Model risk pricing and hedging

de Oliveira Souza, Thiago

Nordea Bank

21 September 2024

Model risk pricing and hedging

Thiago de Oliveira Souza

Nordea

September 21, 2024

Abstract

I use Financial Economics theory to derive a measure of model risk with clear and actionable implications. The resulting "Model Risk Price" is based on the covariance between the payoffs associated with the model and the Stochastic Discount Factor, setting it fundamentally apart from the model accuracy statistics that have been typically used as model risk measures. Given that it is measured in financial terms, the Model Risk Price and its associated hedging strategies can also be intuitively communicated to non-technical audiences, such as investors, CEOs, other C-suite executives, and risk managers. From a practical standpoint, the paper addresses one of the most critical questions posed to risk managers by the firm's investors: What is the precise impact of model risk on their investments, and what concrete actions can be taken to mitigate it. More broadly, the paper makes a seminal contribution to the literature by formally defining and economically measuring the risk of using a model, rather than simply estimating the uncertainty in its output as is currently done.

JEL Codes: G32, G11, G12, G20.

Keywords: Model Risk, Hedging, Equity investors, Asset Pricing, Actionable results

Nordea Bank Abp, Satamaradankatu 5, FI-00020, Helsinki. Email: thiago.souza@nordea.com. Disclaimer: The views expressed in this paper are those of the author. They do not necessarily represent those of Nordea or its staff.

1 Introduction

Consider an insurance policy and a portfolio of stocks. Both have uncertain future payoffs. Hence, both should be considered risky if risk and uncertainty were equivalent. But while a portfolio of stocks is considered risky, the insurance policy is usually considered a hedge. That happens because risk and uncertainty are not equivalent.

Indeed, what determines the risk of a given stream of cash flows are the states in which the payoffs happen. Insurance policies tend to provide positive payoffs exactly when the policy holder experiences a loss. It is the favourable timing of these payments that makes the insurance policy desirable (i.e., a "hedge" and not a "risky" investment).

This analysis also applies to models: Two models with uncertain payoffs may have completely different risk profiles. But evaluating this proposition requires a formal and relevant measure of model risk. It cannot be done using the model uncertainty measures existing in the literature. Providing such a model risk measure is the central contribution of the present paper.

The first step to analyse the risk of a model is to obtain the stream of payoffs associated with using the model. I obtain these payoffs under the assumption that the firms use models to manage the capital injections required from their investors. When the losses are larger than predicted by the models employed by the firm, the investors experience negative cash flows (in the form of required capital injections). When the situation normalizes, the firm returns these injections to the investors as dividends.

As explained, the central question to define the risk of using the model is the timing of the payoffs described above. Therefore, the other fundamental ingredient in the analysis is a theoretical framework describing the economic states in which the payoffs happen. In the present paper, I rely on standard Asset Pricing theory, as in Cochrane (2005), for this task. In this general framework, the value of (an asset that provides) a future payoff, x_{t+1}, is simply

$$p_t = E_t[m_{t+1}x_{t+1}], \tag{1}$$

where p_t is the price of the asset at time t, $E_t[\cdot]$ is the expectation operator conditional on the information available at time t, and m_{t+1} is the Stochastic Discount Factor (SDF, also called "pricing kernel" or "state price deflator") at time $t+1$, which adjusts future cash flows for time and risk.

In the formulation in Eq. (1), risk is given by the covariance between payoffs and the SDF. Hence, a related step is to define a variable that can be used as the SDF for practical purposes, as I discuss in Section 2.2. Essentially, the SDF is the value of a unit payoff in each possible state. As a consequence, the SDF can also be interpreted as a measure of "distress" (of the investors for whom the model risk assessment is being made). Concretely, the time-series of observed realizations of a candidate variable should contain the value of unit payoffs in each date (again, from the perspective of the relevant investors).

Once we have both the stream of payoffs and the SDF proxy, we are ready to measure model risk by essentially applying Eq. (1) to the stream of payoffs associated with the model. The result is the "Model Risk Price", which is the cost – in monetary terms – of using the model in question.[1] In particular, the Model Risk Price is fundamentally distinct from the model accuracy statistics that have been typically used to measure model risk, as the ones in Christoffersen (1998) or Kupiec (1995), for example.

The fact that the Model Risk Price is expressed as a price is another advantage of the framework developed in the present paper because prices can be intuitively communicated to non-technical audiences such as investors, CEOs, other C-suite executives, and risk managers. This cost, unlike model accuracy statistics, is the answer to the fundamental question posed by investors regarding the concrete and measurable impact of model risk on their investments. In addition, the general theoretical framework used to measure model risk also provides concrete actionable results to reduce it, thereby answering the follow up question regarding what can be done in practice to mitigate any eventual problems.

In the empirical part of the paper, I provide several concrete examples of model risk assessments with different SDF proxies, models, periods, and portfolios. One of the

[1] Intuitively similar to the "Value-at-Risk", I consider a negative price as a positive "Model Risk Price", indicating that the result is, in fact, a cost.

key conclusions of this exercise is that the model risk measure that I propose is indeed not reflected by the statistics related to model uncertainty used in the literature so far. Section 3.3 concludes by demonstrating how the models evaluated in the previous section can be effectively hedged using add-ons, as examples of tangible and actionable outcomes derived from the model risk assessment.

The remainder of the paper is organized as follows: In Section 2, I provide the theoretical framework, derive the Model Risk Price measure, and address a few theoretical questions. Section 3 contains the empirical exercises. It includes examples of model risk assessments and of model risk hedging. Section 4 concludes.

2 The asset pricing framework

Assume that the equity investors of a financial institution want to avoid making constant capital injections into the firm. Specifically, they seek assurance that the frequency of capital injections will be at most p.

Let θ be a random variable that represents the monetary value that determines the need for capital injections. Intuitively, θ could be the value of a bank's trading or banking books, including eventual losses associated with defaulted counterparties, loans, and etc.

The realization of θ at time $t+1$ depends on the portfolio chosen by the firm at time t, π_t,

$$\theta_{t+1} \equiv f_t(\pi_t) + \epsilon_{t+1}, \tag{2}$$

where $f_t(\cdot)$ is a give function, with possibly time-varying parameters, and ϵ_t is the error term at time t, representing the uncertainty in the realization of θ. If the realization of θ is too low (e.g., the losses are too high), $\theta_t < \overline{\theta}$, the firm needs a capital injection.

Hence, assuming that the firm managers act to implement the preferences of its shareholders, they choose an optimal portfolio, π_t^*, such that the probability of a capital injection in the next period is less than or equal to p. In this case, if we denote the future realization of θ corresponding to this optimal portfolio by θ^*_{t+1} based on Eq. (2),

$$\theta^*_{t+1} \equiv f_t(\pi^*_t) + \epsilon_{t+1}, \tag{3}$$

the probability of θ^*_{t+1} being smaller than $\overline{\theta}$ must be less than or equal to p,

$$P\left(\theta^*_{t+1} < \overline{\theta}\right) \leq p \quad \implies \quad q_p\left(\theta^*_{t+1}\right) \geq \overline{\theta}, \tag{4}$$

where $P(\cdot)$ is probability and $q_p(\cdot)$ is the p-quantile function.

Given that $q_p\left(\theta^*_{t+1}\right)$ is unknown at time t, the firm employs a model, such as a Value-at-Risk (VaR) at the $(1-p)$ level, to estimate it. The firm then chooses the optimal portfolio, π^*_t in Eq. (3), such that

$$\hat{q}_p\left(\theta^*_{t+1}\right) = \overline{\theta}, \tag{5}$$

where $\hat{q}_p(\cdot)$ is the estimate for $q_p(\cdot)$, and Eq. (4) holds with equality under the assumption that the firm's profits increase in the amount of risk taken (and hence that the firm utilizes its entire risk limit to maximize profits).

2.1 Capital injections and dividends

In this formulation, the payoff to the investors corresponding to capital injections is either zero (when the loss threshold is not breached) or equal to the difference between the loss and the threshold. More formally, we can express these payoffs as

$$x_t = \min\left\{0, \quad \theta_t - \overline{\theta}\right\}, \tag{6}$$

$$= \min\left\{0, \quad \theta_t - \hat{q}_p\left(\theta^*_t\right)\right\}, \tag{7}$$

where the second equation follows from Eq. (5) for the optimal portfolio chosen in the previous period, π^*_{t-1}.

The firm later returns all capital injections to the investors as extraordinary dividends when the conditions normalize. Assuming that this extra capital earns the risk-free rate (for example, it is kept as reserves in the form of Treasury-bills), we obtain

$$x_{pb,t_{pb}} = x_t \left(1 + R_{f,t}\right)^{t_{pb}-t}, \tag{8}$$

where $x_{pb,t_{pb}}$ is the dividend paid back to the investors at time t_{pb}, corresponding to the capital injection x_t, compounded at the risk-free rate at time t, $R_{f,t}$.

From the perspective of the equity investors, the risk of a model used to obtain $\hat{q}_p\left(\theta^*_{t+1}\right)$, is equivalent to the risk of the uncertain streams of payoffs associated with the model: The sum of x_t in Eq. (7), and $x_{pb,t_{pb}}$ in Eq. (8), as I discuss next.

2.2 Model risk pricing

Eq. (1) presented in the introduction states that the price of an asset today is equal to the expected value of the product of the SDF and the asset's payoff in the next period $(t+1)$. The equation applies to any stream of uncertain cash flows because it simply means that current prices should reflect the discounted value of expected future payoffs.

We can also apply Eq. (1) on historical data to obtain the value of a stream of cash flows, as long as we also have a time-series for the SDF, $\{SDF_t\}$. In particular, we can obtain a measure of model risk by pricing the realized payoffs associated with the model, which I define as the Model Risk Price,

$$MRP_t \equiv -\left[\sum_{t'=t_0}^{t} \left(SDF_{t'}\ x_{t'}\right) + \sum_{t'_{pb}=t_0}^{t} \left(SDF_{t'_{pb}}\ x_{pb,t'_{pb}}\right)\right], \tag{9}$$

where t_0 is the beginning of the time interval considered, x_t and $x_{pb,t_{pb}}$ are given by Eq. (7) and Eq. (8), and the negative sign means that the MRP can be interpreted as a cost, similar to a Value-at-Risk (VaR) model. Hence, model risk is directly proportional to MRP, which is the central measure of model risk that I derive in the present paper.

2.2.1 No-drift SDF and the timing of extraordinary dividends

One of the important characteristics of the SDF time-series is that its drift is equal to the risk-free rate. Hence, it is possible to create an economy in which the time value of money is neutral by considering a series without drift, either by de-trending it or by construction. This implies that the cash flows are adjusted for risk, but not simply due to the passage of time. In this case, we can substitute $R_{f,t} = 0$ in Eq. (8) to obtain

$$x_{pb,t_{pb}} = x_t. \tag{10}$$

In addition, assuming that the firm only distributes the capital injections as extraordinary dividends when the SDF returns to its typical, median value, SDF_{med}, we further obtain

$$\forall t'_{pb}, \quad SDF_{t'_{pb}} = SDF_{med}. \tag{11}$$

The assumptions above simplify Eq. (9) to

$$MRP_t \equiv -\left[\sum_{t'=t_0}^{t} (SDF_{t'}\ x_{t'}) + SDF_{med} \sum_{t'=t_0}^{t} x_{t'}\right], \tag{12}$$

where the only difference between the two terms inside the brackets are the SDF values prevailing when the cash flows happen.[2]

The next step is to obtain an appropriate SDF proxy for Eq. (12), which is the subject of the remainder of this section.

[2]This equation also assumes that all capital injections are returned to the investors between t_0 and t, so that the sum can be taken over t' instead of t'_{pb}.

2.2.2 The importance of the SDF to measure risk

As Cochrane (2005) explains in detail, Eq. (1) can be manipulated into

$$p_t = \frac{E_t[x_{t+1}]}{1+R_f} + cov_t(m_{t+1}, x_{t+1}), \tag{13}$$

where R_f is the risk-free rate and $cov_t(\cdot)$ is the conditional covariance at time t. The first term in Eq. (13) is the standard discounted present-value formula. This is the value of the payoff in a risk-neutral world (i.e., discounted at the risk-free rate). The second term is exactly the risk adjustment. Payoffs that covary positively with the SDF have higher values and vice versa.

This highlights a crucial fact: Risk depends on a *co*-variance (with the SDF), not a variance: Quantities that simply reflect uncertainty or inaccuracies about (model) payoffs cannot be measures of (model) risk.

An example by contradiction is the insurance contract mentioned in the introduction, which is the exchange of the insurance premium – a non-random payoff, with zero variance – for the promise of receiving either a given payout if the insured event occurs or nothing otherwise. The buyers are happy with the exchange regardless of the fact that the insurance contract payoff is not only uncertain (i.e., it has non-zero variance across states of nature), but it has lower expected value than the premium on average.[3] Therefore, risk averse individuals must perceive the insurance contract as a hedge (i.e., less risky than the premium) even if the contract has larger variance than the premium (which has no variance). Given that most of the existing literature on "model risk" ignores this fact, such literature would be better classified as addressing model *uncertainty*, instead.

[3] The average premium sold by the insurance company must cover the actuarial expected payout but also the costs associated with running the insurance company. Hence, in equilibrium the premium must higher than the expected payout on average or the firm goes bankrupt.

2.2.3 Selecting an appropriate SDF proxy

The SDF is left undefined in the general form in Eqs. (1) and (13). But in most of the theoretical asset pricing literature, the SDF is derived in terms of consumption. Although not necessary, the consumption-based formulation helps in the interpretation of the SDF and in guiding the selection of appropriate proxies for empirical applications.

Considering a simple two-period model in which the investors choose consumption and asset allocations, Cochrane (2005) obtains Eq. (13) as

$$p_t = \frac{E_t[x_{t+1}]}{1+R_f} + cov_t\left(\delta \frac{u'(c_{t+1})}{u'(c_t)}, x_{t+1}\right), \tag{14}$$

with $m_{t+1} \equiv \delta \frac{u'(c_{t+1})}{u'(c_t)}$, where δ represents time preference (the "subjective discount factor"), $u'(\cdot)$ is the marginal utility of consumption, and c_t is consumption at time t.

This representation of the SDF shows that for two streams of uncertain cash flows which are otherwise exactly the same, the one with higher payoffs in "bad" states of nature is safer (i.e., it has a higher equilibrium price) than the one with higher payoffs in good states. Here, "bad" states are the ones in which (future) consumption is low and marginal utility, $u'(c_{t+1})$, is high because of the diminishing marginal utility of consumption.[4] Indeed, this negative covariance with total consumption is what makes insurance contracts desirable for the insured: The contract pays out exactly in the event of large losses that would have reduced his consumption substantially.

Defining the SDF in terms of consumption provides intuitive guidance in the search for SDF proxies. It highlights that the variable must describe how rich the investors feel in each state. For example, investors feel poor when their (other) assets have done poorly. Therefore, assets that covary positively with large aggregate ("market") portfolios should be riskier. And indeed, this is the traditional Capital Asset Pricing Model (CAPM) of Sharpe (1964) and Lintner (1965).

[4] Marginal utility of consumption is essentially the value of an extra unit of consumption. It decreases as consumption increases (for risk averse individuals) because the most pressing wants and needs are satisfied first, so subsequent units are consumed for less urgent and less satisfying purposes.

Model risk applications: A good SDF proxy candidate to measure model risk from the perspective of the equity investors of a given firm can be created from the return on the stocks of the firm itself if these stocks are traded. Another is the return on the market portfolio (proxy), like in the CAPM. A company employing a model that underestimates losses exactly when the company's stock price is at its lowest levels (and investors are "poor") is a dangerous model for its shareholders: The investors will be heavily diluted as the portfolio losses force equity issues exactly under these dire conditions in which the cost of equity is very large. Alternatively, a model for which large unexpected losses are not correlated with low stock prices is less risky. In this case, the losses would result in less shareholder dilution on average, as the equity issues would happen under more favourable conditions.

Finally, long-term returns should be preferable as state variables as they usually contain more information regarding the level of "distress" that the investors may be experiencing. Examples are the 1- or 5-year returns on the firm's stock or on a proxy portfolio (such as the bank stocks that I consider later in the empirical section),

$$m_t \equiv f\left(R_{1y,t}\right) \quad \text{or} \quad m_t \equiv f\left(R_{5y,t}\right). \tag{15}$$

The asset pricing literature in fact contains a long list of state variables that can be used to derive SDF proxies. For example, Welch and Goyal (2008) analyse several of them in the context of equity premium predictability, while Cochrane (2017) approaches the discussion from a theoretical asset pricing perspective. Naturally, all this literature is also relevant for model risk measurement applications.

2.2.4 A complementary model risk measure

As the previous sections explain, the relation between the cash flows and the SDF is the fundamental measure of risk. However, the covariance with the SDF is difficult to measure in the context of model risk because the stream of payoffs associated with the models is highly non-linear. Instead, as suggested by Eq. (12), we could simply investigate how the

SDF varies from the time when the investors make the capital injections to the time when they receive the capital back. I define this complementary measure of model risk as

$$\Delta SDF \equiv \frac{1}{\sum_t I_{\{x_t \neq 0\}}} \sum_t I_{\{x_t \neq 0\}} \left(SDF_t - SDF_{med} \right), \tag{16}$$

where $\sum_t$ indicates that the sums are taken over the entire sample period and $I_{\{x_t \neq 0\}}$ is the indicator function, which is 1 if x_t is not zero (i.e., when a capital injection happens), and 0 otherwise.

Just like MRP_t in Eq. (12), ΔSDF is proportional to model risk. For example, a positive value implies that, on average, the SDF is larger when the investors make the capital injection payments than when they receive the dividends back. Given that the SDF can be interpreted as a "distress" variable, the timing of these cash flows is unfavourable for the equity investors (i.e., model risk is positive) when ΔSDF is positive.

3 Empirical exercise

We can now explore two types of empirical applications of the model risk pricing framework presented in Section 2 by examining concrete examples of model risk assessments in Section 3.2 and hedging strategies in Section 3.3.

In particular, the exercise in Section 3.2 shows that the model *risk* measure introduced in the present paper is distinct from the model *accuracy* measures existing in the literature. This is specially the case with respect to the statistics used in accuracy tests, like the ones in Kupiec (1995) or Christoffersen (1998), or the statistics related to the Expected Shortfall in Artzner et al. (1999). In fact, the examples in Section 3.2 show that the approach in the present paper consolidates certain aspects of different model accuracy statistics into a proper and unique model risk measure.

Section 3.3 concludes by presenting examples of actionable results in the form of two strategies to hedge model risk (assuming that the model cannot be easily altered otherwise).

Intuitively, the strategy relies on add-ons to adjust the timing of eventual capital raises towards periods when the SDF is lower, thereby ensuring that they occur under more favourable conditions.

3.1 Data description and variables

From Kenneth French's data library, I collect realized daily returns on the market portfolio and the risk-free rate, described in Fama and French (1996), from 1 January 2000 to 31 December 2023. I construct the market premium (MP_t) as the excess returns on the market portfolio over the risk-free rate. I also collect the return on bank stocks, $R_{bank,t}$, from the industry portfolios described in Fama and French (1997) from 1 January 1995 to 31 December 2023.[5]

I obtain the series for "Adjusted Close" of the Vanguard Long-Term Bond Index Fund Investor Shares (VBLTX), which I merge with the one from the Vanguard Long-Term Bond Fund (VBLAX) (for the latest observations), both from Yahoo Finance. I then create a series of daily returns from this merge for the same period of 1 January 2000 to 31 December 2023.

The term spread in monthly frequency from January 2000 to December 2023 is from Amit Goyal's website.[6] I construct the series following the definition presented in Welch and Goyal (2008) in which the Term Spread, tms, is the difference between the long-term yield on government bonds and the Treasury-bill.

[5] These returns are available from the 48 industry portfolios in the data library.
[6] https://sites.google.com/view/agoyal145

3.1.1 SDF proxies

I consider two SDF proxies: SDF_{Banks} and SDF_{term}.[7] The former is based on the 5-year cumulative return on the portfolio of Bank stocks obtained from Kenneth French website. The latter is based on the term spread obtained from Amit Goyal website.

I define SDF_{Banks} as follows. First, I calculate the risk premium on the portfolio of bank stocks mentioned earlier, using the original daily frequency, by subtracting the risk-free rate from the returns on the portfolio,

$$RP_{bank,d} \equiv R_{bank,d} - R_{f,d}, \tag{17}$$

where $RP_{bank,d}$ is the realized risk premium for the bank portfolio at time d. Next, for each time t, I compound these returns over the previous 1260 days (5 years with 252 days),

$$RPcum_{bank,t} = \prod_{d=t-1260}^{t} \left(1 + RP_{bank,d}\right) - 1, \tag{18}$$

where $\prod$ is the product operator. The non-normalized measure of the SDF simply applies the traditional one period discount factor for each time t based on the cumulative returns calculated above,

$$SDF'_{Banks,t} = \frac{1}{1 + RPcum_{bank,t}}. \tag{19}$$

Finally, I normalize all values by the SDF value at the end of the series, T, so that $SDF_{Banks,T}$ is 1,

$$SDF_{Banks,t} = \frac{SDF'_{Banks,t}}{SDF'_{Banks,T}}. \tag{20}$$

SDF_{Banks} discussion: One of the most important aspects in the calculation above is the fact that the SDF only reflects a risk premium, explicitly removing the risk-free rate in

[7] Although the two series that I propose are theoretically motivated, the SDF can be any series that reflects the risk preference of the investors. For example, the series could even be constructed by asking the investors directly. In this case, the SDF series would contain the values of 1 dollar payoffs in each date in the sample as reported by the investors.

Eq. (17). This is important to ensure that the SDF only reflects risks, as discussed in Section 2.2.1. In addition, this series has no drift by construction. Therefore, payoffs that happen at different points in time will only have different values due to risks (and not due to the pure time-value of money otherwise reflected in the risk-free rate).

Another crucial point is that the SDF series must accurately reflect the level of "distress" perceived by the relevant investors at each point in time. If this is not the case, the series must be adjusted before proceeding with the model risk assessment. Hence, whenever possible, the SDF series should be presented to the target investors for whom the model risk assessment is being made, and the series should incorporate their feedback. SDF plots like the ones in Section 3.1.3, Fig. 1 can be helpful in this case. For instance, it is possible that the 5-year window used to calculate returns needs adjustments, or even that the return proxy should be changed (using the returns on the stocks of the bank itself, instead of using a portfolio of banks, for example).

The other SDF proxy is SDF_{term}, which follows a more straightforward calculation. The raw measure of the SDF simply applies a 12-month discount factor conversion at each time t based on the observed term spread at time t, tms_t. Intuitively, the assumption is that the term spread has an impact on the investors proportional to 1 year,

$$SDF'_{term,t} = \frac{1}{(1 + tms_t)^{12}}. \tag{21}$$

As before, I normalize these values by the SDF at the end of the series, T, so that $SDF_{term,T}$ is 1,

$$SDF_{term,t} = \frac{SDF'_{term,t}}{SDF'_{term,T}}. \tag{22}$$

This normalization implies that all values discounted by the SDF will be expressed in "current dollars" (in fact, in terms of the latest date, T, in the sample). Finally, I transform this series from monthly to daily frequency by repeating the same monthly value for all days in the month.

SDF_{term} **discussion:** The main reason to include SDF_{term} as a proxy is to highlight that any type of variable can be used as an SDF. The only restriction is that it must reflect distress for the investors. Indeed, Fama and French (1993) provide some empirical support for considering the term spread as a state variable. And the theoretical justification is that banks tend to be involved in maturity transformation (Diamond and Dybvig, 1983). Hence, a large spread would correspond to "good times", while a low (or even negative spread) would signal less abundant times or difficulties, as reflected by the inverse relation in Eq. (21).

3.1.2 The stream of payoffs ("capital injections")

The only other inputs necessary for the model risk assessment are the series of payoffs (capital injections) required from the investors. I consider four of them: $Garch_{MP}$, $Garch_{VBLTX}$, $Hist_{MP}$, and $Hist_{VBLTX}$. The first two are implied by a Garch (1,1) model, and the last two are implied by a model that uses historical simulations. I apply the two models to both portfolios presented earlier, MP and $VBLTX$.

Each of these four series correspond to a different x_t in Eq. (7), where θ_t is the realized return on the portfolio in question (MP or $VBLTX$), and the VaR obtained by the model in consideration (with the negative sign without loss of generality) replaces $\hat{q}_p\left(\theta_t^*\right)$ in the same equation.

Both the historical simulation ("Hist.") and the Garch models are for the 99 percentile VaR, one-day ahead, and both use rolling windows with 3 years (756 days) of data. For the historical simulation, the prediction for the next period is just the 0.01 quantile of the returns observed in the preceding 3 years,

$$VaR_{\theta,t+1} = q_{0.01}(\{\theta_{t'}\}_{t'=t-756}^{t}), \tag{23}$$

where $\theta_{t'}$ is the return on the portfolio in question at time t'.

For the Garch (1,1) model (with t-student innovations), the parametric VaR is given by

$$VaR_{\theta,t+1} = \mu_t + \sigma_t \cdot t_{\nu,0.01}, \tag{24}$$

where μ_t is the estimated sample mean, σ_t is the conditional standard deviation predicted by the model using information up to time t, and $t_{\nu,0.01}$ is the 0.01-quantile of the t-distribution with ν degrees of freedom, estimated from the data.

3.1.3 Data visualization

Fig. 1 concludes the data description with a visual display of the variables used in the model risk assessment. The graphs are useful for assessing whether the SDF candidates indeed reflect the preferences of the investors. They also help to confirm whether the risk model assumptions are appropriate. For example, the two return series indeed show the volatility clustering that justifies the use of Garch-type models, while the SDF plots show that both series have no drifts, as discussed in Section 2.2.1.

Finally, the graphs may also help to uncover relations between the timing of capital injections and the SDF levels which are, indeed, the central ingredient of the model risk framework in the present paper.

3.2 Model risk assessment

Table 1 shows the outcome of 16 model risk assessments. The first four columns define the assessment formulation, which vary by i) the SDF proxy in column "SDF" (either SDF_{banks} in Eq. (20) or SDF_{term} in Eq. (22)); ii) the sample starting year in column "Start" (2000 or 2010); iii) the portfolio being modelled in column "Portf." (MP or VBLTX described in Section 3.1); and iv) the modelling choice for the risk model in column "Model" (a Garch in Eq. (24) or historical simulation in Eq. (23)).

For each assessment, I calculate the model risk price, MRP_t in Eq. (12), using the entire sample (which varies depending on the starting date considered). This is the most

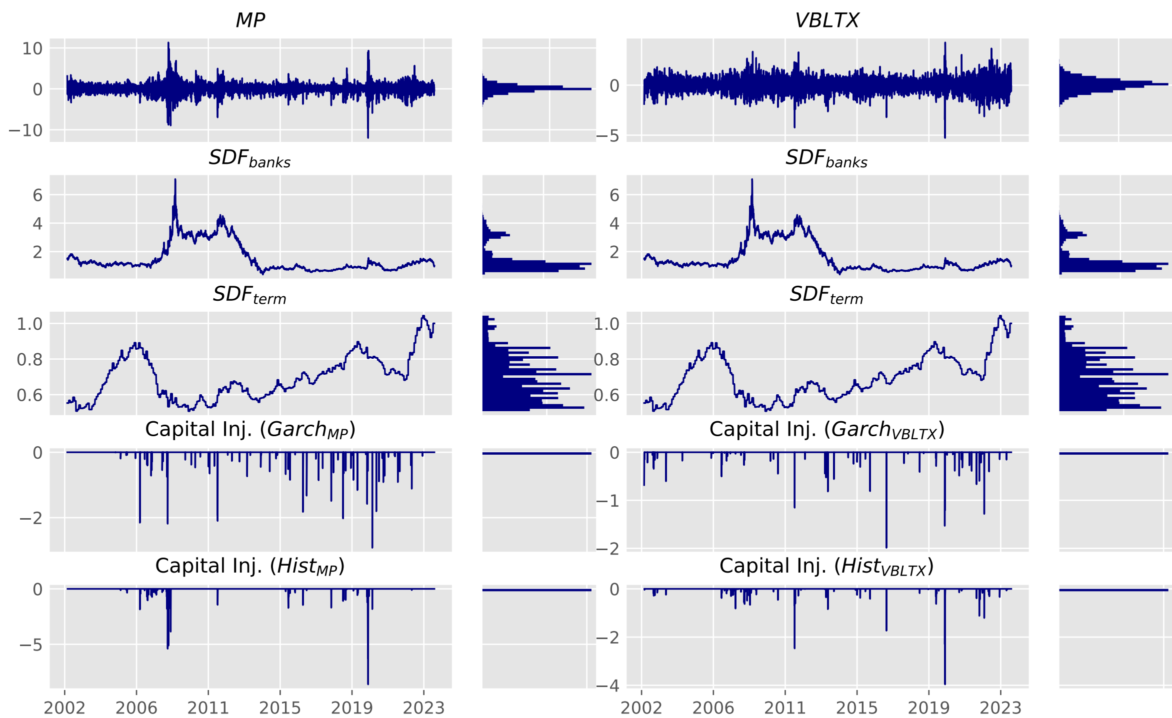

Figure 1: Description of the variables used in the model risk assessment.

MP and VBLTX are the series of daily returns on the respective portfolios; the (duplicated) graphs for SDF_{Banks} and SDF_{term} show the SDF proxies, either based on bank returns or on the term spread, respectively. The four Capital Injection graphs display the difference between realized returns and the VaR (if negative). The ones at the top are for a Garch (1,1) model (i.e., $Garch_{MP}$ or $Garch_{VBLTX}$), the bottom ones are for historical simulation models ($Hist_{MP}$ or $Hist_{VBLTX}$). The smaller graphs on the right-hand side are the histograms corresponding to each time-series. All values are in percent, in daily frequency, and cover the period from Jan 1st 2000 to Dec 31st 2023.
Summary: The most important reason for plotting this graph is to ensure that SDF in fact describes the subjective preferences of the target investors for whom the model risk assessment is being performed. Intuitively, the SDF should be proportional to their perceived "distress". If not, the SDF series should be adjusted before performing the model risk assessment. The magnitude of the capital injections and its covariance with the SDF are crucial to determine the model risk.

important measure of model risk. I divide the result by the number of years in the sample, Y, to obtain

$$MRP_y = \frac{MRT_T}{Y}. \tag{25}$$

The division results in metrics that are more comparable across different sample sizes because they represent the yearly costs of running the models. I estimate one-sided p-values for the significance of MRP_t by bootstrap (column $p(MRP)$ in Table 1).

Column ΔSDF shows the average difference between the SDF values when the investors make capital injections versus when they receive the dividends back (SDF_{med}), as in Eq. (16). The respective t-statistic of the (two-sided) significance test of this difference appears in column $t(\Delta SDF)$.

Table 1 also reports other complementary information, such as the sum of the capital injections in nominal terms,

$$Cap_y = -\frac{1}{Y}\sum_t x_t, \tag{26}$$

and the value of these capital injections (i.e., adjusted by the SDF),

$$C\hat{a}p_y = -\frac{1}{Y}\sum_t (SDF_t \; x_t). \tag{27}$$

Both values are expressed in yearly terms, just like MRP_y. The average capital injection (given that one is needed) is given by

$$\overline{x}|e = -\frac{x_t \; I_{\{x_t \neq 0\}}}{\sum_t I_{\{x_t \neq 0\}}}, \tag{28}$$

which can be loosely interpreted as related to the "Expected Shortfall" described in Artzner et al. (1999).[8] The frequency of capital injections, corresponding to the frequency of VaR breaches in model accuracy tests like the ones in Kupiec (1995) and Christoffersen (1998), is given by

$$P(e) = \frac{I_{\{x_t \neq 0\}}}{N}, \tag{29}$$

[8] In fact, adding the VaR to the capital injection would make the comparison more accurate.

where N is the total number of observations in the sample. Finally, I report the average VaR generated by the model,

$$\overline{VaR} = \frac{\sum_t (VaR_t)}{N}, \tag{30}$$

where VaR_t is the VaR obtained from the model at time t.

The main conclusion from Table 1 is that the model risk price measure, MRP_t in Eq. (12), cannot be systematically obtained from the model accuracy measures that are typically claimed to measure model risk, like the measures related to Eq. (28) or Eq. (29). This is important because it further underscores the distinction between model accuracy and model risk.

In fact, MRP_y encapsulates all the remaining information presented in the table. For example, the Historical Simulation outperforms the Garch in terms of VaR breaches frequency from Eq. (29) in several cases. This suggests that historical simulations are often more accurate than the Garch in the examples that I consider. However, MRP_y tends to imply that the Garch is less risky, which is in line with the evidence in Berkowitz and O'Brien (2002) favouring Garch models for example. Among other reasons, the historical simulations in the exercise tend to require larger capital injections (Cap_y, and often $\overline{x}|e$ in Eq. (28)). This fact is incorporated by MRP_y but ignored in simple tests on VaR breach frequencies, like the ones in Kupiec (1995) and Christoffersen (1998).

3.3 Model risk hedging

Another key advantage of having a complete theoretical framework like the one in the present paper is the possibility of obtaining clear, actionable results to address model risk. This section presents examples of two such results in the form of two strategies to hedge model risk (assuming that the developers cannot easily alter the model to change its risk profile).

Intuitively, the strategies rely on add-ons to adjust the timing of eventual capital raises towards periods when the SDF is lower, thereby ensuring that they occur under more

Table 1: Model risk assessment for different SDF proxies, sample periods, portfolios, and risk models until December 2023.

SDF	Start	Portf.	Model	MRP_y	$p(MRP)$	Cap_y	$\hat{Cap}_y$	$\overline{x}\|e$	$P(e)$	$\overline{VaR}$	SDF_{med}	ΔSDF	$t(\Delta SDF)$
Banks	2000	MP	Garch	7.30	0.02	22.9	33.1	5.80	1.6	25.7	1.13	0.28	2.77
			Hist.	27.60	0.00	43.2	76.3	11.90	1.5	32.8	1.13	0.36	3.99
		VBLTX	Garch	2.20	0.07	10.9	14.5	3.30	1.3	15.1	1.13	0.18	2.20
			Hist.	10.40	0.00	17.4	30.0	4.20	1.7	15.7	1.13	0.60	5.11
	2010	MP	Garch	13.00	0.00	27.5	39.1	6.40	1.7	24.5	0.95	0.43	3.26
			Hist.	8.70	0.01	31.4	38.6	12.90	1.0	34.2	0.95	0.13	1.25
		VBLTX	Garch	4.70	0.01	13.0	17.0	4.00	1.3	15.7	0.95	0.30	2.61
			Hist.	13.00	0.00	20.2	32.2	5.70	1.4	16.3	0.95	0.53	3.78
Term	2000	MP	Garch	1.00	0.01	22.9	16.6	5.80	1.6	25.7	0.68	0.03	2.50
			Hist.	0.50	0.30	43.2	30.0	11.90	1.5	32.8	0.68	0.03	2.61
		VBLTX	Garch	0.00	0.43	10.9	7.5	3.30	1.3	15.1	0.68	0.01	0.71
			Hist.	0.50	0.10	17.4	12.3	4.20	1.7	15.7	0.68	0.00	0.05
	2010	MP	Garch	0.80	0.04	27.5	19.8	6.40	1.7	24.5	0.69	0.01	0.62
			Hist.	2.60	0.00	31.4	24.3	12.90	1.0	34.2	0.69	0.07	4.63
		VBLTX	Garch	0.10	0.29	13.0	9.1	4.00	1.3	15.7	0.69	0.02	1.46
			Hist.	1.00	0.02	20.2	15.0	5.70	1.4	16.3	0.69	0.04	2.42

The main result in the table are the financial measure of model risk, Model Risk Price, MRP_y, and its associated p-value, $p(MRP)$. The other variables are complementary: Cap_y is the observed capital injection requested from equity investors in nominal terms (i.e., the total realized returns minus the VaR); $\hat{Cap}_y$ adjusts these payoffs by their corresponding SDF values. The y subscripts indicate variables on yearly basis (i.e., totals divided by the number of years). $\overline{x}|e$ is the observed average capital injection required given a VaR breach; $P(e)$ is the observed frequency of VaR breaches; and $\overline{VaR}$ is the average VaR generated by the model. All costs are in basis points (and should be multiplied by the portfolio value to convert to actual costs), and $P(e)$ is in percent. SDF_{pb} is the median SDF (when the firm pays back the capital injections to the investors); ΔSDF is the average difference between the SDF values when the investors make the capital injections (i.e., when VaR breaches happen) and the value when they receive their capital back (i.e., SDF_{pb}); $t(\Delta SDF)$ is the t-statistic of testing that this difference is zero. The multi-index shows the SDF choice (based on the return on Bank stocks, Banks, or on the term spread, Term); the year when the sample starts (2000 or 2010), the portfolio being modelled (MP or VBLTX), and the modelling approach (Garch or Historical Simulation). All VaR models are for the 99 percentile loss, and the SDF is normalized to 1 at the last observation (as seen in Fig. 1).
Summary: MRP_y is the fundamental measure of model risk. It encapsulates the remaining information presented in the table. For example, the Historical Simulation outperforms the Garch in terms of VaR breaches frequency in several cases. However, MRP_y still tends to imply that the Garch is less risky. Among other reasons, the historical simulations in the exercise tend to require larger capital injections, Cap_y, and often $\overline{x}|e$. This fact is incorporated by MRP_y but ignored in simple tests on VaR breach frequencies. The table also highlights that ΔSDF is often significant, providing further evidence that the models are risky.

favourable conditions. More specifically, I consider two strategies: i) increasing the estimated VaR when the SDF is above its median or, in addition to that, ii) also decreasing the VaR proportionally when the SDF is below its median.

Strategy (i) is conservative because it only increases the VaR level in some periods, while keeping the VaR unchanged in the remaining periods. I denote this strategy "conservative" and present its results in Table 2. Strategy (ii) aims to keep the VaR level relatively unchanged over time by also reducing the VaR in calmer times. I denote this strategy "balanced", with results in Table 3.

3.3.1 Conservative model risk hedging

I obtain the add-ons to be placed on the estimated VaR by considering variations of Eqs. (23) and (24). Effectively, I replace the original 99 percent VaR level by a higher level when SDF_t is above its payback value (under the assumption that SDF_{t+1} is also likely to be above its payback value).

For the historical simulation, I adjust the VaR level in Eq. (23) via a constant c^*, such that

$$VaR^*_{\theta,t+1} = \begin{cases} q_{0.01-c^*}(\{\theta_{t'}\}^t_{t'=t-756}) & \text{if } SDF_t > SDF_{pb}, \\ q_{0.01}(\{\theta_{t'}\}^t_{t'=t-756}) & \text{if } SDF_t \leq SDF_{pb}, \end{cases} \tag{31}$$

where the equation at the bottom is just Eq. (23). Equivalently for the Garch in Eq. (24), I calculate

$$VaR^*_{\theta,t+1} = \begin{cases} \mu_t + \sigma_t \cdot t_{\nu,0.01-c^*} & \text{if } SDF_t > SDF_{pb}, \\ \mu_t + \sigma_t \cdot t_{\nu,0.01} & \text{if } SDF_t \leq SDF_{pb}, \end{cases} \tag{32}$$

where, again, the bottom equation is unchanged from Eq. (24).

The conservative hedging strategy inevitably generates an average VaR level over time that is higher (i.e., more conservative) than the original one. As Eqs. (31) and (32) show, the hedged VaR level is either higher than or equal to the unhedged VaR level in Eqs. (23) and (24).

3.3.2 Balanced model risk hedging

The balanced hedging strategy aims to keep the average VaR levels of the hedged and unhedged models relatively unchanged. In this case, I compensate the increase in VaR level when SDF_t is above its payback value by proportionally decreasing the VaR level in the remaining periods.

Formally, I replace the unhedged historical simulation VaR in Eq. (23) by

$$VaR^*_{\theta,t+1} = \begin{cases} q_{0.01-c^*}(\{\theta_{t'}\}^t_{t'=t-756}) & \text{if } SDF_t > SDF_{pb}, \\ q_{0.01+c^*}(\{\theta_{t'}\}^t_{t'=t-756}) & \text{if } SDF_t \leq SDF_{pb}. \end{cases} \tag{33}$$

Similarly, I replace the unhedged Garch VaR in Eq. (24) by

$$VaR^*_{\theta,t+1} = \begin{cases} \mu_t + \sigma_t \cdot t_{\nu,0.01-c^*} & \text{if } SDF_t > SDF_{pb}, \\ \mu_t + \sigma_t \cdot t_{\nu,0.01+c^*} & \text{if } SDF_t \leq SDF_{pb}. \end{cases} \tag{34}$$

The addition of c^* to the quantile when SDF_t is not above its payback value in Eqs. (33) and (34) effectively lowers the VaR level of the hedged model in these states relative to the unhedged one. This should result in larger and more frequent breaches. However, the breaches in these states sould be counterbalanced by fewer and smaller breaches in the states when the SDF is above its payback value, thereby keeping the average VaR level of the hedged and unhedged models relatively similar.

3.3.3 The optimal c^*

In all cases, c^* is the optimal constant that minimizes the absolute value of MRP_t in Eq. (12), thereby obtaining the output that results in the model risk price closest to zero.

In the present exercise, I obtain c^* in sample by first calculating MRP_t implied by a grid of 99 candidates

$$c \in \{1, 2, 3..., 99\}, \tag{35}$$

where c is in basis points. Each candidate c replaces c^* in the hedged VaRs in Eqs. (31), (32), (33), and (34), generating its respective model risk price, MRP_t. The candidate c that generates the model risk price closest to zero is the optimal value, c^*.

3.3.4 Results

Tables 2 and 3, respectively, show the results for the conservative and for the balanced hedging strategies. The overall conclusion from the joint analysis of the two strategies is that they succeed in reducing model risk in all cases: The model risk price of all hedged models, MRP_y in Eq. (25), is lower than the unhedged values in Table 1.

The results also indicate that hedging is more efficient for SDF_{Banks}. The reason seems to be twofold: First, the baseline model risk estimates are higher based on this SDF formulation than based on SDF_{term}. Second, the series $SDF_{Banks,t}$ is very persistent by construction. Hence, it is relatively easy to predict its future values from the present ones. Given that timing the (future) SDF values is crucial to reduce model risk, hedging becomes easier in this case.

Finally, increases on the average VaR ($\overline{VaR}$ in the tables) is the common cost associated with the hedging of all models considered (even when using the balanced hedging strategy).

For the conservative hedging strategy in Table 2, the model risk prices decrease but still remain significantly positive for a few models (as given by $p(MRP)$). In addition, conservatively hedging does not always result in lower and insignificant ΔSDF. In terms of hedging costs, the VaR of the hedged models, $\overline{VaR}$, is sometimes substantially larger than the unhedged VaR. And the frequency of VaR breaches, $P(e)$, tends to be lower, which effectively corresponds to increased VaR levels relative to the unhedged models.

Table 2: Model risk assessments for conservatively hedged models until December 2023.

SDF	Start	Portf.	Model	MRP_y	$p(MRP)$	Cap_y	$\hat{Cap}_y$	$\overline{x}\|e$	$P(e)$	$\overline{VaR}$	SDF_{med}	ΔSDF	$t(\Delta SDF)$
Banks	2000	MP	Garch	1.20	0.35	18.8	22.4	6.30	1.2	28.6	1.13	0.06	0.56
			Hist.	1.50	0.25	20.1	24.1	8.80	0.9	47.8	1.13	0.03	0.43
		VBLTX	Garch	0.50	0.39	8.5	10.0	3.40	1.0	16.1	1.13	0.06	0.64
			Hist.	2.40	0.22	7.9	11.3	4.30	0.7	21.5	1.13	0.13	0.87
	2010	MP	Garch	1.00	0.35	19.0	19.1	5.90	1.3	29.6	0.95	0.08	0.72
			Hist.	0.80	0.30	17.5	17.3	9.80	0.7	53.8	0.95	-0.05	-1.40
		VBLTX	Garch	0.40	0.38	7.9	7.9	3.40	0.9	17.9	0.95	0.15	1.17
			Hist.	4.60	0.07	8.0	12.2	4.80	0.7	26.1	0.95	0.29	1.41
Term	2000	MP	Garch	0.20	0.00	1.9	1.5	5.00	0.2	45.6	0.68	0.09	2.56
			Hist.	0.30	0.25	8.8	6.3	11.60	0.3	60.9	0.68	0.01	0.20
		VBLTX	Garch	0.00	0.50	8.7	5.9	3.40	1.0	15.9	0.68	0.01	0.36
			Hist.	0.10	0.37	3.8	2.6	8.80	0.2	29.7	0.68	-0.04	-1.27
	2010	MP	Garch	0.10	0.01	2.0	1.5	4.10	0.2	43.6	0.69	0.06	1.73
			Hist.	0.60	0.03	7.1	5.6	14.20	0.2	68.0	0.69	0.06	1.84
		VBLTX	Garch	0.00	0.49	6.3	4.4	3.70	0.7	18.5	0.69	0.01	0.48
			Hist.	0.10	0.34	4.9	3.5	11.50	0.2	35.2	0.69	-0.01	-0.35

The Model Risk Price is given by MRP_y, and its associated p-value is $p(MRP)$. Cap_y is the observed capital injection requested from equity investors in nominal terms (i.e., the total realized returns minus the VaR); $\hat{Cap}_y$ adjusts these payoffs by their corresponding SDF values. The y subscripts indicate variables on yearly basis (i.e., totals divided by the number of years). $\overline{x}|e$ is the observed average capital injection required given a VaR breach; $P(e)$ is the observed frequency of VaR breaches; and $\overline{VaR}$ is the average VaR generated by the model. All costs are in basis points (and should be multiplied by the portfolio value to convert to actual costs), and $P(e)$ is in percent. SDF_{pb} is the median SDF (when the firm pays back the capital injections to the investors); ΔSDF is the average difference between the SDF values when the investors make the capital injections (i.e., when VaR breaches happen) and the value when they receive their capital back (SDF_{pb}); $t(\Delta SDF)$ is the t-statistic of this difference being different from zero. The multi-index shows the SDF choice (based on the return on Bank stocks, Banks, or on the term spread, Term); the year when the sample starts (2000 or 2010), the portfolio being modelled (MP or VBLTX), and the modelling approach (Garch or Historical Simulation). All VaR models are for the 99 percentile loss, and the SDF is normalized to 1 at the last observation (as seen in Fig. 1). All models are hedged for model risk.

Summary: The hedged models have lower model risk than the unhedged models in most cases. But the hedging costs include increases on the average VaR and VaR levels. Indeed, all hedged models have lower MRP_y than the unhedged ones in Table 1. However, the model risk cost is still significantly positive for a few models ($p(MRP)$). In addition, conservatively hedging does not always result in lower and insignificant ΔSDF. Finally, the average VaR of the hedged models, $\overline{VaR}$, is (sometimes substantially) larger, and the frequency of VaR breaches, $P(e)$, tends to be lower (effectively increasing the VaR level).

Table 3: **Model risk assessments for balanced hedged models until December 2023.**

SDF	Start	Portf.	Model	MRP_y	$p(MRP)$	Cap_y	$\hat{Cap}_y$	$\overline{x}\|e$	$P(e)$	$\overline{VaR}$	SDF_{med}	ΔSDF	$t(\Delta SDF)$
Banks	2000	MP	Garch	-0.10	0.45	24.5	27.5	6.30	1.6	27.6	1.13	-0.01	-0.12
			Hist.	-0.10	0.45	30.4	34.1	8.20	1.5	45.4	1.13	-0.05	-1.06
		VBLTX	Garch	-0.00	0.46	10.5	11.8	3.20	1.3	15.7	1.13	-0.02	-0.30
			Hist.	1.40	0.37	12.1	15.0	3.90	1.2	20.3	1.13	-0.02	-0.17
	2010	MP	Garch	-0.10	0.45	26.4	25.0	6.00	1.8	28.4	0.95	0.03	0.36
			Hist.	-0.10	0.46	25.7	24.3	9.20	1.1	51.4	0.95	-0.08	-3.29
		VBLTX	Garch	-0.10	0.44	10.7	10.1	3.20	1.4	17.2	0.95	0.04	0.48
			Hist.	4.00	0.18	11.9	15.3	4.40	1.1	24.9	0.95	0.11	0.84
Term	2000	MP	Garch	0.20	0.02	2.1	1.6	4.90	0.2	45.6	0.68	0.06	1.52
			Hist.	0.20	0.27	9.0	6.3	11.80	0.3	60.8	0.68	0.01	0.20
		VBLTX	Garch	-0.00	0.50	8.7	5.9	3.30	1.1	15.9	0.68	0.00	0.21
			Hist.	0.00	0.42	3.9	2.7	9.10	0.2	29.7	0.68	-0.04	-1.27
	2010	MP	Garch	0.10	0.08	2.3	1.7	4.60	0.2	43.5	0.69	0.06	1.73
			Hist.	0.60	0.04	7.1	5.6	14.20	0.2	67.9	0.69	0.06	1.84
		VBLTX	Garch	-0.00	0.52	6.4	4.4	3.40	0.7	18.4	0.69	-0.00	-0.24
			Hist.	0.10	0.34	4.9	3.5	11.50	0.2	35.2	0.69	-0.01	-0.35

The Model Risk Price is given by MRP_y, and its associated p-value is $p(MRP)$. Cap_y is the observed capital injection requested from equity investors in nominal terms (i.e., the total realized returns minus the VaR); $\hat{Cap}_y$ adjusts these payoffs by their corresponding SDF values. The y subscripts indicate variables on yearly basis (i.e., totals divided by the number of years). $\overline{x}|e$ is the observed average capital injection required given a VaR breach; $P(e)$ is the observed frequency of VaR breaches; and $\overline{VaR}$ is the average VaR generated by the model. All costs are in basis points (and should be multiplied by the portfolio value to convert to actual costs), and $P(e)$ is in percent. SDF_{pb} is the median SDF (when the firm pays back the capital injections to the investors); ΔSDF is the average difference between the SDF values when the investors make the capital injections (i.e., when VaR breaches happen) and the value when they receive their capital back (SDF_{pb}); $t(\Delta SDF)$ is the t-statistic of this difference being different from zero. The multi-index shows the SDF choice (based on the return on Bank stocks, Banks, or on the term spread, Term); the year when the sample starts (2000 or 2010), the portfolio being modelled (MP or VBLTX), and the modelling approach (Garch or Historical Simulation). All VaR models are for the 99 percentile loss, and the SDF is normalized to 1 at the last observation (as seen in Fig. 1). All models are hedged for model risk.

Summary: The balanced hedged models tend to have even lower model risk, MRP_y, than the conservatively hedged models in Table 2, especially for SDF_{Banks}. Almost none of them has significantly positive model risk ($p(MRP)$). In addition, balanced hedging results in lower and insignificant (or negative) ΔSDF for all models. In terms of hedging costs, the increases in average VaR and VaR levels tend to be notably lower than the ones obtained by conservative hedging (especially based on SDF_{Banks}).

For the balanced hedging strategy in Table 3, the model risk prices, MRP_y, tend to be lower than the ones obtained by the conservative hedges in Table 2, especially for SDF_{Banks}. In this case, almost none of them has significantly positive model risk (in column $p(MRP)$). In addition, balanced hedging results in lower and insignificant (or negative) ΔSDF for all models. In terms of hedging costs, the increases in average VaR and VaR levels tend to be notably lower than the ones obtained by conservative hedging (especially based on SDF_{Banks}). In fact, some balanced hedged models have lower VaR levels than their unhedged counterpart. For example, the frequency of exceptions (which is inversely proportional to the effective VaR level), $P(e)$, is 1.8 for the hedged model under Banks, 2010, MP, Garch in Table 3, compared to 1.7 for the unhedged model in Table 1.

4 Conclusion

In the present paper, I propose that the Model Risk Price is the measure that should be used to evaluate the risk of all models. I argue that the measure is obtained from a formal theoretical framework, and that it generates clear actionable results for each risk assessment. In addition, the Model Risk Price is expressed in financial terms. This implies that it can also be intuitively communicated alongside its associated hedging strategies to non-technical audiences, such as investors, CEOs, other C-suite executives, and risk managers.

The Model Risk Price is also fundamentally distinct from the model accuracy statistics typically used to measure model risk. Hence, the paper provides a seminal contribution to the model *risk* literature by being the first to address risk as defined in Economics and in contrast to *uncertainty*.

From a practical standpoint, the paper addresses one of the most critical questions posed to risk managers by the firm's investors: What is the precise impact of model risk on their investments, and what concrete actions can be taken to mitigate it.

References

Artzner, Philippe, Freddy Delbaen, Jean-Marc Eber, and David Heath, 1999, Coherent measures of risk, *Mathematical Finance* 9, 203–228.

Berkowitz, Jeremy, and James O'Brien, 2002, How accurate are value-at-risk models at commercial banks?, *Journal of Finance* 57, 1093–1111.

Christoffersen, Peter F., 1998, Evaluating interval forecasts, *International Economic Review* 39, 841–862.

Cochrane, John H., 2005, *Asset Pricing, revised edition* (Princeton University Press, Princeton).

Cochrane, John H., 2017, Macro-finance, *Review of Finance* 21, 945–985.

Diamond, Douglas W., and Philip H. Dybvig, 1983, Bank runs, deposit insurance, and liquidity, *Journal of Political Economy* 91, 401–419.

Fama, Eugene F., and Kenneth R. French, 1993, Common risk factors in the returns on stocks and bonds, *Journal of Financial Economics* 33, 3–56.

Fama, Eugene F., and Kenneth R. French, 1996, Multifactor explanations of asset pricing anomalies, *Journal of Finance* 51, 55–84.

Fama, Eugene F., and Kenneth R. French, 1997, Industry costs of equity, *Journal of Financial Economics* 43, 153–193.

Kupiec, Paul H, 1995, Techniques for verifying the accuracy of risk measurement models, *Journal of Derivatives* 3, 73–84.

Lintner, John, 1965, The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets, *Review of Economics and Statistics* 47, 13–37.

Sharpe, William F., 1964, Capital asset prices: A theory of market equilibrium under conditions of risk, *Journal of Finance* 19, 425–442.

Welch, Ivo, and Amit Goyal, 2008, A comprehensive look at the empirical performance of equity premium prediction, *Review of Financial Studies* 21, 1455–1508.