing assets or marketable securities grouped in a portfolio of investments which are easily transformed in cash at the company convenience. The policy implies that when the upper bound 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 zero a positive transfer will be made to again restore the balance to z .

Although Miller and Orr set the lower limit to zero in their work, in practice a real cash manager should set a lower limit above zero for precautionary motives. This lower limit represents a safety cash buffer and its selection will depend on the level of risk the company is willing to accept. This model variation can be found in Ross, Westerfield, and Jordan (2002), which sets a lower limit l , and an upper bound, h . When h is reached a withdrawal transfer is made to restore the balance to a target level of z . In the same way, when the cash balance reaches l , a positive transfer is made to again restore the balance to z . Formally, the transfer occurring at time t , x_t , is elicited by comparing the




Figure 2.3: The Miler-Orr model.

current cash balance, b_{t-1} , to the lower and upper bounds:

$$x_t = \begin{cases} z - b_{t-1}, & \text{if } b_{t-1} > h \\ 0, & \text{if } l < b_{t-1} < h \\ z - b_{t-1}, & \text{if } b_{t-1} < l \end{cases} \quad (2.8)$$

To obtain the limits, once the cash manager has set the lower limit l , the optimal values of the policy parameters h and z are derived from the expected cost per day over any planning horizon of T days given by:

$$E(c) = \gamma \frac{E(N)}{T} + vE(M) \quad (2.9)$$

where $E(c)$ is the expected cost per day, $E(N)$ is the expected number of transfers during the planning period T ; γ is the the cost per transfer; $E(M)$ is the average daily cash balance; and v is the daily rate of interest earned on the portfolio as an opportunity cost of the idle cash. Setting $Z = h - z$, the

problem can be stated in terms of the variance of the net cash flow as:

$$\operatorname{argmin}_{Z,z} E(c) = \operatorname{argmin}_{Z,z} \frac{\gamma\sigma^2}{Zz} + \frac{v(Z+2z)}{3} \quad (2.10)$$

where the first part of the equation is the transfer cost term and the second one is the holding cost term. Note that the average cash balance is $(h+z)/3$. Hence the optimal parameters are given by:

$$z^* = \left(\frac{3\gamma\sigma^2}{4v} \right)^{1/3} \quad (2.11)$$

and

$$Z^* = 2z^*. \quad (2.12)$$

or in terms of the original parameters:

$$h^* = 3z^*. \quad (2.13)$$

The equivalent equations for the case of a lower bound (l) distinct from zero can easily be derived as presented in Ross, Westerfield, and Jordan (2002) to obtain:

$$z^* = l + \left(\frac{3\gamma\sigma^2}{4v} \right)^{1/3} \quad (2.14)$$

and

$$h^* = 3z^* - 2l. \quad (2.15)$$

The reasoning followed was indeed given by Miller and Orr in their original work Miller and Orr (1966) by stating that it would pay to increase z by Δ and to reduce Z by 2Δ since the higher value of the denominator in equation (2.10) implies a lower value for the transfer cost term. The major implication and the main novelty of this model in comparison to the Baumol model is the presence of the observable variance of the net daily cash flow. As in the case of Baumol model, the greater the transfer cost (γ), the higher the target cash balance (z), and the greater the daily rate of interest (v), the lower the target cash balance (z). However the greater the uncertainty of the net daily cash flow, measured by σ^2 , the higher is the target cash balance (z), and the higher the difference between the lower bound (l) and the higher bound (h). This represents a first step towards a more practical approach to the corporate cash management problem since common sense shows that the greater the uncertainty, the greater is the chance that the balance will drop below the lower bound.

A number of extensions of the model have been considered to incorporate systematic drift in the cash balance and to allow for more than one portfolio asset with different transfers and holding costs. Despite the assumption of the totally stochastic mechanism of cash flow, the authors pointed out the presence of both stochastic and deterministic, or at least highly predictable, elements in the cash flow such as payroll disbursements or dividend payments. However, they argued that the gains from exploiting any cash flow pattern are by no means large enough to offset the added costs of model development and implementation. Although a first attempt to a cost sensitive approach to cash optimization models is done (another example is given in Daellenbach (1974)), no proof is given to justify their argument. They will be right when costs of development are higher than savings from optimization models, but they will be wrong otherwise. In this sense, different costs of development and savings will lead to different results depending on the case.

An interesting extra comment is done by the authors regarding the use of forecasts. The transformation of lower and upper bounds into zones rather than single point limits is likely to be produced when forecasts about the near future are used by cash managers. An example is when a firm will not transfer funds to the portfolio when the cash balance is higher than h but important payments will have to be done in the very near future. This procedure would later become the main motivation used by Stone (1972) to develop his cash management model.

Summarizing, the Miller and Orr (1966) model was the first stochastic cash management model proposed in the literature. They introduced the concept of bounds or control limits which are directly linked to the statistical properties of cash flows, which they assumed to be a random walk. Since we here deal with stochastic cash flows and most of recent cash management models stem from this model, we will revisit it several times in this thesis.

*Regarding **Question 1**, Miller and Orr (1966) pointed out the possibility of using forecasts to improve their model. However, they assumed cash flows to be completely unpredictable which is closely related to **Question 2** on the statistical assumptions of daily cash flows.*

2.2.4 *The Archer model*

Archer (1966) proposed a "A model for the determination of firm cash balances". He essentially followed a practical approach without almost any mathematical apparatus. Archer referred to an indication in Baumol (1952) pointing out the tendency of cash flow to be normally distributed. In addition, he proposed that the precautionary cash requirements might also grow as the square root of the volume of its transactions, i.e., its standard deviation. Archer claimed that Baumol's model was incomplete due to the lack of attention paid to this precautionary use of cash.

Archer also suggested that payments for interest, dividends, or capital expenditures are non-operating cash outflows and recommended a separate cash planning for them. This situation will be also covered later by Stone and Miller (1987). Moreover, he pointed out that seasonal activities of companies may result in a different handling of non-normal periods. Another of the useful suggestions made by Archer was the utility of plotting inflows, outflows and net cash flows from operations as a way to study its behavior over the whole planning period. The very likely lack of synchronization between inflows and outflows during the days of the month forces cash managers to maintain a stock of cash sufficient to cover this transaction gap. Assuming that the inflows and outflows are stable for all months this can be easily view by simply plotting inflows and outflows, or simply net cash flows as shown in Figure 2.4. A simple graphical analysis may result in detecting usual patterns such as the common business practice of grouping payments on one or two fixed days in the month, for example, the 10th and 25th.

Archer followed a statistical approach by: (i) plotting the expected variability of net cash flow in an histogram of frequencies; (ii) computing its mean, its standard deviation and extreme values, i.e., minimum and maximum. We entirely agree with Archer in the sense that this graphical exploratory analysis can be of great help in practice as we will see below. As a consequence of this analysis, Archer suggested a basic procedure that cash managers should follow assuming a month as a planning period:

1. Tabulate net cash flow.
2. Compute a mean for each month.
3. Derive its needed cash balance at the beginning of the month for transactions purposes as the mean of monthly net cash flow.




Figure 2.4: Monthly net cash flow example with two fixed days of payment.

4. Compute a measure of variability of the transactions cash balance required such as the standard deviation of the transactions cash balances required for the various months
5. Set an accepted level of risk and derive the needed amount of cash for precautionary purposes
6. Add the precautionary requirement to the transactions requirement of cash from step 2 to obtain the amount of cash needed at the beginning of the month to cover all payments and avoid a shortage of cash for a given level of risk

The implications of the last step of the above mentioned procedure are of great importance. Intuitively one can think that high variability cash flows are likely to produce unexpected cash drains or a cash shortage rather than a more stable pattern of cash flow. Here the concept of probability of running

out of cash is introduced by Archer to show that cash managers may decide the level of risk they are willing to accept. For example, either if we assume or we effectively know that our net cash flow is normally distributed, we can empirically determine that standard deviation of the necessary cash balance at the beginning of the month, e.g., 1000 €. Then, we should add a quantity equivalent to 2 times 1000 €, that is, 2000 € to the initial cash balance if we want to accept a risk of running out of cash of 5%. Recall that values within the band of the mean plus/minus two times the standard deviation of a normal distribution account for 95% of the total set of observations or, in other words, 5% of the observations are left out of the band. Consequently, the cash balance needed at the beginning of the month meets the transaction requirement and the added quantity meet the precautionary requirement of cash.

As a summary, Archer suggested to perform an empirical analysis of daily cash flow to improve its understanding. He did not provide a quantitative technique but a general procedure to manage cash that can be repeated as many times as needed to adjust to any particular situation.

Archer (1966) recommends to explore data to improve the understanding of the cash management problem along the lines of Question 2.

2.2.5 The Eppen-Fama model

A variation of the Miller and Orr (1966) model was introduced by Eppen and Fama (1969) following a dynamic programming approach. However, it was a previous publication (Eppen and Fama, 1968) the one that provided a complete analysis on the effect of variations of transfer, and holding and penalty costs on the optimal policies. The Eppen-Fama model is a generalization of the stochastic Miller-Orr model in which transfer costs contain both fixed and variable components. They showed that if transfer costs have a fixed cost as well as a cost proportional to the amount transferred, the optimal strategy was of the form of two limits (u, d) and two return points (U, D), one for each limit. In this model, when the cash balance reaches the upper bound (d), it is immediately restored to the upper return point (D), and when the cash balance reaches the lower bound (u), it is restored to the lower return point (U), as shown in Figure 2.5.

Following a Markovian approach, they assumed that the probability mass function of the transitions between different possible states was known and station-




Figure 2.5: Eppen-Fama model representation with two return points.

ary. This assumption implies a process of discretization of the cash balance. At any point in time, the cash balance can be in one of N possible states $i = 1, 2, \dots, N$, each one representing a discrete level of the cash balance. The lowest level occurs in state 1 and the highest in state N , and each successive level differ by some constant R , for example 1000 €.

For the general case, two cost functions are defined: firstly, the transfer cost (t_i^k) caused by moving the cash balance from state i to state k :

$$t_i^k = \begin{cases} K_u + c_u(k - i) & \text{if } k > i; K_u, c_u > 0, \\ 0 & \text{if } k = i, \\ K_d + c_d(i - k) & \text{if } k < i; K_d, c_d > 0 \end{cases} \quad (2.16)$$

where K_u and c_u are respectively the fixed and variable components of a positive cash movement, and K_d and c_d are respectively the fixed and variable components of a negative cash movement; secondly, the holding or penalty

cost ($L(k)$) associated to starting a period in state k can be defined as:

$$L(k) = \begin{cases} c_p(M - k) & \text{if } k < M; c_p > 0, \\ c_h(k - M) & \text{if } k > M; c_h > 0 \end{cases} \quad (2.17)$$

where c_p is the marginal penalty cost per period per R unit of cash, c_h is the marginal holding cost per period per R unit of cash, say 1000 €, and M is the minimum cash balance that has to be kept because of any condition required by banks. In absence of this restriction, M is usually set to zero as the minimum cash balance required to be held in the bank account. Recall that Miller and Orr (1966) suggested the use of two or three bounds. In order to account for fixed and variable transaction costs, the Eppen and Fama (1968) BBM model proposed the use of four bounds, resulting in an optimal policy of the following form:

- a) If the cash balance at the beginning of the period is below u , move it up to U .
- b) If the cash balance at the beginning of the period is above d , move it down to D .
- c) If the cash balance at the beginning of the period is between U and D , do nothing.

Further discussion on the properties of the solutions to the corporate cash management problem for different probability distributions and a wide variety of values of the cost parameters can be found in Eppen and Fama (1968). From an experimental perspective, the authors pointed out that higher dispersion in the probability distribution caused the outer bounds, u and d , and the return points, U and D , to be further away from zero. In this sense, in practical applications, it is highly recommended to carefully estimate the probability distribution, specially in the extremes. Moreover, when both the probability distribution and the cost function are symmetric about zero the optimal policies are also symmetrical. In addition, the optimal policy seems to be relatively insensitive to small changes of the cost parameters. This fact may be important when estimating costs since allowing higher errors may be relatively unimportant.

Summarizing, a number of interesting contributions on the practical side of the corporate cash balance problem were made by Eppen and Fama. They considered both fixed and variable transaction costs resulting in a policy based on four bounds. On the utility of the cash balance problem, they pointed out that the cash balance problem captures the essence of the dynamic portfolio

management that has to be performed on a daily basis not only by most of the companies, but also by mutual funds and other financial institutions. They suggested that in many corporate cash balance problems it is reasonable to extend the problem to the consideration of three possible type of assets: 1) cash; 2) an intermediate low return asset such as bonds; and 3) high return assets such as common stocks. The interesting question to be considered concerns the conditions under which the intermediate asset should be treated as a buffer stock between cash and the high-return asset.

Eppen and Fama (1968) provided a general cost structure based on fixed and variable transaction costs, and also linear holding costs which is extensively used in this thesis. They also assumed stochastic cash flows following a random walk, which is related to Question 2.

2.2.6 The Daellenbach model

Daellenbach (1971) proposed an improvement of the Eppen-Fama model in "A stochastic cash balance model with two sources of short-term funds". Daellenbach claimed that his model was a generalization of the Eppen-Fama model to situations where bank account overdrafts are not possible, and using two different sources of short-term funds, namely, marketable securities and short-term loans. Furthermore, in contrast to previous models, the probability distribution of cash flows is not necessarily stationary and, the length of the review periods may vary from period to period. Again, a decision about the adjustment of the cash balance has to be made but, in this model, an allocation decision about either marketable securities or borrowing transactions is also necessary. A dynamic programming approach was proposed labeling periods in the planning horizon as $n = N$ for the first period chronologically, and $n = 1$ for the last period. Then, three state variables were considered to describe the cash balance situation:

1. B_n or the cash balance at the beginning of period n carried forward from $n + 1$.
2. Z_n or the borrowing balance at the beginning of period n carried forward from $n + 1$.
3. S_n or the marketable securities balance at the beginning of period n carried forward from $n + 1$.

If X_n and Y_n denotes transactions in the form of borrowings or marketable securities respectively and, R_n is the sum of uncontrollable cash transactions in period n , with probability density function $f_n(r_n)$, the following balance equation is used to link period $n - 1$ to period n :

$$B_{n-1} = B_n + X_n + Y_n + R_n \quad (2.18)$$

$$Z_{n-1} = Z_n + X_n, \quad Z_n \geq 0 \quad (2.19)$$

$$S_{n-1} = S_n - Y_n, \quad S_n \geq 0 \quad (2.20)$$

subject to:

$$B_n + X_n + Y_n \geq 0 \quad (2.21)$$

$$Z_n + X_n \geq 0 \quad (2.22)$$

$$S_n - Y_n \geq 0 \quad (2.23)$$

meaning that: (i) the beginning cash balance before any adjustment has to be non-negative; (ii) the outstanding borrowing balance cannot be reduced below zero; and (iii) marketable securities cannot be sold short.

According to the previous equations, the state variable set for the cash position at the beginning of period n , prior to any cash balance adjustment, is denoted by $\Omega_n = (B_n, Z_n, S_n)$, the decision variables are (X_n, Y_n) , and the total cost is the sum of: (i) fixed and variable transaction costs for borrowing; (ii) fixed and variable transaction costs for marketable securities; (iii) interest cost on borrowings; (iv) returns on marketable securities (note that this is a negative cost or a benefit); and (v) penalty costs for cash shortages. All these costs can be summarized as follows:

$$T_n(X_n, Y_n; \Omega_n) = H_1(X_n) + H_2(Y_n) + c_{1n}(Z_n + X_n) - c_{2n}(S_n - Y_n) + L_n(B_n + X_n + Y_n) \quad (2.24)$$

where $H_1(X_n)$ is the borrowing cost function computed as:

$$H_1(X_n) = \begin{cases} -b_1^- X_n & \text{if } X_n < 0, \\ b_1^+ X_n & \text{if } X_n \geq 0, \end{cases} \quad (2.25)$$

where b_1^-, b_1^+ are respectively variable borrowing transaction cost for cash increases (+) and cash decreases (-); $H_2(Y_n)$ is the marketable securities cost function computed as:

$$H_2(Y_n) = \begin{cases} -b_2^- Y_n & \text{if } Y_n < 0, \\ b_2^+ Y_n & \text{if } Y_n \geq 0, \end{cases} \quad (2.26)$$

where b_2^-, b_2^+ are respectively variable marketable securities transaction cost for cash increases (+) and cash decreases (-); c_{1n} is the interest cost on ending loan balances; c_{2n} is the return on ending marketable securities holdings; $L_n(B_n)$ is the expected cost of cash shortage incurred at the end of period n computed as:

$$L_n(B_n + X_n + Y_n) = c_{3n} \int_{-\infty}^{-(B_n + X_n + Y_n)} (B_n + X_n + Y_n + r_n) f_n(r_n) dr_n \quad (2.27)$$

where c_{3n} is the penalty on negative ending cash balances in period n .

Considering alternative funding sources such as borrowings and marketable securities, introduces additional considerations on the priorities based on the feasible permutations of the cost coefficients as follows:

- Case 1. If $-b_2^- + c_2 \leq -b_1^- + c_1 \leq b_1^+ + c_1 \leq b_2^+ + c_2$, then borrowing transactions are preferred over marketable securities.
- Case 2. If $-b_1^- + c_1 \leq -b_2^- + c_2 \leq b_2^+ + c_2 \leq b_1^+ + c_1$, then marketable securities transactions are preferred over borrowing.
- Case 3. If $-b_2^- + c_2 \leq -b_1^- + c_1 \leq b_2^+ + c_2 \leq b_1^+ + c_1$, then borrowing transactions are preferred over marketable securities for cash withdrawals, and marketable securities are preferred over borrowings for cash procurements.
- Case 4. If $-b_1^- + c_1 \leq -b_2^- + c_2 \leq b_1^+ + c_1 \leq b_2^+ + c_2$, then marketable securities are preferred over borrowing transactions for cash withdrawals, and borrowings are preferred over marketable securities for cash procurements.
- Case 5. If $-b_2^- + c_2 \leq b_2^+ + c_2 \leq -b_1^- + c_1 \leq b_1^+ + c_1$, then borrowing transactions are preferred over marketable securities for cash withdrawals, and marketable securities are preferred over borrowings for cash procurements.

The Daellenbach model can be regarded as an extension of the Eppen and Fama (1968) and Eppen and Fama (1969) model, but with four return points: $\{U_{1n}, D_{1n}\}$ denoting the use of borrowings as the source of funds, and $\{U_{2n}, D_{2n}\}$ denoting the use of marketable securities as the source of funds. The optimal policy gives preference to the source of funds dictated by the previous five cases based on the cost coefficients. In the case that either constraint (2.22) or (2.23) prevents completion of the transaction, then use the return point relevant for the other source of funds.

Later on, Daellenbach (1974) pointed out an important issue by posing the following general question: *Are cash management models worthwhile?* The objective was to determine upper bounds of the potential savings that can be realized by applying cash management models. In this paper, a variant of the model in Daellenbach (1971) was proposed to consider fixed and variable transaction costs. In addition, a deterministic shortage cost function charging negative cash balances at the end of the day was defined instead of the stochastic previous one. The main criticism to cash management models was done on the assumption of perfectly predictable cash flows. Any cost estimates based on perfect predictions will provide optimistic lower bounds for the actual cost incurred which corresponds to determine what the optimal policy would have been given the actual cash flows. By using random normal simulations, the author estimated the upper bounds obtained by this variant of his cash management model to the performance of a hypothetical cash manager. The author concluded that the benefits from cash management optimization models were in most cases highly uncertain and that they offer a very small economic return.

As a summary, Daellenbach (1971) used dynamic programming to provide a solution to the CMP as a set of control bounds, but considering two available sources of funds, namely, marketable securities and short-term loans. In addition, the usual assumption on stationary cash flows was relaxed. However, in Daellenbach (1974), the author raised doubts about the benefits derived from the deployment of cash management models. In this thesis, we approach this open research question from a general multiobjective perspective.

*Daellenbach (1974) discussed the utility of cash management models under different circumstances, which is closely related to the open research **Question 6**.*

2.2.7 The Stone model

"The use of forecasts and smoothing in control-limit models for cash management" was proposed by Stone (1972). In this work, Stone first reviewed the assumptions of the Baumol and the Miller-Orr models summarized in Table 2.1 and pointed out a series of limitations of these models in real world cash management situations. Stone argued that cash flows are neither completely certain, uniform, and continuous (as they are in the Baumol model) nor are they completely unpredictable (as they are in the Miller-Orr model).

Concept	Baumol Model	Miller-Orr Model
Cash flow	Continuous and known	Stochastic and iid
Transfer cost	Constant and fixed	Constant and fixed
Holding cost	Constant and fixed	Constant and fixed
Assets considered	Cash and securities	Cash and securities
Securities treatment	Reinvested	Reinvested
Minimum cash balance	Yes	Yes

Table 2.1: Assumptions of the Baumol and Miller-Orr models.

Most firms can forecast their cash flows. This is the first time that the concept of forecasting cash flows appears as a key input to any cash management model. The author focuses on the generally attempted tasks performed by cash managers. They usually:

1. Look ahead when buying and selling securities to incorporate data from their cash forecasts.
2. Smooth cash flows by coordinating security maturities with predicted cash needs.
3. Buy the highest yielding securities subject to portfolio and liquidity constraints.
4. Maintain cash balances sufficient to meet banking requirements.

From these generally performed tasks, Stone derived the idea of including both forecasts and maturing securities in his model. Operation of this control-limit model is based on the ability to buy and sell securities of different maturities in order to reduce transaction costs by smoothing cash flows and thereby reducing the number of transactions. It is assumed that the current cash balance, CB_0 , is known and that a forecast of the net cash flow, $E(C_t)$, that will occur on each day t over the next k days is available. 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:

$$E(CB_k) = CB_0 + \sum_{t=1}^k E(C_t). \quad (2.28)$$

Alternatively, if the sum of net cash flows over the next k days are lumped into a single figure, this last equation can be rewritten as:

$$E(CB_k) = CB_0 + E(SC_k). \quad (2.29)$$

Next a number of simple rules are proposed to be followed by cash managers to return to a desired target balance TB , based on two sets of control limits. One set is defined by h_1 and h_0 as the upper and lower control limit for initiating considerations of a 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.




Figure 2.6: Structure of the Stone model with two sets of limits

The set of rules followed by cash managers to operate the model are summarized as follows:

1. If the current cash balance CB_0 is inside 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 balances in the next k days, $E(CB_k)$, exceed the control limits defined by $h_1 - \delta_1$ and $h_0 + \delta_0$, a transaction is made to return the expected cash balance to the target level TB in k days.
- (b) No action is taken otherwise.

The innovation introduced by the Stone model is that 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. Being K_t , the amount of the transaction to be made, these rules can be represented mathematically as follows:

$$K_t = \begin{cases} TB - CB_0 - E(SC_k), & \text{if } CB_0 > h_1 \text{ and } CB_0 + E(SC_k) > h_1 - \delta_1, \\ 0, & \text{if } h_0 < CB_0 < h_1 \\ TB - CB_0 - E(SC_k), & \text{if } CB_0 < h_0 \text{ and } CB_0 + E(SC_k) < h_0 + \delta_0. \end{cases} \quad (2.30)$$

An example of the application of this k -day look-ahead procedure is given by the author where the target balance $TB = 20$, the control limits $h_1 = 27$, and $h_0 = 13$; and $\delta_1 = \delta_0 = 2$, all figures in millions of dollars. A comparison between a 3-day look-ahead procedure and a no-day look-ahead is presented in Table 2.2, where more transaction cash flows are observed in the case of the no-day look-ahead. Surprisingly, no evaluation of the total cost of the policy was performed. From Miller and Orr (1966) and Eppen and Fama (1968) a general cost function considering the transfer cost and the holding cost could be defined to evaluate the average cost of the policy. For instance, assuming the cost of transaction being formed by a fixed part, $\gamma_0 = \$5$, and a variable part, $\gamma_1 = 0.01\%$ of the transaction, and the holding cost per money unit of a positive cash balance at the end of the day, $h = 0.00027$, i.e., 10% p.a. divided by 365 days, the total cost of the policy over the period of 16 days considered can be easily computed using the following expression:

$$Cost = \sum_{t=1}^T (\Gamma(K_t) + hB_t) \quad (2.31)$$

Day	Cash forecast and cash flow			3-day look-ahead		No-day look-ahead	
	Pred CF	Actual CF	Error	Trans CF	Final Balance	Trans CF	Final Balance
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	-17	11	0	26
11	5	3	2	0	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

Table 2.2: Comparison of the Stone 3-day look-ahead with the no-day look-ahead procedure

where $T = 16$ is the considered time horizon, B_t is the cash balance at the end of day t , and $\Gamma(K_t)$ is the transfer cost function that can be defined as:

$$\Gamma(K_t) = \begin{cases} \gamma_0 - \gamma_1 K_t & \text{if } K_t < 0, \\ 0 & \text{if } K_t = 0, \\ \gamma_0 + \gamma_1 K_t & \text{if } K_t > 0. \end{cases} \quad (2.32)$$

Hence, the total cost for the look-ahead procedure of the example is \$86,363, and \$99,941 for the no-day look-ahead procedure. Note that, under the cost structure defined by γ_0 , γ_1 , and h , cost savings obtained by the proposed procedure comes from both, the reduced number of transactions and the smaller average cash balance.

Since the cash policy is fixed for a k -days period, the use of forecasts forces the cash manager to monitor errors for k days after a transaction has occurred. However, the impact of predictive accuracy of the forecasts on the performance of the policy was not evaluated. It is expected that a better prediction will lead to better policies as hypothesized in Gormley and Meade (2007) and, consequently, an evaluation of the impact of predictive accuracy is a mandatory

step. Furthermore, efforts in improving predictive accuracy have associated costs that have to be compared to savings obtained in order to decide if further efforts are worth. The impact of cash flow forecasts is an ongoing issue in cash management that we here address in Question 1, since we consider it a crucial challenge.

On the selection of the model parameters, no particular procedure was specified by Stone although some suggestions were made, namely, not to treat them as fixed parameters, but rather adjust them as necessary. Simulation and the practitioner's judgment were suggested to be the best approach to parameterization. Involvement of cash managers in the process of parameter selection was considered in this case an advantage of this method. An alternative approach to deal with cash flow uncertainty was followed by Hinderer and Waldmann (2001) who developed a rigorous mathematical framework to include in the cash manager decision-making process varying environmental factors.

Furthermore, using this look-ahead procedure, once a transaction is decided, the maturity of the security should be specified. This subpart of the problem represents an additional issue that was not addressed by the previous approaches. The maturity structure decision is a trade-off between returns obtained and the opportunity to smooth cash flows and reduce transaction costs. For instance, when buying securities the logical maturity dates for smoothing cash flow are dates when a future sale of securities is forecast. A heuristic is proposed by the author for specifying quantities and maturities due to the nonlinear dependence of transaction costs and the amount of each maturity that is transacted.

Summarizing, Stone was the first author to formally develop a cash management model using forecasts as a key input. Consequently, they assumed that cash flows are, to some extent, predictable. Several works on daily cash flow prediction (Stone and Wood, 1977; Stone and Miller, 1981; Miller and Stone, 1985; Stone and Miller, 1987) represent an important contribution to cash management literature. However, the lack of a formal procedure to determine the set of parameters (bounds) of the look-ahead procedure rather, apart from the mere suggestion of using simulation, become a serious limitation.

*Stone (1972) is another crucial model that decisively contributed to the use of forecasts as a key input to cash management that we here address in **Question 1**. In addition, Stone pointed out the importance of changing contexts as a motive to adapt cash management models to current circumstances (**Questions 5 and 6**).*

2.2.8 The Constantinides model

Although Neave (1970) showed cases in which the Eppen and Fama (1969) model was not optimal, Constantinides and Richard (1978) proved the existence of optimal simple policies for discounted costs when the net cash flow followed a Wiener process. They studied the case of fixed and variable transaction costs and linear holding and penalty costs, and used impulse control techniques to find sufficient conditions for an optimal policy defined by parameters $d \leq D \leq U \leq u$. In the same way of other bound-based models, control actions are only taken whenever the cash level either rises above u or falls below d money units.

Instead of a discrete time framework which was considered in Eppen and Fama (1968), Eppen and Fama (1969), Girgis (1968), and Neave (1970), Constantinides and Richard assumed that decisions are made continuously in time. Moreover, they assumed that demand over any length of time is generated by a Wiener process, meaning that it is normally distributed with both mean and standard deviation proportional to the length of time considered. On the other hand, they followed the impulse control approach by Bensoussan and Lions (1975) that was later extended by Richard (1977). This control technique is based on control actions taken at stochastic stopping times.

The problem formulation was similar to previous works in cash management. The cash balance at time t is defined as $x = x(t)$ and it is charged with holding/penalty cost $C(x) = \max\{hx, -px\}$, with $h, p > 0$. The transaction cost of changing the cash level from x_0 to x_1 is given by:

$$B(x_1 - x_0) = \begin{cases} K^+ + k^+(x_1 - x_0) & \text{if } x_1 \leq x_0, \\ K^- + k^-(x_0 - x_1) & \text{if } x_1 > x_0, \end{cases} \quad (2.33)$$

with $k^+, k^-, K^+, K^- > 0$, such that a zero control action incurs a fixed cost.

In addition, the cumulative demand for cash in the interval $[t, s]$, denoted by $D(t, s)$, is independent and normally distributed with mean $E[D(t, s)] =$

$\mu(s - t)$ and variance $\text{var}[D(t, s)] = \sigma^2(s - t)$, where μ and σ^2 are constants. Thus, the cumulative demand is given by:

$$D(t, s) = \mu(s - t) + \sigma(w(s) - w(t)) \quad (2.34)$$

where w is a Wiener process in \mathbb{R} with zero drift and diffusion coefficient one. However, the use of diffusion processes to represent cash holding evolution was not new (Miller and Orr, 1966).

Under this framework, cash managers continuously observe the cash level and perform control actions when necessary. At any stopping time τ_i , the control applied ϕ_i , is a random variable that is independent of the future state of the system. An impulse control policy v is represented as a sequence of stopping times and controls: $v = [\tau_1, \phi_1; \tau_2, \phi_2; \dots]$. If $x(\tau_i^-)$ denotes the cash level at stopping time τ_i before the control action ϕ_i is applied and $x(\tau_i)$ denotes the cash level after the control action, then the state equations of the cash level when policy v is applied are given by:

$$dx(t) = -\mu dt - \sigma dw(t) \quad (2.35)$$

when $0 \leq t < \tau_i$, with $x(0^-) = x_0$, and:

$$x(\tau_i) = x(\tau_i^-) + \phi_i, \quad dx(t) = -\mu dt - \sigma dw(t) \quad (2.36)$$

when $\tau_i \leq t < \tau_{i+1}^-$, with $i \geq 1$. From that, given a policy v and an initial cash balance $x(0^-) = x_0$, the expected total cost from time zero to infinity, discounted to time zero is:

$$J_{x_0}(v) = E \left[\sum_{i=1}^{\infty} e^{-\beta \tau_i} B(\phi_i) + \int_0^{\infty} e^{-\beta s} C(x(s)) ds \right] \quad (2.37)$$

where β is the discount rate. The final goal is to choose a policy v^* such that $J_{x_0}(v^*) \leq J_{x_0}(v)$, $\forall v \in \Omega$, being Ω the class of all impulse control policies.

Let $V(x) = J_x(v)$ be the expected total cost from time t to infinity discounted to time t and conditional on the cash level $x(t^-) = x$. Note also that $V(x) \geq 0$ since all costs are non-negative. There are only two possible alternatives for the cash manager: to take no control action or to make the most convenient transaction in terms of future costs. Applying dynamic programming and assuming that the subsequent decisions are also optimal, $V(x)$ must satisfy:

$$V(x(t^-)) = \min \left\{ \begin{array}{l} \inf_{\xi} [B(\xi) + E(C(x(t))dt + e^{-\beta dt}V(x(t) + dx))] \\ E(C(x(t))dt + e^{-\beta dt}V(x(t) + dx)) \end{array} \right\}, \quad (2.38)$$

From that, the authors derive the following theorem:

Theorem 1. *Suppose that $h > \beta k^-$ and $p > \beta k^+$ holds. Then there exists an optimal policy to the cash management problem. This policy is simple and is given by:*

$$y(x) = \begin{cases} D & \text{if } x \leq d, \\ x & \text{if } d < x < u, \\ U & \text{if } u \leq x, \end{cases} \quad (2.39)$$

Note that the previous theorem implies that if $h < \beta k^-$, it will never be optimal to reduce the cash level as long as $K^- > 0$. Similarly, if $p < \beta k^+$, it will never be optimal to increase the cash level as long as $K^+ > 0$. If both conditions, $h < \beta k^-$ and $p < \beta k^+$, hold the optimal policy prescribes no intervention. In the special case of $h < \beta k^-$ and $p > \beta k^+$, it is optimal to increase the cash level but its never optimal to decrease the cash level. They are dealing then with an inventory problem where the control action $\xi(x)$, is constrained to be non-negative.

This model was later extended to the case of quadratic holding-penalty costs in Baccarin (2002), and to a multidimensional cash management system and general cost functions in Baccarin (2009), when cash balances fluctuates as a diffusion process. Premachandra (2004) also used a diffusion process to propose a more generalized version of the Miller-Orr model which relaxes most of its restrictive assumptions. Note that the Wiener process is also a diffusion process (Itô, 1974).

Summarizing, apart from considering continuous time in cash management, the most important contribution of the Constantinides and Richard (1978) model was Theorem 1, which provided the necessary conditions to avoid the triviality of the cash policy. Furthermore, it represents the origin of a number of recent works (Baccarin, 2002; Premachandra, 2004; Baccarin, 2009) on cash management. However, the strong assumption on modeling cash flows as a diffusion process represents a serious limitation in practical applications.

*Theorem 1 by Constantinides and Richard (1978) is a necessary condition that must hold to avoid non-triviality in cash management, which we extend in this thesis to a multiobjective framework when addressing **Questions 6**.*

2.2.9 The Penttinen model

Penttinen (1991) presented myopic and stationary solutions for linear costs using the logistic distribution as the probability density function of random cash demand. Myopic one-period solutions were suggested to avoid computational difficulties in multi-period applications with a large number of discrete states. Contrarily to Constantinides and Richard (1978), Penttinen chose a discrete time framework due to the fact that common planning and control practices in most organizations are typically performed in discrete intervals.

His main goal was to analyze the amount of suboptimality in the myopic solutions. To this end, the problem formulation considered a stochastic cash balance in which demand δ is a random variable. The amount of cash at the beginning of each period n is denoted by x and the cash balance after a control action is taken is denoted by $y(x)$. The author considered transactions costs $a_n(y - x)$ as:

$$a_n(y - x) = \begin{cases} K_n + k_n \cdot (y - x) & \text{if } y - x > 0, \\ 0 & \text{if } y = x, \\ Q_n + q_n \cdot (x - y) & \text{if } y - x < 0, \end{cases} \quad (2.40)$$

where $K_n, Q_n, k_n, q_n \geq 0$. In addition, the retain and penalty cost $m_n(y)$ charges the cash level y at the beginning of each period according to:

$$a_n(y - x) = \begin{cases} r_n(y) & \text{if } y > 0, \\ p_n(-y) & \text{if } y \leq 0. \end{cases} \quad (2.41)$$

Finally, the holding and shortage cost $l_n(z)$ charges the cash level z at the end of each period. Here, the amount of cash remaining is given by $z = y - \delta$ and the optimal balance at this point is zero since any positive balance is subject to a holding cost and any negative to a shortage cost:

$$l_n(z) = \begin{cases} h_n(z) & \text{if } z > 0, \\ s_n(-z) & \text{if } z \leq 0. \end{cases} \quad (2.42)$$

The expected holding and shortage costs are given by the following loss function:

$$L_n(y) = \int_{-\infty}^{\infty} l_n(y - \delta) \phi_n(\delta) d\delta \quad (2.43)$$

which is actually the convolution of $l_n(y - \delta)$ with probability density function $\phi_n(\delta)$. Then, the optimal discounted value of future costs in the beginning of

period n is:

$$C_n(x) = \inf_y \{a_n(y - x) + m_n(y) + L_n(y) + \alpha \phi_n * C_{n+1}(y)\} \quad (2.44)$$

in which α is a discount factor and $*$ stands for convolution. Note that when $\alpha = 0$, the dynamic model is called a myopic model. The optimal policy of this general convex model is given by:

$$L'(T) = -k - m'(T) \quad (2.45)$$

$$L'(U) = q - m'(U) \quad (2.46)$$

$$L(t) - L(T) = K + k(T - t) + m(T) - m(t) \quad (2.47)$$

$$L(u) - L(U) = Q + q(u - U) + m(U) - m(u) \quad (2.48)$$

where $t \leq T \leq U \leq u$ define a transaction rule in the form of a simple policy $y_n(x)$ such that:

$$y_n(x) = \begin{cases} T_n & \text{if } x < t_n, \\ x & \text{if } t_n \leq x \leq u_n, \\ U_n & \text{if } x > u_n. \end{cases} \quad (2.49)$$

Penttinen introduced the logistic distribution as a way to ease calculations. In this case, the optimal myopic policy is given by:

$$T = \mu + \frac{\ln[-(k + r - s)/(k + r + h)]}{d} \quad (2.50)$$

$$U = \mu + \frac{\ln[(q - r + s)/(-q + r + h)]}{d}. \quad (2.51)$$

The reorder point t , and the disposal point u are derived numerically from T and U from equations (2.47) and (2.48). To this end, an iterative procedure was presented to compute solutions which is supposed to achieve rapid convergence. Different empirical results showed the proportionality of policy parameters t , T , U , and u with the shortage cost ratio so that the higher shortage cost the higher the reorder and disposal points.

On the other hand, stationary solutions are based on the assumption that each period possesses the same cost functions and that cash demand is independent and identically distributed. Then, Penttinen presented additional empirical results on the amount of suboptimality between myopic and stationary solutions in the case of no fixed costs. His results show that the stationary model leads to slightly more cautious ordering policies.

As a summary, it is important to highlight the assumption of the logistic distribution within the commonly used family of Gaussian cash flows, as a way to better represent empirical cash flows. He also assumed fixed and linear transaction costs to derive by dynamic programming two kind of optimal policies, namely, myopic (minimizing short-term costs) and stationary (minimizing long-term costs).

*In relation to this thesis, Penttinen (1991) motivated the selection of the logistic distribution to represent cash flows as a way to close the gap between theory and practice in cash management, along the lines of **Question 2**.*

2.2.10 The Gormley-Meade model

Gormley and Meade (2007) claimed "The utility of cash flow forecasts in the management of corporate cash balances" and proposed a Dynamic Simple Policy (DSP) to demonstrate that savings can be obtained by using cash flow forecasts. They suggested the use of an autoregressive model as the key input to their model. However, gains in forecast accuracy over the naive model were scant. Gormley and Meade expected that savings from using a non-naive forecasting model would increase if there were more systematic variation in the cash flow and, consequently, higher forecast accuracy. If this hypothesis is correct, the savings produced by a better forecasting model are expected to be significantly higher than those obtained by the naive forecasting model. However, no evaluation of the impact of predictive accuracy on savings was performed by the authors becoming an open research question as we stated in Section 1.2.

In their approach to the corporate cash management problem, Gormley and Meade used an inventory control stochastic model where cash balances are allowed to move freely between two limits: the lower (D) and the upper balance limit (V). When the cash balance reaches any of these limits a cash transfer is made to return to the corresponding rebalance level (d, v). For illustration purposes, we here reproduce this model using the same notation. The management of the cash balance over a time period T is determined by a set of policy parameters or limits for the instant t that can be extended τ days ahead:

- $D_{t+\tau}$: Lower balance limit at time $t + \tau$.
- $V_{t+\tau}$: Upper balance limit at time $t + \tau$.

- $d_{t+\tau}$: Lower rebalance level at time $t + \tau$.
- $v_{t+\tau}$: Upper rebalance level at time $t + \tau$.

The transfers for any prediction horizon are determined by:

$$K_{t+\tau} = \begin{cases} v_{t+\tau} - \tilde{O}_{t+\tau-1} - \hat{w}_{t+\tau|t}, & \text{if } \tilde{O}_{t+\tau-1} + \hat{w}_{t+\tau|t} > V_{t+\tau}, \\ 0, & \text{otherwise,} \\ d_{t+\tau} - \tilde{O}_{t+\tau-1} - \hat{w}_{t+\tau|t}, & \text{if } \tilde{O}_{t+\tau-1} + \hat{w}_{t+\tau|t} < D_{t+\tau} \end{cases} \quad (2.52)$$

where $\tilde{O}_{t+\tau-1}$ is the predicted opening balance at time $t + \tau - 1$, $\hat{w}_{t+\tau|t}$ is the predicted cash flow for $t + \tau$ using a model which has been trained up to time t . In this model, $D_{t+\tau} \leq d_{t+\tau} \leq v_{t+\tau} \leq V_{t+\tau}$, and the following continuity function holds:

$$\tilde{O}_{t+\tau} = \tilde{O}_{t+\tau-1} + K_{t+\tau} + \tilde{\varepsilon}_{t+\tau|t} \quad (2.53)$$

Notice that the the transfer function reproduced here in equation (2.52) is different to the original one $\tilde{O}_{t+\tau-1} - \hat{w}_{t+\tau|t} < D_{t+\tau}$ in Gormley and Meade (2007). The equation in Gormley and Meade (2007) is mistaken because the reference value for comparison must be $\tilde{O}_{t+\tau-1} + \hat{w}_{t+\tau|t} < D_{t+\tau}$, namely the opening balance at $t + \tau - 1$ plus the predicted cash flow for $t + \tau$.

The expected cost over horizon T is given by the following objective function to be minimized:

$$Cost = \sum_{\tau=1}^T \Gamma(K_{t+\tau}) + \tilde{O}_{t+\tau}(h \cdot I_{\tilde{O}_{t+\tau} > 0} + u \cdot I_{\tilde{O}_{t+\tau} < 0}) \quad (2.54)$$

where the transfer cost function Γ is defined as:

$$\Gamma(K_{t+\tau}) = \begin{cases} \gamma_0^- - \gamma_1^- K_t & \text{if } K_t < 0, \\ 0 & \text{if } K_t = 0, \\ \gamma_0^+ + \gamma_1^+ K_t & \text{if } K_t > 0. \end{cases} \quad (2.55)$$

The notation used by the expected and transfer cost functions is described below:

- h : The holding cost per money unit of a positive cash balance at the end of the day.
- u : The shortage cost per money unit of a negative cash balance at the end oh the day.

- γ_0^+ : Fixed cost of transfer into account.
- γ_0^- : Fixed cost of transfer from account.
- γ_1^+ : Variable cost of transfer into account.
- γ_1^- : Variable cost of transfer from account.
- $I_{\tilde{O}_{t+\tau} > 0}$: Boolean variable that equals 1 if $\tilde{O}_{t+\tau} > 0$ is true, 0 otherwise.
- $I_{\tilde{O}_{t+\tau} < 0}$: Boolean variable that equals 1 if $\tilde{O}_{t+\tau} < 0$ is true, 0 otherwise.

The authors used genetic algorithms to solve the CMP, i.e., to estimate the parameters $\{D_{t+\tau}, d_{t+\tau}, v_{t+\tau}, V_{t+\tau}\}$, from $\tau = 1, \dots, T$. Moreover, since the model accepts forecasts as its main input, a cash flow autoregressive forecasting model was developed. To this end, a Box-Cox transformation (Box and Cox, 1964) was used to achieve normality of the real cash flow data set used in the paper. Stationarity in mean and variance of the transformed data was confirmed by low values of autocorrelation and by Box-Ljung tests. The autoregressive model was fitted to non-holiday dates but the coefficient of determination, R^2 was only of 6.5%, due to high volatility. Normalized mean square errors, for different prediction horizons from 1 up to 100 days were also reported with results ranging in 0.95 to 0.99 showing the low predictive ability of the autoregressive model. An additional non-linear least squares predictive model was tried with even lower predictive accuracy.

A cash management example was given comparing the DSP model using forecasts from the autoregressive forecasting model, the DSP model using forecasts from the constant mean model and the Penttinen (1991) model. Equal holding and shortage costs and only fixed transactions costs were considered but results were so similar that no claim in favor of the cash management model and the forecasting technique could be done. However, the authors concluded that the expected savings from using better forecasting models with higher accuracy would increase.

Summarizing, Gormley and Meade (2007) proposed a cash management model using forecasts as a key input. Surprisingly, they did not refer to the work by Stone (1972) on the use of forecasts in cash management. They proposed evolutionary algorithms to approximately derive cash policies. This solving procedure has been recently followed by Costa Moraes and Nagano (2014).

*Gormley and Meade (2007) hypothesized, but did not prove, that forecasting accuracy leads to cost savings in cash management as we point out in **Question 1**. As a result, we use this model to empirically prove the savings hypothesis. Furthermore, they used an empirical data set providing detailed information about its statistical properties in an first attempt to address **Question 2**. Finally, we use this model as a benchmark to answer **Question 7**.*

2.2.11 The Chen and Simchi-Levi model

The concept of K-convexity was first used by Neave (1970) to show that the Eppen and Fama (1969) model might not be optimal. When there are fixed costs for both inflows and outflows, Chen and Simchi-Levi (2009) used the concept of (K,Q)-convexity by Ye and Duenyas (2007) to provide a characterization of the optimal policy in the stochastic cash balance problem. Their approach was so closely related to inventory control that they used common inventory terminology rather than that usually employed by cash management research. For example, they speak about order and return rather than increase and decrease cash transactions.

They considered a general cost function with holding and transaction costs. At the beginning of each time period, a transaction decision has to be made. Let x be the cash balance at the beginning of time period n before a decision is made, and let y be the cash balance after an transaction was made. The transaction cost is computed as follows:

$$c(x, y) = \begin{cases} K + k(y - x) & \text{if } y > x, \\ 0 & \text{if } y = x, \\ Q + q(x - y) & \text{if } y < x. \end{cases} \quad (2.56)$$

where $K \geq 0$, $Q \geq 0$, and $k + q \geq 0$, assuming that $k \geq q$, i.e., the positive variable transaction cost is greater or equal than the negative variable transaction cost.

On the other hand, the holding cost at time period n is described as a general cost function $l_n(z)$, which depends on the inventory level at the end of the day z which, in turn, depends on the stochastic cash flow ξ_n . Therefore, the expected holding or penalty cost at period n is given by:

$$L_n(y) = E[l_n(z)] = E[l_n(y - \xi_n)] \quad (2.57)$$

In this work, the stochastic cash balance problem is formulated as a dynamic program where $C_n(x)$ is the cost-to-go function at the beginning of a period when there are n periods left in the planning horizon and the initial inventory level is x :

$$C_n(x) = \min_y \{c(y, x) + L_n(y) + \gamma E[C_{n-1}(y - \xi_n)]\} \quad (2.58)$$

where $\gamma \in (0, 1]$ is a discount factor. For clarification, in intervals notation parenthesis or " $()$ " means exclusive and brackets or " $[\]$ " means inclusive. In addition, braces " $\{ \}$ " are used to denote the elements of a set.

They built the process to obtain the optimal policy on the concept of (K, Q) -convexity (Ye and Duenyas, 2007) of the recursive function $C_n(x)$. A real value function is called (K, Q) -convex for $K, Q \geq 0$, if for any x_0, x_1 with $x_0 \leq x_1$ and $\lambda \in [0, 1]$, the following condition holds:

$$f((1 - \lambda)x_0 + \lambda x_1) \leq (1 - \lambda)f(x_0) + \lambda f(x_1) + \lambda K + (1 - \lambda)Q - \min\{\lambda, 1 - \lambda\} \min\{K, Q\}. \quad (2.59)$$

We refer the interested reader to Chen and Simchi-Levi (2009) for both further details about the properties of (K, Q) -convex functions and for a proof that the cost-to-go function $C_n(x)$ is a (K, Q) -convex function. However, a number of additional definitions are required to derive the optimal policy:

$$H_n(x) = L_n(x) + \gamma E[C_{n-1}(x - \xi_n)] \quad (2.60)$$

$$T_n \in \operatorname{argmin}_x \{kx + H_n(x)\} \quad (2.61)$$

$$t_n = \min\{x | kx + H_n(x) = K + kT_n + H_n(T_n)\} \quad (2.62)$$

$$t'_n = \min\{x | kx + H_n(x) = K - Q + kT_n + H_n(T_n)\} \quad (2.63)$$

$$U_n \in \operatorname{argmin}_x \{-qx + H_n(x)\} \quad (2.64)$$

$$u_n = \max\{x | -qx + H_n(x) = Q - qU_n + H_n(U_n)\} \quad (2.65)$$

$$u'_n = \min\{x | -qx + H_n(x) = K - Q - qU_n + H_n(U_n)\} \quad (2.66)$$

where $t_n \leq t'_n \leq T_n$ and $u'_n \leq U_n \leq u_n$. Based on the previous definitions and assuming $K > Q \geq 0$, it is optimal to the set cash level $y_n(x)$, after a decision is made, according to:

$$y_n(x) = \begin{cases} T_n & \text{if } x \leq t_n \\ \in \{x, T_n\} & \text{if } x \in (t_n, t'_n) \\ x & \text{if } x \in [t'_n, u'_n) \\ \in [t'_n, x] & \text{if } x \in [u'_n, u_n) \\ U_n & \text{if } x \geq u_n. \end{cases} \quad (2.67)$$

Summarizing, Chen and Simchi-Levi (2009) followed a sequential decision-making approach using dynamic programming to minimize the total expected costs over the planning horizon. They proposed a model based on bounds, without assuming any particular density function for cash flows but a general one. However, no practical application was given to illustrate the model using a real case.

*Chen and Simchi-Levi (2009) provided a general framework for bound-based models without assuming any particular distribution for cash flows that is related to **Question 2**.*

2.2.12 The Baccarin model

To the best of our knowledge, quadratic holding and penalty costs were considered for the first time in Baccarin (2002). Furthermore, a general multidimensional approach to the cash management problem was also first introduced by Baccarin (2009) using generalized cost functions and providing theoretical results for two bank accounts. Baccarin considered cash management systems with multiple bank accounts where the cash balances fluctuate as a homogeneous diffusion process in \mathbb{R}^n . He formulated the model as an impulse control problem with unbounded cost functions and linear costs.

The optimization problem considering a n -dimensional Wiener cash flow process W_t that determines the dynamics of cash balances $x(t)$ in absence of any control action by means of the following Ito stochastic differential equation:

$$dx(t) = b(x(t))dt + \sigma(x(t))dW_t, \quad x(0) = x \quad (2.68)$$

where $b(x), \sigma(x) \in W^{1,\infty}(\mathbb{R}^n)$. Then, an impulse control strategy within a continuous time framework is a sequence of control actions ξ_i made at time t_i to form policy $V = \{\xi_1, t_1; \dots, \xi_i, t_i; \dots\}$ with $t_i \leq t_{i+1}$. Then, given a policy V , the controlled process $y(t)$ is defined as follows:

$$y(t) = y(0) + \int_0^t b(y(s))ds + \int_0^t \sigma(y(s))dW_s + \xi_1 + \dots + \xi_{\alpha_t}. \quad (2.69)$$

Holding costs are given by some function $f(y)$ and transaction costs by some function $C(\xi)$, which is assumed to be lower semicontinuous and unbounded from above when $|\xi| \rightarrow \infty$. As a result, each control policy V has an associated

cost:

$$J_x(V) = \mathbb{E} \left\{ \sum_{i=1}^{\infty} C(\xi_i) e^{-\gamma t_i} \chi_{t_i < \infty} + \int_0^{\infty} e^{-\gamma s} f(y_x(s)) ds \right\} \quad (2.70)$$

where $\gamma > 0$ is a discount rate and $\chi_{t_i < \infty} = 1$ if $t_i < \infty$, zero otherwise. The problem is then to minimize $J_x(V)$ over the set A of admissible controls V . The optimal control is obtained by dividing \mathbb{R}^n in two complementary regions: a continuation set where the system evolves freely and an intervention set where the system is controlled in an optimal way.

Notice that the n -dimensional cash management system proposed by Baccarin is built by replicating n cash accounts. Control policies are applied to the n cash accounts but no transaction between accounts is allowed in this model avoiding possible control actions within the system itself.

Summarizing, Baccarin (2009) provided a sound theoretical framework for cash management systems with multiple bank accounts within a continuous time framework with general costs functions and an single objective, namely, cost. Cash flows are assumed to follow a Wiener process and the numerical solution to the optimization problem can be obtained by the finite elements method as described in Cortey-Dumont (1985) and Boulbrachene (1998), which consider a discrete approximation of the continuous framework described above. However, transactions between accounts were not considered.

*Baccarin (2009) provided a general framework for cash management systems with multiple bank accounts that is closely related to **Question 9**.*

2.2.13 Alternative cash management models

In this section, we briefly refer to alternative cash management works that received little attention from the research community, in terms of citations, but that we believe they deserve to be mentioned due to interesting characteristics present in their proposal. In addition, Baccarin (2002) and Baccarin (2009) considered the case of cash balance fluctuations following a diffusion process. Smith (1986) and Premachandra (2004) also used a diffusion process to propose generalized versions of the Baumol-Tobin model (Frenkel and Jovanovic, 1980), and the Miller and Orr (1966) model, respectively. When addressing **Question 2**, we empirically test the validity of these assumptions.

On the other hand, Hinderer and Waldmann (2001) formally introduced the concept of environment uncertainty in the CMP by providing a rigorous mathematical framework and by exploring different cases for cash flow processes. Melo and Bilich (2013) proposed an Expectancy Balance Model to minimize combined holding and shortage costs in an attempt to also deal with uncertainty. The authors argued that previous models based on inventory theory or linear programming were not able to represent the complexity of cash management because of their mathematical structure. The procedure of control fluctuations around lower and upper cash balance bounds, namely, a simple policy, does not produce the minimum total cost in the long run due to accuracy failure. This model also considers the existence of both deterministic flows, that are certainly known in advance, and stochastic flows grouped into intervals of occurrence. In this thesis, we address uncertainty introduced by cash flows in a distinct way, namely, by using forecasts (**Question 1**), and also by providing a robust counterpart of cash management models (**Question 5**).

Recently, Costa Moraes and Nagano (2014) proposed the use of genetic algorithms and particle swarm optimization to solve the CMP using the Miller and Orr (1966) model. They provided numerical examples using Gaussian cash flows for both solvers within a structure of a single bank account and two alternative investment accounts. On the contrary, we here aim to provide linear or quadratic programming formulations (**Question 8**) to the CMP from a double multidimensional perspective, namely, multiple bank accounts (**Question 9**) and multiple criteria to incorporate risk analysis (**Question 4**). The underlying motive is to efficiently derive optimal solutions to the CMP by relying on state-of-the-art solvers such as CPLEX or Gurobi, since optimality is not guaranteed by approximate solutions.

2.3 Analysis

One may conclude that all possible situations that cash managers face in their daily tasks have been covered by any of the cash management models that have been proposed so far. However, there exist some rather unexplored topics in the cash management literature as we next point out. In what follows, we summarize the main cash management models presented in the literature according to the six dimensions introduced in Section 1.2, as shown in Table 2.4. By following this approach, we are able to extract useful insights in terms of questions that have received little attention from the research community.

1. **Models.** The use of Bound-Based Models (BBM), whose policies are determined by a set of level or bounds, is a common pattern. From the initial inventory approach to the CMP by Baumol (1952), most models have attempted to derive optimal policies within the framework of some simple policy, typically employing constant cash balance bounds. A slight departure of this framework was considered by Stone (1972) and Gormley and Meade (2007) to introduce forecasts as key inputs to a BBM model. A more practical approach was followed by Archer (1966) to focus on the statistical exploration of data to deal with the lack of synchronization of inflows and outflows.
2. **Cash flow process.** A wide variety of cash flow processes have been considered in the literature, ranging from the uniform and perfectly known cash flow in Baumol (1952) and Tobin (1956), to purely stochastic behavior in Miller and Orr (1966), Eppen and Fama (1969), Constantinides and Richard (1978), Premachandra (2004), Baccarin (2009), and Costa Moraes and Nagano (2014), which usually implies a Gaussian distribution. The selection of any cash flow process implies the assumption of either a continuous time framework (Constantinides and Richard, 1978; Baccarin, 2009) or a discrete time framework (Stone, 1972; Penttinen, 1991; Gormley and Meade, 2007). Since we here follow a data-driven approach based on empirical cash flow data sets, we require a discrete time framework.
3. **Cost functions.** The linear cost assumption is also a common pattern with the exception of Baccarin (2002) and Baccarin (2009), that considered quadratic holding and penalty costs. However, there also exist differences in the linear approach. While Baumol (1952) and Miller and Orr (1966) considered only fixed costs, Tobin (1956) and the rest of subsequent works included fixed and variable costs in their models.
4. **Objectives.** It is also important to note that all models focus on a single objective, namely, cost, neglecting risk analysis. However, the works by Stone (1972), Hinderer and Waldmann (2001), and Gormley and Meade (2007) are remarkable initial attempts to include uncertainty in the analysis of the best policies. The use of forecasts seems to be a sound strategy to reduce uncertainty in the CMP.
5. **Solvers.** There are also differences in the techniques used for solving the CMP. However, three solving techniques stand out: analytic solutions as in Baumol (1952), Tobin (1956), Miller and Orr (1966), Constantinides and Richard (1978), and Hinderer and Waldmann (2001); dynamic pro-

gramming as in Eppen and Fama (1969), Daellenbach (1971), Penttinen (1991), and Chen and Simchi-Levi (2009); and approximate techniques as in Archer (1966), Stone (1972), Gormley and Meade (2007), and Costa Moraes and Nagano (2014).

6. **Bank accounts.** Although cash management systems with multiple bank accounts are the the rule rather than the exception, almost all previous models derive policies for a single bank account and provide no method to extend their results to multiple bank accounts. Only Baccarin (2009) approached the CMP from a multidimensional perspective to deal with multiple bank accounts.

At this point, it is important to recall the set of research questions introduced in Section 1.2 in order to link these questions to previous works. This task ultimately allow us to identify rather unexplored topics in cash management to be covered in this thesis as detailed in Table 2.3.

- **Question 1.** Can cash flow predictive accuracy achieve cost savings in the cash management problem?
- **Question 2.** Are common statistical assumptions of daily cash flow supported by recent empirical data?
- **Question 3.** Is it always possible to achieve a Gaussian, noise-free and linear time-series through data transformations?
- **Question 4.** Can we incorporate risk analysis to the cash management problem?
- **Question 5.** Can we provide a robust counterpart for any cash management model?
- **Question 6.** Under what circumstances or operating conditions a model is better than another?
- **Question 7.** Are control bounds really necessary in cash management?
- **Question 8.** Can we obtain optimal solutions for the multiobjective cash management problem?
- **Question 9.** Can we derive optimal policies for cash management systems with multiple bank accounts?
- **Question 10.** Can we automate decision-making in cash management through the use of specific software?

Question 1 deals with predictive accuracy and with the utility of forecasts in the CMP. Although Gormley and Meade (2007) hypothesized that better forecasts may lead to cost reductions, they did not prove it. We here empirically address this question by considering alternative forecasting methods and a general procedure to estimate cost savings without requiring any forecasting technique. **Question 2** updates the works by Emery (1981), Pindado and Vico (1996), and Gormley and Meade (2007) with recent data from 54 companies in Spain. Gormley and Meade (2007) suggested the use of data transformations to achieve Gaussian cash flows. **Question 3** goes a step further to empirically prove if it is always possible to achieve a Gaussian, noise-free and linear time-series through data transformations. To the best of our knowledge, risk analysis is neglected in cash management literature. Consequently, **Question 4** addresses this question resulting in the main contribution of this thesis. The inherent uncertainty introduced by cash flow variability in the CMP, was approached by Stone (1972) and Gormley and Meade (2007) through the use of forecasts, and by Hinderer and Waldmann (2001) from a theoretical point of view. **Question 5** explores alternative ways of dealing with uncertainty based on data-driven procedures. Daellenbach (1974) raised questions about the utility of cash management models and Constantinides and Richard (1978) provided the necessary conditions to avoid trivial cash policies. **Question 6** extends these works to a multiobjective framework. Since most cash management models present in the literature are BBM, **Question 7** consider Boundless Models in which no restriction is placed on the form of the policy to enlarge the decision variable space. **Question 8** aims to formulate and solve the cash management problem as a linear or quadratic problem ready to be solved by state-of-the-art mathematical programming solvers. Baccarin (2009) approached the multiple bank accounts cash management problem from a theoretical perspective and assuming purely stochastic cash flows. **Question 9** aims to provide a more practical and data-driven approach to derive optimal policies for cash management systems with multiple bank accounts. Finally, **Question 10** intends to fill the gap between theory and practice that is usually neglected in cash management literature.

Research Question	Previous works		
Question 1			
Question 2	Emery (1981)	Pindado(1996)	Gormley (2007)
Question 3	Gormley (2007)		
Question 4			
Question 5	Stone (1972)	Hinderer (2001)	Gormley (2007)
Question 6	Daellenbach (1974)	Constantinides (1978)	
Question 7			
Question 8			
Question 9	Baccarin (2009)		
Question 10			

Table 2.3: Open research questions in cash management

Reference	Model	Cash flow	Cost	Goal	Solver	Accounts
Baumol (1952)	Deterministic	Uniform	Linear (F)	Cost	Analytic	Single
Tobin (1956)	Deterministic	Uniform	Linear (F, V)	Cost	Analytic	Single
Miller-Orr (1966)	BBM	Random walk	Linear (F)	Cost	Analytic	Single
Archer (1966)	Procedimental	Empirical	None	Cost	Simulation	Single
Eppen-Fama (1969)	BBM	Random walk	Linear (F, V)	Cost	DP	Single
Daellenbach (1971)	BBM	Non-stationary	Linear (F, V)	Cost	DP	Single
Stone (1972)	BBM forecast	Predictable	None	Cost	Simulation	Single
Constantinides (1978)	BBM	Diffusion	Linear (F, V)	Cost	Analytic	Single
Penttinen (1991)	BBM	Exponential	Linear (F, V)	Cost	DP	Single
Gormley (2007)	BBM forecast	Empirical	Linear (F, V)	Cost	Evolutionary	Single
Chen-Simchi-Levi (2009)	BBM	Density function	Linear (F, V)	Cost	DP	Single
Baccarin (2009)	Boundless	Diffusion	General	Cost	Finite elements	Multiple

Table 2.4: Characteristics of the main models according to the six dimensions of the problem BBM=Bound-Based Model; F=Fixed; V=Variable; DP=Dynamic Programming.

Part I

Data-driven cash flow analysis

The savings hypothesis

The central question addressed in this chapter is: does a better forecast produce better cash policies in terms of cost? To achieve a deeper understanding of this question, we first explore state-of-the-art machine learning forecasting models such as autoregression, linear regression, radial basis functions, random forests, and a seasonal interaction model. We later rely on time-series cross validation to compare the relative performance of a simple cash management policy derived from each model in terms of savings with respect to a baseline. Finally, we provide a new method to help cash managers estimate if efforts in improving predictive accuracy are proportionally rewarded by cost savings. As a result, cash managers are empowered to test their own savings hypothesis, i.e., if better forecasting models are in place to reduce cash management costs.

3.1 Motivation

In this chapter, we study the performance of cash management models. More precisely we focus on the impact of predictive accuracy on potential cost savings achieved by using cash management models using forecasts as the main input. A measure of quality of any forecasting technique is its predictive accuracy and, under an economic perspective, predictive accuracy must be mapped to estimated cost savings. This analysis assesses how much companies can save by improving predictive models and, consequently, the cost of not predicting, i.e., the missed savings minus the cost of implementing the model. For example, if a reduction of 32% in forecasting error produced 320,000 € in savings per year,

it can be stated that, on average, each percentage point of predictive accuracy is 10,000 € worth.

In this chapter, our discussion seeks to assess the quality of alternative forecasting methods. To this end, we first present and compare different forecasting methods including linear and non-linear models. In this sense, we expect that non-linear models can deal with cash flow time-series in a cost-saving approach. For simplicity reasons, we restrict our analysis to transaction and holding costs for cash balances in a single currency. Using two real data sets from companies in the textile industry in Spain, we develop a proof of concept with respect to the impact of predictive accuracy in the CMP in two steps: first, we show empirically that forecasting accuracy is highly correlated with savings in cash management and, thus, a comparison in terms of accuracy and savings between different forecasting models is performed; second, we propose a method to estimate the effect of forecasting accuracy on cash management in advance.

The rest of this chapter is organized as follows. We firstly describe our two real cash flow data sets in Section 3.2. We later enumerate different forecasting models: linear models such as autoregressive and regression models; and non-linear models, such as radial basis functions, random forests and seasonal indicator models in Section 3.3. These forecasting models will be ranked according to its predictive accuracy in the evaluation Section 3.4. In Section 3.5, we empirically verify that a better forecast produces a better policy. Moreover, we estimate how much savings (if any) can be obtained by the cash policies produced by an improvement in forecasting accuracy. Finally, Section 3.6 summarizes the chapter.

3.2 Description and data preprocessing

In this section, we describe the two real cash flow data sets used in this chapter. Data sets 1 and 2 gather net daily flow on workdays from two different companies in the textile industry. Both sets of observations are in the domain of real numbers and their values' distributions present a bell-like shape but excess kurtosis. Besides, an additional transformation is performed to deal with anomalies. More specifically, following the recommendations in Gormley and Meade (2007) and Hyndman (2016), any observation greater than five times the standard deviation is considered an outlier and it is replaced by a linear interpolation. After time-series cleaning, our cash flow data sets empirical distributions are shown in Figure 3.1. In order to cover a wider range of realistic

industrial company cases, a number of cash flow data sets are derived from the two real data sets as follows:




Figure 3.1: Histogram for data sets 1 and 2 compared to a normal distribution.

- Real cash flow: data sets 1 and 2.
- Stable cash flow: data set 3 is derived from data set 1 and applies to companies in a more stable environment with daily cash flows characterized by a low variance. In this case, observations greater than three times the standard deviation are replaced by values of exactly three times the standard deviation.
- Unstable cash flow: data set 4 also derives from data set 1 and applies for companies with high variances in their daily cash flow due to different reasons such as a reduced number of customers or suppliers as it is the case of small companies. In this case, observations greater than three times the standard deviation are replaced by values of exactly two times the original observation.
- Random shock cash flow: data sets 5 and 6 are derived from the original data sets 1 and 2 respectively, before anomaly detection, and aim to cover the likely occurrence of unexpected changes in industrial markets. In this case, 5% of the observations are randomly chosen and replaced by randomly selected values from the previous set of labeled outliers. This transformation simulates a disturbance that, to some extent, may break existing patterns and increase deviation from expected cash flow

by introducing a number of outliers in the original data. In addition, it helps estimate the impact of outliers in accuracy.

A summary of the characteristics of each data set is presented in Table 3.1.

Data Set	Length	Case	Std.Dev.	Kurtosis
1	2717	Real cash flow	92615	4.08
2	1218	Real cash flow	42514	3.55
3	2717	Stable cash flow	87780	2.40
4	2717	Unstable cash flow	126303	18.05
5	2717	Random shock cash flow	151462	6.49
6	1218	Random shock cash flow	100883	20.33

Table 3.1: Data set summary.

For comparison purposes with Gormley and Meade (2007), here we assume that, apart from daily cash flow data, no other extra features are provided by the company. However, a set of available explanatory variables can be proposed. According to Miller and Stone (1985), and Stone and Wood (1977), we may find seasonal patterns in daily cash flow data. Thus, we consider basic calendar effects such as the day-of-week, the day-of-month and month effect by using categorical or dummy variables. In the latter case, each dummy variable takes a value of one if time t occurs at the corresponding day of the week/month, and zero otherwise. A further step in the search of explanatory power is explored by considering past values of the time-series. From that, a tentative set of explanatory variables is listed below:

- Day-of-month (DOM): Day of month categorical variables or dummy variables (d_{t1}, \dots, d_{t31}) .
- Day-of-week (DOW): Day of week categorical variables or dummy variables for working days (s_{t1}, \dots, s_{t5}) .
- Month: Month dummy variables (m_{t1}, \dots, m_{t12}) .
- Week: Week dummy variables (w_{t1}, \dots, w_{t53}) .
- Past values: Previous observations of the daily cash flow time-series $(y_{t-1}, \dots, y_{t-p})$ where p is the total number of observations considered.

From a combination of these explanatory variables, different predictive models can be built and compared in terms of forecasting accuracy.

3.3 Building forecasting models

The accuracy of any forecasting model depends on its ability to capture the specific characteristics of the data used. According to Stone (1972), real-world cash flows are neither completely known in advance nor are they completely unpredictable. However, a wide set of tools and techniques are available to improve forecasting accuracy. We claim that exploring alternative models to improve forecasting ability is mandatory, specially if improving forecasting accuracy can lead to cost savings.

In this section, we present a number of forecasting models to be evaluated allowing us to identify our best-in-class forecaster that will ultimately be used as the main input to the cash management model. We do not intend to determine the best cash flow forecaster among all methods presented in forecasting research literature. Instead, our final goal is to verify if a better forecaster, in terms of forecasting accuracy, is able to produce a better cash policy in terms of cost savings. In this sense, we expect that non-linear models outperform two of the most usual linear models in cash flow forecasting allowing cash managers to deploy better cash policies.

Then, we here consider five different forecasting models: autoregressive, regression, radial basis functions, random forests and seasonal interaction models. Firstly, we follow the autoregressive approach to daily cash flow forecasting along the lines of Gormley and Meade (2007). In contrast to such approach, we expect that the use of the set of explanatory variables mentioned earlier rather than only a number of previous values of a time-series can help obtain a more accurate prediction. Then, we secondly consider a linear regression model with a set of explanatory variables.

While linear models are often employed in finance due to their simplicity, many non-linear models have been proposed to explain financial phenomena. Perhaps one of the most widely known non-linear model in finance is the Black and Scholes (1973) option pricing model. Moreover, there is a reason for optimism about the use of non-linear models in time-series prediction and finance as stated in Weigend (1994), Kantz and Schreiber (2004), and Small (2005). Firstly, several limitations of linear models were pointed out by Miller and Stone (1985) in daily cash flow forecasting such as interactions and holiday effects. Additionally, statistical hypothesis such as normality and stationarity are required by linear models to produce reliable results. On the other hand, alternative approaches to discover the relationship between time-series observations are also available. In this sense, non-linear models allow to explore

beyond the constraints imposed by linear models through a much wider class of functions.

Although non-linear time-series analysis is not as well established as its linear counterpart (De Gooijer and Hyndman, 2006), works by Teräsvirta (2006), Bradley and Jansen (2004), Clements, Franses, and Swanson (2004), Sarantis (2001), and Conejo et al. (2005) constitute good examples of its application to finance and economics. In this chapter, we consider non-linear models such as radial basis functions and random forest models due to the lack of attention of the research community. Next, we briefly describe our selection of forecasters and provide details on the implementation of non-linear models.

3.3.1 *Autoregressive model*

A widespread linear model in time-series data is the autoregressive (AR) process, where predictions are based on a linear combination of previous observations (Box and Jenkins, 1976). AR models for cash flow forecasting can be found in Gormley and Meade (2007) and Laukaitis (2008). On the other hand, as mentioned in Section 3.2, our cash flow data have a bell-like shape but excess kurtosis. Hence, we follow the recommendations in Gormley and Meade (2007) and use an extension of the Box-Cox transformation described in Box and Cox (1964) to approximate our data to a Gaussian distribution by tuning a parameter λ . Predictions are assessed using the following equation:

$$y_t^{(\lambda)} = \beta_0 + \sum_{i=1}^p \beta_i y_{t-i}^{(\lambda)} + \epsilon \quad (3.1)$$

where $y_t^{(\lambda)}$ is the cash flow forecast at time t , $[y_{t-1}^{(\lambda)}, y_{t-2}^{(\lambda)} \cdots, y_{t-p}^{(\lambda)}]$ are the p -previous observations of a transformed time-series, β_i is the i -th estimation coefficient, and ϵ stands for the prediction error. Superscript (λ) in both forecasts and previous observations denotes data transformation. This transformation is reversible and, therefore, y_t can be derived from $y_t^{(\lambda)}$.

3.3.2 *Regression model*

An autoregressive model is only based on the previous observations of the time-series and misses possible patterns, if any, hidden in the data. When dealing with daily data, these patterns refer to calendar variations such as holidays, day of the month or day of the week. Trying to identify these patterns, here we consider a general regression model based on different explanatory variables. Regression models have been used for cash flow forecasting purposes in Stone

and Wood (1977), Stone and Miller (1987), and Miller and Stone (1985). In this case, it is important to say that the ability of the modeler in the search for the best explanatory variables plays a key role. A general regression model is represented by the following equation:

$$y_t = \sum_{i=0}^n \beta_i x_{ti} + \epsilon. \quad (3.2)$$

In this general model (3.2) we relate y_t , the value of the cash flow at time t to a linear combination of explanatory variables $x_{t1}, x_{t2}, \dots, x_{tn}$ at the same time t , being β_i the i -th regression coefficient, and ϵ the prediction error. From the general model (3.2), a number of particular models can be derived for predictions depending on the different explanatory variables considered. For the implementation of these models we use the *lm* function in R.

3.3.3 Radial basis function model

Financial data are usually originated by complex systems that may include non-linear processes. In order to capture non-linearity in the data, we also consider Radial Basis Function (RBF) models as described in Weigend (1994) and Broomhead and Lowe (1988). To use an RBF model, we first partition the input space by applying the k -medoids algorithm (Park and Jun, 2009) over the training set. Then a scalar Gaussian RBF $\phi(x)$ is used for forecasting:

$$y_t = b_0 + \sum_{k=1}^K b_k \phi(\|x_t - c_k\|) + \epsilon \quad (3.3)$$

where y_t is the value of the target variable at time t , K is the total number of clusters, b_k is the coefficient associated to the k -th cluster, c_k is the k -th cluster medoid, x_t is the input data point at time t , $\| \cdot \|$ is the Euclidean distance and ϵ is the prediction error. Finally, $\phi(x)$ is the following Gaussian function:

$$\phi(x, \alpha) = e^{-x^2/\alpha\rho_k} \quad (3.4)$$

where α is a positive integer parameter and ρ_k is the mean distance between the elements inside the k -th cluster. In this case, predictions are produced using our tentative set of explanatory variables and general matrix functions in R.

Next, we provide an example of predictions obtained using RBF for the last 3 days of data set 1 based on the previous 21 cash flow observations. Two

parameters have to be chosen to produce forecasts using RBF: the total number of clusters K , defining the degree of partition of the input space and the parameter α , determining the contribution of deviate points to the prediction. For the sole purpose of this example, we set $K = 5$ and $\alpha = 10$. Then, we proceed as follows:

1. We create the input space $(2714 - 21) \times 21$ matrix X by embedding in each row 21 consecutive cash flows. Firstly, we transform cash flows to $y_t^{(\lambda)}$, as explained in Section 3.3.1. To avoid high values bias, we later standardize transformed cash flows by demeaning and dividing by the standard deviation.
2. We create a column vector \mathbf{y} of length 2693 with subsequent cash flows.
3. We select each cluster medoids c_k from rows in X using the k-medoids algorithm.
4. We compute ρ_k as the mean Euclidean distance between the elements of the k -th cluster to its medoid c_k .
5. We compute the 2693×5 matrix Φ where each row contains distances computed using the function $\phi(\|x_t - c_k\|)$ for each point in the input space to each cluster.
6. We obtain the column vector \mathbf{b} of weights by solving $\mathbf{b} = (\Phi^T \Phi)^{-1} \Phi^T \mathbf{y}$ using least squares.
7. We produce a 3×21 matrix \hat{X} with the previous 21 observations prior to each of the 3 cash flows to be predicted and a 3×6 matrix $\hat{\Phi}$ where the first column is set to 1 and the rest of elements are distances computed using the function $\phi(\|x_t - c_k\|)$ for each point in \hat{X} to each cluster medoid.
8. We forecast by means of $\hat{y} = \hat{\Phi} \mathbf{b}$, that has to be re-scaled by multiplying by the standard deviation and adding the mean and, finally, λ -transformed.

Now, we are in a position to compare these forecasts to real values and to other forecasts and compute predictive accuracy as we will see below.

3.3.4 Random forest model

Decision trees are non-linear models that split the input space in subsets based on the value of a particular feature. On the other hand, an ensemble methodology is able to construct a predictive model by integrating multiple trees in what is called a decision forest (Dietterich, 2000). Regression forests are used for the non-linear regression of dependent variables given independent inputs based on an ensemble of slightly different trees. Particularly, random forests (RF) are ensembles of randomly trained decision trees (Ho, 1995; Ho, 1998; Criminisi and Shotton, 2013). Recent examples of time-series forecasting using random forests can be found in Kumar and Thenmozhi (2006), Kane et al. (2014), Mei et al. (2014), and Zagorecki (2015).

We make predictions using the R package *randomForest* by Liaw and Wiener (2002) which implements Breiman's random forest algorithm for classification and regression (Breiman, 2001). In this chapter, we limit ourselves to select three parameters: the number (a) of randomly trained trees, the number (b) of variables randomly sampled as candidates at each split, and the node size (c) or the minimum amount of observations in a terminal node used to control overfitting.

For instance, assume that we know there is a strong daily seasonality in our cash flow. One possible way to assess how strong is this seasonality is to produce predictions using two explanatory variables: Day-of-month and Day-of-week. Hence, we aim to create a random forest model and predict the last 100 days of data set 1 based on these two variables. An example on how to proceed is as follows:

1. Create a 2617×2 matrix X containing in each row the Day-of-month and the Day-of-week of past cash flows.
2. Create a column vector \mathbf{y} of length 2617 with the corresponding cash flows.
3. Create a model based on X and \mathbf{y} , with $a = 100$ randomly trained trees, with $b = 2$ randomly sampled variables and node size $c = 50$.
4. Produce a 100×2 matrix \hat{X} with the Day-of-month and the Day-of-week of last 100 cash flows of data set 1.
5. Input matrix \hat{X} to the model to obtain the forecasts.

3.3.5 Seasonal interaction model

Miller and Stone (1985) proposed a method to daily cash flow forecasting by spreading an estimated monthly total over the days of this time period. This distribution approach requires the separation of net cash flow in components, at least, inflows and outflows, that is a different case than the data set in our study. However, we rely on the Miller and Stone (1985) model to forecast cash flow y_t , ranging in $t = 1, 2, \dots, T$, based on seasonal interactions as follows:

$$y_t = \sum_{j=1}^T \delta_j I_{j,t} + \epsilon \quad (3.5)$$

where $I_{j,t}$, is a seasonal indicator (SI) that takes value one, when a given seasonal condition holds, zero otherwise. For instance, we can define $I_{1,t} = 1$, when the day t occurs on the first of January and Monday. To account for all possible combinations of working days and months, we initially consider a full set of $31 \times 5 \times 12 = 1860$ different seasonal indicators. To speed up computations, we estimate δ_j by averaging cash flow grouped by day-of-month, day-of-week and month. When no data is available for estimation purposes, we set $\delta_j = 0$, and forecast using the estimated mean. Simplicity, lack of estimation issues and account for interactions are the main advantages of this model. Now, we are in a position to assess the importance of each explanatory variable or to test the quality of our predictions.

3.4 Forecasting models' comparison

In this section, we aim to evaluate the forecasting accuracy of the presented models for comparison purposes. Alternative models may produce different predictions with different accuracy. The comparison will allow us to determine our best-in-class forecaster to be later used as the input to establish the best cash management policy available. More precisely, we use time-series cross-validation for different prediction horizons (h) from 1 up to 100 days ahead by comparing the mean square error $\varepsilon(h)$ for different models:

$$\varepsilon(h) = \frac{\sum_{t \in \text{test}} (\hat{y}_{t+h} - y_{t+h})^2}{\sum_{t \in \text{test}} (\bar{y} - y_{t+h})^2} \quad (3.6)$$

where h is the prediction horizon in days, \hat{y}_{t+h} is the prediction at time $t + h$, y_{t+h} is the real observation at the time $t + h$, and \bar{y} is the arithmetic mean of the real observations on the training set. Note that the closer ε is to zero, the better the predictive accuracy. If ε is close to one, the performance is

similar to that of the mean as a naive forecast. Values greater than one show that the forecaster has no predictive ability.

In Hyndman and Athanasopoulos (2013), two different time-series cross validation approaches were suggested: one with a fixed origin for the training set, and one with a rolling origin. Algorithm 1 implements these two cross-validation methods.

Algorithm 1: time-series cross validation algorithm

- 1 **Input:** Cash flow data set of T observations, *FixedOrigin*, minimum number g of observations to forecast and prediction horizon h ;
 - 2 **Output:** Forecast accuracy for different prediction horizons;
 - 3 **for** $i = 1, 2, \dots, T - g - h + 1$ **do**
 - 4 Select the observation at time $g + h + i - 1$ for the test set;
 - 5 **if** *FixedOrigin* = *True* **then**
 - 6 Estimate the model with observations at times $1, 2, \dots, g + i - 1$;
 - 7 **else**
 - 8 Estimate the model with observations at times $i, i + 1, \dots, g + i - 1$;
 - 9 **end**
 - 10 Compute the h -step error on the forecast for time $g + h + i - 1$;
 - 11 **end**
 - 12 Compute $\varepsilon(h)$ based on the errors obtained;
-

If binary variable *FixedOrigin* is set to True, the training set is formed by all the observations that occurred prior to the first observation that forms the test set (Method 1). We can get rid of the oldest observations by setting *FixedOrigin* to False (Method 2) and considering only the g most recent values (e.g., the last two or three years) by applying a sliding window of observations. In both methods we assume that the minimum number of g observations required to produce a reliable forecast is the first 65% of the data. In our experiments, high values of g produced almost no difference between Method 1 and Method 2. Using Method 2, smaller values of g in steps of 250, equivalent to 1 year of observations, were also tried with worse results. Because of that, here we only present results for Method 1.

For model selection, we follow an automatic selection method along the lines of Doornik (2008), Hyndman and Athanasopoulos (2013), and Hendry and Doornik (2014). More precisely, we perform a backwards stepwise regression starting with a model containing all potential predictors and removing one predictor at a time. We select the model with the minimum average error $\bar{\varepsilon}$, for prediction horizons ranging in $1, 2, \dots, 100$, computed using Algorithm 1.

For statistically significant differences in performance between models, the U Mann-Whitney test for independent samples with a 95% confidence interval is used (Hollander, Wolfe, and Chicken, 2013). In addition, when parameters selection is necessary, an evaluation of the coefficient of determination, R^2 , over a training set with the oldest 65% of the observations is used to choose the best value for each of the parameters. The final model selection, summarized in Table 3.2, shows that the day-of-month and the day-of-week variables present the best forecasting ability. However, the number of coefficients that are relevant in the case of autoregression and regression using dummy variables are different for data set 1 and 2.

Model	Input variables	Data Set 1	Data Set 2
AR	Past values	12 coefficients	15 coefficients
REG	DOM, DOW	29 coefficients	10 coefficients
RBF	DOM, DOW	$K = 35, \alpha = 10$	$K = 10, \alpha = 10$
RF	DOM, DOW	$a = 20, b = 11, c = 50$	$a = 20, b = 11, c = 50$
SI	DOM-DOW Indicators	155 coefficients	155 coefficients

Table 3.2: Model selection according to average error ratio ($\bar{\varepsilon}$) for horizons up to 100 days.

The relative performance for different prediction horizons using models from Table 3.2 is plotted in Figures 3.2 and 3.3, and results for all data sets using the previously selected models are shown in Table 3.3. Autoregressive models performed no better than the mean as a naive forecaster. A statistically significant difference in average forecasting accuracy was found in favor of the RF model for data sets 1 and 4, whereas it was almost equal to the SI model in the case of data set 3. However, the performance of the regression model was the best in data set 2, suggesting a linear behavior. Interestingly, RBF were less affected than the rest of models by the introduction of outliers in data sets 5 and 6. The poor performance of the SI model on data sets 2 and 6 may be caused by the smaller number of the observations than in the rest of data sets. Summarizing, it is clear that forecasting accuracy of the autoregressive model can be improved by considering alternative models. Next, we measure the savings produced by a better prediction.

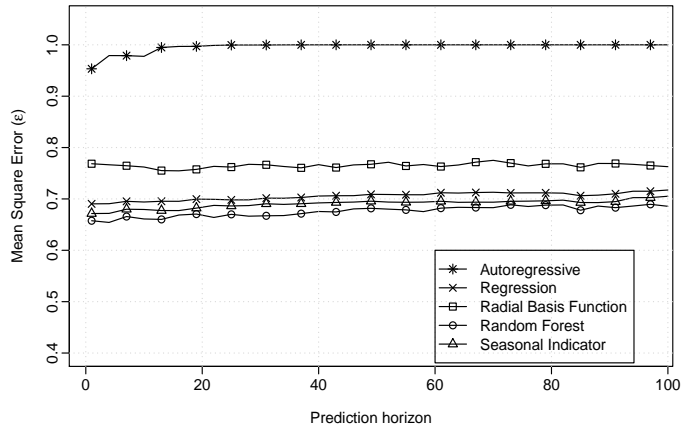


Figure 3.2: Mean square error comparison for different predictive models (Data set 1).

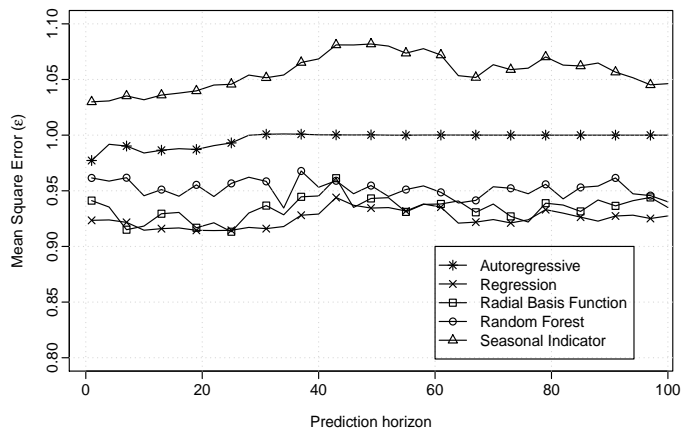


Figure 3.3: Mean square error comparison for different predictive models (Data set 2).

3.5 Does a better forecast produce better policies?

Gormley and Meade (2007) proposed a Dynamic Simple Policy (DSP) to demonstrate the utility of cash flow forecasts in the management of corporate cash balances. They proposed the use of an autoregressive model as the main input to their model. However, gains in forecast accuracy over a naive mean model were scant. Gormley and Meade expected that savings obtained using a non-naive forecasting model would increase if there were more sys-

Set	AR	REG	RBF	RF	SI
1	0,997 (0,008)	0,705 (0,007)	0,765 (0,005)	0,678 (0,010)	0.690 (0.008)
2	0,997 (0,008)	0,925 (0,008)	0,934 (0,009)	0,952 (0,008)	1,057 (0,015)
3	0,997 (0,009)	0,696 (0,007)	0.786 (0.019)	0,661 (0,023)	0.667 [*] (0.023)
4	0,998 (0,006)	0,766 (0,008)	0.785 (0.004)	0,754 (0,010)	0.785 (0.012)
5	1,000 (0,000)	0,920 (0,004)	0.859 (0.070)	0,914 (0,006)	0.942 (0.004)
6	1,001 (0,002)	1,003 (0,003)	0.984 (0.004)	1.003 (0.004)	1.217 (0.020)

Table 3.3: Average predictive error for prediction horizons from 1 to 100 days using Method 1. Standard deviations are shown in parenthesis and best values are bold. Non-significant different models to the best in each row with a 95% confidence interval are marked with ^{*}.

tematic variation in the cash flow and, consequently, higher forecast accuracy. In the previous section, we showed that a better cash flow prediction can be obtained by using different forecast models. In this section, we verify that a better prediction produces a better policy. As a consequence, we find that the savings produced by a better forecasting model are significantly higher than those obtained by a naive forecasting model.

Here, we exploit a simple policy equivalent to that of Gormley and Meade using the best forecasters as detailed in Section 3.4 and compare to the results obtained by a constant mean forecast. This control limit framework is limited to cash balance holding and transaction costs. Other major benefits derived from forecasting accuracy such as short-term investment improvement and prearranged credit lines cost savings are not considered in this chapter. Better forecasting accuracy allows companies to invest in less marketable securities but with higher returns if held to maturity. In addition, companies relying on credit lines can reduce the amount of prearranged credit if better forecasts are available. A good example of cost-benefit analysis considering short-term investment and borrowing costs can be found in Stone (1973).

On the other hand, since the forecast accuracy of the autoregressive model almost equals the mean forecast accuracy (Table 3.3), the comparison to the mean is equivalent to the comparison to the autoregressive model. In what follows, we firstly introduce our empirical settings; secondly, we show empirically that forecasting accuracy leads to cost savings in the corporate cash management problem using a simple policy; and finally, we analyze potential cost savings of improving predictive accuracy of daily cash flow forecasts and a simple policy.

3.5.1 Empirical settings

The corporate cash management problem can be approached from a stochastic point of view by allowing cash balances to wander between two limits: the lower (D) and the upper balance limit (V). When the cash balance reaches any of these limits a cash transfer is made to return to the corresponding rebalance level (d, v). A model of this kind using daily forecasts was proposed by Gormley and Meade (2007) as a trade-off between transaction and holding costs as follows: q is the holding cost per money unit of positive balances at the end of the day; u is the shortage cost per money unit of negative balances at the end of the day; γ_0^+ is the fixed cost of transfer into account; γ_0^- is the fixed cost of transfer from account; γ_1^+ is the variable cost of transfer into account; and γ_1^- is the variable cost of transfer from account.

Recall from section 3.2 that we are dealing with a real business problem. Hence, we focus on current costs charged by banks to industrial companies in Spain. Current bank practices tend to charge a fixed cost for transfers between €1 and €5 and no variable cost so that we set $\gamma_1^+ = 0$ and $\gamma_1^- = 0$. The shortage cost (u) per money unit of a negative cash balance is around 30% which represents a high penalty for negative cash balances. Finally, the holding cost (q) per money unit of a positive cash balance is an opportunity cost of returns not obtained from alternative investments. Since this is not an actual cost but an opportunity cost based on judgmental criteria, we set a range between 3% p.a.¹ and 20% p.a. based on the concept of weighted average cost of capital. This wide range allows us to consider not only the cost of debt but also the cost of equity as an estimate for the holding cost. We firstly try 20 different cost structures considered as the most likely scenario under current costs in Spain, denoted by (1) in Table 3.4. We also consider two additional scenarios, denoted by (2) and (3), to evaluate the effect of changes in particular costs. The second scenario tests the variation of the shortage cost (u) and the third one considers the introduction of variable transfer costs (γ_1).

In our experiments, parameter selection of the cash management model is performed under a business perspective. In Gormley and Meade (2007) the policy parameter values D, d, v, V were chosen to minimize the expected cost over horizon T using a genetic algorithm (Chelouah and Siarry, 2000). Here, since the focus is placed on the comparison between policies obtained from different forecasting models, parameter optimization plays a secondary role. Therefore, these parameters are empirically chosen and kept unaltered in the comparison between savings for each forecasting model and each cost scenario.

¹Per annum.

Cost	Alternative Scenarios		
	Scenario (1)	Scenario (2)	Scenario (3)
Holding cost q	3, 10, 15, 20 %	15 % p.a.	15 % p.a.
Shortage cost u	30 %	10, 20, 40 %	30 %
Fixed into account γ_0^+	1, 2, 3, 4, 5 €	3 €	3 €
Fixed from account γ_0^-	1, 2, 3, 4, 5 €	3 €	3 €
Variable into account γ_1^+	0 %	0 %	0.01, 0.02, 0.04 %
Variable from account γ_1^-	0 %	0 %	0.01, 0.02, 0.04 %

Table 3.4: Cost scenarios. (1) Most likely scenario; (2) Varying the shortage cost u ; (3) Introduction of the variable cost γ_1 .

However, in order to evaluate the influence of these parameters on the utility of the forecast, three different cases are studied based on risk tolerance. Since the cost of a negative balance is very high, common sense leads us to set D to a minimum level so that only a given percentage ($MaxPct$) of expected cash flows can bring the balance from value D to a negative value. The higher the percentage, the higher the probability of an overdraft and, consequently, the riskier the policy under these cost structures. We study three cases with different levels of risk: (i) Low risk or $MaxPct = 5\%$; (ii) Medium risk or $MaxPct = 10\%$; (iii) High risk or $MaxPct = 15\%$.

On the other hand, the use of dynamic simple policy assumes that an unlimited cash buffer is available to transfer into the bank account whenever it is necessary. In practice, this situation is unrealistic. Thus, we restrict high balance levels by setting an upper limit to 1.5 times the lower cash balance limit. Following the recommendations in Gormley and Meade (2007), the positive shift from the lower balance limit (D) of lower rebalance level (d) is proportional to the difference between the higher (V) and the lower balance (D) limits. Finally, the negative shift from the higher balance limit (V) to obtain the higher rebalance level is proportional to the difference between the higher balance limit (V) and the lower rebalance level (d). Here, we chose proportionality constants $\alpha_1 = 0.5$ and $\alpha_2 = 0.5$ to produce an even distance between policy parameters. The entire analysis would remain the same when varying this setting. As a summary, parameters selection is done according to:

- $D = |o_{th}|$ where o_{th} is the $N \cdot MaxPct$ -th element of vector o_t of ascending ordered values of cash flow being N the total number of observations.
- $V = 1.5D$, then $V - D = \frac{D}{2}$
- $d = D + \alpha_1(V - D)$ with $\alpha_1 = 0.5$

- $v = V - \alpha_2(V - d)$ with $\alpha_2 = 0.5$.

Predicted cash flows using different forecasters are used to compare the effect on the total cost over different prediction horizons (h) from 1 up to 100 days ahead. We set g to the minimum number of observations required to estimate the model that is equivalent to 65% of the data set. We proceed as detailed in Algorithm 2.

Algorithm 2: Comparison algorithm

```

1 Input: Cash flow data set of  $T$  observations,  $g, h, MaxPct$ , and a forecaster;
2 Output: Average cost difference between a forecaster and the mean as a
   forecast;
3 for  $i = 1, 2, \dots, T - g - h + 1$  do
4   | Estimate the model with observations at times  $1, 2, \dots, g + i - 1$ ;
5   | Predict for times  $g + i$  up to  $g + h + i$  using the forecaster;
6   | Predict for times  $g + i$  up to  $g + h + i$  using the mean forecaster;
7   | for  $j = 1, 2, \dots, Number\ of\ cost\ structures$  do
8     | | Compute cost for the  $i$ th forecast when using the  $j$ -th structure;
9     | | Compute cost for the  $i$ th mean forecast and the  $j$ -th structure;
10  | end
11 end
12 Compute average cost for each cost structure using the forecaster;
13 Compute average cost for each cost structure using the mean forecaster;
14 Compute difference between average cost of the mean and the forecaster;
```

3.5.2 Impact of predictive accuracy on cost savings

Cost savings are computed as the daily average cost differences between the naive forecast and the best-in-class forecaster for each of the data sets (Table 3.5). Recall that this comparison to the mean is equivalent to the comparison to the autoregressive model. From these results, we can say that, in general, an increase in forecast accuracy leads to significant cost savings using a simple policy. A better forecasting model produces higher savings for either conservative or riskier policies. The effect of forecasting accuracy in daily costs dramatically rises when the policy bounds are reduced as a consequence of a riskier policy. In these cases, forecasting accuracy is much more important in reducing daily cost due to the risk of an overdraft. As expected, cost reductions for the data set 2 are smaller but still significant due to less predictive accuracy.

Data set	Best	Scenario	Low Risk	Medium Risk	High Risk
1	RF	Most likely	183 (71%)	1432 (74%)	2039 (55%)
1	RF	Varying u	143 (55%)	1115 (72%)	1587 (54%)
1	RF	Introducing γ_1	184 (63%)	1434 (73%)	2040 (55%)
2	REG	Most likely	85 (27%)	363 (34%)	448 (25%)
2	REG	Varying u	66 (25%)	282 (33%)	349 (25%)
2	REG	Introducing γ_1	84 (26%)	362 (33%)	448 (25%)
3	RF	Most likely	181 (71%)	1422 (74%)	2025 (55%)
4	RF	Most likely	207 (57%)	1437 (74%)	2156 (53%)
5	RBF	Most likely	455 (25%)	1953 (49%)	3292 (55%)
6	RBF	Most likely	-11 (7%)	173 (19%)	174 (13%)

Table 3.5: Average daily saving for different levels of risk and alternative scenarios. RF=Random Forest; REG=Regression; RBF=Radial Basis function; u = shortage cost, γ_1 = variable transaction cost.

A deeper insight on the different scenarios shows that changes in cost parameters have a reduced impact on cost savings. A scenario of particular importance nowadays is that of low holding costs such as 3% p.a. due to current low interest rates. Our results show that, even in such a scenario, 88%, 75% and 55% savings for the three levels of risk can be obtained in data set 1 using random forests forecasts. Moreover, changes in the variability of cash flow data, studied here by introducing less (data set 3) or more variance (data set 4), produced no major changes. However the effect of random shocks in the cash flow data (data sets 5 and 6), reduced cost savings due to the higher uncertainty of the cash flow data.

3.5.3 Analyzing potential cost savings

Our best-in-class forecasting models, i.e., regression, radial basis functions and random forests models, are attempts to reduce uncertainty in predicting daily cash flow. They represent special cases in which improving predictive accuracy resulted in increasing cost savings over a naive forecast. However, cash managers may be interested in determining how much savings can be achieved by any extra effort in improving predictive accuracy. Since enhancing any forecasting model has a cost in terms of both time and money, it is important to know if this cost is offset by the savings obtained using a better forecasting. We can estimate savings associated to predictive accuracy by obtaining a number of synthetic predictions and evaluate the corresponding policy costs.

Daellenbach (1974) and Costa Moraes and Nagano (2014) synthesized cash flow data for simulation purposes from normal distributions. Here, from a given cash flow time-series (y_{t+h}), a new time-series (\hat{y}_{t+h}) is synthesized by adding a random normal term of mean zero and a variable standard deviation (σ) using the following equation:

$$\hat{y}_{t+h} = y_{t+h} + \mathcal{N}(0, \sigma). \quad (3.7)$$

Increasing the value of σ , a set of time-series with a decreasing degree of similarity to the original time-series can be obtained. This is equivalent to generating a set of synthetic predictions with controlled predictive accuracy that can be evaluated in terms of mean square error ratio $\varepsilon(h)$ for different prediction horizons using equation (3.6). We obtain synthetic predictions covering a range from $\bar{\varepsilon} = 0$ to values greater than 1. Here $\bar{\varepsilon}$ denotes the average of $\varepsilon(h)$ for prediction horizons up to 100 days on a test set formed by the last 35% of the observations of data sets 1 and 2. Later, savings for each of these synthetic forecasts are obtained following Algorithm 2 but using the synthetic forecasts previously generated rather than estimating and predicting.

Results from this simulation for data sets 1 and 2 and three different levels of risk are shown in Figures 3.4 and 3.5. As a reference, the vertical lines locate savings achieved by the best-in-class forecaster for each of the examined levels of risk. For example, using random forests for data Set 1, a value of $\bar{\varepsilon} = 0.68$ (from Table 3.3) was obtained which produced savings of 71, 74 and 55% (from Table 3.5) for the three levels of risk considered.

As expected, improving prediction accuracy, i.e., reducing $\bar{\varepsilon}$, leads to an important increase in cost savings up to 100% in the case of a perfect prediction. Efforts in increasing predictive accuracy are notably rewarded. However, the behavior is different depending on the level of risk chosen by the company.

1. Low risk: the effect of improving predictive accuracy tends to a stable point where any further effort yields no additional saving. In spite of the considerable percentage saved, it seems that improvement potential in predictive accuracy is limited when the risk is low.
2. Medium risk: the effect of limited cost savings when improving predictive accuracy is also present but to a lesser extent.
3. High risk: the behavior is almost linear in the considered $\bar{\varepsilon}$ interval.

It is interesting to point out that the relationship is almost linear in most of the range of the average error $\bar{\varepsilon}$ for each of the three levels of risk. This fact should




Figure 3.4: Savings for different predictive errors and levels of risk for data set 1 in the most likely scenario. S=Synthetic forecasts, RF=Random Forest.




Figure 3.5: Savings for different predictive errors and levels of risk for data set 2 in the most likely scenario. S=Synthetic forecasts, RBF=Radial Basis Function.

encourage practitioners to work hard to obtain a better prediction because they can expect a proportional reward in terms of cost savings. However, in the case of our best-in-class forecasting model using random forests and data set 1, an error $\bar{\epsilon}$ of 0.68 (from Table 3.3) places the savings in the highest value likely to be obtained for the low level of risk. Any effort in improving predictive

accuracy will be useless. This behavior is confirmed by the fact that a perfect prediction was unable to achieve a 100% difference in cost savings.

Summarizing, we propose a new and more comprehensive methodology (shown in Figure 3.6) for the practitioner, i.e., the cash manager, based on the effect of predictive accuracy on cash management cost using daily cash flow forecasts and a simple policy. In order to allow different models to capture patterns, cash managers should consider an additional previous step of feature engineering to obtain a series of extra features. They can also adopt a wider modeling approach that allow them to compare a set of forecasters in terms of forecasting accuracy. At this point, cash managers can easily generate a number of synthetic predictions to cover a wide range of different predictive accuracy by tuning a parameter. These synthetic predictions, and those obtained using our best-in-class forecasters from Table 3.3, are tested in their ability to reduce the cost of the policies by using a simple policy. This step results in a graphical estimation on how much cost savings can be achieved by improving predictive accuracy of our selected forecasters. If estimated savings are greater than the cost of improving the accuracy of the forecasting models, a new modeling process is worth undertaking.

3.6 Summary

From the above-described results, we derive two main findings. First, assessing predictive accuracy is a must in the context of corporate cash management, specially when employing daily forecasts as an input to a cash flow management model. Indeed, we empirically find that cost savings are highly sensitive to improvements on prediction accuracy when using a simple policy, and hence major savings may stem from accurate predictions. Second, from a cost sensitive perspective, cash managers may consider our methodology to decide whether improving the predictive accuracy at hand is financially worthy. These two findings, which we further dissect next, are meant to yield benefits for cash managers.

On the impact of predictive accuracy on cost savings. Gormley and Meade (2007) hypothesized that the more accurate the cash flow forecasting accuracy, the larger the cost savings expected. Here, for the first time in the literature, we have empirically confirmed such hypothesis. Furthermore, we have analyzed the impact of predictive accuracy on average daily cost savings when considering a variety of cost structures (of real-world bank finance con-

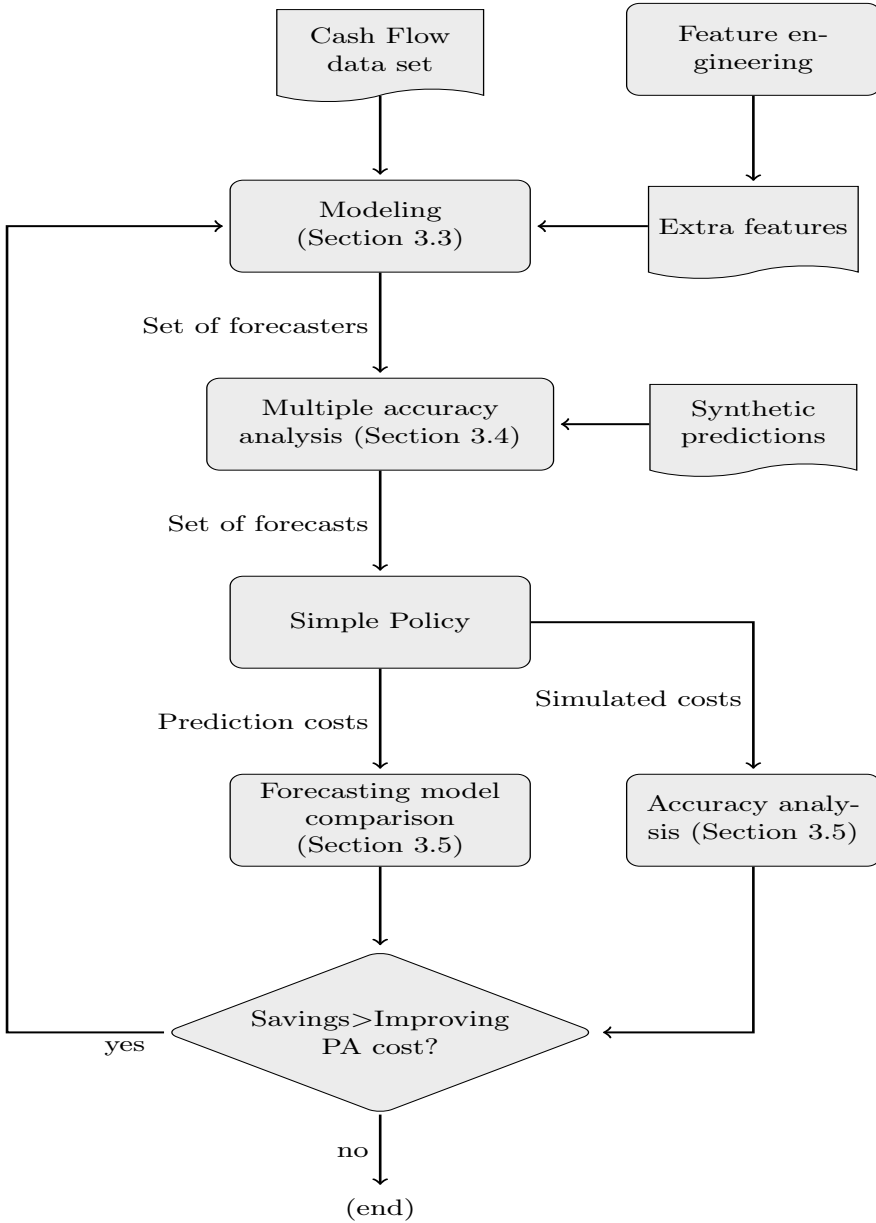


Figure 3.6: Potential cost savings analysis methodology (PA=Predictive Accuracy)

ditions) and cash flow policy parameters. From our analysis we have learned that:

- *Predictive accuracy is strongly correlated with cost savings* when using daily forecasts in cash management models. Thus, cost savings were highly sensitive to improvements on prediction accuracy when using a simple policy and two real-world cash flow data sets.
- *The riskier the cash management policy, the higher the average daily cost reduction in cash.*
- *The realistic cost structures considered in the most likely scenario have little influence on cost savings obtained by the forecasting models.*

What if predictive accuracy increases? Analyzing potential savings.

Cash managers may wonder if efforts on improving forecasting accuracy are expected to be proportionally rewarded by cost savings. Along this direction, we proposed a method for estimating the cost savings potentially delivered by improving predictive accuracy. Independently of the predictive accuracy of the forecaster available to a cash manager, our results help her estimate the cost savings that she might expect. Moreover, even if the cash manager does not count on any forecaster, she can estimate the cost savings that she currently misses. Overall, we learned that different risk levels yield different estimation results so that:

- *When assuming low risk, cost savings are limited* and further efforts in enhancing predictive accuracy are expected to be useless, in terms of both time and money, when a particular point in predictive accuracy is reached; and
- *The higher the risk a cash manager assumes, the higher the expected reward when improving predictive accuracy.*

The analysis of the relationship between predictive accuracy and cost savings leads to confirm the importance of better forecasting models when predictions are used as the main input to cash management models. Some additional intuition can be derived in the sense that this behavior may be caused by a number of reasons: (i) whenever it is possible to reduce uncertainty about the future, better decisions can be made; (ii) improving predictive accuracy is closely linked to discover patterns and an appropriate response to these patterns is necessarily useful; (iii) chances are that cash management models using forecasts as the main input do not work well with low quality forecasts. All of them highlight again the utility of forecasts in cash management.

Besides the above-mentioned benefits, it is important to note that there are additional benefits that companies can derive from improving forecasting accuracy such as short-term investment improvement and credit lines savings. Moreover, since our analysis can be extended to a multiple currency framework, larger multinational companies can also benefit the results presented here.

As an answer to the first research question addressed in this thesis, this chapter confirms the savings hypothesis for the first time in the cash management literature.

Question 1. *Can cash flow predictive accuracy achieve cost savings in the cash management problem?*

Contribution 1. *We empirically confirm the savings hypothesis showing that predictive accuracy is strongly correlated with cost savings.*

This chapter has demonstrated the utility of forecasts in the cash management problem. In the next chapter, we go one step further by deeply exploring the empirical properties of cash flows and its implications for forecasting.

Empirical analysis of daily cash flow

Usual assumptions on the statistical properties of daily cash flow include normality, absence of correlation and stationarity. In this chapter, we provide a comprehensive study based on a real-world cash flow data set from small and medium companies, which is the most common type of companies in Europe. We also propose a new cross-validated test for time-series non-linearity showing that: (i) the usual assumption of normality, absence of correlation and stationarity hardly appear; (ii) non-linearity is often relevant for forecasting; and (iii) typical data transformations have little impact on linearity and normality. Our results provide a forecasting strategy for cash flow management which performs better than classical methods. This evidence may lead to consider a more data-driven approach such as time-series forecasting in an attempt to provide cash managers with expert systems in cash management.

4.1 Motivation

In this chapter, we study the main statistical properties of real-world cash flow data. Zopounidis and Doumpos (2013) highlight the three different levels of analysis that facilitate the decision-making process of operational approaches for financial problems: (i) the focus on building models that describe the characteristics of financial problems; (ii) the development on empirical studies that

seek to test the theoretical models and identify new unobserved explanatory factors; and (iii) the construction of possible solutions and the selection of the most appropriate ways of action. Here, we mainly focus on the second of the previous levels although we also provide possible solutions derived from our findings.

Since Baumol (1952), a number of cash management models have been proposed to control cash balances. These models are based either on the specific statistical properties of cash balances or on cash flow forecasts. A comprehensive review of models, from the first proposals to the most recent contributions, can be found in Gregory (1976) and Srinivasan and Kim (1986), and Costa Moraes, Nagano, and Sobreiro (2015). Most of them are based on assuming a given probability distribution for cash flows such as: (i) a random walk in the form of independent Bernoulli trials as in Miller and Orr (1966); (ii) a Wiener process as in Constantinides and Richard (1978) and Premachandra (2004), and Baccarin (2009); (iii) a double exponential distribution as in Penttinen (1991). From these and other works, we observe that common assumptions on the statistical properties of cash flow time-series include:

- Normality: cash flows follow a Gaussian distribution with observations symmetrically centered around the mean, and with finite variance.
- Absence of correlation: the occurrence of past cash flows does not affect the probability of occurrence of the next ones.
- Stationarity: the probability distribution of cash flows does not change over time and, consequently, its statistical properties such as the mean and variance remain stable.
- Linearity: cash flows are proportional either to another (external) explanatory variable or to a combination of (external) explanatory variables.

Surprisingly, little and/or contradictory empirical evidence on these assumptions has been provided besides individual cases through time. Early on, negative normality tests were reported in Homonoff and Mullins (1975) for the times-series samples of a manufacturing company. Contrastingly, later on, Emery (1981) reported normally distributed cash flow, after data transformation, for two out of three companies, and a small serial dependence for all of them. Pindado and Vico (1996) provided negative normality and independence results on 36 companies, but considering daily cash flow for only a single month. Previous works also reported day-of-week and day-of-month effects on cash flows, in line with the works of Stone and Wood (1977), Stone and Miller

(1981), and Miller and Stone (1985), and Stone and Miller (1987). Recently, Gormley and Meade (2007) described the time-series from a multinational company with a non-normal distribution and serial dependence.

We consider that the evidence derived from these works is inconclusive due to: (i) the disagreement between the conclusions of some of the works; (ii) the limited number of companies analyzed; and (iii) the short time range of the observations. Moreover, none of the previous works considered the presence of non-linear patterns for forecasting purposes. In what follows, we provide an analysis of the statistical properties of 54 real cash flow data sets from small and medium companies in Spain as a representative sample of the most common type of companies in Europe. Indeed, small and medium companies contribute to 99.8% of all enterprises, 57.4% of value added, and 66.8% of employment across the EU28 (Muller et al., 2015). To the best of our knowledge, this is the most comprehensive empirical study on daily cash flow so far. We base this statement on the range of statistical properties considered, and on both the number and length of the data sets, which amount to 58005 observations in total, with a minimum, average and maximum time range of 170, 737, 1508 working days, respectively. A further contribution of the present work is to make all the aforementioned data publicly available online¹. Finally, from a forecasting perspective, we also aim to identify the family of forecasters that best accommodate to our cash flow time-series data sets. To this end, we propose a new and simple cross-validated test for non-linearity that provides further knowledge to cash managers in their search for better forecasting models.

Our results show the unlikely occurrence of normality, absence of correlation and stationarity. These results are consistent with the cited reports of Homonoff and Mullins (1975), based on only one time-series, and Pindado and Vico (1996), based on a very short time range, raising doubts about the claim of independence. However, we also report that normality could not be achieved through removing outliers, contrary to what was reported by Emery (1981), based on only three time-series. Thus, we consider that our results provide stronger evidence against normality, uncorrelatedness and stationarity than previous works. Note that we do not claim that these results can be extrapolated to all kind of companies. On the contrary, we provide further evidence against standard assumptions in cash management. This evidence may lead to consider a more data-driven approach such as time-series forecasting in order to provide cash managers with expert systems in cash management.

¹<http://www.iiia.csic.es/~jar/54datasets3.csv>

In an attempt to achieve Gaussian and stationary time-series, practitioners typically use the Box-Cox transformation (Box and Cox, 1964), and time-series differencing (Makridakis, Wheelwright, and Hyndman, 2008). Furthermore, some kind of outlier treatment is also a recommended practice. Then, we also study the impact of outlier treatment by replacing them with linear interpolations between two consecutive observations. However, in our study, we find little benefit when these methods are applied to our data sets. As a result, we point out the underlying question about data transformation in relation to the properties of a time-series. Is it always possible to achieve a Gaussian and linear time-series through data transformations? We here rely both on common statistical tests and on our novel non-linearity test to answer this question and we find that: (i) outlier treatment and Box-Cox transformation are not always enough to achieve normality; (ii) outlier treatment produces mixed results in terms of noise reduction and information loss; (iii) outlier treatment and Box-Cox transformations do not produce linearity. These results suggest that non-linear models conform a justifiable alternative for cash flow time-series forecasting, beyond the current conjectures of the literature.

4.2 Data summary

The data set contains daily cash flows from 54 different companies from the manufacturing and the service sector in Spain with annual revenue up to €10 million each. No company from the primary sectors is included in the sample. We select only small and medium companies since it is the most common size of companies in both Spain and Europe (Muller et al., 2015). This data set covers a date range of about 8 years and is available online. An instance in the data set contains the following fields or columns:

- Date: standardized YYYY-MM-DD dates from 2009-01-01 to 2016-28-08.
- Company: company identifier from 1 to 54.
- NetCF: daily net cash flow in thousands of €.
- DayMonth: categorical variable with the day of the month from 1 to 31.
- DayWeek: categorical variable with the day of the week from 1 (Monday) to 7 (Sunday).

Table 4.1 shows the statistical summary of daily net cash flow on non-holidays, grouped by company. Small and medium companies are likely to experiment daily null cash flows, meaning that no monetary movement is observed at a

particular working day even under regular activity. As a result, the occurrence of null cash flows is an important characteristic of small and medium companies due to the size of companies. Indeed, almost 30% of the companies in our data set present more than 25% of null cash flow observations. This fact implies a first and important characteristic for forecasting: a null cash flow prediction will be right at least 25% of the times for this group of data sets. Therefore, two good baseline forecasting models for comparative purposes would be an *always-predict-null* or an *always-predict-mean* forecaster.

In addition, the average net cash flow shows that a high percentage of companies present either positive or negative drift with the exception of companies 5 and 28. High positive kurtosis indicates a peaked data distribution in comparison to the normal distribution that has zero kurtosis. The skewness is a measure of the symmetry of the data distribution. Negative skewness indicates a bias to the left, and positive skewness indicates a bias to the right.

4.2.1 Normality

Next, we study if our cash flows follow a Gaussian distribution. The observed kurtosis and skewness can be used as a first normality test of the data distribution for each company. Table 4.1 shows that no company presents zero kurtosis and skewness. Only company 40, with kurtosis 0.58 and skewness 1.02, could be considered close to normality. Two additional tests can be used to either verify or reject the hypothesis of normality: the Shapiro-Wilk test for normality (Royston, 1982) and the Lilliefors (Kolmogorov-Smirnov) test for normality (Lilliefors, 1967). The results from these two tests allow us to reject the hypothesis of normally distributed cash flows for all the companies in our data set (no exception).

As pointed out elsewhere (Emery, 1981; Pindado and Vico, 1996), a possible explanation for non-normality could be the presence of abnormally high values or heavy tails. Thus, we repeated the Shapiro-Wilk test and the Lilliefors (Kolmogorov-Smirnov) test for normality, but using a trimmed version of the net cash flow time-series by deleting observations greater or lower than three times the sample standard deviation. No difference in the results of the tests is observed, confirming the non-normality hypothesis beyond the conjectures of Emery (1981) and Pindado and Vico (1996).

Non-normal residuals may be problematic in the estimation process when using linear models. Data transformations such as the Box and Cox (1964) transformation to normality represent a possible solution. Forecasts are then

Id	Length	Null %	Mean	Std	Kurtosis	Skewness	Min	Max
1	856	35,7	0,01	3,38	594,81	22,37	-9,07	90,27
2	684	29,8	0,26	5,80	58,98	3,69	-56,51	62,66
3	856	8,5	0,36	35,35	163,62	6,28	-303,20	671,04
4	1201	34,9	-0,12	14,32	78,14	-6,30	-223,38	72,76
5	849	19,4	0,00	1,67	56,10	-0,48	-18,26	16,42
6	799	20,7	0,01	6,63	33,21	-2,42	-68,97	56,27
7	772	38,5	0,07	5,36	86,75	6,74	-24,41	82,91
8	695	21,7	0,05	3,15	14,27	-2,57	-24,21	11,31
9	852	18,8	0,73	56,54	18,92	-0,78	-411,41	473,36
10	744	13,2	0,12	6,95	70,63	0,60	-81,13	78,72
11	639	62,6	-0,05	8,56	391,86	-17,65	-191,53	30,74
12	503	2,6	0,48	35,30	449,38	20,70	-47,27	771,38
13	697	24,7	0,52	24,24	18,81	2,06	-99,39	227,45
14	604	4,6	0,10	13,23	8,51	1,05	-63,23	92,71
15	605	4,1	0,68	11,67	4,43	0,33	-54,75	55,61
16	596	6,4	0,01	1,46	107,82	6,68	-8,48	22,61
17	1102	25,1	0,58	13,31	215,97	11,96	-118,01	250,13
18	552	3,1	0,16	2,16	70,23	5,10	-16,14	26,36
19	503	2,4	-0,31	2,58	6,43	0,50	-15,06	15,28
20	848	27,8	0,02	1,07	96,19	3,86	-12,07	16,04
21	829	18,7	-0,06	5,99	33,36	-1,62	-70,00	53,17
22	494	1,6	-0,46	27,28	22,64	-1,96	-244,29	138,87
23	604	9,1	1,63	20,85	79,99	5,41	-124,19	269,27
24	1097	8,4	0,96	20,36	95,45	6,48	-73,33	317,85
25	587	10,9	0,49	13,94	119,60	6,93	-116,01	201,13
26	751	11,6	-0,02	1,77	15,73	0,15	-10,73	15,56
27	332	8,1	0,29	1,64	10,60	2,14	-4,36	11,84
28	855	5,1	0,00	4,64	13,83	1,77	-18,10	39,01
29	609	13,6	0,04	6,07	108,66	-6,35	-90,04	55,89
30	554	8,1	0,03	1,47	68,26	5,47	-4,81	19,82
31	372	29,6	0,37	8,05	31,46	-2,41	-80,44	34,95
32	1103	24,8	0,28	4,03	11,07	0,54	-25,76	24,50
33	854	31,0	-0,19	6,81	115,63	-1,74	-94,33	95,59
34	1508	11,5	-0,06	10,13	19,89	-2,32	-96,82	49,65
35	501	7,4	0,20	5,40	11,41	-0,58	-31,42	29,19
36	359	11,4	0,42	1,85	12,24	2,44	-7,87	11,84
37	361	3,0	-0,69	17,82	139,06	-1,38	-228,88	218,42
38	170	9,4	-1,20	7,10	43,34	-5,73	-61,93	19,66
39	1104	29,0	0,02	0,95	7,95	-0,07	-5,67	6,57
40	198	0,0	0,78	12,38	0,58	1,02	-25,63	36,91
41	341	17,6	-0,25	8,34	15,80	1,22	-44,29	64,34
42	566	11,0	0,01	1,82	308,62	-15,80	-37,02	7,48
43	750	3,2	0,34	13,10	7,66	-0,04	-65,84	73,40
44	287	4,2	0,52	11,46	81,19	-0,05	-118,74	120,34
45	1465	49,8	0,04	9,12	43,51	-2,89	-107,20	75,47
46	565	44,8	0,54	5,58	75,41	2,91	-51,16	73,83
47	503	4,4	1,98	46,81	46,03	1,37	-338,39	478,26
48	605	13,1	0,21	22,71	34,31	-1,68	-207,04	203,09
49	993	50,5	-0,08	1,36	27,18	-2,18	-10,78	12,73
50	605	45,0	-0,01	27,37	43,79	-2,01	-262,52	221,96
51	1225	0,2	15,09	96,96	2,77	0,12	-419,88	481,66
52	1225	0,4	8,94	49,39	36,23	2,81	-325,46	700,66
53	1223	39,7	0,47	9,13	203,12	-10,25	-196,88	38,48
54	1225	52,3	0,46	77,91	151,93	4,28	-1021,36	1532,10

Table 4.1: Data sets statistical summary. Mean, standard deviation, minimum, maximum in thousands of €.

calculated on the transformed data, but we must reverse the transformation to obtain forecasts on the original data, resulting in two additional steps. However, these transformations are not always the solution to the non-normality problem. Using both the original observations and the trimmed version of our data sets, we proceeded to transform the data using a Box-Cox transformation of the type:

$$y^{(\lambda)} = \begin{cases} \frac{(y + \lambda_2)^{\lambda_1} - 1}{\lambda_1} & \text{if } \lambda_1 \neq 0, \\ \log(y + \lambda_2) & \text{if } \lambda_1 = 0, \end{cases} \quad (4.1)$$

where y is the original time-series, and λ_1 and λ_2 are parameters. In these experiments, we first set λ_2 to minus two times the minimum value of the time-series to avoid problems with negative and zero observations. Box and Cox (1964) provided the profile likelihood function for λ_1 and suggested to use this function as a way to tune this parameter. Then, we follow the recommendations in Venables and Ripley (2013) to compute the profile likelihood function for λ_1 , and we later select the value that maximizes the log-likelihood function when applying a linear regression model of the time-series based on day-of-month and day-of-week dummy variables. Then, we repeated the Lilliefors (Kolmogorov-Smirnov) test for normality obtaining again negative results for normality, both for the original and for the trimmed version of our data sets. As a result, we must conclude that, even after Box-Cox transformation, the normality hypothesis does not hold.

4.2.2 Correlation and seasonality

In what follows, we test the correlation of cash flows and we also explore if seasonality is present. Autoregressive Integrated Moving Average (ARIMA) models by Box and Jenkins (1976), have been extensively used for time-series analysis and forecasting. When dealing with time-series, the autocorrelation coefficient, r_k , describes the relationship between observations that are lagged k time periods (Makridakis, Wheelwright, and Hyndman, 2008). We say that a time-series is independent when the r_k values for different lags are close to zero. An example of an independent time-series is the so-called white-noise model where each observation is made by adding a random component to a certain level.

An intuitive plot to assess correlation is the Poincaré map (Kantz and Schreiber, 2004), which is a scatter plot of the original time-series and a k -periods lagged time-series as in Figure 4.1, which shows a lag of 1 day for time-series 1 and 2 from Table 4.1. As a reference, we also include the Poincaré map for a Gaussian noise and for a sinusoidal time-series. A cloud of points suggests absence

of correlation, as for time-series 1 and Gaussian noise, and the presence of any form suggests a more complex relationship, as for time-series 2 and the sinusoidal.

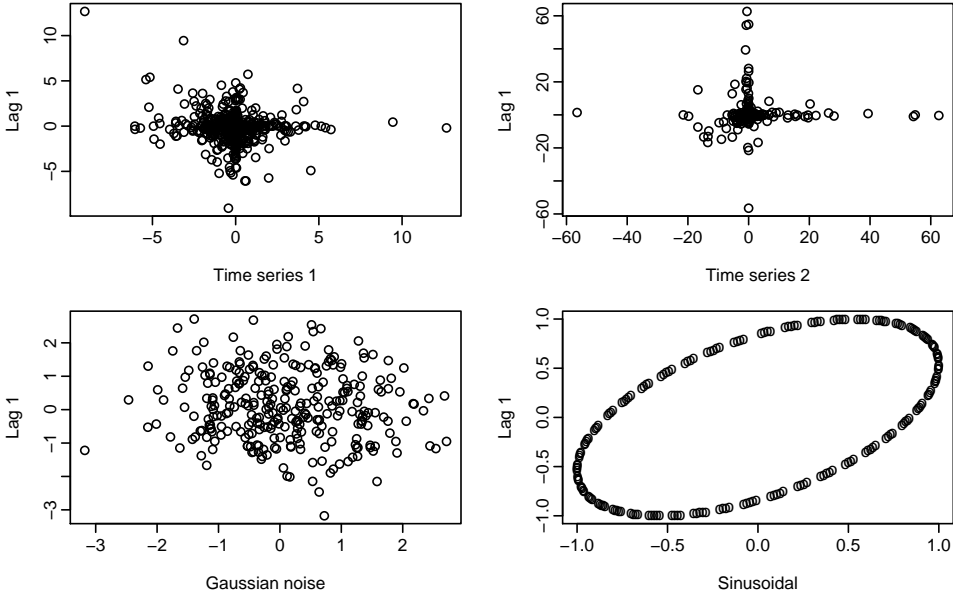


Figure 4.1: Poincaré map with lag 1 for time-series 1 and 2.

A more general approach is to consider a set of the first r_k values as a whole as in the Ljung and Box (1978) test, which we applied to our data and produced mixed results. More precisely, we found that the hypothesis of independence could not be rejected in 24 out of 54 companies as summarized in Table 4.2. These results imply that some kind of serial correlation is likely to be present in the case of companies presenting a certain degree of autocorrelation in the sample. A plausible type of serial correlation is seasonality, that is, the existence of a pattern that repeats itself over fixed time intervals in the data (Makridakis, Wheelwright, and Hyndman, 2008). It can be identified by significant autocorrelation coefficients. Seasonal trend decomposition methods (Cleveland et al., 1990), seasonal ARIMA models (Box and Jenkins, 1976; Franses and Van Dijk, 2005) or linear (and non-linear) regression models based on seasonal variables are available options to deal with seasonality. In cash flow forecasting, the distribution approach by Miller and Stone (1985) also deserves to be mentioned.

As mentioned in the introduction, previous works by Emery (1981), Miller and Stone (1985), and Stone and Miller (1987); and Pindado and Vico (1996), reported the influence of day-of-month and day-of-week effects on cash flow patterns. Here, we test the presence of seasonality by fitting a regression model on daily cash flows using day-of-month and day-of-week dummy variables. Then, we use the F-test to test the significance of the overall regression. Table 4.2 reports, on the one hand, the Ljung-Box independence test and, on the other hand, the F-statistic, the p-value and the coefficient of determination R^2 , derived from the regression. Interestingly, not all the data sets whose independence (absence of correlation) tests were rejected implied a significant regression based on dummy variables. Non-linear patterns, non-periodical temporal correlations, and the effect of outliers become possible explanations as we will see below.

4.2.3 Stationarity

In this section, we analyze if cash flows from our data set can be labeled as stationary. Basically, stationarity means that there is no drift in the time-series behavior over time. We can visually assess stationarity by inspecting a time-series plot. Virtually, every process we find in nature is non-stationary, since its parameters depend on time (Kantz and Schreiber, 2004). However, a minimum requirement is that basic statistical properties of a distribution, such as mean and variance, remain constant over time, when measured through appropriately long time windows.

Emery (1981) studied stationarity by applying the Kolmogorov-Smirnov test for normality of cash flow by months. For comparative purposes, we applied the same procedure and, if any of the monthly tests rejected the hypothesis of normality, the whole time-series was considered non-stationary. Following this procedure, only company 43 could be considered stationary.

Following the recommendations in Kantz and Schreiber (2004), we also perform a stationarity test based on the fluctuations of a sample mean and variance. More precisely, we compute the sample mean and variance of each time-series by months and obtain the standard errors for both. If the observed fluctuations of the running mean and variance are within these errors, then we consider the time-series stationary. The results from this test shows that none of the time-series is stationary.

One way of removing non-stationarity is time-series differencing, which can be defined as the change between two consecutive observations. After differencing,

Id	Ljung Box test	F-statistic	p-value	R^2
1	Non-rejected	1,99	<0,05	0,08
2	Rejected	1,05	0,39	0,05
3	Non-rejected	1,87	<0,05	0,07
4	Rejected	1,51	<0,05	0,04
5	Rejected	1,85	<0,05	0,07
6	Non-rejected	1,12	0,29	0,05
7	Rejected	5,47	<0,05	0,20
8	Rejected	0,79	0,80	0,04
9	Rejected	5,30	<0,05	0,18
10	Rejected	2,04	<0,05	0,09
11	Non-rejected	0,97	0,51	0,05
12	Non-rejected	0,98	0,51	0,07
13	Rejected	5,21	<0,05	0,21
14	Rejected	7,13	<0,05	0,30
15	Rejected	1,92	<0,05	0,10
16	Non-rejected	4,31	<0,05	0,21
17	Non-rejected	4,91	<0,05	0,14
18	Rejected	2,99	<0,05	0,16
19	Rejected	2,58	<0,05	0,16
20	Non-rejected	2,71	<0,05	0,10
21	Non-rejected	1,37	0,08	0,06
22	Non-rejected	1,49	<0,05	0,10
23	Rejected	5,60	<0,05	0,25
24	Non-rejected	15,41	<0,05	0,33
25	Non-rejected	4,23	<0,05	0,21
26	Rejected	1,22	0,18	0,05
27	Non-rejected	1,24	0,18	0,12
28	Rejected	5,64	<0,05	0,19
29	Non-rejected	1,37	0,08	0,08
30	Rejected	6,18	<0,05	0,29
31	Non-rejected	1,25	0,16	0,11
32	Rejected	4,81	<0,05	0,13
33	Rejected	1,57	<0,05	0,06
34	Rejected	11,61	<0,05	0,21
35	Rejected	0,99	0,49	0,07
36	Non-rejected	1,82	<0,05	0,16
37	Rejected	1,58	<0,05	0,14
38	Non-rejected	1,06	0,39	0,21
39	Rejected	6,11	<0,05	0,16
40	Rejected	0,86	0,68	0,15
41	Non-rejected	1,72	<0,05	0,16
42	Non-rejected	3,90	<0,05	0,20
43	Rejected	2,96	<0,05	0,12
44	Non-rejected	1,89	<0,05	0,20
45	Rejected	1,26	0,15	0,03
46	Non-rejected	1,32	0,11	0,08
47	Non-rejected	0,90	0,63	0,06
48	Non-rejected	1,71	<0,05	0,09
49	Rejected	26,15	<0,05	0,48
50	Rejected	1,24	0,17	0,07
51	Rejected	16,66	<0,05	0,32
52	Rejected	5,01	<0,05	0,13
53	Non-rejected	1,59	<0,05	0,04
54	Rejected	0,88	0,67	0,02

Table 4.2: Correlation and seasonality test results.

we repeated our simple test obtaining slightly different results but none of them can be considered stationary. Thus, as a result, we conclude that our cash flow time-series are non-stationary, even after differencing.

4.2.4 Discussion

Our results show that the widely extended hypothesis of cash flow normality is not present in our data sets. The presence of high abnormal values does not explain this behavior since non-normality persisted after removing these abnormal values. Non-linearity could be a possible explanation as we will see below. We also reported mixed results on autocorrelation and the influence of day-of-month and day-of-week effects on cash flow along the lines of the literature. We additionally report that common solutions to non-normality and non-stationarity such as data transformation and differencing produced little benefit when applied to our time-series. Since seasonality and serial correlation are also present in our data set, we further explore the usefulness of alternative forecasting models. More precisely, we next study linearity and data transformation as an additional part of our empirical analysis for cash flow forecasting.

4.3 A cross-validated test for non-linearity

Most forecasting models are linear for computational convenience. However, non-linear patterns are likely to be present in finance and business time-series. A time-series linear model is defined as a variable y_t that depends on an explanatory vector \mathbf{x}_t for any time t as follows:

$$y_t = \boldsymbol{\beta}^T \mathbf{x}_t + e_t \quad (4.2)$$

where $\boldsymbol{\beta}^T$ is a transposed vector of coefficients, and e_t is the error or the residual component. An alternative and more general model can also be considered:

$$y_t = g(\mathbf{x}_t) + \epsilon_t \quad (4.3)$$

where $g(\mathbf{x}_t)$ is any function that aims to describe the underlying time-series. By considering non-linear relationships between the set of predictors and the cash flow dependent variable, more complex patterns such as interactions between the day-of-week and the day-of-month may be captured.

Different tests of linearity can be found in Ramsey (1969), Keenan (1985), Granger, Terasvirta, et al. (1993), and Lee, White, and Granger (1993), and Castle and Hendry (2010). Basically, all of them follow a common approach: first, they choose a regression equation $g(\mathbf{x}_t)$ in (4.3) including linear and non-linear terms and, second, they test for the significance of the non-linear terms. However, these approaches do not accommodate well for forecasting purposes due to the following reasons: (i) the assumption of a specific form

$g(\mathbf{x}_t)$ for the regression equation such as quadratic, cubic or exponential forms; (ii) cross-validation is neglected.

If we relax the assumption of linearity, different non-linear models such as random forests (Breiman, 2001), neural networks (Hornik, Stinchcombe, and White, 1989; Zhang, Patuwo, and Hu, 1998), or radial basis functions (Broomhead and Lowe, 1988), could also be considered. However, the consideration of non-linear functions may lead to overfitting to the original time-series. To prevent this problem, we propose the use of time-series cross-validation. Cross-validation is a method to assess the predictive performance of a forecasting model that circumvents the problem of overfitting the data by testing the accuracy of the model on subset of data not used in the estimation (Hyndman and Athanasopoulos, 2013). As a result, we here propose a simple cross-validated test for non-linearity as described in Figure 4.2 based on the following steps:

1. Estimate two alternative forecasting models, one linear and another one non-linear.
2. Cross-validate the predictive accuracy of both models with respect to a baseline.
3. Label as trivial if both models are significantly worse than the baseline.
4. Label as non-linear if the error of the non-linear model is significantly lower than that of the linear model. Otherwise, label as linear.

Since we do not assume any distribution for the forecasting results, we use the two-sided Wilcoxon rank-based for statistically significant differences in performance between models. More precisely, we test the null hypothesis that the distribution of the difference is symmetric about zero with a 95% confidence interval (Wilcoxon, Katti, and Wilcox, 1970). Approximate p -values are computed based on the asymptotic distribution of the two-sided Wilcoxon test statistic and used to label data sets.

A common practice to assess the usefulness of forecasts derived from any model is to compare its accuracy to that of a baseline forecasting model. The use of a baseline model allows us to label our data sets as trivial if neither the linear model nor the non-linear model are able to improve the accuracy of the baseline. We here report accuracy results with respect to a mean forecaster, meaning that forecasts are always the average of all past observations. We also tried with an additional baseline forecaster using the last observed value

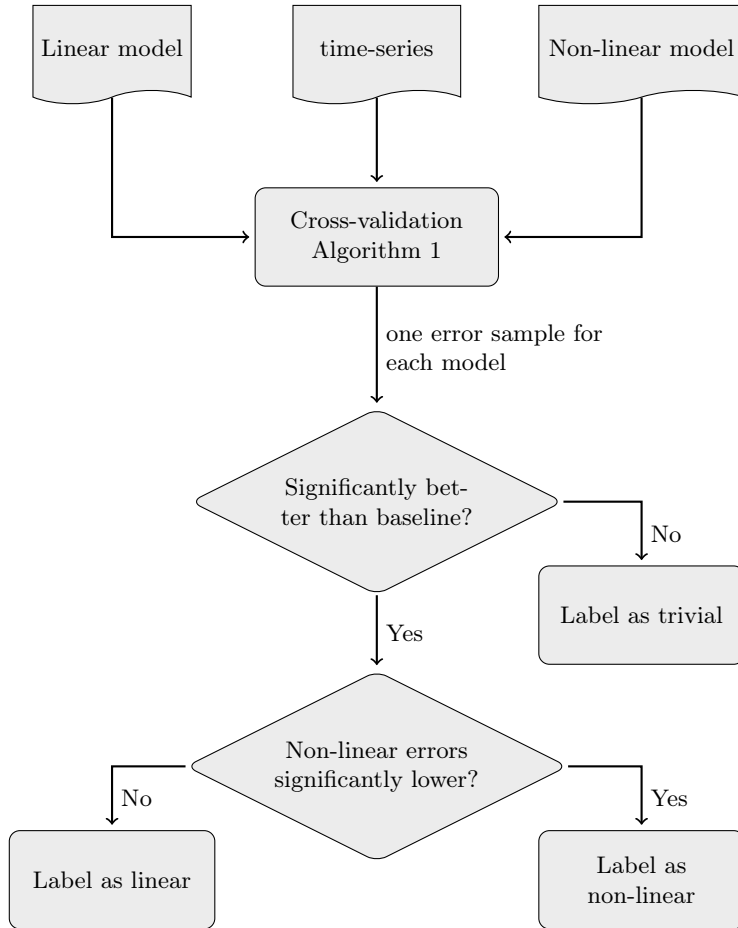


Figure 4.2: Simplified flow chart for our cross-validated test for non-linearity.

as a forecast (persistence model) with much worse results in comparison to the mean forecaster.

As detailed in Algorithm 3, we consider the minimum length k to estimate a model as the 80% of the oldest instances forming the training set. The remaining 20% of the instances form the test set for cross-validation. Initially, both the linear and the non-linear model are estimated using the first 80% of the instances. Then, forecasts for a prediction horizon up to 20 days are computed using the estimated models and squared errors are recorded. Later, we repeat the process by considering an increasing window with an additional

instance in the training set to estimate the models. Again, forecasting errors are recorded for each remaining observation in the test set resulting into two paired error samples, one for the linear model and one for non-linear model.

Algorithm 3: Algorithm for a simple cross-validation test for non-linearity

```
1 Input: Cash flow data set of  $T$  instances, minimum number  $k$  of instances to
   estimate a model, baseline  $m_0$ , linear model  $m_1$ , non-linear model  $m_2$ ,
   prediction horizon  $h$ , level of significance  $\alpha$ ;
2 Output: Average prediction error for different prediction horizons, statistic
   for the difference in means, confidence interval;
3 for  $i = 1, 2, \dots, T - k - h + 1$  do
4   | Select the instances from time  $k + i$  to  $k + h + i - 1$ , for the test set;
5   | Estimate  $m_0$  with instances at times  $1, 2, \dots, k + i - 1$ ;
6   | Estimate  $m_1$  with instances at times  $1, 2, \dots, k + i - 1$ ;
7   | Estimate  $m_2$  with instances at times  $1, 2, \dots, k + i - 1$ ;
8   | Compute test errors  $\varepsilon_0, \varepsilon_1, \varepsilon_2$  from time  $k + i$  to  $k + h + i - 1$ ;
9 end
10 Compute average  $h$ -step errors  $\varepsilon_0(h), \varepsilon_1(h), \varepsilon_2(h)$ ;
11 Test for  $\alpha$  significant differences between  $\varepsilon_0(h), \varepsilon_1(h), \varepsilon_2(h)$ ;
12 if  $\varepsilon_0(h) < \varepsilon_1(h)$  and  $\varepsilon_0(h) < \varepsilon_2(h)$  then
13   | Label as trivial;
14   else if  $\varepsilon_2(h) < \varepsilon_1(h)$  then
15     | Label as non-linear;
16     else
17       | Label as linear;
18     end
19   end
20 end
```

A critical point when using our cross-validated test for non-linearity is the selection of both the linear and the non-linear forecast model. In essence, our test is a comparative tool based on forecasting accuracy as a proxy for non-linearity. Given a set of explanatory variables, a linear label result from our test implies that the non-linear model is not able to capture non-linearity. However, chances are that alternative non-linear models might perform differently. In this sense, if the time-series is not a white-noise process, then the search for a more informative set of features is meant to play a key role. As a result, multiple runs of our test are necessary to discard/assess non-linearity by using alternative linear and non-linear models.

For illustrative purposes, we here restrict ourselves to a linear regression model and a non-linear random forest model, both using day-of-month and day-of-week variables as predictors. In the case of the linear regression model, each

instance contains 34 dummy predictor variables, 30 for day-of-month and 4 for day-of-week, and a cash flow observation. In the case of random forests, each instance contains 2 categorical variables, one for day-of-month and one for day-of-week. Random forests are ensembles of slightly different decision trees (Ho, 1998; Breiman, 2001). An ensemble methodology is able to construct a predictive model by integrating multiple trees in what is called a decision forest (Dietterich, 2000). Decision trees split the input space in subsets based on the value of features such as the day-of-month and day-of-week. In the example in Figure 4.3, for days comprised between the 25th (node S1) and the 29th of each month (node S2) occurring on Friday (node S3), the predicted cash flow is -1.




Figure 4.3: A basic decision tree. DOM = Day-of-month; DOW = Day-of-week

Recent examples of time-series forecasting using random forests can be found in Zagorecki (2015) and Salas-Molina et al. (2017). Summarizing, random forests are used to forecast variables based on an ensemble of different trees. Unlike linear regression, random forests allow to capture (if any) more complex relationships between predictor variables allowing us to identify possible non-linearities in the underlying cash flow process represented by our sample data sets.

Our results, summarized in Table 4.3, show that only about half of the data sets can be labeled as trivial because neither the linear model nor the non-linear model were able to significantly beat the trivial forecaster. From those

time-series in which the absence of correlation could no be rejected (see Table 4.2), 20 out of 24 were labeled as trivial, confirming a stochastic behavior. On the other hand, only 6 of them were labeled as non-linear according to our cross-validated definition. As mentioned above, these results depend on the selected forecasting models. Instead of claiming that random forests are able to better capture non-linear patterns than alternative models, we encourage practitioners to consider additional combinations of both linear and non-linear models.

Id	Reg NSE	RF NSE	Statistic	p-value	Triviality	Linearity
1	0,99	1,00	26	<0,05	Non-Trivial	Linear
3	0,99	1,01	8	<0,05	Non-Trivial	Linear
4	1,00	1,01	0	<0,05	Non-Trivial	Linear
7	0,81	0,83	0	<0,05	Non-Trivial	Linear
9	0,90	0,93	3	<0,05	Non-Trivial	Linear
13	0,86	0,88	13	<0,05	Non-Trivial	Linear
14	0,76	0,77	45	<0,05	Non-Trivial	Linear
16	0,85	0,86	64	0,13	Non-Trivial	Linear
18	0,86	0,88	63	0,12	Non-Trivial	Linear
19	0,96	0,94	182	<0,05	Non-Trivial	Non-linear
20	0,99	0,98	209	<0,05	Non-Trivial	Non-linear
23	0,78	0,79	78	0,33	Non-Trivial	Linear
24	0,73	0,79	0	<0,05	Non-Trivial	Linear
25	0,77	0,81	21	<0,05	Non-Trivial	Linear
28	0,84	0,90	0	<0,05	Non-Trivial	Linear
29	0,99	0,99	30	<0,05	Non-Trivial	Linear
30	0,73	0,80	5	<0,05	Non-Trivial	Linear
33	0,94	0,93	166	<0,05	Non-Trivial	Non-linear
34	0,97	0,95	172	<0,05	Non-Trivial	Non-linear
39	0,96	0,96	36	<0,05	Non-Trivial	Linear
42	0,88	0,87	149	0,11	Non-Trivial	Linear
43	0,99	0,96	210	<0,05	Non-Trivial	Non-linear
48	1,01	0,99	191	<0,05	Non-Trivial	Non-linear
49	0,63	0,65	7	<0,05	Non-Trivial	Linear
51	0,77	0,80	0	<0,05	Non-Trivial	Linear
52	0,94	0,94	116	0,70	Non-Trivial	Linear

Table 4.3: Results of the test for non-linearity. Reg NSE = Regression normalized squared error; RF NSE = Random Forest normalized squared error.

One may assume either linearity or non-linearity from the results of our non-linearity test, but it is important to analyze the robustness of these results to both the presence of outliers and the impact of other data transformations.

4.4 The impact of data transformations

In this section, we aim to analyze the impact of outlier treatments on noise reduction, as intended, and on information loss, as an undesirable effect. We also study the influence of Box-Cox data transformations on the results of our cross-validated non-linearity test. Detection and treatment of outliers is an ongoing issue in data mining (Rousseeuw and Leroy, 1987; Hodge and Austin, 2004). An outlier is an observation that appears to significantly deviate from other members of the sample in which it occurs (Grubbs, 1969). Outliers arise due to changes in systems, measurement errors or simply due to deviations in data. It is also important to note that an outlier may also be the most interesting part of the data.

On the one hand, from the set of cash flow time-series labeled as trivial, some of them may be labeled as non-trivial after removing outliers as a way of noise reduction. On the other hand, from those data sets labeled as non-trivial, some of them may be labeled as trivial due to the information loss produced by the treatment. We here measure the effect of removing outliers on the prediction error using time-series cross-validation for different thresholds of outlier replacement. For each data set, we progressively identify as outliers cash flow observations greater than 5, 4, and 3 times the standard deviation in a training set with the 80% oldest observations. We replace outliers with a linear interpolation and proceed as detailed in Algorithm 3 to cross-validate triviality and linearity. The results from this analysis are summarized in Table 4.4.

By following this procedure, we identify data sets 5, 10, 17, 32, 44 and 54 (6 out of 28), initially labeled as trivial that, after outlier treatment, can be labeled as non-trivial due to noise reduction. Similarly, data sets 4 and 48 that were initially labeled as non-trivial can be labeled as trivial after outlier treatment due to information loss. If we measure noise reduction by the error reduction and information loss by the error increase, then we can assess the impact of outlier treatment. Following this approach, we obtained mixed results for non-trivial data sets after outlier treatment: an average noise reduction of 22%, and an average information loss of 14%. It is important to recall that unexpected observations are often the most interesting part of the data to predict, e.g., when the goal is to forecast unusual but genuine cash flows.

Non-linearity and outliers are closely linked. Indeed, Castle and Hendry (2012) hypothesized that non-linear functions can align with outliers, causing functions to be considered relevant spuriously, which can be detrimental for generalizing and forecasting. If this hypothesis is correct, the relative forecasting ability of a linear model in comparison to a non-linear model would increase

Id	Triviality	After outliers		After outliers and Box-Cox	
		Linearity	Noise reduction	Linearity	Noise reduction
1	Non-Trivial	Linear	0,00	Non-linear*	-0,01
3	Non-Trivial	Linear	0,02	Non-linear*	0,00
5	Non-Trivial	Non-linear	0,40	Non-linear	0,41
7	Non-Trivial	Linear	-0,10	Linear	-0,13
9	Non-Trivial	Linear	-0,04	Linear	-0,04
10	Non-Trivial	Non-linear	0,46	Non-linear	0,47
13	Non-Trivial	Linear	-0,18	Linear	-0,21
14	Non-Trivial	Linear	-0,05	Linear	-0,07
16	Non-Trivial	Linear	-0,18	Linear	-0,17
17	Non-Trivial	Non-linear*	0,71	Non-linear	0,71
18	Non-Trivial	Non-linear*	-0,20	Non-linear	-0,20
19	Non-Trivial	Non-linear	-0,03	Non-linear	-0,04
20	Non-Trivial	Non-linear	-0,02	Non-linear	-0,02
23	Non-Trivial	Non-linear*	-0,22	Non-linear	-0,22
24	Non-Trivial	Linear	-0,20	Linear	-0,06
25	Non-Trivial	Non-linear*	-0,26	Non-linear	-0,25
28	Non-Trivial	Linear	-0,05	Linear	-0,04
29	Non-Trivial	Linear	0,07	Non-linear*	0,00
30	Non-Trivial	Linear	-0,06	Linear	-0,04
32	Non-Trivial	Non-linear	0,18	Non-linear	0,21
33	Non-Trivial	Linear*	-0,12	Linear	-0,11
34	Non-Trivial	Linear*	0,12	Linear	0,09
39	Non-Trivial	Non-linear*	-0,02	Linear*	-0,01
42	Non-Trivial	Linear	-0,23	Linear	-0,14
43	Non-Trivial	Non-linear	0,04	Non-linear	0,03
44	Non-Trivial	Non-linear*	0,48	Non-linear	0,82
49	Non-Trivial	Non-linear*	-0,56	Non-linear	-0,61
51	Non-Trivial	Linear	-0,03	Linear	-0,03
52	Non-Trivial	Linear	0,01	Linear	0,03
54	Non-Trivial	Linear*	0,17	Linear	0,17

Table 4.4: Results of the test for non-linearity after outlier treatment and Box-Cox transformation. Changes in labels are marked with *.

as the presence of outliers in a training set is reduced. From the set of time-series finally labeled as non-trivial, data sets 33, 34 and 54, initially labeled as non-linear changed their labels to linear. Surprisingly, data sets 17, 18, 23, 25, 39, 44 and 49 (7 out of 30), could be labeled as non-linear after outlier treatment. Except for data sets 17 and 44, in all cases there was information loss, i.e., error increase, suggesting that non-linear models can deal better with information loss.

We also considered a Box-Cox transformation to analyze if this kind of data transformation may influence the results from our cross-validated non-linearity test. From the set of non-trivial data sets we compare linearity labels, first, after outlier treatment, and second, after outlier treatment and Box-Cox transformation as described in equation (4.1). In addition, we compare information loss computed as the difference between the sum of errors of the linear and non-linear forecasting models before and after the outlier treatment. A positive value means noise reduction or error reduction while a negative value

means information loss or error increase. These results show that our cross-validated non-linearity test outputs similar results after Box-Cox transformation since the change in labels were produced in data sets with similar linear and non-linear model performance.

A summary of the results of this section is shown in Table 4.5. The high number of trivial data sets may be caused by the general inherent randomness of cash flows. Outlier treatment produced a small improvement in non-triviality but also an outstanding increase in non-linearity. Finally, Box-Cox data transformation yielded similar results but with better results for non-linear models. Thus, we conclude that: (i) common data transformations had little impact on our time-series in terms of linearity; and (ii) outlier treatment and Box-Cox transformation were unable to transform non-linear into linear cash flows.

Label	Raw data	After OT	After OT and DT
Trivial	28	24	24
Non-trivial	26	30	30
-Linear	20	17	15
-Non-linear	6	13	15

Table 4.5: Number of time-series data sets and their labels after transformation. OT=Outlier treatment; DT=Data transformation.

4.5 Summary

Small and medium companies are the most common type of companies in Europe contributing to a high percentage of all enterprises, value added and employment. In this chapter, we provide a complete empirical study of the statistical properties of daily cash flows based on 54 real-world cash flow time-series from small and medium companies. To the best of our knowledge, this study is the most comprehensive empirical study on daily cash flow so far in terms of the range of statistical properties considered, and on both the number and the length of the data sets that we make available online. We focus on the implications for forecasting due to its key role in cash management.

Our results show that the extended hypothesis of normal, stationary and uncorrelated cash flows is hardly present in our cash flow data set. We do not claim that these results can be generalized to all small and medium companies. However, we conclude that the standard assumptions of normality, stationarity and uncorrelatedness that have been extensively used in cash management literature must be verified before the deployment of any cash management model

based on them. We also highlight that common solutions to non-normality and non-stationarity such as data transformation and differencing produce little benefit when applied to our data sets, with the risk of losing important information on extreme cash flows. Alternative and more complex data transformations are nevertheless an option to consider in further research to achieve Gaussian cash flows.

In an attempt to discover the attributes of actual-world cash flows, we also studied the presence of non-linearity. To this end, we proposed a new simple test for non-linearity with two main advantages in comparison to alternative approaches. First, our test does not assume any non-linear function. Second, it is based on time-series cross-validation to increase robustness and avoid overfitting. It is important to note that our cross-validated definition of non-linearity depends on the alternative models considered, one linear and another one non-linear.

Our cross-validated non-linearity test labeled as either trivial, linear or non-linear our cash flow data set after outlier treatment resulting in an important increase the number of data sets labeled as non-linear. After both outlier treatment and Box-Cox transformation, linearity could not be achieved and non-linear models showed more robust.

Our results raise questions about two common assumptions in cash flow time-series since we found that: (i) the usual assumption of normality, absence of correlation and stationarity is hardly present; and (ii) common data transformations such as outlier treatment and Box-Cox transformation have little impact on normality and linearity. Contrary to the rather common assumption in the literature, these results imply that neither it is always possible to achieve a Gaussian, noise-free and linear time-series through data transformation nor it is always desirable due to information loss. Thus, linear models should be considered as an initial step towards more realistic ones which are better adapted to real cash flow situations. The results from our cross-validated test for non-linearity suggest that non-linear models represent a justifiable alternative for time-series forecasting. Moreover, since our test is both model and outlier dependent, a promising line of future work would be the integration of outlier treatment in the test itself in an attempt to assess noise reduction or information loss.

As a result, we claim that a number of preliminary steps are necessary in cash flow forecasting before model selection: (i) statistical summary including normality, correlation and stationarity; (ii) impact of data transformations such as outlier treatment and Box-Cox transformation; and (iii) non-linearity

test to determine the type of model which is expected to deliver a better performance. Finally, this process is not limited to daily cash flow, since it can also be applied to any other time-series data set when cross-validation is required.

To end up, recall from Chapter 1 that research questions 2 and 3 deal with the statistical properties of daily cash flows. The lack of empirical evidence about these properties motivated this chapter by testing the common assumptions of normality, absence of correlation, stationarity and linearity of cash flows. As a result, we next highlight two main contributions of this chapter.

Question 2. *Are common statistical assumptions of daily cash flow supported by recent empirical data?*

Contribution 2. *We demonstrated that normality, absence of correlation and stationarity hardly appear in real-world cash flow data sets. Non-linearity is often relevant and it can be assessed through a new cross-validated test described in this thesis.*

Furthermore, since data transformation is usually considered as a necessary step by time-series analysis techniques to achieve normality and linearity, we also addressed the question of data transformations.

Question 3. *Is it always possible to achieve a Gaussian, noise-free and linear time-series through data transformations?*

Contribution 3. *We showed that data transformations cannot always achieve a Gaussian, noise-free and linear cash-flow time-series suggesting the utility of non-linear forecasting models.*

Chapters 3 and 4 dealt with cash flow forecasting techniques and also with the statistical properties of cash flows and its implications for forecasting. The main focus was the utility of forecasts in cash management. In the next part of this thesis, we use forecasts as a key input to cash management, but we shift the focus to optimization procedures.

Part II

Multiobjective cash management

A multiobjective approach to the cash management problem

The main goal of this chapter is to present a multiobjective approach to the cash management problem. We take the view that cash management decision support systems may benefit by considering not only cost but also risk. To this end, we provide the necessary tools to solve the multiobjective cash management problem. We also provide a framework to assess the utility of cash management models when considering multiple objectives.

5.1 Motivation

It is easy to understand that any real decision-making problem has to take into account multiple objectives. This fact is particularly true in the fields of business and finance in which different, but often conflicting, goals are pursued by decision-makers. Some examples are return and risk in portfolio selection, price and amount in sales planning, and profits and pollution in chemical production. It is a common misconception that most problem solving activities must align with optimizing a single objective, e.g., bringing maximum profit or causing the smallest cost (Branke et al., 2008). As a result, different goals are often translated into economic terms reducing the number of conflicting goals into a single objective. However, different objectives are often measured

using different units. We say then that the objectives are non-commensurable because it is not possible to directly aggregate them.

In addition, the utility obtained from different objectives are very often dependent on the alternatives. Consider the following example of choosing a hotel described in Branke et al. (2008). If the alternatives are a one-star hotel for 70 euros, or a zero-star hotel for 20 euros, the user might prefer the one-star hotel. On the other hand, if the choice is between a five-star hotel for 300 euros, and a four-star hotel for 250 euros, the four-star hotel may be sufficient. That is, stars cannot be simply weighted with money. How much an extra star is valued depends on the alternatives. As a result, it is more convenient to separate the decision-making process in two steps: (i) a learning phase, in which a number of alternatives are presented; and (ii) a selection phase, in which one alternative is chosen among according to the preferences of the decision-maker. Summarizing, it is often easier and natural a decision-making process in which alternative policies are presented before selection, allowing what-if analysis.

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 or Monte Carlo methods. 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, Qi, and Hirschberger, 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.

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

5.2 Formalizing the multiobjective 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, Westerfield, and Jordan, 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 function can be defined as:

$$c(x_t, b_t) = \Gamma(x_t) + H(b_t) \quad (5.1)$$

where $\Gamma(x_t)$ is a transfer cost function, x_t is the transfer made on day t , $H(b_t)$ 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:

$$b_t = b_{t-1} + x_t + f_t \quad (5.2)$$

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

$$\Gamma(x_t) = \begin{cases} \gamma_0^- - \gamma_1^- \cdot x_t & \text{if } x_t < 0, \\ 0 & \text{if } x_t = 0, \\ \gamma_0^+ + \gamma_1^+ \cdot x_t & \text{if } x_t > 0. \end{cases} \quad (5.3)$$

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

$$H(b_t) = \begin{cases} -u \cdot b_t & \text{if } b_t < 0; u > 0, \\ v \cdot b_t & \text{if } b_t > 0; v > 0. \end{cases} \quad (5.4)$$

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 :

$$X^* = \underset{X}{\operatorname{argmin}} E(C(X, T)) \quad (5.5)$$

where:

$$E(C(X, T)) = \frac{1}{T} \sum_{t=1}^T c(x_t, b_t) = \frac{1}{T} \sum_{t=1}^T [\Gamma(x_t) + H(b_t)] \quad (5.6)$$

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, Frey, and Embrechts, 2005). On the other hand, risk management allows to determine the cash buffer that companies need to hold for precautionary purposes (Ross, Westerfield, and Jordan, 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, Frey, and Embrechts, 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 Ballesterio (2005) and 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; 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, Harford, and Li (2013) when defining the determinants of corporate cash policy and Pinkowitz, Stulz, and Williamson (2016) or Bates, Kahle, and Stulz (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 multiobjective 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:

$$X^* = \underset{X}{\operatorname{argmin}} [E(C(X, T)), R(C(X, T))] \quad (5.7)$$

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 $C(X, T)$:

$$R(C(X, T)) = \left(\frac{1}{T} \sum_{t=1}^T (c(x_t) - E(C))^2 \right)^{1/2}, \quad (5.8)$$

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:

$$R(C(X, T)) = \left(E \left[\max(c(x_t) - E(C), 0)^2 \right] \right)^{1/2}. \quad (5.9)$$

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 multiobjective 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 multiobjective function when rewards can be easily translated into monetary value is the use of a linear scalarization function (Roijers et al., 2013). 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; Ballestero and Romero, 1998; Yu, 2013) is applied to optimize the general minimization multiobjective function defined in equation (5.7). 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 (Ballestero and Romero, 1998). First, a cost index θ_1 , defined as:

$$\theta_1(X, T) = \frac{C(X, T) - C_{min}}{C_{max} - C_{min}} \quad (5.10)$$

where $C(X, T)$ is the expected daily cost over T computed for a particular policy X using equation (5.6), 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:

$$\theta_2(X, T) = \frac{R(X, T) - R_{min}}{R_{max} - R_{min}} \quad (5.11)$$

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.

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

In this section, a real case of our multiobjective 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.

5.3.1 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:

$$x_t = \begin{cases} z - b_{t-1}, & \text{if } b_{t-1} > h \\ 0, & \text{if } l < b_{t-1} < h \\ z - b_{t-1}, & \text{if } b_{t-1} < l \end{cases} \quad (5.12)$$

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;

¹Miller 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, Westerfield, and Jordan (2002).

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 5.2 where the policy X is determined by the control limits $\{h \geq z \geq l\}$ according to equation (5.12).

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 5.2 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 (5.12).

In our case study, we use a data set $F \in \mathbb{R}$ 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 our model, we make computations using three different cost scenarios as shown in Table 5.1. Firstly, we consider a Low, Medium and High cost scenario with increasing costs selecting values along the lines of those suggested in 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).

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 %

Table 5.1: Three alternative scenarios with increasing costs including holding/shortage and transaction costs.

5.3.2 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 5.1. 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 5.1 is summarized in Table 5.2, 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.

5.3.3 Selecting policies considering risk preferences

In the context of the CMP, risk preferences of cash managers are worth considering. Recall from Section 5.2, that the goal is minimizing the multiobjective function in equation (5.7). In our illustrative example, maximizing utility amounts to minimizing a weighted loss function $\mathcal{L}(\theta_1, \theta_2)$:

$$\mathcal{L}(\theta_1, \theta_2) = w_1\theta_1 + w_2\theta_2 \quad (5.13)$$

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.

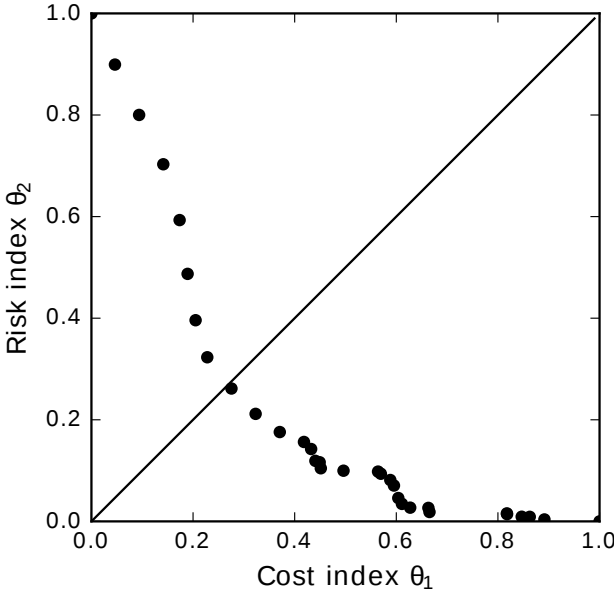


Figure 5.1: Efficient frontier in the cost-risk (θ_1, θ_2) space for the Medium cost scenario using the standard deviation as a measure of risk

According to Ballestero 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 Ballestero (1998) that the utility optimum lies on the normalized efficient frontier between points L and L_∞ , which are defined as follows:

- 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.

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

Table 5.2: Central sample of the efficient frontier from Figure 5.1 (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 $r_0 = 0.5$).

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 5.2 we obtain that bound $L = 16$ when $r_0\theta_1 + \theta_2$ is minimum, and that bound $L_\infty = 9$ when $\theta_1 \approx \theta_2$.

5.3.4 Benchmarking selected policies and risk measures

In this section, we compare in Table 5.3 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, West-erfield, and Jordan (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 multiobjective 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 multiobjective approach.

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

Table 5.3: Relative cost-risk performance of policies for alternative cost scenarios using two measures of risk. The lower the percentage the better.

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 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.

5.3.5 The Sharpe-like ratio

From the analysis of the efficient frontier, a number of insights might be derived. Firstly, in Figure 5.1 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 5.2. 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 5.2. 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.

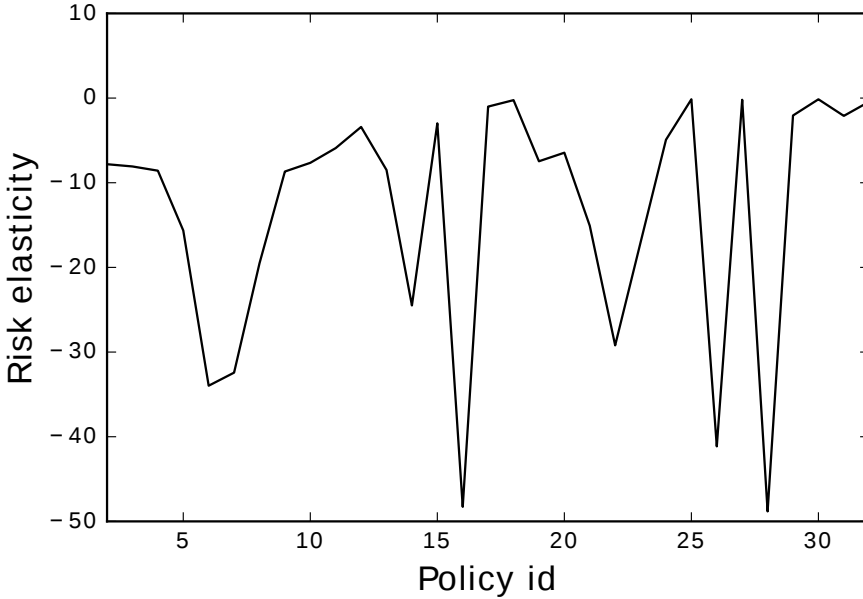


Figure 5.2: Elasticity of the efficient frontier

Moreover, we propose the use of a Sharpe-like ratio (SLR) that can be easily computed for each policy as a synthetic and quick performance measure. The Sharpe ratio (Sharpe, 1966; Sharpe, 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 a single figure. We define the SLR as the geometric average of both cost and risk of each policy:

$$SLR(X) = \sqrt{\theta_1 \cdot \theta_2}. \quad (5.14)$$

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.

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:

$$WSLR(X) = [\theta_1^{w_1} \cdot \theta_2^{w_2}]^{1/(w_1+w_2)}. \quad (5.15)$$

Note that when $w_1 = w_2$, the cash manager has no preference and expressions (5.14) and (5.15) are equivalent. Figure 5.3 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.

An interesting property can be derived from the analysis of the WSLR plot for different weights as shown in Figure 5.3. 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 5.1. 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.

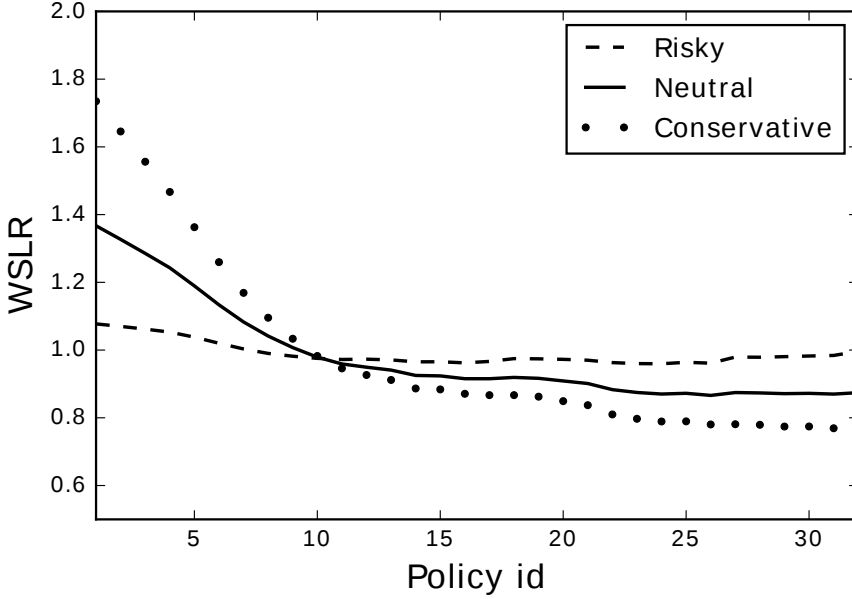


Figure 5.3: Weighted Sharpe-like ratio for different risk preferences and each of the policies. In the Risky case, $w_1 = 0.8, w_2 = 0.2$, in Neutral, $w_1 = 0.5, w_2 = 0.5$, and in Conservative, $w_1 = 0.2, w_2 = 0.8$

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:

$$w_1 \cdot \theta_1 + w_2 \cdot \theta_2 = w_1 \cdot \theta_1 + (1 - w_1) \cdot \theta_1 = w_1 \cdot \theta_1 + \theta_1 - w_1 \cdot \theta_1 = \theta_1. \quad (5.16)$$

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

$$[\theta_1^{w_1} \cdot \theta_1^{w_2}]^{1/(w_1+w_2)} = [\theta_1^{w_1} \cdot \theta_1^{1-w_1}] = \theta_1^{w_1+1-w_1} = \theta_1 \quad (5.17)$$

and for the harmonic mean:

$$\frac{w_1 + w_2}{\frac{w_1}{\theta_1} + \frac{w_2}{\theta_2}} = \frac{1}{\frac{w_1}{\theta_1} + \frac{1-w_1}{\theta_1}} = \frac{\theta_1^2}{w_1 \cdot \theta_1 + \theta_1 - w_1 \cdot \theta_1} = \frac{\theta_1^2}{\theta_1} = \theta_1 \quad (5.18)$$

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

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 the comparison of 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.

5.4 On the utility of cash management models

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

Definition 2. The Cash Management Utility Problem is defined in a multiobjective framework, when less is better, as the problem of determining if policy X is preferred to a baseline policy X_0 , formally expressed as:

$$[g_1(X), \dots, g_q(X)] \leq [g_1(X_0), \dots, g_q(X_0)] \quad (5.19)$$

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

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

Definition 3. A policy X is preferred to a No-Trans policy $X_0 = 0$, in terms of cost and risk indexes θ_1 and θ_2 , when less is better, if:

$$\theta_1(X) + \theta_2(X) \leq \theta_1(X_0) + \theta_2(X_0) \quad (5.20)$$

subject to:

$$X \in S. \quad (5.21)$$

The implications of the CMUP are twofold. First, practitioners may be interested in finding the external conditions that must hold to ensure the utility of a non-trivial policy. An example of this issue was pointed out by Constantinides and Richard (1978), showing that a No-Trans policy is the best alternative in terms of cost when $\gamma_1^+ > u$ and $\gamma_1^- > v$. Second, researchers may be interested in establishing the particular characteristics that both cash management models and alternative cost and risk measures must present in order to avoid non-triviality. As an example, consider the average daily cost as a measure of cost and the daily cost variance as a measure of risk. Let us denote \mathbf{c} as a $\tau \times 1$ vector of daily costs, and \mathbf{d} as a $\tau \times 1$ vector of cost deviations around the average that can be computed as:

$$\mathbf{d} = \mathbf{c} - \frac{\mathbf{1} \cdot \mathbf{1}^T \cdot \mathbf{c}}{\tau} \quad (5.22)$$

where $\mathbf{1}$ is a $\tau \times 1$ vector of ones and superscript T denotes transposition. Cash managers aiming to minimize both cost and variance, or tantamount standard deviation, can derive optimal policies through the following quadratic objective function:

$$\min \mathbf{1}^T \cdot \mathbf{c} \frac{\mathbf{d}^T \cdot \mathbf{d}}{\tau} \quad (5.23)$$

This setting reduces the CMUP to:

$$\mathbf{1}^T \cdot \mathbf{c} + \mathbf{d}^T \cdot \mathbf{d} \leq \mathbf{1}^T \cdot v \cdot \hat{\mathbf{b}}_{t,0} + \mathbf{d}_0^T \cdot \mathbf{d}_0 \quad (5.24)$$

subject to:

$$\mathbf{d} = \mathbf{c} - \frac{\mathbf{1} \cdot \mathbf{1}^T \cdot \mathbf{c}}{n} \quad (5.25)$$

$$\mathbf{d}_0 = v \cdot \hat{\mathbf{b}}_{t,0} - \frac{\mathbf{1} \cdot \mathbf{1}^T v \cdot \hat{\mathbf{b}}_{t,0}}{n}. \quad (5.26)$$

In other words, the utility of a particular cash management model in the previous multiobjective framework is given by the combined ability of the model

to reduce both the cost and risk impact by introducing some control actions summarized in policy X . Furthermore, the CMUP can also be viewed as a precautionary tool to avoid unnecessary efforts in forecasting and mathematical programming tools when some inputs of the problem reduce the utility of the policy.

5.5 Summary

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 multiobjective 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.

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 ratio 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.

We further elaborate on the utility of cash management models by formalizing the problem from a multiobjective perspective in which we compare the

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

Summarizing, 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. A new multiobjective approach to the cash management problem based on compromise programming is now available for cash managers answering the following research question.

***Question 4.** Can we incorporate risk analysis to the cash management problem?*

***Contribution 4.** We formulated the cash management problem from a multiobjective perspective considering both cost and risk.*

The model presented in this chapter is time-invariant, meaning that the context in which the model is deployed does not vary with time. However, chances are that financial circumstances vary with time. This fact makes robustness an additional goal worth considering by cash managers as we do in the following chapter.

On the use of distance indexes to find robust policies

Cash management decision making can be handled from a multiobjective perspective by optimizing not only cost but also risk. Nevertheless, choosing the best policies under a changing context is by no means straightforward. To this end, we rely on compromise programming to incorporate robustness as an additional goal to cost and risk within a multiobjective framework. As a result, we propose to calculate robustness as a multiple criteria distance index that is able to identify the best compromise policies in terms of cost and risk. Such policies are also robust to cash flow regime changes. We show its utility by transforming the Miller and Orr's cash management model into its robust counterpart using real data from an industrial company.

6.1 Motivation

Financial and economic methods are constantly driven by expectations. Decision-makers usually have to deal with a set of relevant variables subject to regime changes because of the economic cycle or other sources of instability. From a macro-economic approach, international economic crises, credit restrictions, or monetary policies are notably influential on decision-makers. Meanwhile, from the micro-economic approach, market changes are usually the rule rather than the exception. Cash flow management is particularly affected by regime

changes. A cash manager must make daily decisions about her firm's optimal cash level for operational and precautionary purposes. The focus is placed on finding the balance between cash holdings and short-term investments. 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 cash management models, from the first proposals to the most recent contributions can be found in Costa Moraes, Nagano, and Sobreiro (2015), Gregory (1976), and Srinivasan and Kim (1986).

The particular characteristics of cash flows (or cash balances) are an important dimension of the CMP. In this sense, cash flows used in the literature range from deterministic (Baumol, 1952) to completely uncertain (Miller and Orr, 1966). The most usual cash flow probability distribution in the literature is the Gaussian in the form of a Wiener process (Constantinides and Richard, 1978; Premachandra, 2004; Baccarin, 2009). However, empirical distributions obtained from real data sets are hardly used with the exception of Gormley and Meade (2007). Beyond the discussion about the most appropriate distribution, cash management models are usually designed from a stationary perspective without considering possible cash flow regime changes. However, economic cycles, market evolution and customer behavior make cash flow a random variable. Thus, optimal solutions to mathematical programming models can be importantly affected by slight perturbations in the input data (Ben-Tal and Nemirovski, 1999). Hence, we say that a system is robust when it is relatively insensitive to changes in environmental factors that can negatively affect its performance (Montgomery, 2013). However, whether cash management models are robust to regime changes remains a question unanswered by previous cash management contributions in the literature.

In order to provide robust solutions to optimization problems, stochastic programming and robust optimization represent alternative approaches. Stochastic programming assumes that randomness in problem parameters is restricted to a particular probability distribution (Sahinidis, 2004; Abdelaziz, Aouni, and El Fayedh, 2007; Aouni and La Torre, 2010). As an alternative approach to stochastic programming, Soyster (1973) introduced the concept of interval uncertainty within the framework of a linear program and its robust counterpart. Later on, Ben-Tal and Nemirovski (1999) and Ben-Tal, El Ghaoui, and Nemirovski (2009), established a sound framework for robust optimization. Robust optimization is a modeling methodology, combined with computational tools, to process optimization problems in which the data are uncertain and is only known to belong to some uncertainty set (Ben-Tal and Nemirovski, 2002). Instead of seeking a probabilistic solution, the decision-maker derives a solution

that is feasible for any realization of a given uncertainty set (Bertsimas, Brown, and Caramanis, 2011). Both approaches, stochastic programming and robust optimization, aim to transform any optimization problem under uncertainty to a collection of deterministic problems that can be solved using state-of-the-art mathematical programming solvers. However, the underlying assumption on the stationarity of probability distributions and uncertainty sets represents an important limitation.

A suitable technique to validate the utility of solutions in a varying (non-stationary) context is the widely used practice in machine learning known as hold-out validation (Mitchell et al., 1997; Hastie et al., 2005; Provost and Fawcett, 2013). Hold-out validation estimates the accuracy of predictive models on a data set different from the original data sample. On the other hand, compromise programming (Ballesterio and Romero, 1998; Yu, 2013), allows a two-step decision making process in which alternative solutions are presented before selection according to their particular preferences.

As a result, we rely on hold-out validation and compromise programming to ensure the robustness of cash management models. More precisely, we follow a data-driven approach to incorporate robustness as an additional goal to cost and risk within a multiobjective decision-making framework. To this end, we assume that regime changes (if any) are contained in a test data set of past cash flows. We later compute a robustness index for cash management policies that is finally used to select the best one. Our procedure departs from previous approaches to deal with uncertainty in the sense that:

1. we learn uncertainty from experience through a data set of past observations;
2. we propose a data-driven approach to test robustness in cash management when possible regime changes are contained in a test data set;
3. we introduce robustness as an additional goal to cost and risk within a multiobjective optimization framework.

For illustrative purposes, we use a real cash flow data set to select the best compromise policies in terms of cost, risk and robustness using an extended version of the one by Miller and Orr (1966). By following our procedure, we show how to transform the simple Miller-Orr's model into its robust counterpart. Summarizing, we support cash managers aiming to analyze cost and risk of cash policies in a changing environment by providing a procedure to ensure robustness in real world applications. More precisely, our contribution:

1. follows a novel data-driven approach to deal with uncertainty in cash management without assuming any theoretical probability distribution;
2. tests robustness of cash management models within a changing environment;
3. transforms a cash management model into its robust counterpart.

Then, under the framework of compromise programming for \mathcal{L}_1 , the goal is to find the policy X that minimizes the weighted Manhattan distance to the ideal point $(0, 0)$:

$$\min [w_1 \cdot \theta_1(X, T) + w_2 \cdot \theta_2(X, T)] \quad (6.1)$$

subject to the following cash balance state equation:

$$\hat{b}_t = \hat{b}_{t-1} + \hat{f}_t + x_t \quad (6.2)$$

where \hat{f}_t and \hat{b}_t denote predicted cash flows and balances, respectively. Since cash managers make decisions in advance, predicted instead of actual values are used for random cash flows and balances resulting in an uncertain optimization problem. However, since neither θ_1 nor θ_2 are able to assess robustness of alternative cash policies, we require a new measure to do so. Next, we consider such a measure.

6.2 On the use of distances as a measure of robustness

A system or a process is robust when it is relatively insensitive to changes in environmental factors, operating conditions and components that can negatively affect its performance (Montgomery, 2013). Therefore, cash management models considering cost and risk objectives must also be evaluated in terms of robustness to deal with changing conditions. Let us illustrate this concept with an example. Consider that an empirical probability distribution $p(b)$ of a cash balance b can be derived from an initial cash balance and a set F of observed past values of cash flows. A usual assumption is that the stochastic cash balance is stationary, i.e., the main attributes of $p(b)$, such as its mean and its standard deviation, do not vary with time. However, chances are that cash flow regime changes lead to variations in both cost and risk of a particular cash policy. For example, assume that some $p(b)$ for a particular company is characterized by a mean value of 80 and a standard deviation of 20. Also assume that, under the framework of the Miller and Orr's cash management

model, the low bound of the best policy is set to 30. This policy clearly becomes suboptimal if the standard deviation rises to 40 due to, for example, a negative market evolution in the last 6 months. In what follows, we first describe our data-driven procedure to ensure robustness of cash policies. Second, we propose the use of a robustness index as a key input to the previous procedure.

6.2.1 A data-driven procedure to ensure robustness

Within the MOCMP in Salas-Molina, Pla-Santamaria, and Rodriguez-Aguilar (2016), policies are evaluated through cost and risk indexes θ_1 and θ_2 , described in Chapter 5. However, these indexes do not provide information about policy robustness. In order to ensure robustness in cash management, we here rely on a data-driven procedure. More precisely, we apply the widely used practice in the field of machine learning consisting of dividing the available data in a training set and a test set (Mitchell et al., 1997; Hastie et al., 2005; Provost and Fawcett, 2013). The utility of predictive models is then assessed by estimating the accuracy of forecasts on a data set (test set) different from the original data sample (training set). In the context of the MOCMP, we propose to use a training data set as a business-as-usual situation, and a test set that may contain changes in the context faced by cash managers. As long as the training set is statistically different to the test set (but possibly many others), we are in a position to measure policy robustness as the change in performance over two different data sets. Thus, we define policy performance as follows:

Definition 4. (Policy performance). The performance P_i of a cash policy i in a normalized cost-risk space is given by the pair of points $P_i = (\theta_{1i}, \theta_{2i})$, computed using (5.10) and (5.11), over a given cash flow data set.

Recall that we do not assume any cash flow probability distribution but we are provided with a data set of past observations. Since we are dealing with cash flow time-series, we may be reasonably interested in evaluating the effect of recent cash flow data on policy performance. Thus, we perform an empirical performance test with the most up-to-date data. If any regime change occurred in recent past, it will be reflected in the performance of policies derived from solving the MOCMP. A feasible solution to the MOCMP is a policy derived from a cash management model (e.g., a policy of the Miller and Orr (1966) model type), that satisfies the cash balance state equation (6.2). Hence, given a feasible policy derived from a particular cash management model, we assess its robustness by computing the difference in performance over a training set (e.g., with the first 80 % of the observations), and over a test set (with the

remaining 20% of the observations). Summarizing, our data-driven procedure to ensure robustness in cash management comprises the following steps:

1. Select a cash management model.
2. Generate a set of alternative feasible policies.
3. Divide the cash flow data set in a training set and a test set.
4. Evaluate policy performance in the training set.
5. Evaluate policy performance in the test set.
6. Measure robustness through a new index that considers policy performance with respect to the training and test sets.
7. Select the best policies using a robustness index.

Next, we propose a collection of distances to measure policy robustness in a multiobjective framework in which cost and risk are goals to minimize.

6.2.2 *Measuring robustness through distance indexes*

In order to derive a robustness index for cash policies, we rely on the concept of distance. A comprehensive work on distances can be found in Deza and Deza (2014). We here define distance as a function in a normalized cost-risk space:

Definition 5. (Distance). In a normalized cost-risk space $S = [0, 1]^2$, a distance is a function $D : S \times S \rightarrow \mathbb{R}^+$, with the following properties:

1. $D(P, P') \geq 0$ (non-negativity);
2. $D(P, P') = D(P', P)$ (symmetry);
3. $D(P, P) = 0$ (reflexivity);
4. $D(P, P'') \leq D(P, P') + D(P', P'')$ (triangle inequality);

for all points $P, P', P'' \in S$.

Hereafter, we propose to test the robustness of policies to cash flow regime changes by computing distances between performance points derived from policies in a training set denoted as $P_i = (\theta_1, \theta_2)$, and performance points derived

from policies in a test set denoted as $P'_i = (\theta'_1, \theta'_2)$. Thus, we propose to use the following collection of distances to measure robustness of policies that we illustrate in Figure 6.1. In the business-as-usual context, policy performance can be synthetically evaluated in terms of cost and risk by means of the (training) distance to the ideal performance $(0, 0)$.

Definition 6. (Training distance). The training distance δ_i of policy i is the Euclidean distance between performance P_i and the ideal policy performance $(0, 0)$.

The shorter the training distance the better the policy within a business-as-usual context. However, regime perturbations may result in performance changes of alternative policies. Then, we introduce the concept of deviation as the difference in performance between policies over the training and test sets in a normalized cost-risk space.

Definition 7. (Deviation). The deviation d_i of policy i is the Euclidean distance between performance P_i and performance P'_i , namely:

$$d_i = \sqrt{(\theta_{1i} - \theta'_{1i})^2 + (\theta_{2i} - \theta'_{2i})^2}. \quad (6.3)$$

According only to the magnitude of deviation, we can initially state that policy 2 in Figure 6.1, with deviation d_2 , is more robust than policy 3, with deviation d_3 , because $d_2 < d_3$. However, although d_1 and d_2 have the same magnitude, a rational cash manager should select policy 2 instead of 1 because the former is closer to the ideal point $(0, 0)$ than the latter. This motivates the need for an additional distance measure.

Definition 8. (Proximity or test distance). The proximity δ'_i of policy i is the Euclidean distance between performance P'_i and the ideal policy performance $(0, 0)$.

In addition, a policy that improves its performance after a context change is preferred to a policy that worsens its performance. For example, policy 2 in Figure 6.1 is an even better policy than policy 1 because, after the evaluation over the test set, performance P'_2 moves closer to the ideal point while P'_1 moves away from the ideal point. Thus, we should also measure movement of policies.

Definition 9. (Movement). The movement of policy i is the difference in distance to the ideal performance between policy performance P'_i and P_i , computed as $\delta'_i - \delta_i$.

We next summarize how each of the aforementioned distances to measure robustness rank policies when the rest of factors (distances) are equal:

- Deviation: policy i is better than policy j if $d_i < d_j$.
- Proximity: policy i is better than policy j if $\delta'_i < \delta'_j$.
- Movement: policy i is better than policy j if $\delta'_i - \delta_i < \delta'_j - \delta_j$.




Figure 6.1: Distances between cash policies evaluated over two data sets.

The collection of distances mentioned above allows us to introduce different measures of robustness according to the particular preferences of cash managers. Next, we introduce two alternative distance functions to measure robustness that we later compare in a numerical example:

$$D_1(P_i, P'_i) = \begin{cases} 2\alpha(\delta_i - \delta'_i) & \text{if } \delta_i > \delta'_i, \\ 2(1 - \alpha)(\delta'_i - \delta_i) & \text{if } \delta_i \leq \delta'_i \end{cases} \quad (6.4)$$

as a linear asymmetric loss function, with $\alpha \in [0, 1]$, and:

$$D_2(P_i, P'_i) = \alpha\delta_i + \beta\delta'_i + (1 - \alpha - \beta)(\sqrt{2} + \delta'_i - \delta_i) \quad (6.5)$$

as a synthetic loss function where $\alpha, \beta \in [0, 1]$ are weights to set preferences objectives, and $\sqrt{2}$ is only added to ensure non-negativity. Notice that (6.4) considers only movement but (6.5) takes into account deviation, proximity and movement. Cash managers can also consider their particular preferences by tuning parameters α and β . In (6.4), parameter α can be tuned to weigh movement. Alternatively, in (6.5), parameter α is linked to deviation in performance between regimes so that the higher the value of α the higher the importance of deviation. Parameter β is linked to the proximity of policies to the ideal point so that the higher the value of β the higher the importance of proximity. Finally, the values given to α and β determine the importance of the movement, since it is weighed by $1 - \alpha - \beta$.

Next, we apply compromise programming (Ballester and Romero, 1998) to find robust policies. Recall that this approach is based on the concept of ideal point where both the minimum cost and the minimum risk occur simultaneously, i.e., $(0, 0)$. This point is usually unfeasible and it is necessary to look for compromise solutions that can be evaluated by computing the distance to the ideal point. Thus, once an entire cash flow data set is divided in a training set and a test set, we first evaluate a given set of feasible policies over the training set to obtain performance points $P_i = (\theta_{1i}, \theta_{2i})$. However, when selecting a policy from the whole cost-risk space, cash managers may be interested in testing the robustness to cash flow regime changes. To this end, we propose the use of a third index θ_3 , namely, a distance index as a measure of robustness when applied to data not considered in the first selection of policies. As a result, given a policy X_i , with expected performance $P_i = (\theta_{1i}, \theta_{2i})$, obtained from the training set, and with empirically tested performance $P'_i = (\theta'_{1i}, \theta'_{2i})$, obtained from the test set, we use a distance function (e.g., (6.5)), to compute the following distance index:

$$\theta_3(X_i) = \frac{D(P_i, P'_i) - D_{min}}{D_{max} - D_{min}} \quad (6.6)$$

where D_{max} and D_{min} are, respectively, the maximum and minimum values of performance distances $D(P_i, P'_i)$ between the training and the test set for each policy i . At this point, we are in a position to derive an efficient frontier with the set of policies that are not dominated by any other policy in terms of cost and robustness. Graphically, the efficient frontier in the normalized cost-robustness space, is built from a set of two-dimensional points (θ_1, θ_3) , as shown in Figure 6.2. Note that since we compute robustness using both cost and risk indexes, we are implicitly considering risk in the selection of policies. As a result, by following our distance-based procedure, cash managers aiming to analyze cost and risk of cash management policies in a changing environment

are able to ensure robustness in real world applications. Next, we illustrate our procedure and its benefits through a numerical example.

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

In this section, we follow the procedure described in Section 6.2.1 to determine the best compromise cash policies in terms of cost and risk that are also robust to cash flow regime changes.

6.3.1 Empirical settings

In what follows, we use the Miller and Orr's cash management model introduced in Chapter 2 to obtain control limits h , z and l , that determine policies according to equation (2.8). In our case study, we use a real data set from a Spanish industrial company composed by 1000 observations of cash flows at non-bank-holidays for a period of about 4 years. We divide the entire data set in a training set with the first 80% of the observations, as the expected conditions, and a test set with the remaining 20%, as representative of changes in current conditions.

For illustrative purposes, we make computations using a cost structure adjusted to current bank practices in Spain, summarized as follows: daily holding cost, $v = 0.1\%$; daily shortage cost, $u = 30\%$, to be used in equation (5.4); fixed transfer cost into account, $\gamma_0^+ = 5 \text{ €}$; fixed transfer from account $\gamma_0^- = 5 \text{ €}$; variable transfer cost into account, $\gamma_1^+ = 0 \text{ €}$, and variable transfer cost from account, $\gamma_1^- = 0 \text{ €}$, to be used in equation (5.3). It is important to note that $\gamma_1^- < v$ and $\gamma_1^+ < u$. Otherwise, no transaction would be made since the unitary costs of transferring money are higher than those of holding the same amount of money.

6.3.2 Selecting robust policies

In order to derive the efficient frontier in the cost-robustness space, shown in Figure 6.2, 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 , and l . Since our cost structure presents a high shortage cost, we expect better policies to be given by high control limits. Consequently, we iterate over feasible combinations of h , z , and l subject to $h \geq z \geq l$, ranging in $[h - \sigma, h + 3\sigma]$, $[z - \sigma, z + 3\sigma]$, $[l - \sigma, l + 3\sigma]$, in steps of 100000 €, where σ is the standard deviation of cash flows in the training set.




Figure 6.2: Efficient frontier in the cost-robustness space (θ_1, θ_3) .

In this chapter, we present alternative policies before selection. Therefore, we evaluate all possible policies in terms of cost and risk. Under the compromise programming framework, the closer to the ideal point $(0, 0)$, the better the policy. A sample of the best policies of the efficient frontier, in terms of Euclidean distance to the ideal is shown in Table 6.1. To compute θ_3 we use distance function (6.5) with parameters $\alpha = 0.33$ and $\beta = 0.33$, to equally weigh deviation, proximity and movement. Then, a cash manager would select either policy 35, 38, 39 or 40 to minimize the distance to the ideal $(0, 0)$. However, the particular preferences of cash managers should be considered to select other policies. A risky cash manager may select one of the policies with larger distances (e.g., policy number 32), in order to reduce daily cost. On the other hand, a conservative cash manager may consider accepting a higher expected cost (e.g., by selecting policy 45), in order to maximize robustness.

<i>Id</i>	High(<i>h</i>)	Target(<i>z</i>)	Low(<i>l</i>)	$D_2(P_i, P'_i)$	θ_1	θ_3	Dist
32	6310	4760	1780	0,05	0,18	0,28	0,33
33	6410	4760	1780	0,05	0,18	0,28	0,33
34	6510	4760	1780	0,05	0,18	0,28	0,33
35	6210	4660	1580	0,05	0,18	0,24	0,30
36	6310	4660	1580	0,04	0,19	0,24	0,31
37	6410	4660	1580	0,04	0,19	0,24	0,31
38	6110	4560	1480	0,04	0,22	0,20	0,30
39	6210	4560	1480	0,04	0,23	0,20	0,30
40	6310	4560	1480	0,04	0,23	0,20	0,30
41	6010	4460	1380	0,03	0,27	0,16	0,31
42	6110	4460	1380	0,03	0,27	0,16	0,31
43	6210	4460	1380	0,03	0,27	0,16	0,31
44	6010	4460	1480	0,03	0,31	0,15	0,34
45	6110	4460	1480	0,03	0,32	0,15	0,35

Table 6.1: Central sample of the efficient frontier (*Id*=policy id; Bounds *h*, *z* and *l* in thousands of euros; $D_2(P_i, P'_i)$ = distance function; θ_1 = cost index; θ_3 = robustness index; Dist=Euclidean distance to the ideal point $(\theta_1, \theta_3) = (0, 0)$).

6.3.3 Analysis

In this example, we use a robustness index to select the best compromise policy in terms of cost and robustness to regime changes. However, an additional evaluation in terms of cost and risk may be useful for comparative purposes. Such an evaluation over the test set in our illustrative example is shown in Table 6.2. As a baseline, we use the best compromise policy obtained using the entire data set but without any robustness index as in Salas-Molina, Pla-Santamaria, and Rodriguez-Aguilar (2016). In addition, we evaluate three different distance measures: (i) the asymmetric loss function in (6.4) with $\alpha = 0.4$ to slightly underweigh the loss of policies that move closer to the ideal point and yield an improvement in cost; (ii) the synthetic loss function in (6.5) considering both deviation and proximity leading to a smaller improvement than the asymmetric function and an increase in risk; (iii) the synthetic loss function in (6.5) but considering deviation, proximity and movement, resulting in a reduction in cost but an increase in risk. Summarizing, the use of a robustness index in our example results in a cost reduction of policies. In addition, considering deviation, proximity and movement produces the best results in terms of cost. However, this cost reduction is only possible by accepting a higher level of risk.

Distance	High(h)	Target(z)	Low(l)	α	β	Cost(%)	Risk(%)
Salas-Molina et al. (2016)	6570	4910	3930	-	-	100	100
Asymmetric loss D_1 in (6.4)	6010	4460	3580	0,4	-	91	99
Synthetic loss D_2 in (6.5)	6610	4960	3380	0,5	0,5	94	112
Synthetic loss D_2 in (6.5)	6110	4560	1480	0,33	0,33	75	132

Table 6.2: Evaluation of policies using different distance measures. Bounds h , z and l in thousands of euros.

6.4 Discussion

In this section, we discuss our distance-based procedure in comparison to alternative approaches to obtain robust solutions to optimization problems such as stochastic programming and robust optimization. In this chapter, we handle uncertainty from an empirical or data-driven approach which considerably departs from previous approaches. More precisely, the use of a distance index as a measure of robustness of cash policies allows to include an additional goal to cost and risk within the framework of multiobjective decision-making. This additional goal can be designed according to the particular preferences of decision-makers. Furthermore, since the selection of the best compromise policies is affected by the uncertainty in cash flows, our approach is specially designed to deal with input data that can be learned from experience through a data set of past observations such as cash flows. At the same time, we follow a hold-out validation procedure based on the comparative analysis of a training data set and a test set that may contain a context change. Finally, we allow practitioners to transform any cash management model into its robust counterpart by following our distance-based approach to find robust cash management policies.

For comparative purposes, let us consider two different approaches to deal with uncertain optimization problems, namely, stochastic programming and robust optimization. Under a general stochastic programming approach (Prékopa, 2013), a deterministic counterpart of the program (6.1)-(6.2) can be considered as follows:

$$\min [w_1 \cdot E(\theta_1(X, T)) + w_2 \cdot E(\theta_2(X, T))] \quad (6.7)$$

subject to:

$$P(\hat{b}_{t-1} + \hat{f}_t + x_t \geq \hat{b}_t) \geq \zeta \quad (6.8)$$

where E is the common expectation operator, P denotes probability and ζ is a safety threshold determined by a cash manager. A typical value for ζ is 95%, meaning that cash balances are above the expected values at least 95% of the times, reducing then the possibility of an overdraft. In other words,

by solving the previous program we are forcing cash balances to be above a certain level determined by the safety threshold ζ . However, the introduction of probabilities implies the assumption of a particular distribution for cash flows resulting in two serious limitations due to: (i) the lack of empirical evidence on the common assumption of normally distributed cash flows (Emery, 1981; Pindado and Vico, 1996); (ii) the additional assumption on the stationarity of cash flow distributions (Constantinides and Richard, 1978; Premachandra, 2004; Baccarin, 2009).

An alternative approach was proposed by Soyster (1973) and, later on, by Ben-Tal and Nemirovski (1999) and Ben-Tal, El Ghaoui, and Nemirovski (2009), who tackled uncertainty by proposing a robust counterpart to an optimization problem when uncertainty is determined either by a system of linear inequality constraints or by a system of conic quadratic inequalities. For instance, a robust counterpart of the program (6.1)-(6.2), when uncertainty in cash flows is determined by $\hat{f} - \varepsilon \leq f \leq \hat{f} + \varepsilon$, being ε the maximum forecasting error, is given by:

$$\min [w_1 \cdot \theta_1(X, T) + w_2 \cdot \theta_2(X, T)] \quad (6.9)$$

subject to:

$$\hat{b}_{t-1} + \hat{f}_t + x_t \geq \hat{b}_t + \varepsilon. \quad (6.10)$$

The assumption of a maximum forecasting error ε may lead either to unrealistic or ultraconservative problem formulations. For instance, consider a data set of empirical errors ranging in $[-100, 100]$, in which the maximum negative error -100 occurred only once out of 1 million times. Setting ε to 100 results in an ultraconservative strategy. In addition, two important issues remain unsolved: (i) the stationarity assumption, meaning that estimation errors keep unaltered; and (ii) the possibility of adapting cash policies to different scenarios according to the particular preferences of decision-makers. These limitations are solved by our data-driven distance procedure as described above.

6.5 Summary

In the cash management problem, decision-making can be enhanced by focusing on optimizing cost, risk and robustness of the available policies rather than just cost. This is particularly true under the realistic assumption of time-varying circumstances. Usual approaches to deal with uncertainty within the framework of optimization problems such as stochastic programming or robust optimization present limitations such as imposing a theoretical probability distribution or the stationarity assumption. To overcome these limitations, we

propose a data-driven procedure to test the robustness of cash policies in a multicriteria decision-making process. Particularly, we introduce the use of a distance-based robustness index to select the best compromise set of cash policies when cash flow regime changes are learned from experience through the use of a data set of past observations. As a result, we argue that our procedure is able to transform a cash management model into its robust counterpart when a data set of past cash flow observations is available through:

- the validation of estimated cost-risk results from cash policies when possible regime changes are contained in a test set;
- the use of distance indexes to measure policy robustness;
- the selection of the best compromise policies in terms of cost and risk that are also robust to regime changes through compromise programming.

An additional remark must be done in the sense that we claim that a new procedure is available to optimize cost and robustness by using two different data sets, a training set and a test set. Our procedure can be replicated as many times as needed by considering alternative (and possibly very different) test sets to incorporate a number of future scenarios. Furthermore, our procedure can be straightforwardly extended to the use of any other cross-validation method.

Summarizing, decision-makers usually have to deal with time-varying financial circumstances. As a result, cash managers may be interested in identifying the best compromise policies in terms of cost and risk that are also robust to cash flow regime changes. In this chapter, we relied on compromise programming to incorporate robustness as an additional goal to cost and risk within a multiobjective framework. This procedure allowed us to provide an answer to research question 5.

Question 5. *Can we provide a robust counterpart for any cash management model?*

Contribution 5. *We proposed a new data-driven multiobjective method to derive the robust counterpart for any cash management model by computing distances in a cost-risk space.*

A closely related topic to the previous research question is the concept of operating condition, which we use in this thesis to refer to any factor that may

influence the performance of a model. In the next chapter, we address the problem of selecting cash management models under different circumstances or operating conditions.

Selecting cash management models from a multiobjective perspective

This chapter addresses the problem of selection of 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. Its utility is shown 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.

7.1 Motivation

A Receiver Operating Characteristic (ROC) curve is a useful graphical tool for classification models that illustrates its performance Metz (1978), Bradley (1997), Flach (2003), and 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), Elkan (2001), Drummond and Holte (2006), and Hernández-Orallo, Flach, and Ferri (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), and Hernández-Orallo (2013a).

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, Westerfield, and Jordan, 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), and Costa Moraes, Nagano, and Sobreiro (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, Qi, and Hirschberger, 2007; Garcia-Bernabeu et al., 2016). Cost and risk are desired but conflicting objectives in cash management as suggested by Salas-Molina, Pla-Santamaria, and Rodriguez-Aguilar (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 chapter, 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 chapter 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.

First, recall from Chapter 2, that the model proposed by Stone (1972) introduced the use of forecasts for cash management. 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. Then, 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. Next, a number of simple rules are proposed to return to a desired target cash balance, z , based on two sets of control limits. 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. 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. 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 representation of models in the cost-risk space.

Second, keep in mind that MOCMP can be viewed as a trade-off between cost and risk. Cost reductions are achieved by reducing cash balances but, at the same time, the risk of an overdraft increases. The set of transactions to be made is determined by the cash management model. Then, the expected cost of any model m , can be assessed by computing its average daily loss over a time horizon T , denoted as $C(m, T)$. From the risk perspective, since cash managers are usually interested in avoiding costs above a target value, e.g., the average daily cost, a suitable risk measure is the upside deviation of the cost function, denoted as $R(m, T)$. Hence, for any particular model m evaluated in a time period of T days, we can express the expected cost and risk as:

$$C(m, T) = E(C) = \frac{1}{T} \sum_{t=1}^T C_t \quad (7.1)$$

$$R(m, T) = (E(\max\{C_t - E(C), 0\}^2))^{1/2}. \quad (7.2)$$

As stated above, the ultimate goal is to find the policy X^* , given by model m , that minimizes both the expected cost $C(m, T)$ and risk $R(m, T)$, over time horizon T . This multiobjective optimization problem requires the definition of a general loss function considering $C(m, T)$ and $R(m, T)$, which will be given later.

7.2 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 chapter, 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 10. 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, Lachiche, and Martínez-Usó (2014) and Hernández-Orallo, Flach, and Ferri (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 11. 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 (7.1) and (7.2).

Definition 12. 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 (7.1) and (7.2), and through the following cost and risk indexes:

$$\theta_1(m, T) = \frac{C(m, T)}{C(m_0, T)} \tag{7.3}$$

$$\theta_2(m, T) = \frac{R(m, T)}{R(m_0, T)}. \tag{7.4}$$

By representing points (θ_1, θ_2) for different models as shown in Figure 7.1, 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 7.1 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.




Figure 7.1: Alternative models in the normalized cost-risk space

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, Flach, and Ferri, 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 7.1 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 (Ballesteros and Romero, 1998; Miettinen, 2012).

Definition 13. A solution is called Pareto optimal if there does not exist any other feasible solution that can achieve the same or better performance for all the criteria while being strictly better for at least one criterion.

As a result, the set of Pareto optimal solutions within S forms the Pareto front as depicted in Figure 7.2. The Pareto front provides valuable information about alternative cash management models. 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.

Furthermore, a family of distances functions \mathcal{L}_h to the ideal point $(0, 0)$ can be computed to determine which model is best in terms of both cost and risk as proposed by compromise programming (Ballesteros and Romero, 1998). Under this framework, cash managers seek to find the best set of solutions achievable by a compromise program, namely, the compromise set. More formally, the concept of compromise set derives from an important result presented by Yu (1973) for bi-criteria problems.

Definition 14. The compromise set is a subset of the Pareto front between points \mathcal{L}_1 and \mathcal{L}_∞ obtained by varying parameter $h \in \mathbb{N}^+$ in the following

family of distance functions:

$$\mathcal{L}_h = [\theta_1^h + \theta_2^h]^{1/h} \tag{7.5}$$

where θ_1 and θ_2 are normalized indexes for each of the two criteria.

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.




Figure 7.2: Pareto front of alternative models in the normalized cost-risk space

7.3 Choosing models according to risk preferences

We mentioned in Section 7.2 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:

$$\mathcal{L}_h(m, \alpha) = [w^h \cdot \theta_1^h + (1 - w)^h \cdot \theta_2^h]^{1/h} \quad (7.6)$$

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, 1992). 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, Flach, and Ferri (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:

$$w^2 \cdot \theta_1^2 + (1 - w)^2 \cdot \theta_2^2 = d^2. \tag{7.7}$$

This equation leads to consider isometric circles and ellipses as a tool to select among alternative cash management models as shown in Figure 7.3.

Definition 15. 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 $w < (1 - w)$.

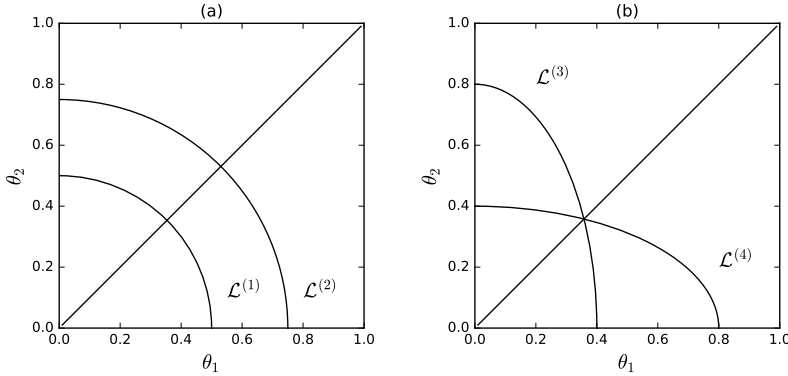


Figure 7.3: (a) Isometrics for neutral cash managers; (b) Isometrics for risky/conservative cash managers

Neutral cash managers may consider a number of circles with increasing radius such as $\mathcal{L}^{(1)}$ and $\mathcal{L}^{(2)}$ in Figure 7.3(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 7.3(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 7.3(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.

7.4 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 (7.6) according to the particular preferences of cash managers.

Recall from Section 7.2 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:

$$\varepsilon = \frac{\sum_{t=1}^N (\hat{f}_t - f_t)^2}{\sum_{t=1}^N (\bar{f} - f_t)^2} \quad (7.8)$$

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 (7.8).

On the other hand, the current cost structure is determined by $\gamma_0^+ = 50 \text{ €}$, $\gamma_0^- = 50 \text{ €}$, $\gamma_1^+ = 0.1\%$, 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 (7.6) 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 4.

Algorithm 4: Algorithm to build context loss curves

- 1 **Inputs:** Model m ; context α ; 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 α_i in α **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 (7.1) and (7.2) 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 7.4(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 predictions 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.




Figure 7.4: (a) Loss curve for two alternative models and different prediction errors; (b) Expected loss after selecting the best model for each prediction error

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 7.4(b) shows the minimum loss that can be achieved after selecting the best model for each prediction error. The first two points 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.

7.5 Summary

This chapter 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. The set of graphical tools and numerical procedures presented in this chapter allow cash managers to analyze the trade-off between cost and risk for different cash management models and operating conditions. As a result, a new contribution can be highlighted.

Question 6. *Under what circumstances or operating conditions a model is better than another?*

Contribution 6. *We adapted ROC analysis to the cash management problem in order to allow cash managers to select cash management models.*

The results presented in this chapter are based on the use of a bound-based model of the Stone (1972) type. However, one may wonder if a new class of cash management models can be designed without using bounds to control cash balances. The next chapter proposes such a new class of cash management models.

Boundless multiobjective models for cash management

In this chapter, we propose a new class of cash management models. Most cash management models are usually based on a set of bounds that complicates the selection of the optimal policies due to non-linearity. We here propose to linearize cash management models to guarantee optimality through linear-quadratic multiobjective compromise programming models. We illustrate our approach through a reformulation of the suboptimal state-of-the-art Gormley-Meade's model to achieve optimality. Furthermore, we introduce a much simpler formulation that we call the Boundless Model that also provides optimal solutions. Results from a sensitivity analysis using real data sets from 54 different companies show that our Boundless Model is highly robust to cash flow prediction errors.

8.1 Motivation

Cash 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 transactions between cash and any other kind of available liquid

asset, resulting in an increase or decrease of the cash level. From those decisions, relevant implications are derived for companies in terms of cost and risk. Indeed, the cash management problem deals with the trade-off between transaction and holding costs, but also with performing operations without taking unnecessary risks.

The ultimate goal of the cash management problem (CMP) is to find the best sequence of control actions over a given planning horizon which is called a policy. The typical framework in the CMP considers two assets, namely, a cash account and an investment account where idle cash balances are allocated for a profit. In addition, most cash management models have a common feature: they are based on setting control limits or bounds. In other words, they can be considered as Bound-Based Models (BBM). In a BBM, cash balance is allowed to wander around between some bounds, usually a high bound and a low bound. When any of these bounds is reached, a control action is made to restore the balance to some target level. For instance, the BBM in Miller and Orr (1966) is based on three bounds and Girgis (1968) and Penttinen (1991) and Gormley and Meade (2007) considered four bounds to account for both fixed and linear transaction costs. Stone (1972) proposed the use of forecasts and five bounds, Neave (1970) and later Chen and Simchi-Levi (2009) proposed optimal policies based on six bounds.

In the aforementioned models, cash managers have to determine the set of bounds which minimizes the sum of expected transaction and holding costs. However, non-linearity introduced by bounds complicates the selection of optimal policies. For instance, Gormley and Meade (2007) and Costa Moraes and Nagano (2014) suggested the use of evolutionary algorithms to solve the CMP. However, this approach does not guarantee the optimality of the cash management solutions. In order to overcome this limitation, we here propose a reformulation of the CMP within a multiobjective framework as proposed by Salas-Molina, Pla-Santamaria, and Rodriguez-Aguilar (2016). More precisely, we encode the CMP as a linear-quadratic compromise programming model (Zeleny, 1982; Yu, 1985; Ballesteros and Romero, 1998). We illustrate our approach through a multiobjective reformulation of the Gormley and Meade (2007) model (MOGM).

Since the ultimate goal of the cash management problem is not to find the best set of bounds, but the best sequence of control actions, we also introduce a straightforward formulation that we call the Boundless Model (BM) that does not impose restrictions on the form of policies. We demonstrate that our BM is a much simpler CMP formulation that guarantees the optimality of the solutions against approximate solutions of a state-of-the-art BBM. Furthermore,

we empirically show that our BM outperforms the MOGM in terms of required run time to solve the CMP.

A critical issue when dealing with cash management models using forecasts as a key input is the influence of forecasting errors on the performance of the model. In order to evaluate this impact, we present a sensitivity analysis of our BM to forecasting errors using 54 real data sets from different companies. The results show that our BM is highly robust to prediction errors. A further advantage of this method is that it allows to estimate the reward that can be obtained by improving predictive accuracy.

Summarizing, the results presented in this chapter imply that cash managers should consider forecasts and our BM as a way to improve decision-making in cash management. In addition, we solve important open research questions in cash management through the following four main contributions:

1. A linear formulation of the CMP to guarantee the optimality of solutions.
2. A linear-quadratic multiobjective reformulation of the Gormley-Meade's model that, unlike the original model, guarantees the optimality of solutions.
3. A novel cash management model, namely, the BM, that provides optimal solutions and generalizes several state-of-the-art models.
4. A sensitivity analysis of the BM to forecasting errors.

In what follows, we first linearize the cash management problem described in Chapter 5. Next, we use this linearization to reformulate the Gormley and Meade (2007) model. Finally, we introduce the BM for multiobjective cash management and we analyze the impact of forecasting errors in the BM's performance.

8.2 Linearizing the cash management problem

The ultimate goal of the MOCMP defined in Chapter 5 is to find the policy X that minimizes the expected cost and risk over the time horizon T . To this end, the expected cost $C(X)$ is measured by the average daily cost:

$$C(X) = \frac{1}{T} \sum_{t=1}^T c(x_t) = \frac{1}{T} \sum_{t=1}^T [\Gamma(x_t) + H(b_t)] \quad (8.1)$$

and the expected risk $R(X)$ is measured by the standard deviation of the daily cost:

$$R(X) = \left(\frac{1}{T} \sum_{t=1}^T (C(X) - c(x_t))^2 \right)^{1/2}. \quad (8.2)$$

Then, under the framework of compromise programming proposed by Balles-tero and Romero (1998), the goal is to find policy X that minimizes a loss function $\mathcal{L}_\alpha(X)$ computing weighted distances to an ideal point where both cost and risk are zero:

$$\min \mathcal{L}_\alpha(X) = \min \left[\left(\frac{w_1}{C_{max}} \cdot C(X) \right)^\alpha + \left(\frac{w_2}{R_{max}} \cdot R(X) \right)^\alpha \right]^{1/\alpha} \quad (8.3)$$

subject to:

$$X \in S \quad (8.4)$$

where α is a positive integer and S is the set of all possible policies. In what follows, we consider only Manhattan distances to the ideal point by setting $\alpha = 1$ in objective function (8.3) to guarantee linearity. Furthermore, C_{max} and R_{max} can be regarded either as budget limitations or as normalization factors for comparative purposes.

Next, we formulate the CMP as a linear program since we aim to guarantee the optimality of solutions. In order to linearize the common two-assets setting of the CMP, consider a company with two bank accounts as depicted in Figure 8.1. Account 1 receives payments from customers (inflows) and it is also used to send payments to suppliers (outflows). Both inflows and outflows are summarized through the net cash flow f_t .




Figure 8.1: The common two-assets setting in the cash management problem.

Let x_t be the difference between inflows and outflows $x_t = x_t^+ - x_t^-$ at bank account 1 with x_t^+ and x_t^- being non-negative real numbers. Then, account 2 represents the amount of alternative investments available to be converted into cash through transaction x_t^+ when needed. In addition, idle cash balances from

account 1 can be allocated in account 2 for a profit through transaction x_t^- . In this setting, the transfer cost function in equation (5.3) can be expressed as follows:

$$\Gamma(x_t) = z_t^+ \cdot \gamma_0^+ + \gamma_1^+ \cdot x_t^+ + z_t^- \cdot \gamma_0^- + \gamma_1^- \cdot x_t^- \quad (8.5)$$

where $z_t^+, z_t^- \in \{0, 1\}$ are binary auxiliary variables satisfying:

$$z_t^+ + z_t^- \leq 1 \quad (8.6)$$

$$m \cdot z_t^+ \leq x_t^+ \leq M \cdot z_t^+ \quad (8.7)$$

$$m \cdot z_t^- \leq x_t^- \leq M \cdot z_t^- \quad (8.8)$$

where $M(m)$ is a very large (small) number. Note that although we restrict x_t^+ and x_t^- to be non-negative numbers, the left-hand side of equations (8.7) and (8.8) are necessary to ensure $z_t^+, z_t^- = 0 \Leftrightarrow x_t^+, x_t^- = 0$. A similar approach can be followed to linearize the holding/penalty cost function (5.4) through:

$$H(b_t) = z_t \cdot h \cdot b_t - (1 - z_t) \cdot u \cdot b_t \quad (8.9)$$

where $z_t \in \{0, 1\}$ is another auxiliary binary variable satisfying:

$$-M(1 - z_t) \leq b_t \leq M \cdot z_t. \quad (8.10)$$

However, since cash managers usually discard policies with negative balances due to high penalty costs, in what follows, we assume $u = \infty$ and $H(b_t) = h \cdot b_t$, which results in the following daily cost function equivalent to equation (5.1):

$$c(x_t) = \gamma_0^+ \cdot z_t^+ + \gamma_1^+ \cdot x_t^+ + z_t^- \cdot \gamma_0^- + \gamma_1^- \cdot x_t^- + h \cdot b_t. \quad (8.11)$$

Since control actions are taken in advance to real cash flow, we use predicted cash flows \hat{f}_t and balances \hat{b}_t instead of actual values in the state transition law in equation (5.2) as follows:

$$\hat{b}_t = \hat{b}_{t-1} + \hat{f}_t + x_t^+ - x_t^- \quad (8.12)$$

Next, we rely on the previous linear functions to derive optimal solutions from a reformulation of the GM model.

8.3 Reformulating the Gormley-Meade’s model as a linear-quadratic program

Recall from Chapter 2 that cash managers can leverage cash flow forecasts to reduce the uncertainty within a short-term planning horizon (Stone, 1972; Stone and Miller, 1987; Salas-Molina et al., 2017). To this end, Gormley and Meade, 2007 proposed a cash management model based on forecasts and four bounds: a low bound D , a high bound V and two target levels v and d . The cash balance is allowed to wander around between these bounds as shown in Figure 8.2.




Figure 8.2: The dynamic simple policy of Gormley-Meade.

When V is reached a withdrawal transfer is made to restore the balance to a target level v . In the same way, when the cash balance reaches D , a positive transfer is made to restore the balance to a target d . Forecasts for a given planning horizon are used as a key input to the model that establishes a dynamic

simple policy of the form:

$$x_t = \begin{cases} v_t - \hat{b}_{t-1} - \hat{f}_t, & \text{if } \hat{b}_{t-1} + \hat{f}_t > V_t, \\ 0, & \text{otherwise,} \\ d_t - \hat{b}_{t-1} - \hat{f}_t, & \text{if } \hat{b}_{t-1} + \hat{f}_t < D_t \end{cases} \quad (8.13)$$

where D_t, d_t, v_t, V_t , for $1 \leq t \leq T$ are the decision variables satisfying $D_t \leq d_t \leq v_t \leq V_t$, and \hat{b}_t and \hat{f}_t are the predicted cash balance and cash flow at time t , respectively. This model is dynamic and the decision variables D_t, d_t, v_t, V_t must be obtained at each time step t .

The form of the policy described in (8.13) within the GM model, implies facing a non-linear problem whose solution is by no means straightforward. The authors suggested the use of evolutionary algorithms to obtain the set of bounds $\{D_t, d_t, v_t, V_t\}$ for each time step t , but this method does not guarantee the optimality of the solutions. Evolutionary algorithms are iterative procedures based on direct search (Branke et al., 2008). A population of feasible solutions is randomly generated and better solutions are found by comparing the fitness of the new population to the current population. As a result, there is no way to test if a given solution is optimal. In order to solve this limitation, we next present a linear-quadratic reformulation of the GM model allowing cash managers to obtain optimal solutions through:

- A linear program by considering only cost as a linear expression.
- A linear-quadratic program by considering a linear cost expression and a quadratic risk expression.

A third option may be considered by solving a linear program with both cost and risk linear expressions, e.g., by defining a linear measure of risk. However, since we are interested in a cost-risk optimization as described in Salas-Molina, Pla-Santamaria, and Rodriguez-Aguilar (2016), we next formalize the second option. Then, within the framework of the GM model, positive transactions x_t^+ occur when low bound D_t is reached. Thus, $z_t^+ = 1$ when $\hat{b}_{t-1} + \hat{f}_t \leq D_t$ and the amount transferred is given by $x_t^+ = d_t - \hat{b}_{t-1} - \hat{f}_t$. This can be expressed by the following linear constraints:

$$\hat{b}_{t-1} + \hat{f}_t - D_t \leq M(1 - z_t^+) \quad (8.14)$$

$$-M(1 - z_t^+) \leq x_t^+ - d_t + \hat{b}_{t-1} + \hat{f}_t \leq M(1 - z_t^+). \quad (8.15)$$

Furthermore, negative transactions x_t^- occur when high bound V_t is reached. Thus, $z_t^- = 1$ when $\hat{b}_{t-1} + \hat{f}_t \geq V_t$ and the amount transferred is given by $x_t^- = \hat{b}_{t-1} + \hat{f}_t - v_t$. This can be expressed by the following linear constraints:

$$V_t - \hat{b}_{t-1} - \hat{f}_t \leq M(1 - z_t^-) \quad (8.16)$$

$$-M(1 - z_t^-) \leq x_t^- + v_t - \hat{b}_{t-1} - \hat{f}_t \leq M(1 - z_t^-). \quad (8.17)$$

A third group of conditions must hold when the cash balance is between bounds D_t and V_t . Thus, when $z_t^+ = 0$ and $z_t^- = 0$ no transaction occurs. This can be expressed by the following linear constraints:

$$D_t - \hat{b}_{t-1} - \hat{f}_t \leq M(z_t^+ + z_t^-) \quad (8.18)$$

$$\hat{b}_{t-1} + \hat{f}_t - V_t \leq M(z_t^+ + z_t^-) \quad (8.19)$$

$$m \cdot z_t^+ \leq x_t^+ \leq M \cdot z_t^+ \quad (8.20)$$

$$m \cdot z_t^- \leq x_t^- \leq M \cdot z_t^- \quad (8.21)$$

As a result, we can reformulate the multiobjective cash management problem encoded in equations (8.3) and (8.4) to accommodate policies of the GM type as the following linear-quadratic multiobjective Gormley-Meade (MOGM) program:

$$\min \left[\frac{w_1}{C_{max}} C(X) + \frac{w_2}{R_{max}} R(X) \right] \quad (8.22)$$

subject to:

$$\hat{b}_t = \hat{b}_{t-1} + \hat{f}_t + x_t^+ - x_t^- \quad (8.23)$$

$$c(x_t) = \gamma_0^+ \cdot z_t^+ + \gamma_1^+ \cdot x_t^+ + \gamma_0^- \cdot z_t^- + \gamma_1^- \cdot x_t^- + h \cdot b_t \quad (8.24)$$

$$M(z_t^+ + z_t^-) \leq \hat{b}_{t-1} + \hat{f}_t - D_t \leq M(1 - z_t^+) \quad (8.25)$$

$$-M(1 - z_t^+) \leq x_t^+ - d_t + \hat{b}_{t-1} + \hat{f}_t \leq M(1 - z_t^+). \quad (8.26)$$

$$M(z_t^+ + z_t^-) \leq V_t - \hat{b}_{t-1} - \hat{f}_t \leq M(1 - z_t^-) \quad (8.27)$$

$$-M(1 - z_t^-) \leq x_t^- + v_t - \hat{b}_{t-1} - \hat{f}_t \leq M(1 - z_t^-) \quad (8.28)$$

$$D_t \leq d_t \leq v_t \leq V_t \quad (8.29)$$

$$z_t^+ + z_t^- \leq 1 \quad (8.30)$$

$$m \cdot z_t^+ \leq x_t^+ \leq M \cdot z_t^+ \quad (8.31)$$

$$m \cdot z_t^- \leq x_t^- \leq M \cdot z_t^- \quad (8.32)$$

$$\hat{b}_t \geq b_{min} \tag{8.33}$$

$$w_1 + w_2 = 1 \tag{8.34}$$

$$1 \leq t \leq T, t \in \mathbb{N}, T \in \mathbb{N} \tag{8.35}$$

where the decision variables are the bounds $\{D_t, d_t, v_t, V_t\}$ that determine control action $x_t \in X$ for each time step t , and $C(X)$ and $R(X)$ are linear-quadratic cost and risk functions such as the ones defined in equations (8.1) and (8.2) that depend on policy X . Note that $\hat{b}_t, x_t^+, x_t^-, z_t^+$ and z_t^- are also decision variables determining policy X as we show in the upcoming section.

Normalization factors C_{max} and R_{max} are used to avoid numerical bias to one of the goals but also for comparative purposes. In this chapter, we propose to set C_{max} and R_{max} to the expected cost and risk of a trivial policy consisting in taking no control action. This procedure allows us to straightforwardly compare the performance of policies through the objective function since values above one indicate a poorer performance than a trivial strategy.

Following the recommendations in Ben-Tal, El Ghaoui, and Nemirovski (2009) for robust optimization, we set a minimum cash balance proportional to the uncertainty introduced by forecasting errors. To this end, we assume that the cumulative forecasting error distribution for planning horizon T is known and presents standard deviation σ_e . Then, we set $b_{min} = \xi \cdot \sigma_e$, where ξ is a positive parameter. Assuming Gaussian forecasting errors, typical values for ξ are 2 or 3 leading, respectively, to an unfeasible program in only 95% or 99% of the realizations of the error distribution.

An important advantage of the the previous linear-quadratic reformulation in comparison to the original GM model is that it can be implemented and optimally solved through the modeling framework provided by mathematical programming solvers such as CPLEX or Gurobi. However, $8 \times T$ decision variables and $23 \times T$ constraints are involved. Hence, one may wonder if a simpler formulation is possible. Next, we present a novel formulation of the CMP, which we call the Boundless Model (BM), which aims to solve this problem.

8.4 Boundless multiobjective models for cash management

In this section, we introduce a novel boundless multiobjective model for cash management. First, we formally describe the model. We next highlight a crucial property of this model and we finally perform a numerical comparison exercise between the MOGM and the BM.

8.4.1 Model formulation

In the GM model described in Section 8.3, cash managers have to determine the set of bounds $\{D_t, d_t, v_t, V_t\}$ which minimizes expected transaction and holding costs for the whole planning horizon T . Indeed, the bounds are previously obtained to the deployment of the policy according to equation (8.13). However, we argue that the ultimate goal of the cash management problem is not to find the best set of bounds, but the best sequence of control actions. As a result, we here propose to change the focus from bounds to the actual policy to simplify computations. To this end, we next present a BM for multiobjective cash management with no bounds:

$$\min \left[\frac{w_1}{C_{max}} C(X) + \frac{w_2}{R_{max}} R(X) \right] \quad (8.36)$$

subject to:

$$\hat{b}_t = \hat{b}_{t-1} + \hat{f}_t + x_t^+ - x_t^- \quad (8.37)$$

$$c(x_t) = \gamma_0^+ \cdot z_t^+ + \gamma_1^+ \cdot x_t^+ + \gamma_0^- \cdot z_t^- + \gamma_1^- \cdot x_t^- + h \cdot b_t \quad (8.38)$$

$$z_t^+ + z_t^- \leq 1 \quad (8.39)$$

$$m \cdot z_t^+ \leq x_t^+ \leq M \cdot z_t^+ \quad (8.40)$$

$$m \cdot z_t^- \leq x_t^- \leq M \cdot z_t^- \quad (8.41)$$

$$\hat{b}_t \geq b_{min} \quad (8.42)$$

$$w_1 + w_2 = 1 \quad (8.43)$$

$$1 \leq t \leq T, t \in \mathbb{N}, T \in \mathbb{N} \quad (8.44)$$

Note that the set of constraints from equation (8.37) to (8.44) are only necessary to linearize the CMP as described in Section 8.2 and to achieve a certain degree of protection against forecasting errors by means of $\hat{b}_t \geq b_{min}$.

8.4.2 The equivalence theorem

The BM model is clearly a much simpler formulation than that of MOGM since only $4 \times T$ decision variables and $8 \times T$ constraints are involved. However, since both models aim to obtain the optimal policy, both formulations must return the same optimal results for the same inputs as we next demonstrate.

As a simple numerical example, consider a predicted cash balance $\hat{b}_{t-1} = 10$ and a forecast $\hat{f}_t = 1$. Let us assume that the optimal policy is $x_t^+ = 4$. Then, the optimal bound d_t is necessarily 15 since $d_t = x_t^+ + \hat{b}_{t-1} + \hat{f}_t = 4 + 10 + 1 = 15$ from equation (8.13). Indeed, the next simplified optimization problem:

$$\min_{x,d} g(x,d) \quad (8.45)$$

subject to:

$$x = d + a \quad (8.46)$$

where x and d are the decision variables and a is a constant, is equivalent to:

$$\min_x g(x, x - a) = \min_x g(x) \quad (8.47)$$

subject to no constraint. This fact leads us to formulate the following theorem.

Theorem 2. *Given a forecast \hat{f}_t for $1 \leq t \leq T$, a minimum balance b_{min} , a cost structure $\beta = \{\gamma_0^+, \gamma_0^-, \gamma_1^+, \gamma_1^-, h\}$, two normalization factors C_{max} and R_{max} , and some weights w_1 and w_2 , the BM encoded from equation (8.36) to (8.44) and the MOGM model encoded from equation (8.22) to (8.35) are equivalent in the sense that both models return the same optimal solutions for the same input parameters.*

Proof. Optimizing a set of bounds for each time step is equivalent to directly optimizing the policy since the former determines the latter and vice versa. Let us assume that $\hat{b}_{t-1} + \hat{f}_t \leq D_t$ holds. Then, finding d_t that minimizes some objective function subject to $x_t^+ = d_t - \hat{b}_{t-1} - \hat{f}_t$ from equation (8.13) is equivalent to finding x_t^+ due to the direct relationship between x_t^+ and d_t . A similar reasoning leads to the same conclusion when $\hat{b}_{t-1} + \hat{f}_t \geq V_t$ or when $D_t \leq \hat{b}_{t-1} + \hat{f}_t \leq V_t$. This fact implies that the set of constraints from (8.25) to (8.29) are redundant and can be removed from the optimization process demonstrating that both models are equivalent. \square

Furthermore, our BM provides a general framework for cash management in several aspects or dimensions of the problem. First, our BM accepts any type

cash flow process from a deterministic one (Baumol, 1952) to a pure stochastic one (Miller and Orr, 1966) and also empirical data sets of cash flows (Gormley and Meade, 2007; Salas-Molina, Pla-Santamaria, and Rodriguez-Aguilar, 2016). Our BM simplifies the Gormley and Meade (2007) cash management model by removing half of the decision variables and almost two thirds of the constraints. In addition, when no forecast is used (Penttinen, 1991), forecasts can be set to zero. Our MOGM model is also a generalization of the Miller and Orr (1966) model that can be obtained by setting forecasts to zero, and by forcing that $v_t = d_t$ and D_t, d_t, v_t and V_t to be constant for the whole planning horizon.

8.4.3 Comparing run times for the MOGM and the BM

In Section 8.4.2, we argue that the BM is an equivalent but simpler formulation to solve the CMP within a multiobjective framework. In what follows, we perform a comparison exercise to empirically confirm that the necessary run time to solve the BM is lower than that required to solve the MOGM.

Consider a typical scenario with a cash management system like the one depicted in Figure 1.2. Temporary idle cash balances can be invested in short-term marketable securities and bonds through an investment account 2 with an average return of 7.2% per annum ($h = 0.02\%$ per day). Inflows and outflows to/from account 1 are charged with fixed costs $\gamma_0^+ = \gamma_0^- = 20$ € and variable costs $\gamma_1^+ = \gamma_1^- = 0.01\%$ defining a cost structure $\beta = \{\gamma_0^+, \gamma_0^-, \gamma_1^+, \gamma_1^-, h\}$ according to current bank practices in Spain. Assume also a Gaussian cash flow process with mean 0.1 and standard deviation 1, figures in millions of euros. Assume also that by applying some forecasting technique, a hypothetical cash manager is able to produce daily forecasts with errors distributed according to $\mathcal{N}(0, 0.5)$. For precautionary purposes, we set a minimum cash balance $b_{min} = 1.5$, equivalent to three standard deviations of the error distribution.

We sample T elements from the cash flow process $\mathcal{N}(0.1, 1)$ that are used as a forecast input to both the GM and BM model. Next, we set an arbitrary initial cash balance 20% above b_{min} and we generate 100 different forecast samples of variable length T . These samples, the initial cash balance, the minimum cash balance and the cost structure β are enough to obtain the optimal policy by solving the MOGM model encoded from equation (8.22) to (8.35) and the BM model encoded from equation (8.36) to (8.44). Then, we implement both models using the Gurobi modeling framework (Gurobi Optimization, Inc, 2016) for Python and the run time results of 100 replicates for planning horizons ranging in $[5, 20]$ in steps of five days are shown in Figure 8.3. Average run

times for our BM are 38% lower with standard deviations 41% smaller than the MOGM confirming that a simpler formulation leads to remarkably lower run times to solve the CMP.

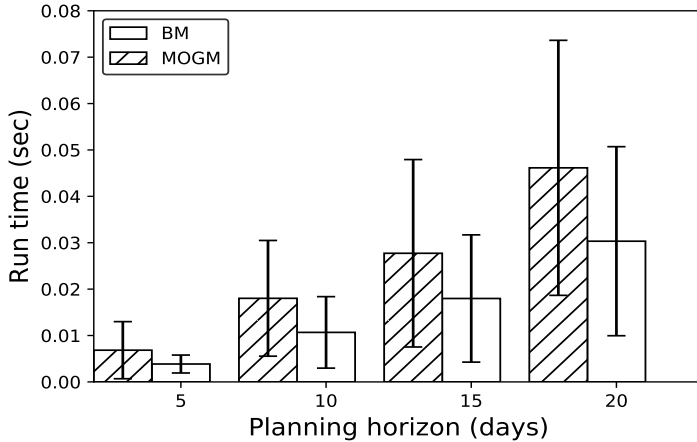


Figure 8.3: Run time for different planning horizons.

8.5 Sensitivity analysis to forecasting errors

The critical step in the deployment of a BM is the procurement of forecasts. However, forecasts are characterized by some prediction error that necessarily impacts the performance of any model using forecasts as a key input. The main goal of this section is to perform a sensitivity analysis of the goodness of the solutions provided by our BM to changes in forecasting accuracy. We first describe our assumptions and our experimental methodology, and we finally discuss the results obtained.

8.5.1 Assumptions

We experiment on 54 real cash flow data sets of variable length, containing from 170 to 1508 observations. As a statistical summary, minimum and maximum values, means and standard deviations, and number of cash flow observations available for each company are shown in Table 8.1. The data set contains daily cash flows from 54 different small and medium companies in Spain with annual revenues of up to 10 million euro each, covering a date range of about 8

years. For illustrative purposes, we consider the following representative cost structure adjusted to current bank practices in Spain and very similar to those used in Costa Moraes and Nagano (2014):

$$\beta = \{\gamma_0^+ = 20\text{€}, \gamma_0^- = 20\text{€}, \gamma_1^+ = 0.01\%, \gamma_1^- = 0.01\%, h = 0.02\%\} \quad (8.48)$$

We set a minimum balance b_{min} equivalent to three standard deviations of the empirical error distribution for a planning horizon of five days to ensure a 99% of feasibility. Finally, we follow a multiobjective approach in which expected cost is defined as the average daily cost as in equation (8.1) and expected risk is defined as the standard deviation of daily cost as in equation (8.2). Then, under a compromise programming framework, the ultimate goal is to find the policy that minimizes equation (8.3).

8.5.2 Methodology

In order to evaluate the impact of the forecasting error in actual cost-risk performance, we propose Algorithm 5. From a cash flow data set stored in vector \mathbf{f} we compute its standard deviation σ_f . Following the method described in Salas-Molina et al. (2017), we obtain forecasts of controllable accuracy by assuming error $e_t \sim \mathcal{N}(0, \sigma_e)$ with $\sigma_e = p \cdot \sigma_f$ and with error proportion $p \in [0.1, 1]$. Sample draws for different values of p in steps of 0.1 results in forecasts ranging from perfect prediction ($\sigma_e \approx 0$) to an always-zero prediction $\sigma_e = \sigma_f$. Then, by computing the actual loss of policy X , namely $\mathcal{L}_1(X)$, when the actual cash balances are not \hat{b}_t but $b_t = \hat{b}_t + e_t$, we can evaluate the impact of the forecasting error in the actual cost-risk performance of the model.

As in the comparison example described in Section 8.4.3, we set an initial cash balance b_0 to a value 20% above b_{min} as a feasible but non-conservative initial condition. Since we compare our BM to a trivial policy with the same initial condition, this setting does not interfere in the conclusions derived from the experiment. Then, given the cost structure described in equation (8.48), we generated 100 replicates with different forecast samples of length five, equivalent to more than two years of total planning horizon for each of the 54 companies. Since we consider ten error proportions, from an almost perfect cash flow prediction to an unpredictable cash flow, we cover the whole range of possible situations. As a result, we perform 54,000 experiments. Note also that, for comparative purposes, we consider that every policy is fixed, i.e., it cannot be modified during the planning horizon. In practice, the optimization process would be repeated either when new information about actual cash flows or when some error tolerances are exceeded.

Company	Min	Max	Mean	Std Dev	Length
1	-90,66	902,69	0,22	39,71	622
2	-565,06	626,55	3,09	65,35	544
3	-6.631,47	6.710,41	-7,44	414,08	935
4	-2.233,81	727,63	-0,88	170,28	893
5	-182,62	164,20	0,01	18,41	709
6	-689,70	562,69	-0,41	72,83	688
7	-300,09	829,05	0,26	65,52	555
8	-242,06	113,14	-0,14	31,89	789
9	-4.703,91	4.733,65	-3,10	658,57	754
10	-1.115,80	787,24	-1,89	83,55	788
11	-1.915,34	307,44	1,33	107,38	428
12	-615,77	7.713,80	-0,11	338,26	555
13	-1.183,62	2.274,46	-0,39	287,26	549
14	-769,28	927,11	-1,00	142,23	606
15	-551,11	556,13	0,39	114,95	696
16	-220,49	226,11	-0,50	18,25	577
17	-2.253,22	2.501,26	0,63	175,38	991
18	-287,58	263,61	-0,09	26,41	610
19	-161,73	154,82	-3,08	25,47	640
20	-150,00	160,38	-0,37	15,40	632
21	-700,00	531,66	-0,54	65,06	730
22	-2.442,94	1.388,74	-2,15	280,20	509
23	-2.898,68	2.898,68	-2,54	336,42	586
24	-3.025,05	3.178,51	-4,05	247,62	1285
25	-1.969,42	2.011,31	-0,39	174,53	600
26	-107,28	155,63	-0,05	18,64	708
27	-70,99	118,38	2,75	16,87	340
28	-324,81	390,08	-0,79	48,56	901
29	-900,41	558,88	-0,34	65,59	574
30	-188,79	198,15	-0,46	17,59	536
31	-1.344,75	349,45	-2,75	119,68	336
32	-359,16	245,04	2,71	48,77	860
33	-943,25	955,89	-1,18	78,27	670
34	-1.149,40	496,55	-1,39	108,36	1490
35	-410,71	291,91	-0,55	57,86	600
36	-78,72	118,40	4,45	18,64	357
37	-2.288,85	2.184,18	-10,16	180,89	497
38	-619,33	196,64	-11,18	67,60	193
39	-64,71	65,67	-0,11	11,66	829
40	-256,27	369,14	0,24	103,05	291
41	-626,65	643,39	-5,55	96,41	300
42	-370,21	368,46	0,47	23,11	749
43	-658,44	733,95	-0,37	131,40	832
44	-1.187,40	1.203,41	-1,83	115,28	378
45	-1.071,96	1.128,00	0,58	127,81	881
46	-511,63	738,32	10,06	75,56	411
47	-10.374,88	4.782,62	-22,94	723,62	532
48	-2.070,38	2.030,93	-5,58	255,32	581
49	-107,84	127,25	-2,07	19,96	573
50	-2.625,18	2.219,57	-2,45	351,19	374
51	-4.198,83	4.816,62	151,28	970,81	1222
52	-3.254,65	7.006,59	89,72	494,93	1220
53	-1.968,77	384,84	7,76	117,51	738
54	-10.213,56	15.321,00	9,61	1.124,10	589

Table 8.1: Data sets description. Figures in thousands of €.

Algorithm 5: Empirical evaluation algorithm

- 1 **Input:** Cash flow data set \mathbf{f} with standard deviation σ_f ; cost structure β ; minimum balance b_{min} ; planning horizon T ; error proportion p ; number of replicates r ;
 - 2 **Output:** Actual $\mathcal{L}_1(X)$ loss of optimal policy X ;
 - 3 **for** each replicate r **do**
 - 4 Set an initial balance b_0 ;
 - 5 Draw a sample of length T from \mathbf{f} as a forecast vector $\hat{\mathbf{f}}$;
 - 6 Obtain policy X that solves the BM model (8.36)-(8.44) for $\hat{\mathbf{f}}$;
 - 7 Set $\sigma_e = p \cdot \sigma_f$ and draw an error sample \mathbf{e} of length T from $\mathcal{N}(0, \sigma_e)$;
 - 8 Set actual balances to $b_t = \hat{b}_t + e_t$;
 - 9 Compute actual loss $\mathcal{L}_1(X)$ of policy X , actual balances b_t and β ;
 - 10 **end**
-

8.5.3 Results and discussion

The results from the evaluation for all 54 companies grouped by error proportion (p) are shown in Figure 8.4. We use median values instead of averages as a way to reduce the impact of extreme values on averages. The results show that actual cost-risk losses derived from the BM policy in comparison to a trivial policy are minimum when near-to-perfect predictions are possible (error proportion is 0.1). However, remarkable savings can be achieved for higher but reasonable error proportions. In practice, one can consider that error proportions below 0.5 are achievable since a high percentage of cash flows are usually known with certainty in the short-term due to payment agreements with customers, suppliers, banks and employees.

It is also worth noting that efforts in improving forecasting accuracy are highly rewarded using the BM. For example, consider a cash manager that is able to achieve forecasting errors of, at most, 40% of the standard deviation of past cash flows ($p = 0.4$), which can be considered a reasonable target. By using predictions and our BM, we observe in Figure 8.4 that more than 20% of the total loss can be saved. Going one step further, if the cost of obtaining predictions is less than that 20%, our hypothetical cash manager should deploy a BM. Summarizing, our BM proves to be significantly robust to forecasting errors since remarkable cost-risk savings with respect to a trivial policy can be obtained even for almost purely random cash flows. This fact must encourage cash managers to produce better cash flow forecasts since improvements in accuracy are highly rewarded.




Figure 8.4: Median cost-risk loss of the BM for different forecasting errors.

8.6 Summary

In this chapter, we solve one serious limitation of cash management models based on bounds: its difficulty to obtain optimal solutions. We here propose to linearize the cash management problem in an attempt to provide a sound mathematical programming framework that returns optimal policies. Once the cash management problem is formulated as a linear program with both continuous and auxiliary binary variables, the bounds that determines the policy of a BBM can be expressed as a set of additional constraints. We illustrate this approach by providing a multiobjective reformulation of the Gormley and Meade (2007) model within a compromise programming framework where cost and risk (but possibly other) are goals to optimize.

Furthermore, we propose a new BM for multiobjective cash management which departs from the state-of-the-art bound-based approach in the fact that is not constrained by any particular form of policy. Since the ultimate goal of the cash management problem is to find the best policy instead of a set of bounds that determines the policy, we show that the constraints imposed by BBMs are redundant and can be removed from the optimization problem.

Since forecasts are characterized by some prediction error, we also analyze the impact of predictive accuracy in cost-risk performance of policies derived from our BM showing that it is highly robust to forecasting errors. Note

also that the methodology used to analyze the impact of forecasting errors also supports decision-making by estimating the reward that can be obtained through the use of better forecasts. Deployment costs can then be compared to estimated savings to decide if further effort in improving predictive accuracy is worthwhile.

In this chapter, we showed that the constraints imposed by these bounds are not necessary. From the experiments performed, we found that significant benefits can be achieved for reasonably low forecasting errors. This fact must encourage cash managers to produce better cash flow forecasts since improvements in accuracy can be converted into important savings as we showed in Chapter 3. As an answer to Question 7, a new class of cash management models is available to cash managers.

Question 7. *Are control bounds really necessary in cash management?*

Contribution 7. *We proposed a new Boundless Model that provides optimal policies using forecasts as a key input without using bounds.*

Recall that the solution to the cash management problem is a policy. However, the process to obtain the best policy is not straightforward. Within a compromise programming framework, we next further elaborate on solving the multiobjective cash management problem.

Compromise programming models for cash management

In this chapter, we address the cash management problem from a multiobjective perspective by considering not only the cost but also the risk of cash policies. We pay attention to the pros and cons of alternative risk measures and its implications for solving the cash management problem. To this end, we rely on compromise programming as a method to minimize the sum of weighted distances to an ideal point where both cost and risk are minimum. Furthermore, we illustrate three different solvers for compromise programming models.

9.1 Motivation

Recall that the multiobjective cash management problem (MOCMP) can be defined as a decision making problem in which, given a set of past cash flow observations determining an initial cash balance, the goal is to find the best sequence of control actions, which is called a policy, in terms of cost and risk. We here rely on compromise programming (CP) (Zeleny, 1982; Yu, 2013; Ballesteros and Romero, 1998; Ballesteros and Pla-Santamaria, 2004) as a suitable technique to derive the best policies by minimizing weighted distances to an ideal point, where both cost and risk independently take minimum values subject to the restrictions of the problem. Under the CP framework, these weights reflect

the particular preferences of cash managers. An important advantage of CP in practical applications is the possibility to specify these preferences in a deployment phase. Then, we follow a two-step decision-making process in which we present alternatives before selection. In the context of cash management, a set of alternative policies is obtained in a learning phase and presented to cash managers. Later, a policy is selected according to their particular preferences in a deployment phase.

From the general formulation of the MOCMP introduced in Chapter 5, we consider a cash balance that starting at an initial value fluctuates according to a particular cash flow process in absence of control actions. At any time, cash managers can take a control action by increasing/decreasing the cash balance but paying a transaction cost. The resulting cash balance at the end of a particular time period is finally determined by the control action and the net cash flow occurred and it is charged with some holding cost. Since risk analysis is incorporated in the MOCMP, we pay particular attention to the pros and cons of different risk measures, such as the variance or the semivariance of daily costs. Moreover, we focus on the problem of estimating large losses as an issue of special concern for cash managers by defining novel risk measures that are able to capture the effect of large losses.

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

Consider again the formulation of the MOCMP described in Chapter 5. Given an initial cash balance b_0 , the solution to the CMP, namely, the policy X , that minimizes the sum of transaction and holding costs, up to time step n , can be obtained by solving the following optimization problem:

$$\min C(X) = \min \sum_{t=1}^n \left(\Gamma(x_t) + H(\hat{b}_t) \right) \quad (9.1)$$

subject to:

$$\hat{b}_t = \hat{b}_{t-1} + \hat{f}_t + x_t \quad (9.2)$$

$$\hat{b}_t \geq 0 \quad (9.3)$$

$$x_t \in S \quad (9.4)$$

$$t = 1, 2, \dots, n. \quad (9.5)$$

Since cash managers usually discard policies including overdrafts, we restrict the feasibility space to non-negative cash balances which is equivalent to set $u = \infty$ in equation (5.4). Set S contains all possible transactions determined by the cash management model, e.g., the cash management model proposed by Miller and Orr (1966), which is based on two bounds and a target level. Nevertheless, not only cost but also risk deserves the attention of cash managers. Cost reductions are achieved by reducing cash balances but, at the same time, the risk of an overdraft increases. As a result, given a cost structure and an initial cash balance, we aim to solve the MOCMP by finding the best policy X , that delivers the best combination in terms of cost and risk over a planning horizon of n time steps:

$$\min_X [C(X), R(X)] \quad (9.6)$$

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

9.2 Risk analysis in cash management

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

9.2.1 Measuring risk in cash management

One may hypothesize that risk is incorporated in the decision-making process of cash management by considering high penalty costs on negative cash balances. This view implies that high cost policies are also high risk policies. However, within the range of low cost policies, decision makers may prefer, for instance, policies with the lowest variability in cost due to the less uncertainty involved. In this section, we aim to answer the question: how can we measure risk

in cash management? Intuitively, risk is associated to uncertainty, danger, chance of loss or damage. It is not the damage itself but the chance of it, the possibility of occurrence. A general definition of risk in a financial context can be found in McNeil, Frey, and Embrechts (2005), who consider risk as any event or action that may adversely affect an organization's ability to achieve its objectives and execute its strategies. To some extent, managers can choose the risks that a business takes (Brealey and Myers, 2003). Quantitatively, risk is also linked to unexpected losses. For example, risk management is an important task in investment because different assets offer different degrees of risk. In the well-known mean-variance model for portfolio selection proposed by Markowitz (1952), profitability is measured by the mean of returns, and risk by the variance of returns over a given period of time in the past.

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

1. Notional-amount approach. For instance, the risk of a portfolio of assets is defined as the sum of the notional values of the individual assets of a portfolio. In this case, the higher the values the higher the risk.
2. Factor-sensitivity measures. These measures provide the change in value associated to a given change in one of the underlying risk factors. For instance, the greeks in portfolios of derivatives.
3. Scenario-based measures. In this approach, a number of future scenarios are considered, e.g., a 10% increase in the USD/EUR exchange rate. Risk is then measured by the the maximum loss produced under all scenarios considered.
4. Risk measures based on loss distributions. These measures are based on statistical quantities that describe the distribution of a random variable over a given period of time. Examples include the Value-at-Risk, the Conditional Value-at-Risk and the variance, which we here accommodate

to a cash management context. All of them summarize in a single value the risk contained in a distribution modeling loss.

9.2.2 Alternative measures of risk

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

$$F_c(c_0) = P(c(x_t) \leq c_0). \quad (9.7)$$

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

$$P(c(x_t) \geq c_0) = 1 - F_c(c_0). \quad (9.8)$$

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

Definition 16. The Cost-at-Risk (CaR) of a cash policy X at a confidence level $\alpha \in [0, 1]$ is given by the smallest number c_0 such that the probability that the cost $c(x_t)$ exceeds c_0 is no larger than $1 - \alpha$, formally:

$$CaR(X, c, \alpha) = \inf\{c_0 \in \mathbb{R} | P(c(x_t) \geq c_0) \leq 1 - \alpha\}, \quad \forall x_t \in X \quad (9.9)$$

or alternatively:

$$CaR(X, c, \alpha) = \inf\{c_0 \in \mathbb{R} | F_c(c_0) \geq \alpha\}, \quad \forall x_t \in X. \quad (9.10)$$

Notice that the CaR of policy X depends on the definition of cost function c and threshold α . Typical values for α are 0.95 or 0.99. Figure 9.1 illustrates the notion of CaR . Say that from a number of experiments, the empirical average daily cost is distributed as shown in the figure. As an example, assume also that a Weibull distribution (Weibull et al., 1951) is the function that best fits

the empirical data. If α is 0.95, we can then expect that the cost exceeds 2200 with probability 0.05.

Probably the major drawback of CaR is that it does not provide information about the severity of losses beyond c_0 . In the usual case of heavy-tailed distributions, the estimation of large losses is an important question to be considered. In the example of Figure 9.1, the last two bars could be located at points 5000 and 5500 and the CaR would remain unaltered. Significant advantages over CaR are provided by the Conditional Cost-at-Risk ($CCaR$) measure, which we define as the conditional excess expectation, similarly to the definition of Conditional Value-at-Risk in Rockafellar and Uryasev (2002).

Definition 17. We define the Conditional Cost-at-Risk ($CCaR$) value of a cash policy X at a confidence level $\alpha \in [0, 1]$ as:

$$CCaR(X, c, \alpha) = E[c(x_t) | c(x_t) > c_0], \quad \forall x_t \in X. \quad (9.11)$$

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

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

$$CCaR(X, c, \alpha) = E(\{c(x_1), \dots, c(x_r)\}) = c_0 + \frac{1}{r} \sum_{t=1}^n \max(c(x_t) - c_0, 0) \quad (9.12)$$

Alternatively, the problem of large losses is also taken into account by using variance as a risk measure. Variance has been extensively used in financial contexts due to the mean-variance portfolio selection model by Markowitz (1952). Moreover, its simplicity and ease of computation in experimental environments makes variance a good risk measure. Common criticism to the use of variance points out that it makes no distinction between positive and negative deviations. However, this drawback can be easily solved by using semivariance or upside/downside deviation as in Ballesterio (2005) and Pla-Santamaria and Bravo (2013). In the context of cash management, we propose to calculate the risk of a policy X by computing the variance of daily costs as follows:

Definition 18. The variance (V) of a cash policy X deployed over n time steps is obtained as:

$$V(X, c) = \frac{1}{n} \sum_{t=1}^n (c(x_t) - E(c))^2, \quad \forall x_t \in X. \quad (9.13)$$




Figure 9.1: Probability density function fitted to the empirical histogram of a cost random variable and its CaR_α quantile.

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

The underlying assumption on the use of variance is that the more disperse the costs within a policy around the expected cost, the higher the risk of the policy. However, since cash managers are probably more interested in upside deviations of cost rather than downside deviations, an upper partial moment such as the semivariance may be considered as an alternative measure of risk. Following McNeil, Frey, and Embrechts (2005), given an exponent k and a reference point c_0 , we here propose an additional measure of risk for a policy X as follows:

Definition 19. The k -Upper Partial Moment (UPM) with respect to c_0 of a cash policy X is obtained as:

$$UPM(X, c, k, c_0) = \int_{c_0}^{\infty} (c - c_0)^k dF_c \quad (9.14)$$

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

Note that if $k = 0$, then $UPM(X, c, 0, c_0) = P(c \geq c_0)$, is the probability that the cost exceeds the reference c_0 . Additionally, if $k = 1$, then $UPM(X, c, 1, c_0)$ is the expected upper deviation of cost from the reference c_0 . Finally, when

$k = 2$ and c_0 is set to the expected cost, then $UPM(X, c, 2, E(c))$ is the upper semivariance of cost. However, since common planning and control practices in most organizations are typically performed in discrete intervals indexed by time step t , the $UPM(X, c, k, c_0)$ can be computed in discrete time as:

$$UPM(X, c, k, c_0) = E((\max\{c(x_t) - c_0, 0\})^k) \tag{9.15}$$

9.2.3 Summary of risk measures

As a summary, the pros and cons for the aforementioned risk measures are presented in Table 9.1. When dealing with risk, cash managers are usually concerned not only with average variation but also with abnormal or extreme values (McNeil, Frey, and Embrechts, 2005; Glasserman, 2003). The risk of large losses must then be considered and, although *CaR* considers heavy tails, it does not provide information about the severity of large losses. Thus *CCaR*, variance or *UPM* can be used instead. The use of standard deviation is preferred to variance because it presents the same units as cost, i.e., money units, and numerical comparisons are then possible. However, a drawback must be pointed out against variance or standard deviation since there is no distinction between positive and negative deviations. This problem is easily solved by considering *UPM*, such as the upper semivariance. Non-linearity is another important aspect to be considered, specially when using this measure as part of an objective function to be minimized. Linear objective functions and linear constraints are usually preferred in mathematical programming. In this sense, the *CCaR* value should be considered as a good risk measure since it can be easily expressed as a linear function in an optimization problem.

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

Table 9.1: Advantages and disadvantages of alternative risk measures.

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

9.3 Compromise models to solve the MOCMP

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

Compromise programming is based on the concept of ideal point and the Zeleny's axiom of choice (Zeleny, 1974), which states that alternatives that are closer to the ideal are better than those that are further. The concept of ideal point is at the core of compromise programming. When less is better, the minimum values for each objective subject to the constraints of the problem determine the ideal point. In the context of the MOCMP, the ideal point in a bidimensional cost-risk space is the point with zero cost and zero risk that simultaneously minimizes $C(X)$ and $R(X)$ in objective function (9.6). Since this ideal point is usually unfeasible, it is necessary to look for compromise solutions by minimizing the distance to this ideal point. A general distance function between two bidimensional points $P_1 = (x_1, y_1)$ and $P_2 = (x_2, y_2)$ is the Minkowski distance of order h , defined as:

$$(|x_1 - x_2|^h + |y_1 - y_2|^h)^{1/h}. \quad (9.16)$$

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

$$\theta_1(X) = \frac{C(X) - C_{min}}{C_{max} - C_{min}} \quad (9.17)$$

$$\theta_2(X) = \frac{R(X) - R_{min}}{R_{max} - R_{min}} \quad (9.18)$$

where C_{max} (R_{max}) and C_{min} (R_{min}) are, respectively, the maximum and minimum values of cost function C (risk function R) subject to the constraints of the problem. Note that due to normalization $\theta_1, \theta_2 \in [0, 1]$. Consequently, the closer to the ideal point $(0, 0)$, the better the solution. Moreover, when considering particular goal preferences, weighted distances must be computed

instead. From that, CP proposes a family of normalized distance functions including weights that determine the decision-maker's risk preferences as:

$$\mathcal{L}_h = [w_1^h \cdot \theta_1^h + w_2^h \cdot \theta_2^h]^{1/h}. \quad (9.19)$$

Note that \mathcal{L}_1 is the Manhattan distance; \mathcal{L}_2 is the Euclidean distance, and \mathcal{L}_∞ is the Chebyshev distance. They are the most used distances in practice for interpretation and computational reasons (Ringuest, 1992; Ballester, 2007). On the other hand, weights w_1 and w_2 in equation (9.19) reflect the particular preferences of cash managers. As a result, considering Manhattan distances to avoid non-linearity, we can formulate the MOCMP as the following CP model:

$$\min [w_1 \cdot \theta_1(X) + w_2 \cdot \theta_2(X)] \quad (9.20)$$

$$X \in S. \quad (9.21)$$

Next, we consider three alternative methods to solve the MOCMP: (i) Monte Carlo methods; (ii) linear programming; and (iii) quadratic programming. We prefer Monte Carlo methods when exploring alternatives within a bounded set, when the specific cost/risk preferences are not known in the learning phase, or when there are reasonable doubts about the minimum/maximum values in (9.17) and (9.18). However, if we can express the objectives and the constraints in (9.20)-(9.21) as linear functions, and we know both weights w_1 and w_2 , and the extreme values in (9.17) and (9.18), we can automate the cash management decision-making process by solving the MOCMP without cash managers' intervention using linear programming. In addition, if any of the objective functions is quadratic (e.g., when using variance as a measure of risk), we can use quadratic programming. As a result, we can solve both linear and quadratic programs using state-of-the-art solvers such as CPLEX or Gurobi.

9.3.1 Solving the MOCMP by Monte Carlo methods

Monte Carlo methods allow for a simulation strategy presenting policies before selection according to the particular risk preferences of cash managers. In Chapter 5, we followed a very similar procedure that we retrieve now for comparative purposes. Recall that we aim to solve the multiobjective problem described in (9.20)-(9.21). Assume also that we want to deploy a policy of the Miller and Orr (1966) type based on three control bounds: a lower bound l_1 , a target level l_2 and an upper bound l_3 . Cash balances are allowed to wander around between bounds l_1 and l_3 , and when any of these bounds is reached,

a control action is taken to restore the cash balance to the target level as described in the following expression:

$$x_t = \begin{cases} l_2 - \hat{b}_{t-1}, & \text{if } \hat{b}_{t-1} > l_3 \\ 0, & \text{if } l_1 \leq \hat{b}_{t-1} \leq l_3 \\ l_2 - \hat{b}_{t-1}, & \text{if } \hat{b}_{t-1} < l_1 \end{cases} \quad (9.22)$$

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

As a result, under the framework of compromise programming for \mathcal{L}_1 in equation (9.19), solving the MOCMP for a given planning horizon of n time steps is equivalent to finding the set $\{l_1, l_2, l_3\}$, which minimizes the weighted Manhattan distance to the ideal point $(0, 0)$:

$$\min \left[w_1 \cdot \frac{C(X) - C_{min}}{C_{max} - C_{min}} + w_2 \cdot \frac{R(X) - R_{min}}{R_{max} - R_{min}} \right] \quad (9.23)$$

subject to:

$$\hat{b}_t = \hat{b}_{t-1} + \hat{f}_t + x_t \quad (9.24)$$

where $X = \{x_t : t = 1, 2, \dots, n\}$ with x_t according to equation (9.22) and bounds satisfying $0 \geq l_1 \geq l_2 \geq l_3$. We here measure cost by the average daily cost and risk by the standard deviation of daily cost as follows:

$$C(X) = E(C) = \frac{1}{n} \sum_{t=1}^n c(x_t) \quad (9.25)$$

$$R(X) = \left(\frac{1}{n} \sum_{t=1}^n (c(x_t) - E(C))^2 \right)^{1/2}. \quad (9.26)$$

An advantage of CP in practical applications is the possibility to specify these preferences in a deployment phase. Then, we follow a two-step decision-making process in which alternatives are presented before selection. Thus, since weights w_1 and w_2 are unknown at this point, we aim to obtain a Pareto efficient set of solutions (Yu, 2013). In other words, we want to derive an efficient frontier with the set of daily policies X , which are not dominated by any other policy in terms of cost and risk. 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. A suitable and simple method to obtain this efficient set is Monte Carlo simulation (Glasserman, 2003). Monte Carlo methods are

based on performing a high number of random experiments that are later evaluated in some outcome domain. The law of large numbers ensures that the estimations derived from this analysis converge to real values as the number of experiments increases. An example of a Monte Carlo method is summarized in the steps detailed in Algorithm 6.

Algorithm 6: Montecarlo method to solve the multiobjective cash management problem

- 1 **Inputs:** Model m ; set of goals indexes θ_i ; cost context β ; cash flow data set F ; time period n ; replicates r ;
 - 2 **Output:** Estimation of the context efficient set of solutions;
 - 3 **for** each replicate **do**
 - 4 Generate a random solution X ;
 - 5 **for** each goal **do**
 - 6 Compute $g_i(X)$;
 - 7 **end**
 - 8 **end**
 - 9 **for** each goal **do**
 - 10 Compute θ_i ;
 - 11 **end**
 - 12 Compute the efficient set;
-

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

$$\hat{F} = [1, 1, 6, -1, -3, -3, -9, 6, 4, 6, 3, 4, 1, -1, -2, 2]. \quad (9.27)$$

After applying Algorithm 6 with 10,000 replicates to our example, we obtain the efficient set summarized in Table 9.2. A visual representation of the efficient frontier derived from the (θ_1, θ_2) values in Table 9.2 is shown in Figure 9.2. As expected, there is a cost-risk tradeoff and lower costs can only be achieved by accepting higher risks. Thus, we can obtain a compromise solution by minimizing, for example, the Manhattan distance to the ideal point $(0, 0)$. In the next step, cash managers have to choose a policy from the efficient frontier according to their risk/cost preferences. In the case of unbiased preference for cost or risk, i.e., $w_1 = w_2$, the best solution to our example is policy 2 with control bounds 1, 11 and 20. However, a conservative cash manager may choose policy 4 in order to reduce risk but accepting a higher cost.

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

Table 9.2: Example efficient set for a Miller-Orr model with three levels.

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




Figure 9.2: Example of efficient set.

Moreover, it is proven that the utility optimum for a decision-maker lies on the normalized efficient frontier between points L and L_∞ . On the one hand, bound L is the point minimizing the linear loss function $r_0\theta_1 + \theta_2$ on the

normalized efficient frontier. On the other hand, bound L_∞ is the intersection of $\theta_1 = \theta_2$ with the efficient frontier. For instance, consider a conservative cash manager with $r_0 = 0.5$. From Table 9.2, bound L , with minimum $r_0\theta_1 + \theta_2$, is policy 3, and bound L_∞ , when cost index θ_1 approximately equals risk index θ_2 , coincides with policy 3. However, when the specific risk preferences of cash managers are known (or known to lie in a given interval), and there is no reasonable estimation doubt about the minimum/maximum values for both cost and risk, a more straightforward procedure can be followed by relying on linear or quadratic programming.

9.3.2 Solving the MOCMP by linear programming

From the set of alternative risk measures considered in Section 9.2.2, consider now $CCaR$ as a measure of risk. Recall that $CCaR$ is defined as the expected cost above a given reference c_0 . Thus, minimizing $CCaR$ is equivalent to minimizing the sum of positive cost deviations from cost reference c_0 , which can be cast as a Goal Programming (GP) model (Abdelaziz, Aouni, and El Fayedh, 2007; Aouni, Colapinto, and La Torre, 2014). GP aggregates multiple objectives to obtain the solution that minimizes the sum of deviations between achievement and the aspiration levels of the goals. Then, we proceed as follows:

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

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

$$\min \sum_{i=1}^q (w_i^+ \delta_i^+ + w_i^- \delta_i^-) \quad (9.28)$$

subject to:

$$g_i + \delta_i^- - \delta_i^+ = \tau_i \quad (9.29)$$

$$\delta_i^-, \delta_i^+ \geq 0, \quad i = 1, 2, \dots, q. \quad (9.30)$$

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

$$\delta_i^+ = \theta_i = \frac{g_i - g_{min,i}}{g_{max,i} - g_{min,i}}. \quad (9.31)$$

Let us consider again our MOCMP with two goals, namely, cost and risk, aggregated through the CP model in (9.20)-(9.21). As described in Chapter 8, note that cost function $c(x_t)$ in equation (8.1) can be linearized by rewriting control action x_t as the difference of two non-negative variables:

$$x_t = x_t^+ - x_t^-. \quad (9.32)$$

subject to the following constraints:

$$m \cdot z_t^+ \leq x_t^+ \leq M \cdot z_t^+ \quad (9.33)$$

$$m \cdot z_t^- \leq x_t^- \leq M \cdot z_t^- \quad (9.34)$$

where $M(m)$ is a very large (small) number in order to ensure that $z_t^+ = 1$ when x_t^+ occurs and that $z_t^- = 1$ when x_t^- occurs. Furthermore, we impose $z_t^+ + z_t^- \leq 1$ to avoid the simultaneous occurrence of x_t^+ and x_t^- . As a result, we can rewrite $c(x_t)$ as follows:

$$c(x_t) = \gamma_0^+ \cdot z_t^+ + \gamma_1^+ \cdot x_t^+ + \gamma_0^- \cdot z_t^- + \gamma_1^- \cdot x_t^- + v \cdot \hat{b}_t. \quad (9.35)$$

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

$$\theta_1(X) = \frac{1}{C_{max}} \sum_{t=1}^n c(x_t) = \frac{1}{C_{max}} \sum_{t=1}^n \left(\Gamma(x_t) + H(\hat{b}_t) \right) \quad (9.36)$$

$$\theta_2(X) = \frac{1}{R_{max}} \sum_{t=1}^n \delta_t^+, \quad (9.37)$$

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

$$\min \left[\frac{w_1}{C_{max}} \sum_{t=1}^n c(x_t) + \frac{w_2}{R_{max}} \sum_{t=1}^n \delta_t^+ \right] \quad (9.38)$$

subject to:

$$\hat{b}_t = \hat{b}_{t-1} + \hat{f}_t + x_t \quad (9.39)$$

$$x_t = x_t^+ - x_t^- \quad (9.40)$$

$$c(x_t) = \gamma_0^+ \cdot z_t^+ + \gamma_1^+ \cdot x_t^+ + \gamma_0^- \cdot z_t^- + \gamma_1^- \cdot x_t^- + v \cdot \hat{b}_t \quad (9.41)$$

$$z_t^+ + z_t^- \leq 1 \quad (9.42)$$

$$m \cdot z_t^+ \leq x_t^+ \leq M \cdot z_t^+ \quad (9.43)$$

$$m \cdot z_t^- \leq x_t^- \leq M \cdot z_t^- \quad (9.44)$$

$$c(x_t) - \delta_t^+ \leq c_0 \quad (9.45)$$

$$\hat{b}_t \geq b_{min} \quad (9.46)$$

$$w_1 + w_2 = 1 \quad (9.47)$$

$$\sum_{t=1}^n c(x_t) \leq C_{max} \quad (9.48)$$

$$\sum_{t=1}^n \delta_t^+ \leq R_{max} \quad (9.49)$$

$$z_t^+, z_t^- \in \{0, 1\} \quad (9.50)$$

$$x_t^+, x_t^-, \hat{b}_t, \delta_t^+ \geq 0 \quad (9.51)$$

where the main decision variables are control actions x_t^+ and x_t^- . In practice, C_{max} and R_{max} can be regarded as budget limitations for both cost and risk, leading to unfeasible policies when these constraints are not satisfied.

Since we use forecasts, cash managers may be interested to protect themselves against forecasting errors. This protection can be achieved through a minimum cash balance b_{min} . For instance, by setting a minimum cash balance equivalent to the maximum forecasting error, we transform an optimization problem affected by uncertainty into its robust counterpart as proposed by Soyster (1973) and Ben-Tal, El Ghaoui, and Nemirovski (2009). It is also important to highlight that we do not impose any additional constraint on the form of the policy $x_t \in X$, apart from non-negativity. We refer to that kind of policies as being produced by a boundless cash management model. When using alternative cash management models, a number of additional constraints must be considered to define the set of feasible policies (S).

As a numerical example, consider again the cost context β , and the set of expected cash flows (\hat{F}) for the next $n = 16$ days detailed in Section 9.3.1.

Using total cost as a measure of cost and $CCaR$ as a measure of risk, we can solve the MOCMP by minimizing (9.38), with 96 decision variables detailed as follows:

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

Assume that a cash manager is biased for cost such that $w_1 = 0.67$ and $w_2 = 0.33$. For precautionary purposes, she sets a minimum cash balance of two standard deviations of the expected cash flow ($b_{min} = 7$). The solution of this MOCMP results in the optimal cash policy and balance shown in Figure 9.2. This policy produces a total cost of 133,600 €, equivalent to an average daily cost of 8,350 €, and a total risk of 10,800 €, in terms of $CCaR$ with respect to cost reference $c_0 = 8,000$ €, representing a combined 62% of the total maximum budget constraints determined by $C_{max} = 0.15$ and $R_{max} = 0.15$, both figures in millions of euros.




Figure 9.3: Policy and balance for the example using LP.

9.3.3 Solving the MOCMP by quadratic programming

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

$$\mathbf{d} = \mathbf{c} - \frac{\mathbf{1} \cdot \mathbf{1}^T \cdot \mathbf{c}}{n} \quad (9.52)$$

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

$$\min \frac{\mathbf{d}^T \cdot \mathbf{d}}{n} \quad (9.53)$$

Similarly to Section 9.3.2, let us consider an $n \times 1$ vector of positive (negative) transactions \mathbf{x}^+ (\mathbf{x}^-) and an $n \times 1$ vector of expected balances $\hat{\mathbf{b}}$. We can then rewrite the state transition law in equation (9.39) in matrix notation as follows:

$$\hat{\mathbf{b}} = \hat{\mathbf{b}}_0 + L \cdot (\hat{\mathbf{f}} + \mathbf{x}^+ - \mathbf{x}^-) \quad (9.54)$$

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

$$\mathbf{c} = \gamma_0^+ \cdot \mathbf{z}^+ + \gamma_1^+ \cdot \mathbf{x}^+ + \gamma_0^- \cdot \mathbf{z}^- + \gamma_1^- \cdot \mathbf{x}^- + v \cdot \hat{\mathbf{b}} \quad (9.55)$$

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

$$\min \left[\frac{w_1}{C_{max}} \frac{\mathbf{1}^T \cdot \mathbf{c}}{n} + \frac{w_2}{R_{max}} \frac{\mathbf{d}^T \cdot \mathbf{d}}{n} \right] \quad (9.56)$$

subject to:

$$\hat{\mathbf{b}} = \hat{\mathbf{b}}_0 + L \cdot (\hat{\mathbf{f}} + \mathbf{x}^+ - \mathbf{x}^-) \quad (9.57)$$

$$\mathbf{c} = \gamma_0^+ \cdot \mathbf{z}^+ + \gamma_1^+ \cdot \mathbf{x}^+ + \gamma_0^- \cdot \mathbf{z}^- + \gamma_1^- \cdot \mathbf{x}^- + v \cdot \hat{\mathbf{b}} \quad (9.58)$$

$$\mathbf{d} = \mathbf{c} - \frac{\mathbf{1} \cdot \mathbf{1}^T \cdot \mathbf{c}}{n} \quad (9.59)$$

$$\mathbf{z}^+ + \mathbf{z}^- \leq \mathbf{1} \tag{9.60}$$

$$m \cdot \mathbf{z}^+ \leq \mathbf{x}^+ \leq M \cdot \mathbf{z}^+ \tag{9.61}$$

$$m \cdot \mathbf{z}^- \leq \mathbf{x}^- \leq M \cdot \mathbf{z}^- \tag{9.62}$$

$$\hat{\mathbf{b}} \geq \mathbf{b}_{min} \tag{9.63}$$

$$\frac{\mathbf{1}^T \cdot \mathbf{c}}{n} \leq C_{max} \tag{9.64}$$

$$\frac{\mathbf{d}^T \cdot \mathbf{d}}{n} \leq R_{max} \tag{9.65}$$

$$\mathbf{x}^+, \mathbf{x}^-, \hat{\mathbf{b}} \in \mathbb{R}_+^n \tag{9.66}$$

$$\mathbf{z}^+, \mathbf{z}^- \in \mathcal{B}^n \tag{9.67}$$

$$w_1 + w_2 = 1 \tag{9.68}$$

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

Following with our running example with $w_1 = 0.67$ and $w_2 = 0.33$ and $b_{min} = 7$, if we set $C_{max} = 0.15$ and $R_{max} = 10$ millions of euros, we obtain the optimal cash policy and balance shown in Figure 9.4. This policy produces a total cost of 140,150 €, equivalent to an average daily cost of 8,787 €, and a total risk in terms of variance of 3,970,029, equivalent to a standard deviation of 1,992 €.




Figure 9.4: Policy and resulting cash balance for the example using QP.

As a summary of the MOCMP solvers, some additional comments must be done on the pros and cons of the previous techniques. First, the selection of R_{max} could be tricky in the LP and QP approaches due to the difference in scale of concepts such as $CCaR$ or variance. When using variance as a measure of risk,

standard deviation is a more known concept that can be used as a proxy to set the maximum accepted risk in terms of variance by squaring the maximum standard deviation value. In addition, when cash managers require to analyze either quantitatively or graphically the impact of h in equation (9.19) as well as of risk preferences, the Monte Carlo approach is a more suitable alternative. Otherwise, a more automated decision-making procedure can be followed by solving the MOCMP using linear or quadratic programming.

9.4 Summary

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

To solve the MOCMP, we rely on a general compromise programming framework to find policies that minimize weighted distances to an ideal (but usually unfeasible) point of zero cost and zero risk. Once the cost and risk objective functions are defined, we propose three different solvers within the framework of compromise programming: (i) Monte Carlo methods; (ii) linear programming; and (iii) quadratic programming. We also make publicly-available the Python code for the three solvers used in the numerical examples ¹. This represents a starting point for practitioners interested in either designing cash management decision support systems or performing their own experiments in multiobjective cash management.

An important feature of the different techniques used to solve the cash management problem is the optimality of solutions. On the one hand, two-stage Monte Carlo methods require intervention of cash managers to choose policies and do not guarantee the optimality of solutions. On the other hand, the linear and quadratic programming counterparts of compromise programming models provide optimal solutions when risk preferences and both cost and risk maximum budgets can be reasonably estimated by cash managers. These results allowed us to answer a new open research question.

¹<https://github.com/PacoSalas/Empowering-cash-managers-CP>

Question 8. *Can we obtain optimal solutions for the multi-objective cash management problem?*

Contribution 8. *Under the framework of compromise programming, we provided linear and quadratic models and solvers of the cash management problem obtaining optimal solutions.*

Within the second part of this thesis, we have considered optimization techniques for the common two-assets setting in which idle balances from a single cash account are allocated in an single investment asset. In the following two chapters of this thesis, we extend the analysis to cash management systems with multiple bank accounts.

Part III

Managing multiple bank accounts

Robust cash management with multiple bank accounts

In this chapter, we approach the cash management problem from a multidimensional perspective. We first provide a formal specification of the cash management problem with multiple bank accounts along with theoretical results characterizing cash management systems. We also adapt our initial formulation to provide two alternative robust counterparts of the problem in order to deal with the uncertainty introduced by cash flow forecasts. In addition, we propose a third and novel two-stage cost-sensitive robust optimization formulation that considers asymmetric estimation costs. We finally illustrate its utility through a case study using real data from an industrial company.

10.1 Motivation

Cash managers usually deal with multiple banks to receive payments from customers and to send payments to suppliers, employees and other creditors. Operating such a cash management system implies a number of transactions between accounts to maintain the system in a state of equilibrium, meaning that there exists enough cash balance to face payments and avoid an overdraft. Hence, a cash management system can be viewed as a set of bank accounts and their relationships. Multidimensional schemes such as cash management systems have received little attention from the research community. Only Baccarin

(2009) proposed an impulse control approach for multiple bank accounts but restricted to continuous fluctuations of cash balances given by homogeneous diffusion processes. However, the assumption of a continuous time diffusion cash flow process and the difficulty to extend the analysis to more complex and realistic cash management systems are serious limitations.

On the other hand, the use of forecasts in the cash management problem was introduced by Stone (1972) and, more recently, Gormley and Meade (2007) and Salas-Molina et al. (2017) showed the utility of cash flow forecasts in the cash management problem. From these results, one may conclude that cash flow forecasts represent a key input in cash management. As a result, we here consider cash flow forecasts to both formalize and solve the cash management problem with multiple bank accounts. To the best of our knowledge, there is a lack of a formal specification for cash management systems with multiple bank accounts including short-term cash flow forecasts as a key input. Hence, we here propose a tool to define, analyze and solve the multiple bank accounts cash management problem (MBACMP).

Since we use cash flow forecasts in the formulation of the problem, we assume that cash managers are able to produce predictions with known accuracy resulting in a certain degree of uncertainty about the near future. The classical approach to deal with uncertainty is stochastic programming (Birge and Louveaux, 2011; Prékopa, 2013), but even in this approach constraints may be violated leading to a soft constrained problem as pointed out by Ben-Tal and Nemirovski (1999). In this chapter, we follow a robust optimization approach, which is a recent methodology to deal with optimization problems in which the data are uncertain and is only known to belong to some uncertainty set (Ben-Tal and Nemirovski, 2002). Instead of seeking a probabilistic solution, the decision-maker derives a solution that is feasible for any realization of a given uncertainty set (Bertsimas, Brown, and Caramanis, 2011).

More precisely, we use the concept of robust counterpart of an uncertain optimization problem (Ben-Tal and Nemirovski, 1999; Ben-Tal and Nemirovski, 2000; Ben-Tal, El Ghaoui, and Nemirovski, 2009) to encode the MBACMP as a linear program with hard constraints, namely, those which must be satisfied whatever the realization of the uncertainty introduced by cash flow forecasting errors. Since we use linear programming to solve the MBACMP, a further advantage of this approach is its computational tractability.

To this end, we first adapt the robust formulations of Soyster (1973) and Ben-Tal and Nemirovski (1999) to the MBACMP. However, the excessive conservatism of the former and the need for parameter tuning along with the

assumption of Gaussian forecasting errors of the latter lead us to propose a novel robust formulation. Consider the usual case of (much) higher penalty costs for negative cash balances than holding costs for positive cash balances. Since the MBACMP is a cost minimization problem, it seems reasonable to derive a robust cash policy based on actual penalty/holding costs. As a result, we here propose a novel two-stage cost-sensitive robust optimization (CSRO) encoded as two linear programs that can be efficiently solved using state-of-the-art solvers such as CPLEX or Gurobi. We illustrate and compare our CSRO approach to the formulations of Soyster and Ben-Tal by means of an empirical case study using real data from an industrial company in Spain. Our case-study results show that the CSRO approach represents a more efficient strategy in terms of cost by assuming reasonable risks.

Summarizing, we here address the MBACMP from a cost-sensitive perspective to provide robust solutions by means of linear programming. As a result, four major contributions derive from our work:

1. A formal specification of the MBACMP which is able to handle cash flow forecasts.
2. Theoretical results that generalize to a multiple bank accounts framework the basic condition for non-trivial cash policies.
3. Two robust formulations of the MBACMP based on state-of-the-art robust optimization techniques.
4. A novel cost-sensitive robust formulation of the MBACMP that achieves significant cost reductions over alternative approaches.

10.2 Formulation of the MBACMP

The purpose of this section is to provide a mathematical formulation of the MBACMP. Within a single bank account framework, consider a cash balance b_t at time t that starts with value b_0 and fluctuates according to a particular cash flow process \mathcal{F} in absence of control actions $x_t \in \mathbb{R}$. At any time $t \geq 0$, a cash manager can take a control action x_t by increasing/decreasing the cash balance from b_t to $b_t + x_t$, paying a cost defined by some transaction cost function $\Gamma(x_t)$. The resulting cash balance, given by $b_t = b_{t-1} + f_t + x_t$, where f_t is the net cash flow that occurred at time t , is charged with holding costs according to some function $H(b_t)$. Any procedure that determines the action x_t at each time step t is called a cash management model.

In practice, cash management systems with multiple bank accounts are the rule rather than the exception. As an example, consider the cash management system depicted in Figure 10.1, where circles stand for bank accounts and directed arrows between circles stand for transactions. This cash management system consists of three bank accounts, including external net cash flows f_t and four possible transactions x_t at each time step t . For instance, $f_{1,t}$ represents the external net cash flow that occurred in account 1, and $x_{1,t}$ represents the amount of the transaction from account 2 to account 1, both at time t . Note that transactions are allowed between accounts 1 and 2, and also between 1 and 3, but not between 2 and 3.




Figure 10.1: A cash management system with three bank accounts.

The state transition of any cash management system with m different bank accounts and n transactions taken at time t , can be represented through the following system of linear equations:

$$\mathbf{b}_{t-1} + \mathbf{f}_t + A \cdot \mathbf{x}_t = \mathbf{b}_t \tag{10.1}$$

where \mathbf{b}_{t-1} and \mathbf{b}_t are $m \times 1$ vectors with previous and current balances for each account, respectively; \mathbf{f}_t is an $m \times 1$ vector with external net cash flows for each account; A is an $m \times n$ incidence matrix with element a_{ij} set to: 1 if transaction j adds cash to account i , -1 if transaction j removes cash from account i , and 0 when no transaction is allowed between accounts; and finally, \mathbf{x}_t is an $n \times 1$ vector with the set of transactions (control actions) occurred at time t . This setting accepts any form of the process \mathcal{F} that generates real cash flows \mathbf{f} such as a probability density function, an empirical data set, or any other cash flow process. However, since the set of control decision of a policy are made in advance to real cash flow, predicted cash flows $\hat{\mathbf{f}}$ are used instead.

In the usual case of linear transaction costs between accounts with a fixed part γ_0 , and a variable part γ_1 , the transaction cost function $\Gamma(\mathbf{x}_t)$ at time t is defined as:

$$\Gamma(\mathbf{x}_t) = \gamma_0^T \cdot \mathbf{z}_t + \gamma_1^T \cdot \mathbf{x}_t \tag{10.2}$$

where \mathbf{z}_t is an $n \times 1$ binary vector with element z_i set to one if the i -th element of \mathbf{x}_t is not null, and zero otherwise; γ_0 is a $n \times 1$ vector of non-negative

fixed transaction costs for each transaction; and γ_1 is a $n \times 1$ vector of non-negative variable transaction costs. On the other hand, the expected holding cost function at time t is usually expressed as:

$$H(\hat{\mathbf{b}}_t) = \mathbf{v}^T \cdot \hat{\mathbf{b}}_t \quad (10.3)$$

where \mathbf{v} is an $m \times 1$ column vector with the j -th element set to the holding cost per money unit for account j , and $\hat{\mathbf{b}}_t$ is the expected cash balance at time step t .

Consider now a cash planning horizon of τ time steps, e.g., the next 5 working days. Given an initial cash distribution \mathbf{b}_0 , the solution to the MBACMP, namely, the policy $X = \langle \mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_\tau \rangle$ that minimizes the sum of transaction and holding costs, up to time step τ , can be obtained by solving the following linear program:

$$\min \sum_{t=1}^{\tau} \left(\Gamma(\mathbf{x}_t) + \mathbf{v}^T \cdot \hat{\mathbf{b}}_t \right) \quad (10.4)$$

subject to:

$$\hat{\mathbf{b}}_{t-1} + \hat{\mathbf{f}}_t + A \cdot \mathbf{x}_t = \hat{\mathbf{b}}_t \quad (10.5)$$

$$\hat{\mathbf{b}}_t \in \mathbb{R}_{\geq 0}^m \quad (10.6)$$

$$\mathbf{x}_t \in \mathbb{R}_{\geq 0}^n \quad (10.7)$$

$$t = 1, 2, \dots, \tau \quad (10.8)$$

The previous linear program is a general model, which we call the Multidimensional Boundless Model (MBM), due to the fact that, apart from non-negativity, no constraint is placed on the minimum/maximum account balances, contrary to what is usually proposed in the literature for a single bank account. In practice, $\tau \times n$ additional binary variables \mathbf{z}_t , to account for the fixed costs of transactions, and $\tau \times m$ balance auxiliary variables $\hat{\mathbf{b}}_t$, for each account at any time, are also necessary. Furthermore, additional constraints of the type $\hat{\mathbf{b}}_t \geq \mathbf{b}_{min}$, and $\hat{\mathbf{b}}_t \leq \mathbf{b}_{max}$, can be included to keep cash balances between certain levels. Summarizing, the MBACMP can be encoded as a mixed integer linear program that is computationally tractable through the use of state-of-the-art commercial solvers even for a large number of bank accounts and transactions.

10.3 Theoretical results

In this section, we characterize the particular structure of the MBACMP that produces non-trivial solutions. To this end, let us first introduce the concept of trivial policy under the context of cash management.

Definition 20. A policy $X_N = \langle \mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_\tau \rangle$ is said to be trivial when some naive method N is used to obtain each of its elements.

The notion of trivial policy is similar to that of naive forecasts within the context of fitting predictive models to existing data (see e.g. Makridakis, Wheelwright, and Hyndman (2008)). In order to assess the goodness of fit of a given model to a dataset, it is customary to compare the accuracy of the model to that obtained by using a naive method, for instance, by predicting always the mean of the data used to fit the model. Similarly, an example of a trivial policy in cash management is taking no control action.

We here aim to characterize non-trivial policies that are able to outperform a given trivial policy. Within the cash management problem for a single bank account, Constantinides and Richard (1978) pointed out the conditions for non-triviality. Next, we generalize these conditions for cash management systems with multiple bank accounts. To this end, let us consider a linear cost scenario determined by vectors γ_0 , γ_1 and \mathbf{v} , in which the following theorem holds:

Theorem 3. *Within a cash management system defined by an incidence matrix A , if the linear program of the MBACMP is feasible, then the necessary condition for non-triviality in a linear cost scenario is given by:*

$$\gamma_1^T + \mathbf{v}^T A < 0 \quad (10.9)$$

Proof. Let us consider a general policy $X = \langle \mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_\tau \rangle$ and a particular trivial policy $X_0 = \langle \mathbf{0}, \mathbf{0}, \dots, \mathbf{0} \rangle$, where $\mathbf{0}$ is an $n \times 1$ vector of zeros, equivalent to taking no control action. A policy X is non-trivial with respect to X_0 if the sum of transaction and holding costs derived from X are smaller than the holding costs derived from the trivial policy X_0 :

$$\gamma_0^T \sum_{t=1}^{\tau} z_t + \gamma_1^T \sum_{t=1}^{\tau} \mathbf{x}_t + \sum_{t=1}^{\tau} \mathbf{v}^T \cdot \hat{\mathbf{b}}_t(X) < \sum_{t=1}^{\tau} \mathbf{v}^T \cdot \hat{\mathbf{b}}_t(X_0) \quad (10.10)$$

$$\gamma_0^T \sum_{t=1}^{\tau} z_t + \gamma_1^T \sum_{t=1}^{\tau} \mathbf{x}_t < \sum_{t=1}^{\tau} \mathbf{v}^T \cdot \hat{\mathbf{b}}_t(X_0) - \sum_{t=1}^{\tau} \mathbf{v}^T \cdot \hat{\mathbf{b}}_t(X) \quad (10.11)$$

$$\gamma_0^T \sum_{t=1}^{\tau} z_t + \gamma_1^T \sum_{t=1}^{\tau} x_t < v^T \left(b_0 + \sum_{t=1}^{\tau} \hat{f}_t - b_0 - \sum_{t=1}^{\tau} \hat{f}_t - A \sum_{t=1}^{\tau} x_t \right) \quad (10.12)$$

$$\gamma_0^T \sum_{t=1}^{\tau} z_t + \gamma_1^T \sum_{t=1}^{\tau} x_t < -v^T A \sum_{t=1}^{\tau} x_t. \quad (10.13)$$

Then, since $z_t \neq 0$, when $x_t \neq 0$, it is never optimal to transfer money through x_t , unless $\gamma_1^T + v^T A < 0$. \square

In other words, the cost of transferring one money unit through any transaction must be smaller than the difference $(v_l - v_j)$ in holding costs between the source account (l) and the target account (j) with $v_l > v_j$.

Next, we discuss an important result that stems from theorem 3. Before that, we introduce an important concept regarding the structure of a cash management system.

Definition 21. Given a cash management system with incidence matrix A , we say that there is a loop between accounts j and l iff there is a pair of transactions (x_i, x_k) such that $a_{ji} = 1$, $a_{li} = -1$, $a_{jk} = -1$ and $a_{lk} = 1$.

The definition above tells us that there is a transaction x_i from l to j , and another transaction x_k from j to l . An example of such a loop is depicted in Figure 10.2. Notice that a loop between accounts j and l indicates that transactions can eventually occur in both directions, namely from j to l and from l to j . In other words, loops may eventually involve bidirectional transactions. Nonetheless, the following theorem characterizes the conditions under which such bidirectional transactions cannot occur. Furthermore, it does indicate which transaction is actually preferred.




Figure 10.2: An example of a loop between accounts.

Theorem 4. Given a cash management system with incidence matrix A that satisfies the necessary condition for non-triviality from Theorem 3, bidirectional transactions within loops cannot occur. The preferred transaction is either x_i when $\gamma_{1,i} < v_l - v_j$ or x_k when $\gamma_{1,k} < v_j - v_l$.

Proof. Given a cash management system defined by incidence matrix A , with elements a_{ij} , for any pair of transactions (x_i, x_k) bidirectionally connecting accounts (j, l) such that $a_{ji} = 1$, $a_{li} = -1$, $a_{jk} = -1$ and $a_{lk} = 1$, the condition in equation (10.9) is equivalent to the next double element-wise comparison:

$$\gamma_{1,i} + v_j - v_l < 0 \quad (10.14)$$

$$\gamma_{1,k} + v_l - v_j < 0 \quad (10.15)$$

which can only hold either when $v_l > v_j$ in (10.14), or when $v_l < v_j$ in (10.15), but never simultaneously in (10.14) and (10.15), provided that $\gamma_{1,i}, \gamma_{1,k} < |v_l - v_j|$, for any $\gamma_{1,i}, \gamma_{1,k} \geq 0$. \square

10.4 Robust formulations of the MBACMP

Recall from Section 10.2 that we use cash flow forecasts with a certain accuracy as a key input to the MBACMP. In order to deal with the uncertainty introduced by forecasting errors, we here consider two alternative robust formulations of the MBACMP that can be ultimately encoded as linear programs. More precisely, we adapt the seminal formulation of Soyster (1973) and the more recent formulation of Ben-Tal and Nemirovski (1999) and Ben-Tal and Nemirovski (2000) and Ben-Tal, El Ghaoui, and Nemirovski (2009). By choosing these two formulations, we aim to confirm (or discard) within the context of cash management the claim of Ben-Tal of excessive conservatism of the Soyster's approach.

In an attempt to connect (big) data with optimization, we represent uncertainty introduced by cash flow forecasts by means of a data set containing (not-necessarily Gaussian) forecasting errors for each bank account. To this end, we assume that a data set of length N with past cash flow forecasting errors $\mathbf{e}_j = \hat{\mathbf{f}}_j - \mathbf{f}_j$, for each account j is available. Note that \mathbf{e}_j can be equivalently defined as $\mathbf{e}_j = \hat{\mathbf{b}}_j - \mathbf{b}_j$, because the only difference between cash flows and cash balances for any bank account is its initial balance. Furthermore, since we are dealing with policies deployed over a planning horizon of τ time steps, we require to define cumulative errors of dimension τ as follows.

Definition 22. Given a vector of forecasting errors \mathbf{e}_j of length N , its τ -cumulative forecasting error vector $\bar{\mathbf{e}}_j(\tau)$ of length $N - \tau + 1$ contains at each element the rolling sum of τ consecutive errors in \mathbf{e}_j .

Next, we propose a method to compute cumulative errors. To this end, we rely on the concept of time delay embedding (Serra et al., 2012) to compute

empirical τ -cumulative errors for each account j . From each vector \mathbf{e}_j of length N , we can sample τ consecutive elements to construct a $\tau \times 1$ error vector:

$$\mathbf{e}_{j,i}(\tau) = [e_{i-\tau+1}, \dots, e_{i-1}, e_i]^T \quad (10.16)$$

where subscript i denotes the i -th element of the original sampled vector \mathbf{e}_j . For instance, $\mathbf{e}_{j,N}(\tau)$ contains the last τ elements of \mathbf{e}_j . Through vector concatenation, we next construct an $(N - \tau + 1) \times \tau$ error matrix $\mathbf{E}_j(\tau)$ for each account j , as follows:

$$\mathbf{E}_j(\tau) = [\mathbf{e}_{j,\tau}, \dots, \mathbf{e}_{j,N-1}, \mathbf{e}_{j,N}]^T. \quad (10.17)$$

Finally, we obtain an $(N - \tau + 1) \times 1$ vector $\bar{\mathbf{e}}_j(\tau)$ of τ -cumulative errors for each account j by summing row-wise elements $E_{i,k}$ of matrix $\mathbf{E}_j(\tau)$:

$$\bar{\mathbf{e}}_j(\tau) = \left[\sum_{k=1}^{\tau} E_{1,k}, \sum_{k=1}^{\tau} E_{2,k}, \dots, \sum_{k=1}^{\tau} E_{N-\tau+1,k} \right]^T \quad (10.18)$$

In what follows, we adapt the formulation of Soyster and Ben-Tal to the MBACMP. Later, we compare these formulations to our novel cost-sensitive robust approach.

10.4.1 Adapting the robust formulation of Soyster

Since we introduce forecasts in the formulation of the MBACMP, we are dealing with a linear program with uncertain data. Soyster (1973) provided a general robust formulation for linear optimization problems that we here apply to propose the robust formulation of Soyster to the MBACMP as follows:

$$\min \sum_{t=1}^{\tau} \left(\Gamma(\mathbf{x}_t) + \mathbf{v}^T \cdot \hat{\mathbf{b}}_t \right) \quad (10.19)$$

subject to (10.5)-(10.7)-(10.8) and:

$$\hat{\mathbf{b}}_t \geq \boldsymbol{\varepsilon} \quad (10.20)$$

where $\boldsymbol{\varepsilon} \in \mathbb{R}^m$ is an $m \times 1$ vector with each j -th element set to:

$$\varepsilon_j = |\min(\bar{\mathbf{e}}_j(\tau))|. \quad (10.21)$$

The operator \min computes the minimum value among the elements of a given vector. The Soyster's method provides the highest protection against forecasting errors due to the fact that the maximum forecasting error is used to

introduce a minimum balance constraint for each bank account. This method is considered extremely conservative by Ben-Tal and Nemirovski (1999) and Bertsimas and Sim (2004). However, we aim to test if this claim is correct also within the context of the MBACMP.

10.4.2 Adapting the robust formulation of Ben-Tal

Ben-Tal and Nemirovski (1999) and Ben-Tal and Nemirovski (2000) and Ben-Tal, El Ghaoui, and Nemirovski (2009) suggested the use of ellipsoidal uncertainty sets to overcome the conservatism of the Soyster's method and hence better capture the particular risk preferences of decision-makers. Within the particular framework of the MBACMP, the robust formulation of Ben-Tal is equivalent to solving the following program:

$$\min \sum_{t=1}^{\tau} \left(\Gamma(\mathbf{x}_t) + \mathbf{v}^T \cdot \hat{\mathbf{b}}_t \right) \quad (10.22)$$

subject to (10.5)-(10.7)-(10.8) and:

$$\hat{\mathbf{b}}_t \geq \boldsymbol{\varepsilon} \quad (10.23)$$

where $\boldsymbol{\varepsilon} \in \mathbb{R}^m$ is an $m \times 1$ vector with each j -th element set to:

$$\varepsilon_j = \xi \cdot \text{std}(\bar{e}_j(\tau)). \quad (10.24)$$

The operator std computes the standard deviation of the elements of a given vector, and parameter $\xi \in \mathbb{R}_+$ is a subjective value chosen by the cash manager to reflect her attitude towards risk. The larger the value of ξ , the more averse to risk she is. For instance, if we assume normally distributed forecasting errors, setting $\xi = 2$ would be approximately equivalent to ensure a positive cash balance with probability 0.95. However, the selection of ξ may be problematic when dealing with non-Gaussian forecasting errors, for instance, when non-linear techniques are used to obtain forecasts (Salas-Molina et al., 2017).

In addition, cash managers may be much more concerned about negative forecasting errors than positive ones due to the fact that penalty costs on negative cash balances are usually higher than holding costs on positive balances. In a first attempt to incorporate the preferences of cost-sensitive cash managers in the optimization process, we suggest to replace the operator standard deviation in equation (10.24) with the lower semideviation considering only negative errors. Next, we go one step further by providing a new cost-sensitive robust formulation to account for cost asymmetry in forecasting errors.

10.5 A new two-stage cost-sensitive robust formulation of the MBACMP

In this section, we approach robust optimization of the MBACMP from a cost-sensitive perspective. More precisely, we assume that cash flow forecasting errors have a different cost for a positive deviation (overestimation) than for a negative one (underestimation). We here import the definitions for total overestimation, underestimation and optimal forecasting bias presented in Hernández-Orallo (2013b). Next, we derive our cost-sensitive robust counterpart of the MBACMP by: (i) determining the optimal shift or bias to minimize an asymmetric loss function; and (ii) using this shift as a key input to allow a cost-sensitive robust optimization (CSRO).

In order to illustrate our approach, consider again the MBACMP formulation in (10.4)-(10.8) detailed in Section 10.2. Consider the common situation in which penalty costs on negative cash balances are much higher than holding costs on positive cash balances. With perfect cash balance predictions, the optimal cash balance would be zero for the whole planning horizon. However, this situation is far from being possible in real-world applications. Instead, assume that we are provided with a data set of N past cash flow observations (f_i), and forecasts (\hat{f}_i), resulting in forecasting errors $e_i = \hat{f}_i - f_i$, computed for each instance $i \in [1, N]$ in the data set. Let us further consider the following asymmetric absolute error function:

Definition 23. The asymmetric absolute error is a function $\ell : \mathbb{F} \times \mathbb{F} \rightarrow \mathbb{R}$ which compares elements in the forecasting domain and it is defined as follows:

$$\ell_\alpha(\hat{f}_i, f_i) = \begin{cases} 2\alpha(f_i - \hat{f}_i) & \text{if } (\hat{f}_i < f_i), \\ 2(1 - \alpha)(\hat{f}_i - f_i) & \text{if } (\hat{f}_i \geq f_i), \end{cases} \quad (10.25)$$

where $\alpha \in [0, 1]$ is the cost asymmetry.

For instance, by setting $\alpha = 0.66$ we consider underestimation twice as harmful as overestimation. As a result, given a prediction vector $\hat{\mathbf{f}}$, and actual observations vector \mathbf{f} for any bank account, we can determine the $N \times 1$ error vector $\mathbf{e} = \hat{\mathbf{f}} - \mathbf{f}$, and the total overestimation (O) and underestimation (U) function values as follows:

$$O(\hat{\mathbf{f}}, \mathbf{f}) \equiv \sum_i e_i \mid e_i \geq 0 \quad (10.26)$$

$$U(\hat{\mathbf{f}}, \mathbf{f}) \equiv \sum_i e_i \mid e_i < 0 \quad (10.27)$$

By incorporating both overestimation and underestimation in the asymmetric absolute error function (10.25), we are able to compute the total loss as:

$$L(\hat{\mathbf{f}}, \mathbf{f}) = \sum_i \ell_\alpha(\hat{f}_i, f_i) = 2(1 - \alpha)O - 2\alpha U. \quad (10.28)$$

Since forecasting cash flows $\hat{\mathbf{f}}_t$ is equivalent to forecasting cash balances $\hat{\mathbf{b}}_t$, given a predicted cash balance vector $\hat{\mathbf{b}}$ for any bank account, we can compute a shifted prediction vector $\hat{\mathbf{b}}(s)$ by adding the same shift $s \in \mathbb{R}$ to each element within the vector:

$$\hat{b}_j(s) \leftarrow \hat{b}_j + s_j. \quad (10.29)$$

Different shifts lead to different values of overestimation O_s and underestimation U_s . Consequently, the optimal shift can be obtained by minimizing the total loss:

$$\min_s L(\hat{\mathbf{b}}, \mathbf{b}) = \min_s \sum_i \ell_\alpha(\hat{b}_j, b_j) = \min_s (2(1 - \alpha)O_s - 2\alpha U_s) \quad (10.30)$$

Non-linearity in the optimization problem represented by equation (10.30) is an important limitation. Next, we use goal programming (Aouni, Colapinto, and La Torre, 2014) as a method to derive an equivalent linear program. Goal programming aggregates multiple objectives to obtain the solution that minimizes the sum of deviations between achievement and the aspiration levels of the goals. Since overestimation in (10.26) and underestimation in (10.27) are defined as a sum of deviations (errors), we here rely on goal programming to reformulate the problem (10.30) as a linear program by means of the following multiple criteria weighted goal programming model to be solved for each account j :

$$\min \left[2(1 - \alpha) \sum_{i=1}^N \delta_{ij}^+ + 2\alpha \sum_{i=1}^N \delta_{ij}^- \right] \quad (10.31)$$

subject to:

$$\bar{e}_{ij} + s_j + \delta_{ij}^- - \delta_{ij}^+ = 0 \quad (10.32)$$

$$\delta_{ij}^+, \delta_{ij}^- \in \mathbb{R}_+ \quad (10.33)$$

$$i = 1, \dots, N - \tau + 1 \quad (10.34)$$

where $s_j \in \mathbb{R}$ is the only decision variable, \bar{e}_{ij} is the i -th element of τ -cumulative error vector $\bar{\mathbf{e}}_j(\tau)$, and $\delta_{ij}^+(\delta_{ij}^-)$ is the overestimation (underestimation) for each

instance i in vector $\bar{e}_j(\tau)$ for account j . By transforming the non-linear program in (10.30) into the linear program (10.31)-(10.34) we are able to use commercial linear programming solvers such as CPLEX or Gurobi to obtain the optimal shift $s_j^*(\tau)$, which ultimately depends on the planning horizon τ since we consider cumulative forecasting errors.

This shift avoids the problem of: (i) selecting a conservative precautionary cash balance based on the maximum forecasting error, as it was the case in the Soyster's robust formulation described in Section 10.4.1; and (ii) selecting a conservatism parameter ξ , as in the robust formulation of Ben-Tal from Section 10.4.2. Instead, the selection is performed through the asymmetry of underestimation and overestimation costs, allowing a cost-sensitive selection based, for example, on the current holding and penalty costs charged by banks to cash managers.

Summarizing, we propose a two-stage cost-sensitive robust optimization of the MBACMP that comprises the following steps:

1. Given an asymmetry error factor α derived from the current penalty/holding costs and a data set with forecasting errors for each bank account j , determine the optimal shift $s_j^*(\tau)$ by minimizing an asymmetric loss function.
2. Set a minimum cash balance constraint $\hat{\mathbf{b}}_t \geq \mathbf{s}^*(\tau)$, where each j -th element of $\mathbf{s}^*(\tau)$ is set to $s_j^*(\tau)$, and solve the following linear program:

$$\min \sum_{t=1}^{\tau} \left(\Gamma(\mathbf{x}_t) + \mathbf{v}^T \cdot \hat{\mathbf{b}}_t \right) \tag{10.35}$$

subject to:

$$\hat{\mathbf{b}}_{t-1} + \hat{\mathbf{f}}_t + A \cdot \mathbf{x}_t = \hat{\mathbf{b}}_t \tag{10.36}$$

$$\hat{\mathbf{b}}_t \geq \mathbf{s}^*(\tau) \tag{10.37}$$

$$\mathbf{x}_t \in \mathbb{R}_{\geq 0}^n \tag{10.38}$$

$$t = 1, 2, \dots, \tau. \tag{10.39}$$

10.5.1 An illustrative example

Consider a simple instance of the single objective program (10.4)-(10.8), for a system of two bank accounts such as the one shown in Figure 10.3. Collections from customers and payments to suppliers are managed only through account 1 ($\hat{f}_{2,t} = 0$) and there is no transaction cost between accounts ($\Gamma(\mathbf{x}_t) = 0$). Assume also that the holding cost for the second bank account is zero. The resulting optimization problem determines the best cash balance in account 1 by minimizing:

$$\min \sum_{t=1}^{\tau} v_1 \cdot \hat{b}_{1,t} \tag{10.40}$$

subject to:

$$\hat{b}_{1,t-1} + \hat{f}_{1,t} + x_{1,t} - x_{2,t} = \hat{b}_{1,t} \tag{10.41}$$

$$\hat{b}_{2,t-1} - x_{1,t} + x_{2,t} = \hat{b}_{2,t} \tag{10.42}$$

$$\hat{b}_{1,t} \geq 0 \tag{10.43}$$

$$t = 1, 2, \dots, \tau. \tag{10.44}$$

If we had perfect cash balance predictions, the optimal solution to the previous program with holding cost $v_1 > 0$ would be $\hat{b}_{1,t} = 0$ for all t . However, perfect predictions are not available in real world applications. Assume also that the current cost structure charges much higher penalty costs on negative cash balances than holding costs on positive cash balances, e.g., 10% versus 2%. Thus, we set $\alpha = 10/(10 + 2)$ as an asymmetric error factor for equation (10.25).




Figure 10.3: A cash management system with two bank accounts.

Assume now that we are provided with a data set of past cash balance predictions as shown in Table 10.1, with maximum absolute error of 1 and error's standard deviation of 0.5. Both Soyster's and Ben-Tal's formulation with $\xi = 2$ lead to replace constraint (10.43) with:

$$\hat{b}_{1,t} \geq 1 \tag{10.45}$$

leading to an optimal cash balance solution of $\hat{b}_{1,t} = 1$ million euros, for all t . Alternatively, by incorporating our cost-sensitive robust optimization approach, we first obtain an optimal shift to produce biased forecasts as follows:

$$\hat{b}_{1,t}(s) \leftarrow \hat{b}_{1,t} + s \quad (10.46)$$

by solving the linear program represented by equations (10.31) to (10.34). In our example, the optimal shift s^* is 0.5, as shown in Figure 10.4. Thus, the cost-sensitive robust counterpart of the program represented by equations (10.40) to (10.44) is obtained by replacing constraint (10.43) with:

$$\hat{b}_t \geq 0.5 \quad (10.47)$$

leading to a less conservative and more efficient solution in terms of cost by keeping a constant cash balance of 0.5 million euros.

Time	1	2	3	4	5	6	7	8	9	10
Forecast (\hat{b}_t)	6	3,5	7	6	8	9,6	5,5	7,2	4	6,8
Actual (b_t)	7	4	6,5	6,2	7,8	8,8	6	7	3,8	6,5
Error (ε)	-1	-0,5	0,5	-0,2	0,2	0,8	-0,5	0,2	0,2	0,3
Overestimation (O_t)	0	0	0,5	0	0,2	0,8	0	0,2	0,2	0,3
Underestimation (U_t)	-1,0	-0,5	0,0	-0,2	0,0	0,0	-0,5	0,0	0,0	0,0

Table 10.1: Example error data set in millions of €.

10.6 Empirical case study

In this section, by means of an empirical case study using real data from an industrial company in Spain, we show that our cost-sensitive robust formulation of the MBACMP can robustly deal with cash management systems with multiple bank accounts. More precisely, we compare the robust formulations of Soyster described in Section 10.4.1 and Ben-Tal described in Section 10.4.2 to our two-stage cost-sensitive approach from Section 10.5.

10.6.1 Assumptions

Consider a cash management system with two bank accounts and an investment account like the one in Figure 10.5. A hypothetical cash manager receives payments from customers and manages payments to suppliers through bank accounts 1 and 2. Daily net cash flows are summarized in variables $f_{1,t}$ and




Figure 10.4: Total loss due to estimation errors.

$f_{2,t}$. Temporary idle cash balances can be invested in short-term marketable securities and bonds through an investment account 3 with an average return of 3.6% per annum, equivalent to 0.01% per day. This is equivalent to set a holding cost 0.01% per day for both accounts 1 and 2. Transactions are allowed between all three accounts and charged with fixed (γ_0) and variable (γ_1) costs determining the cost structure (β) detailed in Table 10.2.

In this case study, we are provided with a data set of both past cash flow observations for two different bank accounts, with 1249 daily observations for account 1 and 1468 observations for account 2. In order to extract (if any) non-linear linear patterns from time-series (Kantz and Schreiber, 2004; Clements, Franses, and Swanson, 2004), we here use random forests to obtain forecasts. By selecting a non-linear forecasting technique, we also relax the assumption of Gaussian forecasting errors. Random forests (Ho, 1995; Ho, 1998; Breiman, 2001) have been recently used to forecast finance and economic time-series (Conejo et al., 2005), including cash flow data (Salas-Molina et al., 2017). Since we are not provided with any other explanatory variable, we use date-related variables such as day-of-week and day-of-month to account for time dependent patterns. From the set of forecasts obtained for each account, we first compute forecasting error vectors e_1 and e_2 . We also assume that our




Figure 10.5: A cash management system with two bank accounts and an investment account.

cash manager is interested determining robust optimal policies for a planning horizon $\tau = 5$ days. As a result, following the time delay embedding procedure described in Section 10.4.1, we obtain two 5-cumulative empirical error vectors \bar{e}_1 and \bar{e}_2 .

Transaction	γ_0 (€)	γ_1 (%)	Account	v (%)
1	50	0	1	0.01
2	50	0	2	0.01
3	100	0.01	3	0
4	50	0.001		
5	100	0.01		
6	50	0.001		

Table 10.2: Cost structure data (β) for our case study.

Next, from the alternative robust formulations of Soyster (Section 10.4.1), Ben-Tal (Section 10.4.2), and our cost-sensitive robust optimization approach (Section 10.5), we derive the set of minimum cash balance constraints shown in Table 10.3. No parameter selection is required using the formulation of Soyster. On the other hand, we use a typical value $\xi = 2$ to approximately cover a 95% of the empirical error variability. Finally, we set $\alpha = 0.95$ to account for the likely situation of much higher penalty costs (on negative cash balances) than holding costs (on non-negative cash balances).

In the cash management literature for a single bank account, the common two-assets setting (cash and alternative investment) is extensively used (Miller and Orr, 1966; Gormley and Meade, 2007). It is also assumed that the alternative investment account acts as an infinite source of funds. Although such an

Formulation	Soyster	Ben-Tal ($\xi = 2$)	CSRO ($\alpha = 0.95$)
ε_j	$ \min(\bar{e}_j(\tau)) $	$\xi \cdot \text{std}(\bar{e}_j(\tau))$	$s^*(\bar{e}_j(\tau))$
$b_{min,1}$	606,376	357,921	247,705
$b_{min,2}$	784,101	318,318	281,874
$b_{min,3}$	0	0	0

Table 10.3: Minimum balances in euros used the case study.

assumption may seem quite strong at first glance, it is actually realistic since cash managers can obtain, at least in theory, unlimited funds in financial markets at a given transaction cost (see Table 10.2). Thus, we set the initial cash balance for investment account 3 to a large value with respect to the rest of accounts. Furthermore, we set an initial state for accounts 1 and 2 to a small percentage (20%) above minimum values from Table 10.3. Note that this arbitrary initial state selection has no influence in the comparative experiments since the optimal policy immediately adjusts balances to the optimal values.

10.6.2 Empirical comparison

In order to compare the three robust formulations, we propose Algorithm 7. From a cash management system defined by incidence matrix A and given a cost structure β , we compute the total expected and actual cost for the optimal policy obtained by solving each of the three robust formulations.

The total expected cost is the value of the objective function (10.4) when deploying the optimal policy of the formulation using forecasts. The actual cost is the the value of the objective function (10.4) when deploying the optimal policy of the robust formulation but evaluated over a set of the corresponding actual values of cash flows instead of forecasts. For comparative purposes, we consider that the policy is fixed, i.e., it cannot be modified during the planning horizon. As a result, by comparing expected and actual costs we aim to assess the impact of forecasting accuracy on costs. If any of the policies evaluated using actual values leads to negative balances due to forecasting errors, we consider this policy infeasible and it is not included in overall computations.

In order to evaluate the three robust formulations, we generate 1000 replicates of the MBACMP for a planning horizon of $\tau = 5$ working days, equivalent to approximately 20 years. We consider five performance indicators: (i) expected average cost; (ii) actual average cost; (iii) standard deviation of expected cost; (iv) standard deviation of actual cost; and (v) number of infeasible policies with

Algorithm 7: Empirical comparison algorithm

-
- 1 **Input:** Cash management system A ; cost structure β ; vector of minimum balances \mathbf{b}_{min} ; vector of N forecasts $\hat{\mathbf{f}}$; planning horizon τ ; vector of N actual values \mathbf{f} ; number of replicates r ;
 - 2 **Output:** Total expected and actual cost of optimal policy X ;
 - 3 **for** each replicate r **do**
 - 4 Set a initial balance \mathbf{b}_0 ;
 - 5 Draw a sample $\hat{\mathbf{f}}(\tau)$ of length τ from $\hat{\mathbf{f}}$;
 - 6 Set $\mathbf{f}(\tau)$ to the corresponding actual values of $\hat{\mathbf{f}}(\tau)$;
 - 7 Obtain policy X that minimizes equation (10.4) s.t. $\hat{\mathbf{f}}$, A , \mathbf{b}_{min} and β ;
 - 8 Compute expected cost of policy X for $\hat{\mathbf{f}}(\tau)$ and β ;
 - 9 Compute actual cost of policy X for $\mathbf{f}(\tau)$ and β ;
 - 10 **end**
 - 11 Return expected and actual cost of optimal policy X ;
-

infinite cost due to forecasting errors. The results of this empirical comparison are summarized in Table 10.4.

Formulation	Soyster	Ben-Tal ($\xi = 2$)	CSRO ($\alpha = 0.95$)
Expected average cost	839	470	391
Actual average cost	844	483	408
Std. dev. expected cost	23	34	39
Std. dev. actual cost	84	86	87
Infeasible policies	0.1%	4.5%	7.6%

Table 10.4: Results of the empirical case study. Cost figures in euros.

10.6.3 Analysis

The results presented in Table 10.4 are consistent with the common criticism of ultraconservatism of the Soyster formulation. The higher the minimum balances the less risky the policy in terms of expected standard deviation of expected cost and number of infeasible policies. However, a more efficient strategy in terms of cost can be followed by assuming reasonable risks. In this sense, Ben-Tal's formulation allows a considerable reduction of expected and actual cost of around 43%. Moreover, our novel CSRO approach produces an even higher cost reduction of around 52%. However, these cost reductions are only

possible by assuming a higher level of risk. We here consider three measures of risk, the standard deviation for both expected and actual cost and the number of infeasible policies. The expected deviation in cost notably increases in the Ben-Tal and CSRO formulations with respect to Soyster. However, the actual deviation in cost remains almost unaltered. In terms of percentage of infeasible policies, only 1 out of 1000 replicates resulted in negative cash balances using the Soyster formulation. In the case of Ben-Tal, the 4.5% result is close to the expected value of 5% derived from the setting $\xi = 2$ when assuming normally distributed forecasting errors. Finally, a higher risk is taken in the case of the CSRO. However, it is important to highlight that in our experiments we impose a fixed policy over the planning horizon. In practice, this assumption can be relaxed by monitoring errors on a daily basis and reformulating the policy by solving a new MBACMP if a given threshold is reached.

10.7 Summary

Most cash management models in the literature try to solve the cash management problem for a single bank account. In this chapter, we present a formal specification for cash management systems with multiple bank accounts that is also able to handle short-term cash flow forecasts. Our formulation provides a flexible framework to define cash management systems with variable complexity according to the particular needs of cash managers. We provide further insight by means of Theorem 3, which shows the necessary conditions to avoid trivial solutions to the problem, and Theorem 4, which shows that loops cannot occur within cash management systems in absence of a stronger constraint.

We consider cash flow forecasts as a factor of central importance in cash management. Therefore, forecasts become a key input to the MBACMP and their inherent uncertainty has to be appropriately handled to provide robust solutions. In this sense, we consider three alternative robust formulations of the MBACMP. First, we adapt the formulations of Soyster and Ben-Tal through the use of cumulative forecasting errors as a data-driven approach to obtain robust policies. Furthermore, we provide a novel cost-sensitive robust formulation of the MBACMP that allows to: (i) consider less conservative solutions; (ii) avoid the problem of selecting a conservatism parameter; and (iii) select minimum cash balance constraints through a cost-sensitive linear optimization problem. The results from the empirical case study comparing the three robust formulations (Soyster, Ben-Tal and CSRO) suggest that important cost reductions can be achieved through a less conservative strategy.

Summarizing, the optimization framework described in this chapter contributes to: (i) providing a general mathematical programming framework that allows to deal with cash management systems with multiple accounts; (ii) handling the inherent uncertainty of cash flow forecasts in a cost-sensitive manner such that cash managers can derive a robust counterpart of their particular optimization problem according to their risk preferences. Cash management systems with multiple bank accounts are the rule rather than the exception. As a result, an important contribution of this thesis is connected to the next research question.

Question 9. *Can we derive optimal policies for cash management systems with multiple bank accounts?*

Contribution 9. *After formalizing the multiple bank accounts cash management problem, we provided useful theoretical results on cash management systems and we also proposed a novel data-driven procedure to derive optimal policies.*

Selecting the best policy in cash management systems with multiple bank accounts can be facilitated with the help of specific software tools. In the next chapter, we provide a software library containing a tool to derive optimal policies for multidimensional cash management systems.

PyCaMa: Python for cash management

Selecting the best policy to keep the balance between what a company holds in cash and what is placed in alternative investments in cash management systems with multiple bank accounts is by no means straightforward. We here introduce PyCaMa, a Python module for multiobjective cash management based on linear programming that allows to derive optimal policies for cash management with multiple bank accounts in terms of both cost and risk of policies.

11.1 Motivation

Cash managers usually deal with multiple banks to receive payments from customers and to send payments to suppliers. Operating such a cash management system implies a number of transactions between accounts, what is called a policy, to maintain the system in a safe state, meaning that there exists enough cash balance to face payments and avoid an overdraft. In addition, optimal policies allow to keep the sum of both transaction and holding costs at a minimum. However, cash managers may be interested not only in cost but also in the risk of policies. Hence, risk analysis can also be incorporated as an additional goal to be minimized in cash management. As a result, deriving optimal policies in terms of both cost and risk within systems with multiple

bank accounts is not an easy task. To this end, we here introduce PyCaMa, a software tool to provide such optimal policies.

Despite the recent advances in cash management (Costa Moraes, Nagano, and Sobreiro, 2015), there is a lack of supporting software to aid the transition from theory to practice. In order to fill this gap, we provide a cash management module in Python for practitioners interested in building decision support systems for cash management. In addition, this software allows to tackle open research questions such as: (i) managing multiple bank accounts (Baccarin, 2009); (ii) the impact of cash flow forecasting accuracy in the cost of policies (Gormley and Meade, 2007; Salas-Molina et al., 2017); (iii) the utility of different risk measures in cash management (Salas-Molina, Pla-Santamaria, and Rodriguez-Aguilar, 2016); and (iv) robust optimization (Soyster, 1973; Ben-Tal and Nemirovski, 2002).

In practice, cash management systems can be represented as a set of bank accounts and a set of transactions between them. These systems can be introduced in PyCaMa by means of an incidence matrix establishing the relationship between allowed transactions and bank accounts. Once a cash management system is defined, cash managers should describe the current cost structure including fixed and variable costs for each transaction, and holding costs for each bank account. If available, PyCaMa also accepts cash flow forecasts to reduce the uncertainty about the future (Stone, 1972; Gormley and Meade, 2007). In addition, minimum balances for each account can be set for precautionary purposes. Finally, PyCaMa provides optimal policies for a given planning horizon by solving a linear program using a state-of-the-art mathematical programming solver such as Gurobi (Gurobi Optimization, Inc, 2016).

Summarizing, PyCaMa is a Python-Gurobi tool aimed to automate multiobjective decision-making in cash management. To the best of our knowledge, PyCaMa is the first software tool to solve the multiobjective cash management problem with multiple bank accounts. PyCaMa contributes to support cash management decision-making: (i) by empowering cash managers to derive optimal cash policies within a real-world context in which cash management systems with multiple bank accounts are the rule rather than the exception; and (ii) by providing a computational finance framework that can be used either as a tool for empirical research or as a benchmarking for further research in cash management. Next, we describe the optimization problem that PyCaMa solves.

11.2 Problem formulation

In order to formulate the problem, we first define a cash management system as a set of bank accounts and their relationship such as the one depicted in Figure 11.1. Any cash management system with m bank accounts and n allowed transactions can be represented by an $m \times n$ incidence matrix A , with element $a_{ij} = 1$ if transaction j adds cash to account i , $a_{ij} = -1$ if transaction j removes cash from account i , and $a_{ij} = 0$ when no transaction is allowed between accounts. In the usual case of linear transaction costs between accounts with a fixed part γ_0 , and a variable part γ_1 , the transaction cost function $\Gamma(\mathbf{x}_t)$ at time t is defined as:

$$\Gamma(\mathbf{x}_t) = \boldsymbol{\gamma}_0^T \cdot \mathbf{z}_t + \boldsymbol{\gamma}_1^T \cdot \mathbf{x}_t \quad (11.1)$$

where \mathbf{z}_t is an $n \times 1$ binary vector with element z_i set to one if the i -th element of \mathbf{x}_t is not null, and zero otherwise; $\boldsymbol{\gamma}_0$ is a $n \times 1$ vector of fixed transaction costs for each transaction; and $\boldsymbol{\gamma}_1$ is a $n \times 1$ vector of variable transaction costs. On the other hand, the expected holding cost function at time t is usually expressed as:

$$H(\hat{\mathbf{b}}_t) = \mathbf{v}^T \cdot \hat{\mathbf{b}}_t \quad (11.2)$$

where \mathbf{v} is an $m \times 1$ column vector with the i -th element set to the holding cost per money unit for account i .




Figure 11.1: A cash management system with three accounts.

11.2.1 Single objective optimization: cost

Consider now a cash planning horizon of τ time steps, e.g., the next 5 working days. Within a single cost objective, given an initial cash balance \mathbf{b}_0 , the solution to the problem, namely, the policy $X = \langle \mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_\tau \rangle$ that minimizes the sum of transaction and holding costs, up to the time step τ , can be obtained by solving the following linear program:

$$\min \sum_{t=1}^{\tau} c(\mathbf{x}_t) = \sum_{t=1}^{\tau} \left(\Gamma(\mathbf{x}_t) + \mathbf{v}^T \cdot \hat{\mathbf{b}}_t \right) \quad (11.3)$$

subject to:

$$\hat{\mathbf{b}}_{t-1} + \hat{\mathbf{f}}_t + A \cdot \mathbf{x}_t = \hat{\mathbf{b}}_t \quad (11.4)$$

$$\hat{\mathbf{b}}_t \geq \hat{\mathbf{b}}_{min} \quad (11.5)$$

$$\mathbf{x}_t \in \mathbb{R}_{\geq 0}^n \quad (11.6)$$

$$t = 1, 2, \dots, \tau \quad (11.7)$$

where $\hat{\mathbf{b}}_{t-1}$ and $\hat{\mathbf{b}}_t$ are $m \times 1$ vectors with previous and current balances for each account, respectively; $\hat{\mathbf{f}}_t$ is an $m \times 1$ vector with expected net cash flows for each account; and finally, \mathbf{x}_t is an $n \times 1$ vector with the set of transactions (control actions) occurred at time t ; and $\hat{\mathbf{b}}_{min}$ is a $m \times 1$ vector of minimum cash balances.

11.2.2 Multiobjective optimization: cost and risk

However, cash managers may also be interested in the risk of cash management policies (Salas-Molina, Pla-Santamaria, and Rodriguez-Aguilar, 2016). Similarly to the definition of Conditional Value-at-Risk in (Rockafellar and Uryasev, 2002), we consider the Conditional Cost-at-Risk ($CCaR$) measure of policy X , which we define as the conditional excess expectation above a particular cost reference c_0 as follows:

$$CCaR(X, c) = E[c(\mathbf{x}_t) | c(\mathbf{x}_t) > c_0], \quad \forall \mathbf{x}_t \in X. \quad (11.8)$$

An additional advantage of $CCaR$ is that it is a coherent measure of risk in the sense of Artzner et al. (1999). Minimizing $CCaR$ is equivalent to minimize the sum of the positive deviation of cost above reference c_0 . As a result, we next incorporate $CCaR$ as an additional goal to be optimized through the following multiobjective linear program in which cost-risk preferences are expressed by means of weights w_1 and w_2 :

$$\min \left[\frac{w_1}{C_{max}} \sum_{t=1}^n c(\mathbf{x}_t) + \frac{w_2}{R_{max}} \sum_{t=1}^n \delta_t^+ \right] \quad (11.9)$$

subject to:

$$\hat{\mathbf{b}}_{t-1} + \hat{\mathbf{f}}_t + A \cdot \mathbf{x}_t = \hat{\mathbf{b}}_t \quad (11.10)$$

$$\hat{\mathbf{b}}_t \geq \hat{\mathbf{b}}_{min} \quad (11.11)$$

$$c(\mathbf{x}_t) - \delta_t^+ \leq c_0 \quad (11.12)$$

$$\sum_{t=1}^n c(\mathbf{x}_t) \leq C_{max} \quad (11.13)$$

$$\sum_{t=1}^n \delta_t^+ \leq R_{max} \quad (11.14)$$

$$\mathbf{x}_t \in \mathbb{R}_{\geq 0}^n \quad (11.15)$$

$$w_1 + w_2 = 1 \quad (11.16)$$

$$t = 1, 2, \dots, \tau \quad (11.17)$$

where δ_t^+ is an auxiliary variable used to measure deviations from a cost reference as in goal programming (Aouni, Colapinto, and La Torre, 2014). C_{max} and R_{max} can be regarded as budget limitations for both cost and risk, leading to unfeasible policies when these constraints are not satisfied.

11.3 Software description

Since the cash management problem is an optimization problem, PyCaMa is based on linear programming to provide optimal policies. However, it is important to highlight that the linear programs described in Section 11.2.1 (only for cost) and in Section 11.2.2 (for both cost and risk) should be considered as baseline models that can be used for benchmarking purposes.

A cash management system, a cost structure and a set of minimum balances are sufficient to create an instance of the *multibank* class. It is assumed that cash managers are able to produce cash flow forecasts for each bank account as an additional input to the problem. Otherwise, forecasts must be set to zero. Next, given an initial condition and a set of cash flow forecasts for the immediate future, cash managers can derive optimal policies either in terms of only cost or in terms of both cost and risk.

11.3.1 Software architecture

PyCaMa is organized around the *multibank* class. An instance of this class is initialized by means of different data structures in Python such as: (i) a list of m banks; (ii) a list of n transactions; (iii) an $m \times n$ incidence matrix A , establishing the cash management system; (iv) two dictionaries linking transactions and both fixed (γ_0) and variable (γ_1) transaction costs; (v) a dictionary linking each bank account to holding costs in vector \mathbf{v} ; (vi) a list with m minimum cash balances in $\hat{\mathbf{b}}_{min}$; as follows:

```
1 myproblem = multibank(banks, trans, A, g0, g1, v, b_min)
```

Once a *multibank* object is created, a number of methods are implemented to provide cash managers with useful functionalities that we next describe.

11.3.2 Software functionalities

Cash managers can retrieve the main characteristics of the cash management system they are dealing with by using function *describe*. All the input data is then shown for descriptive purposes. The main functionality of PyCaMa is function *solvecost*, which provides a solution (if any) for the linear program encoded by equations (11.3)-(11.7). Given a list of length m with an initial condition \mathbf{b}_0 , and a $\tau \times m$ matrix F of forecasts (with elements set to zero if not available) obtained by concatenating vectors $\hat{\mathbf{f}}_t$ with t ranging in $1, 2, \dots, \tau$, the optimal policy is obtained by executing:

```
1 solution_1 = myproblem.solvecost(b0, F)
```

If the linear program has a feasible solution, function *solvecost* returns its optimal policy for each transaction and time step. Otherwise, *solvecost* warns the user that it was unable to find a solution. In addition, a $\tau \times n$ matrix with the optimal policy, and a $\tau \times m$ matrix with optimal balances derived from the last optimization can be retrieved by means of functions *policy* and *balance*, respectively, and the last objective value by calling the attribute *objval* of the *multibank* object.

Furthermore, cash managers interested in minimizing not only cost but also the risk of policies measured by the *CCaR*, can call the function *solverisk*, which provides a solution (if any) for the linear program encoded by equations (11.9)-(11.17), given \mathbf{b}_0 and F , a cost reference c_0 , C_{max} and R_{max} budget limitations, and weights w_1 and w_2 , by executing:

```
1 solution_2 = myproblem.solverisk(b0, F, c0, Cmax, Rmax, w1, w2)
```

Summarizing, the Python input and outputs of PyCaMa are shown in Table 11.1.

Inputs	Outputs	Function or attribute
List of banks	Description of the system	<i>describe</i>
List of transactions	Cost optimal policy list	<i>solvecost</i>
Incidence matrix	Cost-Risk optimal policy list	<i>solverisk</i>
Dictionary of transaction costs	Last optimal policy matrix	<i>policy</i>
Dictionary of holding costs	Last optimal balance matrix	<i>balance</i>
List of minimum balances	Last objective value	<i>objval</i>
Matrix of forecasts	Cost budget	<i>costmax</i>
List of initial cash balances	Risk budget	<i>riskmax</i>
Cost reference	Cost reference	<i>costref</i>
Cost and risk maximum budgets	Cost weight	<i>costweight</i>
Cost and risk weights	Risk weight	<i>riskweight</i>

Table 11.1: Python inputs and outputs of PyCaMa

11.4 An illustrative example

Consider again the cash management system of Figure 11.1 with two current bank accounts 1 and 2, and an investment account 3. Temporary idle cash balances can be invested in short-term marketable securities and bonds through an investment account 3 with an average return of 3.6% per annum, equivalent to 0.01% per day. This is equivalent to set a holding cost 0.01% per day for both accounts 1 and 2. Transactions are allowed between all three accounts and charged with fixed (γ_0) and variable (γ_1) costs determining the cost structure detailed in Table 11.2.

Transaction	γ_0 (€)	γ_1 (%)	Account	v (%)
1	50	0	1	0.01
2	50	0	2	0.01
3	100	0.01	3	0
4	50	0.001		
5	100	0.01		
6	50	0.001		

Table 11.2: Cost structure data for the example.

Assume also that a hypothetical cash manager can obtain forecasts with a maximum cumulative error of 2 million euros for a planning horizon of five days. As a result, she sets a minimum cash balance of 2 million for accounts 1 and 2. Given an initial cash balance $b_0 = [5, 8, 12]$, for accounts 1, 2 and 3, she aims to derive the optimal policy for the next five days. To this end, let us assume that she obtains the next matrix of forecasts (with figures in millions of euros) by applying some predictive method as in Salas-Molina et al. (2017).

$$F = \begin{bmatrix} 1 & -3 & 0 \\ 1 & -9 & 0 \\ 6 & 6 & 0 \\ -1 & 4 & 0 \\ -3 & 6 & 0 \end{bmatrix}. \quad (11.18)$$

An instance of the *multibank* class is created by introducing all the required input data as follows:

```

1 from PyCaMa import * # Import module
2 banks = [1, 2, 3] # Bank accounts
3 trans = [1, 2, 3, 4, 5, 6] # Transactions
4 g0 = {1:50, 2:50, 3:100, 4:50, 5:100, 6:50} # Fixed costs
5 g1 = {1:0, 2:0, 3:100, 4:10, 5:100, 6:10} # Variable costs
6 bmin = [2, 2, 0] # Minimum balances
7 v = {1:100, 2:100, 3:0} # Holding costs
8 A = np.array([[ 1, -1, 0, 0, 1, -1],
9               [-1, 1, 1, -1, 0, 0],
10              [ 0, 0, -1, 1, -1, 1]]) # Incidence matrix
11 myproblem = multibank(banks, trans, A, g0, g1, v, bmin)

```

Then, cash managers can derive the optimal policy by executing function *solvecost*(b_0, F):

```

1 b0 = [5, 8, 12] # Initial balance
2 F = np.array([[ 1, -3, 0],
3               [ 1, -9, 0],
4               [ 6, 6, 0],
5               [-1, -4, 0],
6               [-1, 6, 0]]) # Forecast matrix
7 solution_1 = myproblem.solvecost(b0, F) # Solution

```

A more compact representation of the optimal policy and balances can be obtained by calling functions *policy* and *balance*, which is ready to be visualized through common plotting commands in Python as shown in Figure 11.2. The same instance of the *multibank* class problem can now be solved in terms of cost and risk by setting a cost reference $c_0 = 3,000$ €, budget limits $C_{max} = 5,000$ €, $R_{max} = 5,000$ €, and weights $w_1 = w_2 = 0.5$, and by executing *solverisk*($b_0, F, c_0, C_{max}, R_{max}, w_1, w_2$), resulting in a slightly different policy.




Figure 11.2: Optimal balances for the cost minimization example.

11.5 Impact

PyCaMa is a cash management tool that can be used either to automate decision-making in cash management or to support scientific discovery in the context of computational finance. More precisely, PyCaMa is a promising tool to tackle the following set of open research questions:

- The cash management problem with multiple bank accounts. There is a lack of research about multidimensional cash management systems with the exception of Baccarin (2009), who followed a rigorous theoretical approach. PyCaMa offers a suitable way to follow a more practical approach by providing support for research on actual-world scenarios.
- The impact of cash flow forecasting accuracy. Although the utility of forecasts in cash management was initially demonstrated by Gormley and Meade (2007) and confirmed by Salas-Molina et al. (2017), both approaches were restricted to a single bank account. PyCaMa allows to extend this analysis to cash management systems with multiple bank accounts. As a result, PyCaMa represents a tool to find the best forecasting

models and their potential benefits derived from improving predictive accuracy.

- The utility of different risk measures. A multiobjective approach to the cash management problem has been recently proposed in Salas-Molina, Pla-Santamaria, and Rodriguez-Aguilar (2016) in which the risk of alternative policies is measured by the standard deviation of daily costs. The utility of alternative risk measures can be evaluated by easily extending PyCaMa to consider different risk measures.
- Robust optimization. Two robust approaches to optimization problems were proposed by Soyster (1973) and Ben-Tal and Nemirovski (2002) to deal with uncertainty. PyCaMa can be used to help researchers compare existing and new robust approaches to the cash management problem in terms of both cost and risk.

Moreover, under the realistic assumption of time-varying circumstances, cash managers and researchers are allowed to modify the cost structure to analyze to what extent a change in any of the parameters of the cost structure leads to different optimal policies and, ultimately, to a variation in the cost (and risk) of managing cash. It is also important to highlight that, since Miller and Orr (1966) different optimization models have been proposed based on a set of bounds. Determining such bounds may be problematic in practice due to the strong assumptions made on the probability distribution of cash flows. PyCaMa do not assume any particular form of the cash flow generating process allowing a boundless optimization procedure in the sense that no restriction is placed neither on the form of the policy nor on the distribution of cash flows. As a result, we expect that PyCaMa will be progressively adopted by cash managers and researchers as a more efficient tool to automate decision-making in cash management.

11.6 Summary

In this chapter, we have introduced PyCaMa, a Python module for cost and risk multiobjective optimization within a context of cash management systems with multiple bank accounts. PyCaMa solves the cash management problem when it is formulated as a linear program that aims to minimize either only cost or both cost and risk of cash policies. PyCaMa is implemented through the Gurobi Python modeling environment as a powerful and flexible way to allow an easy integration of its functionality in a more general application. Through an illustrative example, we have shown the key features of PyCaMa, and we have

demonstrated how PyCaMa allows users to model complex cash management systems in an intuitive manner transforming a graphical representation in an optimization model ready to find a solution and to further experimentation. We firmly believe that PyCaMa can be a helpful tool for academic research and financial decision-support software development in the field of short-term financial planning. Natural extensions of PyCaMa include the implementation of different forecasting techniques and additional measures of risk to be added to the current functionality.

Selecting the best policy in cash management systems with multiple bank accounts is by no means straightforward. As a result, a software tool that supports cash managers' decision-making is also a helpful contribution.

Question 10. *Can we automate decision-making in cash management through the use of specific software?*

Contribution 10. *We provided a Python module for multiobjective cash management with multiple bank accounts.*

Conclusions and future work

This chapter summarizes the contributions, discusses the challenges and highlights the benefits of the data-driven approach to multiobjective optimization proposed in this thesis. We followed two mayor directions to address the cash management problem: (i) a data-driven approach based on forecasting cash flows; and (ii) a multiobjective approach to the cash management problem considering not only cost but also the risk of alternative policies. The main available resource for decision-making is data. As a result, we have established a link between machine learning and multiobjective decision-making within the context of cash management.

12.1 Summary of challenges and contributions

In what follows, we first summarize the challenges addressed in this thesis that we ultimately map to the contributions of this thesis. In the first part of this thesis, we focused on cash flow data analysis. Decision-making in cash flow management has been supported by different models assuming different cash flow data generating process such as a deterministic cash flows (Baumol, 1952), a Bernouilli process (Miller and Orr, 1966), a Wiener process (Baccarin, 2009; Constantinides and Richard, 1978; Premachandra, 2004), a double exponential distribution (Penttinen, 1991), or empirical distributions obtained from real data sets (Gormley and Meade, 2007). The assumption of any particular distribution modelling a cash flow process implies the additional assumption of its predictability. Stochastic processes such as the Bernouilli or the Wiener

process consider cash flows totally unpredictable. In practice, however, several useful features from cash flow can be used to predict cash flows resulting in cost savings in cash management. Gormley and Meade (2007) hypothesized that the more accurate the cash flow forecasting accuracy, the larger the expected cost savings. In Chapter 3, we empirically confirmed such hypothesis for the first time in the cash management literature. Furthermore, we analyzed the impact of predictive accuracy on average daily cost savings when considering a variety of cost structures.

Question 1. *Can cash flow predictive accuracy achieve cost savings in the cash management problem?*

Contribution 1. *We empirically confirmed the savings hypothesis showing that predictive accuracy is strongly correlated with cost savings.*

The lack of empirical evidence of the statistical properties of cash flows in the literature (with the exception of Mullins and Homonoff (1976), Emery (1981), and Pindado and Vico (1996)) led us to test the common assumptions of normality, absence of correlation, stationarity and linearity of cash flows. In Chapter 4, we performed a comprehensive comparative study based on 54 real-world daily cash flow data sets showing that the usual assumption of normality, absence of correlation and stationarity hardly appear. We also proposed a new cross-validated test for time series non-linearity that we later used to derive further insight on the implications for forecasting showing that non-linearity is often relevant.

Question 2. *Are common statistical assumptions of daily cash flow supported by recent empirical data?*

Contribution 2. *We demonstrated that normality, absence of correlation and stationarity hardly appear in real-world cash flow data sets. Non-linearity is often relevant and it can be assessed through a new cross-validated test described in this thesis.*

Since data transformation is usually considered as a necessary step by time-series analysis techniques to achieve normality and linearity, we also address

the question of data transformations in Chapter 4. Using the same 54 real-world daily cash flow data sets, we empirically show that data transformations such as Box-Cox (Box and Cox, 1964) or outlier treatment have little impact on linearity and normality.

Question 3. *Is it always possible to achieve a Gaussian, noise-free and linear time-series through data transformations?*

Contribution 3. *We showed that data transformations cannot always achieve a Gaussian, noise-free and linear cash-flow time-series suggesting the utility of non-linear forecasting models.*

It is easy to understand that most decision-making problems have to take into account multiple objectives. Thus, in the second part of this thesis, we followed a multiobjective strategy to solve the CMP. In order to allow cash managers to control the amount of risk that their company take, in Chapter 5, we proposed a new multiobjective approach to the cash management problem based on compromise programming that differs from previous models by considering risk and cash managers' risk preferences. Along this direction, we proposed a multiobjective model to minimize cost and risk in cash management that can employ alternative measures of risk. We further elaborated on the utility of cash management models by formalizing the cash management utility problem from a multiobjective perspective in which we compare the loss derived from a given policy to the loss derived from a trivial policy. Cash managers have now new management tools to control the amount of risk they take in their decision-making processes.

Question 4. *Can we incorporate risk analysis to the cash management problem?*

Contribution 4. *We formulated the cash management problem from a multiobjective perspective considering both cost and risk.*

Decision-makers usually have to deal with time-varying financial circumstances. As a result, cash managers may be interested in identifying the best compromise policies in terms of cost and risk, that are also robust to cash flow regime

changes. Nevertheless, choosing the best policies under a changing context is by no means straightforward. In Chapter 6, we relied on compromise programming to incorporate robustness as an additional goal to cost and risk within a multiobjective framework. More precisely, we proposed to calculate robustness as a multiple criteria distance index that is able to identify the best compromise policies in terms of cost and risk that are also robust to cash flow regime changes.

Question 5. *Can we provide a robust counterpart for any cash management model?*

Contribution 5. *We proposed a new data-driven multiobjective method to derive the robust counterpart for any cash management model by computing distances in a cost-risk space.*

A closely related topic to the previous research question is the concept of operating condition. Chapter 7 demonstrated that ROC analysis can be adapted to enhance the understanding of the multiobjective cash management problem. More precisely, we presented three procedures for: (i) graphically representing models in the cost-risk space; (ii) choosing models according to the risk preferences of cash managers; (iii) deriving cost-risk curves for different operating conditions. These procedures allow cash managers to graphically analyze the tradeoff between cost and risk for different models and operating conditions.

Question 6. *Under what circumstances or operating conditions a model is better than another?*

Contribution 6. *We adapted ROC analysis to the cash management problem in order to allow cash managers to select cash management models.*

From the seminal work by Baumol (1952), all cash management models were based on a set of control bounds. Thus, we call them Bound-Based Models in which cash balance is allowed to wander around between some bounds. Since the ultimate goal of the cash management problem is not to find the best set of bounds but the best sequence of control actions, in Chapter 8, we showed that the constraints imposed by these bounds are not necessary. We

proposed a linear formulation of the CMP that guarantees optimal solutions. We also experimented on 54 real data sets from different companies in Spain and found significant benefits for reasonably low forecasting errors. This fact must encourage cash managers to produce better cash flow forecasts since improvements in accuracy can be converted into important savings as we pointed out in Chapter 3 by means of the savings hypothesis.

Question 7. *Are control bounds really necessary in cash management?*

Contribution 7. *We proposed a new Boundless Model that provides optimal policies using forecasts as a key input without using bounds.*

We know that the solution to the cash management problem is a policy. However, the process to obtain the best policy is not straightforward. Different approaches considering a single objective have been provided in the literature including mathematical programming and other heuristics. Within a multi-objective framework in which both cost and risk are desired but conflicting objectives, we proposed in Chapter 9 to derive the best policies by means of solvers such as linear programming and quadratic programming within the framework of compromise programming. We also made publicly-available the Python code for the the solvers as a starting point for practitioners interested in either designing cash management decision support systems or performing their own experiments in multiobjective cash management.

Question 8. *Can we obtain optimal solutions for the multi-objective cash management problem?*

Contribution 8. *Under the framework of compromise programming, we provided linear and quadratic models and solvers of the cash management problem obtaining optimal solutions.*

In the third part of this thesis, we followed a novel strategy that departs from existing approaches in the cash management literature. Cash management systems with multiple bank accounts are the rule rather than the exception. Thus, we provided a formal specification of the cash management problem

with multiple bank accounts along with theoretical results characterizing cash management systems. We also adapted our initial formulation to provide two alternative robust counterparts of the problem in order to deal with the uncertainty introduced by cash flow forecasts. Finally, we proposed a novel two-stage cost-sensitive robust optimization formulation that considers asymmetric estimation costs that overcomes existing limitations in previous approaches.

Question 9. *Can we derive optimal policies for cash management systems with multiple bank accounts?*

Contribution 9. *After formalizing the multiple bank accounts cash management problem, we provided useful theoretical results on cash management systems and we also proposed a novel data-driven procedure to derive optimal policies.*

Selecting the best policy in cash management systems with multiple bank accounts is by no means straightforward. Thus, a software tool that supports cash managers' decision-making is also a helpful contribution. As a result, we provide the software foundations to develop commercial software products for cash managers.

Question 10. *Can we automate decision-making in cash management through the use of specific software?*

Contribution 10. *We provided a Python module for multiobjective cash management with multiple bank accounts.*

Summarizing, we have contributed to improve decision-making within the context of cash management. We have empirically demonstrated that predictive accuracy in cash flow forecasting leads to cost savings. We have empirically analyzed the main statistical properties of real-world cash flows. We have covered the most important dimensions of the cash management problem extending the state-of-the-art single objective analysis present in the literature to a multiobjective approach in which both cost and risk are desired objectives. We have designed new tools to account for time-varying circumstances and operating conditions. We have proposed a new class of cash management boundless models using forecasts as a key input. We have provided new interesting theoretical results and we have implemented new solvers ready to be

embedded in decision support systems for cash management. We have formalized the cash management problem for multiple bank accounts and we have proposed a new method to deal with cash flow uncertainty. Finally, we have provided a Python module for multiobjective cash management based on linear programming that allows to derive optimal policies for cash management with multiple bank accounts in terms of both cost and risk of policies. In all the aforementioned contributions, we have followed a data-driven decision-making strategy, which we further dissect next, combining useful knowledge derived from data and advanced optimization techniques as the main result of this thesis.

12.2 Data-driven multiobjective decision-making

There is no doubt that the main available resource for decision-making is data. Data contains information and useful information is out there for those who need it, for those who want to use it to make better decisions. Disregarding this information would be equivalent to leaving the tap on wasting water. We firmly believe that data must be incorporated in cash management decision-making as a key input to increase both the understanding and the solutions to business and finance problems. A number of examples of the utility of data in the context of cash management have been described in this thesis. As a result, we have established a strong link between machine learning and multiobjective decision-making that is ready to be used for profit by cash managers.

We mentioned in the introduction that the size of business data bases may continue to increase on a daily basis as a result of transaction recording. Big data is a recent term used to describe the exponential growth in volume, availability and use of information (Doumpos and Grigoroudis, 2013). In this sense, a data-driven approach in which decision-making is based on the analysis of data, rather than purely on intuition should result in better decisions in cash management. This thesis has followed an integrated approach in which machine learning best practices (Chapter 3) have been used in combination to multiple criteria decision-making to solve the cash management problem (Chapter 5). Within this context, we payed special attention to time-series forecasting techniques as a way to reduce cash flow uncertainty.

In cash management research, little attention has been placed on the utility of cash flow forecasts with the exception of Stone (1972) and Gormley and Meade (2007). Both works showed the utility of forecasting in cash management, but none of them researched the importance of the predictive accuracy

and neglected the notion of risk. The underlying value of machine learning in cash management is reflected in two important lessons learned from this thesis: (i) the confirmation of the savings hypothesis (Chapter 3); and (ii) the performance improvements achieved by our Boundless Models for low forecasting errors (Chapter 8). Both lessons must encourage cash managers to produce better cash flow forecasts since improvements in accuracy can be converted into important benefits in terms of cost and risk. In addition, choosing the best cash policies under a changing context is by no means straightforward. In Chapter 6, we also propose a data-driven approach to test robustness in cash management when possible regime changes are contained in a test data set. As a result, we incorporated robustness as an additional goal to cost and risk within a multiobjective framework.

Cash flow forecasting is also crucial to managing multiple bank accounts. Despite the recent advances in cash management, there is a surprising lack of research to aid cash managers to deal with multiple bank accounts. In order to fill this gap, we have provided a multidimensional framework to cover the transition from a single bank account to a more realistic environment in which cash management systems include several bank accounts (Chapter 10). Our definition of cash management system is critical to understand both benefits and opportunities for enhancing decision-making in the sense that any possible configuration can be formulated and solved using state-of-the-art mathematical programming solvers such as CPLEX or Gurobi. Furthermore, we rely again on a data-driven procedure to determine the minimum cash balances required to derive robust cash management policies.

Summarizing, machine learning represents a sound framework to integrate data in the cash management problem through cash flow forecasting to achieve intelligent decision support systems in the sense of Doumpos and Grigoroudis (2013). In addition, multiobjective optimization have been incorporated to field of cash management as a way to improve decision-making through considering additional goals to cost such as risk that may be interesting for cash managers. In this thesis, we offered an integrated approach of machine learning and multiple criteria decision-making which departs from existing cash management research on the following aspects: (i) it can be adapted to cash management models accepting forecasts as a key input; (ii) it estimates the utility of forecasts; (iii) it incorporates risk in decision-making; (iv) it is based on time series cross-validation to suggest the use of either linear or non-linear forecasting models; and (v) it can be extended to consider multiple bank accounts.

12.3 Future work

This thesis has integrated a data-driven procedure such as forecasting in the multiobjective cash management problem in which cost and risk are objectives to optimize by means of state-of-the-art algorithms. We have shown that forecasting accuracy has an important impact on cost savings. Thus, we consider that further research on better forecasting techniques is worthy. On the other hand, alternative measures of risk have been proposed to account for the risk of cash policies within a multiobjective optimization framework. New measures of risk may lead to non-linear problems requiring the need for alternative optimization algorithms different to linear or quadratic programming. Finally, although we have proposed a software tool as a first step towards the construction of decision support systems, there is an evident lack of supporting technology to aid both cash managers in daily decision-making and researchers in the pursuit of new research questions. Next, we further elaborate on these three interesting lines of future work.

12.3.1 *Forecasting techniques*

Different forecasting techniques have been explored in Chapter 3 to be used as a key input to cash management optimization models and we have shown that forecasting accuracy has an important impact on cost savings. This analysis of the relationship between predictive accuracy and cost savings confirmed the importance of better forecasting models when predictions are used as the main input to cash management models. It is worth mentioning again the specific works on cash flow forecasting by Stone (1972), Stone and Wood (1977), Miller and Stone (1985), and Stone and Miller (1987). Some additional intuition have been added in this thesis through the introduction of linear and non-linear forecasting techniques such as regression, random forests or radial basis functions. Non-linearity seems to be an important factor when considering alternative forecasting techniques for cash flow forecasting such as neural networks, support vector machines or even deep learning.

On the other hand, cash flows are usually recorded and stored in the form of time-series. However, since cash flows come from (or go to) known agents within the set of relationships of a firm, alternative approaches to cash flow forecasting can be studied by considering a number of key features for each of these agents. As a result, future work is in place to search for a more informative set of features in the cash management problem by incorporating useful data from customers, suppliers, employees, tax offices and other stakeholders.

In this sense, feature engineering is meant to play a key role to help improve decision-making in cash management.

To end up, it is important to highlight that whenever it is possible to reduce uncertainty about the future, better decisions can be made. The task of building forecasting models is an interesting research discipline in itself, but chances are that cash management optimization models using forecasts as the main input may benefit both from uncertainty reduction and better optimization algorithms.

12.3.2 Evaluating alternative measures of risk

In Chapter 9, we discussed on alternative measures of risk from a theoretical point of view. Advantages and disadvantages of different measures were considered as a first step to select between them. However, an interesting line of future work would be the development of an empirical study comparing alternative measures of risk in terms of a given performance indicator. Note that the selection of this indicator is not a trivial task. Since we aim to compare risk measures, we should evaluate risk as a performance indicator. However, this implies the selection of a measure of risk returning us to the starting point.

Apart from evaluating available measures of risk, this future line of work may produce additional benefits by encouraging researchers to design new measures that are more appropriate for possibly different cash management contexts.

12.3.3 Non-linear multiobjective optimization

As in the case of forecasting techniques, multiobjective optimization is a vast and interesting research field in itself. Here, we do not pursue covering all relevant future challenges of multiobjective optimization but concentrate on how recent research can contribute to improve multiobjective decision-making in cash management.

In Chapter 9, we introduced a number of alternative measures to incorporate risk analysis in cash management. However, defining and measuring risk is not an easy task (Szegö, 2002; McNeil, Frey, and Embrechts, 2005). As a result, a number of alternative measures of risk can be considered according to the particular preferences of cash managers. Apart from coherence in the sense of Artzner et al. (1999), a desired property of any particular risk measure is linearity to be minimized in a linear program. Even quadratic measures of risk such as the variance can be considered to be solved by state-of-the-

art solvers such as CPLEX or Gurobi. However, chances are that new non-linear risk measures require the need for non-linear multiobjective optimization algorithms to solve the cash management problem.

Interactive and non-interactive non-linear multiobjective approaches (Miettinen, 2012), or evolutionary multiobjective optimization (EMO) (Branke et al., 2008), represent a good starting point to incorporate non-linear risk measures in cash management. However, recall that approximate heuristics such as EMO do not guarantee optimality. Thus, the trade-off between optimality and the utility of a non-linear risk measure must be carefully analyzed.

Besides cost and risk, the consideration of additional goals such as stability and robustness constitutes a promising line of work since it may enhance the multiobjective decision-making process. Furthermore, recent results by Qi, Steuer, and Wimmer (2017) suggest an analytic derivation of the efficient surface in portfolio selection with three criteria. The study of the applicability of these results to the field of multiobjective cash management is also a worth tackling question.

12.3.4 Software

Despite the recent advances in cash management, there is a lack of supporting technology to aid the transition from theory to practice. Both cash managers and researchers can benefit from available cash management software for either designing decision-support systems or performing new experiments.

On the one hand, cash managers usually deal with multiple bank accounts to receive payments from customers and to send payments to suppliers. Multiple bank accounts and probably thousands of customers and suppliers makes cash management systems a complex network of financial relationships. Operating such a cash management system is by no means straightforward. Hence, cash managers require decision-support systems to make better decisions in a timely manner within an ever changing context. Software development should necessarily help cash managers to face their daily challenges.

On the other hand, cash management software also contributes to the process of scientific discovery by allowing further, better and faster experiments. In this thesis, we experimented on the utility of forecasts, on the statistical properties of cash flows, on the introduction of risk analysis in cash management, on the robustness of cash management models, on the performance of new cash management models, and on the management of systems with multiple bank accounts. All this experimentation would not have been possible without

the development of specific software. New researchers can benefit from this software but there is a long way ahead to improve our software proposals and to suggest new ones.

Finally, in Chapter 7, we suggested the use of graphical tools to enhance both the understanding and the selection of alternative cash management models. Along the lines of this proposal, we firmly believe that new software developments should include not only numerical optimization algorithms but also graphical tools allowing practitioners to easily visualize the results of alternative strategies.

Appendix A

Terminology and acronyms

In this appendix, we provide useful terminology and acronyms used in this thesis. Cash management deals with receipts and collections received from customers or any other creditor of the company, and payments issued to suppliers or any other debtor of the company. Consequently, cash managers handle the balance between what the company keeps in cash and what has been placed in short-term investments. Thus, the first term we are interested in is cash which can be defined as money, usually in the form of coins or banknotes. However, specially in business practice, coins and banknotes are hardly used and we should better refer to cash as the money in the form of a bank account. Next we itemize useful terminology about cash management and some other related concepts:

- Cash: money in the form of coins, banknotes or bank account.
- Cash balance: readily available amount of cash at any moment in time.
- Cash flow: movement of money into and out of any bank account.
- Inflow, receipt, collection: incoming cash flow, from any other agent into the company.
- Outflow, disbursement, payment: outgoing cash flow, from the company to any other agent.
- Net cash flow: sum of inflows minus sum of outflows.

- Transfer, transaction: all flow of cash into and out of a bank account, usually for control purposes.
- Liquid or short-term asset: any investment which can be easily converted in cash such as deposit accounts, treasury bills and marketable securities.
- Borrowing, loans: funds obtained by paying an interest rate on them.
- Policy: sequence of control actions, or transactions, over a time period.
- Holding cost: cost per money unit of a positive cash balance at the end of the day. Equivalent to the opportunity cost of maintaining a certain amount of idle cash.
- Opportunity cost: the interest that could have been earned in an alternative use of cash.
- Shortage cost: cost per money unit of a negative cash balance at the end of the day.
- Transfer cost: cost of transferring funds from one account to another. It can be either linear, with a fixed and variable part, or non-linear, defined by a more complex cost function.
- Cost function: is a relation between a set of inputs such as a policy or a cash balance and a cost output.
- Cash flow process: a system that generates cash flows, such as a probability density function or any data set with real past cash flows.
- Model: a general function that accepts an initial cash balance condition and a cash flow process and outputs a policy.
- Solver: under the framework of cash management, an algorithm used to find a sufficiently good policy.
- Objective: goal or required attribute, usually expressed as a value function, to be fulfilled by a policy such as reduced cost.
- Maturity: period of time for which a financial instrument remains outstanding at the end of which the financial instrument will cease to exist.
- Cash concentration, cash pooling: the practice of moving cash from multiple bank accounts into a main account.

-
- Saving: cost reduction from a given reference.
 - Time series: a set of observations chronologically ordered.
 - Random variable: quantity subject to changes due to chance.
 - Probability distribution: a theoretical or empirical function that assigns a probability to each value of a random variable.
 - Drift: trend of a time series usually measured by the average change within a given time range.
 - Variability, volatility: dispersion of observations usually measured by the standard deviation.
 - Heteroskedasticity: it is said of a data set with subsets that have different variabilities.
 - Normality: it is said of data set following Gaussian distribution with observations symmetrically centered around the mean, and with finite variance.
 - Correlation: it is said of a data set with ordered observations when the occurrence of past observations affects the probability of occurrence of the next ones.
 - Stationarity: it is said of a data set whose probability distribution does not change over time and, consequently, its statistical properties such as the mean and variance remain stable.
 - Linearity: proportionality either to another (external) explanatory variable or to a combination of (external) explanatory variables.

In Table A.1, we provide a list of the acronyms and abbreviations used in this thesis for quick referencing:

Table A.1: Acronyms and abbreviations

Acronym	Meaning
AR	Autoregression
ARIMA	Autoregression Integrated Moving Average
BM	Boundless Model
BBM	Bound-Based Model
CaR	Cost-at-Risk
CCaR	Conditional Cost-at-Risk
CMP	Cash Management Problem
CMUP	Cash Management Utility Problem
CP	Compromise Programming
CR	Cost-Risk
CSRO	Cost-Sensitive Robust Optimization
DP	Dynamic Programming
DSP	Dynamic Simple Policy
DT	Data Transformation
GA	Genetic Algorithm
GM	Gormley-Meade
GP	Goal Programming
LP	Linear Programming
MBACMP	Multiple Bank Accounts Cash Management Problem
MCDM	Multiple Criteria Decision-Making
MOCMP	Multiobjective Cash Management Problem
MOGM	Multiobjective Gormley-Meade
NSE	Normalized Squared Error
OT	Outlier Treatment
p.a.	Per annum
PSO	Particle Swarm Optimization
QP	Quadratic Programming
RBF	Radial Basis Function
RF	Random Forest
ROC	Receiver Operating Characteristic
SI	Seasonal Indicator
SLR	Sharpe-Like Ratio
SP	Stochastic Programming
UPM	Upper Partial Moment
WGP	Weighted Goal Programming
WSLR	Weighted Sharpe-Like Ratio

Bibliography

- Abdelaziz, Fouad Ben, Belaid Aouni, and Rimeh El Fayedh (2007). “Multi-objective stochastic programming for portfolio selection”. In: *European Journal of Operational Research* 177.3, pp. 1811–1823.
- Aouni, Belaid, Cinzia Colapinto, and Davide La Torre (2014). “Financial portfolio management through the goal programming model: Current state-of-the-art”. In: *European Journal of Operational Research* 234.2, pp. 536–545.
- Aouni, Belaid and Davide La Torre (2010). “A generalized stochastic goal programming model”. In: *Applied Mathematics and Computation* 215.12, pp. 4347–4357.
- Archer, Stephen H (1966). “A model for the determination of firm cash balances”. In: *Journal of Financial and Quantitative Analysis* 1.01, pp. 1–11.
- Artzner, Philippe et al. (1999). “Coherent measures of risk”. In: *Mathematical finance* 9.3, pp. 203–228.
- Baccarin, Stefano (2002). “Optimal impulse control for cash management with quadratic holding-penalty costs”. In: *Decisions in Economics and Finance* 25.1, pp. 19–32.

- Baccarin, Stefano (2009). “Optimal impulse control for a multidimensional cash management system with generalized cost functions”. In: *European Journal of Operational Research* 196.1, pp. 198–206.
- Ballester, E (1998). “Approximating the optimum portfolio for an investor with particular preferences”. In: *Journal of the Operational Research Society* 49, pp. 998–1000.
- Ballester, Enrique (2005). “Mean-Semivariance Efficient Frontier: A Downside Risk Model for Portfolio Selection”. In: *Applied Mathematical Finance* 12.1, pp. 1–15.
- (2007). “Compromise programming: A utility-based linear-quadratic composite metric from the trade-off between achievement and balanced (non-corner) solutions”. In: *European journal of operational research* 182.3, pp. 1369–1382.
- Ballester, Enrique and David Pla-Santamaria (2004). “Selecting portfolios for mutual funds”. In: *Omega* 32.5, pp. 385–394.
- Ballester, Enrique and Carlos Romero (1998). *Multiple criteria decision making and its applications to economic problems*. Springer Science & Business Media.
- Bates, Thomas W, Kathleen M Kahle, and René M Stulz (2009). “Why do US firms hold so much more cash than they used to?” In: *The journal of finance* 64.5, pp. 1985–2021.
- Baumol, William J (1952). “The transactions demand for cash: An inventory theoretic approach”. In: *The Quarterly Journal of Economics* 66.4, pp. 545–556.
- Ben-Tal, Aharon, Laurent El Ghaoui, and Arkadi Nemirovski (2009). *Robust optimization*. Princeton University Press.
- Ben-Tal, Aharon and Arkadi Nemirovski (1999). “Robust solutions of uncertain linear programs”. In: *Operations research letters* 25.1, pp. 1–13.

-
- (2000). “Robust solutions of linear programming problems contaminated with uncertain data”. In: *Mathematical programming* 88.3, pp. 411–424.
- (2002). “Robust optimization—methodology and applications”. In: *Mathematical Programming* 92.3, pp. 453–480.
- Bensoussan, A and JL Lions (1975). “Nouvelles méthodes en contrôle impulsionnel”. In: *Applied Mathematics and Optimization* 1.4, pp. 289–312.
- Bertsimas, Dimitris, David B Brown, and Constantine Caramanis (2011). “Theory and applications of robust optimization”. In: *SIAM review* 53.3, pp. 464–501.
- Bertsimas, Dimitris and Melvyn Sim (2004). “The price of robustness”. In: *Operations research* 52.1, pp. 35–53.
- Bi, Jinbo and Kristin P Bennett (2003). “Regression error characteristic curves”. In: *Proceedings of the 20th International Conference on Machine Learning (ICML-03)*, pp. 43–50.
- Birge, John R and Francois Louveaux (2011). *Introduction to stochastic programming*. Springer Science & Business Media.
- Black, Fischer and Myron Scholes (1973). “The pricing of options and corporate liabilities”. In: *The Journal of Political Economy* 81.3, pp. 637–654.
- Boulbrachene, Messaoud (1998). “The noncoercive quasi-variational inequalities related to impulse control problems”. In: *Computers & Mathematics with Applications* 35.12, pp. 101–108.
- Box, George EP and David R Cox (1964). “An analysis of transformations”. In: *Journal of the Royal Statistical Society. Series B (Methodological)*, pp. 211–252.
- Box, George EP and Gwilym M Jenkins (1976). *Time series analysis: forecasting and control*. Holden-Day.

- Bradley, Andrew P (1997). “The use of the area under the ROC curve in the evaluation of machine learning algorithms”. In: *Pattern recognition* 30.7, pp. 1145–1159.
- Bradley, Michael D and Dennis W Jansen (2004). “Forecasting with a nonlinear dynamic model of stock returns and industrial production”. In: *International Journal of Forecasting* 20.2, pp. 321–342.
- Branke, Jürgen et al. (2008). *Multiobjective optimization: Interactive and evolutionary approaches*. Vol. 5252. Springer.
- Bravo, Mila, Enrique Ballesterro, and David Pla-Santamaria (2012). “Evaluating Fund Performance by Compromise Programming with Linear-Quadratic Composite Metric: An Actual Case on The CaixaBank in Spain”. In: *Journal of Multi-Criteria Decision Analysis* 19.5-6, pp. 247–255.
- Brealey, Richard A and Myers (2003). *Principles of corporate finance*. McGraw-Hill.
- Breiman, Leo (2001). “Random forests”. In: *Machine learning* 45.1, pp. 5–32.
- Broomhead, D and D Lowe (1988). “Multivariable functional interpolation and adaptive networks”. In: *Complex Systems* 2, pp. 321–355.
- Castle, Jennifer L and David F Hendry (2010). “A low-dimension portmanteau test for non-linearity”. In: *Journal of Econometrics* 158.2, pp. 231–245.
- (2012). “Automatic selection for non-linear models”. In: *System Identification, Environmental Modelling, and Control System Design*. Springer, pp. 229–250.
- Chelouah, Rachid and Patrick Siarry (2000). “A continuous genetic algorithm designed for the global optimization of multimodal functions”. In: *Journal of Heuristics* 6.2, pp. 191–213.
- Chen, Xin and David Simchi-Levi (2009). “A new approach for the stochastic cash balance problem with fixed costs”. In: *Probability in the Engineering and Informational Sciences* 23.04, pp. 545–562.

- Clements, Michael P, Philip Hans Franses, and Norman R Swanson (2004). "Forecasting economic and financial time-series with non-linear models". In: *International Journal of Forecasting* 20.2, pp. 169–183.
- Cleveland, Robert B et al. (1990). "STL: A seasonal-trend decomposition procedure based on loess". In: *Journal of Official Statistics* 6.1, pp. 3–73.
- Conejo, Antonio J et al. (2005). "Forecasting electricity prices for a day-ahead pool-based electric energy market". In: *International Journal of Forecasting* 21.3, pp. 435–462.
- Constantinides, George M and Scott F Richard (1978). "Existence of optimal simple policies for discounted-cost inventory and cash management in continuous time". In: *Operations Research* 26.4, pp. 620–636.
- Cortey-Dumont, Ph (1985). "On finite element approximation in the L-infinite norm of variational inequalities". In: *Numerische Mathematik* 47.1, pp. 45–57.
- Costa Moraes, Marcelo Botelho da and Marcelo Seido Nagano (2014). "Evolutionary models in cash management policies with multiple assets". In: *Economic Modelling* 39, pp. 1–7.
- Costa Moraes, Marcelo Botelho da, Marcelo Seido Nagano, and Vinicius Amorim Sobreiro (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.
- Criminisi, Antonio and Jamie Shotton (2013). *Decision forests for computer vision and medical image analysis*. Springer.
- Daellenbach, Hans G (1971). "A stochastic cash balance model with two sources of short-term funds". In: *International Economic Review* 12.2, pp. 250–256.
- (1974). "Are cash management optimization models worthwhile?" In: *Journal of Financial and Quantitative Analysis* 9.04, pp. 607–626.
- De Gooijer, Jan G and Rob J Hyndman (2006). "25 years of time series forecasting". In: *International journal of forecasting* 22.3, pp. 443–473.

- Deza, Michel Marie and Elena Deza (2014). *Encyclopedia of distances*. Third Edition. Springer.
- Dietterich, Thomas G (2000). “Ensemble methods in machine learning”. In: *International workshop on multiple classifier systems*. Springer, pp. 1–15.
- Doornik, Jurgen A (2008). “Encompassing and automatic model selection”. In: *Oxford Bulletin of Economics and Statistics* 70.1, pp. 915–925.
- Doumpos, Michael and Evangelos Grigoroudis (2013). *Multicriteria decision aid and artificial intelligence: links, theory and applications*. John Wiley & Sons.
- Doumpos, Michael and Constantin Zopounidis (2007). “Model combination for credit risk assessment: A stacked generalization approach”. In: *Annals of Operations Research* 151.1, pp. 289–306.
- Drummond, Chris and Robert C Holte (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*. ACM, pp. 198–207.
- (2006). “Cost curves: An improved method for visualizing classifier performance”. In: *Machine learning* 65.1, pp. 95–130.
- Elkan, Charles (2001). “The foundations of cost-sensitive learning”. In: *International joint conference on artificial intelligence*. Vol. 17. 1. Lawrence Erlbaum Associates Ltd, pp. 973–978.
- Emery, Gary W (1981). “Some empirical evidence on the properties of daily cash flow”. In: *Financial management* 10.1, pp. 21–28.
- Eppen, Gary D and Eugene F Fama (1968). “Solutions for cash-balance and simple dynamic-portfolio problems”. In: *Journal of Business* 41.1, pp. 94–112.
- (1969). “Cash balance and simple dynamic portfolio problems with proportional costs”. In: *International Economic Review* 10.2, pp. 119–133.

- Fabozzi, Frank J and Leslie N Masonson (1985). *Corporate Cash Management: Techniques and Analysis*. Irwin Professional Publishing.
- Fawcett, Tom (2006). “An introduction to ROC analysis”. In: *Pattern recognition letters* 27.8, pp. 861–874.
- Flach, Peter A (2003). “The geometry of ROC space: understanding machine learning metrics through ROC isometrics”. In: *ICML*, pp. 194–201.
- Franses, Philip Hans and Dick Van Dijk (2005). “The forecasting performance of various models for seasonality and nonlinearity for quarterly industrial production”. In: *International Journal of Forecasting* 21.1, pp. 87–102.
- Frenkel, Jacob and Boyan Jovanovic (1980). “On transactions and precautionary demand for money”. In: *The Quarterly Journal of Economics* 94, 24–43.
- Gao, Huasheng, Jarrad Harford, and Kai Li (2013). “Determinants of corporate cash policy: Insights from private firms”. In: *Journal of Financial Economics* 109.3, pp. 623–639.
- Garcia-Bernabeu, Ana et al. (2016). “Photovoltaic power plants: a multicriteria approach to investment decisions and a case study in western Spain”. In: *Annals of Operations Research* 245.1-2, pp. 163–175.
- Girgis, Nadia Makary (1968). “Optimal cash balance levels”. In: *Management Science* 15.3, pp. 130–140.
- Glasserman, Paul (2003). *Monte Carlo methods in financial engineering*. Springer Science & Business Media.
- Gormley, Fionnuala M and Nigel Meade (2007). “The utility of cash flow forecasts in the management of corporate cash balances”. In: *European journal of operational research* 182.2, pp. 923–935.
- Granger, Clive WJ, Timo Terasvirta, et al. (1993). “Modelling non-linear economic relationships”. In: *OUP Catalogue*. Oxford University Press.

- Gregory, Geoffrey (1976). “Cash flow models: a review”. In: *Omega* 4.6, pp. 643–656.
- Grubbs, Frank E (1969). “Procedures for detecting outlying observations in samples”. In: *Technometrics* 11.1, pp. 1–21.
- Gurobi Optimization, Inc (2016). *Gurobi Optimizer Reference Manual*.
- Hastie, Trevor et al. (2005). “The elements of statistical learning: data mining, inference and prediction”. In: *The Mathematical Intelligencer* 27.2, pp. 83–85.
- Hendry, David F and Jurgen A Doornik (2014). *Empirical model discovery and theory evaluation: automatic selection methods in econometrics*. MIT Press.
- Hernández-Orallo, José (2013a). “ROC curves for regression”. In: *Pattern Recognition* 46.12, pp. 3395–3411.
- (2013b). “ROC curves for regression”. In: *Pattern Recognition* 46.12, pp. 3395–3411.
- Hernández-Orallo, José, Peter Flach, and César Ferri (2013). “ROC curves in cost space”. In: *Machine learning* 93.1, pp. 71–91.
- Hernández-Orallo, José, Nicolas Lachiche, and Adolfo Martínez-Usó (2014). “Predictive models for multidimensional data when the resolution context changes”. In: *Workshop on Learning over Multiple Contexts at ECML*. Vol. 2014.
- Hinderer, Karl and K-H Waldmann (2001). “Cash management in a randomly varying environment”. In: *European Journal of Operational Research* 130.3, pp. 468–485.
- Ho, Tin Kam (1995). “Random decision forests”. In: *Document Analysis and Recognition, 1995., Proceedings of the Third International Conference on*. Vol. 1. IEEE, pp. 278–282.

- (1998). “The random subspace method for constructing decision forests”. In: *Pattern Analysis and Machine Intelligence, IEEE Transactions on* 20.8, pp. 832–844.
- Hodge, Victoria J and Jim Austin (2004). “A survey of outlier detection methodologies”. In: *Artificial intelligence review* 22.2, pp. 85–126.
- Hollander, Myles, Douglas A Wolfe, and Eric Chicken (2013). *Nonparametric statistical methods*. John Wiley & Sons.
- Homonoff, Richard and David Wiley Mullins (1975). *Cash management: an inventory control limit approach*. Lexington Books.
- Hornik, Kurt, Maxwell Stinchcombe, and Halbert White (1989). “Multilayer feedforward networks are universal approximators”. In: *Neural networks* 2.5, pp. 359–366.
- Hyndman, RJ and G Athanasopoulos (2013). *Forecasting: principles and practice*.
- Hyndman, Rob J (2016). *forecast: Forecasting functions for time series and linear models*. R package version 7.1.
- Itô, Kiyosi (1974). *Diffusion Processes*. Wiley Online Library.
- Kane, Michael J et al. (2014). “Comparison of ARIMA and Random Forest time series models for prediction of avian influenza H5N1 outbreaks”. In: *BMC bioinformatics* 15.1, p. 1.
- Kantz, Holger and Thomas Schreiber (2004). *Nonlinear time series analysis*. Vol. 7. Cambridge University Press.
- Keenan, Daniel MacRae (1985). “A Tukey nonadditivity-type test for time series nonlinearity”. In: *Biometrika* 72.1, pp. 39–44.
- Keynes, John Maynard (1936). *General theory of employment, interest and money*. Macmillan Cambridge University Press.

- Kumar, Manish and M Thenmozhi (2006). “Forecasting stock index movement: A comparison of support vector machines and random forest”. In: *Indian Institute of Capital Markets 9th Capital Markets Conference Paper*.
- Laukaitis, Algirdas (2008). “Functional data analysis for cash flow and transactions intensity continuous-time prediction using Hilbert-valued autoregressive processes”. In: *European Journal of Operational Research* 185.3, pp. 1607–1614.
- Lee, Tae-Hwy, Halbert White, and Clive WJ Granger (1993). “Testing for neglected nonlinearity in time series models: A comparison of neural network methods and alternative tests”. In: *Journal of Econometrics* 56.3, pp. 269–290.
- Liaw, Andy and Matthew Wiener (2002). “Classification and Regression by random forest”. In: *R News* 2.3, pp. 18–22.
- Lilliefors, Hubert W (1967). “On the Kolmogorov-Smirnov test for normality with mean and variance unknown”. In: *Journal of the American Statistical Association* 62.318, pp. 399–402.
- Ljung, Greta M and George EP Box (1978). “On a measure of lack of fit in time series models”. In: *Biometrika* 65.2, pp. 297–303.
- Makridakis, Spyros, Steven C Wheelwright, and Rob J Hyndman (2008). *Forecasting methods and applications*. John Wiley & Sons.
- Markowitz, Harry (1952). “Portfolio selection”. In: *The journal of finance* 7.1, pp. 77–91.
- Marsland, Stephen (2009). *Machine learning: an algorithmic perspective*. CRC press.
- McNeil, Alexander J, Rüdiger Frey, and Paul Embrechts (2005). *Quantitative risk management: Concepts, techniques and tools*. Princeton University Press.

- Mei, Jie et al. (2014). “A random forest method for real-time price forecasting in New York electricity market”. In: *PES General Meeting/ Conference & Exposition, 2014 IEEE*. IEEE, pp. 1–5.
- Melo, Marcos AS and Feruccio Bilich (2013). “Expectancy balance model for cash flow”. In: *Journal of Economics and Finance* 37.2, pp. 240–252.
- Metz, Charles E (1978). “Basic principles of ROC analysis”. In: *Seminars in nuclear medicine*. Vol. 8. 4. Elsevier, pp. 283–298.
- Mierzejewski, Fernando (2011). “A model of equilibrium in markets of cash balances”. In: *IMA Journal of Management Mathematics* 22.3, pp. 253–270.
- Miettinen, Kaisa (2012). *Nonlinear multiobjective optimization*. Vol. 12. Springer Science & Business Media.
- Miller, Merton H and Daniel Orr (1966). “A Model of the Demand for Money by Firms”. In: *The Quarterly Journal of Economics* 80.3, pp. 413–435.
- Miller, Tom W and Bernell K Stone (1985). “Daily Cash Forecasting and Seasonal Resolution: Alternative Models and Techniques for Using the Distribution Approach”. In: *Journal of Financial and Quantitative Analysis* 20.03, pp. 335–351.
- Mitchell, Tom M et al. (1997). *Machine learning*. McGraw-Hill Boston.
- Montgomery, Douglas C (2013). *Design and analysis of experiments*. Eighth Edition. Wiley.
- Muller, Patrice et al. (2015). *Annual report on european SMEs*.
- Mullins, David and Richard Homonoff (1976). *Applications of inventory cash management models*. SC Myers.
- Myers, Stewart C and Richard Arthur Brealey (2003). *Principles of corporate finance*. Seventh Edition. McGraw-Hill.

- Neave, Edwin H (1970). “The stochastic cash balance problem with fixed costs for increases and decreases”. In: *Management Science* 16.7, pp. 472–490.
- Park, Hae-Sang and Chi-Hyuck Jun (2009). “A simple and fast algorithm for K-medoids clustering”. In: *Expert Systems with Applications* 36.2, pp. 3336–3341.
- Penttinen, Markku J (1991). “Myopic and stationary solutions for stochastic cash balance problems”. In: *European journal of operational research* 52.2, pp. 155–166.
- Pindado, J and J Vico (1996). “Evidencia empírica sobre los flujos de caja. Un nuevo enfoque en su tratamiento”. In: *Revista Española de Financiación y Contabilidad* 25.87, pp. 497–517.
- Pinkowitz, Lee, René M Stulz, and Rohan Williamson (2016). “Do US Firms Hold More Cash than Foreign Firms Do?” In: *Review of Financial Studies* 29.2, pp. 309–348.
- Pla-Santamaria, David and Mila Bravo (2013). “Portfolio optimization based on downside risk: a mean-semivariance efficient frontier from Dow Jones blue chips”. In: *Annals of Operations Research* 205.1, pp. 189–201.
- Prékopa, András (2013). *Stochastic programming*. Vol. 324. Springer Science & Business Media.
- Premachandra, IM (2004). “A diffusion approximation model for managing cash in firms: An alternative approach to the Miller–Orr model”. In: *European Journal of Operational Research* 157.1, pp. 218–226.
- Provost, Foster and Tom Fawcett (2013). *Data Science for Business: What you need to know about data mining and data-analytic thinking*. O’Reilly Media, Inc.
- Qi, Yue, Ralph E Steuer, and Maximilian Wimmer (2017). “An analytical derivation of the efficient surface in portfolio selection with three criteria”. In: *Annals of Operations Research* 251.1-2, pp. 161–177.

- Ramsey, James Bernard (1969). “Tests for specification errors in classical linear least-squares regression analysis”. In: *Journal of the Royal Statistical Society. Series B (Methodological)* 31.2, pp. 350–371.
- Richard, Scott F (1977). “Optimal impulse control of a diffusion process with both fixed and proportional costs of control”. In: *SIAM Journal on Control and Optimization* 15.1, pp. 79–91.
- Ringuest, Jeffrey L (1992). *Multiobjective optimization: behavioral and computational considerations*. Springer Science & Business Media.
- Rockafellar, R Tyrrell and Stanislav Uryasev (2002). “Conditional value-at-risk for general loss distributions”. In: *Journal of Banking & Finance* 26.7, pp. 1443–1471.
- Roijers, Diederik Marijn et al. (2013). “A survey of multi-objective sequential decision-making”. In: *Journal of Artificial Intelligence Research*.
- Ross, Stephen A, Randolph Westerfield, and Bradford D Jordan (2002). *Fundamentals of corporate finance*. Sixth. McGraw-Hill.
- Rousseeuw, Peter J and Annick M Leroy (1987). *Robust regression and outlier detection*. Vol. 589. John Wiley & Sons.
- Royston, JP (1982). “An extension of Shapiro and Wilk’s W test for normality to large samples”. In: *Applied Statistics*, pp. 115–124.
- Sahinidis, Nikolaos V (2004). “Optimization under uncertainty: state-of-the-art and opportunities”. In: *Computers & Chemical Engineering* 28.6, pp. 971–983.
- Salas-Molina, Francisco, David Pla-Santamaria, and Juan A Rodriguez-Aguilar (2016). “A multi-objective approach to the cash management problem”. In: *Annals of Operations Research*, pp. 1–15.
- Salas-Molina, Francisco et al. (2017). “Empowering cash managers to achieve cost savings by improving predictive accuracy”. In: *International Journal of Forecasting* 33.2, pp. 403–415.

- Sarantis, Nicholas (2001). “Nonlinearities, cyclical behaviour and predictability in stock markets: international evidence”. In: *International Journal of Forecasting* 17.3, pp. 459–482.
- Serra, Joan et al. (2012). “Predictability of music descriptor time series and its application to cover song detection”. In: *IEEE Transactions on Audio, Speech, and Language Processing* 20.2, pp. 514–525.
- Sharpe, William F (1966). “Mutual fund performance”. In: *The Journal of Business* 39.1, pp. 119–138.
- (1994). “The sharpe ratio”. In: *The journal of Portfolio Management* 21.1, pp. 49–58.
- Small, Michael (2005). *Applied nonlinear time series analysis: applications in physics, physiology and finance*. Vol. 52. World Scientific.
- Smith, Gregor W (1986). “A dynamic Baumol-Tobin model of money demand”. In: *The Review of Economic Studies* 53.3, pp. 465–469.
- Soyster, Allen L (1973). “Technical note—convex programming with set-inclusive constraints and applications to inexact linear programming”. In: *Operations research* 21.5, pp. 1154–1157.
- Srinivasan, Venkat and Yong H Kim (1986). “Deterministic cash flow management: state of the art and research directions”. In: *Omega* 14.2, pp. 145–166.
- Steuer, Ralph E, Yue Qi, and Markus Hirschberger (2007). “Suitable-portfolio investors, nondominated frontier sensitivity, and the effect of multiple objectives on standard portfolio selection”. In: *Annals of Operations Research* 152.1, pp. 297–317.
- Stone, Bernell K (1972). “The use of forecasts and smoothing in control-limit models for cash management”. In: *Financial Management* 1.1, pp. 72–84.
- (1973). “Cash planning and credit-line determination with a financial statement simulator: A case report on short-term financial planning”. In: *Journal of Financial and Quantitative Analysis* 8.05, pp. 711–729.

- Stone, Bernell K and Tom W Miller (1981). “Daily cash forecasting: a structuring framework”. In: *Journal of Cash Management* 1.1, pp. 35–50.
- (1987). “Daily cash forecasting with multiplicative models of cash flow patterns”. In: *Financial Management* 16.4, pp. 45–54.
- Stone, Bernell K and Robert A Wood (1977). “Daily cash forecasting: a simple method for implementing the distribution approach”. In: *Financial Management* 6.3, pp. 40–50.
- Szegö, Giorgio (2002). “Measures of risk”. In: *Journal of Banking & Finance* 26.7, pp. 1253–1272.
- Teräsvirta, Timo (2006). “Forecasting economic variables with nonlinear models”. In: vol. 1. Elsevier, pp. 413–457.
- Tobin, James (1956). “The interest elasticity of transactions demand for cash”. In: *Review of Economic and Statistics* 38.3, pp. 241–247.
- Torgo, Luís (2005). “Regression error characteristic surfaces”. In: *Proceedings of the eleventh ACM SIGKDD international conference on Knowledge discovery in data mining*. ACM, pp. 697–702.
- Venables, William N and Brian D Ripley (2013). *Modern applied statistics with S-PLUS*. Springer Science & Business Media.
- Weibull, Waloddi et al. (1951). “A statistical distribution function of wide applicability”. In: *Journal of Applied Mechanics* 18.3, pp. 293–297.
- Weigend, Andreas S (1994). “Time series prediction: forecasting the future and understanding the past”. In: *Santa Fe Institute Studies in the Sciences of Complexity*.
- Whalen, Edward L (1966). “A rationalization of the precautionary demand for cash”. In: *The Quarterly Journal of Economics* 80.2, pp. 314–324.
- Wilcoxon, Frank, SK Katti, and Roberta A Wilcox (1970). *Critical values and probability levels for the Wilcoxon rank sum test and the Wilcoxon signed rank test*.

- Ye, Qing and Izak Duenyas (2007). “Optimal capacity investment decisions with two-sided fixed-capacity adjustment costs”. In: *Operations research* 55.2, pp. 272–283.
- Yu, Po-Lung (1973). “A class of solutions for group decision problems”. In: *Management Science* 19.8, pp. 936–946.
- (1985). *Multiple criteria decision making: concepts, techniques and extensions*. Plenum Press.
- (2013). *Multiple-criteria decision making: concepts, techniques, and extensions*. Springer Science & Business Media.
- Zagorecki, Adam (2015). “Prediction of Methane Outbreaks in Coal Mines from Multivariate Time Series Using Random Forest”. In: *Rough Sets, Fuzzy Sets, Data Mining, and Granular Computing*. Springer, pp. 494–500.
- Zeleny, Milan (1974). “A concept of compromise solutions and the method of the displaced ideal”. In: *Computers & Operations Research* 1.3, pp. 479–496.
- (1982). *Multiple criteria decision making*. McGraw-Hill.
- Zhang, Guoqiang, B Eddy Patuwo, and Michael Y Hu (1998). “Forecasting with artificial neural networks:: The state of the art”. In: *International journal of forecasting* 14.1, pp. 35–62.
- Zopounidis, Constantin (1999). “Multicriteria decision aid in financial management”. In: *European Journal of Operational Research* 119.2, pp. 404–415.
- Zopounidis, Constantin and Michael Doumpos (2013). “Multicriteria decision systems for financial problems”. In: *Top* 21.2, pp. 241–261.