

Risk Control in Asset Management: Motives and Concepts

Thomas Dangl, Otto Randl and Josef Zechner

Abstract In traditional portfolio theory, risk management is limited to the choice of the relative weights of the riskless asset and a diversified basket of risky securities, respectively. Yet in industry, risk management represents a central aspect of asset management, with distinct responsibilities and organizational structures. We identify frictions that lead to increased importance of risk management and describe three major challenges to be met by the risk manager. First, we derive a framework to determine a portfolio position's marginal risk contribution and to decide on optimal portfolio weights of active managers. Second, we survey methods to control downside risk and unwanted risks since investors frequently have nonstandard preferences, which make them seek protection against excessive losses. Third, we point out that quantitative portfolio management usually requires the selection and parametrization of stylized models of financial markets. We, therefore, discuss risk management approaches to deal with parameter uncertainty, such as shrinkage procedures or resampling procedures, and techniques of dealing with model uncertainty via methods of Bayesian model averaging.

We thank Victor DeMiguel, Matthias Scherer, Neal Stoughton, Raman Uppal, Arne Westerkamp, Rudi Zagst and an anonymous referee for helpful comments.

T. Dangl
Vienna University of Technology, Vienna, Austria
e-mail: thomas.dangl@tuwien.ac.at

O. Randl · J. Zechner(✉)
WU Vienna University of Economics and Business, Vienna, Austria
e-mail: otto.randl@wu.ac.at

J. Zechner
e-mail: josef.zechner@wu.ac.at

1 Introduction

In traditional portfolio theory the scope for risk management is limited. Wilson [63] showed that in the absence of frictions the consumption allocation of each agent in an efficient equilibrium satisfies a linear sharing rule as long as agents have equi-cautious HARA utilities. This implies that investors are indifferent between the universe of securities and having access to only two appropriately defined portfolio positions, a result that is usually referred to as the Two-Fund Separation Theorem. If a riskless asset exists, then these two portfolios can be identified as the riskless asset and the tangency portfolio. Risk management in this traditional portfolio theory is, therefore, trivial: the portfolio manager only needs to choose the optimal location on the line that combines the riskless asset with the tangency portfolio, i.e., on the capital market line. Risk management is thus equivalent to choosing the relative weights that should be given to the tangency portfolio and to the riskless asset, respectively.

In a more realistic model that allows for frictions, risk management in asset management becomes a much more central and complex component of asset management. First, a world with costly information acquisition will feature informational asymmetries regarding the return moments, as analyzed in the seminal paper by Grossman and Stiglitz [29]. In this setup, investors generally do not hold the same portfolio of risky assets and the two fund separation theorem breaks down (see, e.g., Admati [1]). We will refer to such portfolios as active portfolios. In such a setup, risk management differs from the simple structure described above for the traditional portfolio theory. Second, frictions such as costly information acquisition frequently require delegated portfolio management, whereby an investor transfers decision power to a portfolio manager. This gives rise to principal-agent conflicts that may be mitigated by risk monitoring and portfolio risk control. Third, investors may have nonstandard objective functions. For example, the investor may exhibit large costs if the end-of-period portfolio value falls below a critical level. This may be the case, for example, because investors are subject to their own principal-agent conflicts. Alternatively, investors may be faced with model risk, and thus be unable to derive probability distributions over possible portfolio outcomes. In such a setting investors may have nonstandard preferences, such as ambiguity aversion. We will now discuss each of these deviations from the classical frictionless paradigm and analyze how it affects portfolio risk management.

2 Risk Management for Active Portfolios

If the optimal portfolio differs from the market portfolio, portfolio risk management becomes a much more complicated and important task for the portfolio manager. For active portfolios individual positions' risk contributions are no longer fully determined by their exposures to systematic risk factors that affect the overall market portfolio. A position's contribution to overall portfolio risk must not be measured

by the sensitivity to the systematic risk factors, but instead by the sensitivity to the investor's portfolio return. For active portfolios the manager must, therefore, correctly measure each asset's risk-contribution to the overall portfolio risk and ensure that it corresponds to the expected return contribution of the asset. We will now derive a simple framework that a portfolio manager may use to achieve this.

We consider an investor who wishes to maximize his expected utility, $E[\tilde{u}]$. In this section, we consider the case where the investor exhibits constant absolute risk aversion with the coefficient of absolute risk aversion denoted by Γ . In the following derivations, we borrow ideas from Sharpe [61] and assume for convenience that investment returns and their dispersions are small relative to initial wealth, V_0 . Thus, we can approximate $\Gamma \cong \gamma/V_0$ with γ denoting the investor's relative risk aversion. This allows for easy translation of the results into the context of later sections, where we focus on relative risk aversion.¹ An expected-utility maximizer with constant absolute risk aversion solves

$$\max_w E[\tilde{u}] = \max_w E[-\exp(-\Gamma(V_0(1 + w'\tilde{r})))] = \max_w E[-\exp(-\gamma(1 + w'\tilde{r}))], \quad (1)$$

where w represents the $(N \times 1)$ vector of portfolio weights and \tilde{r} is the $(N \times 1)$ vector of securities returns. We make standard assumptions of mean-variance analysis, and denote μ^e as the $(N \times 1)$ vector of securities' expected returns in excess of the risk free rate r_f , $\sigma_p^2(w)$ the portfolio's return variance given weights w , and Σ the covariance matrix of excess returns. $\mathbf{MR} = 2\Sigma w$ constitutes the vector of marginal risk contributions resulting from a marginal increase in portfolio weight of the respective asset, i.e., $\mathbf{MR} = \partial(\sigma_p^2)/\partial w$, financed against the riskless asset. For each asset i in the portfolio we must, therefore, have

$$\mu_i^e = \gamma e_i' \Sigma w = \frac{1}{2} \gamma \mathbf{MR}_i, \quad (2)$$

which implies

$$\frac{\mu_i^e}{\mathbf{MR}_i} = \frac{\mu_j^e}{\mathbf{MR}_j} = \frac{1}{2} \gamma \quad \forall i, j. \quad (3)$$

These results show the fundamental difference between risk management for active and passive portfolios. While in the traditional world of portfolio theory, each asset's risk contribution was easily measured by a constant (vector of) beta coefficient(s) to the systematic risk factor(s), the active investor must measure a security's risk contribution by the sensitivity of the asset to the specific portfolio return, expressed by $2e_i' \Sigma w$. This expression makes clear that each position's marginal risk contribution depends not only on the covariance matrix Σ , but also on the portfolio weights, i.e., the chosen vector w . It actually converges to the portfolio variance, σ_p^2 , as the security's weight approaches one. In the case of active portfolios, these weights are likely to change over time, and so will each position's marginal risk

¹ See, e.g., Pennacchi [55] for more details on this assumption.

contribution. The portfolio manager can no longer observe a position's relevant risk characteristics from readily available data providers such as the stock's beta reported by Bloomberg, but must calculate the marginal risk contributions based on the portfolio characteristics. As shown in Eq. (3), a major responsibility of the portfolio risk manager now is to ensure that the ratios of securities' expected excess returns over their marginal risk contribution are equated.

2.1 Factor Structure and Portfolio Risk

A prevalent model of investment management in practice features a CIO who decides on the portfolio's asset allocation and on the allocation between passively or actively managed mandates within each asset class. The actual management of the positions within each asset class is then delegated to external managers. In the following we provide a consistent framework within which such a problem can be analyzed. We hereby assume a linear return generating process so that the vector of asset excess returns, r^e can be written as

$$r^e = \alpha + Bf^e + \epsilon, \quad (4)$$

where

- r^e is the $(N \times 1)$ vector of fund or manager returns in excess of the risk free return
- B is a $(N \times K)$ matrix that denotes the exposure of each of the N assets to the K return factors
- f^e is a $(K \times 1)$ vector of factor excess returns and
- ϵ is the error term (independent of f^e).

Let Σ_f denote the covariance matrix of factor excess returns and Ω the covariance matrix of residuals, ϵ . Then, the covariance matrix of managers' excess returns Σ is given by

$$\begin{aligned} \Sigma &= E(r^e r^{e'}) \\ &= E([Bf^e + \epsilon][Bf^e + \epsilon]') \\ &= E(Bf^e f^{e'} B' + \epsilon\epsilon') \\ &= BE(f^e f^{e'})B' + E(\epsilon\epsilon') \\ &= B\Sigma_f B' + \Omega. \end{aligned}$$

Let w denote the $N \times 1$ vector of weights assigned to managers by the CIO, then the portfolio excess return r_p^e is given by

$$r_p^e = w' r^e.$$

If e_i is the i th column of the $(N \times N)$ identity matrix then

$$\begin{aligned}\text{Cov}(r_i^e, r_p^e) &= \text{Cov}(e_i' B f^e + e_i' \epsilon, w' B f^e + w' \epsilon) \\ &= \text{Cov}(e_i' B f^e, w' B f^e) + \text{Cov}(e_i' \epsilon, w' \epsilon) \\ &= e_i' B \Sigma_f B' w + e_i' \Omega w.\end{aligned}$$

The beta of manager i 's return with respect to the portfolio is then

$$\tilde{\beta}_i = \frac{e_i' B \Sigma_f B' w + e_i' \Omega w}{w' (B \Sigma_f B' + \Omega) w}.$$

Thus, we have an orthogonal decomposition of the vector of betas, $\tilde{\beta}$, into a part that is due to factor exposure, $\tilde{\beta}^S$, and a part that is due to the residuals of active managers (tracking error), $\tilde{\beta}^I$

$$\tilde{\beta} = \frac{B \Sigma_f B' w + \Omega w}{w' (B \Sigma_f B' + \Omega) w} = \underbrace{\frac{B \Sigma_f B' w}{w' (B \Sigma_f B' + \Omega) w}}_{\tilde{\beta}^S} + \underbrace{\frac{\Omega w}{w' (B \Sigma_f B' + \Omega) w}}_{\tilde{\beta}^I}.$$

We can now determine the beta of a pure factor excess return f_k^e to the portfolio. With e_k^F denoting the k th column of the $(K \times K)$ identity matrix, the covariance between the factor excess return and the portfolio excess return is

$$\begin{aligned}\text{Cov}(f_k^e, r_p^e) &= \text{Cov}(e_k^{F'} f^e, w' B f^e + w' \epsilon) \\ &= e_k^{F'} \Sigma_f B' w.\end{aligned}$$

The vector of pure factor betas, $\tilde{\beta}^F$, to the portfolio is therefore

$$\tilde{\beta}^F = \frac{\Sigma_f B' w}{w' (B \Sigma_f B' + \Omega) w}$$

We thus have $\tilde{\beta}^S = B \tilde{\beta}^F$. Consequently, a position's beta to the portfolio can be written as

$$\tilde{\beta} = B \tilde{\beta}^F + \tilde{\beta}^I.$$

i.e., we can decompose the position's beta into the exposure-weighted betas of the pure factor returns plus the beta of the position's residual return.

Next we can derive the vector of marginal risk contributions of the portfolio positions. Given the factor structure above, the effect of a small change in portfolio weights, w , on portfolio risk, σ_p^2 is given by MR:

$$\begin{aligned} \frac{1}{2}\text{MR} &= \frac{1}{2} \frac{\partial}{\partial w} w' \Sigma w = \Sigma w \\ &= \sigma_p^2 \tilde{\beta} = \sigma_p^2 (B\tilde{\beta}^F + \tilde{\beta}^I) \end{aligned}$$

Thus, an individual portfolio position i 's marginal risk contribution, MR_i , is given by

$$\begin{aligned} \frac{1}{2}\text{MR}_i &= \frac{1}{2} \frac{\partial}{\partial w_i} w' \Sigma w \\ &= e_i' \Sigma w = \sigma_p^2 \tilde{\beta}_i \\ &= \sigma_p^2 (B\tilde{\beta}_i^F + \tilde{\beta}_i^I). \end{aligned} \tag{5}$$

2.2 Allocation to Active and Passive Funds

One important objective of risk control in a world with active investment strategies is to ensure that an active portfolio manager's contribution to the portfolio return justifies his idiosyncratic risk or "tracking error". If this is not the case, then it is better to replace the active manager with a passive position that only provides a pure factor exposures but no idiosyncratic risks. To analyze this question we define ν^e as the vector of expected excess returns of the factor-portfolios and assume without loss of generality $\nu^e > 0$. Then, the vector of expected portfolio excess returns can be written as

$$\mu^e = E(\alpha + Bf^e + \epsilon) = \alpha + B\nu^e. \tag{6}$$

The first order optimality condition (3) states that the portfolio weight assigned to manager i should not be reduced as long as it holds that:

$$\frac{\mu_i^e}{\text{MR}_i} \geq \frac{\mu_j^e}{\text{MR}_j}, \quad \forall j.$$

Substituting marginal risk contribution from (5) and expected return from (6) into the above relation, we conclude that a manager i with $\text{MR}_i > 0$ justifies her portfolio weight relative to a pure factor investment in factor k iff

$$\frac{\mu_i^e}{E(f_k^e)} = \frac{e_i' B E(f^e) + \alpha_i}{E(f_k^e)} = \frac{e_i' B \nu^e + \alpha_i}{\nu_k^e} \geq \frac{\tilde{\beta}_i}{\tilde{\beta}_k^F} = \frac{e_i' B \tilde{\beta}^F + e_i' \tilde{\beta}^I}{\tilde{\beta}_k^F}.$$

Consider the case where asset manager i has exposure only to factor k , denoted by $B_{i,k}$. Then, this manager justifies her capital allocation iff

$$\frac{B_{i,k}\nu_k^e + \alpha_i}{\nu_k^e} \geq \frac{B_{i,k}\tilde{\beta}_k^F + \tilde{\beta}_i^I}{\tilde{\beta}_k^F}$$

$$B_{i,k} + \frac{\alpha_i}{\nu_k^e} \geq B_{i,k} + \frac{\tilde{\beta}_i^I}{\tilde{\beta}_k^F}$$

$$\alpha_i \geq \frac{\tilde{\beta}_i^I}{\tilde{\beta}_k^F} \nu_k^e.$$

Note that in general this condition depends on the portfolio weight. For sufficiently small weights w_i , manager i 's tracking error risk will be “non-systematic” in the portfolio context, i.e., $\tilde{\beta}_i^I = 0$. However, as manager i 's weight in the portfolio increases, his tracking error becomes “systematic” in the portfolio context. Therefore, the manager’s hurdle rate increases with the portfolio weight. This is illustrated in Example 1.

Example 1 Consider the special case where there is only one single factor and a portfolio, which consists of a passive factor-investment and a single active fund. The portfolio weight of the passive investment is denoted by w_1 and that of the active fund by w_2 . The active fund is assumed to have a beta with respect to the factor denoted by β and idiosyncratic volatility of σ_I .²

The covariance of factor returns is then a simple scalar equal to the factor return variance, the matrix of factor exposures B has dimension (2×1) and the idiosyncratic covariance matrix is (2×2)

$$w = \begin{pmatrix} 1 - w_2 \\ w_2 \end{pmatrix}, \quad \Sigma_f = \sigma_\nu^2, \quad B = \begin{pmatrix} 1 \\ \beta \end{pmatrix}, \quad \Omega = \begin{pmatrix} 0 & 0 \\ 0 & \sigma_I^2 \end{pmatrix}.$$

The usual assumption $\nu^e > 0$, $\sigma_I^2 > 0$ applies. The hurdle to be met by the alpha of the active fund is accordingly given by

$$\alpha \geq H(w_2) = \frac{\tilde{\beta}_2^I}{\tilde{\beta}_F} \nu^e = \frac{\sigma_I^2 w_2}{\sigma_\nu^2 (1 - (1 - \beta)w_2)} \nu^e.$$

The derivative of this hurdle with respect to the weight of the active fund w_2 is

$$\frac{dH}{dw_2} = \frac{\sigma_I^2}{\sigma_\nu^2} \nu^e \frac{1}{(1 - (1 - \beta)w_2)^2} > 0,$$

i.e., the hurdle $H(w_2)$ has a strictly positive slope, thus, the higher the portfolio weight w_2 of an active fund, the higher is the required α it must deliver. This is so because with low portfolio weight, the active fund’s idiosyncratic volatility is almost orthogonal to

² Note that β is the linear exposure of the fund to the factor. It is a constant and independent of portfolio weights. In contrast, betas of portfolio constituents relative to the portfolio, $\tilde{\beta}^F$ and $\tilde{\beta}^I$, depend on weights.

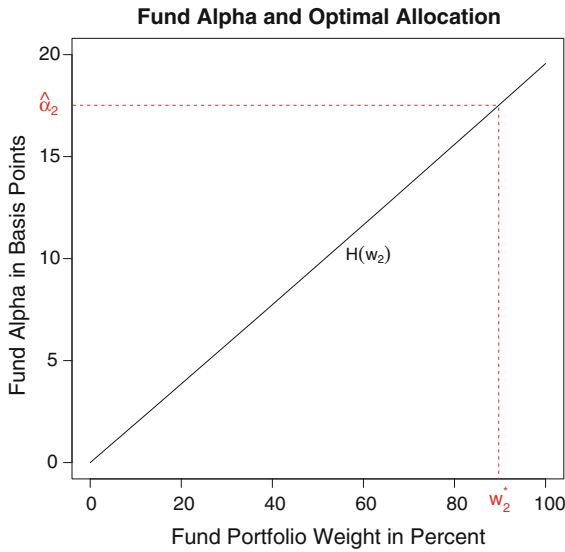


Fig. 1 Minimum alpha justifying portfolio weights

the portfolio return, and so its contribution to the overall portfolio risk is low. When in contrast the active fund has a high portfolio weight, its idiosyncratic volatility already co-determines the portfolio return and is—in the portfolio’s context—a systematic component. The marginal risk contribution of the fund is then larger and consequently demands a higher compensation, translating into an upward-sloping α -hurdle.

Take as an example JPMorgan Funds—Highbridge US STEEP, an open-end fund incorporated in Luxembourg that has exposure primarily to U.S. companies, through the use of derivatives. Using monthly data from 12/2008 to 12/2013 (data source: Bloomberg), we estimate

$$\hat{\Sigma}_f = \hat{\sigma}_v^2 = 0.002069, \quad \hat{B} = \begin{pmatrix} 1 \\ 0.9821 \end{pmatrix}, \quad \hat{\Omega} = \begin{pmatrix} 0 & 0 \\ 0 & 0.000303 \end{pmatrix}.$$

Furthermore, we use the historical average of the market risk premium $\bar{v} = 0.013127$, and the fund’s estimated alpha $\hat{\alpha} = 0.001751$. The optimal allocation is the vector of weights w^* such that the marginal excess return divided by the marginal risk contribution is equal for both assets in the portfolio. The increasing relationship between alpha and optimal fund weight is illustrated in Fig. 1. At the estimated alpha of 17.51 basis points, the optimal weights are given by

$$w^* = \begin{pmatrix} 0.1029 \\ 0.8971 \end{pmatrix}.$$

3 Dealing with Investors Downside-Risk Aversion

When discussing investor's utility optimization in Sect. 2, we referred to literature showing that under fairly general assumptions optimal static sharing rules are linear in the investment's payoff, i.e., optimal risk sharing implies holding a certain fraction of a risky investment rather than negotiating contracts with nonlinear payoffs. In a dynamic context, Merton [51] derives an optimal savings-consumption rule that is also in accordance with this finding. Consider a continuous-time framework with a single risky and a riskless asset, where the investor can change the allocation w_t to the risky asset over time. When the risky asset follows a geometric Brownian motion with drift μ and volatility σ , and utility exhibits a constant relative risk aversion γ , then the optimal allocation to the risky asset is constant over time and can be described as $w_t = \mu^e / (\gamma\sigma^2)$. This means with constant investment opportunities (μ and σ constant over time) investors keep the proportions of the risky and risk-free assets in the portfolio unchanged over time. To keep weights constant, portfolio rebalancing requires buying the risky asset when it decreases in value and selling it with increasing prices.

While these theoretical results suggest that an investor should not avoid exposure to risky investments even after sharp draw-downs of her portfolio's value, financial intermediaries face strong demand for products that provide portfolio insurance. That is, investors seem to have considerable downside-risk aversion. Rebalancing to constant portfolio weights is in clear contrast to portfolio insurance strategies, where the allocation to the risky asset has to be decreased if it falls in value, and the risky asset will be purchased in response to price increases. Perold and Sharpe [56] note that these opposing rebalancing rules lead to different shapes of strategy payoff curves. Buying stocks as they fall (as in the Merton model) leads to concave payoff curves. Such strategies do well in flat but oscillating markets, as assets are bought cheaply and sold at higher prices. However, in persistent downmarkets losses are aggravated from buying ever more stocks as they fall. Portfolio insurance rebalancing rules prescribe the opposite: selling stocks as they fall. This limits the impact of persistent down markets on the final portfolio value and at the same time keeps the potential of upmarkets intact, leading to a convex payoff profile. Yet if markets turn out flat but oscillating, convex strategies perform poorly.

3.1 Portfolio Insurance

In this paper, we define portfolio insurance as a dynamic investment strategy that is designed to limit downside risk. The variants of portfolio insurance are, therefore, popular examples of convex strategies. The widespread use of portfolio insurance strategies among both individual and institutional investors indicates that not all market participants are equally capable of bearing the downside risk associated with their average holding of risky assets. Individual investors might be subject to habit

formation or recognition of subsistence levels that define a minimum level of wealth required. For corporations, limited debt capacity makes it impossible to benefit from profitable investment projects if wealth falls below a critical value. Furthermore, kinks in the utility function could originate in agency problems, e.g., career concerns of portfolio managers, who see fund flows and pay respond in an asymmetric way to performance. In the literature on portfolio insurance, Leland [47] has stated the prevalence of convex over concave strategies for an investor whose risk aversion decreases in wealth more rapidly than for the representative agent. Alternatively, portfolio insurance strategies should be demanded by investors with average risk tolerance, but above average return expectations. Leland argues that insured strategies allow such an optimistic investor to more fully exploit positive alpha situations through greater levels of risky investment, while still keeping risk within manageable bounds.

Brennan and Solanki [14] contrast this analysis and derive a formal condition for optimality of an option like payoff that is typical for portfolio insurance. It can be shown that a payoff function where the investor receives the maximum of the reference portfolio's value and a guaranteed amount is optimal only under the stringent conditions of a zero risk premium and linear utility for wealth levels in excess of the guaranteed amount. Similarly, Benninga and Blume [9] argue that in complete markets utility functions consistent with optimality of portfolio insurance would have to exhibit unrealistic features, like unbounded risk aversion at some wealth level. However, they make the point that portfolio insurance can be optimal if markets are not complete. An extreme example of market incompleteness in this context, which makes portfolio insurance attractive, is the impossibility for an investor to allocate funds into the risk-free asset. Grossman and Vila [30] discuss portfolio insurance in complete markets, noting that the solution of an investor's constrained portfolio optimization problem (subject to a minimum wealth constraint $V_T > K$) can be characterized by the solution of the unconstrained problem plus a put option with exercise price K . More recently, Dichtl and Drobetz [19] provide empirical evidence that portfolio insurance is consistent with prospect theory, introduced by Kahneman and Tversky [41]. Loss-averse investors seem to use a reference point to evaluate portfolio gains and losses. They experience an asymmetric response to increasing versus decreasing wealth, in being more sensitive to losses than to gains. In addition, risk aversion also depends on the current wealth level relative to the reference point. The model by Gomes [27] shows that the optimal dynamic strategy followed by loss-averse investors can be consistent with portfolio insurance.³

³ It is interesting to study the potential effects of portfolio insurance on the aggregate market. As our focus is the perspective of a risk-manager who does not take into account such market-wide effects of his actions, we do not cover this literature. We refer the interested reader to Leland and Rubinstein [46], Brennan and Schwartz [13], Grossman and Zhou [32] and Basak [6] as a starting point.

3.2 Popular Portfolio Insurance Strategies

The main portfolio insurance strategies used in practice are stop-loss strategies, option-based portfolio insurance, constant proportion portfolio insurance, ratcheting strategies with adjustments to the minimum wealth target, and value-at-risk based portfolio insurance.

3.2.1 Stop-Loss Strategies

The simplest dynamic strategy for an investor to limit downside risk is to protect his investment using a stop-loss strategy. In this case, the investor sets a minimum wealth target or floor F_T , that must be exceeded by the portfolio value V_T at the investment horizon T . He then monitors if the current value of the portfolio V_t exceeds the present value of the floor $\exp(-r_f(T-t))F_T$, where r_f is the riskless rate of interest. When the portfolio value reaches the present value of the floor, the investor sells the risky and buys the riskfree asset. While this strategy has the benefit of simplicity, there are several disadvantages. First, due to discreteness of trading or illiquidity of assets, the transaction price might be undesirably far below the price triggering portfolio reallocation. Second, once the allocation has switched into the riskfree asset the portfolio will grow deterministically at the riskfree rate, making it impossible to even partially participate in a possible recovery in the price of the risky asset.

3.2.2 Option-Based Portfolio Insurance

Brennan and Schwartz [12] and Leland [47] describe that portfolio insurance can be implemented in two equivalent ways: (1) holding the reference portfolio plus a put option, or (2) holding the riskfree asset plus a call option. When splitting his portfolio into a position S_0 in the risky asset and P_0 in a protective put option at time $t = 0$, the investor has to take into account the purchase price of the option when setting the exercise price K , solving $(S_0 + P_0(K)) \cdot (F_T/V_0) = K$ for K . The ratio F_T/V_0 is the minimum wealth target expressed as a fraction of initial wealth. If such an option is available on the market it can be purchased and no further action is needed over the investment horizon; alternatively such an option can be synthetically replicated as popularized by Rubinstein and Leland [58]. Again, the risky asset will be bought on price increases and sold on falling prices, but in contrast to the stop-loss strategy, changes in the portfolio allocation will now be implemented smoothly. Even after a fall in the risky asset's price there is scope to partially participate in an eventual recovery as long as Delta is strictly positive. Toward the end of the investment horizon, Delta will generally be very close to either zero or one, potentially leading to undesired portfolio switching if the risky asset fluctuates around the present value of the exercise price.

3.2.3 Constant Proportion Portfolio Insurance

In order to provide a simpler alternative to the option replication approach described above, Black and Jones [10] propose CPPI for equity portfolios. Black and Perold [11] describe properties of CPPI and propose a kinked utility function for which CPPI is the optimal strategy. Implementation of CPPI starts with calculation of the *cushion* $C_t = V_t - F_t$, which is the amount by which the current portfolio value V_t exceeds the present value of the minimum wealth target ($F_t = \exp(-r_f(T - t))F_T$). Thus, the cushion can be interpreted as the risk capital available at time t . The exposure E_t to the risky asset is determined as a constant multiple m of the cushion C_t , while the remainder is invested risk free. To avoid excessive leverage, exposure will typically be determined subject to the constraint of a maximum leverage ratio l , hence $E_t = \min\{m \cdot C_t, l \cdot V_t\}$. If the portfolio is monitored in continuous time, the portfolio value at time T cannot fall below F_T . However, discrete trading in combination with sudden price jumps could lead to a breach of the minimum wealth target (gap risk).

3.2.4 Ratcheting Strategies

The portfolio insurance strategies discussed so far limit the potential shortfall from the start of the investment period to its end, frequently a calendar year. But investors may also be concerned with losing unrealized profits that have been earned within the year. Estep and Kritzman [23] propose a technique called TIPP (time invariant portfolio insurance) as a simple way of achieving (partial) protection of interim gains in addition to the protection offered by CPPI. Their methodology adjusts the floor F_t used to calculate the cushion C_t over time. The TIPP floor is set as the maximum of last period's floor and a fraction k of the current portfolio value: $F_t = \max(F_{t-1}, kV_t)$. This method of *ratcheting the floor up* is time invariant in the sense that the notion of a target date T is lost. However, if the percentage protection is required with respect to a specific target date, the method can be easily adjusted by setting a target date floor F_T proportional to current portfolio value V_t , which is then discounted. Grossman and Zhou [31] provide a formal analysis of portfolio insurance with a rolling floor, while Brennan and Schwartz [13] characterize a complete class of time-invariant portfolio insurance strategies, where asset allocation is allowed to depend on current portfolio value, but is independent of time.

3.2.5 Value-at-Risk-Based Portfolio Insurance

In a broader context, Value-at-Risk (VaR) has emerged as a standard for measurement and management of financial market risk. VaR has to be specified with confidence a and horizon Δt and is the loss amount that will be exceeded only with probability $(1 - a)$ over the time span Δt . It is, therefore, a natural measure to control portfolio drawdown risk. The typical definition of VaR assumes that over the time horizon no adjustments are made to the portfolio. Yet, if under adverse market movements

risk reducing transactions are implemented, VaR is likely to overestimate actual losses, making portfolio insurance even more effective. On the other hand, poor estimation of the return distribution will lead to bad quality of the VaR estimate. Herold et al. [35, 36] describe a VaR-based method for controlling shortfall risk. The allocation to the risky asset is chosen such that the VaR equals the prespecified minimum return. They note that their method can be seen as a generalized version of CPPI with a dynamic multiplier $m_t = 1/(\Phi^{-1}(a)\sqrt{\Delta t}\sigma_t)$, where $\Phi^{-1}(a)$ is the a -percentile of the standard normal distribution, and σ_t is the volatility of the reference portfolio. Typically, market volatility increases when markets crash, leading to a more pronounced reduction of the allocation to the risky asset as both the cushion and the multiplier shrink. This offers the potential advantage of VaR-based risk control that if markets calm, the allocation to the risky asset will increase again, allowing the portfolio to benefit from a recovery. Basak and Shapiro [7] take a critical view on VaR-based risk management: Strictly interpreting VaR as a risk quantile, managers could be inclined to deliberately assume extreme risks if they are not penalized for the severity of losses that occur with a probability less than $1 - a$. However, in a portfolio insurance context this could be easily fixed, e.g., by restrictions on assuming tail risks.

3.3 Performance Comparison

Benninga [8] uses Monte Carlo simulation techniques to compare stop-loss, OBPI, and CPPI. Surprisingly, he finds that stop-loss dominates with respect to terminal wealth and Sharpe ratio. Dybvig [21] considers asset allocation and portfolio payouts in the context of endowment management. If payouts are not allowed to decrease, CPPI exhibits more desirable properties than constant mix strategies. Balder et al. [4] analyze risks associated with implementation of CPPI under discrete-time trading and transaction costs. Zagst and Kraus [64] compare OBPI and CPPI with respect to stochastic dominance. Taking into account that implied volatility—which is relevant for OBPI—is usually higher than realized volatility—relevant for CPPI—they find that under specific parametrizations CPPI dominates. Recently, Dockner [20] compares buy-and-hold, OBPI and CPPI concluding that there does not exist a clear ranking of the alternatives. Dichtl and Drobetz [19] consider prospect theory (Kahneman and Tversky [41]) as framework to evaluate portfolio insurance strategies. They use a twofold methodological approach: Monte Carlo simulation and historical simulation with data for the German stock market. Within the behavioral finance context chosen, their findings provide clear support for the justification of downside protection strategies. Interestingly, in their study stop-loss, OBPI and CPPI turn out attractive while the high protection level of TIPP associated with opportunity costs in terms of reduced upside potential turns out to be suboptimal. Finally, they recommend to implement CPPI aggressively by using the highest multiplier m consistent with tolerance for overnight or gap risk.

Example 2 In 4 out of the 18 calendar years from 1995 to 2013, the S&P 500 total return index lost more than 5%. For investors with limited risk capacity it was not helpful that these losses happened three times in a row (2000, 2001, and 2002), or were severe (2008). The following example illustrates how simple versions of common techniques to control downside risk have performed over these 18 years. We assume investment opportunities in the S&P 500 index and a risk-free asset, an investment horizon equal to the calendar year, and a frictionless market (no transaction costs). Each calendar year the investment starts with a January 1st portfolio value of 100. Rebalancing is possible with daily frequency. For the portfolio insurance strategies investigated, the desired minimum wealth is given with 95, and free parameters are set in a way to make the strategies comparable, by ensuring equal equity allocations at portfolio start. This is achieved by resetting the multiples m for CPPI and TIPP each January 1st according to the Delta of the OBPI strategy. Similarly, the VaR confidence level is set to achieve this same equity proportion at the start of the calendar year. OBPI Delta also governs the initial equity portion of the buy-and-hold portfolio. Table 1 reports the main results, and Fig. 2 summarizes the distribution of year-end portfolio values in a box plot.

The achieved minimum wealth levels show that for CPPI, TIPP, OBPI, and VaR-based portfolio insurance even in the worst year the desired minimum wealth has been missed just slightly, while in the case of the stop loss strategy there is a considerable gap. This can be partly explained by the simple setup of the example (e.g., rebalancing using daily closing prices only, while in practice intraday decision-making and trading will happen). But a possibly large gap between desired and achieved minimum wealth is also systematic of stop loss strategies because of the mechanics of stop-loss orders. The moment the stop limit is reached, a market order to sell the entire portfolio is executed. The trading price, therefore, can and frequently will be lower than the limit. This can pose considerable problems in highly volatile and illiquid market environments. Option replication comes next in missing desired wealth protection.

Table 1 Portfolio insurance strategies

	Mean	Median	SD all	SD lower	Min	Max	Turnover
Long only	110.92	113.91	19.29	15.55	63.91	137.59	0.00
Buy & hold	108.32	107.15	12.52	9.18	79.99	132.11	0.00
Stop loss	108.77	105.77	16.09	6.52	89.13	137.59	0.42
CPPI	107.71	104.77	12.52	2.82	94.62	136.89	4.58
TIPP	105.31	104.20	7.45	3.29	94.63	122.75	1.03
OBPI	108.50	105.21	12.43	4.29	95.00	135.58	3.63
Option repl.	108.84	107.07	11.93	4.72	92.04	132.59	3.64
VaR	108.21	104.15	13.21	2.64	94.79	137.59	8.16

Comparison of portfolio insurance strategies, annual horizon, S&P 500, 1995–2013. We report end-of-year wealth levels per investment of 100 (mean, median, min, max); standard deviations calculated both over the whole sample (SD all) and for the subsample where the annual S&P 500 total return is below its mean (SD lower); turnover in the annual turnover ratio




Fig. 2 Comparison of portfolio insurance strategies, annual horizon, S&P 500, 1995–2013. For each strategy, the shaded area indicates the observations from the 25th to the 75th percentile, the median is shown as the line across the box and the mean as a diamond within the box. The whiskers denote the lowest datum still within 1.5 interquartile range of the lower quartile, and the highest datum still within 1.5 interquartile range of the upper quartile. If there are more extreme observations they are shown separately by a circle. The semitransparent horizontal line indicates the desired minimum wealth level

In the example, this might be due the simplified setup, where the exercise price of the option to be replicated is determined only once per year (at year start), and then daily Delta is calculated for this option and used for allocation into the risky and the riskless asset. In practice, new information on volatility and the level of interest rates will also lead to a reset of the strike used for calculation of the Delta. Another observation is that the standard deviation of annual returns is lowest for TIPP, which comes at the price of the lowest average return. If the cross-sectional standard deviation is computed only for the years with below-average S&P 500 returns, it is lowest for VaR-based risk control. For all methods shown, practical implementation will typically use higher levels of sophistication. For example, trading filters will be applied

to avoid adjusting portfolios as frequently as in the example leading to high turnover values.

3.4 Other Risks

In the previous discussion, shortfall risk was seen from the perspective of an investor holding assets only. However, many institutional investors simultaneously optimize a portfolio of assets A and liabilities L . Sharpe and Tint [62] describe a flexible approach to systematically incorporate liabilities into pension fund asset allocation, by optimizing over a surplus measure $S = A - kL$, where $k \in [0, 1]$ is a factor denoting the relative weight attached to liabilities. In the context of asset liability management, Ang et al. [2] analyze the effect of downside risk aversion, and offer an explanation why risk aversion tends to be high when the value of the assets approaches the value of the liabilities. Ang et al. [2] specify the objective function of the fund as mean-variance over asset returns plus a downside risk penalty on the liability shortfall that is proportional to the value of an option to exchange the optimal portfolio for the random value of the liabilities. An investor following their advice tends to be more risk averse than a portfolio manager implementing the Sharpe and Tint [62] model. For very high funding ratios, the impact of downside risk on risk taking, and therefore the asset allocation of the pension fund manager is small. For deeply underfunded plans, the value of the option is also relatively insensitive to changes in volatility, again leading to a small impact on asset allocation. The effect of liabilities on asset allocation is strongest when the portfolio value is close to the value of liabilities. In this case, lower volatility reduces the value of the exchange option, leading to a smaller penalty.

Another hedging motive arises if investors wish to bear only specific risks. This might be due to specialization of the investor in a certain asset class, making it desirable to hedge against risks not primarily driving the returns of this asset class. A popular example is currency risk, which has been recently analyzed by Campbell et al. [15] who find full currency hedging to be optimal for a variance-minimizing bond investor, but discuss the potential for overall risk reduction from keeping foreign exchange exposure partly unhedged in the case of equity portfolios.

4 Parameter Uncertainty and Model Uncertainty

Quantitative portfolio management builds on optimization output of stylized models, which (i) need to be carefully chosen to capture relevant features of the market framework and (ii) must be calibrated and parameterized. These choices, model selection, as well as model calibration, bear the risk of misspecification, which might have severely negative consequences on the desired out-of-sample properties of the

portfolio. Thus, a main application of risk management in asset management is controlling the risk inherent in model specification and parameter selection. In this section, we distinguish between parameter uncertainty and model uncertainty in the following way. With *parameter uncertainty* we refer to the case where we know the structure of the data generating process that lies behind the observed set of data, but the parameters of the process must be empirically determined.⁴ Finite data history is the only limiting factor, which prevents us from deriving the true values of the model parameters. Under the assumption of the null hypothesis, we can derive the joint distribution of the estimated parameters relative to the true values, and finally the joint predictive distribution of asset returns under full consideration of estimation problems. Thus, we can treat parameter uncertainty simply as an additional source of variability in returns. It is noncontroversial to assume that a decision-maker does not distinguish between uncertainty in returns caused by the general variability of returns and uncertainty that has its origin in estimation problems, and hence the portfolio optimization paradigm is not affected.

In contrast, with *model uncertainty* we refer to the case where a decision-maker is not sure, which model is the correct formulation that describes the underlying dynamics of asset returns. In such a case, it is generally not possible to specify probabilities for the models considered as feasible. Thus, model uncertainty increases uncertainty about asset returns, but we are not able to state a definite probability distribution of returns, which incorporates model uncertainty. That is, model uncertainty is a prototypical case of Knightian uncertainty, referring to Knight [42], where it is not possible to characterize the uncertain entity (in our case the asset return) by means of a probability distribution. Consequently, model uncertainty fundamentally changes the decision-making framework and we have to make assumptions regarding a decision-maker's preferences concerning situations of ambiguity.

4.1 Parameter Uncertainty

The most obvious estimation problem in a traditional minimum-variance portfolio optimization task arises when determining the covariance structure of asset returns. This is so because estimates of the sample covariance matrix turn out to be weakly conditioned in general and—as soon as the number of assets is larger than the number of periods considered in the return history—the sample covariance matrix is singular by construction.

Example 3 Consider as a broad asset universe the S&P 500 with $N = 500$ constituents. It is common practice to estimate the covariance structure of stock returns from two years of weekly returns. The argument for a restriction of the history to $T = 104$ weeks is a reaction to the fact that there is apparently some time-variation

⁴ We assume in general, that the model has a structure, which ensures that parameters are identifiable. For example, it is assumed that log-returns are normally distributed, but mean and variance must be estimated from observed data.

in the covariance structure, which the estimate is able to capture only if one restricts the used history.⁵

Let r denote the $(T \times N)$ matrix containing weekly returns, then the sample covariance matrix $\hat{\Sigma}_S$ is determined by

$$\hat{\Sigma}_S = \frac{1}{T-1} r' M r, \quad (7)$$

where the symmetric and idempotent matrix M is the residual maker with respect to a regression onto a constant,

$$M = \mathbb{I} - \mathbf{1}(\mathbf{1}'\mathbf{1})^{-1}\mathbf{1}',$$

with \mathbb{I} the $(T \times T)$ identity matrix and $\mathbf{1}$ a column vector containing T times the constant 1.

In the assumed setup, the sample covariance matrix is singular by construction. This is so because from (7) it follows that the rank of $\hat{\Sigma}_S$ is bounded from above by $\min\{N, T-1\}$.⁶ And even in the case where the number of return observations per asset exceeds the number of assets ($T > N+1$) the sample covariance matrix is weakly determined, hence, subject to large estimation errors since one has to estimate $N(N+1)/2$ elements of $\hat{\Sigma}_S$ from $T \cdot N$ observations.

Since a simple Markowitz optimization, see Markowitz [49], needs to invert the covariance matrix, matrix singularity prohibits any attempt of advanced portfolio optimization, and is thus the most evident estimation problem in portfolio management. Elton and Gruber [22] is an early contribution, which proposes the use of structural estimators of the covariance matrix. Jobson and Korkie [38] provide a rigorous analysis of the small sample properties of estimates of the covariance structure of returns.

Less evident are the problems caused by errors in the estimates of return expectations, whereas it turns out that they are economically much more critical. Jorion [39] shows in the context of international equity portfolio selection that the errors in the estimates of return expectations have a severe impact on the out-of-sample

⁵ Such an approach is typical for dealing with inadequate model specification. The formal estimate is based on the assumption that the covariance structure is stable. Since data show that the covariance structure is not stable, an ad-hoc adaptation—the limitation of the data history—is used to capture the recent covariance structure. The optimal amount of historical data that should be used cannot be derived within the model, but must be roughly calibrated to some measure of goodness-of-fit, which balances estimation error against timely response to time variations.

⁶ The residual maker M has at most rank $T-1$ because it generates residuals from a projection onto a one-dimensional subspace of \mathbb{R}^T . Since r has at most rank N , we have

$$\text{rank}(\hat{\Sigma}_S) \leq \min\{N, T-1\}.$$

For example, the sample covariance matrix estimated from two years of weekly returns of the 500 constituents of the S&P500 (104 observations per stock) has at most rank 103. Hence, it is not positive definite and not invertible, because at least 397 of its 500 eigenvalues are exactly equal 0.

performance of optimized portfolios. He further shows that the Bayes-Stein shrinkage approach introduced in Jorion [40] helps mitigate errors and at the same time improves out-of-sample properties of the portfolio.

Structural Estimators Means and covariances of asset returns are the most basic inputs into a portfolio optimization model. However, estimation errors in further model parameters like some measure of risk aversion, speed of reversion to long-term averages, etc., must be estimated from empirical data and are, thus, equally likely inflicted with estimation errors. While sample estimates of distribution means, (co-)variances and higher moments are generally unbiased and efficient, they tend to be noisy. This can be improved by imposing some sort of structure on the estimated parameters. Such structural estimates are less prone to estimation errors at the expense of ignoring part of the information inherent in the observed data sample. When determining the covariance structure of asset returns, Elton and Gruber [22] analyze a set of different structural assumptions, e.g., what they call the single index model (assuming that the pairwise covariance of asset returns is only generated by the assets individual correlation to a market index), the mean model (pairwise correlations between assets are assumed constant across the asset universe), and models that assume that the correlation structure of asset returns is determined by within industry averages or across industry averages or by a (small) number of principal components of the sample covariance matrix. They show that especially the particularly restrictive estimates (single index model and mean model) deliver forecasts of future correlation that are more accurate than the simple historical sample estimates.⁷

Shrinkage Estimators When determining model parameters θ , it is very popular to apply some shrinkage approach. This approach aims to combine the advantages of a sample estimate $\hat{\theta}_S$ (pure reliance on sample data) and a structural estimate $\hat{\theta}_{\text{struct}}$ (robustness) by computing some sort of weighted average⁸

$$\hat{\theta} = \lambda \hat{\theta}_S + (1 - \lambda) \hat{\theta}_{\text{struct}}.$$

While practitioners often use ad hoc weighting schemes, the literature provides a powerful Bayesian interpretation of shrinkage, which allows for the computation of *optimal* weights. In this Bayesian view, the structural estimator serves as the prior, which anchors the location of model parameters θ and the sample estimate acts as the conditioning signal. Bayes' rule then gives a stringent advice of how to combine prior and signal in order to compute the updated posterior that is used as an input for the portfolio optimization. The abovementioned Bayes-Stein shrinkage used in Jorion [39, 40] focuses on estimates of the expected returns. In the context of covariance estimation, an early contribution is Frost and Savarino [25]. More recently, Ledoit and Wolf [43] determine a more general Bayesian framework to optimize the shrinkage intensity, in which the authors explicitly correct for the fact

⁷ See, e.g., Dangl and Kashofer [18] for an overview of structural estimates of the covariance structure of large equity portfolios—including shrinkage estimates.

⁸ Shrinkage is usually a multivariate concept, i.e., λ is in general not a fixed scalar, but it depends on the observed data in some nonlinear fashion.

that the prior (i.e., the structural estimate of the covariance structure) as well as the updating information (i.e., the sample covariance matrix) are determined from the same data. Consequently, errors in these two inputs are not independent and the Bayesian estimate must control for the interdependence.⁹

Weight Restrictions A commonly observed reaction to parameter uncertainty in portfolio management is imposing ad hoc restrictions on portfolio weights. That is, the discretion of a portfolio optimizer is limited by maximum as well as minimum constraints on the weights of portfolio constituents.¹⁰ In sample, weight restrictions clearly reduce portfolio performance (as measured by the objective function used in the optimization approach).¹¹ Nevertheless, out of sample studies show, that in many cases weight restrictions improve the risk-return trade-off of portfolios. Jagannathan and Ma [37] provide evidence why weight restrictions might be an efficient response to estimation errors in the covariance structure. Analyzing minimum-variance portfolios they show that binding long only constraints are equivalent to shrinking extreme covariance estimates toward more moderate levels.

Robust Optimization A more systematic approach to parameter uncertainty than weight restrictions is robust optimization. After determining the uncertainty set \mathcal{S} for the relevant parameter vector p , robust portfolio optimization is usually formulated as a max-min problem where the vector w of portfolio weights solves the equation

$$w \in \operatorname{argmax}_w \{ \min_{p \in \mathcal{S}} f(w; p) \},$$

with $f(w; p)$ being the planner's objective function that she seeks to maximize. This is a conservative or worst-case approach, which in many real-world applications shows favorable out-of-sample properties (see Fabozzi et al. [24], or for more details on robust and convex optimization problems and its applications in finance see Lobo et al. [48]). Provided a distribution of the parameters is available, the rather extreme max-min approach could be relaxed by applying convex risk measures. In the context of derivatives pricing, Bannoe and Scherer [5] develop the concept of risk-capturing functionals and exemplify risk averse pricing using an average Value-at-Risk measure.

Resampling A different approach to deal with parameter uncertainty in asset management is resampling. This technique does not attempt to produce more robust parameter estimates or to build a portfolio-optimization model, which directly regards parameter uncertainty in portfolio optimization. Resampling is a simulation-based approach that was first described in the portfolio-optimization context by Michaud [52] and exists in different specifications. It takes the sample estimates of mean

⁹ See also Ledoit and Wolf [44, 45] for more on shrinkage estimates of the covariance structure.

¹⁰ Weight restrictions are frequently part of regulatory measures targeting the fund industry aimed to control the risk characteristics of investment funds.

¹¹ Green and Hollifield [28] argue that in the apparent presence of a strong factor structure in the cross section of equity returns, mean-variance optimal portfolios should take large short positions in selected assets. Hence, a restriction to a long-only portfolio is expected to negatively influence portfolio performance.

returns as well as of the covariance matrix and resamples a number R of return ‘histories’ (where R is typically between 1,000 and 10,000). From each of these return histories, an estimate of the vector of mean returns as well as of the covariance matrix is derived. These estimates form the ingredients to calculate R different versions of the mean-variance frontier. Resampling approaches differ in the set of restrictions used to determine the mean-variance frontiers and in the way how the frontiers are averaged to get the definite portfolio weights. Some authors criticize that the unconditionally optimal portfolio does not simply follow from an average over R vectors of conditionally optimal portfolio weights (see, e.g., Scherer [59] or Markowitz and Usmen [50]), others point out that the ad-hoc approach of resampling could be improved by using a Bayesian approach (see, e.g., Scherer [60], or Harvey et al. [33, 34]). Despite the critique, all those studies appreciate the out-of-sample characteristics of resampled portfolios.

Example 4 This simple example builds on Example 1 which discusses the optimal weight of an active fund relative to a passive factor investment. An index-investment in the S&P 500 serves as the passive factor investment and an active fund with the constituents of the S&P 500 as its investment universe is the delegated active investment strategy. In Example 1 we take a history of five years of monthly log-returns (60 observations) to estimate mean returns as well as the covariance structure and the alpha, which the fund generates relative to the passive investment. We use these estimates to conclude that the optimal portfolio weight of the fund should be roughly 90 % and only 10 % of wealth should be held as a passive investment.

Being concerned about the quality of our parameter estimation that feeds into the optimization, we first examine the regression, which was performed to come up with these estimates. Assuming that log-returns are normally distributed, we conclude from the regression in Example 1 that our best estimates of the parameters α , β and ν are

$$\hat{\alpha} = 17.51 \text{ bp/month}, \quad \hat{\beta} = 0.9821, \quad \hat{\nu} = 131.27 \text{ bp/month},$$

and that the estimation errors are t-distributed with a standard deviation¹²

$$\sigma_{58}(\hat{\alpha}) = 23.40 \text{ bp/month}, \quad \sigma_{58}(\hat{\beta}) = 0.0498, \quad \sigma_{59}(\hat{\nu}) = 454.91 \text{ bp/month}.$$

Furthermore, estimation errors in $\hat{\alpha}$ and $\hat{\beta}$ are negatively correlated with a correlation coefficient $\rho = -27.93 \%$ and errors in the estimate of the market risk premium $\hat{\nu}$ are uncorrelated to the errors in $\hat{\alpha}$ and $\hat{\beta}$.

A statistician would now conclude that neither the fund’s α nor the factor’s risk premium ν is significantly different from zero, and thus an investor should seek exposure to none of the two. Another approach is to extend the optimization problem and include parameter uncertainty as an additional source of variability in the final outcome.

¹² Subscripts denote degrees of freedom.

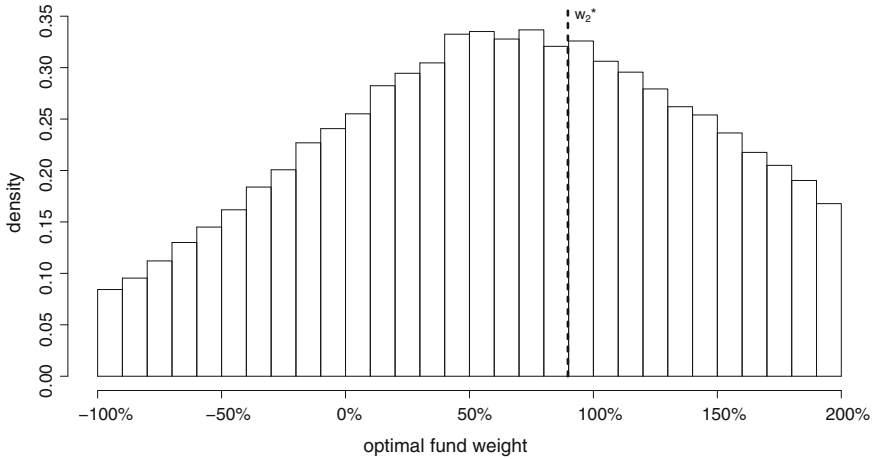


Fig. 3 Distribution of optimal portfolio weight in the interval $[-100\%, 200\%]$ of the active investment over 100,000 resampled histories. Approximately 29 % of weights lie outside the stated interval

In contrast to a full consideration of parameter uncertainty, we use a resampling approach, which addresses this issue in a more ad hoc manner. We take the empirical estimates as the true moments of the joint distribution of factor returns and active returns, and resample 100,000 histories.¹³ Then, we perform the optimization discussed in Example 1 on each of the simulated histories. Figure 3 illustrates the distribution of optimal active weights across these 100,000 histories. Given the null hypothesis that returns are normally distributed with the estimated moments, resampling gives a good and reliable overview of the joint distribution of model parameters we estimate and—finally—an overview of the distribution of optimal weights. We can conclude that in the present setup, optimal active weights are not well determined since the estimation of the optimization model from only 60 observations per time series is too noisy to get a well-determined outcome. While resampling generates a good picture of the overall effects of parameter uncertainty, it provides no natural advice for the optimal portfolio decision beyond this illustrative insight.¹⁴

¹³ This is the simplest version of resampling, mostly used in portfolio optimization. Given the null hypothesis that returns are normally distributed, we know that the empirical estimates of distribution moments are t -distributed around the true parameters, see Jobson and Korkie [38] for a detailed derivation of the small sample properties of these estimates. Thus, a more advanced approach samples for each of the histories, first the model parameters from their joint distribution, and then—given the selected moments—the history of normally distributed returns. Harvey et al. [33] is an example that uses advanced resampling to compare Bayesian inference with simple resampling.

¹⁴ Some authors do propose schemes how to generate portfolio decisions from the cross section of the simulation results, see, e.g., Michaud and Michaud [53]. These schemes are, however, criticized by other authors for not being well-founded in decision theory, e.g., Markowitz and Usmen [50] and others mentioned in the text above.

Finally, a study that perfectly illustrates the strong implications of parameter uncertainty on optimal portfolio decisions is Pastor and Stambaugh [54]. The authors question the paradigm that due to mean reverting returns, stocks are less risky in the long run than over short horizons. This proposition is true if we know the parameters of the underlying mean reverting process with certainty. Pastor and Stambaugh [54] show that as soon as we properly regard estimation errors in model parameters, additional uncertainty from estimation errors dominates the variance reduction due to mean reversion, and thus they provide strong evidence against time diversification in equity returns.

4.2 Model Uncertainty

Qualitatively different from dealing with parameter uncertainty is the issue of model uncertainty. Since it is not at all clear what the exact characteristics of the data-generating process, which underlies asset returns are, it is not obvious which attributes a model must feature in order to capture all economically relevant effects of the portfolio selection process. Hence, every model of optimal portfolio choice bears the risk of being misspecified. In Sect. 4.1 we already mention the fact that traditional portfolio models assume that mean returns and the covariance structure of returns are constant over time. This is in contrast to empirical evidence that the moments of the return distribution are time varying. Limiting the history, which is used to estimate distribution parameters, is a frequently used procedure to get a more actual estimate. The correct length of historical data that shall be used is, however, only rarely determined in a systematic manner.

Bayesian Model Averaging A systematic approach to estimation under model uncertainty is Bayesian model averaging. It builds on the concept of a Bayesian decision-maker that has a prior about the probability weights of competing models that are constructed to predict relevant variables (e.g., asset returns) one period ahead. Observed returns are then used to determine posterior probability weights for each of the models considered applying Bayes rule.¹⁵ Each of the competing models generates a predictive density for the next period's return. After observing the return, models which have assigned a high likelihood to the observed value (compared to others) experience an upward revision of their probability weight. In contrast, models that have assigned a low likelihood to the observed value experience a downward revision of their weight. Finally, the overall predictive density is calculated as a probability-weighted sum of all models' predictive densities. This Bayesian model averaging is an elegant way to approach a problem of model uncertainty to transform it into a standard portfolio problem to find the optimal risk-return trade-off under the derived predictive return distribution. This approach can, however, only be applied

¹⁵ The posterior probability that a certain model is the correct model is proportional to the product of the model's prior probability weight and the realized likelihood of the observed return.

under the assumption that the decision-maker has a single prior and that she shows no aversion against the ambiguity inherent in the model uncertainty.¹⁶

Raftery et al. [57] provide the technical details of Bayesian model averaging and Avramov [3], Cremers [16], and Dangl and Halling [17] are applications to return prediction. Bayesian model averaging treats model uncertainty just as an additional source of variation. The predictive density for next period's returns becomes more disperse the higher the uncertainty about models, which differ in their prediction. The optimal portfolio selection is then unchanged, but regards the additional contribution to uncertainty.

Ambiguity Aversion If it is not possible to explicitly assess the probability that a certain model correctly mirrors the portfolio selection problem and investors are averse to this form of ambiguity, alternative portfolio selection approaches are needed. Garlappi et al. [26] develop a portfolio selection approach for investors who have multiple priors over return expectations and show ambiguity aversion. The authors prove that the portfolio selection problem of such an ambiguity-averse investor can be formulated by imposing two modifications to the standard mean-variance model, (i) an additional constraint that guarantees that the expected return lies in a specified confidence region (the way how multiple priors are modeled) and (ii) an additional minimization over all expected returns that conform to the priors (mirroring ambiguity aversion). This model gives an intuitive illustration of the fact that ambiguity averse investors show explicit desire for robustness.

5 Conclusion

The asset management industry has substantial influence on financial markets and on the welfare of many citizens. Increasingly, citizens are saving for retirement via delegated portfolio managers such as pension funds or mutual funds. In many cases there are multiple layers of delegation. It is, therefore, crucial for the welfare of modern societies that portfolio managers manage and control their portfolio risks. This article provides an eagle's perspective on risk management in asset management.

In traditional portfolio theory, the scope for risk control in portfolio management is limited. Risk management is essentially equivalent to determining the fraction of capital that the manager invests in a broadly and well diversified basket of risky securities. Thus, the "risk manager" only needs to find the optimal location on the securities market line. By contrast, in a more realistic model of the world that accounts for frictions, risk management becomes a central and important module in asset management that is frequently separate from other divisions of an asset manager. We identify several major frictions that require risk management that goes beyond choosing the weight of the riskless asset in the portfolio. First, in a world with costly information acquisition, investors do not hold the same mix of risky assets. This

¹⁶ As explained in the introduction to this section, ambiguity aversion refers to preferences that express discomfort with uncertainty in the sense of Knight [42].

requires measuring a position's risk contribution relative to the specific portfolio. Thus, risk management requires constant measurement of each portfolio position's marginal risk contribution and comparing it to its marginal return contribution. This article derives a framework to calculate the marginal risk contributions and to decide on optimal portfolio weights of active managers.

In many realistic instances, investors have nonstandard preferences, which make them particularly sensitive to downside risks. We, therefore, review the main portfolio insurance concepts to achieve protection against downside risk. Stop-loss strategies, option-based portfolio insurance, constant proportion portfolio insurance, ratcheting strategies, and value-at-risk-based portfolio insurance. Using data for the S&P 500 since 1995 we simulate these alternative risk management concepts and demonstrate their risk and return characteristics.

Finally, we point out that quantitative portfolio management usually builds on the output from rather stylized models, which must be chosen to capture the relevant market environment, and which must be calibrated and parameterized. Both these choices, i.e., model selection and model calibration, contain the risk of misspecification, and thus the risk of negative effects on out-of-sample portfolio performance. We survey and discuss risk management approaches to deal with parameter uncertainty, such as shrinkage procedures or resampling procedures. Qualitatively different from parameter uncertainty is the effect of model uncertainty. Different ways of dealing with model uncertainty via methods of Bayesian model averaging and the consideration of ambiguity aversion are, therefore, surveyed and discussed.

The increased risk during the financial crisis and the following sovereign debt crisis has led to a substantially increased focus on risk control in the asset management industry. At the same time these market episodes have also demonstrated the limitations of risk management in asset management. For example that volatile markets without strong trends make existing downside protection strategies very expensive for investors. Furthermore, risk management concepts for long-term investors are still in their infancy. Scenario-based approaches, possibly combined with min-max strategies may be more useful in this context than standard risk management tools.

Open Access This chapter is distributed under the terms of the Creative Commons Attribution Noncommercial License, which permits any noncommercial use, distribution, and reproduction in any medium, provided the original author(s) and source are credited.

References

1. Admati, A.: A noisy rational expectations equilibrium for multi-asset securities markets. *Econometrica* **53**(3), 629–657 (1985)
2. Ang, A., Chen, B., Sundaresan, S.: Liability-driven investment with downside risk. *J. Portf. Manag.* **40**(1), 71–87 (2013)
3. Avramov, D.: Stock return predictability and model uncertainty. *J. Financ. Econ.* **64**, 423–458 (2002)

4. Balder, S., Brandl, M., Mahayni, A.: Effectiveness of CPPI strategies under discrete-time trading. *J. Econ. Dyn. Control* **33**(1), 204–220 (2009)
5. Bannör, K.F., Scherer, M.: Capturing parameter risk with convex risk measures. *Eur. Actuar. J.* **3**(1), 97–132 (2013)
6. Basak, S.: A comparative study of portfolio insurance. *J. Econ. Dyn. Control* **26**, 1217–1241 (2002)
7. Basak, S., Shapiro, A.: Value-at-risk-based risk management: optimal policies and asset prices. *Rev. Financ. Stud.* **14**(2), 371–405 (2001)
8. Benninga, S.: Comparing portfolio insurance strategies. *Finanzmarkt Portf. Manag.* **4**(1), 20–30 (1990)
9. Benninga, S., Blume, M.E.: On the optimality of portfolio insurance. *J. Financ.* **40**(5), 1341–1352 (1985)
10. Black, F., Jones, R.: Simplifying portfolio insurance. *J. Portf. Manag.* **14**(1), 48–51 (1987)
11. Black, F., Perold, A.F.: Theory of constant proportion portfolio insurance. *J. Econ. Dyn. Control* **16**(3–4), 403–426 (1992)
12. Brennan, M.J., Schwartz, E.S.: The pricing of equity-linked life insurance policies with an asset value guarantee. *J. Financ. Econ.* **3**, 195–213 (1976)
13. Brennan, M.J., Schwartz, E.S.: Time-invariant portfolio insurance strategies. *J. Financ.* **43**(2), 283–299 (1988)
14. Brennan, M.J., Solanki, R.: Optimal portfolio insurance. *J. Financ. Quant. Anal.* **16**(3), 279–300 (1981)
15. Campbell, J.Y., Medeiros, K.S.-D., Viceira, L.M.: Global currency hedging. *J. Financ.* **65**(1), 87–121 (2010)
16. Cremers, M.K.J.: Stock return predictability: a Bayesian model selection perspective. *Rev. Financ. Stud.* **15**, 1223–1249 (2002)
17. Dangl, T., Halling, M.: Predictive regressions with time-varying coefficients. *J. Financ. Econ.* **106**, 157–181 (2012)
18. Dangl, T., Kashofer, M.: Minimum-variance stock picking—a shift in preferences for minimum-variance portfolio constituents. Working paper (2013)
19. Dichtl, H., Drobotz, W.: Portfolio insurance and prospect theory investors: popularity and optimal design of capital protected financial products. *J. Bank. Financ.* **35**(7), 1683–1697 (2011)
20. Dockner, E.: Sind Finanzprodukte mit Kapitalgarantie eine attraktive Anlageform? In: Frick, R., Gantenbein, P., Reichling, P. (eds.) *Asset Management*, pp. 271–284. Haupt, Bern Stuttgart Wien (2012)
21. Dybvig, P.H.: Using asset allocation to protect spending. *Financ. Anal. J.* **55**(1), 49–62 (1999)
22. Elton, E.J., Gruber, M.J.: Estimating the dependence structure of share prices—implications for portfolio selection. *J. Financ.* **28**(5), 1203–1232 (1973)
23. Estep, T., Kritzman, M.: TIPP: insurance without complexity. *J. Portf. Manag.* **14**(4), 38–42 (1988)
24. Fabozzi, F.J., Kolm, P.N., Pachamanova, D., Focardi, S.M.: *Robust Portfolio Optimization and Management*. Wiley, Hoboken (2007)
25. Frost, P.A., Savarino, J.E.: An empirical Bayes approach to efficient portfolio selection. *J. Financ. Quant. Anal.* **21**(3), 293–305 (1986)
26. Garlappi, L., Uppal, R., Wang, T.: Portfolio selection with parameter and model uncertainty: a multi-prior approach. *Rev. Financ. Stud.* **20**(1), 41–81 (2007)
27. Gomes, F.J.: Portfolio choice and trading volume with loss-averse investors. *J. Bus.* **78**(2), 675–706 (2005)
28. Green, R.C., Hollifield, B.: When will mean-variance efficient portfolios be well diversified? *J. Financ.* **47**(5), 1785–1809 (1992)
29. Grossman, S.J., Stiglitz, J.E.: On the impossibility of informationally efficient markets. *Am. Econ. Rev.* **70**(3), 393–408 (1980)
30. Grossman, S.J., Vila, J.-L.: Portfolio insurance in complete markets: a note. *J. Bus.* **62**(4), 473–476 (1989)

31. Grossman, S.J., Zhou, Z.: Optimal investment strategies for controlling drawdowns. *Math. Financ.* **3**(3), 241–276 (1993)
32. Grossman, S.J., Zhou, Z.: Equilibrium analysis of portfolio insurance. *J. Financ.* **51**(4), 1379–1403 (1996)
33. Harvey, C.R., Liechty, J.C., Liechty, M.W.: Bayes vs. resampling: a rematch. *J. Invest. Manag.* **6**(1), 1–17 (2008)
34. Harvey, C.R., Liechty, J.C., Liechty, M.W., Müller, P.: Portfolio selection with higher moments. *Quant. Financ.* **10**(5), 469–485 (2010)
35. Herold, U., Maurer, R., Porschaker, N.: Total return fixed-income portfolio management. A risk-based dynamic strategy. *J. Portf. Manag.* Spring **31**, 32–43 (2005)
36. Herold, U., Maurer, R., Stamos, M., Vo, H.T.: Total return strategies for multi-asset portfolios: dynamically managing portfolio risk. *J. Portf. Manag.* **33**(2), 60–76 (2007)
37. Jagannathan, R., Ma, T.: Risk reduction in large portfolios: why imposing the wrong constraints helps. *J. Financ.* **58**(4), 1651–1684 (2003)
38. Jobson, J., Korkie, B.: Estimation for Markowitz efficient portfolios. *J. Am. Stat. Assoc.* **75**(371), 544–554 (1980)
39. Jorion, P.: International portfolio diversification with estimation risk. *J. Bus.* **58**(3), 259–278 (1985)
40. Jorion, P.: Bayes–Stein estimation for portfolio analysis. *J. Financ. Quant. Anal.* **21**(3), 279–292 (1986)
41. Kahneman, D., Tversky, A.: Prospect theory: an analysis of decision under risk. *Econometrica* **47**(2), 263–291 (1979)
42. Knight, F.H.: Risk, uncertainty and profit. Sentry press, Reprinted 1956 (1921)
43. Ledoit, O., Wolf, M.: Improved estimation of the covariance matrix of stock returns with an application to portfolio selection. *J. Empir. Financ.* **10**, 603–621 (2003)
44. Ledoit, O., Wolf, M.: Honey, I shrunk the sample covariance matrix. *J. Portf. Manag.* **30**(4), 110–119 (2004a)
45. Ledoit, O., Wolf, M.: A well-conditioned estimator for large-dimensional covariance matrices. *J. Multivar. Anal.* **88**, 365–411 (2004b)
46. Leland, H., Rubinstein, M.: Comments on the market crash: six months after. *J. Econ. Perspect.* **2**(3), 45–50 (1988)
47. Leland, H.E.: Who should buy portfolio insurance? *J. Financ.* **35**(2), 581–594 (1980)
48. Lobo, M.S., Vandenberghe, L., Boyd, S., Lebret, H.: Applications of second-order cone programming. *Linear Algebra Appl.* **284**(1–2), 193–228 (1998)
49. Markowitz, H.: Portfolio selection. *J. Financ.* **7**, 77–91 (1952)
50. Markowitz, H., Usmen, N.: Resampled frontiers versus diffuse Bayes: an experiment. *J. Invest. Manag.* **4**(1), 9–25 (2003)
51. Merton, R.: Optimum consumption and portfolio rules in a continuous-time model. *J. Econ. Theory* **3**(4), 373–413 (1971)
52. Michaud, R.O.: *Efficient Asset Management*. Oxford University Press, Oxford (1998)
53. Michaud, R.O., Michaud, R.O.: *Efficient Asset Management*, 2nd edn. Oxford University Press, Oxford (2008)
54. Pastor, L., Stambaugh, R.F.: Are stocks really less volatile in the long run? *J. Financ.* **67**(2), 431–477 (2012)
55. Pennacchi, G.: *Theory of Asset Pricing*. The Addison-Wesley series in finance. Addison-Wesley, Boston (2008)
56. Perold, A.F., Sharpe, W.F.: Dynamic strategies for asset allocation. *Financ. Anal. J.* **44**(1), 16–27 (1988)
57. Raftery, A.E., Madigan, D., Hoeting, J.A.: Bayesian model averaging for linear regression models. *J. Am. Stat. Assoc.* **92**, 179–191 (1997)
58. Rubinstein, M., Leland, H.E.: Replicating options with positions in stock and cash. *Financ. Anal. J.* **37**(4), 63–72 (1981)
59. Scherer, B.: Portfolio resampling: review and critique. *Financ. Anal. J.* **58**(6), 98–109 (2002)

60. Scherer, B.: A note on the out-of-sample performance of resampled efficiency. *J. Asset Manag.* **7**(3/4), 170–178 (2006)
61. Sharpe, W.F.: Decentralized investment management. *J. Financ.* **36**(2), 217–234 (1981)
62. Sharpe, W.F., Tint, L.G.: Liabilities—a new approach. *J. Portf. Manag.* **16**(2), 5–10 (1990)
63. Wilson, R.: The theory of syndicates. *Econometrica* **36**(1), 119–132 (1968)
64. Zagst, R., Kraus, J.: Stochastic dominance of portfolio insurance strategies—OBPI versus CPPI. *Ann. Op. Res.* **185**(1), 75–103 (2011)