

# Chapter 11

## Optimal Financial Decision Making Under Uncertainty

Giorgio Consigli, Daniel Kuhn, and Paolo Brandimarte

**Abstract** We use a fairly general framework to analyze a rich variety of financial optimization models presented in the literature, with emphasis on contributions included in this volume and a related special issue of *OR Spectrum*. We do not aim at providing readers with an exhaustive survey, rather we focus on a limited but significant set of modeling and methodological issues. The framework is based on a benchmark discrete-time stochastic control optimization framework, and a benchmark financial problem, asset-liability management, whose generality is considered in this chapter. A wide set of financial problems, ranging from asset allocation to financial engineering problems, is outlined, in terms of objectives, risk models, solution methods, and model users. We pay special attention to the interplay between alternative uncertainty representations and solution methods, which have an impact on the kind of solution which is obtained. Finally, we outline relevant directions for further research and optimization paradigms integration.

**Keywords** Stochastic control • Dynamic programming • Multistage stochastic programming • Robust optimization • Distributionally robust optimization • Decision rules • Asset-liability management • Pension fund management

---

G. Consigli (✉)

Department of Management, Economics and Quantitative Methods, University of Bergamo,  
Via dei Caniana 2, 24127 Bergamo, Italy  
e-mail: [giorgio.consigli@unibg.it](mailto:giorgio.consigli@unibg.it)

D. Kuhn

Risk Analytics and Optimization Chair, École Polytechnique Fédérale de Lausanne, CH-1015  
Lausanne, Switzerland  
e-mail: [daniel.kuhn@epfl.ch](mailto:daniel.kuhn@epfl.ch)

P. Brandimarte

Dipartimento di Scienze Matematiche, Politecnico di Torino, Corso Duca degli Abruzzi 24,  
10129 Torino, Italy  
e-mail: [paolo.brandimarte@polito.it](mailto:paolo.brandimarte@polito.it)

© Springer International Publishing Switzerland 2017

G. Consigli et al. (eds.), *Optimal Financial Decision Making under Uncertainty*,  
International Series in Operations Research & Management Science 245,  
DOI 10.1007/978-3-319-41613-7\_11

255

## 11.1 The Domain of Financial Optimization

Since the seminal work of Markowitz [71], the literature on the application of optimization models to financial decision problems has witnessed an astonishing growth. The contributions presented in this volume and the companion special issue (SI) of OR Spectrum [38] offer a broad picture illustrating a variety of problems and solution approaches that have been the subject of recent research, from both a theoretical and an applied perspective. The main purpose of this chapter is to review the building blocks of recent research on optimization models in finance. We do not aim at giving an exhaustive literature survey, though, and due emphasis is given to contributions to this volume and the SI. The (less ambitious) aim is to reconsider the contributions within a common framework in order to spot research directions and integration opportunities. This should be especially valuable to practitioners and newcomers, possibly Ph.D. students, who may find the heterogeneity of literature somewhat confusing.

Papers dealing with optimization in finance may be characterized according to different features:

- The kind of financial problem that is addressed, such as portfolio selection, asset pricing, hedging, or asset-liability management. These are sometimes considered as different problems but, actually, there is an interplay among them and we see in what follows that indeed *asset-liability management* (ALM) models provide a rather general modeling framework. The specific practical problem tackled restricts the choice of decisions to be made, the constraints that decisions must satisfy, and the criteria to evaluate solution quality. For a given financial problem, alternative modeling and solution approaches may be available.
- We have to define a risk model, which consists of a set of relevant risk factors and a representation of their uncertainty. Classical models rely on a probabilistic characterization of uncertainty, but there is a growing stream of contributions dealing with robustness and ambiguity, as well as model-free, data-driven approaches. Some of these paradigms are indeed non-stochastic, and the choice of a risk model must be compatible with the available information used for estimation and fitting purposes.
- The approach taken to manage a risk-reward tradeoff. Modern Portfolio Theory (MPT) relies on standard deviation of return as a risk measure, which is traded off against expected return within a static model. This is an example of the mean-risk approach, which may be generalized by using other risk measures. However, alternative approaches may be taken, relying on classical utility functions, or on stochastic dominance concepts. MPT revolves around a static portfolio choice, and the extension of risk measures to multiple periods turns out to be rather tricky. Mostly due to evolving regulatory frameworks and increasing tail risk, the set of possible risk-reward trade-offs has increased dramatically over the last two decades.
- All of the previous choices yield an optimization model formulation, which, among other things, may be static or dynamic, discrete-time or continuous-time.

Model selection and problem formulation, however, cannot be made without paying due attention to computational viability. The essential trade-off is in this context between finding the exact solution of a maybe wrong model, or an approximate solution of the right model.

- The set of available solution approaches has widened considerably over the years, due to increasing hardware power and the development of convex optimization techniques. The Markowitz model for instance boils down to a simple and deterministic convex optimization problem and, from a strictly algorithmic viewpoint, is not quite challenging. Sophisticated methods for multistage stochastic programming with recourse, dynamic programming, and robust optimization are called for when tackling challenging problems, and they determine the exact kind of solution that we find and the way in which it can be used.
- Last, but not least, the solution must somehow be validated, and the overall modeling approach must be questioned and critically assessed.

By combining the above features, a wide variety of optimization models can be developed, which should be viewed within a common framework, as far as possible. Finding a generic formulation that can be instantiated to yield any conceivable model is a hopeless endeavour, but we believe that it is useful to set a benchmark model as a reference, in order to compare different model instances and methodological challenges. From a formal viewpoint, a discrete-time stochastic control model is arguably a good reference framework, whereas, from a financial viewpoint, an asset-liability management model is a suitably general problem.

No optimization model can clearly be considered without reference to a real-world financial context that practitioners have to face in their day-by-day activity. Hence, we start by summarizing a few relevant trends that affect financial decision making in Sect. 11.2. Then, in Sect. 11.3 we describe a benchmark modeling framework and the essential elements in building a financial decision model under uncertainty. There, we emphasize the interplay between the model building and the model solving approaches, which is particularly critical in the context of multistage decision problems. Then, in Sect. 11.4 we take a financial view, and consider how several classes of financial optimization problems may be considered as specific ALM cases. We introduce a pension fund ALM problem as representative of a variety of financial management problems. In Sect. 11.5 we follow up discussing a limited set of solution approaches suitable for dynamic problems: stochastic programming with recourse, distributionally robust optimization and learning decision rules. Finally, we outline possible research directions in Sect. 11.6, and we draw some conclusions in Sect. 11.7.

## 11.2 A Changing Financial Landscape

The contributions in this volume and the SI, to a certain extent stimulated by ongoing and continuously refined risk management regulatory frameworks in the banking and insurance industry, reflect a structural transformation of agents'

management options and decision paradigms within a new financial environment [38]. Such transition may be analyzed from different viewpoints: here we consider primarily practitioners' and modelers' perspectives, the latter with associated computational and numerical implications. When considered in terms of *supply* and *demand* sides, as is common in financial economics, we see that market evolution can be conveniently linked to agents' decision processes: long positions in the market are associated with the demand side and lead to an optimal asset management problem, in which agents' risk attitude need to be considered and a given representation of risk is required. On the contrary short positions are associated with the supply side and lead to an optimal liability management problem and a related pricing issue. The two indeed embed two ALM problems, as we shall see here below. A distinctive element of the financial markets' growth is associated with increasing product diversification, which calls for modeling approaches able to accommodate non-trivial decision spaces, see [1, 46, 54, 66]. The dimension of the investment universe in real applications has an impact on both the adopted uncertainty model and the available optimization options, particularly in dynamic problems [42, 78, 81].

The recent literature on financial decision making reflects the emergence of new market features that call for a revision of traditional assumptions:

- The persistence of an unprecedented period of interest rate curves flattening at almost zero level, which leads to a quest for increased sources of return, and the consequent emergence of alternative investment opportunities, possibly involving illiquid assets.
- A revision of the set of relevant risk factors (risk as exposure to uncertain outcomes that can be assigned probabilities) like sovereign risk, even in OECD countries, longevity risk, systemic risk, as well as model risk, and ambiguity (interpreted as exposure to uncertain outcomes that admit no probabilistic description) [25, 46].
- The heterogeneity of agents' planning horizons, as different agents may be concerned with long vs. short time horizons, and the need to balance long-term objectives (as typical of pension funds), with short-term performance.

The above features have clear implications on the formulation of the associated optimization problem:

- Financial markets need not be consistent with canonical information efficiency assumptions and financial optimization approaches may very well adapt to alternative assumptions. The evolution of risk premia in equity, bond and alternative markets carrying different liquidity is complex to model and forecast but can hardly be avoided [39, 70, 78, 93].
- Asset pricing models need not be based on the assumption of market completeness [1, 46, 54].
- Decision horizons and rebalancing frequency matter: when moving from a single-period, myopic setting to a dynamic one, the risk modelling effort increases substantially and the issue of the optimal policies' and risk measures'

dynamic time consistency becomes central [34, 49]. Further the issue of an increasing dependence of the optimal policy on functions of *risk-adjusted performance* measures [21, 44, 78] becomes relevant over increasing planning horizons. On the other hand extending the horizon requires short-term risk management and profitability not to be jeopardized [1].

Modeling implications are relevant. Indeed, the contributions in this volume and the SI [38] may be distinguished as carrying scenario-based [42] or parametric [21, 58] or set-based [81] representations of uncertainty; model-driven [39, 42, 58, 70, 93] or data-driven approaches [25, 32, 55, 81]; reflecting equilibrium conditions—*arbitrage free* [1, 54, 70] or just based on statistical criteria [25, 31, 32].

From a computational viewpoint, the central approximation issue when dealing with dynamics typically carrying a continuous probability space must be addressed. A trade-off problem lies at the heart of such approximation effort, where an increasingly accurate characterization of the uncertainty may not be consistent with a realistic problem description, particularly in real-world applications [10]. A *sufficient* approximation, achieved through a computationally efficient approach, is regarded as a necessary condition to determine accurate risk control strategies. A substantial literature in stochastic programming addresses scenario generation issues [37, 48, 82]. The following evidences emerge in this respect in this volume:

- robust approaches [25, 81] rely on an approximation scheme which reduces the computational burden in terms of scenario management at the cost of increasing computational cost in the optimization phase [81],
- stochastic programming approaches rely very much on scenario *reduction* methods and the trade-off between approximation quality and stability of the optimal solution drives the computational effort [42].

To a certain extent independently from the adopted problem formulation and specific features, the definition of an effective decision process in contemporary markets relies very much on a sequence of steps: from the analysis of the decision universe, the definition of decision criteria driving the optimization process, the introduction of a statistical model and the derivation of internal (to the agents' economy) and external constraints.

### 11.3 The Elements of a Decision Model

The original mean-variance model is static and aims at managing an asset portfolio. As a result, decision variables are just portfolio weights and the distributional inputs are (seemingly) modest: a vector of expected returns and a covariance matrix. As it turns out, even this distributional information is not trivial to give in a robust way, but the matter gets much more complicated in a more general dynamic framework. In this section we describe a reference optimization model, namely, a generic discrete-time stochastic control problem. This cannot be regarded as an abstract model from

which any relevant financial optimization model can be instantiated, but it allows us to spot the key elements of an optimization model, which have been specified in a wide variety of ways in the literature. From an applied financial perspective, a fairly general model framework can be identified in the *asset-liability* management (ALM) domain.

### 11.3.1 Discrete-Time Stochastic Control

We consider the formulation of an optimal financial planning problem as the optimal control of a dynamic stochastic system whose essential elements in discrete-time are:

- A discrete and finite sequence of time instants,  $\mathcal{T} := \{0, \dots, T\}$ . The discretization may arise from a suitable partition of a continuous time interval, or as a consequence of a decision process where choices are only made at specific time instants.
- A sequence of control/decision variables denoted by  $\mathbf{x}_t$ ,  $t = 0, \dots, T - 1$ . In general no decision is allowed at the end of the time horizon  $T$ , where we only check the last outcome.
- A sequence of state variables  $\mathbf{s}_t$ ,  $t = 0, \dots, T$ ;  $\mathbf{s}_T$  is the terminal state.
- A sequence of random variables  $\xi_t(\omega)$ ,  $t = 1, \dots, T$ . This is the stochastic process followed by risk factors, where  $\omega \in \Omega$  corresponds to a sample path. The stochastic process may be discrete-state, continuous-state, or a hybrid. In most financial optimization problems it is assumed that uncertainty is purely exogenous, i.e., risk factors are not influenced by control actions.

Controlling the system means making a decision, i.e., choosing a control  $\mathbf{x}_t$  at each time  $t$ ,  $t = 0, \dots, T - 1$ , after observing the state  $\mathbf{s}_t$ . At  $t = 0$  only  $\mathbf{s}_0$  is known and the first control  $\mathbf{x}_0$  is taken under complete uncertainty. The next state  $\mathbf{s}_{t+1}$  may depend in a possibly complex way on the state and control trajectory up to  $t$ , as well as on the realization  $\xi_{t+1}(\omega)$  of the risk factors. A relatively simple case applies when the Markov property holds, which implies that we may introduce a state transition function  $\Phi_t$ , at time  $t$ , such that

$$\mathbf{s}_{t+1} = \Phi_t(\mathbf{s}_t, \mathbf{x}_t, \xi_{t+1}(\omega)). \quad (11.1)$$

The last control action  $\mathbf{x}_{T-1}$  will result in the last state  $\mathbf{s}_T$ . The first requirement of a control action is feasibility, and we may denote the feasible set at time  $t$  as  $A_t(\mathbf{s}_t)$ , emphasizing dependence on the current state. Note that the feasible set is random, as it depends on risk factors through the state variable. The sequence of controls must be a *good* one, where quality may be measured by introducing a sequence of functions  $f_t(\mathbf{s}_t, \mathbf{x}_t)$ ,  $t = 1, \dots, T - 1$ , and a function  $F_T(\mathbf{s}_T)$  to evaluate the terminal state. These functions may be maximized or minimized depending on their nature, and we end up with a somewhat loose and abstract formulation:

$$\min_{\mathbf{x}_t} \mathbb{E} \left[ \sum_{t=0}^{T-1} f_t(\mathbf{s}_t, \mathbf{x}_t) + F_T(\mathbf{s}_T) \right] \quad (11.2)$$

$$\text{s.t. } \mathbf{x}_t \in A_t(\mathbf{s}_t) \quad (11.3)$$

$$\mathbf{s}_{t+1} = \Phi_t(\mathbf{s}_t, \mathbf{x}_t, \xi_{t+1}(\omega)). \quad (11.4)$$

This formulation is quite intuitive, but loose, as it does not clarify a few important features (see also [34]).

- The nature of the solution. In a deterministic multiperiod problem, we fix the whole set of control actions  $\mathbf{x}_t$  at time  $t = 0$ . Thus, decisions are a sequence of vectors. On the contrary, in a stochastic *multistage* problem the sequence of decisions is a random process, as it depends on the unfolding of uncertainty. They may be thought of as a function of the current state, which gives a solution in feedback form, or as a function of the underlying stochastic process. In the last case, care must be taken to enforce a sensible non-anticipativity condition. We further discuss the nature of the solution in Sect. 11.3.2.
- The objective function (11.2) is additive with respect to time, and might be specified in a way in which only the terminal state matters, or the whole trajectory. The involved functions could be used to capture the risk-reward tradeoff, which may be expressed by risk measures or utility functions. However, we cannot take for granted that a simple additive structure will be able to capture complex trade-offs often driving decision processes.
- The satisfaction of constraints (11.3) must be further qualified. If the representation of uncertainty relies on a discrete tree process constraints must be always satisfied, but when dealing with continuous factors, *almost sure* (a.s., with probability 1) feasibility is required. This condition may be relaxed by requiring that the constraint is satisfied with a suitable probability, leading to chance-constrained problems.
- The state transition function (11.4) captures the uncertain evolution of the state variable, but we leave the evolution of the risk factors implicit. A critical modeling choice is the selection of a suitable set of driving risk factors that affect, e.g., prices (interest rate and credit spreads affect bond prices). The underlying stochastic process ranges from a simple sequence of i.i.d. variables to a complex process exhibiting path dependencies, going through the intermediate case of a Markov process. The expectation in the objective function is taken with respect to a probability measure associated with this model. However, an increasingly relevant amount of research questions our ability to pinpoint a probability measure reliably. Distributional ambiguity and non-stochastic characterizations of uncertainty have been proposed.

In the following we will delve more deeply into these issues, but before doing so, a simple but relevant example is necessary.

### 11.3.2 *The Interplay Between Model Building and Solution Method*

When looking at the model formulation (11.2), it is not quite apparent which kind of solution we may get and how it can be used. This depends on how uncertainty is expressed and often *approximated* to make the problem computationally viable. The dynamics of  $\mathbf{s}_t$  and  $\boldsymbol{\xi}(\omega)$  will depend on such assumptions. Alternative approaches to express uncertainty include:

- Scenario trees, which we may think as discrete characterization of a continuous-time and/or continuous-state model as well as generated by discrete processes [37];
- Model-free, data-driven approaches [25, 32, 46];
- Uncertainty sets maybe associated with probability distribution supports.

These alternatives, discussed below in more detail, may lead to different solution strategies in the stochastic control problem:

- Multistage stochastic programming. Emphasis is given to the first stage or *implementable* decision  $\mathbf{x}_0$ .
- Dynamic programming, where we typically, recover by backward recursion an optimal policy, whose time and state evolution will in general depend on the statistical assumptions of the underlying random process,
- Robust optimization has been extended recently to multistage problems [81], called adjustable robust optimization. Robust approaches yield optimal strategies with respect to uncertainty sets. Distributional robustness instead originate from uncertainty over probability domains.
- The adoption of decision rules as shown below has received increasing interest to address dynamic optimization problems [90].

It is useful to clarify first the key distinctive elements of those modeling and solution options with respect to the adopted representation of uncertainty and the assumptions on the underlying data process.

#### 11.3.2.1 Scenario Tree, Non-anticipativity and Information

The discrete evolution of a dynamic stochastic system is often described by a *tree* process, which captures the random dynamics of  $\boldsymbol{\xi}_t(\omega)$ . Scenario trees reflect the dynamic interaction between control actions—the first taken under full uncertainty, the others always under residual uncertainty—and revelation of uncertainty over a maybe very long time horizon [1, 27, 42, 66, 81]. The tree clearly shows the information to which the decision process  $\mathbf{x}_t(\omega)$  must be adapted. A filtered probability space is associated with the tree. Consider the simple example in Fig. 11.1. At time  $t = 1$  we essentially see a  $\sigma$ -algebra generated by the sets  $\{\omega_1, \omega_2, \omega_3, \omega_4\}$  and  $\{\omega_5, \omega_6, \omega_7, \omega_8\}$ , whereas at time  $t = 2$  random variables



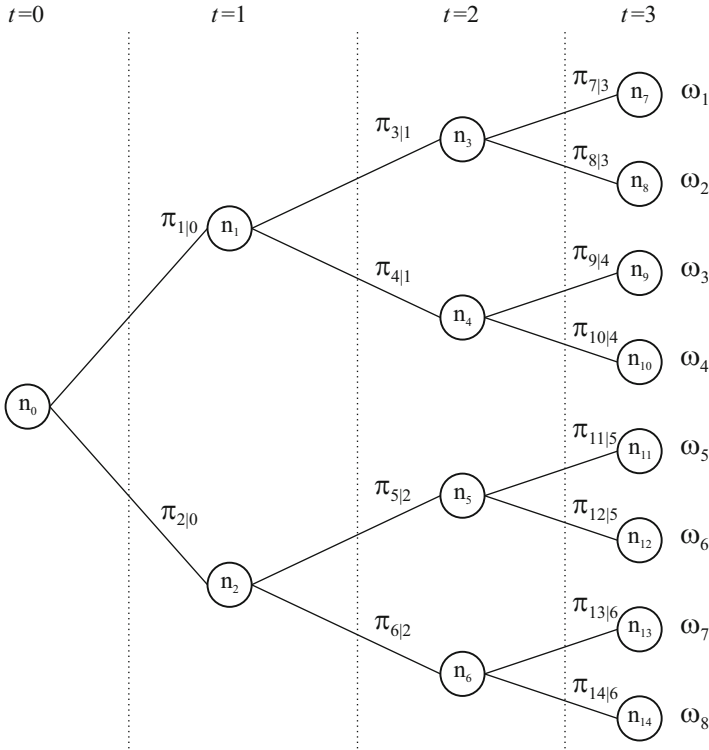


Fig. 11.1 A typical scenario tree

must be measurable with respect to the  $\sigma$ -algebra generated by  $\{\omega_1, \omega_2\}$ ,  $\{\omega_3, \omega_4\}$ ,  $\{\omega_5, \omega_6\}$ , and  $\{\omega_7, \omega_8\}$ . As such, the sequence of  $\sigma$ -algebras provides an information model and for consistency, we require decisions to be adapted, measurable, with respect to those information flows.

The decision process is non-anticipative, which means that a decision at any node in the tree must be the same for any scenario that is indistinguishable up to that time instant. As we shall see in Sect. 11.4.2, this requirement may be built in the formulation of the model by an appropriate choice of variables, which are associated with nodes. A noticeable feature of multistage stochastic programming is that the solution is a stochastic decision process, namely a tree, where the decision at the root node, the implementable decision, is what matters [10]. On the one hand, this gives stochastic programming an operational flavor. On the other hand, this may also be a disadvantage with respect to other approaches, due to decreasing computational efficiency. The conditional behavior of a (non-recombining) scenario tree process is reflected in a specific labeling convention which is useful to summarize. We adopt a particularly simple convention and denote with  $\mathcal{N}$  the set of nodes in the tree;  $n_0$  is the root node. The (unique) predecessor of node  $n \in \mathcal{N} \setminus \{n_0\}$  is denoted by  $a(n)$ , while the set of terminal nodes is denoted by  $\mathcal{S}$ . Each node belongs to a scenario, which is the sequence of event nodes along the unique paths leading from

$n_0$  to  $s \in \mathcal{S}$ , with probability  $\pi^s$ .  $\mathcal{T} = \mathcal{N} \setminus (\{n_0\} \cup \mathcal{S})$  is the set of intermediate trading nodes. We refer to [54] for a more structured labeling scheme. The product of conditional probabilities in Fig. 11.1 will determine the scenario probabilities. In presence of vector tree processes, as common in financial applications, path probabilities are assumed to be equal with conditional probabilities at each stage equal to 1 over the number of branches departing from that node.


The above stochastic structure is common for multistage stochastic programs [10] and provides an intuitive model to understand the interplay between underlying random process evolution, the associated information process and the resulting optimal decision process, defined by a root decision and a sequence of *recourse* decisions. An extended set of financial models carrying specific properties has been adopted over a wide range of financial applications [1, 38, 39, 42, 70]. Validation by stability analysis is necessary and an additional complication in finance: scenarios must be arbitrage-free.

Estimating a full-fledged uncertainty model is certainly not trivial, and one should question its validity. One possible alternative, is to resort to a data-driven approach. Then, a more natural stochastic dynamics can be built according to the following model without any conditional structure, as such referred to as *linear scenarios* or *scenario fan*.

This is the canonical output of risk measurement applications based on Monte Carlo methods [18] in which, for given initial input portfolio, the evaluation of relevant statistical percentiles may be performed. Here we are primarily interested in this model representation in connection with the solution of a stochastic control problem. In a dynamic setting, the model is indeed also associated with a stochastic program where a condition of *perfect foresight* is assumed, corresponding to the so-called wait-and-see problem [15]. The solution of problem (11.2) under the wait-and-see assumption provides a lower bound to the stochastic solution: such difference reflects the expected value of perfect information [15]. As the set of  $\mathcal{S}$  increases, the associated solution  $\mathbf{x}_{n_0}$  is expected to converge to a stable solution. In financial optimization problems it is common to enforce a *non-arbitrage* condition [42, 54] to require a branching degree at each stage bounded below by the cardinality of the investment universe.

Alternatively assume a simple one-period tree following the dynamics in Fig. 11.2. This would be a possible representation of a discrete probability space with finite and countable support and typically a probability measure giving to each path the same probability. Such model would be stage-wise consistent with a (stochastic) dynamic programming solution approach [34] within a backward recursion algorithm [18].

Robust models and distributionally robust models do not require a specification of an underlying process' sample paths and neither of them will depend on such discrete approximation: instead they will focus either on an uncertainty domain associated with  $\omega$  or, more precisely  $\xi(\omega)$ , or on uncertainty affecting the probability measure to be adopted to describe the problem stochasticity. When, still under the (collapsed 1-period) process characterization in Fig. 11.2, a data-driven approach is adopted then as common in risk management applications relying on *historical*



**Fig. 11.2** Linear scenarios

*simulation*, past data realizations are used, typically within a bootstrapping approach [18, 25, 32], to populate the sample space in the problem definition. The associated probability distribution is in this case also referred to as the *empirical* distribution. Empirical distributions are generated by mapping into future random events a given data history, relying on a discrete probability distribution.

Summarizing, either probabilistic models are adopted in financial optimization problems, but the model based on what we introduced as linear scenarios, may be considered either in one period control problems or in dynamic problems but in such case relying on backward recursion approaches. Accordingly, not to violate the required non-anticipativity condition or measurability condition of the optimal strategy (with respect to the current information structure), any decision must be taken in face of residual uncertainty. Here next we consider how different assumptions on the probability space translate into two very popular optimization paradigms, eventually leading, jointly, to an increasingly popular problem formulation.

### 11.3.2.2 Stochastic, Robust and Distributionally Robust Optimization

A compact representation of problem (11.2)—where for sake of simplicity all constraints are embedded in the decision space  $\mathcal{X}$  definition, and a random process  $\xi$ , defined in an appropriate probability space, is assumed to characterize the problem’s overall uncertainty—is the following stochastic program:

$$\min_{\mathbf{x} \in \mathcal{X}} \mathbb{E}_{\mathbb{P}} [f(\mathbf{x}, \xi)]. \tag{11.5}$$

Expectation is taken under the probability measure  $\mathbb{P}$  and  $f$  is a functional, or risk function, mapping the interaction between decisions and random events into a given payoff or risk estimate, we wish to control. We refer to [34] in this volume for a thorough analysis of dynamic risk measures. The formulation (11.5) is indeed consistent with a dynamic stochastic program where the decision space is described as a product space and non-anticipativity is enforced implicitly. The probability space is endowed with a filtration, that we may assume generated by the process  $\xi$ , and expectation is taken at the end of the planning horizon [49]. For a planning horizon  $T = 1$ , the above problem formulation collapses to the one-period static case [21, 25, 32, 55]. We consider here general properties of those formulations.

It is widely recognized that multistage stochastic programs have represented now for quite sometime an effective mathematical framework for decision making under uncertainty in several application domains with noticeable examples specifically in finance and energy [10]. A distinctive element of this approach may be found in its ability through a scenario tree representation to combine an optimal decision tree process with specific assumptions on the underlying random process evolution. A discrete framework allows an accurate and rich mathematical representation of the decision problem with a sufficient description of the underlying sources of uncertainty [37]. In presence of a risk exposure generated by continuously evolving market conditions (e.g. commodity prices for energy problems, financial returns in portfolio management and so forth), randomness evolves continuously in time and an approximation issue arises, as well known. The adoption of sampling methods and in general scenario reduction and generation methods aimed at minimizing the *cost* associated with such discrete approximation has represented a relevant research focus in this context [10]. The trade-off between computational viability and scenario tree expansion is central to the modeling effort.

Statistical modeling, particularly in the case of financial applications, has been typically welcome by practitioners as a longly established way to incorporate economic and financial stylised evidences and solid quantitative analysis in day-by-day decision making and within optimization models. Applied research in risk management, based on advanced statistical models had been indeed extremely successful to capture financial portfolios' risk exposure for years [91] and it was stimulated by regulatory institutions.

A fundamental challenge in stochastic programming formulations, however, is that the distribution of  $\xi$  is in general not directly observable but must be estimated from historical time series data. In fact, not even the simplest moments such as the means and covariances can be estimated to within an acceptable precision [69, 72]. Estimation errors are problematic because financial optimization problems tend to be highly sensitive to the distributional input parameters. Consequently, estimation errors in input parameters are amplified in the optimization (e.g., assets with overestimated mean returns are given too much weight and assets with underestimated means are given too little weight in the optimal portfolio), which results in unstable portfolios that perform poorly in out-of-sample experiments [20, 35, 73]. This phenomenon is akin to overfitting in statistics: a model that is perfectly optimized for in-sample data has little explanatory power and displays poor generalizability to out-of-sample data.

A popular approach to combat estimation errors in input parameters of financial optimization models is to adopt a robust approach. Here, the uncertain input parameters are assumed to reside within an uncertainty set  $\mathcal{U}$  that captures the agent's prior knowledge of the uncertainty. The objective of robust optimization models is to identify decisions that are optimal under the worst possible parameter realizations within the uncertainty set [9, 12]. This gives rise to worst-case optimization problems of the form

$$\min_{\mathbf{x} \in \mathcal{X}} \max_{\boldsymbol{\xi} \in \mathcal{U}} f(\mathbf{x}, \boldsymbol{\xi}). \quad (11.6)$$

The resulting decisions display an attractive non-inferiority property, that is, their out-of-sample performance is necessarily better (lower) than the optimal value of the robust optimization problem provided that the uncertain parameters materialize within the prescribed uncertainty set. The shape of the uncertainty set in problems of the type (11.6) should be chosen judiciously as it can heavily impact the quality of the resulting optimal decisions and the tractability of the robust optimization problem. Popular uncertainty sets that have been studied in a financial context are the box uncertainty set [92], the ellipsoidal uncertainty set [50, 52], the budget uncertainty set [4, 5], the factor model uncertainty set [57] as well as a class of data-driven uncertainty sets that are constructed using statistical hypothesis tests [14]. In robust portfolio optimization, for instance, there is strong evidence that robust portfolios are less susceptible to overfitting effects than classical Markowitz portfolios and therefore display an improved out-of-sample performance [30]. Generalized robust portfolio optimization models that include both stocks and European-style options have been discussed in [98]. This model offers two layers of robustness guarantees: a weak guarantee that holds whenever the asset returns materialize within a prescribed confidence set reflecting normal market conditions and a strong guarantee that becomes effective when the asset returns materialize outside of the confidence set. The model is therefore akin to the comprehensive robust counterpart model [8], which allows for a controlled deterioration in performance when the data falls outside of the uncertainty set. Robust models with objective functions that are linear in the decisions but convex nonlinear in the uncertain parameters have been proposed in [62]. This versatile model can capture nonlinear dependencies between prices and returns as they are common in classical stochastic stock price models.

Despite the ostensible simplicity of modeling uncertainty through sets, robust optimization has been exceptionally successful in providing high-quality and efficiently computable solutions for a broad spectrum of decision problems ranging from engineering design, finance, and machine learning to policy making and business analytics [9]. Nevertheless, it has been observed that robust optimization models can lead to an under-specification of uncertainty as they fail to exploit prior distributional information that may be available. In these situations, robust optimization models may lead to over-conservative decisions. By also exploiting properties of stochastic programming models, distributionally robust models address this issue.

Since the pioneering work of Keynes [63] and Knight [65], it is common in decision theory to distinguish the concepts of risk and ambiguity. Classical stochastic programming can only be used in a risky environment. Indeed, the objective of stochastic programming is to minimize the expectation or some risk measure of the cost  $f(\mathbf{x}, \xi)$ , where the expectation is taken with respect to the distribution  $\mathbb{P}$  of  $\xi$ , which is assumed to be known. If we identify the uncertainty set  $\mathcal{U}$  in the robust optimization problem (11.6) with the support of the probability distribution  $\mathbb{P}$ , then it becomes apparent that stochastic programming and robust optimization offer complementary models for the decision maker's risk attitude: however, when viewed through a decision theory lens, neither stochastic programming nor robust optimization address ambiguity.

Distributionally robust optimization is a natural generalization of both stochastic programming and robust optimization, which accounts both for the decision maker's attitude towards risk and ambiguity. In particular, a generic distributionally robust optimization problem can be formulated as

$$\min_{\mathbf{x} \in \mathcal{X}} \max_{\mathbb{P} \in \mathcal{P}} \mathbb{E}_{\mathbb{P}} [f(\mathbf{x}, \xi)], \quad (11.7)$$

where the probability distribution  $\mathbb{P}$  of  $\xi$  is only known to belong to a prescribed ambiguity set  $\mathcal{P}$ , that is, a family of (possibly infinitely many) probability distributions consistent with the available raw data or prior structural information. The distributionally robust optimization model can be interpreted as a game against 'nature'. In this game, the agent first selects a decision with the goal to minimize expected costs, in response to which 'nature' selects a distribution from within the ambiguity set with the goal to inflict maximum harm to the agent. This setup prompts the agent to select worst-case optimal decisions that offer performance guarantees valid for all distributions in the ambiguity set.

Notice that the introduced min-max problem (11.7) belongs indeed also to the tradition of stochastic programming problems from the seminal work in [47]. Distributionally robust optimization leads to less conservative decisions than classical robust optimization, and it enables modelers to incorporate information about estimation errors into optimization problems. Therefore, it results in a more realistic account of uncertainty. Moreover, maybe surprisingly, distributionally robust optimization problems can often be solved exactly and in polynomial time—very much like the simpler robust optimization models. For a general introduction to distributionally robust optimization we refer to [40, 56, 95].

Many innovations in distributionally robust optimization have originated from the study of financial decision problem. For instance, one of the earliest modern distributionally robust optimization models in the literature studies the construction of portfolios that are optimal in terms of worst-case Value-at-Risk, where the worst case is taken over all asset return distributions with given means and covariances [50]. Worst-case expected utility maximization models are addressed in [80], while worst-case Value-at-Risk minimization models for nonlinear portfolios containing stocks and options are described in [97]. Moreover, distributionally robust portfolio

optimization models using an ambiguity set in which some marginal distributions are known, while the global dependency structure or copula is uncertain, are studied in [45]. While many papers focus on ambiguity sets described by first- and second-order moment, possibly complemented by support information, asymmetric distributional information in the form of forward- and backward-deviation measures are described in [79]. Distributionally robust growth-optimal portfolios that offer attractive performance guarantees across several investment periods are described in [86], where the asset returns are assumed to follow a weak sense white noise process, which means that the ambiguity set contains all distributions under which the asset returns are serially uncorrelated and have period-wise identical first and second-order moments. Besides moment-based ambiguity sets, which abound in the current literature and offer attractive tractability properties, a different stream of research has focused on the construction of ambiguity sets using probability metrics. Here, the idea is to construct ambiguity sets that can be viewed as balls in the space of probability distributions with respect to a probability distance function such as the Prohorov metric [51], the Kullback-Leibler divergence [60], or the Wasserstein metric [75, 83] etc. Such metric-based ambiguity sets contain all distributions that are sufficiently close to a prescribed nominal distribution with respect to the prescribed probability metric. This setup allows the modeler to control the degree of conservatism of the underlying optimization problem by tuning the radius of the ambiguity set. In particular, if the radius is set to zero, the ambiguity set collapses to a singleton that contains only the nominal distribution. Then, the distributionally robust optimization problem reduces to the classical stochastic program (11.5).

In the final section the above general concepts are considered in the domain of a financial application. We complete this short methodological summary by considering data-driven approaches as in [25, 32, 81] and their compatibility with robust and stochastic programming formulations.

### 11.3.2.3 Data-Driven Optimization

Most classical stochastic programs as well as the vast majority of robust and distributionally robust optimization models are constructed with a parametric uncertainty model in mind. Thus, the distribution of  $\xi$  is estimated statistically, or it is constructed on the basis of expert information or known structural properties. The estimated distribution  $\hat{\mathbb{P}}$  may then be directly used in the stochastic program (11.5). Alternatively, the support or an appropriate confidence set of  $\hat{\mathbb{P}}$  can be viewed as an uncertainty set  $\mathcal{U}$  for the robust optimization problem (11.6). Yet another possibility is to use  $\hat{\mathbb{P}}$  to construct an ambiguity set  $\mathcal{P}$  for the distributionally robust optimization problem (11.7) (e.g., by defining a moment ambiguity set via some standard or generalized moments of  $\hat{\mathbb{P}}$  or by designating  $\hat{\mathbb{P}}$  as the center of a spherical ambiguity set with respect to a probability metric). In contrast, data-driven optimization uses any available data directly in the optimization model—without the detour of calibrating a statistical model.

Assume, for instance, that we have access to a set of independent samples  $\xi^{(1)}, \dots, \xi^{(N)}$  drawn uniformly from the (unknown) distribution  $\mathbb{P}$ . We now outline how one can use such a time series to construct data-driven variants of the stochastic, robust, and distributionally robust optimization problems (11.5)–(11.7), respectively. First, the celebrated sample average approximation [89, Chap. 5]

$$\min_{\mathbf{x} \in \mathcal{X}} \frac{1}{N} \sum_{i=1}^N f(\mathbf{x}, \xi^{(i)}) \quad (11.8)$$

provides a data-driven approximation of the stochastic program (11.5). Here, the unknown true distribution  $\mathbb{P}$  is effectively replaced by the uniform distribution on the data points  $\xi^{(1)}, \dots, \xi^{(N)}$ , which is the empirical distribution. The sample average approximation has been used with great success in financial engineering, risk management and economics [96]. For a modern textbook treatment of the broader area of Monte Carlo simulation we refer to [18].

A data-driven counterpart of the robust optimization problem (11.6) is given by the scenario program

$$\min_{\mathbf{x} \in \mathcal{X}} \max_{i=1, \dots, N} f(\mathbf{x}, \xi^{(i)}), \quad (11.9)$$

which is in fact a variant of (11.6) where the uncertainty set  $\mathcal{U}$  has been replaced by the support of the empirical distribution. Even though it is conceptually simple, the scenario program has many desirable theoretical properties. On the one hand, the scenario program (11.9) is efficiently solvable even in situations where the robust problem (11.6) with a polyhedral or ellipsoidal uncertainty set  $\mathcal{U}$  is intractable. Due to the inner approximation entailed by the sampling, the solution of (11.9) underestimates the optimal value of (11.6), which typically involves an infinite number of scenarios. However, it has been shown that the cost of any optimal solution of (11.9) under a new data point  $\xi^{(N+1)}$  is bounded by the optimal value of (11.9) with high probability provided the number of samples  $N$  is sufficiently large [26, 28]. This result is remarkable as it holds independently of the distribution  $\mathbb{P}$  of the samples and therefore is applicable even in situations where  $\mathbb{P}$  is unknown—as is typically the case in financial applications. Thus, we can interpret (11.9) as an optimization problem that minimizes the cost threshold that can be exceeded only with a certain prescribed probability, which implies that (11.9) is closely related to a value-at-risk minimization problem. For a more detailed discussion of this connection we refer to [24]; applications in portfolio analysis and design are described in [25]. Another modern approach to data-driven robust optimization seeks decisions that are robust with respect to the set of all parameters that pass a prescribed statistical hypothesis test [14].

There are different approaches to deriving data-driven counterparts of the distributionally robust optimization problem (11.7). In [40] it is shown how time series data can be used to construct confidence sets for the first two moments of  $\xi$ . The ambiguity set of all distributions whose first two moments reside within



these confidence sets is guaranteed to contain the unknown true data-generating distribution  $\mathbb{P}$  at the prescribed confidence level. The time series data can also be used in more direct ways to construct ambiguity sets. Specifically, several authors have proposed to study the following data-driven distributionally robust optimization problem

$$\min_{\mathbf{x} \in \mathcal{X}} \max_{\mathbb{P} \in \mathbb{B}_r(\hat{\mathbb{P}})} \mathbb{E}_{\mathbb{P}} [f(\mathbf{x}, \xi)], \quad (11.10)$$

where the ambiguity set  $\mathbb{B}_r(\hat{\mathbb{P}})$  is defined as the ball of radius  $r$  around the empirical distribution  $\hat{\mathbb{P}}$  with respect to some probability metric. Note that problem (11.10) reduces to the sample average approximation (11.8) when the radius  $r$  drops to zero. In practice, the crux is to select the radius  $r$  judiciously such that the unknown data-generating distribution  $\mathbb{P}$  belongs to the ambiguity set  $\mathbb{B}_r(\hat{\mathbb{P}})$  with high confidence. The optimal solution of (11.10) thus displays attractive out-of-sample guarantees, but it may be difficult to compute. In order to ease the computational burden, it has been shown that a robust problem (11.6) whose uncertainty set  $\mathcal{U}$  is a union of norm balls centered at the  $N$  sample points may provide a close and computationally tractable approximation for (11.10). In [75] it has been shown that the data-driven distributionally robust optimization problem (11.10) is in fact tractable when the Wasserstein metric is used in the definition of  $\mathbb{B}_r(\hat{\mathbb{P}})$ . In this case no approximation is required.

The use of the Wasserstein metric is also attractive conceptually. Indeed, if (11.10) is a distributionally robust portfolio selection problem where  $\mathbb{B}_r(\hat{\mathbb{P}})$  represents the Wasserstein ball of radius  $r$  around some nominal distribution  $\hat{\mathbb{P}}$ , then the optimal solution of (11.10) converges to the equally-weighted portfolio as  $r$  tends to infinity (i.e., in the case of extreme ambiguity) [84]. This result, within a financial context, is appealing because it has been observed in empirical studies that the equally-weighted portfolio consistently outperforms many classical Markowitz-type portfolios in terms of Sharpe ratio, certainty-equivalent return or turnover [41]. The fact that the equally-weighted portfolio is optimal under extreme ambiguity therefore provides a solid theoretical justification for a surprising empirical observation.

The above concepts are now considered in the specific domain of a simple, though rather general specification of an asset-liability management problem.

## 11.4 Asset-Liability Management

In this section we describe a simple *asset-liability management* (ALM) problem, as an example of a stochastic control problem. This class of models subsumes asset management models where, for given market prices, we want to select the best investment policy under budget, turnover, inventory balance constraints, according to some utility function or a risk-reward tradeoff, over a given time horizon.

As well as liability management problems or liability-driven investment models [1] and genuine asset-liability models [10, 42, 81]. Asset management problems are relevant in fund management applications, active and passive portfolio management with or without index tracking [21, 58], allocation among liquid or illiquid asset classes including [78], and so forth. In full generality and in consideration of the contributions in this volume and the special issue, we formulate a generic pension fund problem. Here the pension fund manager seeks an optimal portfolio allocation in such a way as to meet a stream of (uncertain) liabilities. Alternatively for given payoff promised upon retirement by a certain retirement product, she/he might just search an optimal replication strategy [1, 54]. Such replication problem is equivalent to a hedging problem in financial engineering, and its solution depends on the assumptions made on market completeness. This class of problems can thus also be formulated as ALM problems. We provide next a set of examples canonical in finance to convey such generality.

### 11.4.1 An Overview of Financial Planning Problems

Different assumptions on the specification of problem (11.2) will result in different types of financial optimization problems and qualify most of the collected contributions.

*Example 1 (Portfolio Management).* The definition of an optimal portfolio strategy under a set of constraints and within different methodological assumptions is considered by Calafiore [25], Györfi and Ottucsak [58], Gilli and Schumann [55], Mulvey et al. [78] in this volume and by Pagliarani and Vargiolu [93], Desmettre et al. [44], Kopa and Post [67], Bruni et al. [21] and Cetinkaya and Thiele [32] in the special issue.

Here, as time  $t \in \mathcal{T}$  evolves, the state of the system  $\mathbf{s}_t$  is captured by the dynamics of portfolio losses or returns, the control  $\mathbf{x}_t$  reflects the dynamic portfolio rebalancing decisions within the feasibility region determined by  $A_t(\mathbf{s}_t)$ . Under parametric assumptions on the sample space or model driven approaches, the characterization of the probability measure  $\mathbb{P} \in \mathcal{P}$  is given and no min-max approach is needed. In the more general case of data-driven, model-free approaches or within an uncertainty set characterization, typical of robust formulations, we allow the minimum loss to be associated with a distributionally worst case characterization of the underlying stochastics. The problem collapses to one-period static optimization if  $T = 1$ : under which case, the optimality of a myopic policy will depend primarily on the statistical properties of the underlying return process. It is sufficient to allow for short cash positions or explicitly for borrowing choices that this turns into an ALM problem. Increasingly in the fund management industry the cost function includes an indexation scheme [21] resulting into an optimal tracking problem.

*Example 2 (Asset Pricing).* We reconsider here pricing problems, in the sense of asset pricing and derivative contracts valuation as well as the structuring of new financial payoffs, specifically from the market makers, protection seller, viewpoints. MacLean and Zhao [70], Aro and Pennanen [1], Giandomenico and Pinar [54], Mulvey et al. [78] in this volume, Koniecz et al. [66], Trigeorgis et al. [46] in the SI, provide different approaches to address such issue. The *fair* valuation of a contract can be conveniently regarded as the solution of an optimization problem leading today to a minimal investment cost of a portfolio that at the maturity of the contract will deliver the contract payoff, independently from the market condition at expiry. A characterizing condition of this problem is that the portfolio strategy leading to such payoff is actually self-financed.

For  $t \in \mathcal{T}$ , the state of the system  $\mathbf{s}_t$  will in this case be represented by say the derivatives' *moneyness*, the control  $\mathbf{x}_t$  will determine the dynamic portfolio hedging strategy with  $A_t(\mathbf{s}_t)$  to denote the self-financing condition and other technical constraints. The state transition operator  $\Phi_t$  will reflect the contract complexity and  $f_t$  the portfolio replication cost function to be minimized. The time space of the problem here may typically be described by a sufficiently thinned discretization. In general such hedging problem admits a unique solution under Black and Scholes assumptions, whereas in real markets and operational contexts hardly payoffs can be perfectly replicated giving raise to what is referred to as the cost induced by market incompleteness [1, 54]. As a by-product of such optimization approach, the probability measure governing the stochastic dynamics of the problem needs not be an input, rather an output of the dual problem solution. A rather interesting application is considered in Aro and Pennanen [1] and Pachamanova et al. [81] where retirement products are considered. Here assets are used to hedge a given liability whose payoff needs to be replicated. The stochastic program solution leads to both a minimal hedging cost trajectory and a fair value of the contract. The characterization of the underlying stochasticity is typically attained through some analytical description of the underlying risk process but ambiguity over models selection is accommodated in Trigeorgis et al. [46]. In rather general terms the numerical approximation of time and states needs to be carefully determined and this is an issue.

*Example 3 (Canonical Asset-Liability Management).* A decision maker is here confronted with an optimization problem involving both assets and liabilities: the universe of possible decisions including their dynamics over time is of primary concern to the decision maker. Historically ALM problems have gained popularity as natural modeling frameworks for the description of enterprise-wide long-term management problems, in which the dynamic interplay between asset returns and liability costs plays a central role in the definition of a strategic policy. Increasingly however and from what was said above, from a mathematical standpoint this is possibly the most general modeling framework available today. Individuals' consumption-investment problems [66] and pension fund ALM problems [1, 81] provide different perspectives falling in this application domain. Contributions by Pachamanova et al [81], Aro and Pennanen [1], Dempster et al. [42] in this volume, Dupacova and Kozmik [49], Koniecz et al. [66], Davis and Lleo [39] fall in this area.

Here  $t \in \mathcal{T}$ , is typically discrete and the state of the system  $\mathbf{s}_t$  will be represented by an A-L ratio, such as a funding or solvency ratio for an insurance company, the fund return for a Pension fund and alike, the control  $\mathbf{x}_t$  will include in the most general case asset and liability decisions to be taken over time.  $A_t(\mathbf{s}_t)$  a maybe extended set of conditions on the optimal policy, with regulatory constrains and limitations on the investment and borrowing policies. The chapter by Pachamanova et al. [81] as well as several contributions in [10] provide good examples of such set-up. In [36] an ALM problem for a large property and casualty (P/C) division is considered based on an extended set of both liquid as well as illiquid assets and liability streams associated with the P/C activity. The model provides a good example of the impact on classical ALM applications of the current integration of a capital adequacy regulatory framework.

### 11.4.2 A Simple ALM Model

Examples 1–3 in Sect. 11.4.1 may be given, with possibly a few additional constraints, a mathematical formulation similar to the one proposed in (11.11)–(11.16) and be interpreted as specific ALM instances. The area of pension fund management provides an interesting case where most of the above theoretical and modeling issues can be traced within a unified framework. Indeed, depending on the pension scheme, we have here relevant regulatory constraints, long-term asset and liabilities, short-term solvency and liquidity constraints, as well as financial engineering applications due to the growing role of *defined contribution* (DC) and mixed DC-*defined benefit* (DB) or *hybrid* schemes. From the perspective of the stochastic system evolution in this segment we also see how indeed policy makers, individuals and pension providers (public and private) interact to determine jointly the system dynamics and influencing its evolution. Konicz et al. [66] consider the individual perspective, Aro and Pennanen [1] as well as Pachamanova et al. [81] the Pension fund perspective.

The model is based on a scenario tree, the standard representation of uncertainty used in stochastic programming, as illustrated in Fig. 11.1. This should be contrasted against the linear scenario arrangement in Fig. 11.2. Consider the following variables specification under the introduced nodal labeling convention.

- $L^n$  are pension payments in node  $n \in \mathcal{N}$ ,
  - $\Lambda^n$  is the pension fund liability in node  $n \in \mathcal{N}$ : this is the discounted value of all future pension payments,
  - $c$  is a (symmetric) percentage transaction cost,
  - $\bar{h}_i^{n_0}$  is the initial holding for asset  $i = 1, \dots, I$  at the root node.
  - $r_i^n$  is the (price) return of asset  $i$  in node  $n$ .
  - $z_i^n$  is the amount of asset  $i$  purchased in  $n$ .
  - $y_i^n$  is the amount of asset  $i$  sold in  $n$ .
  - $x_i^n$  is the amount of asset  $i$  we hold at node  $n$ , after rebalancing. Accordingly
- $$X^n = \sum_{i=1}^I x_i^n$$

- $\phi^n$  is the pension fund's funding ratio in node  $n$ :  $\phi^n = \frac{X^n}{\Lambda^n}$
- $W^s$  is the pension fund terminal deficit  $W^s = \Lambda^s - X^s$   $s \in \mathcal{S}$ .
- $\rho(w)$  is a terminal functional on  $w$ . It is a risk measure that in (11.2) is considered in additive form while here is just applied to the terminal funding gap.

The pension fund manager seeks the minimization of a function of the fund's deficit:

$$\min_{x_i^n, y_i^n, z_i^n} \sum_{s \in \mathcal{S}} \pi^s \rho(W^s) \quad (11.11)$$

$$\text{s.t. } x_i^{n_0} = \bar{h}_i^{n_0} + z_i^{n_0} - y_i^{n_0}, \quad \forall i \quad (11.12)$$

$$x_i^n = x_i^{a(n)}(1 + r_i^n) + z_i^n - y_i^n, \quad \forall i, \forall n \in \mathcal{T} \quad (11.13)$$

$$(1 - c) \sum_{i=1}^I y_i^n - (1 + c) \sum_{i=1}^I z_i^n = L^n, \quad \forall n \in \mathcal{T} \quad (11.14)$$

$$X^n \geq \phi \Lambda^n \quad \forall n \in \mathcal{T} \quad (11.15)$$

$$W^s = \sum_{i=1}^I x_i^{a(s)}(1 + r_i^s) - L^s, \quad \forall s \in \mathcal{S} \quad (11.16)$$

$$x_i^n, z_i^n, y_i^n, W^s \geq 0, y_i^s = z_i^s = 0 \quad (11.17)$$

Equation (11.12) expresses the initial asset balance, taking the current holdings into account; the asset balance at intermediate trading dates is taken into account by Eq. (11.13). Equation (11.14) ensures that enough cash is generated by selling assets in order to meet current liabilities. Reinvestment at each stage is allowed until the start of the last stage. No buying or selling decisions are possible at the horizon where the pension fund surplus is computed. Upon selling or buying the pension fund manager faces transaction costs as indicated in the cash balance constraint (11.14). Equation (11.16) is used to evaluate terminal surplus at leaf nodes. Pension payments in this simple model are net of received contributions, which for this reason do not appear in the problem formulation.

In this rather simple ALM problem formulation, no cash is actually generated by current holdings and at each decision stage liabilities are funded through portfolio re-balancing decisions. Nor is it possible to borrow. This is typically the case when one faces an investment universe based on total return indices with no cash flows over the planning horizon. A minimal funding ratio  $\phi$  is allowed over the horizon. Such a constraint may reflect pending regulatory conditions in the market. With no borrowing and no cash account it is possible to satisfy all liabilities only by selling assets, and thus an issue of the fund's solvency actually arises. The problem resembles very much a fund manager's problem when issuing an annuity with given random cash flows (that may depend on exogenous elements) and resulting in an optimal self-financing portfolio strategy over a given horizon. From a financial standpoint this problem falls in the class of *liability-driven* investment (LDI) problems [1].

Formulation (11.11)–(11.16) is in general consistent with a random model for asset returns and pension payments. Liabilities are stochastic, and they depend on both future pension dynamics as well as on interest rate dynamics. Pension payments on the other hand will typically depend on future inflation and employees' careers. All these random elements must be considered when generating a scenario tree instance, and they will determine the computational effort in the problem solution [1, 42, 81]. The constraints are linear and assumed to be satisfied *almost surely*, as is common in stochastic programming models.

From a modeling viewpoint, a key distinction is between a (DB) pension scheme as the one above and a (DC) scheme: in the latter, pension payments and the accumulation of pension rights by an active member of the fund will depend on the fund's return year-by-year. Under such assumption, payments  $L^n$  will depend on the strategy to date inducing a non-linearity in the problem and a non-convex funding ratio control problem, with both assets and liabilities which are now decision-dependent. This is one of the challenges when modeling a *pure* DC pension problem. Under such scheme, furthermore, the pension fund manager does not carry any market risk unless a minimum guarantee is attached to the problem [81], in which case the problem would be associated with a DC *protected* scheme.

A continuously evolving market of retirement products includes hybrid pension schemes and life insurance contracts combining optionalities with classical annuities. The problem instance (11.11)–(11.16) may be adapted to the associated ALM problems. Let in particular  $\rho(W^s)$  represent the payoff at the terminal horizon of an annuity. Then the life insurer will seek the least expensive *replicating* portfolio to attain such payoff. The present cost of such portfolio, under self-financing conditions, will also provide the current fair value of the contract. As in [1, 54] this is the case of a pricing problem in incomplete markets. Market incompleteness implies the absence of a sufficient set of market instruments to perfectly replicate the contract payoff. A financial cost will arise whose minimization may be included in the objective function formulation.

We take the above pension fund management problem as sufficiently general specification to focus on modeling and solution issues specifically associated with the problem's dynamic formulation. We are primarily interested in the way in which the different concepts of risk are considered in this case, limiting ourselves to how the PF ALM problem is handled from a computational viewpoint. It will also provide a way to identify a set of open numerical and modeling issues. A discrete, multi-stage problem is considered and as specified from the beginning, we maintain a dynamic system risk control perspective. The following risk sources need to be taken into account:

- On the asset side, market risk factors affecting the price evolution of assets' total returns and their probability space characterization including the case of associated distributional uncertainty. Here we go from factors such as interest rates and credit spreads for fixed-income assets to risk premia for equity or autoregressive models for market indices or factor models in robust approaches and so forth [37, 39, 81].

- Liability risk also depends on factors such as interest rates and inflation, but more specifically on the future behavior of pension members survival intensities and resulting longevity risk [1] as well as on the contributors to pensioners ratios, whose dynamics may be difficult to capture.
- Joint A-L risks focus on the correlation between the above risk factors: a sudden divergence of asset values from liability values resulting into an unexpected fall of the fund's solvency ratio may compromise from 1 day to the next a market's systemic equilibrium and for such reason is constantly monitored by regulatory institutions. It is very much in the actuarial tradition to focus on asset-liability duration mismatching, and this is easily accommodated within a dynamic control problem such as (11.11)–(11.16).
- Increasingly however all the above elements are captured within a critical modeling approach in which model risk and thus ambiguity issues are properly considered [46, 68]. Such effort has been motivated by both the serious deficiencies reported by canonical modeling approaches in the recent 2007–2010 global financial crisis and by the emerging possibility to achieve a more effective overall risk control with a far less expensive computational effort.

In Sect. 11.5 we relate the above points to those methodological approaches currently attracting most interest also in the fund management industry and the only ones able to handle effectively a dynamic management problem: dynamic stochastic programming (DSP) and distributionally robust optimization (DRO) based on decision rules. The reference ALM problem is the one introduced above in this Sect. 11.4.2.

In the final Sect. 11.6 we summarize a set of open modeling issues as well as desirable developments aimed at facilitating the practical adoption of dynamic approaches now-a-days still limited to few, even if rather relevant industry cases.

## 11.5 Solution Methods and Decision Support

The ALM problem formulation (11.11)–(11.16) is preliminary to the adoption of one or another solution approach. In an operational context, by solving the problem, a pension fund manager aims at identifying an effective strategic asset allocation able to preserve the fund's solvency, payout all liabilities and hedge appropriately all asset and liability risks. In a liability-driven approach she/he would be primarily concerned with a cost-effective liability replication and the generation of an excess return through active asset management. From a methodological viewpoint, depending on the adopted optimization approach, Eqs. (11.11)–(11.16) will be subject to further refinements. Increasingly in market practice the solution of such an optimization problem is embedded in the definition of a decision support system [36] whose key modules are represented by a user interface, a problem instance generator, a solution algorithm and a set of output analyses aimed at supporting the decision process. Hardly any such solution will translate straight-away into an

actual decision: it will rather provide one of the few fundamental informations for that. Here below, without going into the complexities associated with the definition of a decision support tool, we focus specifically on key methodological issues. The four basic *building blocks* of a quantitative ALM approach—problem formulation, stochastics, solution and decisions—must be kept in mind [42].

A DSP approach preserves all four elements, puts a clear emphasis on each such block and however, specifically with respect to *stochastics* remains subject to model risk. Increasingly, validation through appropriate statistical measures and out-of-sample backtesting of stochastic models and decisions is part of the process. A robust approach, instead aims at avoiding model risk by incorporating the stochastics within a reformulation of the original optimization problem. Either approach is aimed at generating an effective decision process and, in a dynamic set-up key to a consistent problem solution, it focuses on the interaction between decisions and sequential revelation of uncertainty as expressed by the formulation (11.2)–(11.4).

### 11.5.1 Stochastic Programming

Consider the optimization problem introduced in Sect. 11.4.2. Let  $\mathbf{x} \in \mathcal{X}$  be the control vector or portfolio allocation process and  $[\rho_T(\cdot)|\Sigma]$  be a terminal risk measure evaluated at the end of the planning horizon, conditional on the information  $\Sigma$ . The focus of the pension fund manager is on the risk associated with possibly deteriorating funding conditions as captured by the *funding ratio*  $\frac{X_n}{\Lambda_n}$  along the scenario tree. All random elements of the problem enter the definition of the coefficient tree process  $\xi(\omega)$  defined on a probability space  $(\Omega, \Sigma, \mathbb{P})$ . The specification of a dynamic risk measures  $\rho : \mathcal{X} \times \Omega \mapsto \mathbb{R}$  as well as the specification of conditional expectations  $\mathbb{E}[\cdot|\Sigma_t]$  for increasing  $t$  implies the nesting of the associated functionals [34]. Any decision is taken relying on current information and facing a residual uncertainty, as captured by the subtree originating from that node. Under Markovian assumptions on the decision problem and convexity of the risk measure  $\rho$  with respect to the decisions, it is possible to solve the problem recursively relying on a sequence of nested dynamic programs or Benders decomposition [42, 49]. In the tree representation there is always a one-to-one relationship between each state of the process—a node in the tree—and its history: the path from the root to the current node.

The decision space  $\mathcal{X}$  may accommodate an extended set of asset and liability classes, accordingly it will increase the dimension of  $\xi$ . The specification of the dependence of  $\xi$  on  $\omega$ —the stochastics—depends on the adopted statistical assumptions. Specifically in financial management and pricing applications, the tree structure is constrained to carry a sufficient branching structure to accommodate equilibrium conditions such as arbitrage-free conditions [54, 64] leading to the well-known *curse-of-dimensionality* problem. An approximation issue arises also considering the relationship between  $\xi$  and  $\omega$  and their underlying, maybe, continuous time companions.



Under this framework the scenario tree structure [37, 54] will determine both the stochastic *input* information for the optimization problem and the optimal decision process or contingency plan *output* [42]. Along a given scenario, carrying a specific probability of occurrence, the pension fund manager may collect a remarkable amount of information, in terms of:

- asset and liability returns, pension payments and contribution flows, funding ratio and funding gap dynamics,
- optimal investment portfolio rebalancing to replicate a given liability,
- liquidity shortages and need of extraordinary sponsors' contributions,
- evolving risk exposure as determined by inflation or interest rates,
- sensitivity to exogenous market or liability shocks and so forth.

The optimal decision process will provide the best possible strategy in the face of the uncertainty captured by the scenario tree process. The discrete framework enables an extended and thorough analysis of all relationships within the given ALM model specification and those relationships may be analyzed and validated from a quantitative as well as qualitative viewpoint by the decision maker [36].

Once the model risks related to a maybe rough analytical specification of the underlying risk process  $\omega$  and derivation of  $\xi(\omega)$  have been identified, several solution algorithms may be adopted on decomposed problems or deterministic equivalent instances:

- Under assumptions of stage-wise independence of the coefficient process, [49] present an application of *stochastic dual dynamic programming* (SDDP) to solve a set of sample average approximations of multiperiod risk measure minimization problems such as a multiperiod CVaR model [34].
- Under sufficiently general assumptions on the underlying process and convex constraints [36] employ CPLEX's quadratically constrained quadratic programming solver, on its own and combined with a conic solver, on the *deterministic equivalent* of an insurance ALM problem.
- If the objective function is separable and displays a nested structure with a Markovian constraint region, nested Benders decomposition may still provide an efficient solution approach [43].
- Within a rather general modeling framework for a pension planning problem with power utility, [66] propose a combined stochastic programming and optimal control approach. As in [36] the authors rely here on GAMS and Matlab as algebraic language and modeling tool and on MOSEK as conic solver to handle the non-linearities.

Under different assumptions on the problem specification, the above may be employed to generate an optimal decision tree process from a given input information. The mixed approach proposed in [66] assumes a partition of the decision horizon between a first horizon for the multistage stochastic program and a second horizon for the optimal control: in this second period an optimal HJB approach is adopted relying on a set of constraint relaxations.

A deterministic equivalent instance of the PF ALM problem is commonly derived for convex programming problems of large dimensions without sacrificing any modeling detail but facing the complexities associated with effective sampling methods [37, 42] and in any case the costs that may be induced by a statistical model misspecification and high sensitivity of the optimal strategy to the input coefficient tree process. Still for an ALM problem, departing from a MSP formulation and related solution, [78] suggest a further direction to minimize the costs associated with the tree approximation, through the introduction of *policy rules*: those are specified in coherence with financial management practice and *evaluated* by simulation over maybe complex investment universes. The adoption of Monte Carlo simulation prior to optimization, for scenario construction, and after for policy rules evaluation provides a way to control the approximation costs associated with a scenario-based optimization.

The estimation of the in-sample and out-of-sample stability of a stochastic program also goes in the same direction [42]. We have in-sample stability when relying on a given sampling method we generate and solve an instance of the optimization problem and then repeating such process several times, for different scenario tree instances we report a relatively stable optimal value of the problem maybe generated by different optimal implementable decisions. Lack of this type of stability (which is easy to detect) would result for a given statistical model and sampling method that different problem instances would result into different objective optimal values and controls. Out-of-sample stability implies that for different scenario trees and associated optimal decisions, given each such decision we employ a procedure to evaluate what would be the solution of the optimization problem when reducing asymptotically the sampling error (increasing the number of scenarios): if different input solution vectors result into similar *true* objective values we will have out-of-sample stability.

Overall the adoption of a combined set of optimization and simulation techniques appears highly desirable in presence of financial planning problems, such as a pension fund ALM problem, specified over long horizons (up to 30 years) and with a long experience on simulation techniques for risk assessment but not yet fully accustomed to dynamic optimization approaches.

### 11.5.2 *Dynamic Optimization Via Decision Rules*

The ALM model introduced in Sect. 11.4.2 can be viewed as a scenario-tree approximation of a fundamental control model of the type (11.2)–(11.4). Instead of using a scenario-tree approximation, however, such control models can be rendered tractable by using a decision rule approximation. As decision rules are most frequently used in robust optimization, we will explain the mechanics of the decision rule approximation on the example of the worst-case optimization problem (11.6).

To capture the flow of information mathematically, we assume that the components of the uncertain parameter vector  $\xi$  are revealed sequentially as time

progresses. As in multistage stochastic programming, future decisions are then modeled as functions of the observable data. For ease of exposition, assume that *all* decisions must be chosen after a linear transformation  $P\xi$  of the uncertain parameter  $\xi$  (e.g., a projection on some of the components of  $\xi$ ) has been observed. Mathematically, this means that the decision must be modeled as an element of the function space  $\mathcal{L}^0(P\Xi; \mathcal{X})$ , which contains all measurable functions mapping  $P\Xi$  to  $\mathcal{X}$ . In other words, the decision becomes a function that assigns to each possible data observation  $P\xi \in P\Xi$  a feasible action  $\mathbf{x}(P\xi) \in \mathcal{X}$ . In this situation, the robust program (11.6) becomes an infinite-dimensional functional optimization problem of the form

$$\min_{\mathbf{x} \in \mathcal{L}^0(P\Xi; \mathcal{X})} \max_{\xi \in \mathcal{U}} f(\mathbf{x}(P\xi), \xi). \quad (11.18)$$

The dynamic versions of the stochastic program (11.5) and the distributionally robust model (11.7) are constructed similarly in the obvious way. All of these models easily extend to more general decision-making situations where different transformations  $P_1\xi, \dots, P_T\xi$  of the uncertain parameter  $\xi$  are observed at times  $1, \dots, T$  and where  $\mathbf{x}$  admits a temporal decomposition of the form  $(\mathbf{x}_1, \dots, \mathbf{x}_T)$ , with  $\mathbf{x}_t$  capturing the subvector of the decisions that must be taken at time  $t$ , respectively. This is exactly what happens in the ALM model of Sect. 11.4.2, where  $P_t\xi$  coincides with the collection of all market and liability risk factors revealed at time  $t$ . To keep the presentation simple, however, we will focus on the generic model (11.6) in the remainder.

Dynamic models are often far beyond the reach of analytical methods or classical numerical techniques plagued by the notorious curse of dimensionality. Rigorous complexity results indicate that, for fundamental reasons, dynamic models need to be approximated in order to become computationally tractable [90]. For instance, one could approximate the adaptive model (11.18) by the corresponding static robust optimization problem (11.6). This approximation is conservative—in the sense that it artificially limits the decision maker’s flexibility and therefore leads to an upper bound on the optimal value of (11.18). While this approximation may seem overly crude, it has been shown to yield nearly optimal portfolios in various multi-period asset allocation problems and has distinct computational advantages over scenario-tree-based stochastic programming models [6]. Static robust formulations of the multiperiod portfolio selection problem with transaction costs are also considered in [11], where it is shown empirically that robust polyhedral optimization can enhance the performance of single period and deterministic multiperiod portfolio optimization methods.

Instead of naively approximating adaptive by static models, one can alternatively impose a linear structure on the recourse decisions, that is, one may approximate the measurable function  $\mathbf{x}(P\xi)$  by a linear function  $XP\xi$ , which is completely determined by the matrix  $X$ . The sensitivity matrix  $X$  determines the rate of change of each scalar decision with respect to changes in the uncertain parameters. Under this *linear decision rule approximation*, the adaptive robust optimization

problem (11.18) reduces to

$$\begin{aligned} \min_X \max_{\xi \in \mathcal{U}} f(XP\xi, \xi) \\ \text{s.t. } XP\xi \in \mathcal{X} \quad \forall \xi \in \mathcal{U}. \end{aligned} \quad (11.19)$$

Linear decision rules have attracted considerable interest in recent years because they inherit the favorable scalability properties of static robust optimization models. Indeed, (11.19) is formally equivalent to a static robust optimization problem where the sensitivity matrix  $X$  represents the static (here-and-now) decision.

The linear decision rule approximation has been popularized with the seminal paper [7]. Fuelled by their success in robust optimization, linear decision rules have also gained renewed interest from the stochastic programming community; see, e.g., [90]. To improve the approximation quality, several authors have proposed more flexible *non-linear* decision rules such as deflected and segregated decision rules, which can be interpreted as linear decision rules on an augmented probability space and therefore display favorable scalability properties [33, 56]. An efficient procedure to quantify the degree of suboptimality of linear and non-linear decision rules has been proposed in [68] and [53], respectively.

Linear decision rules have been used with great success in dynamic asset allocation with transaction costs [22, 23]. They have also been used in the context of portfolio execution, where they were shown to achieve near optimal performance [74]. Both linear and non-linear decision rules lend themselves ideally to integration with constraint sampling techniques from data-driven optimization [94]. Interestingly, there is ample evidence suggesting that simple decision rules with few degrees of freedom may systematically outperform more versatile decision rules which are prone to overfitting and error-maximization phenomena. For instance, in [19] the portfolio weight of each stock is modeled as a function of the firm's market capitalization, book-to-market ratio, or lagged return. As in data-driven optimization, the coefficients of this simple portfolio rule are found by maximizing the average utility an investor would have obtained by implementing the policy over a historical sample period. This approach is simple to implement and produces outstanding results in and out-of-sample experiments. Similarly, attractive out-of-sample results for an index tracking application were obtained with the robust data-driven dynamic programming approach proposed in [59].

## 11.6 Open Issues

We have considered a rich set of modeling and theoretical issues from an extended set of contributions to clarify key diverging points that may emerge when similar financial optimisation problems are addressed relying on different methodological assumptions. Without going into a rigorous treatment, the generality of ALM approaches to accommodate a wide range of financial management problems, from

portfolio selection to derivatives pricing in incomplete markets, has been emphasised. Specifically in relation with a pension fund ALM problem, we analysed the decision processes as a sequence of steps, each one carrying its own modeling and methodological complexity, but jointly leading to the definition of quantitatively-based optimal decisions. Indeed we have also remarked that efficient decision making does not necessarily require excessive modeling complexity and a *friendly* approach to computational management, based on simple probabilistic assumptions, may be sufficient. By proposing a unified modeling framework, our interest was primarily on the potentials offered by the integration within a given method of results and advances may be achieved under other modeling and optimization paradigms [38].

Without ambition to be exhaustive in this respect, in what follows we would like to indicate a set of relevant open issues which may emerge from the introduced financial domains but also carrying implications to other application areas. We focus on the two fundamental elements of decision making under uncertainty: stochastics as featured by  $\xi(\omega)$  and solutions  $\mathbf{x} \in \mathcal{X}$  in (11.2)–(11.4).

### 11.6.1 Probability Distributions and Optimization

A compact formulation of a stochastic program under distributional uncertainty has been introduced as

$$\min_{\mathbf{x} \in \mathcal{X}} \max_{\mathbb{P} \in \mathcal{P}} \mathbb{E}_{\mathbb{P}}[f(\mathbf{x}, \xi)]. \quad (11.20)$$

If the ambiguity set  $\mathcal{P}$  is given by a singleton  $\{\mathbb{P}\}$ , we remain within a canonical stochastic programming framework: here under maybe an extremely challenging computational environment, relying on the discrete scenario tree representation key issues such as *arbitrage free* market conditions, market *incompleteness*, accurate stress-testing and sensitivity analysis (along a set of scenarios), as well as a rich sequence of I/O analyses can be conducted in a sufficiently straight-forward way. The optimal risk control  $x$  takes the form of an optimal tree decision process and contingency plan. This is done at an approximation cost and somehow aware of the potential damages that might be induced by estimation errors in parametric models or the adoption of generating processes for  $\xi$  that may be contradicted by market dynamics. Such risks would persist under in-sample and out-of-sample stable programs.

For general ambiguity sets that may contain more than one distribution, the worst-case expectation with respect to  $\mathcal{P}$  corresponds to a coherent risk measure, that is,  $\max_{\mathbb{P} \in \mathcal{P}} \mathbb{E}_{\mathbb{P}}[f(\mathbf{x}, \xi)]$  can be viewed as the *risk* of the random loss  $f(\mathbf{x}, \xi)$  under the coherent risk measure induced by  $\mathcal{P}$  [2]. In a distributionally robust setting, arbitrage-freeness of the market is usually no serious concern. In fact, the existence of arbitrage portfolios is ruled out if the underlying financial market is

arbitrage-free under at least one distribution  $\mathbb{P} \in \mathcal{P}$ ; and this is generically the case if  $\mathcal{P}$  is defined through moment constraints or as a ball in the space of probability distributions. To our best knowledge, market incompleteness has not yet been studied systematically in a distributionally robust framework. However, problem (11.20) lends itself to various stress-tests and sensitivity experiments. For example, it is well-known that the extremal distribution achieving the worst-case expectation in (11.20) for a fixed decision  $\mathbf{x}$  can be computed systematically and efficiently for most commonly used ambiguity sets; see e.g. [86] for an explicit example. This worst-case distribution can be used to conduct contamination or stress-test experiments. Specifically, the performance of a fixed decision  $\mathbf{x}$  can be evaluated under a nominal distribution contaminated with an extremal distribution. The sensitivity of the resulting out-of-sample performance with respect to the contamination parameter provides valuable information about the robustness of the decision  $\mathbf{x}$ . Another striking benefit of the distributionally robust model (11.20) is that it can sometimes be solved more efficiently than the corresponding stochastic program with a nominal distribution [40].

In principle data-driven approaches may be used to construct  $\mathcal{P}$  [14, 40]. It is also possible to use the *empirical* distribution directly in the construction of the ambiguity set; see e.g. [75]. Serial independence of the training data may be needed in this case, however, to establish rigorous out-of-sample performance guarantees. But such an assumption may not be justified in reality. An interesting avenue for future research is to develop out-of-sample guarantees for serially dependent training data; e.g., for data generated by a Markov or an autoregressive process.

### 11.6.2 Dynamic Time Consistency

The concept of time-consistent dynamic risk measures [34, 49] is key to the definition of a dynamic risk control problem and may be analysed with respect to the evolution of the process  $\xi(\omega)$  as well as to the optimal control  $\mathbf{x}$  [34]: in this latter case loosely speaking time consistency implies that as time evolves  $t = 0, 1, 2, \dots, T$  by solving a sequence of stochastic control problems leading to optimal controls  $\{x\}_{s=t}^T$  optimality would be preserved throughout for each  $t$ . In a discrete setting for scenario tree processes such condition should be considered conditionally with respect to actual realizations. As for time-consistency of a risk measure with respect to  $\xi(\omega)$  this would imply the persistence as time evolves of previously assessed risk rankings conditionally on the sequence of associated  $\sigma$ -fields. In this case we have a form of monotonicity since we require that as information accumulates, successive evaluations of a random process by the risk measure remain homogenous.

In either cases, the evolution of information plays a central role: this is captured for risk measures by the evolution of the process filtration, while in optimal control problems as in (11.2)–(11.4) the key information would be associated with the

evolution of the state variable  $\mathbf{s}_t$ . In [34] we discuss the conditions to connect the two concepts within a recursive approach based on dynamic programming principles. Recursivity of risk measures is not always attainable [34].

In [49] conditions for time-consistency of an optimization problem based on nested CVaR measures is studied, leading to a solution approach based on SDDP.

In this context, the concept of *cost* of time-inconsistency appears key to the discussion as introduced in [34] and may lead to an effective evaluation of its implications on the optimization problem. In a discrete scenario tree setting the discussion should be specialised with respect to conditional risk measures. In this direction, across the different dynamic optimization paradigms, research is continuing.

### 11.6.3 Practical Financial Optimization

The sheer amount of literature on financial optimization shows a huge academic interest. From the practical point of view, there are in fact some interesting real-life applications documented [3, 13, 29, 36, 77], but it may be argued that the full potential has not been achieved, yet, at least as far as dynamic portfolio optimization is concerned. What are the main issues potentially hindering the application of financial optimization models in practice? Computational efficiency is certainly a concern in multistage stochastic programming models, but it is arguably not the main issue in other cases. One question we should ask is what would be the main sources of concern for a practitioner considering the use of an optimization model as a decision support tool:

- How will I be held accountable for the decision I am going to make?
- Are there any misalignments between the model's objective and how my performance is assessed?
- How can I trust the model and understand the reasons behind its recommendations?
- How can I compare the model output against alternatives?
- How can I test the robustness of the model's recommendation?
- How can I incorporate my own views?

Some of these concerns, being actually related to a *principal-agent* type of problem, may look nontechnical, but they should be addressed somehow in practical optimization. There is little point in considering model objectives that not aligned with the actual incentives of the decision maker. For instance, it is a nontrivial task to align long-term objectives with potential short-term performance evaluation affecting a portfolio manager. Furthermore, in statistical learning there is a well-known trade-off between model's sophistication, which may contribute to its predictive power, and its interpretability [61]. Relying on the output of a sophisticated model may be hard for a manager, if she gets a supposed optimal solution, but it is difficult to understand why that is the optimal solution. Furthermore, sophisticated models

may be fragile. They are subject to overfitting issues in statistical learning, and regularized approaches (like lasso regression) may be helpful. We note that the contribution of [42] has a similar flavor. By a similar token, as suggested in [16], under significant uncertainty simpler rules may be preferable [17]. This does not imply that sophisticated models must be given up, but assessing their performance under stress is certainly no easy task. As suggested in [76], when model risk is an issue, even stress-testing per se may be hard to grasp under certain asset pricing models.

Recent trends in decision rules for dynamic optimization offer solutions that may be easily tested under different scenarios, without requiring extensive rolling horizon simulations, which are computationally demanding. Approximate dynamic programming offers similar benefits for a skeptical manager. When choosing a modeling framework, these considerations may be relevant, and in this way trust may be built, possibly leading to an easier adoption and, maybe, the development of more refined approaches over time.

Personal views are relevant to a portfolio manager, and the increasing role of Bayesian-learning approaches in finance and risk management is well documented [85]. The Black–Litterman model is a well-known approach that incorporates Bayesian principles [87]. They are also relevant for risk management [88], as they provide one possible answer to model risk issues. Model robustness and ambiguity are the subject of active research, and an integration of available approaches is called for from a practitioner's point of view, with an emphasis of proof-of-concept methods to evaluate actual out-of-sample performance under stressed market conditions.

## 11.7 Conclusions

The variety of contributions in this volume and the companion SI shows that, after several years, financial optimization is still a field in flux and a very active research field. New challenges have emerged from the financial side, and new algorithmic frameworks have been developed to address issues in previous models and solution methods. Integration opportunities between solution methods, like stochastic and robust optimization, or dynamic programming and stochastic programming with recourse, are rather evident, as well as the interplay of different financial problems like pricing, hedging, and asset management. We did not cover every possible application of optimization modeling to finance; for instance, we did not consider pricing model calibration and optimal order execution. These and other applications testify the interest and the practical relevance of this research field.



## References

1. H. Aro, T. Pennanen, Liability driven investment in longevity risk management. This volume
2. P. Artzner, F. Delbaen, J.-M. Eber, D. Heath, Coherent measures of risk. *Math. Financ.* **9**, 203–228 (1999)
3. M. Avriel, H. Pri-Zan, R. Meiri, A. Peretz, Opti-money at Bank Hapoalim: a model-based investment decision-support system for individual customers *Interfaces* **34**(1), 39–50 (2004)
4. C. Bandi, D. Bertsimas, Tractable stochastic analysis in high dimensions via robust optimization. *Math. Program.* **134**(1), 23–70 (2012)
5. C. Bandi, D. Bertsimas, Robust option pricing. *Eur. J. Oper. Res.* **239**(3), 842–853 (2014)
6. A. Ben-Tal, T. Margalit, A. Nemirovski, Robust modeling of multi-stage portfolio problems, in *High Performance Optimization*, ed. by H. Frenk, K. Roos, T. Terlaky, S. Zhang (Springer US, New York, 2000), pp. 303–328
7. A. Ben-Tal, A. Goryashko, E. Guslitzer, A. Nemirovski, Adjustable robust solutions of uncertain linear programs. *Math. Program.* **99**(2), 351–376 (2004)
8. A. Ben-Tal, S. Boyd, A. Nemirovski, Extending scope of robust optimization: comprehensive robust counterparts of uncertain problems. *Math. Program.* **107**(2), 63–89 (2006)
9. A. Ben-Tal, L. El Ghaoui, A. Nemirovski, *Robust Optimization* (Princeton University Press, Princeton, 2009)
10. M. Bertocchi, G. Consigli, M.A.H. Dempster (eds.), *Stochastic Optimization Methods in Finance and Energy. New Financial Products and Energy Market Strategies*. Fred Hillier International Series in Operations Research and Management Science (Springer US, New York, 2011)
11. D. Bertsimas, D. Pachamanova, Robust multiperiod portfolio management in the presence of transaction costs. *Comput. Oper. Res.* **35**(1), 3–17 (2008)
12. D. Bertsimas, M. Sim, The price of robustness. *Oper. Res.* **52**(1), 35–53 (2004)
13. D. Bertsimas, C. Darnell, R. Soucy, Portfolio construction through mixed-integer programming at Grantham, Mayo, Van Otterloo and company. *Interfaces* **29**(1), 49–66 (1999)
14. D. Bertsimas, V. Gupta, N. Kallus, Data-driven robust optimization (2013). Available from arXiv:1401.0212
15. J.R. Birge, F. Louveaux, *Introduction to Stochastic Programming*, 2nd edn. Springer Series in Operations Research and Financial Engineering (Springer, New York, 2011)
16. R. Bookstaber, *A Demon of Our Own Design* (Wiley, New York, 2007)
17. R. Bookstaber, J. Langsam, On the optimality of coarse behavior rules. *J. Theor. Biol.* **116**(2), 161–193 (1985)
18. P. Brandimarte, *Handbook in Monte Carlo Simulation: Applications in Financial Engineering, Risk Management, and Economics* (Wiley, New York, 2014)
19. M.W. Brandt, P. Santa-Clara, R. Valkanov, Parametric portfolio policies: exploiting characteristics in the cross-section of equity returns. *Rev. Financ. Stud.* **22**(9), 3411–3447 (2009)
20. M. Broadie, Computing efficient frontiers using estimated parameters. *Ann. Oper. Res.* **45**(1), 21–58 (1993)
21. R. Bruni, F. Cesarone, A. Scozzari, F. Tardella, A linear risk-return model for enhanced indexation in portfolio optimization. *OR Spectr.* **37**(3), 735–759 (2015)
22. G.C. Calafiore, Multi-period portfolio optimization with linear control policies. *Automatica* **44**(10), 2463–2473 (2008)
23. G.C. Calafiore, An affine control method for optimal dynamic asset allocation with transaction costs. *SIAM J. Control. Optim.* **48**(4), 2254–2274 (2009)
24. G.C. Calafiore, Random convex programs. *SIAM J. Optim.* **20**(6), 3427–3464 (2010)
25. G.C. Calafiore, Scenario optimization methods in portfolio analysis and design. This volume
26. G.C. Calafiore, M.C. Campi, The scenario approach to robust control design. *IEEE Trans. Autom. Control* **51**(5), 742–753 (2006)
27. J.Y. Campbell, L.M. Viceira, *Strategic Asset Allocation* (Oxford University Press, Oxford, 2002)

28. M.C. Campi, S. Garatti, The exact feasibility of randomized solutions of uncertain convex programs. *SIAM J. Optim.* **19**(3), 1211–1230 (2008)
29. D.R. Cariño, T. Kent, D.H. Myers, C. Stacy, M. Sylvanus, A.L. Turner, K. Watanabe, W.T. Ziemba, The Russell-Yasuda Kasai Model: an asset/liability model for a Japanese insurance company using multistage stochastic programming. *Interfaces* **24**(1), 29–49 (1994)
30. S. Ceria, R. Stubbs, Incorporating estimation errors into portfolio selection: Robust portfolio construction. *J. Asset Manag.* **7**(2), 109–127 (2006)
31. R. Cerqueti, P. Falbo, G. Guastaroba, C. Pelizzari, Approximating multivariate Markov chains for bootstrapping through contiguous partitions. *OR Spectr.* **37**(3), 803–841 (2015)
32. E. Çetinkaya, A. Thiele, Data-driven portfolio management with quantile constraints. *OR Spectr.* **37**(3), 761–786 (2015)
33. X. Chen, M. Sim, P. Sun, A robust optimization perspective on stochastic programming. *Oper. Res.* **55**(6), 1058–1071 (2007)
34. Z. Chen, G. Consigli, J. Liu, G. Li, T. Fu, Q. Hu, Multi-period risk measures and dynamic risk control. This volume
35. V.K. Chopra, W.T. Ziemba, The effect of errors in means, variances and covariances on optimal portfolio choice. *J. Portf. Manag.* **19**(2), 6–11 (1993)
36. G. Consigli, V. Moriggia, Applying stochastic programming to insurance portfolios stress-testing. *Quant. Financ. Lett.* **2**(1), 7–13 (2014)
37. G. Consigli, G. Iaquina, V. Moriggia, Path-dependent scenario trees for multistage stochastic programs in finance. *Quant. Financ.* **12**(8), 1265–1281 (2012)
38. G. Consigli, P. Brandimarte, D. Kuhn, Editorial: financial optimization: optimization paradigms and financial planning under uncertainty. *OR Spectr.* **37**(3), 553–557 (2015)
39. M.H.A. Davis, S. Lleo, Jump-diffusion asset-liability management via risk-sensitive control. *OR Spectr.* **37**(3), 655–675 (2015)
40. E. Delage, Y. Ye, Distributionally robust optimization under moment uncertainty with application to data-driven problems. *Oper. Res.* **58**(3), 595–612 (2010)
41. V. DeMiguel, L. Garlappi, R. Uppal, Optimal versus naive diversification: how inefficient is the 1/N portfolio strategy? *Rev. Financ. Stud.* **22**(5), 1915–1953 (2009)
42. M.A.H. Dempster, E.A. Medova, Y.S. Yong, Stabilizing implementable decisions in dynamic stochastic programming. This volume
43. M.A.H. Dempster, M. Germano, E.A. Medova, M. Villaverde, Global asset liability management. *Br. Actuar. J.* **9**(1), 137–195 (2003)
44. S. Desmettre, R. Korn, P. Ruckdeschel, F.T. Seifried, Robust worst-case optimal investment. *OR Spectr.* **37**(3), 677–701 (2015)
45. X. Doan, X. Li, K. Natarajan, Robustness to dependency in portfolio optimization using overlapping marginals. Available on Optimization Online (2013)
46. T. Driouchi, L. Trigeorgis, Y. Gao, Choquet-based European option pricing with stochastic (and fixed) strikes. *OR Spectr.* **37**(3), 787–802 (2015)
47. J. Dupačová, The minimax approach to stochastic programming and an illustrative application. *Stochastics* **20**, 73–88 (1987)
48. J. Dupacova, G. Consigli, S.W. Wallace, Scenarios for multistage stochastic programmes. *Ann. Oper. Res.* **100**, 25–53 (2001)
49. J. Dupačová, V. Kozmík, Structure of risk-averse multistage stochastic programs. *OR Spectr.* **37**(3), 559–582 (2015)
50. L. El Ghaoui, M. Oks, F. Outstry, Worst-case value-at-risk and robust portfolio optimization: a conic programming approach. *Oper. Res.* **51**(4), 543–556 (2003)
51. E. Erdoğan, G. Iyengar, Ambiguous chance constrained problems and robust optimization. *Math. Program. B* **107**(1–2), 37–61 (2006)
52. D. Fabozzi, P. Kolm, D. Pachamanova, *Robust Portfolio Optimization and Management* (Wiley, New York, 2007)
53. A. Georghiou, W. Wiesemann, D. Kuhn, Generalized decision rule approximations for stochastic programming via liftings. *Math. Program.* **152**(1–2), 301–338 (2015)

54. M. Giandomenico, M. Pinar, Pricing American options: binomial and trinomial non-recombinant trees and two exercise rights. This volume
55. M. Gilli, E. Schumann, Constructing portfolios with heuristics. This volume
56. J. Goh, M. Sim, Distributionally robust optimization and its tractable approximations. *Oper. Res.* **58**(4), 902–917 (2010)
57. D. Goldfarb, G. Iyengar, Robust portfolio selection problems. *Math. Oper. Res.* **28**(1), 1–38 (2003)
58. L. Györfi, G. Ottucsak, H. Walk, The growth optimal strategy is secure too. This volume
59. G.A. Hanasusanto, D. Kuhn, Robust data-driven dynamic programming, in *Advances in Neural Information Processing Systems*, pp. 827–835 (2013)
60. Z. Hu, L.J. Hong, Kullback-Leibler divergence constrained distributionally robust optimization. Available from Optimization Online (2013)
61. G. James, D. Witten, T. Hastie, R. Tibshirani, *An Introduction to Statistical Learning: with Applications in R* (Springer, Berlin, 2013)
62. B. Kavas, A. Thiele, A log-robust optimization approach to portfolio management. *OR Spectr.* **33**(1), 207–233 (2011)
63. J.M. Keynes, *A Treatise on Probability* (MacMillan, London, 1921)
64. A.J. King, Duality and martingales: a stochastic programming perspective on contingent claims. *Math. Program. B* **91**, 543–562 (2002)
65. F.H. Knight, *Risk, Uncertainty and Profit* (Hart, Schaffner and Marx, Boston, 1921)
66. A.K. Koniecz, D. Pisinger, K.M. Rasmussen, A combined stochastic programming and optimal control approach to personal finance and pensions. *OR Spectr.* **37**(3), 583–616 (2015)
67. M. Kopa, T. Post, A general test for SSD portfolio efficiency. *OR Spectr.* **37**(3), 703–734 (2015)
68. D. Kuhn, W. Wiesemann, A. Georghiou, Primal and dual linear decision rules in stochastic and robust optimization. *Math. Program.* **130**(1), 177–209 (2011)
69. D.G. Luenberger, *Investment Science* (Oxford University Press, New York, 1998)
70. L.C. MacLean, Y. Zhao, Asset price dynamics: shocks and regimes. This volume
71. H.M. Markowitz, *Portfolio Selection: Efficient Diversification of Investments*, 2nd edn. (Wiley, New York, 1991)
72. R.C. Merton, On estimating the expected return on the market: an exploratory investigation. *J. Financ. Econ.* **8**(4), 323–361 (1980)
73. R.O. Michaud, *Efficient Asset Management: A Practical Guide to Stock Portfolio Management and Asset Allocation* (Oxford University Press, Oxford, 2001)
74. C.C. Moallemi, M. Saglam, Dynamic portfolio choice with linear rebalancing rules (2015). Available at SSRN: <http://ssrn.com/abstract=2011605>
75. P. Mohajerin Esfahani, D. Kuhn, Data-driven distributionally robust optimization using the Wasserstein metric: performance guarantees and tractable reformulations. Available from Optimization Online (2015)
76. M. Morini, *Understanding and Managing Model Risk: A Practical Guide for Quants, Traders and Validators* (Wiley, New York, 2011)
77. J.M. Mulvey, G. Gould, C. Morgan, An asset and liability management system for towers Perrin-Tillinghast. *Interfaces* **30**(1), 96–114 (2000)
78. J.M. Mulvey, W.C. Kim, C. Lin, Optimizing a portfolio of liquid and illiquid assets. This volume
79. K. Natarajan, D. Pachamanova, M. Sim, Incorporating asymmetric distributional information in robust value-at-risk optimization. *Manag. Sci.* **54**(3), 573–585 (2008)
80. K. Natarajan, M. Sim, J. Uichanco, Tractable robust expected utility and risk models for portfolio optimisation. *Math. Financ.* **20**(4), 695–731 (2010)
81. D. Pachamanova, N. Gulpinar, E. Canakoglu, Robust data-driven approaches to pension fund asset-liability management under uncertainty. This volume
82. G.C. Pflug, W. Römisch, *Modeling, Measuring and Managing Risk* (World Scientific Publishing, Singapore, 2007)
83. G.C. Pflug, D. Wozabal, Ambiguity in portfolio selection. *Quant. Finan.* **7**(4), 435–442 (2007)

84. G.C. Pflug, A. Pichler, D. Wozabal, The  $1/n$  investment strategy is optimal under high model ambiguity. *J. Bank. Financ.* **36**(2), 410–417 (2012)
85. S.T. Rachev, J.S.J. Hsu, B.S. Bagasheva, F.J. Fabozzi, *Bayesian Methods in Finance* (Wiley, New York, 2008)
86. N. Rujeerapaiboon, D. Kuhn, W. Wiesemann, Robust growth-optimal portfolios. *Manag. Sci.* **62**(7), 2090–2109 (2016)
87. S. Satchell, A. Scowcroft, A demystification of the Black–Litterman model: managing quantitative and traditional portfolio construction. *J. Asset Manag.* **1**, 138–150 (2000)
88. M. Sekerke, *Bayesian Risk Management: A Guide to Model Risk and Sequential Learning in Financial Markets* (Wiley, New York, 2015)
89. A. Shapiro, D. Dentcheva, A. Ruszczyński, *Lectures on Stochastic Programming: Modeling and Theory* (SIAM, Philadelphia, 2009)
90. A. Shapiro, A. Nemirovski, On complexity of stochastic programming problems, in *Continuous Optimization: Current Trends and Applications*, ed. by V. Jeyakumar, A.M. Rubinov (Springer, Berlin, 2005), pp. 111–144
91. G. Szego (ed.), *Risk Measures for the 21st Century* (Wiley, New York, 2004)
92. R.H. Tütüncü, M. Koenig, Robust asset allocation. *Ann. Oper. Res.* **132**(1–4), 157–187 (2004)
93. S. Pagliarani, T. Vargiolu, Portfolio optimization in a defaultable Lévy-driven market model. *OR Spectr.* **37**(3), 617–654 (2015)
94. P. Vayanos, D. Kuhn, B. Rustem, A constraint sampling approach for multi-stage robust optimization. *Automatica* **48**(3), 459–471 (2012)
95. W. Wiesemann, D. Kuhn, M. Sim, Distributionally robust convex optimization. *Oper. Res.* **62**(6), 1358–1376 (2014)
96. H. Xu, D. Zhang, Monte Carlo methods for mean-risk optimization and portfolio selection. *Comput. Manag. Sci.* **9**(1), 3–29 (2012)
97. S. Zymmler, D. Kuhn, B. Rustem, Worst-case value at risk of nonlinear portfolios. *Manag. Sci.* **59**(1), 172–188 (2013)
98. S. Zymmler, B. Rustem, D. Kuhn, Robust portfolio optimization with derivative insurance guarantees. *Eur. J. Oper. Res.* **210**(2), 410–424 (2011)