Universidad del Rosario | Facultad de Economía

SERIE DOCUMENTOS DE TRABAJO

No. 316

Febrero de 2024

Worst-Case Higher Moment Risk Measure: Addressing Distributional Shifts and Procyclicality

Carlos Castro-Iragorri

Fabio Gómez

Nancy Quiceno

Universidad del Rosario, Facultad de Economía
Cl. 12c #4-59. Bogotá, D. C., Colombia
PBX: (031) 2970200 ext. 4163

facultadeconomia@urosario.edu.co
https://www.urosario.edu.co/Facultad-de-Economia/Inicio/
https://repository.urosario.edu.co/handle/10336/20291

Worst-Case Higher Moment Risk Measure: Addressing Distributional Shifts and Procyclicality

Carlos Castro-Iragorri[a]*, Fabio Gómez[a],[b], and Nancy Quiceno[c]

[a]School of Economics, Universidad del Rosario

[b]School of Risk and Actuarial Studies, UNSW Sydney

[c]Cámara de Riesgo Central de Contraparte, CRCC

February 28, 2024

Abstract

This paper addresses the inherent procyclicality in widely adopted financial risk measures, such as Expected Shortfall (ES). We propose an innovative approach utilizing the Higher Moment (HM) risk measure, which offers a robust solution to distributional shifts by incorporating adaptive features. Empirical results using historical S&P500 returns indicate that worst-case HM risk measures significantly reduce the underestimation of risk and provide more stable risk assessments throughout the financial cycle compared to traditional ES predictions. These results suggest that HM risk measures represent a viable alternative to regulatory add-ons for stress testing and procyclicality mitigation in financial risk management.

Keywords: procyclicality, higher moment risk, stress testing, expected shortfall

JEL: G32, G17, C58

1 Introduction

Accurately measuring risk is crucial in the financial industry. It plays a pivotal role in determining the risk capital allocation at financial institutions and establishing margin and collateral requirements for derivative products. Metrics such as Value at Risk (VaR) and Expected Shortfall (ES) are widely adopted and have become indispensable tools for quantifying and managing financial risk. The Basel Committee on Banking Supervision emphasizes the significance of ES as a recommended risk measure in regulatory frameworks (BCBS, 2016). Nevertheless, relying on statistical estimation methods for these risk measures introduces a trade-off between achieving the desired sensitivity to economic cycles and avoiding negative feedback loops (Brunnermeier and Pedersen, 2009). This inherent procyclicality arises from both regulatory practices, such as the VaR rule (Adrian and Shin, 2014), and internal judgment within institutions (Gordy and Howells, 2006), as well as methodological approaches (Bräutigam et al., 2021).

Procyclicality in risk measures manifests as a systematic bias that amplifies the proclivity of these measures to follow the ups and downs of the financial markets. This phenomenon presents a significant challenge, as risk capital requirements are susceptible to proving inadequate in pre-crisis periods and excessive during crises. The delayed responsiveness of risk measures to changes

*Corresponding author: Carlos Castro-Iragorri

in financial market conditions undermines their efficacy as anticipatory tools. In response to this challenge, regulators have historically turned to the incorporation of add-ons, such as buffers, stress tests, and floors, to recalibrate risk measures (Cominetta et al., 2019). While these adjustments aim to bolster the robustness of risk estimates, they introduce complexities and may not eliminate procyclicality.

To address the limitations of add-on adjusted risk measures, we propose an innovative approach using the higher moment (HM) risk measure. The HM risk measure was introduced by Krokhmal (2007) to extend ES, and it has since attracted significant interest in risk management. Researchers including Krokhmal (2007), Matmoura and Penev (2013), Tang and Yang (2012), and Gómez et al. (2022) have highlighted its viability in areas such as stochastic programming, portfolio optimization, and risk management. Additionally, Gómez et al. (2022) have demonstrated its compatibility with higher-order stochastic dominance, a well-established concept in decision theory with applications to various important economic and financial issues. In this way, the HM risk measure shares similarities with the upper partial moment risk measure, which is acknowledged in BCBS (2011) as one of the promising alternatives to ES. From a machine-learning perspective, Duchi and Namkoong (2021) and Duchi et al. (2023) have recently investigated this risk measure in the context of distributionally robust stochastic optimization.

When estimating risk measures, the empirical distribution function is commonly used. The rationale for this approach is based on the assumption that past and future observations are drawn from the same population. However, it is crucial to acknowledge that this assumption rarely holds in practice. Practitioners frequently encounter distributional shifts, wherein the characteristics of the data distribution change. These shifts can be triggered by factors such as economic events, regulatory changes, or uncertainties in market conditions. These distributional shifts challenge the reliability of risk estimates, necessitating risk management practices to adapt to the dynamic nature of underlying data distributions. Robust risk models should integrate adaptive features and scenario analyses to effectively account for the impact of distributional shifts on risk assessments. A common strategy from robust statistics and machine learning is to build a robust model for distributional perturbations. This idea emerges from the robust control approach in control theory, as illustrated, for example, by Petersen et al. (2000). In a systematic exploration, Hansen and Sargent (2008) demonstrate how robust control principles are applied to economic models. This idea is in line with the argument of Bloom (2014) that agents may struggle to form precise probability distributions for future outcomes under heightened uncertainty, often resorting to worst-case scenario determination. Our formulation uses Wasserstein ambiguity sets to define a worst-case risk measure. We opted for this approach because it allows for easy and intuitive exploration of out-of-sample scenarios, thanks to the Wasserstein distance. Furthermore, our approach eliminates the need for external add-ons, often seen as esoteric.

Utilizing historical daily data of S&P 500 returns, our empirical investigation indicates that risk predictions based on worst-case HM risk measures show a significant reduction in underestimating risk. While employing ES to predict risk over a one-year horizon provides adequate coverage only 50% of the time; worst-case HM risk measures offer coverage for 67% of the days in the historical sample (as targeted). Simultaneously, these measures effectively mitigate procyclicality, providing a more stable risk assessment throughout the financial cycle. Predicted risk indicators, adjusted for performance and based on HM risk measures, exhibit a significantly lower negative correlation with financial cycle indicators than ES predictions. This suggests a reduced lagging relationship to the economic cycle. The stability observed in these measures is attributed to their adaptability to changing market conditions, promoting dynamic stability in risk forecasts. These properties position HM risk measures as a practical alternative to the add-on adjustments regulators recommend for stress testing and procyclicality mitigation.

Our contributions can be summarized as follows: Firstly, we introduce a worst-case risk measure designed to effectively accommodate shifts in the risk distribution caused by changing economic conditions. Secondly, our empirical findings demonstrate that HM risk measures offer sufficient coverage in risk forecasting and can mitigate the procyclicality inherent in using historical data. Thirdly, we establish that HM risk measures represent a superior alternative to the add-ons proposed by regulators: they are theoretically supported and achieve an empirical balance between risk sensitivity and procyclical mitigation.

The paper is organized as follows. Section 2 provides an overview of the discussion and various approaches to mitigate procyclicality, introducing the methodology for comparing the one-year ahead prediction performance of the considered risk measures and add-ons. Section 3 presents higher moment risk measures and outlines the methodological approach for estimation. In Section 4, we elucidate how model uncertainty in HM can be viewed as a generalized form of an add-on to a standard risk measure like ES. This section also outlines a methodological approach for estimating add-ons based on the moments of the historical sample. Section 5 offers an empirical illustration of HM measures using historical daily returns of the S&P 500. Finally, Section 6 presents the conclusions drawn from the study.

2 Mitigating procyclicality

Prevailing methodologies estimating risk measures based on historical data face a formidable challenge posed by procyclicality. Adrian and Shin (2014) demonstrate that capital charges exacerbate the already procyclical nature of credit supply. The VaR rule is designed to maintain a constant ratio of the risk measure to the required capital. However, adjusting the risk exposure, which is sensitive to the economic cycle, necessitates a robust de-leveraging process during periods of market stress, thereby amplifying the impact of an economic downturn. Their model is constructed such that the risk measure remains invariant to the cycle, providing institutions with incentives to align their book equity with the risk measure and to adjust their exposures through the contraction and expansion of assets. Empirically, their findings indicate that assets in a sample of large US banks are responsive to changes in debt levels while remaining insensitive to adjustments in equity, underscoring the significance of the VaR rule.

Adjustment through the economic cycle in institutional leverage is closely linked to the escalation of margin and collateral requirements during economic downturns. Consequently, the procyclicality of margin requirements has become a growing interest for central counterparties (CCPs). Gurrola-Perez (2021) argues that central counterparties face a trilemma when measuring margin requirements. They must balance the need for risk sensitivity, mitigating procyclicality and promoting economic efficiency for market participants. While regulators have decisively resolved the dilemma in favour of sufficiently sensitive margin requirements, it does not absolve them from addressing the issue of procyclicality.

The European Market Infrastructure Regulation, specifically in Article 41 part 5 introduces three well-known anticyclical measures, also known as procyclical add-ons. The first add-on, commonly referred to as a buffer, mandates a 25% increase in estimated margins, with this increase dynamically adjusted over time. The second add-on, often called stress, assigns a 25% weight to stressed observations in constructing the sample to estimate margins. This sample, known as the look-back period, relies on historical data. The third add-on, known as the floor, stipulates that margin requirements should not fall below those estimated using a 10-year look-back period.

No consensus exists regarding the implementation and efficacy of these add-ons. Gurrola-Perez (2021) contends that these add-ons overlook certain crucial aspects related to the margin model. For instance, a false sense of security may be instilled by a larger look-back period, as different risk measures are susceptible to sampling errors that are not necessarily rectified by employing larger samples. Additionally, the extent to which non-independent or non-identically distributed financial data may impact the risk measure in terms of risk coverage and procyclicality is often underestimated. Various studies have reviewed the empirical performance of these procyclical add-ons, yielding mixed results. Gurrola-Perez (2021) observes that longer look-back periods yield lower procyclical values but simultaneously underestimate the level of risk, leading to increased losses for market participants. Murphy et al. (2016), utilizing a simulation framework, demonstrate that the add-ons effectively mitigate procyclicality, yet an optimal calibration could offer a superior trade-off. Boudiaf et al. (2023) survey using these tools, along with others like SPAN, employed at European CCPs, revealing diverse calibrations adopted in practical applications.

Despite a consensus on the characterization of procyclicality, no universally accepted method exists for gauging this phenomenon's extent. Gurrola-Perez (2021) investigates diverse tools central counterparties utilise to mitigate procyclicality in estimating margins. Various measures, such as the Peak-to-Trough ratio, are employed to quantify the extent of procyclical behaviour, capturing the ratio of the maximum initial margin to the minimum within a fixed observation period. Another measure, denoted as the "n-day," quantifies the largest increase in margin over an n-day period. In a related vein, Bräutigam et al. (2021) examine the correlation coefficient between the financial cycle measure and the look-forward ratio. This ratio delineates the disparity between the one-year ahead predicted risk and the realized future risk, relying on standard risk measures such as VaR and ES.

We follow Bräutigam et al. (2021) in using the look-forward measure to determine the performance of the risk measure. The look-forward ratio based on ES will be our benchmark measure,

$$R_{q,T_{lb}}(t) = \frac{\mathrm{ES}_q(t+1y)}{\mathrm{ES}_{q,T_{lb}}(t)}, \tag{2.1}$$

where the numerator is the ex-post future risk estimated with data one year ahead of the current period t. The denominator represents the predicted (ex-ante) risk estimated with historical data available at time t using different look-back periods T_{lb} ranging from 6 months to 2 years of data. Note that if $R_{q,T_{lb}}(t)$ is close to 1 then the predicted risk, mainly determined by the look-back period and the chosen risk measure (ES in this case), correctly assesses future risk. If $R_{q,T_{lb}}(t) < 1$, then the risk measure underestimates risk and overestimates risk if it is significantly larger than 1.

We define the regulatory add-ons as modified versions of the predicted risk measure. Recall that, as documented by Boudiaf et al. (2023), there is no unique approach to define these add-ons. For example, in the case of the buffer, the predicted risk incorporates a naive adjustment that increases the level of predicted risk and reduces the look-forward ratio,

$$R_{q,T_{lb}}^{buffer}(t) = \frac{\mathrm{ES}_q(t+1y)}{\mathrm{ES}_{q,T_{lb}}(t)(1+0.25)}. \tag{2.2}$$

In this case, the buffer will necessarily increase the performance; the look-forward ratio will be closer to or above 1. In the case of the stress add-on, a dynamic look-back period is employed, contingent upon whether the sample returns utilized for estimating the predicted risk encompass more than 10% of observations subjected to stress. Should this criterion not be met, the sample is expanded to integrate stressed historical returns. The identification of stressed returns begins with estimating return volatility using an Exponentially Weighted Moving Average (EWMA). Volatility is estimated at each time point based on a one-year sample and a smoothing parameter of 0.95. Daily returns are classified as stressed if the volatility estimated at that particular time point exceeds 1.5 times the average volatility estimate across the current sample.

$$R_{q,T_{lb}}^{stress}(t) = \frac{\mathrm{ES}_q(t+1y)}{\mathrm{ES}_{q,\hat{f}}(t)}, \tag{2.3}$$

where $\hat{f}$ represents the subset of returns utilized for estimating the loss distribution. If $\hat{f}$ is deemed to lack a sufficient number of stressed returns, the subset of returns is expanded to encompass historically stressed returns.

In the case of the floor add-on, we choose the maximum predicted risk from ES estimated using a look-back period of 1 year and a look-back period of 5 years,

$$R_{q,T_{lb}}^{floor}(t) = \frac{\mathrm{ES}_q(t+1y)}{\max(\mathrm{ES}_{q,T_{lb}}(t), \mathrm{ES}_{q,T_{lb}^*}(t))}, \tag{2.4}$$

where $T_{lb}^* > T_{lb}$. As suggested by Gurrola-Perez (2021), a longer look-back period does not necessarily imply a larger predicted risk measure. Therefore, its effect on performance will depend on the sample.

Consider any other risk measure ρ_χ, where $\chi = \chi(q)$ is a parameter that makes ES_q and ρ_χ comparable in some sense. Let ρ_{χ,δ_n} represent a worst-case risk measure corresponding to ρ_χ, where δ_n quantifies the degree of pessimism; see, for instance, (4.2) and (5.1) regarding the HM risk measure. In this case, the look-forward ratio is defined as follows,

$$R_{q,T_{lb}}^{HM}(t) = \frac{\rho_\chi(t+1y)}{\rho_{\chi,\delta_n,T_{lb}}(t)}. \tag{2.5}$$

In Section 3, we will introduce HM risk measures as a feasible alternative to standard risk measures such as VaR and their add-on modified counterparts that we have exemplified using ES. We chose ES because it is a special case of the HM risk measure, as shown in Section 3.

Bräutigam et al. (2021) findings reveal a negative correlation between the look-forward ratio using VaR, derived from a dataset encompassing financial indexes like the S&P 500 over a monthly interval spanning 1987 to 2018. According to the authors, this negative correlation suggests that, on average, one-year-ahead losses have been overestimated with measures based on available data at the time of estimation. An alternative interpretation posits that the look-forward ratio exhibits a significant lag concerning the financial cycle indicator, with the former being assessed as the annualized one-year moving average of the empirical standard deviation or mean absolute deviation.

3 The higher moment risk measure

To begin, we establish some notation. For a random variable S, we represent its distribution function by F_S. For a distribution function F, denote by $F^{\leftarrow}$ and $F^{\rightarrow}$ its left and right generalized inverses, respectively, namely, for $q \in [0,1]$, $F^{\leftarrow}(q) = \inf\{s \in \mathbb{R} : F(s) \geq q\}$ and $F^{\rightarrow}(q) = \sup\{s \in \mathbb{R} : F(s) \leq q\}$, with $\inf \emptyset = \infty$ and $\sup \emptyset = -\infty$ by convention. For an event $(\cdot)$, the notation $1_{(\cdot)}$ denotes its indicator function, equal to 1 if the event occurs, and 0 otherwise. For a real number a, write $a_+ = a1_{(a>0)}$, $\operatorname{sign} a = 1_{(a>0)} - 1_{(a<0)}$, and $|a| = a \operatorname{sign} a$.

Assume an atomless probability space $(\Omega, \mathcal{F}, \mu)$. Let S be a real-valued random variable, representing a risk variable in a loss-profit style. Denote by $\hat{s} = F_S^{\leftarrow}(1) \leq \infty$ the essential supremum of S and by $\hat{\pi} = F_S(\hat{s}) - F_S(\hat{s}-)$ the probability mass at $\hat{s}$. Assume that $S \in L^p := L^p(\Omega, \mathcal{F}, \mu)$ for some $p \geq 1$, that is,

$$\|S\|_{L^p} = \left(\int_{\mathbb{R}} |t|^p \, dF_S(t)\right)^{\frac{1}{p}} < \infty.$$

Now we define the HM risk measure. According to Krokhmal (2007), for some $0 < q < 1$, the HM risk measure of S is defined to be

$$\rho(S) = \inf_{s \in \mathbb{R}} \left\{ s + \frac{1}{1-q} \left(E\left[(S-s)_+^p\right]\right)^{\frac{1}{p}} \right\}. \tag{3.1}$$

It is apparent that the effective region over which the infimum of (3.1) is taken is $s \leq \hat{s}$.

When $p = 1$ and $q \in (0,1)$, the HM risk measure for S simplifies to

$$\rho(S) = s_* + \frac{1}{1-q} E\left[(S - s_*)_+\right],$$

with all possible minimizers s_* forming a closed set $[F_S^{\leftarrow}(q), F_S^{\rightarrow}(q)]$, which is a proper interval when q corresponds to the level of a flat piece of the distribution F and is a single point otherwise. When $p = 1$, $\rho(S)$ is an alternative expression for the well-known ES risk measure with confidence level q:

$$\mathrm{ES}_q(S) = \frac{1}{1-q} \int_q^1 F_S^{\leftarrow}(\tilde{q}) d\tilde{q}. \tag{3.2}$$

When $p > 1$ and $q \in (0, 1 - \hat{\pi}^{\frac{1}{p}})$, the HM risk measure for S is given by

$$\rho(S) = s_* + \frac{1}{1-q} \left(E\left[(S - s_*)_+^p\right]\right)^{\frac{1}{p}}, \tag{3.3}$$

where the minimizer $s = s_* < \hat{s}$ is the unique solution to the equation

$$\frac{\left(E\left[(S-s)_+^{p-1}\right]\right)^p}{\left(E\left[(S-s)_+^p\right]\right)^{p-1}} = (1-q)^p; \tag{3.4}$$

see Example 2.2 of Krokhmal (2007) or Theorem 2.1 of Tang and Yang (2012). When $p > 1$ and $q \in [1 - \hat{\pi}^{\frac{1}{p}}, 1)$, the HM risk measure for S is simply $\hat{s}$; see Lemma 2.3(b) of Gómez et al. (2022).

4 Add-on from model uncertainty

Let $\mathcal{P}(\mathbb{R}\times\mathbb{R})$ be the set of all probability measures on the Borel sigma algebra $\mathcal{B}(\mathbb{R}\times\mathbb{R})$. An element $\pi\in\mathcal{P}(\mathbb{R}\times\mathbb{R})$ represents the joint distribution of a random vector (U,V), where $U\in\mathbb{R}$ and $V\in\mathbb{R}$. We denote by π_U and π_V as the respective marginal distributions of U and V under π. For $p\geq 1$, let $\mathcal{P}_p(\mathbb{R})$ represent the class of all probability measures Q on the Borel sigma algebra $\mathcal{B}(\mathbb{R})$ with finite moment of order p, namely,

$$\int_{\mathbb{R}} |t|^p\, Q(dt) < \infty.$$

To measure the discrepancy between two probability measures $Q_1, Q_2 \in \mathcal{P}_p$, we utilize the Wasserstein metric, defined by

$$W_p(Q_1,Q_2) = \inf\left\{ \left(E_\pi[\ |U-V|^p]\right)^{\frac{1}{p}} :\ \pi\in\mathcal{P}(\mathbb{R}\times\mathbb{R}), \pi_U = Q_1, \pi_V = Q_2 \right\}. \tag{4.1}$$

Let S be a random variable representing the aggregate loss in a loss-profit style and having an induced probability measure P. We are interested in measuring S according to ρ, that is, $\rho_P(S)$. In order to account for model uncertainty, we adopt a worst-case risk measure approach, where we consider a range of possible probability distributions that encompass P. Specifically, we consider

$$\sup_{Q:\ W_p(Q,P)\leq\delta} \rho_Q(S). \tag{4.2}$$

The infinite-dimensional problem (4.2) is typically challenging to solve explicitly. However, in some cases, it can be transformed into a form of risk measure under the reference probability distribution P and a regularization term. Blanchet et al. (2019) and Blanchet et al. (2022) have successfully transformed such problems in their work on machine learning and mean-variance portfolio optimization, respectively. Regarding the HM risk measure (3.1), for $p\geq 1$ and $q\in(0,1)$, Blanchet et al. (2024) show that

$$\sup_{Q:\ W_p(Q,P)\leq\delta} \rho_Q(S) = \rho_P(S) + \frac{1}{1-q}\delta. \tag{4.3}$$

Let $S^1, \ldots, S^n$ be a sample of S, whose empirical probability measure is denoted by P_n. To approximate $\rho_P(S)$, the conventional approach involves the sample-average approximation, namely, $\rho_{P_n}(S)$. However, the sample-average approximation exhibits suboptimal out-of-sample performance and does not guarantee proper risk coverage. Recognizing these limitations, we explore a robustified version of the problem:

$$\begin{aligned} \hat{J}_n(\delta) &= \sup_{Q:\ W_p(Q,P_n)\leq\delta} \rho_Q(S). \\ &= \inf_{s\in\mathbb{R}} \left\{ s + \frac{1}{1-q}\left(\frac{1}{n}\sum_{i=1}^{n}\left(S^i - s\right)_+^p\right)^{\frac{1}{p}} \right\} + \frac{1}{1-q}\delta. \end{aligned} \tag{4.4}$$

Now, the question arises: how can we estimate δ? To answer this question, we employ the following limit result due to Blanchet et al. (2024). Assume that the condition $q\in(0,1-\hat{\pi}^{\frac{1}{p}})$

holds, where $\hat{\pi}$ represents the probability mass at the essential supremum of S. For any $P \in \mathcal{P}_{2p}$ and $\delta_n = \kappa n^{-\frac{1}{2}}$, it holds that

$$n^{\frac{1}{2}} \left(\hat{J}_n(\delta_n) - J^* \right) \xrightarrow{d} \kappa\mu_P + \sigma_P N(0,1),$$

where

$$\begin{aligned} \mu_P &= \frac{1}{1-q}, & (4.5) \\ \sigma_P &= \frac{1}{p(1-q)} E_P[(S-s_*)_+^p]^{\frac{1}{p}-1} \left(E_P[S-s_*]_+^{2p} - (E_P[S-s_*]_+^p)^2 \right)^{\frac{1}{2}}. & (4.6) \end{aligned}$$

For practical proposes, when n is sufficiently large,

$$n^{\frac{1}{2}} \left(\hat{J}_n(\delta_n) - J^* \right) \approx \kappa\mu_P + \sigma_P N(0,1),$$

which, in turns, implies

$$\begin{aligned} P\left(\hat{J}_n(\delta_n) \geq J^* \right) &\approx P\left(\kappa\mu_P + \sigma_P N(0,1) \geq 0 \right) \\ &= P\left(N(0,1) \geq -\kappa \frac{\mu_P}{\sigma_P} \right). \end{aligned}$$

Thus, to estimate δ_n aiming for performance guarantees with a confidence level of $1-\alpha$, we may set $\delta_n = \kappa(\alpha) n^{-\frac{1}{2}}$, where

$$\kappa(\alpha) = \frac{\sigma_P}{\mu_P} \phi^{-1}(1-\alpha).$$

Here $\phi^{-1}(\alpha)$ denotes the $\alpha \times 100\%$ quantile of the standard normal distribution. To estimate μ_P and σ_P, we may employ bootstrapping.

5 Empirical results

We examine daily returns on the SP&500 spanning the period from January 2, 1987, to December 29, 2023, to estimate both the look-forward ratio and the financial cycle indicator [1].

In the context of the HM risk measure, there is no predetermined reference for a confidence level compared to ES. However, we aim to modify the risk measure while keeping the value relatively stable. For that purpose, we need to decide on a new confidence level $\chi = \chi(p, q; F_S) \in [0, 1)$ such that

$$\rho_{1,q}(S) = \rho_{p,\chi}(S), \tag{5.1}$$

for some random variables S. The existence and uniqueness of such confidence level are ensured by Lemma 4.1 of Gómez et al. (2022).

We estimate predicted and future (ex-post) risk measures using a confidence level of 95% ($q = 0.95$). For the predicted HM risk, we set $p = 2$; thus, we suggest using a value of $\chi = 0.746$, which ensures that (5.1) holds when S is normally distributed. See, for instance, Section 4 of Bellini and Di Bernardino (2017) for a similar procedure in the context of expectiles. In addition, to obtain

[1] We are extending the sample used by Bräutigam et al. (2021).

the estimated measure 4.4, we use the CVXR package in R for disciplined convex optimization. Finally, we use a 90% ($\alpha = 0.9$) confidence level to estimate δ_n.

In Figure 1, we present the evolution of the look-forward ratio and the financial cycle. The former is derived from a one-year standard deviation of observed returns σ_{r_t}. For robustness, we also use the VIX as an external cycle indicator because it does not depend on the returns, and it is well documented that it increases during US recessions (Bloom, 2014).

However, we have to shorten the historical sample because VIX has only been available since the early 1990's. The sample average of the look-forward ratio is 1.09, suggesting that, on average, the predicted risk adequately represents future risk. However, this average conceals significant deviations from the optimal value of 1 across the sample. To assess the performance of the risk measure used for predicting risk, we calculate the percentage of days in the sample where $R_{q,T_{lb}}(t) \leq 1$ indicating when the risk forecast offers adequate coverage of ex-post risk. We find this measure to be 49%, signifying that, in approximately half of the sample, the risk measure fails to provide sufficient coverage for the level of exposure. These results broadly align with the findings of Bräutigam et al. (2021).

Another noteworthy aspect of figures 1 and 2 is the leading relationship between the financial cycle indicator and the look-forward ratio. Bräutigam et al. (2021) quantify this aspect by the negative correlation between both indicators. The extended sample shows correlations of -0.64 and -0.37 depending on the financial cycle measure, standard deviation of returns, or VIX, respectively. In summary, our study replicates and complements the work of Bräutigam et al. (2021), suggesting that risk measures are inherently procyclical due to their reliance on historical data. Simultaneously, we highlight the limited risk coverage provided by these measures.

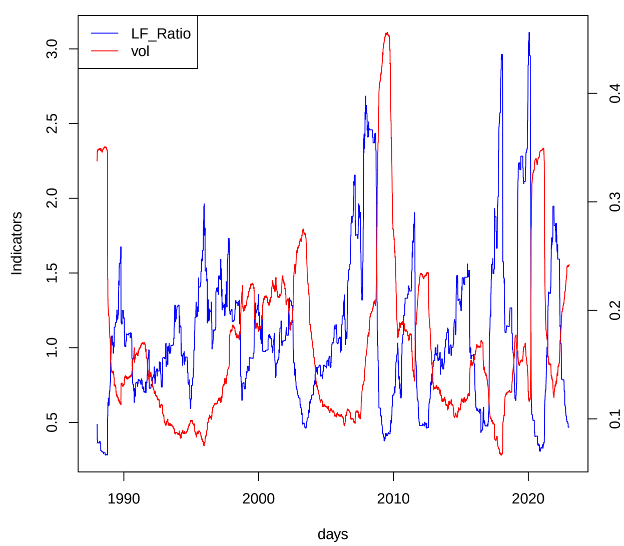

Figure 1: Financial cycle based on returns.

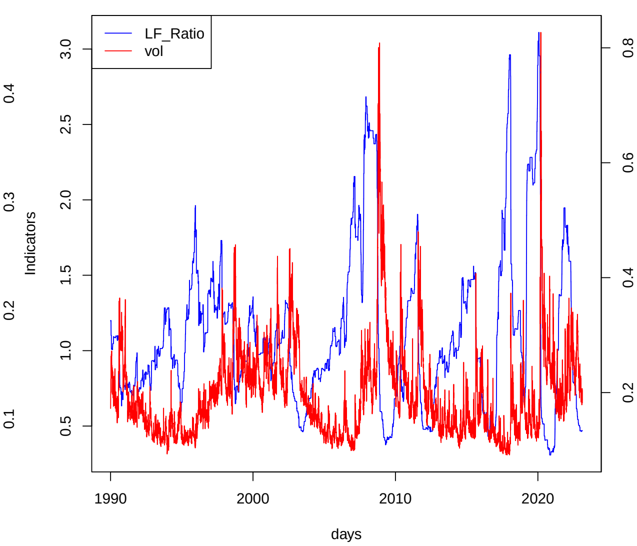

Figure 2: Financial cycle based on VIX.

The figures depicts the time series of the look-forward ratio (blue line on the left-axis), estimated using ES as the predicted risk measure, accompanied by a 95% confidence interval. Additionally, the financial cycle indicator is illustrated on the left-axis in red.

	Look-back(years)		performance		procyclicality	
Model	T_{lb}	δ_n	Av. LFR	% Overestimation	σ_{r_t}	VIX
$R_{q,T_{lb}}(t)$	1		1.09	0.49	-0.64	-0.37
	2		1.05	0.55	-0.36	-0.11
	0.5		1.21	0.42	-0.61	-0.37
$R_{q,T_{lb}}^{buffer}(t)$	1		0.87	0.71	-0.64	-0.27
	2		0.83	0.74	-0.36	-0.11
	0.5		0.97	0.66	-0.61	-0.37
$R_{q,T_{lb}}^{stress}(t)$	1		3.79	0.32	-0.58	-0.35
$R_{q,T_{lb}}^{floor}(t)$	$\max(1,5)$		0.9	0.65	-0.26	0.01
$R_{q,T_{lb}}^{HM}(t)$	1	1	0.96	0.67	-0.64	-0.28
	2	3	0.95	0.67	-0.39	-0.14
	2	5	0.97	0.65	-0.36	-0.10

Table 1: The table presents performance metrics, including the average look-forward ratio across the sample, the percentage of overestimation, and procyclicality (correlation with the financial cycle indicator). We use the one-year moving average of the standard deviation of returns or the VIX for the financial cycle. These metrics assess the estimation of predicted risk through various methodologies, encompassing ES, buffer and floor add-on adjusted ES, and HM risk measures. The predicted risk measures are derived from different look-back periods, ranging from 6 months to 2 years. For HM risk measures, the parameter δ_n is estimated using a look-back period ranging from 1 to 5 years.

We undertake a comparative analysis of the look-forward ratio for Expected Shortfall (ES), integrating buffer, stress, and floor add-ons into predicted risk assessments. The initial three rows of Table 1 present results using ES with varying look-back periods. Increasing the look-back period from 1 to 2 years has a marginal effect, reducing the correlation from -0.64 to -0.36 (with VIX, the correlations decrease from -0.37 to -0.11), without a significant change in performance. Conversely, employing a six-month look-back period yields no notable impact.

When incorporating the buffer add-on into the ES predicted risk (rows 3 to 6 in the table), we observe an expected enhancement in performance. This add-on implies a location shift in the distribution, reducing the percentage of days where risk is underestimated from 51% to 29%. However, there is no discernible impact on the correlation with the financial cycle indicator. The inclusion of the stress-type add-on, complementing one-year sampled past returns with historically stressed returns, does not improve performance or mitigate procyclicality. These results suggest that arbitrarily including historical stress returns does not necessarily enhance risk coverage throughout the sample period. Perhaps a more nuanced analysis of which stressed returns to include, along with the most recent returns, is required instead of a rule-based approach as described in section 2.

We further introduce a floor-type add-on that considers the maximum between ES estimated with look-back periods of 1 or 5 years to estimate the predicted risk (row 8 in the table). This addition leads to a slight performance improvement, reducing the percentage of days where risk is underestimated from 51% to 35%. Regarding procyclicality, we observe a significant decrease in correlation to -0.26 or 0.01 compared to the non-adjusted ES.

Among the three add-ons evaluated, we find that the buffer and floor add-ons improve risk coverage while stress does not. In terms of mitigating procyclicality, only the floor add-on provides a significant reduction in the negative correlation, thus mitigating the lead-lag relationship with the financial cycle indicator.

The worst-case Higher Moment (HM) risk measure, as introduced in Section 3, provides a mechanism to address the inherent uncertainty in the empirical risk distribution utilized for risk measure estimation. This approach yields a robust risk measure capable of significantly reducing the risk of underestimating the level of risk. In the ninth row of Table 1, employing a one-year look-back period already demonstrates a notable performance improvement, with the percentage of days where risk is underestimated dropping from 51% to 33%. Unlike the buffer add-on, the HM risk measure introduces a mechanism to induce a stressed version of the risk measure that dynamically evolves, rather than implementing a fixed incremental adjustment to the chosen risk measure (in this case, ES).

The interpretation provided in Section 4 underscores the additional adjustment capability of the worst-case HM risk measure, accommodating a sample-sensitive adjustment mechanism by allowing the optimal δ_n to depend on the first and second moments of the historical distribution. By incorporating different look-back periods, for instance, using a 2-year look-back period for the risk distribution and 3 to 5 years for the optimal adjustment δ_n, a significant reduction in mitigating procyclicality is achieved.

The last two rows in Table 1 demonstrate that increasing the look-back period in the dynamic estimation of the optimal δ_n reduces the correlation to -0.36 and -0.10. This reduction is further corroborated by a decrease in the lagging relationship between the look-forward ratio and the financial cycle indicator, as observed in Figures 3 and 4.

6 Conclusions

The worst-case HM risk measures contribute to a nuanced understanding of risk, accommodating the expected trade-off between risk sensibility and dynamic stability. This introduces a novel perspective on risk management, where uncertainty is not merely addressed through external add-ons but is fundamentally integrated into the risk measurement process. As we delve into further research, optimal calibration of these HM risk measures can be explored to strike a balance that enhances economic efficiency for market participants while maintaining the resilience and stability of risk forecasts.

7 Declaration of generative AI and AI-assisted technologies in the writing process

During the preparation of this work the author(s) used ChatGPT-3.5 and Grammarly in order to improve language and readability. After using this tool/service, the author(s) reviewed and edited the content as needed and take(s) full responsibility for the content of the publication.

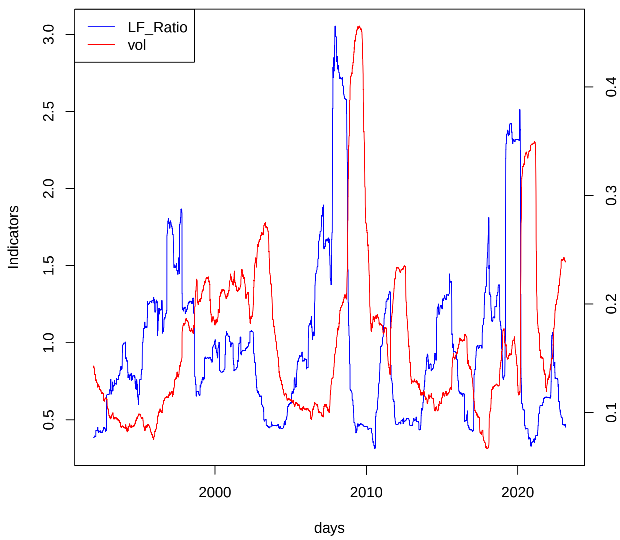

Figure 3: Financial cycle based on returns.

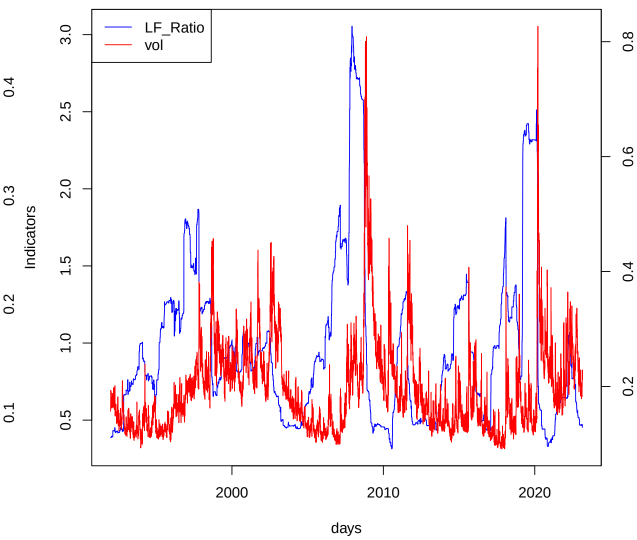

Figure 4: Financial cycle based on VIX.

The figure illustrates the time series of the look-forward ratio (depicted on the left-axis in blue), estimated using HM risk measure as the predicted risk measure, and is accompanied by a 95% confidence interval. Additionally, the figure presents the financial cycle indicator illustrated on the left-axis in red.

References

[1] Adrian, T. and Shin, H.S. 2014. Procyclical leverage and value-at-risk. The Review of Financial Studies, 27(2), pp.373-403

[2] BCBS, 2011. Messages from the academic literature on risk measurement for the trading book. Bank for International Settlements. Working Paper, No. 19. Available at: https://www.bis.org/publ/bcbs_wp19.pdf.

[3] BCBS, 2016. Minimum capital requirements for market risk. Standards, No. d352. Available at: https://www.bis.org/bcbs/publ/d352.htm.

[4] Bellini, F. and Di Bernardino, E., 2017. Risk management with expectiles. The European Journal of Finance, 23(6), pp.487-506.

[5] Blanchet, J., Chen, L. and Zhou, X.Y., 2022. Distributionally robust mean-variance portfolio selection with Wasserstein distances. Management Science, 68(9), pp.6382-6410.

[6] Blanchet, J., Kang, Y. and Murthy, K., 2019. Robust Wasserstein profile inference and applications to machine learning. Journal of Applied Probability, 56(3), pp.830-857.

[7] Bloom, N., 2014. Fluctuations in uncertainty. Journal of Economic Perspectives, 28(2), pp.153-176.

[8] Boudiaf, I.A., Scheicher, M. and Vacirca, F., 2023. CCP initial margin models in Europe. ECB Occasional Paper No. 2023/314. Available at SSRN: https://ssrn.com/abstract=4411243.

[9] Bräutigam, M., Dacorogna, M. and Kratz, M., 2023. Pro-cyclicality beyond business cycle. Mathematical Finance, 33(2), pp.308-341.

[10] Brunnermeier, M.K. and Pedersen, L.H., 2008. Market liquidity and funding liquidity. The Review of Financial Studies, 22(6), pp.2201-2238.

[11] Cominetta, M., Grill, M. and Jukonis, A., 2019. Investigating initial margin procyclicality and corrective tools using EMIR data. Macroprudential Bulletin, European Central Bank.

[12] Duchi, J., Hashimoto, T. and Namkoong, H., 2023. Distributionally robust losses for latent covariate mixtures. Operations Research, 71(2), pp.649-664.

[13] Duchi, J.C. and Namkoong, H., 2021. Learning models with uniform performance via distributionally robust optimization. The Annals of Statistics, 49(3), pp.1378-1406.

[14] Gómez, F., Tang, Q. and Tong, Z., 2022. The gradient allocation principle based on the higher moment risk measure. Journal of Banking & Finance, 143, p.106544.

[15] Gordy, M.B. and Howells, B., 2006. Procyclicality in Basel II: Can we treat the disease without killing the patient? Journal of Financial Intermediation 15, pp.395-41.

[16] Gurrola-Perez, P., 2021. Procyclicality of central counterparty margin models: Systemic problems need systemic approaches. Journal of Financial Market Infrastructures, 10(1), pp.23-55.

[17] Hansen, L.P. and Sargent, T.J., 2008. Robustness. Princeton University Press.

[18] Krokhmal, P., 2007. Higher moment coherent risk measures. Quantitative Finance, 7(4), pp. 373–387.

[19] Matmoura,Y. and Penev, S., 2013. Multistage optimization of option portfolio using higher order coherent risk measures. European Journal of Operational Research 227 (1), 190–198.

[20] Murphy, D., Vasios, M. and Vause, N., 2016. A Comparative analysis of tools to limit the procyclicality of initial margin requirements. Bank of England Working Paper No. 597.

[21] Petersen, I.R., James, M.R. and Dupuis, P., 2000. Minimax optimal control of stochastic uncertain systems with relative entropy constraints. IEEE Transactions on Automatic Control, 45(3), pp.398-412.

[22] Tang, Q. and Yang, F., 2012. On the Haezendonck–Goovaerts risk measure for extreme risks. Insurance: Mathematics and Economics, 50(1), pp.217-227.

Carlos Castro-Iragorri
School of Economics
Universidad del Rosario
Bogotá, Bogotá 111711, Colombia
Email: carlos.castro@urosario.edu.co

Fabio Gómez
School of Risk and Actuarial Studies
UNSW Sydney
Sydney, NSW 2052, Australia
School of Economics
Universidad del Rosario
Bogotá, Bogotá 111711, Colombia
Email: fabio.gomez@unsw.edu.au

Nancy Quinceno
Cámara de Riesgo Central de Contraparte
Bogotá, Bogotá 111711, Colombia
Email: nquiceno@camaraderiesgo.com.co