Economics-Aware Machine Learning for Option-Implied Risk Metrics

Heqing Shi[1], Yi Cao[2], and Zexun Chen[1]

[1]University of Edinburgh Business School, Edinburgh, UK

[2]Xi'an Jiaotong-Liverpool University, Suzhou, China

October 22, 2024

Version of September 2024

Abstract

Machine learning models are predominantly data-driven and often lack embedded domain knowledge. This limitation is particularly significant in the field of finance, where certain asset conditions must be maintained. To address this, we propose a novel constrained Gaussian Process model (consGP) that simultaneously minimises interpolation loss and satisfies encoded linear inequalities representing economic constraints. This approach enables the consGP to learn from market data whilst adhering to fundamental economic principles. We apply this model to the estimation of option-implied risk metrics, where the consGP demonstrates robust performance in estimating risk-neutral density (RND) across sparse and noisy option observations. This model has been demonstrated to be particularly suitable for modeling stock options with limited sample sizes due to insufficient liquidity. Our comprehensive empirical studies, conducted using a cross-section of S&P 500 stocks, reveal that the consGP model outperforms traditional structural models in recovering stock-level RND. This improved performance translates into enhanced predictive information and tangible economic benefits for investors. The consGP model thus represents a significant advancement in integrating machine learning techniques with domain-specific financial constraints, offering a more robust and economics-aware approach to option pricing and risk assessment.

Keywords: Risk-neutral density; Gaussian process; Option-implied risk metrics

JEL code: C63, F47

1 Introduction

Machine learning models, whilst containing robust function-fitting capabilities and adaptability, are fundamentally data-driven. These models lack inherent understanding of the domain-specific knowledge underpinning the data, potentially leading to paradoxical or inconsistent predictions. This limitation is particularly significant in the financial industry, where certain axioms must be upheld, such as the non-negativity of asset prices. Over the past couple of years, a key challenge in finance has been the development of models that can effectively learn information from data with the "awareness" of the pertinent domain knowledge. The seminal work of Chen et al. (2023) establishes a comprehensive framework of "transfer learning", facilitating the transition from theoretical to market data. In this paper, we investigate this topic from the perspective of a small sample of daily data, which is particularly important in risk metrics estimation. To maintain adequate sensitivity to the most recent market information, risk metrics are often estimated by the historical data in a short time period. Consequently, it remains a challenge to learn meaningful insights from a limited data sample over a short time period, while aware of the domain-specific knowledge. This challenge is particularly acute for transfer learning frameworks that typically rely on substantial volumes of data.

In this paper, we introduce a novel Gaussian Process model that integrates established economic theories as constraints. Our approach commences with the training of a standard Gaussian Process (GP) model using market data. Owing to their ability to incorporate prior knowledge through kernel functions and provide a measure of uncertainty in predictions, GPs are often well-suited for small sample sizes. We then encode the economic constraints as a set of linear inequalities. To constrain the GP effectively, we simultaneously solve two conditions: firstly, the interpolation condition for minimising loss, and secondly, the inequality conditions that enforce the constraints. By interpolating the observed data whilst being "aware" of these constraints, the model is able to avoid overfitting to noisy data and generate predictions that are "aware" of economic principles.

By virtue of this, our model can be applied to a wide range of small sample learning topics in finance and economics. In this paper, we demonstrate its application to the topic of option-implied risk metrics, which we have selected for two reasons. Firstly, learning the embedded information for ad-hoc risk measurement constitutes a complex economic problem, typically reliant on a relatively small dataset (usually option data from a single trading day). To accurately compute risk metrics, one must first extract the risk-neutral density (RND) of the underlying returns from the learned option prices, a process in which even minor deviations can result in significantly distorted RNDs and, by extension, risk metrics. Secondly, the field of option-implied risk has a substantial body of well-established research that can serve as a benchmark for our model, enabling a rigorous evaluation of its performance.

Sepcifically, option prices encode investor expectations and risk preferences for the underlying asset returns, which is vital for the estimation of risk metrics in ex-ante. To decode this predictive information, one could use the model-free calculations of implied risk metrics directly with traded options, for example the implied moments including the implied volatility, skewness and kurtosis as proposed in Bakshi et al. (2003) (henceforth BKM), and the implied tail risk measures including the implied Value-at-Risk (VaR) and the implied Expected Shortfall (ES) as proposed in Barone-Adesi (2016) (henceforth BA). However, these methods of model-free calculations largely depend on the observed option prices which can be noisy as they are subject to pricing and recording errors, especially for less liquid option contracts. In addition, options are usually observed with limited and not equidistant strike prices, which introduces estimation bias to implied risk metrics. One could instead turn to recover the risk-neutral density (RND) of the underlying asset returns from option prices. The implied RND can be recovered parametrically by assuming a parametric distribution for the underlying asset returns. One flexible choice is the mixture of normals as used in Huggenberger et al. (2018), but it may still fail to capture the important information in option prices due to a fixed distributional form. To be free of parametric assumptions, the RND is usually recovered in a model-free fashion according to Breeden and Litzenberger (1978). Recovering a model-free RND from options requires no-arbitrage prices and deviations from the no-arbitrage could lead to a less-informative recovered RND. Shimko (1993), Figlewski (2010) address this problem by interpolating the implied volatilities (IVs) with polynomial spline and transforming the interpolated IVs to no-arbitrage prices, relying on the Black-Scholes formula. Nevertheless, as the same for the observed prices, the observed IVs contain noise, and interpolation could easily overfit to those noisy and erroneous IVs, in which case the no-arbitrage conditions can no longer be guaranteed by the Black-Scholes transformation. Aït-Sahalia and Lo (2000), Aït-Sahalia and Duarte (2003) employ the nonparametric kernel regression to construct an estimator for IVs, which is less influenced by outliers but is still intensive in sample size and data regularity. Aït-Sahalia and Duarte (2003) use a combination of constrained least squares and nonparametric kernel regression to consider the no-arbitrage constraints, but this constraining-then-smoothing procedure could sacrifice too much information from the option data and introduce biases at the first step of solving the constrained least squared regression.

Decoding the predictive information from options hence faces a dilemma. The model-free calculations of implied risk metrics directly with option prices allow only a limited number of metrics to be estimated (at least up to the current literature), and the recovery of implied RND has to either accept parametric assumptions or deal with the noisy and irregular option data in order to recover properly in a model-free way. This dilemma is even more severe for stock options, as for individual stocks, their options are traded actively only at strikes around the spot price and the noise in prices is larger because of more speculation. This paper provides insights to recover stock-level RND with individual stock options and decode the predictive information for stock returns from the recovered RND, by using the economics-aware machine learning that enforces the prior knowledge about the shape of the latent option pricing function from the no-arbitrage principles. This is made possible because of the recent developments in machine learning and optimization procedures which give rise to nonlinear models with constraints. We achieve this by enforcing multiple linear no-arbitrage inequalities to a Gaussian process (GP) which is a nonparametric data-driven learning model. Therefore, the no-arbitrage representation of option prices can be learned over an extended range of strikes given only the observed data and the priory known constraints, which exhibits robustness to noisy and erroneous prices and helps to solve the dilemma by deriving model-free stock-level RND with improved regularity and economic informativeness.

This paper contributes to three strands of literature. First, we contribute to the studies that attempt to model option prices nonparametrically. One challenge in these studies is reserving no-arbitrage properties for option prices while being independent of parametric assumptions. We propose to employ the framework of constrained GP as proposed in López-Lopera et al. (2018) to encode the full set of no-arbitrage constraints into a GP. The no-arbitrage constraints require the option prices to be monotonic and convex with respect to strikes, and the prices need to be non-negative at the same time, which can be formulated as a set of linear inequalities. We train an unconstrained GP with observed option prices at the initial step, and then the no-arbitrage inequalities are simultaneously enforced to the model outputs by solving a quadratic problem conditioning on the fine-tuned hyperparameters. The choice of GP bypasses the selection of parametric option pricing models, and the inclusion of no-arbitrage constraints guarantees economic-aware outputs of option prices, particularly outside of the observed range of strikes. There is no need to model the IVs and transform them to Black-Scholes prices as in Shimko (1993), Aït-Sahalia and Lo (2000), Figlewski (2010), which still has the reliance on the Black-Scholes model and can not guarantee no-arbitrage when the IVs are noisy. Additionally, the no-arbitrage-constrained GP can be extended to higher dimensions to model option prices if high-dimensional option characteristics need to be incorporated. The smoothness of the learned option pricing function is also ensured by the kernel function inherited in the GP.

This paper also contributes to the estimation of stock-level RND using stock options. Based on the no-arbitrage-constrained GP, we therefore propose a novel model-free RND-recovery method. The no-arbitrage constraints are encoded globally into the fine-tuned GP, which outputs the economics-aware option prices both within and outside of the observed range of strikes, and the Breeden and Litzenberger (1978) theorem can therefore be directly applied. We do not need to choose an ad-hoc parametric distribution to extrapolate the tails of the RND, for example, the normal or the GEV tails as suggested in Figlewski (2010). In addition, we learn the no-arbitrage representation of option prices separately for call and put options using the constrained GP, and we blend the densities implied in call and put options into one aggregated density. That is, we use the out-of-the-money call (put) options to form the RHS (LHS) of the RND respectively while ensuring the recovered RND integrals to one. This allows the information from two sides to be encoded jointly, which is particularly useful when using the RND to gauge the asymmetry of investor sentiments. We show in the simulations that our proposed constrained-GP-based RND-recovery achieves much lower density loss compared with the conventional IV-based approaches, especially when the option sample is small and noisy as is common to observe for stock options.

We contribute to the empirical literature of option-implied information for stock returns as well. The information implied in stock-level RND is rarely discussed as there is no suitable estimation tool. Our proposed stock-level RND-recovery with no-arbitrage-constrained GP allows us to comprehensively evaluate the predictive information encoded in the option-implied RND on the cross-section of stocks that are the constituents of the S&P 500 index. We estimate an array of implied risk metrics using the recovered stock-level RND to decode the implied information. We focus on the implied tail risk measures including VaR, ES and left partial moment (LPM) to examine the information for the left tail risk and use the implied moment-based return predictors including the innovation of volatility, skewness and the asymmetry of variance to examine the stock return predictability unlocked by the RND. To the best of our knowledge, our paper is among the first to study the stock-level RND implied in options with such a large stock cross-section. Our empirical analysis covers a long time period from 04 Jan 1996 to 31 Dec 2022. Estimating the constrained GP to recover the RND for each stock from call and put options, with each available time-to-maturity, on each trading day requires a vast amount of computational power, as there are around 10 million runs of model training in total. We use parallel computing procedures to fasten the computation. We find, using stock-level time-series and cross-sectional regressions, that the implied risk metrics estimated from the constrained-GP-based stock-level RND present significantly improved predictive power for stock returns, in terms of both the left tail and the expectation. When using the constrained-GP-based RND to estimate tail risk measures, the out-of-sample R^1 is 4.92% higher on average across all stocks than using the IV-based approaches. When using the constrained-GP-based RND to construct moment-based predictors for stock returns, the Sharpe ratio of the long-short portfolio is 5.30 times higher on average.

The remainder of this paper is organised as follows. Section 2 describes the process to enforce the full-set of no-arbitrage constraints simultaneously into a GP, and explains how we recover a hybrid RND using call and put options jointly. Section 3 validates the accuracy of the RND estimated with the no-arbitraged-constrained GP in a simulated economy, under various option sample sparsity and irregularity conditions. 4 describes the data. Section 5 discusses our empirical investigation and findings, and Section 6 concludes.

2 Methodology

In this section, we outline our methodology of encoding the economics-motivated no-arbitrage option pricing constraints into a Gaussian process. The no-arbitrage constraints are **hard**-encoded globally so that the constrained GP is economics-aware while learning from observed option prices. This is a nonparametric and data-driven method to obtain no-arbitrage instead of relying on any parametric option pricing models.

2.1 Economics-Aware Machine: The Constrained GP

Given the function of option prices $C, P : [0, +\infty) \times [0, +\infty) \to \mathbb{R}^+$, the objective is to learn the function from noisy observations of call and put prices $\tilde{C}, \tilde{P}$, on traded strike prices $\tilde{K}$, with the following no-arbitrage constraints as in Härdle and Hlávka (2009) to satisfy for all strike prices K:

(i) positivity: $C(K) \geq 0, P(K) \geq 0$

(ii) monotonicity: $\partial_K C \leq 0, \partial_K P \geq 0$

(iii) convexity: $\partial_K^2 C \geq 0, \partial_K^2 P \geq 0$

Consider a conventional GP prior on the option prices $C(K), P(K), K \in \mathbb{R}^+$, with a covariance function $\mathcal{K}_{\boldsymbol{\theta}}(K, K')$:

$$C(\boldsymbol{K}), P(\boldsymbol{K}) = \eta(\boldsymbol{K}) + y(\boldsymbol{K}), \quad \forall \boldsymbol{K} \in \mathbb{R}^+, \tag{1}$$

$$y(\boldsymbol{K}) \sim \mathcal{GP}\left(\mathbf{0}, \mathcal{K}_{\boldsymbol{\theta}}(\boldsymbol{K}, \boldsymbol{K}')\right), \tag{2}$$

where $\eta(\boldsymbol{K})$ is the mean process and $y(\boldsymbol{K})$ is a zero-mean GP. For convenience, we consider the case where $\eta(\boldsymbol{K}) = \mathbf{0}$, and choose $\mathcal{K}_{\boldsymbol{\theta}}$ as a squared exponential covariance function.[1] The model performance and economic implications in Section 5 are robust to different choices of covariance function $\mathcal{K}_{\boldsymbol{\theta}}$. [$\cdots$][2] As we observe option prices with noise from the option market, we add homogenous random noise $\epsilon \sim \mathcal{N}(0, \sigma^2_{noise}\mathbf{1})$ to the zero-mean GP:

$$\tilde{C}(K_i), \tilde{P}(K_i) = \tilde{y}_i(K_i) + \epsilon_i, \quad i = 1, \cdots, n \tag{3}$$

$$\tilde{y}(\boldsymbol{K}) \sim \mathcal{GP}\left(\mathbf{0}, \mathcal{K}_{\boldsymbol{\theta}}(\boldsymbol{K}, \boldsymbol{K}') + \sigma_{noise}\mathbf{1}\right). \tag{4}$$

To simultaneously encode the positivity, monotonicity and convexity constraint into $C(K), P(K)$ over the entire domain of $K \in \mathbb{R}^+$, we first follow Maatouk and Bay (2017) to construct a finite-dimensional approximation of the zero-mean GP:

$$y_m(x) = \sum_{j=0}^{m} y(t_j)\phi_j(x), \tag{5}$$

where $x \in \mathcal{D}$ is a *min-max*-standardized compact input space such that $\mathcal{D} = [0, 1]$,[3] m is the number of knots $t_1, \cdots t_m$ over the support $\mathcal{D}$,[4] $\phi_1, \cdots, \phi_m$ are hat basis functions given by:

$$\phi_j(x) := \begin{cases} 1 - \left|\frac{x - t_j}{\Delta_m}\right|, & \text{if } \left|\frac{x - t_j}{\Delta_m}\right| \leq 1, \\ 0, & \text{otherwise.} \end{cases} \tag{6}$$

For simplicity, we consider equally spaced knots $t_j = (j-1)\Delta_m$ with $\Delta_m = 1/(m-1)$ as in López-Lopera et al. (2018). Nonuniform knots or even knot-free design is also possible according

[1] Squared exponential covariance function: $\mathcal{K}_{\boldsymbol{\theta}}(K, K') = \sigma^2 \exp\left\{-\frac{(K-K')^2}{2l^2}\right\}$ with $\boldsymbol{\theta} = (\sigma, l)$.

[2] We may consider other $\mathcal{K}_{\boldsymbol{\theta}}$ somewhere in Appendix.

[3] $x = (K - K_{min})/(K_{max} - K_{min})$. We set $K_{min} = 0, K_{max} = 2 \times S_t e^{rT}$ to allow the left and right extrapolation.

[4] A larger m gives improved approximation and smoother y_m, but requires higher computational cost. We choose $m = 201$ to achieve a good trade-off between the two targets.

to López-Lopera et al. (2018).

Denote the option price at the knot t_j of the finite-dimensional GP y_m as ξ_j, for $j = 1, \cdots, m$. Encoding the no-arbitrage conditions (i) - (iii) from the option pricing theory is equivalent to computing the distribution of y_m conditionally on $y_m \in \mathcal{E}_{no-arbitrage}$, where $\mathcal{E}_{no-arbitrage}$ is a convex set of functions defined by no-arbitrage inequality constraints:

$$\mathcal{E}_{no-arbitrage} := \begin{cases} y \in C(\mathcal{D}, \mathbb{R}) \;\; s.t. \;\; y(x) \geq 0 \;\; \forall x \in \mathcal{D}, \\ y \in C(\mathcal{D}, \mathbb{R}) \;\; s.t. \;\; \partial_x y \leq 0 \;\; \text{(call)}, \partial_x y \geq 0 \;\; \text{(put)} \;\; \forall x \in \mathcal{D}, \\ y \in C(\mathcal{D}, \mathbb{R}) \;\; s.t. \;\; \partial_x^2 y \geq 0 \;\; \forall x \in \mathcal{D}. \end{cases} \tag{7}$$

The benefit of using the hat basis ϕ_j and the finite-dimensional approximation of GP y_m, according to Maatouk and Bay (2017), is that we have $y_m \in \mathcal{E}_{no-arbitrage} \Leftrightarrow \xi \in \mathcal{C}_{no-arbitrage}$, which means that encoding the no-arbitrage constraints to the entire domain of y_m is equivalent to encoding the same set of constraints to the designed knot points $t_1, \cdots, t_m$. Thus, the conditional distribution of y_m on infinite-dimensional inequality constraints $\mathcal{E}_{no-arbitrage}$ can be solved with only finite-dimensional inequality constraints:

$$\mathcal{C}_{no-arbitrage} := \begin{cases} \forall j = 1, \ldots, m : \overbrace{c_j \geq 0}^{\text{positivity}}, \\ \forall j = 2, \ldots, m : \overbrace{c_j - c_{j-1} \leq 0}^{\text{non-increasing}} \text{ (call)}, \; \overbrace{c_j - c_{j-1} \geq 0}^{\text{non-decreasing}} \text{ (put)}, \\ \forall j = 3, \ldots, m : \overbrace{c_j - c_{j-1} \geq c_{j-1} - c_{j-2}}^{\text{convexity}}. \end{cases} \tag{8}$$

Given the observed option prices $\tilde{C}(K), \tilde{P}(K)$, the finite-dimensional GP y_m can be written as the following in order to learn from observed prices with no-arbitrage constraints:

$$\underbrace{y_m(x_i) = \sum_{j=1}^{m} \xi_j \phi_j(x_i) + \epsilon_i,}_{\text{finite-dimensional GP with noise}} \;\; \text{s.t.} \;\; \underbrace{\begin{cases} y_m(x_i) = \tilde{y}_i & \text{(interpolation conditions)}, \\ y_m \in \mathcal{C}_{no-arbitrage} & \text{(inequality conditions)}. \end{cases}}_{\text{learn from the observed prices + no-arbitrage constraints}} \tag{9}$$

Since the vector $\boldsymbol{\xi} = [\xi_1, \cdots, \xi_m]'$ is a vector of realizations of the zero-mean GP $\boldsymbol{y}_m$ at the knots $t_1, \cdots, t_m$, and $\boldsymbol{\Phi}\boldsymbol{\xi} = \boldsymbol{y}_m = [y_m(x_1), \cdots, y_m(x_n)]'$ is the vector of the zero-mean GP realizations at available training points $x_1, \cdots, x_n$, we then have $\boldsymbol{\xi} \sim (\mathbf{0}, \boldsymbol{\Gamma}) \;\; s.t. \;\; \boldsymbol{\Phi}\boldsymbol{\xi} = \boldsymbol{y}, \boldsymbol{x} \in \mathcal{C}_{no-arbitrage}$, where $\boldsymbol{\Gamma} = \mathcal{K}_{\boldsymbol{\theta}}(t_i, t_j)_{1 \leq i,j \leq m}$ is the covariance function of y_m evaluated at knots t_i, t_j.

The no-arbitrage constraints in (8) can be encoded with the following linear operations:

$$\mathcal{C}_{no-arbitrage} = \left\{ \forall k = 1, 2, 3 : l_k \leq \sum_{j=1}^{m} \lambda_{k,j} c_j \leq u_k \right\}, \tag{10}$$

where $k = 1, 2, 3$ corresponds to positivity, monotonicity and convexity constraint respectively, l_k and u_k are the lower bound and upper bound of the inequalities. Therefore, the no-arbitrage constrained GP can be solved with:

$$\boldsymbol{\xi} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Gamma}) \quad \text{s.t.} \quad \begin{cases} \boldsymbol{\Phi\xi} = \mathbf{y} & \text{(interpolation conditions)}, \\ \boldsymbol{l} \leq \boldsymbol{\Lambda\xi} \leq \boldsymbol{u} & \text{(inequality conditions)}, \end{cases} \tag{11}$$

where $\boldsymbol{\Lambda} = (\lambda_{k,j})_{1 \leq k \leq 3, 1 \leq j \leq m}$ is an auxiliary matrix to encode the desired set of no-arbitrage inequalities, $\boldsymbol{l} = (l_k)_{1 \leq k \leq 3}$ and $\boldsymbol{u} = (u_k)_{1 \leq k \leq 3}$ are vectors of lower and upper bounds of the inequalities.[5] That is, we are interested in solving the conditional distribution:

$$\boldsymbol{\xi} | \{ \boldsymbol{\Phi\xi} = \boldsymbol{y}, \boldsymbol{l} \leq \boldsymbol{\Lambda\xi} \leq \boldsymbol{u} \}, \quad \boldsymbol{\xi} \sim (\mathbf{0}, \boldsymbol{\Gamma}), \tag{12}$$

where $p(\boldsymbol{\xi}) \propto \frac{1}{2} \boldsymbol{\xi}' \boldsymbol{\Gamma}^{-1} \boldsymbol{\xi}$ is proportional to the probability density of the unconditional $\boldsymbol{\xi}$ up to a scaling constant. Therefore, encoding no-arbitrage constraints to the finite-dimensional GP y_m is equivalent to solving $p(\boldsymbol{\xi})$ with restrictions $\boldsymbol{\Phi\xi} = \boldsymbol{y}$ and $\boldsymbol{l} \leq \boldsymbol{\Lambda\xi} \leq \boldsymbol{u}$.

To reduce the computational complexity, we do not sample from the posterior distribution of $\boldsymbol{\xi}$ according to López-Lopera et al. (2018), but rather focus on the *maximum-a-posteriori* of it. As in Chataigner et al. (2021), the MAP estimation $\boldsymbol{\xi}^*$ can be interpreted as the most probable predictions of $\boldsymbol{y}_m$ given the observed noisy option prices (interpolation conditions) and the no-arbitrage constraints (inequality conditions). $\boldsymbol{\xi}^*$ can be estimated by solving a quadratic programming:[6]

$$\boldsymbol{\xi}^* = \min_{\boldsymbol{\xi} \in \mathbb{R}^m} \left\{ \boldsymbol{\xi}' \boldsymbol{\Gamma} \boldsymbol{\xi} \middle| \boldsymbol{\Phi\xi} = \tilde{\boldsymbol{y}}, \boldsymbol{l} \leq \boldsymbol{\Lambda\xi} \leq \boldsymbol{u} \right\}, \tag{13}$$

and we then plug $\boldsymbol{\xi}^*$ back to (9) to obtain the predictions of no-arbitrage option prices:

$$y_m(x_i^{test}) = \sum_{j=1}^{m} \xi_j^* \phi_j(x_i^{test}). \tag{14}$$

[5] We can provide an example of the matrix $\boldsymbol{\Lambda}$ in Appendix.

[6] We use the open-source software package CVXOPT based on Python to solve the convex optimization.

2.2 Hybrid RND With Call and Put Options

We recover model-free RND with the predicted option prices from a no-arbitrage encoded finite-dimensional GP. We define $x := K/S_t$ as the moneyness of the observed option with strike price K and spot underlying price S_t at time t, and train the model with observations $(\tilde{x}, \tilde{y})$. One of the convenience of using $x = K/S_t$ is that $x - 1$ represents simple return of the underlying asset, and therefore $\partial_x y_m, \partial_x^2 y_m$ imply the RND of simple asset returns at the terminal time T, according to Breeden and Litzenberger (1978).

Denote y_m^{call}, y_m^{put} the learned call option and put option prices with no-arbitrage constraints. Since for the monotonicity constraint, we impose the weak form of no-arbitrage $\partial_x y_m^{call} \leq 0$ and $\partial_x y_m^{put} \geq 0$, which means the learned price is not bounded from above. We get around this issue by using only the predicted out-of-the-money option prices to recover the model-free RND. Consider a vector of test points $x^{test} = [0.0, 0.01, \cdots, 1.99, 2.0]'$, we calculate the predictions of out-of-the-money call option prices $y_m^{call}(x^{test})\ \forall x^{test} \in [1.0, 2.0]$, and out-of-the-money put option prices $y_m^{put}(x^{test})\ \forall x^{test} \in [0.0, 1.0)$. Since the encoded monotonicity constraints guarantee that $\lim_{x \to 2.0} y_m^{call}(x) = 0$ and $\lim_{x \to 0} y_m^{put}(x) = 0$, and the sufficient traded options around $x = 1.0$ to train the GP ensure the regularity of $y_m^{call}(x = 1.0)$ and $y_m^{put}(x = 1.0)$,[7] the regularity of the recovered $F_{call}^{\mathbb{Q}} = \partial_x y_m^{call}(x)\ \forall x \in [1.0, 2.0]$ and $F_{put}^{\mathbb{Q}} = \partial_x y_m^{put}(x)\ \forall x \in [0.0, 1.0)$ can therefore be maintained.

We proceed to append the recovered RND $f_{call}^{\mathbb{Q}}$ with predicted out-of-the-money call option prices y_m^{call} to the recovered RND $f_{put}^{\mathbb{Q}}$ with predicted out-of-the-money put option prices y_m^{put}. This means that the entire model-free RND $f_{hybrid}^{\mathbb{Q}}$ of the underlying asset returns is composed by the information from both the call and put options:

$$f_{hybrid}^{\mathbb{Q}}(x) := \begin{cases} e^{r\tau} \frac{\partial^2 y_m^{put}(x)}{\partial x^2} & \text{if } 0 \leq x \leq 1.0, \\ e^{r\tau} \frac{\partial^2 y_m^{call}(x)}{\partial x^2} & \text{if } 1.0 < x \leq 2.0, \end{cases} \tag{15}$$

where the left (right) portion of the hybrid RND is governed by the out-of-the-money put (call) option prices from the trained finite-dimensional GP $y_m(x)$ with no-arbitrage constraints.

[7] Otherwise, the regularity of the predicted option prices might not be guaranteed when $x \to 0$ for y_m^{call} and $x \to 2.0$ for y_m^{put}.

3 Simulations

We simulate the Heston-type stochastic volatility model according to Heston (1993). The Heston (1993) model describes the dynamics of the asset price S_t and its variance v_t under the risk-neutral probability measure $\mathbb{Q}$ with the following system of stochastic differential equations:

$$dS_t = rS_t dt + \sqrt{v_t} S_t dW_t^S, \tag{16}$$

$$dv_t = \kappa\left(\theta - v_t\right) dt + \sigma\sqrt{v_t} dW_t^v, \tag{17}$$

where S_t is the price of the underlying asset at time t, v_t is the instantaneous variance of the asset price, r is the dirt rate of the asset price, k is the speed of mean reversion of the variance, θ is the long-term mean of the variance, σ is the volatility of the variance. W_t^S and W_t^σ are two Brownian motions with correlation ρ that add stochasticity to the asset price and the variance process respectively.

Synthetic option prices are generated from the Heston (1993) model. We rely on the characteristic function of the Heston (1993) model to calculate the associated theoretical RND by using the Fourier inversion. We follow Gil-Pelaez (1951) to evaluate the Fourier inversion of the characteristic function more efficiently. The technical details of the implementation of the Heston (1993) model RND computation can be found in Appendix A.

We simulate the Heston (1993) model with the following set of parameters: $r = 0.05, \rho = 0.8, \kappa = 3, \theta = 0.2, \sigma = 0.1$, with the instantaneous variance $v_0 = 0.15$ and the spot asset price $S_t = 100$. We denote the theoretical RND implied in this model as $f^{\mathbb{Q}}_{Heston}$. To capture the real characteristics of traded option prices, we consider a range of different sample size levels l and noise conditions σ. Given $N = 100$ synthetic option prices from the perfect Heston (1993) model, the sample size l controls how many option prices are sampled from the N prices, and the noise condition σ defines the Gaussian random noise added to the sampled prices. Different l and σ attempt to simulate different levels of option liquidity and pricing errors respectively. In our simulation study, we consider $l = \{1.0, 0.5, 0.25, 0.05\}$ and $\sigma = \{0.0, 0.3, 0.6, 0.9, 1.2\}$. We compare the theoretical RND $f^{\mathbb{Q}}_{Heston}$ with the CGP-based RND $f^{\mathbb{Q}}_{CGP}$ derived from the GP with no-arbitrage constraints. We also consider other model-free RND from the IV-based approaches including IVSPL-, IVGEV- and IVSVI-based RND as the benchmark recovered RND.

To evaluate the entire RND recovered by different approaches, we first perform a quantile-by-quantile comparison. The quantile loss at a given quantile is defined as the absolute deviation of the quantile from the recovered RND to the theoretical quantile from the Heston (1993) model. For each scenario with a sample size l and a noise level σ, the simulation is repeated 100 times and we then average the losses. From Figure 1, it is clear that the CGP-based RND $f^{\mathbb{Q}}_{CGP}$ is closest to the "true" RND $f^{\mathbb{Q}}_{Heston}$ in terms of the considered quantiles, and is superior to other IV-based model-free RND particularly when the sample size is small and the pricing noise is large.

[Insert Figure 1 about here]

We proceed to compare the recovered RND with the theoretical RND using probability density divergence measures. The divergence measures calculate the distance of the recovered RND to the "true" RND from the Heston (1993) model. We consider the following divergence measures: L^2 (or Euclidean), Kullback-Leibler (KL) as in Kullback and Leibler (1951), Jensen-Shannon (JS) as in Tishby et al. (2000), Wasserstein as in Rubner et al. (2000) and Hellinger as in Thomas and Joy (2006). These measure the information loss when using the recovered RND from respective approaches to approximate the theoretical RND. The calculation of each divergence measure can be found in Appendix, and we again average the divergence across the 100 simulations. Figure 2 shows the results of each RND divergence measure across different RND-recovery approaches. The information loss of the CGP-based RND is the smallest among all the considered RND and is also robust to different definitions of divergence.

[Insert Figure 2 about here]

4 Data

In this section, we describe our data on stocks and stock options that we employ in our empirical study. Stock and stock option data are from OptionMetrics through Wharton Research Data Services (WRDS). The risk-free rate data are obtained from OptionMetrics as well with linear interpolation of the zero coupon yield curve. Our sample period is from 04 January 1996 to 31 December 2022. We focus on stocks that are in the S&P 500 index on the first trading day of each year, and the stock sample is held fixed throughout the year.

For stocks in our sample, we collect the stock options written on them. We collect options with time-to-maturity from 7 days (one week) to 360 days (one year) as these stock options are actively traded and hence more informative. Most of the traded stock options are American style (with an "A" exercise style flag in OptionMetrics), and we therefore collect only American stock options from OptionMetrics. We do not adjust for the early-exercise premium in American options, as this premium should be tiny for the less than one-year option terms we consider. We apply a series of filters to option data to discard thinly traded stock options and possible record errors. We eliminate options with zero volume or open interest, and options with zero best bid price are excluded. To ensure meaningful stock-level RND can be recovered from traded options, we require for both put and call that there are at least five observed options, with more than two of them being out-of-the-money. The filters we apply are the least restrictive ones among the stock option literature. The standard option data filters commonly used as suggested in Goyal and Saretto (2009) may substantially miss useful information implied in options as only the liquid and noiseless enough options are left. The less restrictive filters allow us to test whether our proposed stock-level RND-recovery with constrained GP is able to unlock economic information from illiquid and noisy options that are traded. We also discard stocks whose options are inactively traded. The complete list of stock and stock option selection criteria can be found in Appendix B.

We recover stock-level RND using the filtered stock options and then calculate implied tail risk measures and implied moment-based return predictors. The implied tail risk measures are calculated for each stock and for each week on Wednesday. The implied moment-based return predictors are calculated for each stock and for each month on the last trading date of the month. We also calculate stock-based measures by only using observed stock returns. Since the calculation of stock-based measures requires historical stock data, the comparison of option-implied measures with stock-based measures in our empirical analysis starts at the beginning of 1997 instead of 1996 which is the beginning of the stock option sample, as we require one year of stock data in history to calculate stock-based measures.

We report in Table 1 the cross-sectional and time-series average of some selected option and stock characteristics, including the number of traded option contracts, the time-to-maturity, the moneyness, the trading volume, and the open interest of options. We follow Barone-Adesi (2016) to calculate the minimum $\alpha^{\mathbb{Q}}$ for each stock and for each available time-to-maturity on a given day. The minimum $\alpha^{\mathbb{Q}}$ implies the minimum cumulative probability that can be recovered from raw options without any further estimation techniques. We use only put options to calculate the minimum $\alpha^{\mathbb{Q}}$ as the left tail of the implied distribution is concerned most of the time. In Table 1, we divide stocks into four categories according to their firm size which is computed as the stock's close price times its share outstanding. The "Small" category includes small-sized stocks whose size is below the 25% quantile, the "Small-Medium" category includes small-to-medium-sized stocks whose size is between the 25% and the 50% quantile, the "Medium-Large" category includes medium-large-sized stocks whose size is between the 50% and 75% quantile, and the "Large" category includes stocks whose size is above the 75% quantile. Stock and option characteristics are calculated monthly on the end date of each month.

[Insert Table 1 about here]

5 Empirical Analysis

In this section, we extensively evaluate the predictive information encoded in the stock-level RND estimated with our CGP-based RND-recovery method and compare it with the commonly-used RND-recovery with implied volatilities in the literature. We compress the information of the recovered RND using a comprehensive list of risk metrics as introduced in Section 5.2, including tail risk measures and moment-based cross-sectional stock return predictors. The trapezium rule is used to evaluate the integral when using the recovered RND to calculate the risk metrics.[8]

Section 5.1 nonparametrically evaluates the regularity of the recovered stock-level RND, using the stock options with one-month time-to-maturity as an example. Section 5.3 investigates the informativeness of the RND-based implied tail risk measures for the left-tail risk of stock returns. Section 5.4 tests the additional economic information of moment-based cross-sectional stock return predictors when the predictors are constructed with the CGP-based RND. Section 5.5 explores the reactions of stock-level RND around the stock's earning announcement dates to the stock returns in this period.

5.1 RND Estimation Evaluation

We follow Diebold et al. (1997) to evaluate the model-free RND recovered from option prices. Evaluating the RND estimation is a nontrivial task as the realization of the density estimation is not observable in complete.

Given a series of observed stock returns $r_{i,t}$ for stock i at time t, and a sequence of the RND estimation $f^{\mathbb{Q}}_{i,t-1}(r_t)$ formed at time $t-1$, we rely on the probability integral transform (PIT) z_t,

[8]The trapezium rule with discretized integrand: $\int_a^b f(x)dx \approx \sum_{k=1}^{N} \frac{f(x_{k-1})+f(x_k)}{2}\Delta x_k$.

of the realization of the return process taken with respect to the RND estimation, as suggested in Diebold et al. (1997), to test the specification of the density. The PIT at time t is defined as $z_{i,t} = \int_{-\infty}^{r_{i,t+1}} f_{i,t}^{\mathbb{Q}}(r_{t+1})dr$, with respect to the RND estimation for stock i formed at time $t-1$ for stock return at time t. According to Diebold et al. (1997) and Berkowitz (2001), if $f_{i,t}^{\mathbb{Q}}$ is correctly specified, then $z_{i,t} \overset{i.i.d.}{\sim} \mathcal{U}(0,1)$.

Denote $z_{i,t}^{Q}, z_{i,t}^{F}$ the quantiles and probabilities from the empirical cumulative density function of $z_{i,t}$, we test the uniformity of $z_{i,t}$ by regressing $z_{i,t}^{F}$ on $z_{i,t}^{Q}$. The $z_{i,t}$ from a sequence of well-specified RND estimations will exhibit a coefficient that is close to one and an intercept that is close to zero.[9]

[Insert Figure 3 about here]

Figure 3 illustrates the RND evaluation results for stock-level RND recovered from options with one-month time-to-maturity as an example. We select 16 different representative stocks across the industry of technology, healthcare, financial and industrial to evaluate the process of the recovered one-month stock-level RND. The upper row of each subplot in Figure 3 shows the distribution of $z_{i,t}$, and the lower row shows the line of $\left(z_{i,t}^{Q}, z_{i,t}^{F}\right)$. The dashed grey line is a 45° straight line which represents the case of a well-specified process of RND estimations. We observe that only the line of $\left(z_{i,t}^{Q}, z_{i,t}^{F}\right)$ generated from the processes of the CGP-based stock-level RND estimations are consistently close to the dashed grey line, and the lines of $\left(z_{i,t}^{Q}, z_{i,t}^{F}\right)$ generated from the IV-based RND estimations present obvious departure from the dashed grey line, which highlights the superiority of the CGP-based method for the stock-level RND estimation.

5.2 Generalized Implied Risk Metrics Estimation

This section provides a generalized approach to estimate risk metrics of the underlying asset returns implied in options. Unlike the existing model-free estimation of some implied risk metrics, such as the implied moments of Bakshi et al. (2003), and the implied VaR and ES in Barone-Adesi (2016), the generalized estimation we use relies only on the RND recovered. [$\cdots$][10]

Given that the model-free RND $f^{\mathbb{Q}}$ is recovered "exactly" over the support $x \in [0, 2.0]$, a wide variety of risk metrics can be evaluated numerically without the estimation of further parameters.

[9] We do not use the popular Kolmogorov-Smirnov (KS) test, as this method is less sensitive to "irregular" $z_{i,t}$ that is outside of $\mathcal{U}(0,1)$.

[10] We might need a short discussion of the difference between our generalized estimation and the model-free approximation of Bakshi et al. (2003) and Barone-Adesi (2016).

We focus on the estimation of the following implied risk metrics:

$$\text{Implied VaR:} \quad VaR^{\mathbb{Q}}_{t,T}(R;\alpha) := -\inf\left\{z \in \mathbb{R} | F^{\mathbb{Q}}_{t,T}(z) \leq \alpha\right\}, \tag{18}$$

$$\text{Implied ES:} \quad ES^{\mathbb{Q}}_{t,T}(R;\alpha) := 1/\alpha \int_{-\infty}^{VaR^{\mathbb{Q}}} R f^{\mathbb{Q}}_{t,T}(R) dR, \tag{19}$$

$$\text{Implied Volatility:} \quad Vol^{\mathbb{Q}}_{t,T}(R) := \int (R-\mu^*)^2 f^{\mathbb{Q}}_{t,T}(R) dR, \tag{20}$$

$$\text{Implied Skewness:} \quad Skew^{\mathbb{Q}}_{t,T}(R) := \int \left(\frac{R-\mu^*}{\sigma^*}\right)^3 f^{\mathbb{Q}}_{t,T}(R) dR, \tag{21}$$

$$\text{Implied LPM:} \quad LPM^{\mathbb{Q}}_{t,T}(R;\bar{h}) := \int_{-\infty}^{\bar{h}} \left(R-\bar{h}\right)^2 f^{\mathbb{Q}}_{t,T}(R) dR, \tag{22}$$

$$\text{Implied UPM:} \quad UPM^{\mathbb{Q}}_{t,T}(R;\bar{h}) := \int_{\bar{h}}^{\infty} \left(R-\bar{h}\right)^2 f^{\mathbb{Q}}_{t,T}(R) dR. \tag{23}$$

The selected implied risk metrics from (18) to (23) emphasis different aspects of the recovered stock-level RND $f^{\mathbb{Q}}$. The implied VaR and ES focus on the extreme left tail of the RND, and the implied volatility and skewness focus on the entire shape of the RND. Compared with the implied volatility, the estimation of implied skewness demands a more well-specified stock-level RND as the calculation of the third-order central moment is sensitive to outliers in the estimated RND. The implied LPM and UPM measure the partial moments of the return distribution instead of the complete moments, which also builds on the notion of semivariance as in Markowitz (1952). In our empirical analysis, we construct tail risk measures and moment-based predictors for stock returns using the above definitions of risk metrics.

5.3 Implied Tail Risk Measures

Different from the standard RND-recovery as in Aït-Sahalia and Lo (2000) and Figlewski (2010) where the tails of the RND are extrapolated assuming a particular distribution with heavy tails, such as GEV, student-t and skewed-t distribution, as no-arbitrage option prices are not available outside of the observed range of moneyness, our method does not append any ad-hoc distribution to the tails. Instead, since we learn the economics-aware representation of the option pricing relation from observed options, no-arbitrage prices in the tails can be predicted from the constrained GP without additional assumptions. This allows us to directly estimate the model-free RND tails together with the median part of the RND, conditional on the observed option prices only. As we are interested in the extreme loss measures of stock returns implied in their options, we use the recovered left RND tail to calculate the implied tail risk measures.

5.3.1 Predict The Quantile of Stock Returns

The implied VaR and ES measure the magnitude of the extreme loss of stock returns implied in stock options. The implied LPM measures the left semi-variance of stock returns, that is how volatile the stock return will be conditional on the stock return being negative. These three tail risk measures have different dependencies on the recovered RND. The implied VaR requires only a point x^{α} at which the implied cumulative probability is α from the left tail. The implied ES requires the left tail below x^{α}. The implied LPM requires the entire left half of the RND (that is the RND below $x^{50\%}$).

Stock-level quantile regressions in time-series. We first analyze the predictive information provided by the tail risk measures estimated from the CGP-based stock-level RND. As we are interested in the relation of the conditional left quantile of stock returns with the estimated implied tail risk measures, we employ the stock-level quantile regression model as follows:

$$Q_{r_{i,t+1wk}}(\theta|\mathcal{M}_{i,t}^{tail}) = \beta_0(\theta) + \beta_1(\theta)\mathcal{M}_{i,t}^{tail} + \boldsymbol{Z}_{i,t}^{\top}\boldsymbol{\beta}_Z(\theta), \tag{24}$$

where $Q_{r_{i,t+h}}(\theta \mid \mathcal{M}_{i,t}^{\text{tail}})$ represent the conditional θ-quantile of the weekly returns for stock i at one week after the implied tail risk measure estimation time t. The set $\mathcal{M}_{i,t}^{\text{tail}} = \{\text{VaR}_{i,t}^{\mathbb{Q}}, \text{ES}_{i,t}^{\mathbb{Q}}, \text{LPM}_{i,t}^{\mathbb{Q}}\}$ comprises the implied tail risk measures that we consider. $\boldsymbol{Z}_{i,t} = \{\text{VaR}^{\mathbb{P}}, \text{Beta}_{\text{mkt}}, \text{Coskew}, \text{Size}, \text{Ivol}, \text{Ivol}^{\perp}, \text{Iskew}, \text{Ikurt}, \text{Idiovol}, \text{Idiovol}^{\perp}, \text{Idioskew}\}$ denotes the list of control variables. Both the implied tail risk measures and the control variables in the model (24) are estimated at time t. We choose $\theta = 10\%$ to identify losses as extreme as the 10% left quantile of the weekly stock returns.

The regression coefficients of the predictive quantile model in (24) are estimated by minimizing the quantile-weighted absolute value of errors according to Koenker and Bassett (1978):

$$\hat{\boldsymbol{\beta}}_{\theta} = \underset{\boldsymbol{\beta}_{\theta} \in \mathbb{R}^p}{\text{argmin}} \sum_{t=1}^{T_i - 1wk} \rho_{\theta}\left(r_{i,t+1wk} - \boldsymbol{x}_{i,t}^{\top}\boldsymbol{\beta}\right), \tag{25}$$

where T_i is the number of available weekly return observations for stock i, $r_{i,t+1wk}$ is the realized weekly stock return at time $t + 1wk$, and $\boldsymbol{x}_{i,t}$ is the vector of included variables.

The quantile regression model provides a more rigorous evaluation for the performance of implied tail risk measures in anticipating the extreme left-tail risk of stock returns, as suggested by Gaglianone et al. (2011), compared with the violation-based backtesting as in Christoffersen (1998) and Engle and Manganelli (2004). Since the inclusion of the investor's subjective risk aversion in the implied tail risk measures, they are not directly comparable to the realized stock returns under the physical probability measure in magnitude, which therefore makes the violation-based tests less indicative for the informativeness of the considered tail risk measure.

[Insert Table 2 about here]

Table 2 reports the percentage of stock-level time-series quantile regressions where $VaR^{\mathbb{Q}}$ is statistically significant for predicting the 10% quantile of weekly returns, at the 10%, 5% and 1% significance level, across the entire cross-section of stocks in our sample. We observe that the implied VaR derived from the CGP-based RND is a stronger predictor for the left tail of the stock's weekly returns than the VaR derived from historical stock price data with a GARCH-GJR (1,1) model. The $VaR^{\mathbb{Q}}$ from the CGP-based RND unlocks statistically significant predictive information for about 17.57% more stocks in average than the stock-based $VaR^{\mathbb{P}}$, when considering a univariate regression model of (24) with $VaR^{\mathbb{Q}}$ and $VaR^{\mathbb{P}}$ separately as reported in the "Option" and "Stock" column. The orthogonal information provided by $VaR^{\mathbb{Q}}$ is more evident when a bivariate model of (24) where $VaR^{\mathbb{Q}}$ and $VaR^{\mathbb{P}}$ are included jointly, as the percentage of regressions where $VaR^{\mathbb{Q}}$ is statistically significant is 52.17% higher than the case of $VaR^{\mathbb{P}}$ in average. The percentage of regressions where $VaR^{\mathbb{Q}}$ is significant remains considerably high when other control variables are considered in the quantile model. Similar observations can be found for the $ES^{\mathbb{Q}}$ and $LPM^{\mathbb{Q}}$ from the CGP-based RND, as presented in Table 3 and Table 4. This implies that the extrapolated option prices over small moneyness with no-arbitrage constraints encode useful information which is superior to the information in historical stock returns for the left tail of stock returns.

[Insert Table 3 about here]

[Insert Table 4 about here]

We proceed to compare the informativeness of the implied tail risk measures generated from different RND-recovery methods. We use the IV-based methods as proposed in Figlewski (2010) but with different extrapolation techniques as the benchmark models. Specifically, we interpolate the observed IVs with cubic spline, and consider constant extrapolation—extrapolating the left and right part of the IV smirk with the nearest observed IV, and GEV extrapolation—extrapolating the left-tail repriced Black-Scholes put option prices according to Hamidieh (2017). In addition, we also consider the RND-recovery based on parametric no-arbitrage IV smirk. That is, we calibrate a stochastic volatility inspired (SVI) model proposed by Gatheral (2004) to the observed IVs.

[Insert Figure 4 about here]

Figure 4 shows the informativeness of the implied tail risk measures estimated from the RND-recovery methods we consider. We calculate the out-of-sample R^1 measure to compare the performance of the implied tail risk measure estimation. The out-of-sample R^1 measure is calculated from the stock-level univariate quantile model where the implied tail risk measure to test is the only predictor. The averaged R^1 is the average of out-of-sample R^1 across all stocks in the sample. For each stock-level univariate quantile regression, we use the first 70% of the weekly observations $\left(r_{i,t+1wk}, \mathcal{M}_{i,t}^{tail}\right)$ as the in-sample data to estimate the coefficients $\beta_0(\theta = 10\%)$ and $\beta_1(\theta = 10\%)$, and keep them unchanged to calculate the out-of-sample R^1 in the remaining 30% of the observations according to Koenker and Machado (1999).[11]

The informativeness of the implied tail risk measure from the recovered RND depends on the predictive information encoded in the extrapolated left tail of the RND. We observe from Figure 4 that the implied tail risk measures estimated from the left RND tail recovered with the no-arbitrage constrained GP provide the highest averaged R^1 across full-sample, in-sample and as well as out-of-sample of the weekly observations. Notably, for the implied ES and implied LPM which can be regarded as nonlinear transformations of the RND, the outperformance of the CGP-based RND left tail compared with the constant, GEV and SVI-parameterized extrapolated RND left tail is more apparent. For the implied ES and implied LPM, we observe nearly zero or even negative averaged out-of-sample R^1 from the GEV and SVI-parameterized left RND tail, which implies that extrapolating the left-tail of the stock-level RND with parametric component does not help to capture predictive information for the tail risk of stock returns in out-of-sample.

Table 5 reports the summary statistics of the entire distribution of the in-sample and out-of-sample R^1 across all considered stocks. It confirms that the outperformance of the CGP-based left RND tail is not only observed in the mean of stock-level R^1 as shown in Figure 4, but is also observed in other quartiles of the stock-level R^1. This ensures that the CGP-based left RND tail unlocks more predictive information than other IV-based left RND tails for most of the stocks, regardless of whether the options provide tiny or considerable information in nature.

[11] The calculation of out-of-sample R^1 can be put in Appendix.

[Insert Table 5 about here]

Panel quantile regression with fixed effects. To examine whether the superior predictive information for the tail risk of stock returns provided by the implied tail risk measures persists in a large panel of stocks, we estimate panel quantile regression models with fixed effects. By considering the fixed effects, unobserved characteristics that are invariant across stocks over time can be controlled. [$\cdots$][12] While the stock-level quantile regressions in time-series investigate the predictive power of the implied tail risk measures for each individual stock, the panel quantile regressions test whether they are still significant predictors for the tail risk of stock returns in cross-section. We follow Machado and Silva (2019) to estimate panel quantile regressions with fixed effects:

$$Q_{i,t+1wk}\left(\theta \middle| \mathcal{M}_{i,t}^{tail}\right) = (\alpha_i + \delta_i q(\theta)) + \beta^{tail}\mathcal{M}_{i,t}^{tail} + \boldsymbol{Z}_{i,t}^{\top}\boldsymbol{\beta}^{Z} + \boldsymbol{Z}_{i,t}^{\top}\boldsymbol{\gamma} q(\tau), \tag{26}$$

given observations $\left(\boldsymbol{r}_{i,t+1wk}, \boldsymbol{X}_{i,t}^{\top}\right)^{\top}$ from the panel of N individual stocks $i = 1, \cdots, n$ over T time periods, $t = 1, \cdots, T$. $\boldsymbol{X}_{i,t}^{\top}$ is the vector of included predictors with control variables $\left(\boldsymbol{X}_{i,t} = \left[\mathcal{M}_{i,t}^{tail}, Z_{i,t}\right]\right)$. The parameters (α_i, δ_i) capture the stock i fixed effects, and $\boldsymbol{Z}^{\top}$ is a vector of transformation of $\boldsymbol{X}^{\top}$. The scalar coefficient $\alpha_i(\theta) := \alpha_i + \delta_i q(\theta)$ is the θ-quantile fixed effects for stock i, with $\mathbb{P}\left\{\delta_i + \boldsymbol{Z}_{i,t}^{\top} > 0\right\}$.

Table 6 presents the results of the fixed effects panel quantile regressions. We observe that both the option-implied and the stock-based VaR are statistically significant predictors for the 10%-quantile of the stock return in the next week, although they are in different magnitudes as they are estimated from different probability measures. When we consider a bivariate case with $VaR^{\mathbb{Q}}$ and $VaR^{\mathbb{P}}$, the coefficient of $VaR^{\mathbb{P}}$ decreases to near zero and its absolute value of t-statistic drops to 0.51, while the coefficient of $VaR^{\mathbb{Q}}$ is almost not changed with its absolute value of t-statistic boosted to 29.33. This observation highlights the fact that there is extra predictive information encoded in the $VaR^{\mathbb{Q}}$ from the CGP-based RND and the information is strongly orthogonal to the information encoded in $VaR^{\mathbb{P}}$. From Table 6, the predictive power of $VaR^{\mathbb{Q}}$ is statistically significant in the panel data of cross-sectional weekly stock returns over the considered period of time, and the predictive power is also robust to other control variables. The same panel quantile regressions with fixed effects for the $ES^{\mathbb{Q}}$ and $LPM^{\mathbb{Q}}$ from the CGP-based stock-level RND can

[12] More detailed explanation of why using panel quantile regression with fixed effects? How it shows the tail risk predictability of implied tail risk metrics in panel?

be found in Table 7 and Table 8.

[Insert Table 6 about here]

[Insert Table 7 about here]

[Insert Table 8 about here]

5.3.2 Predict Crash Risk of Stock Returns

The above analysis of tail risk predictability uses the stock's raw returns. To make sure that the tail risk predictability provided by the implied tail risk measures is not distorted by market-wide factors captured in raw returns, we test whether the implied tail risk measures can still predict the crash risk of stock returns. The estimation of the stock's crash risk depends on its residual returns instead of raw returns. We follow Hutton et al. (2009) to estimate the stock's residual returns from an expanded index model:

$$r_{i,t} = \alpha_0 + \alpha_1 r_{mkt,t-2} + \alpha_2 r_{mkt,t-1} + \alpha_3 r_{mkt,t} + \alpha_4 r_{mkt,t+1} + \alpha_5 r_{mkt,t+2} + \epsilon_{i,t}, \tag{27}$$

where $r_{i,t}$ is the raw return of stock i at time t, $r_{mkt,t}$ is the raw return of the S&P 500 index at time t. The nonsynchronous trading is captured by including lead and lag terms for the S&P 500 index. To avoid the estimation error for crash risk due to the highly skewed residuals in (27), we apply the log transformation according to Chen et al. (2001) and Hutton et al. (2009). That is, we calculate $\tilde{r}_{i,t} := \log(1+\epsilon_{i,t})$ as the residual return of stock i at time t. For each stock in the sample, we calculate $\tilde{r}_{i,t}$ in a yearly basis, with the condition that there is at least 50 day-observation for raw stock returns and the index returns in the year.

We use the negative coefficient of skewness (NCSKEW), and the down-to-up volatility (DUVOL) as in Chen et al. (2001) to capture the stock's crash risk. These two measures are calculated using the daily residual stock returns within a specified investment horizon from time t. For example, we calculate the $NCSKEW$ of stock i over the horizon from time t to T as the follows:

$$NCSKEW_{i,t:t+T} = -\left(n(n-1)^{3/2}\sum \tilde{r}_{i,t}^3\right) \Big/ \left((n-1)(n-2)\left(\sum \tilde{r}_{i,t}^2\right)^{3/2}\right). \tag{28}$$

A higher $NCSKEW_{i,t:T}$ corresponds to a more left-skewed residual stock return distribution. We calculate the DUVOL of stock i from time t to T as follows:

$$DUVOL_{i,t:T} = \log \left\{ (n_u - 1) \sum_{down} \tilde{r}_{i,t}^2 \middle/ \left((n_d - 1) \sum_{up} \tilde{r}_{i,t}^2 \right) \right\}, \tag{29}$$

where n_u and n_d are the number of up and down days, and up (down) days are the days where the residual return is above (below) the period mean. Again, a higher $DUVOL_{i,t:T}$ corresponds to a more left-skewed residual return distribution within the period.

Cross-sectional regressions of stock's crash risk. We employ the Fama and MacBeth (1973) two-step procedure to test whether the implied tail risk metrics can predict the stock's crash risk in cross-section. Implied tail risk measures are sampled on the end date t of each month, and the stock's crash risk is calculated with the returns within a looking-ahead period that starts from time t and ends at $t+h$. We estimate the cross-sectional regression models as follows:

$$CRASH_{i,t+h} = \beta_t^0 + \beta_t^1 \mathcal{M}_{i,t}^{tail} + \boldsymbol{Z}_{i,t}^\top \boldsymbol{\beta}_t^Z + \epsilon_{i,t}, \tag{30}$$

where $CRASH_{i,t+h} = \{NCSKEW_{i,t:t+h}, DUVOL_{i,t:t+h}\}$, $\mathcal{M}_{i,t}^{tail}$ and $\boldsymbol{Z}_{i,t}$ are the tested implied tail risk measure and the list of control variables as defined previously.

[Insert Table 11 about here]

Table 11 presents the regression results using $NCSKEW_{i,t+h}$ as the stock crash risk proxy, with h equals 6 months. The choice of 6 months is to include more extreme realizations of residual stock returns to make the crash proxy be able to capture extreme crash risk. We observe that the $VaR^{\mathbb{Q}}$ from the CGP-based RND is a statistically significant predictor for the stock's $NCSKEW_{i,t+6mo}$ in cross-section, with t-statistic of -5.90. It is interesting to observe a strong negative association between $VaR^{\mathbb{Q}}$ and $NCSKEW_{i,t+6mo}$. One possible interpretation is that $VaR^{\mathbb{Q}}$ increases with the thickness of the left tail for the stock returns, which is opposed to the crash risk proxy that measures the left-skewness for the stock returns by construction. The statistically significant predictive power of $VaR^{\mathbb{Q}}$ is not affected after controlling for the firm-characteristics, the BKM's option-characteristics and the stock's idiosyncratic risks.

[Insert Table 14 about here]

Table 14 presents the regression results using $DUVOL_{i,t+6mo}$ as the stock crash risk proxy. The results are similar to what can be observed from Table 11. The $VaR^{\mathbb{Q}}$ from the CGP-based RND is a strong predictor for $DUVOL_{i,t+6mo}$ in the stock's cross-section as well. The adjusted R^2 from the $VaR^{\mathbb{Q}}$-only regression is at 2.50% for $DUVOL_{i,t+6mo}$, which is slightly higher than the case of $NCSKEW_{i,t+6mo}$. This is because that the calculation of $NCSKEW$ is more sensitive to the noisy outlier residual stock returns in the considered horizon than $DUVOL$.

In Table 17, we benchmark the implied tail risk measures estimated with our CGP-based RND to those from the RND recovered by the IV-based methods, in terms of predicting the stock's crash risk. From Model 1 to Model 5 in Table 17, control variables are added progressively. In this test, we consider more looking-ahead horizons, including 1-month, 3-month and 6-month period, to investigate whether the stock's crash risk predictability of the implied tail risk measures is robust to the different calculation horizon of the crash risk proxy. IV-based methods with parametric left RND tail extrapolations perform closely in terms of the t-statistics of $VaR^{\mathbb{Q}}$. The exception is the SVI-parameterized left extrapolation. We observe that in model 3 and model 5 in Panel A, and model 5 in Panel B, with the crash proxy calculation horizon of 1 month, the $VaR^{\mathbb{Q}}$ from the SVI-parameterized extrapolated left RND tail is not statistically significant. In addition, we also report the averaged R^2 across 5 models for each horizon in Table 17. It is notable that the no-arbitrage constrained GP extrapolation provides the highest averaged R^2 over the 1-month, 3-month and 6-month horizon.

[Insert Table 17 about here]

5.4 Implied Moment-Based Stock Return Predictors

The estimated implied tail risk measures focus on the left part of the recovered stock-level RND. In this section, we evaluate the predictive information encoded in the complete RND, by using the moment-based measures including the implied volatility, skewness, LPM and UPM that are introduced in Section 5.2. Since the complete stock-level RND is used, the informativeness of the moment-based measures from the RND majorly depends on two aspects of the RND-recovery method: the capability of encoding forward-looking information from observed stock option prices and the regularity of the recovered RND. The regularity of the RND is important as too much irregularities will distrot the estimation of moment-based predictors based on the RND.

We estimate the option-implied volatility and skewness based on the recovered stock-level RND. The estimation of the implied volatility and skewness can also be calculated according to the model-free approximations as proposed in Bakshi et al. (2003). We start with analizing the estimations of the implied volatility and skewness from different approaches. Table 18 summarizes the distribution and the correlations of the monthly stock-level implied volatility and skewness across the cross-section of stocks in our sample.

The Panel A shows that the distribution of the stock-level implied volatility does not vary significantly across different estimation techniques, but this is not the case for the distribution of the stock-level implied skewness. We observe notable differences in the mean, the median, the 5^{th} and 95^{th} percentile, from the distribution of the implied skewness by different RND-recovery approaches. Among all the considered techniques, the CGP-based RND-recovery provides the estimation of implied skewness with the smallest standard deviation, which is a hint that this technique generates more stable and noiseless third-moment estimations.

[. . .][13]

The Panel B shows that the cross-sectional correlations of the monthly stock-level implied volatility across different approaches are high. Although the CGP-based method learns directly from the observed stock option prices with no-arbitrage constraints, which is methodologically different from the other IV-based approaches, the average cross-sectional correlation of the CGP-based implied skewness with others is still as high as 0.83. This is actually expected and also confirms that the RND-based estimation of implied moment is correctly specified. However, the cross-sectional correlations of the implied skewness are notably smaller. The highest implied skewness correlation is observed between the CGP-based and the BKM method at 0.02.

[Insert Table 18 about here]

5.4.1 Portfolio Sorting

We evaluate the economic information provided by the implied moment-based measures from the recovered stock-level RND using portfolio analysis. Stocks are sorted into quintile portfolios based on the magnitude of the estimated moment-based measures. The first quintile corresponds

[13] We might need a follow-up regression to show the robust and stability of implied skewness estimation using the CGP-based RND indeed help to predict the skewness of stock returns. That is, we may need the evidence to show such skewness estimations from the CGP-based RND are more close to the realized ones.

to the stocks with the moment-based measures in the lowest magnitude, and the fifth quintile corresponds to the stocks with the moment-based measures in the highest magnitude. We then form long-short portfolios by buying the stocks in the fifth quintile and selling the stocks in the first quintile. The quintile portfolios are formed on the end date of each month, from Jan 1997 to Dec 2022, and the performance is evaluated using the realized stock returns by holding the portfolios untill the end of the next month.

Stock-level RND-based skewness. Figure 5 shows the monthly cumulative returns of the long-short $Skew^{\mathbb{Q}}$ portfolios by using different RND-recovery methods. We observe that prior to the financial crisis of 2007-2008, the performance of the CGP portfolio is outperformed by the IVGEV portfolio, ranking at the second place among all the recovery approaches considered. However, the performance of the CGP portfolio is notably improved after the financial crisis and outperforms all other portfolios significantly. This is because that the trading volume of stock options explodes after the financial crisis, which means that there is more informative information which can be learned directly from the stock option prices, and the nonparametric data-driven approach of no-arbitrage constrained GP is most suitable in this scenario.

[Insert Figure 5 about here]

The Sharpe ratio of the CGP portfolio is the highest at 0.20, over the period from Jan 1997 to Dec 2022, and is nearly as double as the IVGEV portfolio which is the second best. The constant and the SVI-paramterized extrapolated RND generate Sharpe ratios that are not significantly different from zero, with 0.02 and 0.05 respectively. We also calculate the t-statistic for each portfolio. Only the CGP-based long-short $Skew^{\mathbb{Q}}$ portfolio gives a notable t-statistic at 2.9, which implies strong statistically nonzero monthly portfolio returns.

Stock-level RND-based volatility innovation. In the literature of empirical asset pricing with option-implied predictors, according to Dennis and Mayhew (2002), it is acknowledged that it is the innovation of implied volatility, instead of the implied volatility itself, that relates to stock returns. Dennis et al. (2006) relies on the stock-level standardized implied volatilities from stock options with one-month expiration, in the Berkeley Option Data Base, to calculate the implied volatility innovation in a daily frequency. We use the RND-based estimation of the implied volatility innovation as explained in Section 5.2 in monthly frequency.

Figure 6 shows the monthly cumulative returns of the long-short $\Delta Vol^{\mathbb{Q}}$ portfolios by using different RND-recovery methods. Similar to the case of $Skew^{\mathbb{Q}}$, the CGP-based long-short portfolio notably outperforms other portfolios after the financial crisis of 2007-2008. Prior to the crisis, the CGP-based portfolio performs closely to the IVSVI portfolio and better than IVSPL and IVGEV portfolios, but is slightly outperformed by the BKM portfolio.

The Sharpe ratio of the CGP portfolio is the highest at 0.04, while all other competing portfolios generate zero or negative Sharpe ratios, which highlights the additional economic information unlocked by the estimation of stock-level implied volatilities from CGP-based RND. Though the monthly returns of the CGP-based $\Delta Vol^{\mathbb{Q}}$ long-short portfolio are not statistically different from zero, with t-statistic of 0.65, it still presents the highest t-statistic for monthly returns among all considered portfolios.

[Insert Figure 6 about here]

Stock-level RND-based asymmetry of variance. We proceed to investigate whether the inclusion of the information from the out-of-the-money call and out-of-the-money put stock options in the stock-level RND, by using the CGP-based recovery, provides additional economic value. Bollerslev et al. (2019) separates the so-called up and down semi-variances from positive and negative high-frequency price increments. We follow a similar intuition, but instead of using observed stock returns, we attempt to construct the semi-variances with the option-implied return distribution. The $UPM^{\mathbb{Q}}$ and $LPM^{\mathbb{Q}}$ in Section 5.2 correspond to up and down implied variances for stock returns. We define the difference $AVAR^{\mathbb{Q}} := UPM^{\mathbb{Q}} - LPM^{\mathbb{Q}}$ as a moment-based measure to capture the asymmetry in stock return variance.

[Insert Figure 7 about here]

Figure 7 shows the monthly cumulative returns of the long-short $AVAR^{\mathbb{Q}}$ portfolios by using different RND-recovery methods. Again, we observe that the performance of the CGP-based portfolio is closely tracked by the IVSPL and the IVSVI portfolios before the financial crisis of 2007-2008. The IVGEV-based $AVAR^{\mathbb{Q}}$ long-short portfolio performs the best in this pre-crisis period. But the performance of the CGP-based portfolio is notably improved after the crisis. The CGP-based long-short portfolio outperforms others at all times in the post-crisis period.

The Sharpe ratio of the CGP-based portfolio is again the highest at 0.12. We calculate the t-statistics of the monthly returns for all considered portfolios. Our construction of the stock-level asymmetry of variance with the CGP-based recovery generates monthly long-short portfolio returns that are statistically nonzero with t-statistic of 1.80, which is also the highest among the tested portfolios. The t-statistics of other $AVAR^{\mathbb{Q}}$ long-short portfolios do not imply significantly nonzero monthly returns.

[Insert Table 19 about here]

We compare the performance of the moment-based long-short portfolios formed with different RND-recovery techniques. Table 19 presents the Sharpe ratios and t-statistics of the monthly portfolio returns. For the three moment-based stock return predictors we consider, the CGP-based portfolios provide the highest Sharpe ratios and t-statistics, which highlights the additional economic information unlocked by the no-arbitrage constrained GP.

5.4.2 Cross-Sectional Regressions

We proceed to evaluate the pricing effects of the moment-based return predictors using the Fama and MacBeth (1973) cross-sectional regressions. The cross-sectional regressions serve as further scrutiny of the economic significance provided by the moment-based predictors in addition to the portfolio analysis in Section 5.4, in which a set of control variables can be included.

Univariate setting. We first consider a univariate case of the cross-sectional regressions where only the estimated moment-based predictor $\mathcal{M}_t^{mom}$ from the stock-level RND at time t is included, with the regression model as follows:

$$ret_{i,t+1mo} = \beta_t^0 + \beta_t^1 \mathcal{M}_{i,t}^{mom} + \epsilon_{i,t}, \tag{31}$$

where $ret_{i,t+1mo}$ is the return for stock i from t to one month after t, $\mathcal{M}_{i,t}^{mom} = \left\{Skew_{i,t}^{\mathbb{Q}}, \Delta Vol_{i,t}^{\mathbb{Q}}, AVAR_{i,t}^{\mathbb{Q}}\right\}$. This univariate setting directly investigates the economic information of the moment-based predictor estimation itself. The regression result should give implications for the informativeness of the RND-recovery methods. Table 20 shows the results of the univariate cross-sectional regressions. We report the coefficients of each moment-based predictor and the intercepts, with their corresponding t-statistics. We observe that for all the moment-based predictors $Skew^{\mathbb{Q}}, \Delta Vol^{\mathbb{Q}}, AVAR^{\mathbb{Q}}$ we consider, the estimation from the CGP-based RND is always statistically significant. The t-statistics of the estimation form the CGP-based RND are 1.97, 2.01 and 2.50 respectively for the predictor of $Skew^{\mathbb{Q}}, \Delta Vol$ and $AVAR^{\mathbb{Q}}$. Other considered RND-recovery approaches can not provide statistically significant moment-based predictor estimations in cross-section, except that the IVSPL method gives the estimations of $AVAR^{\mathbb{Q}}$ with t-statistic of 1.72 for cross-sectional stock returns.

[Insert Table 20 about here]

Multivariate setting with control variables. We then consider a multivariate case of the cross-sectional regressions where a set of control variables are included to test the robustness of the moment-based predictor estimation using each RND-recovery method. The multivariate cross-ectional regression model is specifed as follows:

$$ret_{i,t+1mo} = \beta_t^0 + \beta_t^1 \mathcal{M}_{i,t}^{mom} + \boldsymbol{Z}_{i,t}^{\top} \boldsymbol{\beta}_t^Z + \epsilon_{i,t}, \tag{32}$$

where $\boldsymbol{Z}_{i,t}^{\top}$ is a set of contemporaneous control variables for stock i as we defined in the previous sections. Table 21 shows the results of the multivariate cross-sectional regressions. We report the coefficients and the t-statistics of the moment-based predictors and the intercepts in multivariate models. In this test, we only consider the predictor estimations from the CGP-based RND. From column (1) to column (4) for each moment-based predictor in Table 21, firm-characteristics, BKM's option-characteristics and stock's idiosyncratic risk factors are progressively added into the cross-sectional regression model. We observe that in most cases, including control variables does not affect the statistical significance of the moment-based predictor estimated from the CGP-based RND. The only exception is column (4) for $Skew^{\mathbb{Q}}$ where all control variables are included jointly. This indicates that estimating moment-based predictors using the CGP-based RND unlocks significant information for cross-sectional stock returns, which is orthogonal to the stock information from historical data and the BKM's model-free risk-neutral moments from option prices.

[Insert Table 21 about here]

5.5 Capturing Earning Announcement Information

Earning announcement of a firm is considered as an information-intensive event for its stock returns. Empirical studies such as Diavatopoulos et al. (2012), Atilgan (2014), Dubinsky et al. (2018) and Lei et al. (2020) document strong return predictability during the earning announcement dates (EAD) due to informed trading of individual stock options. In this section, we inspect whether such predictive information for stock returns can be captured by the stock-level model-free RND. We employ $AVAR^{\mathbb{Q}}$ as a predictor to capture the option-based predictive information around EAD. The use of $AVAR^{\mathbb{Q}}$ shares the rationale of the stock's IV spread as in Atilgan (2014). If investors have private information about future stock price increases (decreases), then call (put) options will be increasingly demanded, which will drive the IVs from call (put) options up. Since the CGP-based RND has the advantage of sourcing information from both the call and the put options for a stock, the $AVAR^{\mathbb{Q}}$ is able to encode the spread of investor's positive and negative sentiment by construction.

5.5.1 OLS With Selected Stocks

We investigate whether there is anticipated information from options for the stock return r_i^{EAD} around EAD for stock i with the OLS model as follows:

$$r_{i,t}^{EAD} = \alpha_0 + \alpha_1 \Delta_{-30}^{t-d} AVAR_i^{\mathbb{Q}} + \epsilon_t, \tag{33}$$

where $r_{i,t}^{EAD}$ is the cumulative stock return from time $t-d+1$ to time $t+1$ which is around the earning announcement date at t of stock i, and $\Delta_{-30}^{t-d} AVAR_i^{\mathbb{Q}}$ denotes the innovation of the stock's asymmetry of implied variance from time $t-30$ to time $t-d$. d is a constant that decides the time window to calculate the innovation of the asymmetry of implied variance. We consider $d = 20$ and $d = 5$. When $d = 20$, we calculate the innovation of the asymmetry of implied variance $\Delta_{-30}^{t-20} AVAR_i^{\mathbb{Q}} = AVAR_{i,t-20}^{\mathbb{Q}} - AVAR_{i,t-30}^{\mathbb{Q}}$, which uses more remote implied information in stock options from the earning announcement time t. When $d = 5$ we calculate $\Delta_{-30}^{t-5} AVAR_i^{\mathbb{Q}} = AVAR_{i,t-5}^{\mathbb{Q}} - AVAR_{i,t-30}^{\mathbb{Q}}$, which instead uses implied information more near to the earning announcement time t. The cumulative stock return $r_{i,t}^{EAD}$ is measured immediately after the formation of $\Delta_{-30}^{t-d} AVAR_i^{\mathbb{Q}}$ to avoid the overlapping of information.

[Insert Table 22 about here]

We first test whether the RND-based asymmetry of variance estimations using the no-arbitrage constrained GP can predict the EAD return of on a selected subsample of stocks. The stocks are chosen randomly, with a focus on diversified industries including technology, healthcare, financial and industrial. For each industry group, we choose 4 stocks. There are 16 different stocks in total in the subsample. The time-series OLS model in (33) is estimated for each of the 16 individual stocks.

Table 22 shows the result for stock-level time-series OLS regressions. We report the coefficient and the associated t-statistic for the innovation of the asymmetry of implied variance estimations for each stock-level regression. The left column and the right column under the section of each RND-recovery method correspond to the case of $d = 20$ and $d = 5$. For the CGP-based RND recovery, we observe the most individual OLS regressions where $\Delta_{-30}^{t-d} AVAR_i^{\mathbb{Q}}$ is statistically significant at least at the 10% significance level. When $d = 20$, there are 16 out of 16 individual regressions that are statistically significant, and when $d = 5$, there are 9 out of 16 individual regressions that are statistically significant. This percentage of statistically significant regressions from the CGP-based stock-level RND is much higher than the percentages from other RND-recovery methods we consider, which highlights the superior capability of the CGP-based RND-recovery in encoding predictive information for stock returns, particularly when there is important information to release to the market.

5.5.2 OLS With Panel of Stocks

We continue to investigate whether the predictive information for stock returns around the EAD captured by the innovation of the asymmetry of implied variance still exists in a wider panel of stocks. We estimate a panel OLS model with a cross-section of stocks from stock i to stock N, spanning the observed earning announcement time $t = \{t_0, t_1, \cdots, t_q\}$ for each stock, where t_0 is the first EAD and t_q is the last EAD within the considered time period. The panel OLS model is as follows:

$$r_{it}^{EAD} = \alpha_0 + \alpha_1 \Delta_{-30}^{t-d} AVAR_{it} + \epsilon_{it}. \tag{34}$$

[Insert Table 23 about here]

The same cross-section of stocks as in previous sections is used to estimate model (34). Table 23 shows the result for the panel OLS regressions. Again, we report the coefficient and the associated t-statistic for the innovation of the asymmetry of implied variance estimations across the cross-section of stocks. In Panel A, we present the result of a full-time-period panel OLS with the time period covering from Jan 1996 to Dec 2022. We observe that the estimated $\Delta_{-30}^{t-d} AVAR_{it}^{\mathbb{Q}}$ from the CGP-based RND is a statistically significant predictor for EAD stock returns, for both the case of $d = 20$ and $d = 5$. At the same time, the estimated $\Delta_{-30}^{t-d} AVAR_{it}^{\mathbb{Q}}$ from the IVGEV-based RND exhibits the closest performance as it is also statistically significant to predict EAD stock returns.

Other considered RND-recovery methods fail to capture this predictive information with the RND-based $\Delta_{-30}^{t-d}AVAR_{it}^{\mathbb{Q}}$ estimations. We further split the entire time period into a pre-crisis sub-period covering the time from Jan 1996 to Jan 2007 and a post-crisis sub-period covering the time from Jan 2007 to Dec 2022, to examine whether the informativeness of the IVGEV-based $\Delta_{-30}^{t-d}AVAR_{it}^{\mathbb{Q}}$ and the CGP-based $\Delta_{-30}^{t-d}AVAR_{it}^{\mathbb{Q}}$ diverges in the two sub-periods. The Panel B shows the panel OLS regression result for the pre-crisis sub-period. Only the IVGEV-based $\Delta_{-30}^{t-d}AVAR_{it}^{\mathbb{Q}}$ captures the predictive information for the EAD stock returns when $d = -5$. Panel C shows the panel OLS regression result for the post-crisis period, where we observe that the informativeness of the CGP-based $\Delta_{-30}^{t-d}AVAR_{it}^{\mathbb{Q}}$ outperforms that of the IVGEV-based $\Delta_{-30}^{t-d}AVAR_{it}^{\mathbb{Q}}$. Since a statistically significant $\Delta_{-30}^{t-d}AVAR_{it}^{\mathbb{Q}}$ to predict the EAD stock returns is observed for $d = -20$ and $d = -5$ when using the CGP-based RND, while a strong predictor for the EAD stock returns is only observed for $d = 20$ when using the IVGEV-based RND.

6 Conclusion

Using the empirical context of risk assessment and return prediction as the proving ground, we perform a comprehensive analysis of the capability of our proposed consGP model, as an economics-aware machine, in learning the predictive information from stock options. Learning from stock options is challenging due to the limitations of small and noisy observations. Our findings demonstrate that when embedded with domain knowledge, data-driven machine learning method can help to decode economic information when the financial data is small-sized and noisy, in which case structural models would typically fail as model parameters can not be accurately calibrated.

Our paper also provides a general framework to estimate implied risk metrics. Based on the nonparametric no-arbitrage option pricing relation learned with the consGP model, we further recover a hybrid model-free RND implied in options that incorporates the information from calls and puts jointly. The recovered RND allows investors to estimate a wider range of risk metrics for returns than in the literature. Enforced with the no-arbitrage constraints from the domain knowledge of option pricing, the consGP-based RND shows lower density loss while ensuring the regularity conditions of probability density, which is essential for reliable risk metric estimations.

We evaluate our method with a series of risk forecasting and asset pricing tests. The out-of-sample predictive power of the tail risk measures estimated from the consGP-based RND is in average 3.92 times higher than those estimated from traditional RND-recovery approaches. When using the recovered stock-level RND to estimate moment-based factors for stock returns and form long-short portfolios, our consGP-based method provides 4.30 times higher Sharpe ratio in average than others. The enhanced predictive information and the tangible economic benefits unlocked by the economics-aware consGP model support that both fitting to observations and adhering to fundamental economic principles are necessary to learn from financial data under small-sample and noisy conditions.

The overall success of enforcing option pricing domain knowledge to GPs for nonparametrically learning the no-arbitrage pricing relation and deriving the model-free RND brings promise for option pricing when there is no luxury of noiseless option observations to estimate parametric models which, in addition, also suffer from model-specification errors across the time-series and the cross-section of assets. Potential further research includes extending the machine learning model to high-dimensional so that the prior knowledge with respect to different features can be addressed and multivariate option characteristics can be incorporated.

Appendix A Heston (1993) Model Implementation

The characteristic function. To derive the theoretical RND $f^{\mathbb{Q}}_{Heston}(S_T|S_t)$ from the Heston (1993) model, the first step is to obtain the characteristics function for the log-price $\log(S_T)$. Throughout the paper, we use the modified characteristic function of $\log(S_T)$ for Heston (1993) model as proposed in Schoutens et al. (2003) (henceforth, WEJ characteristic function). This characteristic function is equivalent to the original one as in Heston (1993) but is easier to evaluate. The WEJ characteristic function is defined as follows:

$$
\begin{aligned}
\phi\left(u; S_t, v_0, t, T\right) =& \mathbb{E}^{\mathbb{Q}}\left[e^{iu\log(S_T)}\middle| S_t, v_0\right] && (35)\\
=& \exp\left(iu\left(\log(S_t) + r\tau\right)\right) && (36)\\
&\times \exp\left(\theta\kappa\sigma^{-2}\left(\left(\kappa - \rho\sigma ui - d\right)\tau - 2\log\left(\left(1 - ge^{-d\tau}\right)\Big/\left(1-g\right)\right)\right)\right) && (37)\\
&\times \exp\left(v_0\sigma^{-2}\left(\kappa - \rho\sigma iu - d\right)\left(1 - e^{-d\tau}\right)\Big/\left(1 - ge^{-d\tau}\right)\right), && (38)
\end{aligned}
$$

where

$$
\begin{aligned}
d &= \left(\left(\rho\sigma ui - \kappa\right)^2 - \sigma^2\left(-iu - u^2\right)\right)^{1/2} && (39)\\
g &= \left(\kappa - \rho\sigma ui - d\right)\Big/\left(\kappa - \rho\sigma ui + d\right) && (40)\\
\tau &= T - t. && (41)
\end{aligned}
$$

Inverse Fourier transform. Given the characteristic function of $\log(S_T)$, the density function can be then calculated by taking the inverse Fourier transform:

$$
f_X(x) = \frac{1}{2\pi}\int_{\mathbb{R}} e^{-iux\phi_X(u)du}, \tag{42}
$$

where $x := \log(S_T)$. We rely on Gil-Pelaez (1951) formula to calculate the inversion of the Fourier transform more efficiently. The cumulative density function can be obtained using the following Fourier inversion:

$$
\begin{aligned}
F_X(x) = \mathbb{P}(X < x) &= \int_{-\infty}^{x} f_X(t)dt && (43)\\
&= \frac{1}{2} - \frac{1}{2\pi}\int_{\mathbb{R}} \frac{e^{-iux}\phi_X(u)}{iu}du, && (44)
\end{aligned}
$$

which can be written as follows according to Gil-Pelaez (1951):

$$
\begin{aligned}
F_X(x) &= \frac{1}{2} - \frac{1}{2\pi}\int_{\mathbb{R}} e^{-iux}\phi_X(u)\cdot\frac{1}{iu}du && (45)\\
&= \frac{1}{2} - \left(\frac{1}{2\pi}\int_{-\infty}^{0} e^{-iux}\phi_X(u)\frac{1}{iu}du + \frac{1}{2\pi}\int_0^{\infty} e^{-iux}\phi_X(u)\frac{1}{iu}du\right) && (46)\\
&= \frac{1}{2} - \frac{1}{2\pi}\int_0^{\infty}\left(-e^{iux}\phi_X(-u) + e^{-iux}\phi_X(u)\right)\frac{1}{iu}du && (47)\\
&= \frac{1}{2} - \frac{1}{\pi}\int_0^{\infty}\mathcal{R}e\left[\frac{e^{-iux}\phi_X(u)}{iu}\right]du. && (48)
\end{aligned}
$$

Taking the derivative leads to the Gil-Pelaez (1951) representation of the density function:

$$
f_X(x) = \frac{1}{\pi}\int_0^{\infty}\mathcal{R}e\left[e^{-iux}\phi_X(u)\right]du. \tag{49}
$$

Option pricing via Fourier inversion. The pricing formula for an European call option with strike K and maturity T can be written as:

$$
C(S_t, r, K, t, T) = S_t\tilde{\mathbb{Q}}\left(S_T > K\right) - e^{-r(T-t)}K\mathbb{Q}\left(S_T > K\right), \tag{50}
$$

where $\tilde{\mathbb{Q}}$ is the probability under the stock numéraire and $\mathbb{Q}$ is the probability under the money market numéraire. The probability $\tilde{\mathbb{Q}}$ and $\mathbb{Q}$ can be expressed as follows in terms of the Gil-Pelaez (1951) formula:

$$
\begin{aligned}
\tilde{\mathbb{Q}}(S_T > K) &= \frac{1}{2} + \frac{1}{\pi}\int_0^{\infty}\mathcal{R}e\left[\frac{e^{-iuk}\tilde{\phi}_X(u)}{iu}\right]du && (51)\\
&= \frac{1}{2} + \frac{1}{\pi}\int_0^{\infty}\mathcal{R}e\left[\frac{e^{-iuk}\phi_X(u-i)}{iu\phi_X(-i)}\right]du, && (52)
\end{aligned}
$$

and

$$
\begin{aligned}
\mathbb{Q}(S_T > K) &= 1 - \mathbb{Q}(S_T < K) = 1 - \mathbb{Q}(X_T < k) && (53)\\
&= \frac{1}{2} + \frac{1}{\pi}\int_0^{\infty}\mathcal{R}e\left[\frac{e^{-iuk}\phi_X(u)}{iu}\right]du, && (54)
\end{aligned}
$$

where $k := \log(S_T/S_t)$ and

$$
\tilde{\phi}_X(u) := \mathbb{E}^{\tilde{\mathbb{Q}}}\left[e^{iuX_T}\right] = \mathbb{E}^{\mathbb{Q}}\left[\frac{d\tilde{\mathbb{Q}}}{d\mathbb{Q}}e^{iuX_T}\right] \tag{55}
$$

$$= \mathbb{E}^{\mathbb{Q}}\left[\frac{S_t e^{X_T}}{S_t \mathbb{E}^{\mathbb{Q}}[e^{X_T}]} e^{iuX_T}\right] \tag{56}$$

$$= \mathbb{E}^{\mathbb{Q}}\left[\frac{e^{(iu+1)X_T}}{\phi_X(-i)}\right] \tag{57}$$

$$= \frac{\phi_X(u-i)}{\phi_X(-i)}. \tag{58}$$

The corresponding European put option prices can then be calculated with the well-known put-call parity given the call prices obtained using the option pricing via Fourier inversion with the Gil-Pelaez (1951) formula.

Appendix B Stock and Option Data Selection Criteria

We select stocks that are included in the S&P 500 index on the first trading date of each year from 1996 to 2022. The selected stocks are held fixed throughout each year. We then discard stocks whose options are inactively traded, that is we only include stocks for which there are options traded for more than 200 days in a year.

For each stock on each day and for each time-to-maturity, we include stock options with strike prices less than $2 \times S_t e^{r\tau}$. This is because that we estimate the RND over the range $K \in [0, 2 \times S_t e^{r\tau}]$ for all stocks. We exclude stock options with zero trading volume and open interest at the same time. We exclude stock options with zero best-bid price. We do not estimate the RND when the number of call and put options is less than 5.

Appendix C Goodness-of-fit of Quantile Regression

We follow Koenker and Machado (1999) to calculate R^1 which is a goodness-of-fit measure for quantile regression analogous to the conventional R^2 for ordinary least squares regression. Consider a linear conditional quantile model:

$$\mathcal{Q}_{y_i}(\tau|\boldsymbol{x}) = \boldsymbol{x}_{i1}^{\top}\boldsymbol{\beta}_1(\tau) + \boldsymbol{x}_{i2}^{\top}\boldsymbol{\beta}_2(\tau), \tag{59}$$

where $\boldsymbol{x}_{i,1}^{\top}$ is a vector of ones and $\boldsymbol{x}_{i2}^{\top}$ denotes all other regressors. Let $\hat{\boldsymbol{\beta}}(\tau)$ denote the minimizer of the full quantile model

$$\hat{V}(\tau) = \underset{b \in \mathbb{R}^p}{\operatorname{argmin}} \sum \rho_\tau \left(y_i - \boldsymbol{x}_i^{\top}\boldsymbol{b}\right), \tag{60}$$

and let $\tilde{\boldsymbol{\beta}}(\tau) = \left[\tilde{\boldsymbol{\beta}}_1(\tau)^\top, \mathbf{0}^\top\right]$ denotes the minimizer of the corresponding restricted quantile model

$$\tilde{V}(\tau) = \underset{\boldsymbol{b}_1 \in \mathbb{R}^{p-q}}{\operatorname{argmin}} \sum \rho_\tau \left(y_i - \boldsymbol{x}_{1i}^\top \boldsymbol{b}_1 \right). \tag{61}$$

$\rho_\tau(\cdot)$ is the quantile-weighted absolute value of errors according to Koenker and Bassett (1978). That is, $\hat{\boldsymbol{\beta}}(\tau)$ and $\tilde{\boldsymbol{\beta}}(\tau)$ denote the estimated coefficients of the full and the restricted quantile regression model. $\hat{V}(\tau)$ and $\tilde{V}(\tau)$ denote the quantile-weighted absolute value of errors calculated from the full and the restricted model respectively. The goodness-of-fit criterion is then defined as

$$R^1(\tau) = 1 - \hat{V}(\tau)/\tilde{V}(\tau). \tag{62}$$

References

Atilgan, Y., 2014. Volatility spreads and earnings announcement returns. Journal of Banking & Finance 38, 205–215. URL: http://dx.doi.org/10.1016/j.jbankfin.2013.10.007, doi:10.1016/j.jbankfin.2013.10.007.

Aït-Sahalia, Y., Duarte, J., 2003. Nonparametric option pricing under shape restrictions. Journal of Econometrics 116, 9–47.

Aït-Sahalia, Y., Lo, A.W., 2000. Nonparametric risk management and implied risk aversion. Journal of Econometrics 94, 9–51. URL: http://dx.doi.org/10.1016/S0304-4076(99)00016-0, doi:10.1016/s0304-4076(99)00016-0.

Bakshi, G., Kapadia, N., Madan, D., 2003. Stock return characteristics, skew laws, and the differential pricing of individual equity options. Review of Financial Studies 16, 101–143. URL: http://dx.doi.org/10.1093/rfs/16.1.0101, doi:10.1093/rfs/16.1.0101.

Barone-Adesi, G., 2016. Var and cvar implied in option prices. Journal of Risk and Financial Management 9, 2. URL: http://dx.doi.org/10.3390/jrfm9010002, doi:10.3390/jrfm9010002.

Berkowitz, J., 2001. Testing density forecasts, with applications to risk management. Journal of Business & Economic Statistics 19, 465–474. URL: http://dx.doi.org/10.1198/07350010152596718, doi:10.1198/07350010152596718.

Bollerslev, T., Li, S.Z., Zhao, B., 2019. Good volatility, bad volatility, and the cross section of stock returns. Journal of Financial and Quantitative Analysis 55, 751–781. URL: http://dx.doi.org/10.1017/S0022109019000097, doi:10.1017/s0022109019000097.

Breeden, D.T., Litzenberger, R.H., 1978. Prices of state-contingent claims implicit in option prices. Journal of business , 621–651.

Chataigner, M., Cousin, A., Crépey, S., Dixon, M., Gueye, D., 2021. Short communication: Beyond surrogate modeling: Learning the local volatility via shape constraints. SIAM Journal on Financial Mathematics 12, SC58–SC69. URL: http://dx.doi.org/10.1137/20M1381538, doi:10.1137/20m1381538.

Chen, H., Cheng, Y., Liu, Y., Tang, K., 2023. Teaching economics to the machine, in: Proceedings of the China International Conference in Finance (CICF2023), Sloan School of Management, Massachusetts Institute of Technology. p. Session 32 Fintech and Machine Learning.

Chen, J., Hong, H., Stein, J.C., 2001. Forecasting crashes: trading volume, past returns, and conditional skewness in stock prices. Journal of Financial Economics 61, 345–381. URL: http://dx.doi.org/10.1016/S0304-405X(01)00066-6, doi:10.1016/s0304-405x(01)00066-6.

Christoffersen, P.F., 1998. Evaluating interval forecasts. International economic review , 841–862.

Dennis, P., Mayhew, S., 2002. Risk-neutral skewness: Evidence from stock options. The Journal of Financial and Quantitative Analysis 37, 471. URL: http://dx.doi.org/10.2307/3594989, doi:10.2307/3594989.

Dennis, P., Mayhew, S., Stivers, C., 2006. Stock returns, implied volatility innovations, and the asymmetric volatility phenomenon. Journal of Financial and Quantitative Analysis 41, 381–406. URL: http://dx.doi.org/10.1017/S0022109000002118, doi:10.1017/s0022109000002118.

Diavatopoulos, D., Doran, J.S., Fodor, A., Peterson, D.R., 2012. The information content of implied skewness and kurtosis changes prior to earnings announcements for stock and option returns. Journal of Banking & Finance 36, 786–802. URL: http://dx.doi.org/10.1016/j.jbankfin.2011.09.012, doi:10.1016/j.jbankfin.2011.09.012.

Diebold, F.X., Gunther, T.A., Tay, A., 1997. Evaluating density forecasts.

Dubinsky, A., Johannes, M., Kaeck, A., Seeger, N.J., 2018. Option pricing of earnings announcement risks. The Review of Financial Studies 32, 646–687. URL: http://dx.doi.org/10.1093/rfs/hhy060, doi:10.1093/rfs/hhy060.

Engle, R.F., Manganelli, S., 2004. Caviar: Conditional autoregressive value at risk by regression quantiles. Journal of Business & Economic Statistics 22, 367–381. URL: http://dx.doi.org/10.1198/073500104000000370, doi:10.1198/073500104000000370.

Fama, E.F., MacBeth, J.D., 1973. Risk, return, and equilibrium: Empirical tests. Journal of political economy 81, 607–636.

Figlewski, S., 2010. Estimating the Implied Risk-Neutral Density for the US Market Portfolio*. Oxford University Press. p. 323–353. URL: http://dx.doi.org/10.1093/acprof:oso/9780199549498.003.0015, doi:10.1093/acprof:oso/9780199549498.003.0015.

Gaglianone, W.P., Lima, L.R., Linton, O., Smith, D.R., 2011. Evaluating value-at-risk models via quantile regression. Journal of Business & Economic Statistics 29, 150–160.

Gatheral, J., 2004. A parsimonious arbitrage-free implied volatility parameterization with application to the valuation of volatility derivatives. Presentation at Global Derivatives & Risk Management, Madrid , 0.

Gil-Pelaez, J., 1951. Note on the inversion theorem. Biometrika 38, 481–482.

Goyal, A., Saretto, A., 2009. Cross-section of option returns and volatility. Journal of Financial Economics 94, 310–326.

Hamidieh, K., 2017. Estimating the tail shape parameter from option prices. The Journal of Risk 19, 85–110. URL: http://dx.doi.org/10.21314/JOR.2017.366, doi:10.21314/jor.2017.366.

Härdle, W., Hlávka, Z., 2009. Dynamics of state price densities. Journal of Econometrics 150, 1–15. URL: http://dx.doi.org/10.1016/j.jeconom.2009.01.005, doi:10.1016/j.jeconom.2009.01.005.

Heston, S.L., 1993. A closed-form solution for options with stochastic volatility with applications to bond and currency options. The review of financial studies 6, 327–343.

Huggenberger, M., Zhang, C., Zhou, T., 2018. Forward-looking tail risk measures. Available at SSRN 2909808 .

Hutton, A.P., Marcus, A.J., Tehranian, H., 2009. Opaque financial reports, r2, and crash risk. Journal of financial Economics 94, 67–86.

Koenker, R., Bassett, G., 1978. Regression quantiles. Econometrica 46, 33. URL: http://dx.doi.org/10.2307/1913643, doi:10.2307/1913643.

Koenker, R., Machado, J.A.F., 1999. Goodness of fit and related inference processes for quantile regression. Journal of the American Statistical Association 94, 1296–1310. URL: http://dx.doi.org/10.1080/01621459.1999.10473882, doi:10.1080/01621459.1999.10473882.

Kullback, S., Leibler, R.A., 1951. On information and sufficiency. The annals of mathematical statistics 22, 79–86.

Lei, Q., Wang, X.W., Yan, Z., 2020. Volatility spread and stock market response to earnings announcements. Journal of Banking & Finance 119, 105126. URL: http://dx.doi.org/10.1016/j.jbankfin.2017.04.002, doi:10.1016/j.jbankfin.2017.04.002.

López-Lopera, A.F., Bachoc, F., Durrande, N., Roustant, O., 2018. Finite-dimensional gaussian approximation with linear inequality constraints. SIAM/ASA Journal on Uncertainty Quantification 6, 1224–1255.

Maatouk, H., Bay, X., 2017. Gaussian process emulators for computer experiments with inequality constraints. Mathematical Geosciences 49, 557–582.

Machado, J.A., Silva, J.S., 2019. Quantiles via moments. Journal of econometrics 213, 145–173.

Markowitz, H., 1952. Portfolio selection. The Journal of Finance 7, 77. URL: http://dx.doi.org/10.2307/2975974, doi:10.2307/2975974.

Newey, W.K., West, K.D., 1987. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. Econometrica 55, 703–708. URL: http://www.jstor.org/stable/1913610.

Rubner, Y., Tomasi, C., Guibas, L.J., 2000. The earth mover's distance as a metric for image retrieval. International journal of computer vision 40, 99–121.

Schoutens, W., Simons, E., Tistaert, J., 2003. A perfect calibration! now what? The best of Wilmott , 281.

Shimko, D., 1993. Bounds of probability. Risk 6, 33–37.

Thomas, M., Joy, A.T., 2006. Elements of information theory. Wiley-Interscience.

Tishby, N., Pereira, F.C., Bialek, W., 2000. The information bottleneck method. arXiv preprint physics/0004057 .

Tables

Table 1: **Summary statistics for stock options of S&P 500 constituent stocks.** This table describes the stock options that we use in our empirical analysis. We report the average number of traded option contracts, the average time-to-maturity (in calendar days), the minimum, median, maximum and mean option moneyness (moneyness is defined as K/S_t), the mean option trading volume, and the mean option open interest. To understand the degree of tail risk encoded in raw stock options, we follow Barone-Adesi (2016) to calculate $\alpha^{\mathbb{Q}}$ which is the implied cumulative probability at a given strike level. We calculate $\alpha^{\mathbb{Q}}$ at the second to the lowest strike price, which is reported as the minimum $\alpha^{\mathbb{Q}}$ in this table. A larger minimum $\alpha^{\mathbb{Q}}$ means a less-complete RND implied in raw options, and a smaller minimum $\alpha^{\mathbb{Q}}$ means a more-complete RND implied in raw options. The mean stock trading volume and the mean firm's capitalization are also reported. All statistics are calculated per stock-day at the end of each month from Jan 1996 to Dec 2022 by averaging across all time-to-maturity available.

	Small			Small-Medium			Medium-Large			Large		
	Call	Put	Overall	Call	Put	Overall	Call	Put	Overall	Call	Put	Overall
Average number of contracts	8.74	8.95	17.69	11.58	11.52	23.10	14.04	13.97	28.01	20.77	20.74	41.51
Average time-to-maturity	/	/	105.75	/	/	101.24	/	/	96.40	/	/	93.17
Minimum moneyness	0.71	0.73	0.66	0.73	0.74	0.68	0.72	0.72	0.66	0.68	0.69	0.63
Median moneyness	1.00	1.01	1.00	0.99	0.97	0.97	0.98	0.96	0.96	0.97	0.95	0.96
Maximum moneyness	1.29	1.33	1.38	1.24	1.24	1.29	1.23	1.23	1.28	1.26	1.26	1.31
Mean moneyness	1.00	1.01	1.01	0.99	0.97	0.98	0.98	0.96	0.97	0.98	0.96	0.96
Mean option volume	43.75	30.81	37.28	40.78	28.62	34.70	61.87	41.62	51.75	150.06	94.61	122.33
Mean option open interest	835.98	708.31	772.14	797.12	656.01	726.57	1076.87	866.86	971.87	2531.04	2036.63	2283.83
Minimum $\alpha^{\mathbb{Q}}$	/	23.30%	/	/	14.85%	/	/	10.37%	/	/	6.63%	/
Mean stock volume	/	/	14.56	/	/	14.67	/	/	15.02	/	/	15.81
Mean firm size	/	/	15.60	/	/	16.57	/	/	17.35	/	/	18.67

Table 2: **Stock-level predictive quantile regressions in time-series.** Time-series stock-level quantile regressions of weekly stock return over the lagged implied VaR ($\mathcal{M}_t = VaR_t^{\mathbb{Q}}$), stock-based VaR and other control variables. Figures in the upper panel represent the percentage of firm-regressions where the VaR metric is significant, across different significance levels. Volatility measures in BKM* and Idiosyn* are orthogonalized to the risk metric to address the multicollinearity issue.

$p-$value	Option	Stock	Option+Stock	Firm-chara	BKM*	Idiosyn*	Raw
$\leq 10\%$	82.67	66.79	76.53/25.99	77.26	77.44	71.84	45.13
$\leq 5\%$	79.42	61.37	71.30/18.59	73.10	72.38	66.43	36.64
$\leq 1\%$	70.58	51.81	62.82/9.57	65.34	60.65	58.30	23.83
No. of firms	554	554	554	551	551	551	551
Average firm-week obs.	532	532	532	531	531	531	531
Average R^1	4.99%	2.59%	5.62%	7.04%	8.89%	10.14%	10.15%

Table 3: **Stock-level predictive quantile regressions in time-series.** $\mathcal{M}_t = ES_t^{\mathbb{Q}}$.

$p-$value	Option	Stock	Option+Stock	Firm-chara	BKM*	Idiosyn*	Raw
$\leq 10\%$	83.03	65.52	75.63/27.62	77.62	77.98	72.92	42.42
$\leq 5\%$	78.89	60.29	70.76/20.94	73.10	73.83	68.95	35.56
$\leq 1\%$	71.12	51.62	63.18/9.75	65.52	63.54	59.93	22.56
No. of firms	554	554	554	551	551	551	551
Average firm-week obs.	532	532	532	531	531	531	531
Average R^1	5.05%	2.53%	5.70%	7.10%	8.88%	10.17%	10.18%

Table 4: **Stock-level predictive quantile regressions in time-series.** $\mathcal{M}_t = LPM_t^{\mathbb{Q}}$.

$p-$value	Option	Stock	Option+Stock	Firm-chara	BKM*	Idiosyn*	Raw
$\leq 10\%$	81.77	60.29	77.44/27.08	79.06	76.53	71.30	37.00
$\leq 5\%$	75.99	54.51	72.56/18.77	73.29	71.66	65.88	29.96
$\leq 1\%$	68.23	44.40	64.98/9.21	65.16	61.01	57.58	17.87
No. of firms	554	554	554	551	551	551	551
Average firm-week obs.	532	532	532	531	531	531	531
Average R^1	4.79%	2.21%	5.45%	6.93%	8.79%	10.05%	10.07%

Table 5: **The Summary Statistics of the In-Sample and Out-Of-Sample R^1 Distribution for Implied Tail Risk Metrics Across Stocks.** The in-sample and out-of-sample R^1 of stock-level univariate quantile regressions across all stocks using different RND-recovery methods. There are 516 different stocks in this test. All figures are in percentage. The conclusion will not change with winsorization.

	Constant		GEV		SVI		CGP	
	In-Sample	Out-Of-Sample	In-Sample	Out-Of-Sample	In-Sample	Out-Of-Sample	In-Sample	Out-Of-Sample
Panel A: Implied VaR								
Min	**0.0**	**-38.05**	**0.0**	-36.15	**0.0**	-63.31	**0.0**	-58.86
25%	1.75	-0.5	1.69	-0.53	1.09	-0.95	**2.11**	**-0.15**
50%	4.55	2.16	4.27	1.56	2.94	0.91	**5.02**	**2.41**
75%	7.32	5.25	6.9	5.38	5.61	3.33	**7.81**	**5.71**
Max	25.59	40.19	26.77	37.09	23.86	36.02	**34.18**	**44.81**
Std	4.26	6.83	4.34	6.73	3.93	6.35	**4.71**	**7.38**
Mean	5.21	2.69	5.02	2.34	4.03	1.42	**5.78**	**3.14**
Panel B: Implied ES								
Min	**0.0**	-97.43	**0.0**	-533.8	**0.0**	-253.99	**0.0**	**-67.89**
25%	1.72	-1.61	0.94	-1.35	0.09	-1.62	**2.37**	**-0.72**
50%	4.39	0.99	3.21	0.68	0.32	-0.23	**5.07**	**2.14**
75%	7.07	4.44	5.6	3.71	0.77	0.41	**7.83**	**5.5**
Max	**37.04**	37.9	15.23	28.78	5.64	30.16	30.87	**43.12**
Std	4.32	10.43	3.1	**26.11**	0.84	12.9	**4.55**	8.3
Mean	5.01	0.56	3.71	-1.38	0.63	-1.87	**5.74**	**2.52**
Panel C: Implied LPM								
Min	**0.0**	-55.62	**0.0**	**-30.0**	**0.0**	-49.97	**0.0**	-56.55
25%	1.66	-0.79	1.51	-0.75	0.09	-1.27	**2.14**	**-0.22**
50%	4.0	2.1	3.77	1.86	0.34	-0.22	**4.47**	**2.31**
75%	6.78	5.05	6.87	4.8	0.89	0.36	**7.41**	**5.36**
Max	30.29	45.35	24.72	29.11	7.62	10.31	**40.75**	**47.54**
Std	4.57	**7.64**	4.31	6.17	0.91	4.18	**4.84**	7.08
Mean	5.06	2.43	4.84	2.04	0.65	-1.06	**5.46**	**2.91**

Table 6: **Fixed effects panel quantile regressions of Machado and Silva (2019): stock returns with the implied VaR from the CGP-based stock-level RND.** Panel quantile regression with fixed effects of weekly stock returns over the lagged implied VaR, stock-based VaR and other control variables. $\text{Ivol}^{\perp}$ and $\text{Idiovol}^{\perp}$ are orthogonalized variables to the implied VaR.

	Option	Stock	Option+Stock	Firm-chara	BKM*	Idiosyn*	Raw
$VaR^{\mathbb{Q}}$	-2.854***		-2.826***	-3.143***	-3.092***	-3.049***	-2.254***
	(-18.79)		(-29.33)	(-36.01)	(-36.59)	(-4.84)	(-14.63)
$VaR^{\mathbb{P}}$		-0.915***	-0.023				
		(-18.26)	(0.51)				
Beta_{mkt}				0.003***	0.004***	0.002	0.003***
				(2.61)	(3.00)	(0.22)	(2.87)
Coskew				0.012***	0.011***	0.010	0.010***
				(9.79)	(8.87)	(1.00)	(8.46)
Size				-0.006***	-0.004***	-0.003	-0.003***
				(-9.08)	(-6.04)	(-0.69)	(-5.60)
Ivol							-0.065***
							(-8.54)
$\text{Ivol}^{\perp}$					-0.059***	-0.062	
					(-7.54)	(-1.01)	
Iskew					0.000	0.000	0.000
					(-1.07)	(-0.13)	(-1.09)
Ikurt					0.000	0.000	0.000
					(0.94)	(0.12)	(1.00)
Idiovol							0.022***
							(3.74)
$\text{Idiovol}^{\perp}$						0.022	
						(0.75)	
Idioskew						0.000	0.000***
						(0.30)	(2.84)
Firm-Week Obs.	301,040	301,040	301,040	300,727	300,72	300,727	300,727

Table 7: **Fixed effects panel quantile regressions of Machado and Silva (2019): stock returns with the implied ES from the CGP-based stock-level RND.** Panel quantile regression with fixed effects of weekly stock returns over the lagged implied ES, stock-based ES and other control variables. $\text{Ivol}^{\perp}$ and $\text{Idiovol}^{\perp}$ are orthogonalized variables to the implied ES.

	Option	Stock	Option+Stock	Firm-chara	BKM*	Idiosyn*	Raw
$ES^{\mathbb{Q}}$	-2.315***		-2.234***	-2.483***	-2.448***	-2.421***	-1.579***
	(-35.88)		(-21.04)	(-29.50)	(-49.38)	(-15.68)	(-16.81)
$ES^{\mathbb{P}}$		-0.586***	-0.0556				
		(-16.41)	(-1.18)				
Beta_{mkt}				0.0020	0.0026***	0.0014	0.0029***
				(1.26)	(2.77)	(0.45)	(3.30)
Cskew				0.0118***	0.0106***	0.0095***	0.0096***
				(7.20)	(10.78)	(2.93)	(10.15)
Size				-0.0041***	-0.0027***	-0.0022	-0.0023***
				(-5.10)	(-5.32)	(-1.30)	(-4.66)
Ivol							-0.0736***
							(-11.09)
$\text{Ivol}^{\perp}$					-0.0634***	-0.0674***	
					(-9.87)	(-3.10)	
Iskew					-0.0002	-0.0002	-0.0002
					(-1.15)	(-0.34)	(-1.15)
Ikurt					0.0000	0.0000	0.0000
					(1.10)	(0.34)	(1.13)
Idiovol							0.0098**
							(2.12)
$\text{Idiovol}^{\perp}$						0.0277*	
						(1.70)	
Idioskew						0.0004	0.0005***
						(0.91)	(3.49)
Firm-Week Obs.	301,040	301,040	301,040	300,727	300,727	300,727	300,727

Table 8: **Fixed effects panel quantile regressions of Machado and Silva (2019): stock returns with the implied LPM from the CGP-based stock-level RND.** Panel quantile regression with fixed effects of weekly stock returns over the lagged implied LPM, stock-based LPM and other control variables. $\text{Ivol}^{\perp}$ and $\text{Idiovol}^{\perp}$ are orthogonalized variables to the implied LPM.

	Option	Stock	Option+Stock	Firm-chara	BKM*	Idiosyn*	Raw
$LPM^{\mathbb{Q}}$	-18.28***		-18.97***	-19.78***	-19.45***	-19.35***	-12.05***
	(-43.55)		(-36.40)	(-42.17)	(-42.59)	(-41.88)	(-14.83)
$LPM^{\mathbb{P}}$		-0.0033***	0.0004**				
		(-23.50)	(2.49)				
Beta_{mkt}				0.0026***	0.0033***	0.0026***	0.0036***
				(3.00)	(3.80)	(2.97)	(4.20)
Coskew				0.0122***	0.0108***	0.0103***	0.0103***
				(13.50)	(11.86)	(10.95)	(10.94)
Size				-0.0044***	-0.0029***	-0.0026***	-0.0027***
				(-9.89)	(-5.91)	(-5.26)	(-5.33)
Ivol							-0.0710***
$\text{Ivol}^{\perp}$					-0.0671***	-0.0698***	
					(-11.64)	(-11.25)	
Iskew					-0.0002	-0.0002	-0.0002
					(-1.16)	(-1.12)	(-1.11)
Ikurt					0.0000	0.0000	0.0000
					(1.19)	(1.17)	(1.15)
Idiovol							-0.0010
							(-0.22)
$\text{Idiovol}^{\perp}$						0.0120**	
						(2.55)	
Idioskew						0.0005***	0.0005***
						(3.40)	(3.63)
Firm-Week Obs.	301,040	301,040	301,040	300,727	300,727	300,727	300,727

Table 9: **Fama and MacBeth (1973) Cross-Sectional Regressions:** $NCSKEW_{i,t:t+6mo}$**.** Fama and MacBeth (1973) cross-sectional regressions of $\text{NCSKEW}_{t,t:t+6mo}$. t-statistics are adjusted according to Newey and West (1987) with 12 lags.

	Option	Firm-chara	BKM*	Idiosyn*	Raw
$VaR^{\mathbb{Q}}$	-13.90***	-14.27***	-12.00***	-10.51**	-8.089*
	(-5.90)	(-3.49)	(-3.04)	(-2.46)	(-1.69)
Beta_{mkt}		0.008	0.006	-0.001	0.009
		(0.18)	(0.14)	(-0.03)	(0.19)
Coskew		0.053	0.086	0.146	0.136
		(0.46)	(0.76)	(1.22)	(1.16)
Size		-0.002	-0.007	-0.005	0.003
		(-0.10)	(-0.38)	(-0.24)	(0.17)
Ivol					-0.302*
					(-1.97)
$\text{Ivol}^{\perp}$			-0.013	0.052	
			(-0.08)	(0.35)	
Iskew			-0.035***	-0.030**	-0.031**
			(-3.09)	(-2.47)	(-2.36)
Ikurt			0.001	0.001	0.000
			(0.47)	(0.31)	(0.06)
Idiovol					0.200
					(0.65)
$\text{Idiovol}^{\perp}$				-0.284	
				(-0.69)	
Idioskew				-0.003	-0.003
				(-0.27)	(-0.26)
α_0	0.174***	0.170	0.138	0.107	0.007
	(3.17)	(0.65)	(0.52)	(0.39)	(0.03)
Firm-Month Obs.	47,861	47,813	47,813	47,813	47,813
adj. R^2	1.10%	2.30%	2.40%	3.00%	3.10%

Table 10: **Fama and MacBeth (1973) Cross-Sectional Regressions:** $NCSKEW_{i,t:t+6mo}$**.** Fama and MacBeth (1973) cross-sectional regressions of $\text{NCSKEW}_{t,t:t+6mo}$. t-statistics are adjusted according to Newey and West (1987) with 12 lags.

	Option	Firm-chara	BKM*	Idiosyn*	Raw
$ES^{\mathbb{Q}}$	-10.5500***	-9.7500***	-8.1780***	-7.3130**	-3.4990
	(-5.57)	(-3.08)	(-2.64)	(-2.17)	(-1.07)
Beta_{mkt}		-0.0065	0.0011	-0.0038	-0.0002
		(-0.15)	(0.03)	(-0.08)	(-0.00)
Coskew		0.0497	0.0831	0.1490	0.1420
		(0.45)	(0.75)	(1.27)	(1.22)
Size		0.0018	-0.0046	-0.0033	0.0055
		(0.10)	(-0.24)	(-0.17)	(0.27)
Ivol					-0.3920**
					(-2.45)
$\text{Ivol}^{\perp}$			-0.1090	-0.0214	
			(-0.62)	(-0.13)	
Iskew			-0.0409***	-0.0352***	-0.0328**
			(-3.55)	(-2.85)	(-2.43)
Ikurt			0.0010	0.0006	0.0000
			(0.31)	(0.17)	(0.00)
Idiovol					0.1280
					(0.44)
$\text{Idiovol}^{\perp}$				-0.3270	
				(-0.84)	
Idioskew				-0.0026	-0.0028
				(-0.22)	(-0.24)
α_0	0.1980***	0.1370	0.1020	0.0890	-0.0256
	(3.26)	(0.52)	(0.39)	(0.32)	(-0.09)
Firm-Month Obs.	47,861	47,813	47,813	47,813	47,813
adj. R^2	1.10%	2.20%	2.40%	3.00%	3.00%

Table 11: **Fama and MacBeth (1973) Cross-Sectional Regressions:** $NCSKEW_{i,t:t+6mo}$**.** Fama and MacBeth (1973) cross-sectional regressions of $\text{NCSKEW}_{t,t:t+6mo}$. t-statistics are adjusted according to Newey and West (1987) with 12 lags.

	Option	Firm-chara	BKM*	Idiosyn*	Raw
$LPM^{\mathbb{Q}}$	-96.75***	-87.82***	-80.68***	-69.36**	-37.69
	(-5.29)	(-3.10)	(-2.92)	(-2.36)	(-1.25)
Beta_{mkt}		-0.0125	-0.0007	-0.0047	0.0030
		(-0.29)	(-0.02)	(-0.10)	(0.07)
Coskew		0.0512	0.0831	0.147	0.140
		(0.46)	(0.73)	(1.22)	(1.21)
Size		0.0048	-0.0030	-0.0012	0.0063
		(0.27)	(-0.16)	(-0.06)	(0.31)
Ivol					-0.386**
					(-2.31)
$\text{Ivol}^{\perp}$			-0.203	-0.125	
			(-1.56)	(-1.01)	
Iskew			-0.0377***	-0.0322***	-0.0302**
			(-3.20)	(-2.63)	(-2.32)
Ikurt			0.0017	0.0014	0.0008
			(0.55)	(0.43)	(0.23)
Idiovol					0.144
					(0.52)
$\text{Idiovol}^{\perp}$				-0.381	
				(-1.06)	
Idioskew				-0.0044	-0.0027
				(-0.39)	(-0.24)
α_0	0.0494	-0.0199	-0.0155	-0.0330	-0.0865
	(1.13)	(-0.09)	(-0.07)	(-0.14)	(-0.31)
Firm-Month Obs.	47,861	47,813	47,813	47,813	47,813
adj. R^2	1.20%	2.40%	2.50%	3.00%	3.30%

Table 12: **Fama and MacBeth (1973) Cross-Sectional Regressions:** $DUVOL_{i,t:t+6mo}$Fama and MacBeth (1973) cross-sectional regressions of DUVOL$_{t,t:t+6mo}$. t-statistics are adjusted according to Newey and West (1987) with 12 lags.

	Option	Firm-chara	BKM*	Idiosyn*	Raw
$VaR^{\mathbb{Q}}$	-6.824***	-6.270***	-5.414***	-4.908***	-3.806***
	(-8.02)	(-4.54)	(-4.04)	(-3.28)	(-3.45)
Beta$_{mkt}$		-0.003	-0.003	-0.006	-0.003
		(-0.16)	(-0.16)	(-0.29)	(-0.16)
Coskew		-0.023	-0.017	0.002	0.000
		(-0.52)	(-0.39)	(0.04)	(0.00)
Size		0.001	-0.001	0.000	0.003
		(0.17)	(-0.18)	(-0.07)	(0.44)
Ivol					-0.127**
					(-2.56)
Ivol$^{\perp}$			-0.035	-0.019	
			(-0.69)	(-0.42)	
Iskew			-0.016***	-0.015***	-0.015***
			(-4.48)	(-4.20)	(-4.12)
Ikurt			0.000	0.000	0.000
			(0.28)	(0.14)	(-0.15)
Idiovol					0.075
					(0.71)
Idiovol$^{\perp}$				-0.063	
				(-0.55)	
Idioskew				-0.001	0.000
				(-0.21)	(-0.09)
α_0	0.122***	0.095	0.080	0.070	0.032
	(4.60)	(1.07)	(0.91)	(0.74)	(0.32)
Firm-Month Obs.	47,861	47,813	47,813	47,813	47,813
adj. R^2	2.50%	4.80%	5.10%	5.70%	6.10%

Table 13: **Fama and MacBeth (1973) Cross-Sectional Regressions:** $DUVOL_{i,t:t+6mo}$Fama and MacBeth (1973) cross-sectional regressions of $\text{DUVOL}_{t,t:t+6mo}$. t-statistics are adjusted according to Newey and West (1987) with 12 lags.

	Option	Firm-chara	BKM*	Idiosyn*	Raw
$ES^{\mathbb{Q}}$	-5.497***	-4.887***	-4.289***	-3.962***	-2.773***
	(-8.26)	(-4.61)	(-4.13)	(-3.42)	(-3.37)
Beta_{mkt}		-0.0037	-0.0010	-0.0029	-0.0033
		(-0.21)	(-0.05)	(-0.15)	(-0.17)
Coskew		-0.0232	-0.0174	0.0025	0.0027
		(-0.54)	(-0.40)	(0.06)	(0.06)
Size		0.0018	-0.0007	-0.0004	0.0032
		(0.33)	(-0.12)	(-0.06)	(0.49)
Ivol					-0.136***
					(-2.62)
$\text{Ivol}^{\perp}$			-0.0546	-0.0329	
			(-0.89)	(-0.59)	
Iskew			-0.0177***	-0.0166***	-0.0155***
			(-4.96)	(-4.68)	(-4.36)
Ikurt			0.0001	-0.0000	-0.0002
			(0.12)	(-0.00)	(-0.18)
Idiovol					0.0633
					(0.61)
$\text{Idiovol}^{\perp}$				-0.0705	
				(-0.66)	
Idioskew				-0.0005	-0.0002
				(-0.15)	(-0.06)
α_0	0.142***	0.0986	0.0851	0.0785	0.0350
	(5.03)	(1.12)	(0.97)	(0.83)	(0.36)
Firm-Month Obs.	47,861	47,813	47,813	47,813	47,813
adj. R^2	2.40%	4.60%	5.10%	5.60%	6.00%

Table 14: **Fama and MacBeth (1973) Cross-Sectional Regressions:** $DUVOL_{i,t:t+6mo}$Fama and MacBeth (1973) cross-sectional regressions of $\text{DUVOL}_{t,t:t+6mo}$. t-statistics are adjusted according to Newey and West (1987) with 12 lags.

	Option	Firm-chara	BKM*	Idiosyn*	Raw
$LPM^{\mathbb{Q}}$	-52.43***	-45.94***	-42.24***	-38.49***	-29.88***
	(-7.84)	(-4.68)	(-4.35)	(-3.70)	(-3.38)
Beta_{mkt}		-0.0056	-0.0016	-0.0036	-0.0038
		(-0.31)	(-0.08)	(-0.19)	(-0.20)
Coskew		-0.0234	-0.0170	0.0021	0.0024
		(-0.54)	(-0.38)	(0.05)	(0.05)
Size		0.0032	0.0002	0.0007	0.0038
		(0.59)	(0.03)	(0.12)	(0.59)
Ivol					-0.124**
					(-2.34)
$\text{Ivol}^{\perp}$			-0.0760	-0.0567	
			(-1.56)	(-1.40)	
Iskew			-0.0167***	-0.0157***	-0.0146***
			(-4.47)	(-4.30)	(-4.16)
Ikurt			0.0003	0.0002	0.0001
			(0.41)	(0.29)	(0.15)
Idiovol					0.0780
					(0.75)
$\text{Idiovol}^{\perp}$				-0.0832	
				(-0.86)	
Idioskew				-0.0010	0.00002
				(-0.30)	(0.01)
α_0	0.0665***	0.0204	0.0203	0.0136	-0.0118
	(3.18)	(0.26)	(0.26)	(0.16)	(-0.12)
Firm-Month Obs.	47,861	47,813	47,813	47,813	47,813
adj. R^2	2.60%	4.90%	5.20%	5.70%	6.20%

Table 15: **Fama and MacBeth (1973) Cross-Sectional Regressions:** $CRASH_{i,t:t+T}$ **with** $VaR^{\mathbb{Q}}$ **Generated From Different RND-recovery Methods and Across Multiple Horizons.** The results of Fama and MacBeth (1973) cross-sectional regressions of $CRASH_{t,t:t+T}$ across 1-month, 3-month and 6-month horizons, over the lagged implied VaR from different RND-recovery methods.

	Constant			GEV			SVI			CGP		
	1mo	3mo	6mo	1mo	3mo	6mo	1mo	3mo	6mo	1mo	3mo	6mo
Panel A: NCSKEW												
Model 1	**-2.19**	**-5.87**	**-6.71**	**-2.21**	**-5.81**	**-6.29**	**-2.27**	**-5.79**	**-6.03**	**-2.02**	**-5.35**	**-5.90**
Model 2	**-2.75**	**-3.63**	**-3.45**	**-2.74**	**-3.83**	**-3.82**	**-1.80**	**-3.60**	**-3.01**	**-3.22**	**-3.72**	**-3.49**
Model 3	**-1.81**	**-2.36**	**-2.81**	**-1.82**	**-2.73**	**-3.30**	-1.60	**-2.81**	**-2.71**	**-2.27**	**-2.83**	**-3.04**
Model 4	**-2.15**	**-2.40**	**-2.34**	**-2.16**	**-2.91**	**-2.96**	**-2.17**	**-3.22**	**-2.68**	**-2.42**	**-2.68**	**2.46**
Model 5	**-2.00**	**-2.28**	**-2.19**	**-2.48**	**-2.25**	**-1.84**	-1.41	**-2.71**	**-1.67**	**-1.98**	**-1.78**	**-1.69**
$\overline{R}^2$	2.36%	2.40%	2.28%	3.68%	2.50%	**2.38%**	3.38%	2.20%	2.14%	**3.98%**	**2.62%**	**2.38%**
Panel B: DUVOL												
Model 1	**-4.74**	**-7.39**	**-8.35**	**-5.52**	**-8.33**	**-9.26**	**-4.60**	**-6.72**	**-7.71**	**-5.44**	**-7.61**	**-8.02**
Model 2	**-2.93**	**-4.24**	**-4.40**	**-3.62**	**-4.88**	**-5.29**	**-2.28**	**-3.41**	**-3.53**	**-4.07**	**-4.72**	**-4.54**
Model 3	**-2.60**	**-3.40**	**-3.67**	**-2.88**	**-4.00**	**-4.46**	**-2.23**	**-3.08**	**-3.12**	**-3.31**	**-3.99**	**-4.04**
Model 4	**-2.66**	**-3.21**	**-3.08**	**-2.94**	**-3.80**	**-3.78**	**-2.45**	**-3.10**	**-2.86**	**-3.10**	**-3.60**	**-3.28**
Model 5	**-1.70**	**-2.80**	**-3.40**	**-2.76**	**-2.97**	**-3.11**	1.09	**-2.21**	**-1.77**	**-2.63**	**-2.93**	**-3.45**
$\overline{R}^2$	2.48%	3.66%	4.66%	2.48%	3.72%	4.74%	2.30%	3.40%	4.42%	**2.76%**	**3.94%**	**4.84%**

Table 16: **Fama and MacBeth (1973) Cross-Sectional Regressions:** $CRASH_{i,t:t+T}$ **with** $ES^{\mathbb{Q}}$ **Generated From Different RND-recovery Methods and Across Multiple Horizons.** The results of Fama and MacBeth (1973) cross-sectional regressions of $CRASH_{t,t:t+T}$ across 1-month, 3-month and 6-month horizons, over the lagged implied ES from different RND-recovery methods.

	Constant			GEV			SVI			CGP		
	1mo	3mo	6mo	1mo	3mo	6mo	1mo	3mo	6mo	1mo	3mo	6mo
Panel A: NCSKEW												
Model 1	**-2.04**	**-4.87**	**-6.05**	**-2.28**	**-4.07**	**-5.28**	-0.66	0.22	-0.11	**-1.7**	**-4.88**	**-5.57**
Model 2	**-2.46**	**-2.77**	**-2.96**	**-2.73**	**-2.58**	**-2.96**	-0.57	-0.36	-0.71	**-2.46**	**-3.03**	**-3.08**
Model 3	**-1.76**	**-2.17**	**-2.6**	**-2.23**	**-1.96**	**-2.78**	-0.78	-0.12	-0.41	-1.58	**-2.19**	**-2.64**
Model 4	**-2.03**	**-2.28**	**-2.28**	**-2.52**	**-2.3**	**-2.85**	-0.58	0.09	-0.53	**-1.8**	**-2.15**	**-2.17**
Model 5	-1.55	**-1.78**	**-1.75**	**-2.18**	-1.32	-1.59	-0.67	-0.29	-0.95	-1.36	-1.03	-1.07
$\overline{R^2}$	3.756%	2.394%	2.278%	3.602%	2.336%	2.14%	3.19%	2.084%	1.886%	**3.842%**	**2.562%**	**2.332%**
Panel B: DUVOL												
Model 1	**-5.09**	**-7.02**	**-8.43**	**-5.15**	**-6.88**	**-7.78**	-0.95	-0.27	-0.72	**-5.41**	**-7.61**	**-8.26**
Model 2	**-3.23**	**-3.94**	**-4.35**	**-3.36**	**-3.82**	**-4.13**	-1.1	-1.04	-1.64	**-3.96**	**-4.5**	**-4.61**
Model 3	**-2.93**	**-3.57**	**-3.88**	**-3.17**	**-3.11**	**-3.61**	-1.24	-0.82	-1.5	**-3.32**	**-3.86**	**-4.13**
Model 4	**-2.79**	**-3.37**	**-3.34**	**-3.02**	**-3.1**	**-3.41**	-1.15	-0.55	-1.41	**-3.05**	**-3.5**	**-3.42**
Model 5	**-1.77**	**-2.47**	**-2.99**	**-1.91**	**-2.09**	**-2.32**	-1.23	-1.05	**-2.15**	**-2.74**	**-2.58**	**-3.37**
$\overline{R^2}$	2.482%	3.634%	4.65%	2.362%	3.482%	4.36%	2.062%	3.124%	3.856%	**2.74%**	**3.878%**	**4.736%**

Table 17: **Fama and MacBeth (1973) Cross-Sectional Regressions: $CRASH_{i,t:t+T}$ with $LPM^{\mathbb{Q}}$ Generated From Different RND-recovery Methods and Across Multiple Horizons.** The results of Fama and MacBeth (1973) cross-sectional regressions of $CRASH_{t,t:t+T}$ across 1-month, 3-month and 6-month horizons, over the lagged implied LPM from different RND-recovery methods.

	Constant			GEV			SVI			CGP		
	1mo	3mo	6mo	1mo	3mo	6mo	1mo	3mo	6mo	1mo	3mo	6mo
Panel A: NCSKEW												
Model 1	**-2.23**	**-5.19**	**-5.65**	**-2.38**	**-5.07**	**-5.49**	-0.65	0.17	-0.29	**-1.98**	**-5.01**	**-5.29**
Model 2	**-2.66**	**-2.98**	**-2.96**	**-3.07**	**-3.44**	**-3.6**	-0.02	0.19	**-0.01**	**-3**	**-3.38**	**-3.1**
Model 3	**-2.5**	**-2.97**	**-2.9**	**-2.51**	**-2.94**	**-3.52**	-0.41	-0.17	-0.35	**-2.42**	**-2.88**	**-2.92**
Model 4	**-2.57**	**-2.85**	**-2.5**	**-2.67**	**-2.7**	**-2.9**	-0.18	0.02	-0.47	**-2.5**	**-2.73**	**-2.36**
Model 5	**-2.73**	**-2.62**	**-2.28**	**-3.2**	**-1.94**	**-2.03**	-0.1	0.43	-0.16	**-2.07**	**-1.65**	-1.25
$\overline{R^2}$	3.72%	2.494%	2.254%	3.54%	2.376%	2.172%	3.134%	2.152%	1.934%	**3.818%**	**2.598%**	**2.476%**
Panel B: DUVOL												
Model 1	**-5.13**	**-7.06**	**-7.94**	**-5.71**	**-7.91**	**-8.55**	**-1.79**	-1.05	-1.5	**-5.9**	**-7.63**	**-7.84**
Model 2	**-3.53**	**-4.06**	**-4.34**	**-3.87**	**-4.56**	**-5.05**	-1.36	-1.15	-1.48	**-4.61**	**-4.82**	**-4.68**
Model 3	**-3.52**	**-4.09**	**-4.12**	**-3.54**	**-4.24**	**-4.75**	**-1.88**	-1.41	**-1.82**	**-3.9**	**-4.36**	**-4.35**
Model 4	**-3.4**	**-3.88**	**-3.71**	**-3.46**	**-3.87**	**-4.05**	**-1.71**	-1.19	**-1.79**	**-3.58**	**-3.98**	**-3.7**
Model 5	**-2.88**	**-3.15**	**-3.46**	**-2.91**	**-2.81**	**-3.26**	-1.48	-0.87	**-1.75**	**-2.64**	**-2.96**	**-3.38**
$\overline{R^2}$	2.454%	3.708%	4.614%	2.376%	3.534%	4.486%	1.958%	3.17%	3.934%	**2.646%**	**3.978%**	**4.926%**

Table 18: **Descriptive Statistics & Cross-sectional Correlations of Implied Moments From Recovered RND.** The descriptive statistics and cross-sectional correlations of $Vol^{\mathbb{Q}}$ and $Skew^{\mathbb{Q}}$ estimated by using the recovered RND with different recovery methods, based on the monthly sample.

	Constant		GEV		SVI		CGP		BKM	
	$Vol^{\mathbb{Q}}$	$Skew^{\mathbb{Q}}$	$Vol^{\mathbb{Q}}$	$Skew^{\mathbb{Q}}$	$Vol^{\mathbb{Q}}$	$Skew^{\mathbb{Q}}$	$Vol^{\mathbb{Q}}$	$Skew^{\mathbb{Q}}$	$Vol^{\mathbb{Q}}$	$Skew^{\mathbb{Q}}$
Panel A: Descriptive statistics										
Mean	0.33	-4.21	0.33	-11.68	0.33	1.36×10^6	0.35	0.030	0.28	-5.29
Std	0.12	383.23	0.12	586.05	0.13	3.79×10^8	0.13	8.69	0.15	511.19
Median	0.30	-2.44	0.31	-4.37	0.30	-2.69	0.33	-0.26	0.25	-1.19
5^{th} Percentile	0.18	-29.49	0.18	-22.20	0.18	-67.80	0.20	-7.68	0.13	-3.29
95^{th} Percentile	0.57	3.34	0.56	1.29	0.57	6.55	0.60	10.39	0.54	1.04
Firm-Month Obs.	73,152		69,683		77,391		61,089		76,845	
Panel B: Cross-sectional correlations										
Constant	1.0	1.0	0.96	0.00	0.91	0.00	0.88	0.01	0.84	0.01
GEV			1.0	1.0	0.88	0.00	0.88	-0.01	0.80	-0.004
SVI					1.0	1.0	0.82	0.00	0.74	0.00
CGP							1.0	1.0	0.75	0.02
BKM									1.0	1.0

Table 19: **Sharpe Ratio & t-Statistics of The High-Low Portfolios.** Economic performance of the High-Low portfolios of moment-based predictors constructed by different RND-recovery methods.

	Constant		GEV		SVI		CGP	
	SR	t-Statistics	SR	t-Statistics	SR	t-Statistics	SR	t-Statistics
$Skew^{\mathbb{Q}}$	0.02	0.29	0.12	1.76	0.05	0.71	**0.20**	**2.9*****
$\Delta Vol^{\mathbb{Q}}$	0.00	0.02	-0.023	-0.34	-0.00	-0.07	**0.04**	**0.65**
$AVAR^{\mathbb{Q}}$	0.10	1.43	0.10	1.52	0.11	1.61	**0.12**	**1.80***

Table 20: **Fama and MacBeth (1973) Cross-sectional Regressions: Stock Monthly Returns with The Lagged $\mathcal{M}_t$ Only.** Univariate Fama and MacBeth (1973) cross-sectional regressions of monthly stock returns over the lagged $\mathcal{M}_t$.

	$Skew^{\mathbb{Q}}$				$\Delta Vol^{\mathbb{Q}}$				$AVAR^{\mathbb{Q}}$			
	Constant	GEV	SVI	CGP	Constant	GEV	SVI	CGP	Constant	GEV	SVI	CGP
$\mathcal{M}_t$	0.000	0.000	0.000	0.001*	0.014	0.000	-0.005	0.021**	3.364*	1.103	0.124	5.328**
	(1.02)	(1.44)	(0.04)	(1.97)	(0.81)	(-0.01)	(-0.39)	(2.01)	(1.72)	(1.23)	(0.95)	(2.50)
α_0	0.01***	0.01***	0.01***	0.01***	0.009***	0.010***	0.009***	0.009***	0.009***	0.009***	0.009***	0.009***
	(2.94)	(3.00)	(3.10)	(2.72)	(3.08)	(3.23)	(2.78)	(3.03)	(3.00)	(2.79)	(2.82)	(2.81)
Firm-Month Obs.	71,631	68,548	74,307	60,597	48,639	48,639	48,639	48,639	59,626	58,821	60,584	60,282
adj. R^2	0.7%	0.4%	0.0%	0.3%	1.2%	1.1%	0.7%	0.8%	1.94%	1.31%	0.71%	3.0%

Table 21: **Fama and MacBeth (1973) Cross-sectional Regressions: Stock Monthly Returns with The Lagged $\mathcal{M}_t$ Generated From The CGP-based RND-recovery**Fama and MacBeth (1973) cross-sectional regressions of monthly stock returns over the lagged option-implied predictor $\mathcal{M}_t$ constructed from the CGP-based RND.

	$Skew^{\mathbb{Q}}$				$\Delta Vol^{\mathbb{Q}}$				$AVAR^{\mathbb{Q}}$			
	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
$\mathcal{M}_t$	0.001*	0.001 **	0.001**	0.001	0.021**	0.019**	0.019*	0.020*	5.328**	5.566***	5.936***	5.459***
	(1.97)	(2.12)	(2.34)	(1.25)	(2.01)	(2.18)	(1.84)	(1.75)	(2.50)	(3.29)	(3.51)	(2.91)
α_0	0.01***	0.003	0.005	0.001	0.009***	0.002	0.004	0.001	0.009***	0.004	0.006	0.004
	(2.72)	(0.79)	(1.15)	(0.23)	(3.03)	(0.46)	(0.73)	(0.09)	(2.81)	(1.29)	(1.61)	(0.64)
Firm-chara	No	Yes	Yes	Yes	No	Yes	Yes	Yes	No	Yes	Yes	Yes
BKM	No	No	Yes	Yes	No	No	Yes	Yes	No	No	Yes	Yes
Idiosyn	No	No	No	Yes	No	No	No	Yes	No	No	No	Yes
Firm-Month Obs.	60,597	60,442	60,401	50,207	48,639	48,609	48,609	48,609	60,282	60,128	60,090	50,076
adj. R^2	0.3%	8.1%	10.4%	12.0%	0.8%	8.6%	11.0%	12.2%	3.0%	9.8%	11.5%	13.1%

Table 22: **Stock-level Predictive OLS Regressions: Stock Return Around EAD with The $AVAR^{\mathbb{Q}}$ Innovation.** Stock-level OLS regressions of stock's cumulative return around EAD over the changes in $AVAR$ estimated from stock's implied RND in ex-ante. The variable $AVAR(-30,-20)$ and $AVAR(-30,-5)$ are used to predict the stock return around EAD, $ret(-19,1)$ and $ret(-4,1)$, respectively.

	Constant		GEV		SVI		CGP	
	$ret(-19,1)$	$ret(-4,1)$	$ret(-19,1)$	$ret(-4,1)$	$ret(-19,1)$	$ret(-4,1)$	$ret(-19,1)$	$ret(-4,1)$
Panel A: Technology								
AAPL	6.0477	9.9047	5.5107	6.6367	-0.9069	1.6402**	-40.4151**	-8.6537
	0.2927	(0.73)	(0.35)	(0.88)	(-0.87)	(2.05)	(-2.08)	(-0.81)
CSCO	-27.9290*	-10.7046	24.4054**	1.0285	1.8093	1.0429	57.1950**	14.9497
	(-1.96)	(-0.97)	(2.21)	(0.15)	(1.18)	(1.33)	(2.45)	(0.77)
NVDA	-28.7480*	5.2046	3.5595	-13.3027***	-0.4089	1.7725	-28.8055*	-7.1988
	(-1.79)	(0.39)	(0.44)	(-3.18)	(-0.14)	(0.87)	(-1.70)	(-0.46)
EA	-14.5192	16.2953	-1.9196	-21.1370	-1.0703	-0.3245	-138.5293*	-65.9735
	(-0.74)	(0.79)	(-0.16)	(-1.07)	(-0.80)	(-0.37)	(-1.74)	(-1.45)
Panel B: Healthcare								
ABT	14.0254	16.4094	24.1011**	6.0253	-0.2470	-0.6577	63.0966***	34.6365**
	(1.27)	(2.47)**	(2.28)	(0.88)	(-0.25)	(-1.05)	(3.42)	(2.53)
GILD	15.9568	5.1441	20.1782	-41.5540**	2.2453	0.8665	-66.6414*	3.9827
	(0.69)	(0.33)	0.81	(-2.19)	(1.61)	(1.00)	(-1.74)	(0.15)
PFE	35.8608	-0.2810	-19.6735	2.8070	0.5086	-0.4690	-63.8612*	-56.7734**
	(1.40)	(-0.01)	(-0.60)	(0.17)	(0.49)	(-0.49)	(-1.70)	(-2.23)
LH	3.1087	-3.7077	22.2431	10.9793	2.5118*	1.2036	71.9068*	47.7638*
	(0.11)	(-0.17)	(0.91)	(0.52)	(1.87)	(1.18)	(1.88)	(1.90)

Table 22 (Cont'd): **Stock-level Predictive OLS Regressions: Stock Return Around EAD with The $AVAR^{\mathbb{Q}}$ Innovation.** Table 11 (Cont'd): Firm-level OLS regressions of stock's cumulative return around EAD over the changes in $AVAR$ estimated from stock's implied RND in ex-ante. The variable $AVAR(-30,-20)$ and $AVAR(-30,-5)$ are used to predict the stock return around EAD, $ret(-19,1)$ and $ret(-4,1)$, respectively.

	Constant		GEV		SVI		CGP	
	$ret(-19,1)$	$ret(-4,1)$	$ret(-19,1)$	$ret(-4,1)$	$ret(-19,1)$	$ret(-4,1)$	$ret(-19,1)$	$ret(-4,1)$
Panel C: Financial								
AIG	35.7875***	30.4688***	-5.2388	-3.7521	4.1450*	-1.9666	-135.1817***	-109.4042***
	(5.88)	(5.62)	(-0.31)	(-0.38)	(1.94)	(-0.93)	(-5.90)	(-5.43)
C	-2.3817	6.5626	2.1785	10.3519***	-0.8545	0.6996	-75.3655***	11.5406
	(-0.29)	(1.14)	(0.52)	(3.80)	(-0.78)	(1.24)	(-3.59)	(0.86)
CB	3.5973	-22.1985***	11.0458	-5.5043	0.9135	0.0132	140.1492***	-48.7901***
	(0.34)	(-3.14)	(0.74)	(-0.52)	(0.66)	(0.02)	(7.06)	(-3.85)
MSCI	8.6096	-14.1011	28.4621*	11.6441	4.7760	3.9019	228.8783**	-162.1783*
	(0.94)	(-0.71)	(1.91)	(0.38)	(0.70)	(1.02)	(2.68)	(-2.10)
Panel D: Industrial								
MMM	-24.4914	-33.5160**	22.5130	-11.0100	-0.6168	0.0704	86.1972**	13.2952
	(-1.09)	(-2.58)	(0.85)	(-0.80)	(-0.56)	(0.11)	(2.21)	(0.72)
EMN	9.6860	-1.6533	51.0754**	-0.7098	-2.6231	0.3163	-124.3675**	-57.5852*
	(0.69)	(-0.17)	(2.11)	(-0.05)	(-1.01)	(0.30)	(-2.21)	(-1.78)
CTAS	28.5855	-7.3219	-14.0730	-12.0222	6.5470*	-0.1897	97.2485***	65.1856***
	(1.50)	(-0.53)	(-1.23)	(-1.57)	(1.79)	(-0.16)	(3.34)	(2.88)
UAL	-13.8635	0.6806	18.6255	-1.8257	3.5791	-0.2066	126.9446***	61.5061*
	(-0.55)	(0.13)	(1.21)	(-0.21)	(1.48)	(-0.12)	(3.39)	(1.71)

Table 23: **Panel Predictive OLS Regressions with Fixed Effects: Stock Return Around EAD with The $AVAR^{\mathbb{Q}}$ Innovation.** Panel OLS regressions (with fixed effects) of stock's cumulative return around EAD over the changes in $AVAR$ estimated from stock's implied RND in ex-ante. Panel A reports the result for the full-sample stock data from 1996 to 2022. Panel B reports the result for stock data prior to the 07-08 financial crisis. Panel C reports the result for stock data after the crisis.

	Constant		GEV		SVI		CGP	
	$ret(-19,1)$	$ret(-4,1)$	$ret(-19,1)$	$ret(-4,1)$	$ret(-19,1)$	$ret(-4,1)$	$ret(-19,1)$	$ret(-4,1)$
Panel A: 1996-2022								
$AVAR(-30,-20)$	-0.2768		1.1773***		-0.0103		2.3662**	
	(-0.54)		(2.61)		(-0.14)		(2.15)	
$AVAR(-30,-5)$		0.0664		-0.6244**		0.0461		1.4534**
		(0.18)		(-2.01)		(0.92)		(2.01)
No. Firm-EAD Obs.	26,080	25,731	25,512	25,214	26,538	26,190	26,311	25,979
Panel B: 1996-2007								
$AVAR(-30,-20)$	-0.6014		1.3235		-0.1034		2.0841	
	(-0.29)		(1.15)		(-0.58)		(1.00)	
$AVAR(-30,-5)$		1.3276		-1.2887**		0.1890		0.1421
		(1.02)		(-2.00)		(1.54)		(0.10)
No. Firm-EAD Obs.	5,239	5,078	5,134	4,983	5,445	5,276	5,435	5,268
Panel C: 2007-2022								
$AVAR(-30,-20)$	-0.2314		1.1332**		0.0273		2.5837**	
	(-0.44)		(2.30)		(0.35)		(1.97)	
$AVAR(-30,-5)$		-0.2636		-0.2911		0.0243		1.9285**
		(-0.07)		(-0.81)		(0.44)		(2.25)
No. Firm-EAD Obs.	20,841	20,653	20,378	20,231	21,093	20,914	20,876	20,711

Figures

Figure 1: **The quantile-to-quantile loss of each recovered RND.** From the top row to the bottom row, the sample size level l decreases from 1.0 to 0.05, and from the left column to the right column, the noise level σ increases from 0.0 to 1.2. The letters $a, b, c, d, e, f, g, h, i$ denote the absolute deviation of the $10\%, 20\%, 30\%, 40\%, 50\%, 60\%, 70\%, 80\%, 90\%$ quantile respectively.

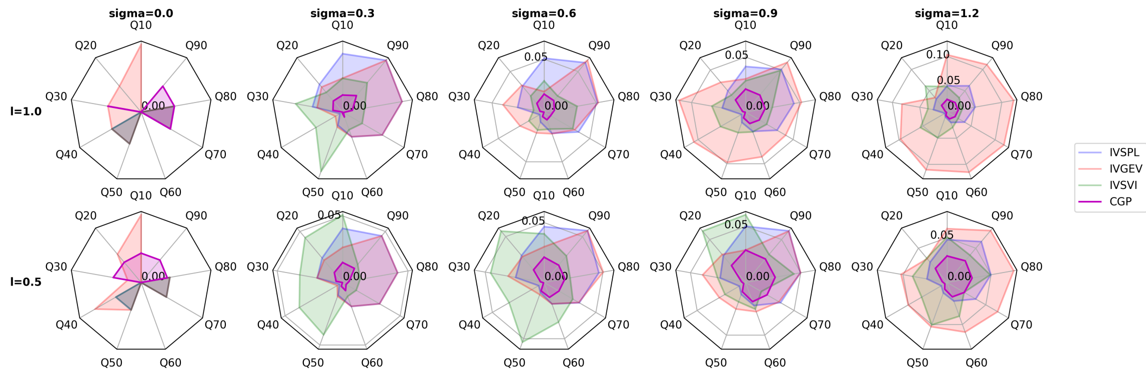

(a) RND quantile loss when the sample size level l is set to 1.0 and 0.5.

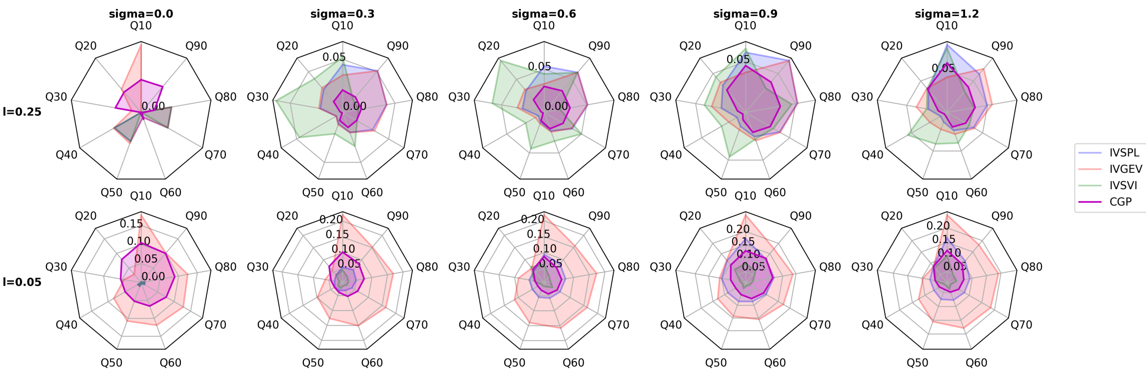

(b) RND quantile loss when the sample size level l is set to 0.25 and 0.05.

Figure 2: **The density divergence loss of each recovered RND.** From the top row to the bottom row, the sample size level l decreases from 1.0 to 0.05, and from the left column to the right column, the noise level σ increases from 0.0 to 1.2. The letters j, k, l, m, n denote the L^2, KL, JS, Wasserstein and Hellinger divergence measurement.

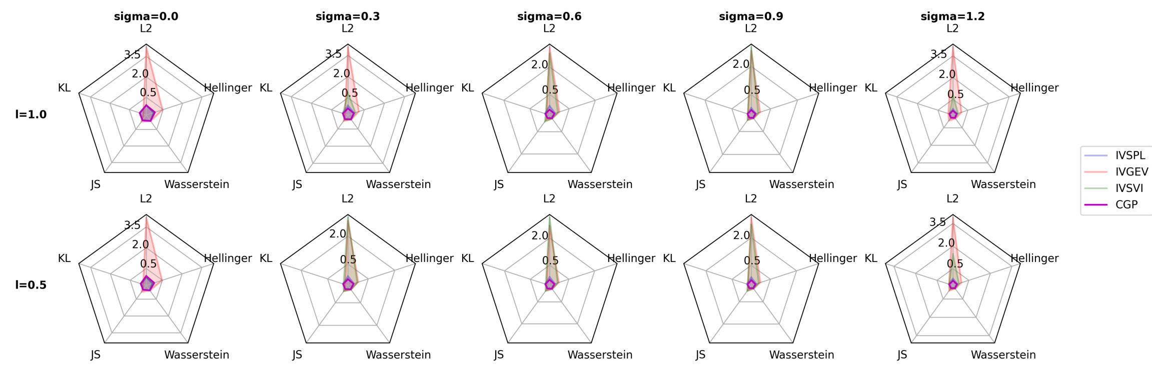

(a) RND divergence loss when the sample size level l is set to 1.0 and 0.5.

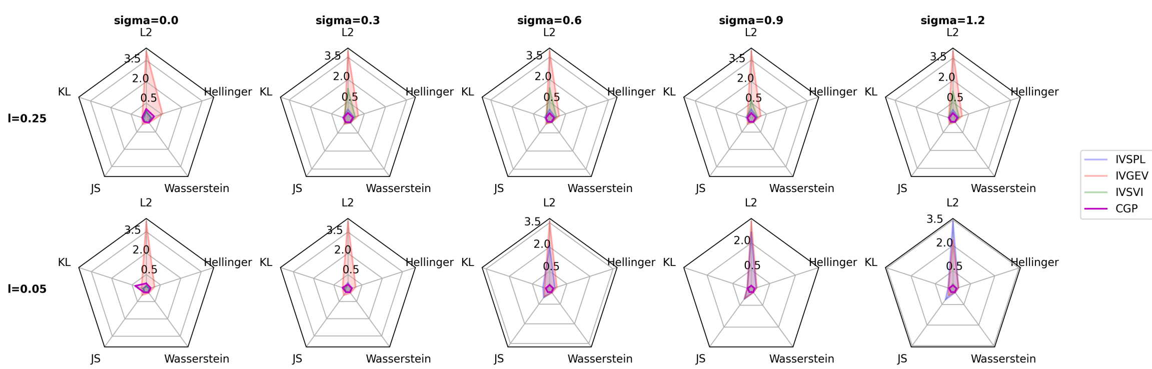

(b) RND divergence loss when the sample size level l is set to 0.25 and 0.05.

Figure 3: **RND estimation evaluation on slected stocks.** We calculate the PIT transform z_t according to Diebold et al. (1997) using the stock options with one-month time-to-maturity and the realized stock returns one month after the estimation of the stock-level RND. This evaluation is done on the 16 selected stocks for demonstration purpose. We also compare the RND evaluation across the recovery methods of IVSPL, IVGEV, IVSVI and our proposed CGP-based technique. Each of the 16 subplots in this figure shows the RND evaluation for one of the stocks in the selected sample. The first row of each subplot shows the histogram of z_t, and the second row shows the quantile z_t^Q versus the empirical cumulative probabilities z_t^F of z_t. Since $z_{i,t} \overset{i.i.d.}{\sim} \mathcal{U}(0,1)$ if the series of z_t is from a well-specified RND, as supposed in Diebold et al. (1997) and Berkowitz (2001), then the line $\left(z_t^Q, z_t^F\right)$ should be close to a 45° line if the RND is well-specified. The grey dotted line is the 45° line in this figure. We can observe the obvious departure of $\left(z_t^Q, z_t^F\right)$ from the 45° line on some stocks by using IV-based RND-recovery and the minimum departures by using our proposed CGP-based RND-recovery.

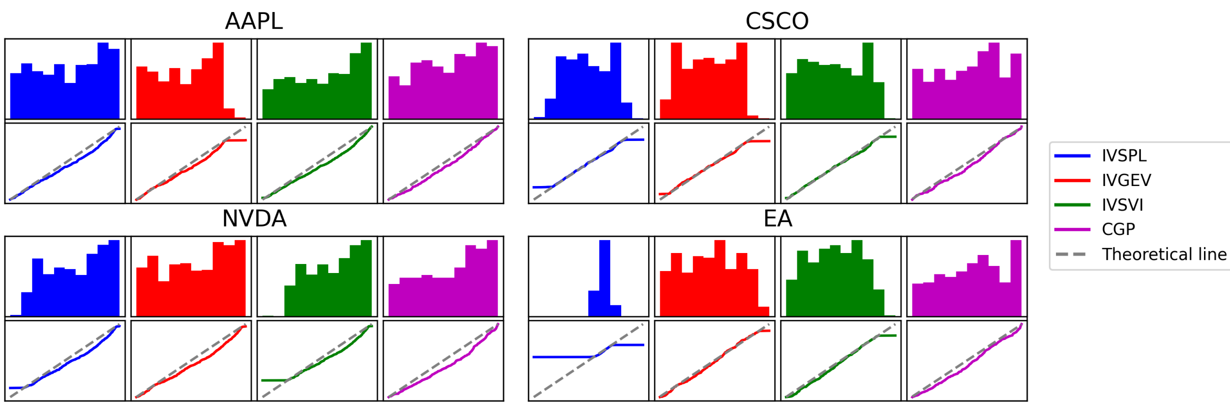

(a) Technology

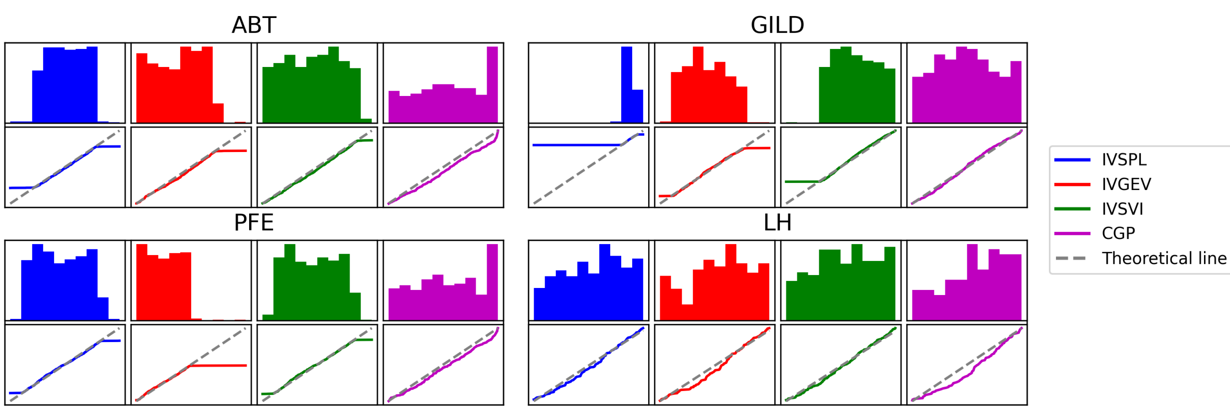

(b) Healthcare

Figure 3 (Cont'd): **RND estimation evaluation on real stocks.** The panel (a), (b) in the above figure present the RND evaluation results on selected stocks in the technology and healthcare industry. The panel (c), (d) in the figure below present the RND evaluation results on selected stocks in the financial and industrial industry.

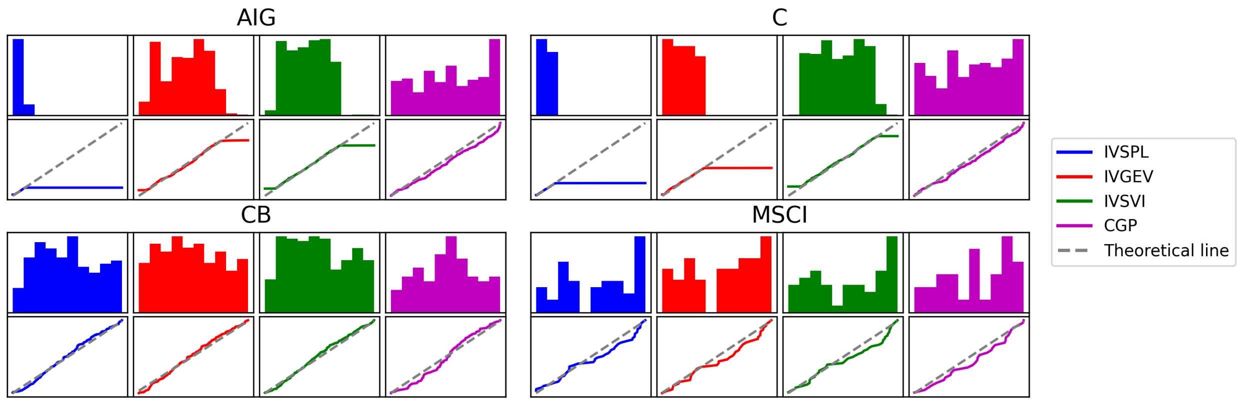

(c) Financial

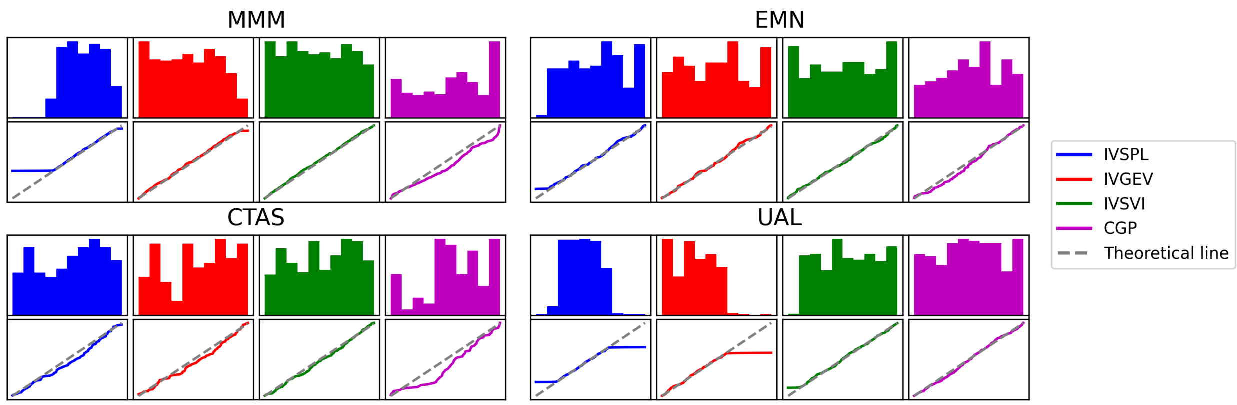

(d) Industrial

Figure 4: **The comparison of the informativeness of implied tail risk measures across different RND-recovery methods.** The averaged R^1 across all stock-level quantile regressions. From left ($VaR^{\mathbb{Q}}$) to right ($LPM^{\mathbb{Q}}$), more RND information is used, and the measure becomes more sensitive to the accuracy and the shape of the RND.

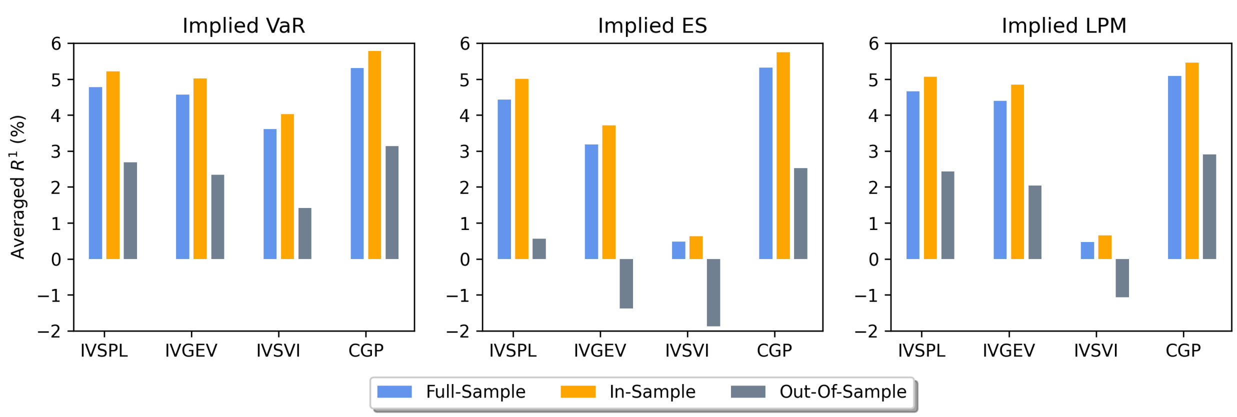

Figure 5: **The implied skewness portfolios.** The cumulative return of High-Low portfolios, sorted on $Skew^{\mathbb{Q}}$ estimated from the recovered RND with different methods.

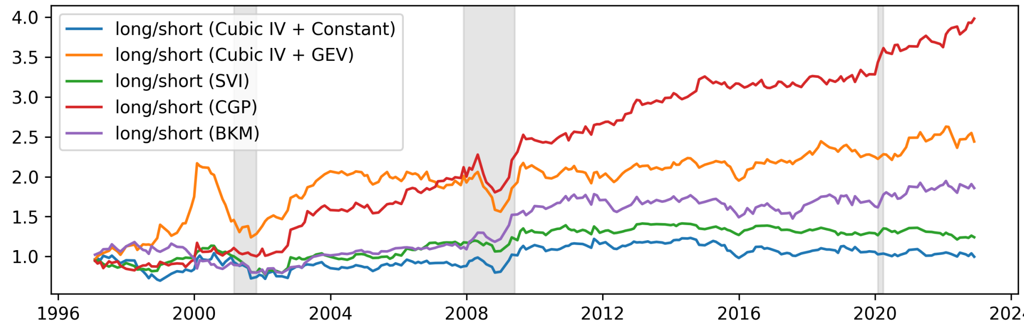

Figure 6: **The implied volatility innovation portfolios.** The cumulative return of High-Low portfolios, sorted on $\Delta Vol^{\mathbb{Q}}$ estimated from the recovered RND with different methods.

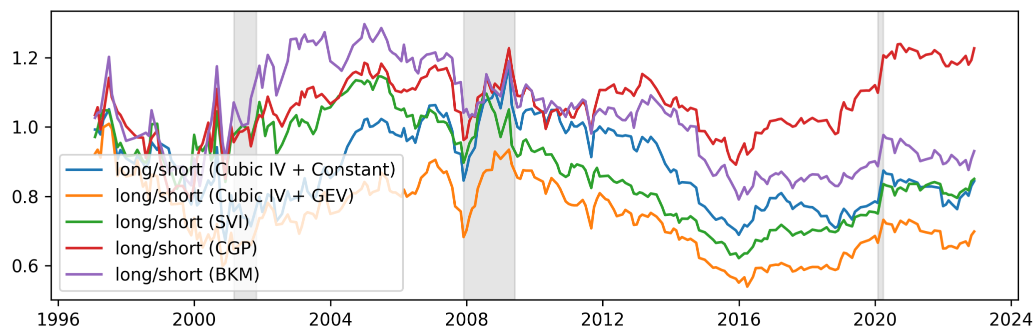

Figure 7: **The implied asymmetry of variance portfolios.** The cumulative return of High-Low portfolios, sorted on $AVAR^{\mathbb{Q}}$ estimated from the recovered RND with different methods.