Model Risk: A Conceptual Framework for Risk Measurement and Hedging

Nicole Branger[‡] Christian Schlag[‡]

This version: January 15, 2004

[‡]Both authors are from the Faculty of Economics and Business Administration, Goethe University, Mertonstr. 17/Uni-Pf 77, D-60054 Frankfurt am Main, Germany. E-mail: branger@wiwi.uni-frankfurt.de, schlag@wiwi.uni-frankfurt.de

Model Risk: A Conceptual Framework for Risk Measurement and Hedging

Abstract

The vast majority of approaches to risk management, hedging, or portfolio planning assume that some model is given. However, under model risk, the true data generating process is not known. The focus of this paper is on problems related to the hedging of derivative contracts. We explain the main general concepts, provide economic interpretations, and illustrate our arguments by simple and straightforward examples.

Model risk can be dealt with in several ways, e.g. by not taking model risk into account at all ('naive' approach), by relying on worst case approaches, or by using Bayesian techniques. The integration of market risk and model risk turns out be crucial, since a risk measure that is used to find an optimal, risk-minimizing hedging strategy should capture the overall gains and losses of a position, irrespective of their reason. Since the full problem may be much too complicated to solve in realistic model setups, we furthermore discuss robust hedging strategies determined under simplifying assumptions. Our examples show that model risk is relevant and that the choice of risk measure matters, so that it should be based on sound economic arguments.

JEL: G12, G13

Keywords: Model Risk, Risk Measures, Hedging, Expected Shortfall, Bayesian statistics

1 Introduction and Motivation

Under model risk, the true data generating process for the state variables in a financial valuation or hedging problem is unknown. However, most approaches to risk management, risk measurement, hedging, pricing, or portfolio planning start from the assumption of a given model, which is basically equivalent to ignoring model risk. Up to now there has been no conceptual analysis of model risk and of possible strategies to cope with this (seemingly) new type of risk.

In this paper we focus on the hedging of contingent claims under model risk. We consider risk-minimizing hedging strategies, which requires as a first step to discuss different approaches to risk measurement in the presence of model risk. That model risk indeed matters for hedging can be easily seen by comparing delta hedging strategies based on the correct model to those generated by the (wrong) Black Scholes (BS) model. Both for stochastic volatility models and models with stochastic jumps, the performance of BS strategies deviates significantly from the performance of the correct strategies, as shown for the stochastic volatility case in Branger and Schlag (2003).

As a key contribution of our paper we discuss how to define risk measures in case of model risk, and we analyze the resulting hedging strategies and their performance under model risk in detail. In a nutshell, our results show that both model risk and the way we cope with model risk matter. First, the additional effort which is needed to take model risk into account for purposes of risk measurement and hedging can improve the efficiency of risk limiting systems and the performance of hedging strategies significantly. Second, different approaches to cope with model risk result in different hedging strategies, and a hedging strategy that is optimal for one measure may perform badly for another. We provide detailed economic interpretations for the different approaches, and these economic considerations have to precede any effort to implement sophisticated computational algorithms.

We start with a definition of model risk in Section 2. In increasing order of complexity, we may be faced with estimation risk (when the true parameters of a given true model are not known), with uncertainty about the true model within a given model class (e.g., whether Heston (1993) or Schöbel and Zhu (1999) is the correct stochastic volatility model), and finally with uncertainty about the true model class (e.g., whether the true model includes stochastic jumps or stochastic volatility or both).

The problem posed by model risk is similar, but certainly not identical to the problem of market incompleteness. Under market incompleteness, the true model (or probability measure) is assumed to be known, but the equivalent martingale measure is not unique. Under model risk, even the physical measure is unknown, and the market can be either incomplete or complete.

Various approaches to define risk measures in case of model risk and their economic interpretations are discussed in Section 3. They all start from some basic risk measure which would also be used in a world without model risk. Examples include expected shortfall, shortfall probability, or the payoff variance. Taken this choice of the basic market risk measure as given, we discuss several approaches to cope with model risk. In the 'naive' approach, we simply choose one model and ignore model risk altogether. In the worst case approach, we use the maximal risk measure over all models. And in the Bayesian approach, model probabilities are used either to define an aggregate model which integrates market and model risk, or to calculate the expectation over the risk measures in the candidate models.

For all these approaches, we provide their economic motivation and a detailed discussion of their advantages and disadvantages. Only when an approach has been chosen based on sound economic arguments, the sophisticated mathematics needed to find the optimal hedge pay off in terms of a better hedging performance. We also show that there are several pitfalls in the comparison of risk measures across the candidate models, i.e. the models considered by the decision maker. The way this comparison should be done depends on the objective for which risk is measured. If we want to set a risk limit, the total risk in a model matters. On the other hand, if we want to judge the performance of a hedging strategy, then we have to look at the additional risk of a given strategy relative to the optimal one.

Hedging approaches under model risk are discussed in Section 4. The straightforward approach is to minimize the aggregate risk measure over all trading strategies. These strategies are, of course, subject to the same criticism as the risk measures they are based on. Furthermore, in many realistic model setups the resulting optimization problem is much too difficult or too time-consuming to solve. There is trade-off between finding the optimal solution of the hedging problem and using simple strategies that are more easy to implement. These robust strategies may be determined in some hedge model, or they may be found as a combination of the optimal hedging strategies in some candidate models. One main question for future research is how to set up the simplified approach such that the strategies perform acceptably well in a variety of models, even if they are not optimal in any of these models.

We illustrate our arguments by means of simple examples. Our objective is to highlight the impact of model risk on risk measurement and hedging. We use three one-period trinomial models as candidate models, the traded assets are the stock and the money market account, and the claim to be hedged is a call option on the stock. As our basic risk measure we use the expected shortfall, which is one possibility to measure the downside risk of a position.

2 Model Setup

2.1 Definition of Model Risk

Under model risk, the true data generating process is not known. Instead, we are faced with several 'candidate models', where each candidate model may represent the true process. In a more general interpretation, the model is not known, and there are even no candidate models. However, this definition seems too general to work with, and we go with the assumption that we are at least able to specify a set of candidate models. Note that the true process is not necessarily in this set of candidate models.

To clarify the definition of model risk, consider a case where we increase the uncertainty about the model continuously. In the simplest case, we know the true model, but there is uncertainty about the parameters. For example, the true model may be the stochastic volatility model suggested by Heston (1993), but the structural parameters as well as current volatility may not be known and may have to be estimated. This is what we call *estimation risk*. Model risk also arises when we know the class of models, but not which model from this class is the true one. In the class of stochastic volatility models, e.g., we may not know whether the model of Heston (1993) or the model of Schöbel and Zhu (1999) describes the true process. When we are even uncertain about the model class itself, we are faced with the maximal amount of model risk. In this case, we cannot ultimately decide whether the stock is, e.g., exposed to stochastic volatility, to stochastic jumps, or to both of these risk factors.

One might try to solve the problem of model risk by going to the data in order to identify the true process. However, the problem of estimation risk remains. Furthermore, since any conditioning on a model would just re-introduce model risk, some non-parametric method should be applied. From the literature on implied distributions we know that there are in general not enough option prices to identify the risk-neutral probability distribution uniquely. The problem is even worse when we want to additionally identify the dynamics of the process, but it is exactly the dynamics which are needed to determine a meaningful hedging strategy. Despite these problems estimation can still be used to find probability distributions for parameters and to find model probabilities. These probability distributions are then used in the Bayesian approach to model risk.

At a first glance, model risk seems to be quite similar to market incompleteness. However, market incompleteness does not induce model risk, and on the other hand, model risk can even exist in a complete market, a case for which we will give an example below. It is therefore important to carefully explain the differences between model risk and market incompleteness.

The main difference is that in an incomplete market, the true data generating process is still known. The number of risk factors is just too high relative to the number of linearly independent traded assets, so that not all contingent claims are attainable, and the equivalent martingale measure is not unique. One standard approach to hedging in an incomplete market is to decide on some risk measure like the expected shortfall, and then to find the risk-minimizing hedge. Obviously, this optimal hedge depends on the data generating process. Under model risk we cannot condition the hedging strategy on a fixed model any longer. We simply do not know the probability distribution of the hedging errors.

Model risk does not imply that the candidate models are incomplete. Assume that model risk arises from the combination of several deterministic volatility models which differ with respect to the volatility function. Each of these models is complete, but each one implies a different price and a different hedging strategy for, say, a call option. As long as we do not know the true model, we do not know which of the 'replicating' strategies we have to implement. Therefore, under model risk, there is no perfect hedge for the call.

We can also find an example where the market is complete, but there is still model risk. Assume that the stock price follows a geometric Brownian motion with known volatility, but unknown drift. Then, the model is complete, and any claim can be replicated. However, hedging strategies that depend on the true drift (like quantile hedging strategies proposed by Föllmer and Leukert (1999)) as well as portfolio planning problems suffer from model risk.

2.2 Model Used for Numerical Examples

To illustrate our arguments we rely on simple numerical examples. Our objective is to highlight the impact of model risk and the different ways to cope with model risk. Therefore we consider a model setup that is as simple as possible to avoid complicated mathematics, which would distract from the basic economic questions.

We capture model risk by assuming three different candidate models. Each of these candidate models is a one-period model with three states at time $t = T$. We denote the physical probability measure in model i by P_i. We assume that the money market account and a stock with initial price $S_0 = 1$ are traded. To keep the setup simple, we set the risk free rate equal to zero ($r = 0$) in all models. The claim to be hedged is an at-the-money call option with strike price $K = 1$. Details on the candidate models and the physical probability measures for each of these models can be found in Table 1.

3 Risk Measures

Our objective is to find the optimal, i.e. risk-minimizing hedge in case of model risk. For this purpose, we first have to define the risk measure. In this section, we discuss general problems and approaches to risk measurement in case of model risk.

It is important to be clear about the notation and to distinguish between risk measures in case of model risk and pure model risk measures (i.e. measures for the amount of model risk). First, a *risk measure in case of model risk* is used to measure the overall amount of risk a derivative position is exposed to. It takes market risk (like stock price risk, interest rate risk, or volatility risk) and additionally model risk into account. Concepts and approaches used for this purpose are discussed in this section.

Second, there are *model risk measures* which capture model risk itself, i.e. the amount of model risk a given portfolio is exposed to. They do not incorporate the impact of market risk in the candidate models, but they focus solely on model risk. The discussion of these model risk measures is subject to further research, as well as questions like the premium paid for model risk, or the use of model risk measures to assess the difference between models.

The claim to be hedged is given by its terminal payoff C_T, and the value of the hedge portfolio at time T is denoted by H_T. Both these values depend on the model and on the state in the respective model. The difference between the value of the claim and the value of the hedge portfolio is denoted by $X = C_T - H_T$. More generally, X is the payoff for which we want to measure the risk. The risk measure of the payoff X is denoted by $\rho(X)$, and if we measure the risk given some probability measure P, we write $\rho^P(X)$.

3.1 Basic Risk Measures

Our objective is to measure both market risk and model risk. As we want to focus on concepts how to cope with model risk, we do not discuss the problems related to market risk measurement. Rather, we take some basic risk measure as given which we apply whenever there is no model risk but only market risk. This basic risk measure allows us to quantify the risk in each candidate model.

There are many different concepts to measure market risk like the expected squared hedging error (see, e.g., Schweizer (1999)), the Value at Risk (VaR), the expected shortfall or the shortfall probability (see, e.g., Föllmer and Leukert (2000)), or coherent measures of risk (see, e.g., Artzner, Delbaen, Eber, and Heath (1999)), to name just a few. We use the expected shortfall as basic risk measure in the examples. Given some model P, the risk measure can be calculated as

$$\rho^P(C_T, H_T) = E^P[(C_T - H_T)^+] = E^P[X^+].$$

This choice is mainly motivated by its simple form. The following analysis could basically be done for any other market risk measure, too.

3.2 Naive approach

The naive approach simply ignores model risk. We choose one candidate model, say model i, and behave as if the chosen model was the true one. The risk measure is therefore

$$\rho(X) = \rho^{P_i}(X).$$

This approach is very easy. But it suffers from severe short-comings as it does not even take a look at the risk measure of the position in all the other models.

If we want to limit the risk of a position, the naive approach can perform very badly. Imagine a situation where we choose, just by chance, a model i where the risk of the position is moderate and well acceptable. Now assume that in some other candidate model the same position has a very high risk far beyond the threshold of acceptability. If this model happens to be the true model, the decision to accept the position is wrong, and the risk limit is violated. The naive approach can also be bad when it comes to hedging, as we discuss in more detail in Section 4.2. The hedging strategy is simply dictated by the chosen model. If this model is an outlier, then the hedging strategy performs bad in most other models. Again, if we do not consider the other models, we are not able to correct for this problem.

3.3 Worst Case Approach

In the worst case approach, the risk measure for some position X is the supremum over the basic risk measures in the candidate models:

$$\rho(X) \quad = \quad \sup_i \left\{ \rho^{P_i}(X) \right\}.$$

The main problem of the worst case approach is how to choose the set of models. The first possibility is to consider literally all models. The second possibility is to restrict the analysis to all models in some compact set. If the models are described by some parameters, this amounts to restricting these parameters to some interval. And the third possibility is to consider only a finite number of models.

The difference between these possibilities can best be explained if we consider the objective of limiting risk. The idea of the worst case approach is that we are on the safe side if the worst case risk measure is below a given risk limit. However, this only holds if the true model is included in the set of candidate models. The implicit assumption of the worst case approach is thus that we are able to include all possible models. If we use all models — where 'all' is to be taken literally — this implicit assumption holds. Then, the worst case risk measure is indeed an upper bound on the risk of the claim. And if the worst case risk measure is below some limit, the risk of the claim is limited irrespective of the true model.

If we use a compact set of models or a finite number of models, we can never be sure that the true model is among the candidate models. The real worst case is not that the model with the highest risk measure within this set is the true one, but that some model outside this set with an even higher risk measure is the true one. The name worst case is therefore misleading, as it falsely suggests that we do not have to care about even worse cases any more. Due to this wrong suggestion, the problem of forgotten models – that arises for all approaches dealing with model risk – is particularly severe for the worst case approach.

On the other hand, even if it was possible, the use of all models would also suffer from serious problems. For many claims, the worst case risk measure is either equal to some trivial upper bound or even infinity. In our example where the basic risk measure is expected shortfall, the risk of a put short for example would be equal to the strike price, while for a short call it would simply be infinite, assuming there is no finite upper bound for the stock price. This observation motivates the use of a compact set of models, ensuring that the worst case loss exists and preventing us from always ending up with the trivial bounds. At the same time, however, it is clear that in this case we can always find a model outside the compact set in which the claim has a higher risk than the maximum over the models in our compact set.

Furthermore, the worst case approach obviously represents an extreme, since the result is entirely dictated by some catastrophe model or scenario, while all other models are ignored. This is similar to judging a stock by the lowest possible price, which is, of course, zero. Stated differently, the approach is based on the assumption that investors are infinitely averse towards model risk.

Due to all these problems, a worst case risk measure should not enter the objective function when we look for an optimal strategy. Rather, the worst case risk should be a restriction saying that the risk measure must not exceed some pre-specified boundary. The most important motivation for such a restriction is perhaps a situation where a position has to be acceptable for several decision makers which all use different models, including the regulating agency.

For the reader with a deeper interest in the foundations of worst case risk measures and strategies the papers by Kirch (2002), Gilboa and Schmeidler (1989), Föllmer and

Schied (2002), Kerkhof, Melenberg, and Schumacher (2002), and Klibanoff, Marinacci, and Mukerji (2003) will offer a further discussion of this issue.

3.4 Bayesian Approach

The Bayesian approach relies on a probability distribution over the set of candidate models. We denote the probability of model i by $p(M_i)$. The model probabilities will be used to integrate market and model risk. The probability distribution over the set of models may be given exogeneously, or it may be estimated. In the latter case, the use of Bayesian estimation seems appropriate. Starting from some (possibly non-informative) prior for the parameters and for the models, we obtain posterior distributions over the parameters and the models by combining prior knowledge and information contained in the data. The estimation can, for example, be based on Markov Chain Monte Carlo methods, and it can use information from both time series and cross sections of prices.

The critical assumption in the Bayesian approach is that we know the model probabilities and can actually capture model risk just by these probabilities. It is sometimes argued that model risk is different in that, per definition, probabilities for the models are simply unknown. However, the same could be objected to the use of probabilities for the assessment of market risk. This discussion relates to the general distinction between risk and uncertainty, where risk is a situation for which probabilities exist, while uncertainty denotes a situation where probabilities are not known. This distinction is not specific to model risk or model uncertainty. The use of probabilities for models is just one way to deal with model risk. We claim that it is a sound approach, but we do not at all claim that it is the only approach, or that other approaches are questionable.

Like in the worst case approach, we cannot be sure that the true model is within the set of candidate models. However, this problem is less important here since the Bayesian approach does not make the implicit assumption that the true model is in this set and because one may assume that in most situations a neglected model will have a rather low model probability.

3.4.1 Model Integration

Model integration amounts to integrating market risk, which is captured in the candidate models, and model risk into one big aggregate model. In our simple example with three trinomial candidate models we end up with a one-period model with nine states at time T. The probabilities of the states are just the probabilities in the candidate models, multiplied by the model probabilities. In a multi-period context, the integration of models is more complicated. The simplest solution is to calculate the distribution of the payoff X at time T in each model, and then to aggregate these probability distributions into one big distribution:

$$\rho(X) \quad = \quad \rho^P(X) \qquad \text{where} \quad P = \sum_i p(M_i) P_i.$$

An extension would be to introduce learning into the model. The investor would use the given model probabilities to integrate market and model risk over the next time period, after which he would update the model probabilities, before again moving forward one time step.

In the case of model integration, market risk and model risk are treated exactly alike. Only the probabilities of the states and of the models are important. Economically, one can argue that for deciding on the optimal strategy, the person responsible for a position is only interested in the overall gain or loss of its portfolio and the corresponding probabilities, but not in whether losses can be attributed to market risk or model risk.

If we use a Bayesian approach and integrate the models, we have to keep in mind the difference between the true model and the model used by the investor. The investor acts as if the aggregate model was the true one. However, the world does not behave according to this aggregate model, but according to the true model. The investor knows that his model is not correct, but, given his information, it is the best model he can follow.

3.4.2 Risk Integration

In case of risk integration, the risk measure is defined as the expectation of some function ϕ over the risk measures in the candidate models

$$\rho(X) \quad = \quad \sum_i p(M_i)\phi\left(\rho^{P_i}(X)\right).$$

The function ϕ is increasing and describes the attitude of the decision maker towards model risk. To express model risk aversion one would choose ϕ to be convex, whereas a linear or concave ϕ would represent the case of model risk neutrality or even a preference for model risk.

We start the discussion assuming ϕ to be linear. The decision maker only cares about the average basic risk measure, not about its variance. When the basic risk measure is the expectation of some function of the payoff, it can be shown that the expected risk measure coincides with the risk measure in case of simple model integration. For basic risk measures like the expected shortfall or the expected squared hedging error, model integration therefore amounts to the assumption of model risk neutrality.

For a convex function ϕ, there is a trade-off between a low expected risk measure and a low variation of the risk measure across models. The decision maker is model risk averse. He does not only prefer the individual risk measures to be as low as possible, but he also wants them to be as equal as possible across models, so that it is less important which of the candidate models is the true one.

For $\phi(x) = x^p$,

$$\{\rho(X)\}^{1/p} \quad = \quad \left\{\sum_i p(M_i)\left(\rho^{P_i}(X)\right)^p\right\}^{1/p}$$

converges for $p \to \infty$ towards the worst case risk measure. Thus, the more model risk averse the investor becomes, the more his decision is biased towards the decision implied by the worst case risk measure.

4 Hedging

4.1 Basics

Like in many approaches for the case of incomplete markets, we consider risk-minimizing hedging strategies. The first step is to choose a risk measure, like, for example, expected shortfall or expected squared hedging error. Under model risk, one additionally has to decide on an approach for taking model risk into account, as discussed in Section 3. In the second step, we have to find the optimal, risk-minimizing hedging strategy. If the initial capital for the hedge is not given exogenously, it also has to be determined.

In our simple example, we assume that the candidate models are given exogenously. In general, all candidate models should be calibrated to the same data, that is to the same time series of stock and/or option prices, and to the same cross section of option prices. Since prices are noisy, these data will in general not allow to identify the true model uniquely, so this calibration cannot eliminate model risk. If the market price of the claim to be hedged is known, then the models should also be calibrated to this market price. The most prominent example for the latter is the use of the BS model based on the implied volatility of the claim.

It is important to note here that even a perfect fit to given prices does not imply a perfect hedge. This can be seen from the literature on implied trees, where for a given set of option prices, there are many implied trees that give exactly these prices, but give rise to quite different hedging strategies (see Rubinstein (1994) or Franke, Stapleton, and Subrahmanyam (1999)).

In a multiperiod or even a continuous-time framework, the candidate models should be recalibrated at every point in time. The investor therefore knows that his model does not describe the true data generating process. By recalibrating, however, he tries to adjust his hedging decision to the available information as well as possible, which also relates to the remarks at the end of Section 3.4.1.

A further question is whether the initial capital is given exogenuously or determined as part of the optimal hedging strategy. If a market price for the contingent claim to be hedged is known and the claim was actually sold for this price, then the initial capital should be equal to this market price. The initial capital is also given exogenously if the amount of money avaible for the hedge is determined by some third party like a central risk management unit which assigns hedging capital to the different trading desks. For an illiquid derivative contract, a market price may not exist. In this case, the initial capital is not given exogenously, but can be determined as part of the solution to the problem of finding the optimal hedge. The objective is either to find the risk-minimizing hedge, or to find the cheapest hedge for which the risk measure does not exceed some given value.

Once we have found the optimal hedging strategy, we assess its performance in each of the candidate models and with respect to the other aggregate risk measures discussed in Section 3. The objective is to show whether strategies that take model risk into account are indeed better than naive ones, and to analyze the differences between the various aggregate risk measures. We want to show that both model risk and the choice of the aggregate risk measure matter.

4.2 Naive Approach: Hedging Based on one Specific Model

The naive approach ignores model risk. The investor starts by choosing just one model, and then he behaves as if this model described the true data generating process. The main advantage of the naive approach is its simplicity. It is much easier to find an optimal hedge within some simple candidate model than to find the optimal hedge if all candidate models have to be considered simultaneously and an aggregate risk measure is used.

The objective of this section is to analyze naive hedges and their performance. In the discussion, we refer to two numerical examples, where we hedge a call option on the stock with strike price $K = 1.0$. The initial capital for the hedge is fixed at 0.075. Figures 1 and 2 show the expected shortfall within the three candidate models as a function of the hedge ratio, based on scenarios 1 and 2 from Table 1. When we fix one model, the total risk measure can be used to decide whether a given position is acceptable or not, and to determine the charge on the risk limit. Furthermore, the strategy with the minimal total risk measure is the optimal hedging strategy within the respective model.

The figures show that the total risk measures differ across the models, and that the risk-minimizing hedge ratios also differ. However, the differences in total risk measures are of no use if the objective is to find the optimal hedge. As an example, consider Figure 1 and start from the optimal hedge in model 1. The risk-minimizing hedge ratio is equal to 0.1875. If model 2 were the true one, the total risk of this hedging strategy is even lower. Nevertheless, the strategy is not optimal in model 2, where the risk-minimizing hedge ratio is 0.375.

For hedging purposes, it is not the total risk but the additional risk generated by the choice of a wrong model that matters. The additional risk of a hedging strategy in a candidate model is defined as the difference between the risk generated by this strategy in a given model minus the risk of the optimal hedge in this model. As an example, consider Figure 1. If model 1 is the correct one, then the additional risk of using model 2 is moderate, while that of using model 3 is considerable.

4.3 Aggregate Approach: Taking Model Risk into Account

4.3.1 Worst-Case Approach

In the worst case approach, the investor measures the risk of a given payoff by the maximal risk measure over all candidate models. The objective is to find the hedging strategy for which this maximal risk is minimal. For the optimal strategy under the worst case approach, two scenarios are possible. In the first case, the worst case hedge coincides with the optimal hedge in one of the candidate models. As an example, consider Figure 2, where the worst case hedge coincides with the optimal hedge in model 3. In contrast to this, Figure 1 shows the case when the worst case hedge is determined by the intersection of the risk curves for models 1 and 3.

Worst-case hedging strategies suffer from the same shortcomings as the worst case risk measure discussed in Section 3.3. First, the name suggests that the investor is on the safe side, irrespective of the true model. But this only holds under the additional assumption that the true model is in the set of candidate models. As an example, consider Figure 2 and assume that just the models 1 and 2 form the set of candidate models. The optimal hedge ratio is then 0.375, and the worst case risk of this strategy is 0.035. If model 3 is the true one, then the risk of the hedge is 0.18, more than five times higher than assumed.

Second, the hedge is determined by a catastrophe scenario. The amount of risk in any other model does not matter, and it does not matter whether there exists some other strategy for which the risk in nearly all other models is significantly lower. Since they ignore a considerable amount of information, the general efficiency of worst case strategies is questionable. Consider again the example in Figure 2, and assume now that model 3 is also in the set of candidate models. The optimal hedge is then determined by model 3, even if its model probability is extremely low.

Finally, note that naive hedging strategies may be very bad when the performance is measured by the worst case risk measure. In Figures 1 and 2, the use of a naive hedge may increase the risk by more than 50% compared to the worst case hedge.

4.3.2 Bayesian Approach

In case of model integration, market risk and model risk are aggregated in one big model without any further distinction. The left panels of Figures 3 and 4 show the Bayesian risk measure as a function of the hedge ratio, where we have used three different choices for the model probabilities. Additionally, the vertical lines give the optimal hedge ratios in the candidate models and the optimal hedge ratio for the Bayesian risk measure. In most cases, the latter does not coincide with any of the candidate hedge ratios. This also implies that irrespective of the choice of the hedge model, a naive hedge will not be optimal.

The optimal hedge depends on the candidate models and on the model probabilities. The higher the probability of a model, the more the Bayesian hedge tends towards the hedge in this model. In the limit, a model probability of one just results in the naive hedge, although the choice of a hedge model is now based on a sound statistical criterion.

The problem posed by outlier models in the worst case approach is at least mitigated in the Bayesian approach, when the outlier models have only low model probabilites.

For a risk measure that is based on expectations like the expected shortfall, model integration is a special case of risk integration. In case of model risk neutrality, risk integration just reduces to model integration. If the aggregation function ϕ applied to the risk measures from the candidate models is not linear but convex, investors are model risk averse. Then the optimal hedging strategy not only depends on the mean risk measure, but on some trade-off between the mean and the variance across the models. Furthermore, if model risk aversion goes to infinity, then, as shown above, the Bayesian risk measure and the associated hedge converge towards the results for the worst case approach.

The right panels of Figures 3 and 4 show the Bayesian risk measure for $\phi(x) = x^2$ as a function of the hedge ratio. Again, the vertical lines denote the optimal hedge ratios. A comparison with the left graphs where the investor is model risk neutral confirms that the optimal hedge strategy tends towards the worst case hedge. The figures also show that the hedge ratio not only depends on the candidate models but also on the model probabilities, so that for a moderate model risk aversion, catastrophe scenarios with a low probability do not have such a strong impact on the results.

4.4 Between Naive Hedging and Full Aggregation

Hedging strategies that take model risk into account are conceptually superior to naive hedging strategies. In a realistic model setup, however, the solution to the hedging problem in case of model risk may not be known or may be much to complicated to implement. For example, in a stochastic volatility model, the strategy that minimizes the expected shortfall is hard to find even under model certainty, and the additional introduction of model risk complicates the resulting optimization problem considerably.

These problems motivate the use of *robust hedging strategies*, which are easier to implement. They should be set up in such a way that they perform well for the whole set of candidate models, and they should not be too sensitive with respect to model risk.

The question of how to determine such hedging strategies is left to future research. In the simplest case, robust hedging strategies are optimal in some well chosen hedge model, or they arise as an appropriate combination of the optimal hedges in some selected candidate models.

Once the hedge model is found, model risk is again ignored in what follows. At a first glance, this looks very similar to the naive hedging approach. However, in the latter, the hedge model is chosen in an ad-hoc way whereas here the hedge model is chosen based on some economic criterion, e.g. an aggregate risk measure. As an example, consider Figure 3 again. In the upper graph in the left panel (where investors are model risk neutral) model 2 would be the optimal hedge model. The additional risk compared to the optimal strategy in case of model integration is quite low, so that the robust hedge performs well.

Another possibility is to combine the hedging strategies from different hedge models. In the Bayesian approach, the most natural way to do this is to use the mean of the hedge ratios in the candidate models. In the worst case approach the risk-minimizing hedge in a given candidate model is applied to all the other models, the resulting risk measures are computed, and the strategy is assigned the maximum of these numbers. The hedge model will then be the model which minimizes this quantity. For example, Figure 1 shows that for the risk-minimizing strategy in model 2 the worst case would be that model 3 is the true model, and this strategy is at the same time the one minimizing the modified worst case measure just described.

For the Bayesian approach, Figure 5 shows that there may be scenarios where the optimal hedging strategy under risk integration and the mean hedging strategy differ considerably. One example is the left graph in the lower row. The optimal hedge ratio under model integration is 0.8125, while the mean hedge ratio is 0.5281. However, the risk measure does not increase much. An example where not only the hedge ratios are different, but where the wrong hedge ratio also implies an economically significant increase in the risk measure can be found in the right graph in the lower row. Here, the optimal hedge ratio is 0.875 with an aggregate risk measure of 0.0796, whereas the mean hedge ratio is 0.6200 with an aggregate risk measure of 0.0915. There are also other examples where both the hedge ratios and the risk measures are quite similar, and in these situations, the approximate approaches work in a satisfactory manner.

5 Conclusion

The paper gives a roadmap for the integration of model risk into risk measurement and hedging. For both worst case and Bayesian approaches we show how to define aggregate risk measures taking model risk into account. The optimal hedging strategies perform significantly better than the naive ones, which simply ignore model risk. We provide economic interpretations and motivations for the approaches presented. We argue that worst-case approaches suffer from severe shortcomings, and that, on a conceptual level, the integration of market risk and model risk is crucial.

The main point of this paper is that the introduction of model risk raises many new questions beyond those already posed by market risk or by incomplete markets. We think that the identification and discussion of the economic problems and the derivation of solution approaches based on economic arguments has to precede any further analysis of the solution methods themselves.

Further research on model risk could proceed in several directions. First, the analysis for the simple model setup in our paper should be extended to more sophisticated models which include factors like stochastic volatility and stochastic jumps. A second topic is the degree of difference between models. The more different two models are, the more different are their implications for pricing, hedging, or portfolio planning, and the more we have to care about model risk and the identification of models or model probabilities. We thus need concepts to measure the difference between models. Third, it is still an open question how to set up a robust hedge, i.e. how to choose the hedge model, which elements (like SV or SJ) to include in the hedge model, and how to combine strategies from different hedge models.

References

Artzner, P., F. Delbaen, J.-M. Eber, and D. Heath, 1999, Coherent Measures of Risk, *Mathematical Finance* 9, 203–228.

Branger, N., and C. Schlag, 2003, Can Tests Based on Option Hedging Errors Correctly Identify Volatility Risk Premia?, Working Paper.

Föllmer, H., and P. Leukert, 1999, Quantile Hedging, *Finance and Stochastics* 3, 251–273.

Föllmer, H., and P. Leukert, 2000, Efficient hedges: cost versus shortfall risk, *Finance and Stochastics* 4, 117–146.

Föllmer, H., and A. Schied, 2002, Robust preferences and convex measures of risk, Working Paper.

Franke, G., R. Stapleton, and M. G. Subrahmanyam, 1999, Why are Options Overpriced? The Black-Scholes Model and Alternative Characterizations of the Pricing Kernel, *European Finance Review* pp. 79–102.

Gilboa, I., and D. Schmeidler, 1989, Maxmin Expected Utility with a Non-Unique Prior, *Journal of Mathematical Economics* 18, 141–153.

Heston, S. L., 1993, A closed-form solution for options with stochastic volatility with applications to bond and currency options, *Review of Financial Studies* 6, 327–343.

Kerkhof, J., B. Melenberg, and H. Schumacher, 2002, Model Risk and Regulatory Capital, Working Paper.

Kirch, M., 2002, Efficient Hedging In Incomplete Markets Under Model Uncertainty, Working Paper, TU Wien.

Klibanoff, P., M. Marinacci, and S. Mukerji, 2003, A Smooth Model of Decision Making Under Ambiguity, Working Paper.

Rubinstein, M., 1994, Implied Binomial Trees, *Journal of Finance* 49, 771–818.

Schöbel, R., and J. Zhu, 1999, Stochastic Volatility With an Ornstein-Uhlenbeck Process: An Extension, *European Finance Review* 3, 23–46.

Schweizer, M., 1999, A Guided Tour through Quadratic Hedging Approaches, Working Paper, Technische Universität Berlin.

Scenario 1						
	S_1^{up}	S_1^{mid}	S_1^{down}	$P(S_1^{up})$	$P(S_1^{mid})$	$P(S_1^{down})$
Model 1	1.40	1.10	0.60	0.20	0.35	0.45
Model 2	1.20	1.00	0.80	0.30	0.10	0.60
Model 3	1.60	1.15	0.40	0.30	0.60	0.10

Scenario 2						
	S_1^{up}	S_1^{mid}	S_1^{down}	$P(S_1^{up})$	$P(S_1^{mid})$	$P(S_1^{down})$
Model 1	1.40	1.10	0.80	0.20	0.10	0.70
Model 2	1.20	1.00	0.80	0.30	0.10	0.60
Model 3	1.60	1.10	0.40	0.50	0.30	0.20

Scenario 3						
	S_1^{up}	S_1^{mid}	S_1^{down}	$P(S_1^{up})$	$P(S_1^{mid})$	$P(S_1^{down})$
Model 1	1.40	1.10	0.60	0.50	0.30	0.20
Model 2	1.20	1.00	0.80	0.30	0.10	0.60
Model 3	1.60	1.15	0.40	0.10	0.70	0.20

Table 1: Scenarios

In each scenario, there are three candidate models. The table gives the stock prices S in these candidate models and the corresponding probability distributions P.

	Model 1	Model 2	Model 3
Model Probabilities 1	1/3	1/3	1/3
Model Probabilities 2	0.49	0.49	0.02
Model Probabilities 3	0.49	0.02	0.49

Table 2: Model Probabilities

The Bayesian approach relies on model probabilities. The three different probability measures over the candidate models are given in the table.

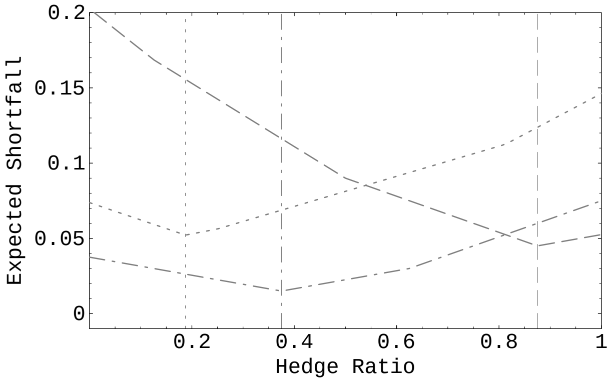

Figure 1: Expected Shortfall as a function of the number of stocks (Scenario 1)

The expected shortfall is shown as a function of the number of stocks in the hedge portfolio. We calculate the naive risk measure in model 1 (dotted line), in model 2 (dotted-dashed line), and in model 3 (dashed line). The claim to be hedged is a call on the stock with strike 1, the initial capital is equal to 0.075. The model setup is given by Scenario 1 from Table 1.

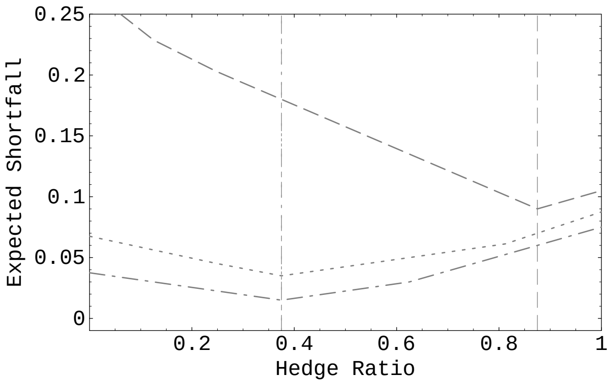

Figure 2: Expected Shortfall as a function of the number of stocks (Scenario 2)

The expected shortfall is shown as a function of the number of stocks in the hedge portfolio. We calculate the naive risk measure in model 1 (dotted line), in model 2 (dotted-dashed line), and in model 3 (dashed line). The claim to be hedged is a call on the stock with strike 1, the initial capital is equal to 0.075. The model setup is given by Scenario 2 from Table 1.

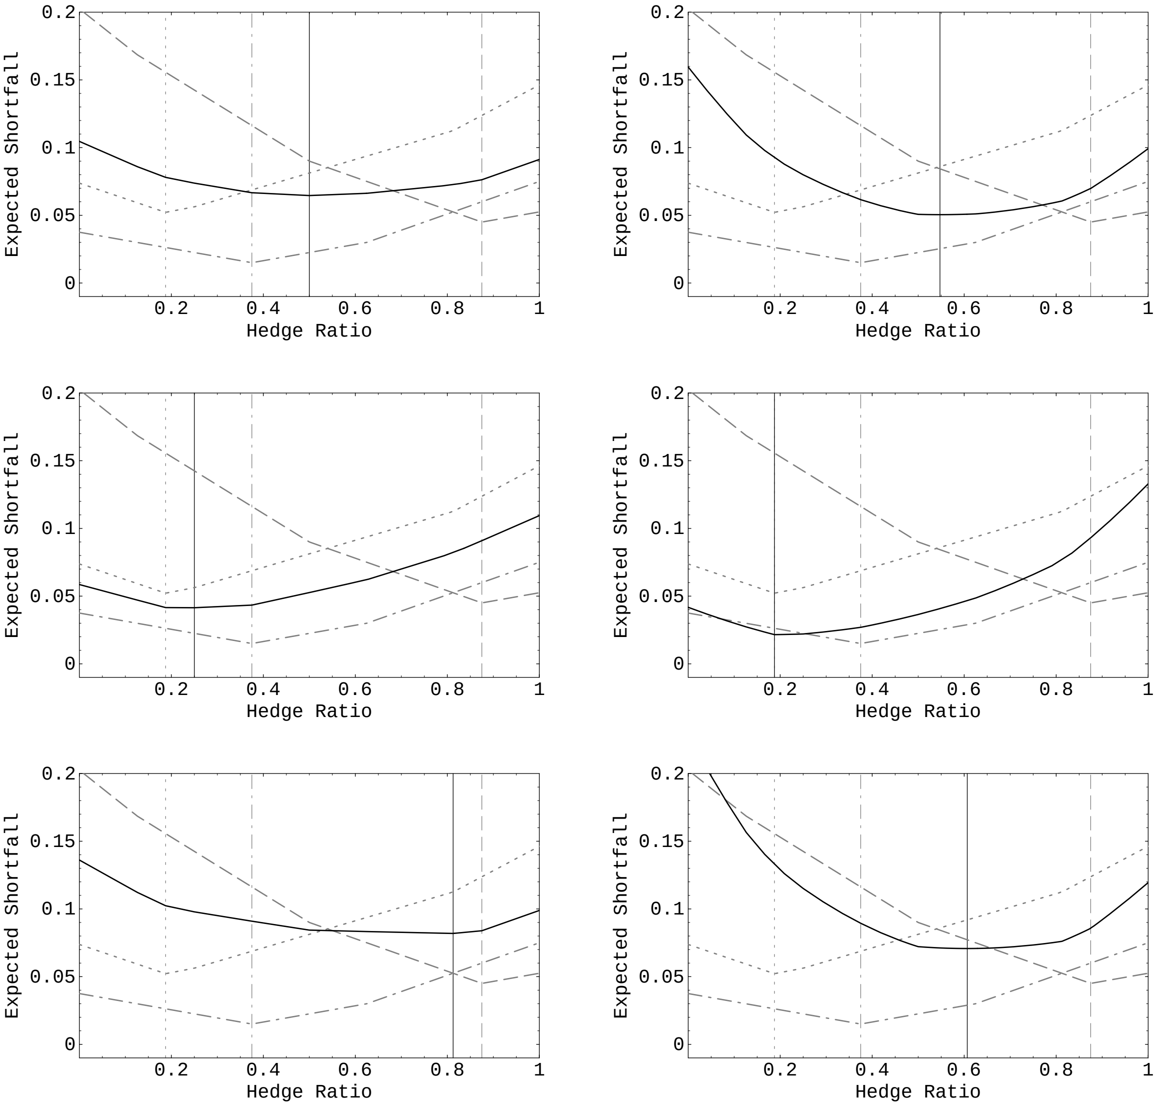

Figure 3: Bayes Approach: linear or convex function (Scenario 1)

The expected shortfall is shown as a function of the number of stocks in the hedge portfolio. We calculate the naive risk measure in model 1 (dotted line), in model 2 (dotted-dashed line), and in model 3 (dashed line). The claim to be hedged is a call on the stock with strike 1, the initial capital is equal to 0.075. The model setup is given by Scenario 1 from Table 1. We furthermore calculate the risk measure for Bayes risk integration (solid black line) in case of model risk neutrality (left graphs) and model risk aversion, captured by $\phi(x) = x^2$ (right graphs). The model probabilities are $(1/3, 1/3, 1/3)$, $(0.49, 0.49, 0.02)$, and $(0.49, 0.02, 0.49)$ in the upper, middle, and lower row, respectively.

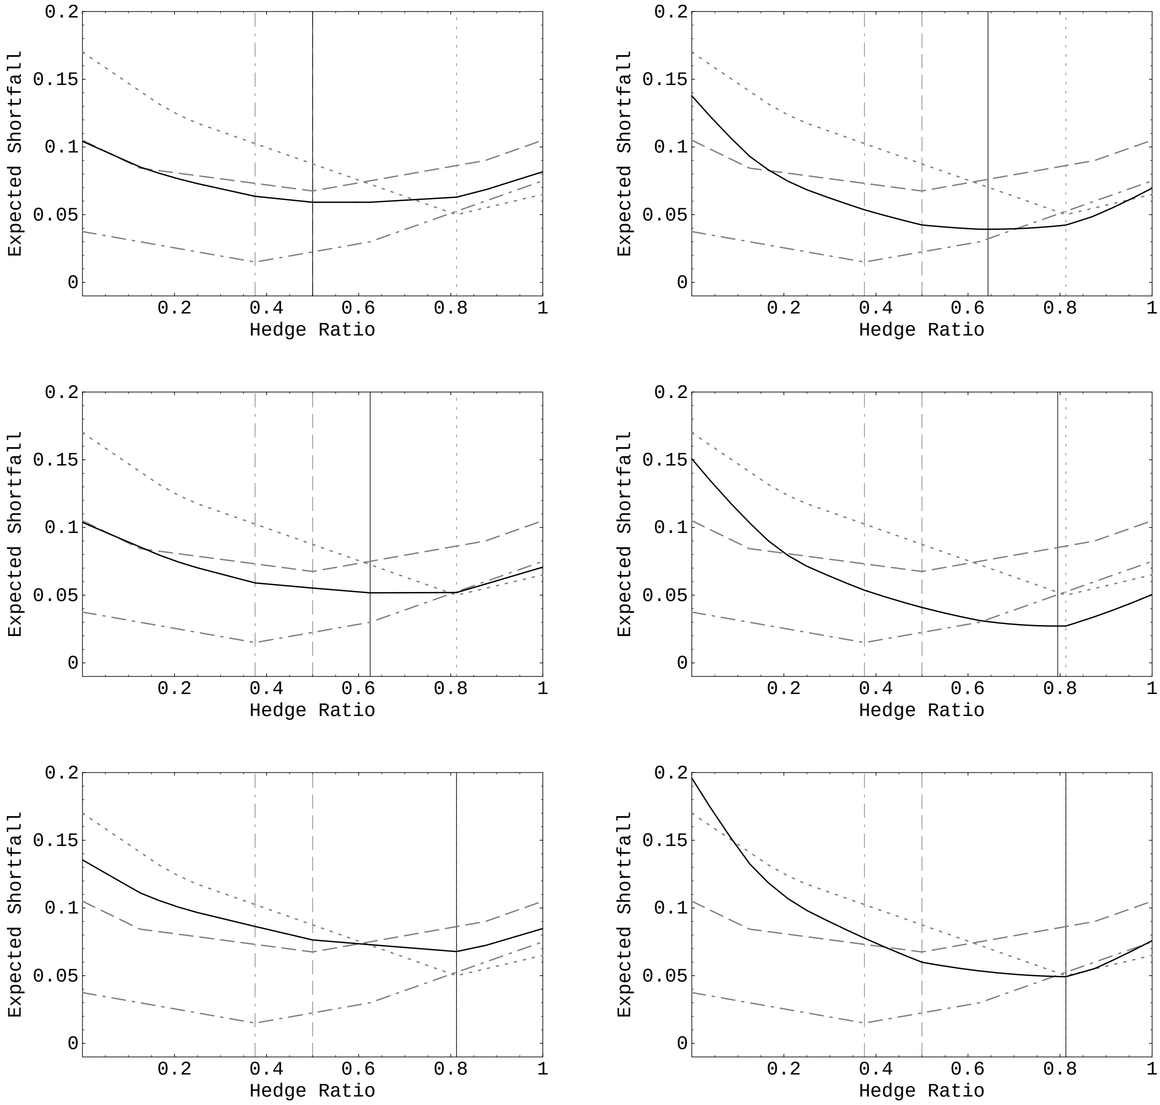

Figure 4: Bayes Approach: linear or convex function (Scenario 3)

The expected shortfall is shown as a function of the number of stocks in the hedge portfolio. We calculate the naive risk measure in model 1 (dotted line), in model 2 (dotted-dashed line), and in model 3 (dashed line). The claim to be hedged is a call on the stock with strike 1, the initial capital is equal to 0.075. The model setup is given by Scenario 3 from Table 1. We furthermore calculate the risk measure for Bayes risk integration (solid black line) in case of model risk neutrality (left graphs) and model risk aversion, captured by $\phi(x) = x^2$ (right graphs). The model probabilities are $(1/3, 1/3, 1/3)$ in the upper row, $(0.49, 0.49, 0.02)$ in the middle row, and $(0.49, 0.02, 0.49)$ in the lower row.

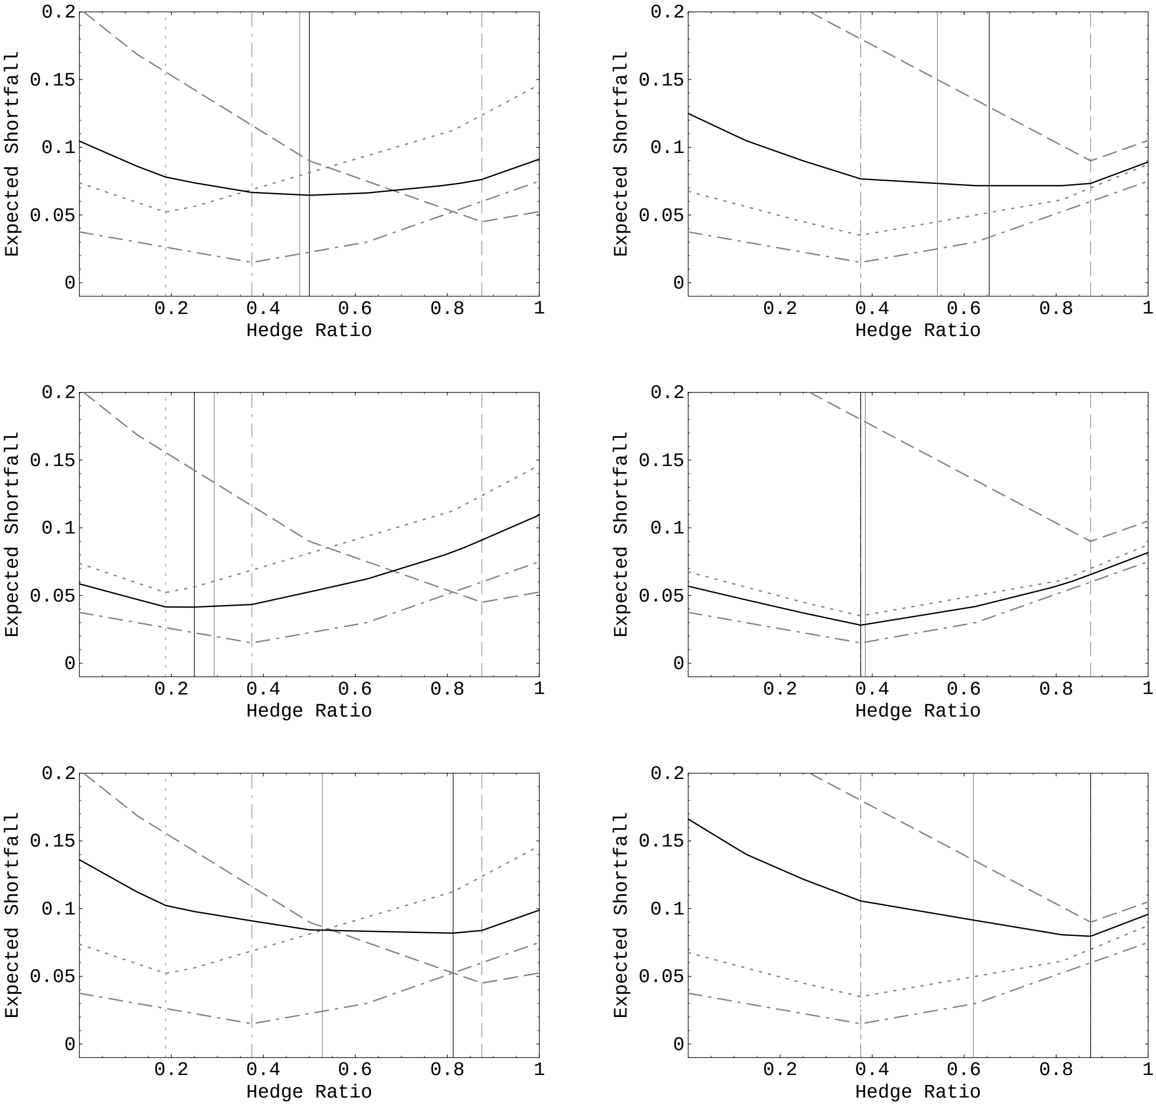

Figure 5: Bayes Approach: model integration versus mean hedge ratio (Scenario 1 and 2)

The expected shortfall is shown as a function of the number of stocks in the hedge portfolio. We calculate the naive risk measure in model 1 (dotted line), in model 2 (dotted-dashed line), and in model 3 (dashed line). The claim to be hedged is a call on the stock with strike 1, the initial capital is equal to 0.075. The model setup is given by Scenarios 1 and 2 from Table 1. We furthermore calculate the risk measure for Bayes risk integration (solid black line) in case of model risk neutrality and give the optimal hedge ratio for this measure (solid black vertical line) and the mean hedge ratio over the candidate models (grey solid vertical line). The left graphs are based on scenario 1, the right graphs are based on scenario 2. The model probabilities are $(1/3, 1/3, 1/3)$ in the upper row, $(0.49, 0.49, 0.02)$ in the middle row, and $(0.49, 0.02, 0.49)$ in the lower row.