University of Southampton

Faculty of Social Sciences

School of Economic, Social and Political Sciences

Essays in Macro Asset Pricing Models with News Sentiment

by

Xuefeng Yang

ORCID: 0009-0005-7902-9854

Thesis for the degree of *Doctor of Philosophy*

November 2025

University of Southampton

Abstract

Faculty of Social Sciences

Department of Economics

Thesis for the degree of *Doctor of Philosophy*

Essays in Macro Asset Pricing Models with News Sentiment

by

Xuefeng Yang

This thesis investigates the role of investor sentiment in asset pricing through the lens of Shefrin's behavioral pricing kernel, quantile preferences, and probability weighting functions from non-expected utility theory. Each chapter examines a distinct dimention of how sentiment influences asset prices, offering both theoretical insights and empirical validation across different model frameworks.

Chapter 2 conducts an empirical test of the consumption-based CAPM by approximating the stochastic discount factor (SDF) as a linear function of consumption growth and news sentiment. This approach aims to reconcile the tension between the high cross-sectional explanatory power of conditional asset pricing models and the underlying assumptions. By explicitly incorporating sentiment into the pricing kernel, the model offers a more behaviorally consistent representation of investor preferences. Notably, the inclusion of sentiment results in a less negative SDF and a more

moderate estimate of risk aversion, thereby aligning empirical estimates more closely with economically plausible values.

Chapter 3 introduces a quantile-preferences asset pricing model that incorporates a sentiment index constructed using Principal Component Analysis (PCA) on a comprehensive set of survey-based, market-based, and news-based sentiment indicators. The framework leverages a Panel Quantile Regression (PQR) model to investigate how sentiment influences the cross-sectional distribution of asset returns, with particular emphasis on extreme tail events. By controlling for unobserved heterogeneity across financial assets, the PQR approach enables the identification of systematic patterns in the tail behavior of returns. A comparison between PQR estimates and those from individual Univariate Quantile Regressions (UQR) reveals that, once heterogeneity is accounted for, the first principal component (PC1)—interpreted as a proxy for aggregate sentiment—exerts a more pronounced negative effect on the upper quantiles of returns than indicated by most UQR estimates.

Chapter 4 investigates the relationship between tail-overweighting parameters embedded in the probability weighting functions of non-expected utility models and an external sentiment measure, proxied by the first principal component (PC1) extracted via PCA in Chapter 3. Among the parametric weighting functions analyzed, the Prelec probability weighting function exhibits the strongest empirical alignment, with its curvature parameter α showing a notable 80% correlation with PC1. The findings offer empirical evidence of a significant correlation between probability weighting parameters and an external sentiment measure in real-world, non-experimental data. This contributes to the broader behavioral asset pricing literature by shedding new light on the dynamic nature of probability weighting and its implications for resolving the pricing kernel puzzle.

Contents

Research Thesis: Declaration of Authorship

Print name: Xuefeng Yang

Title of thesis: Essays in Macro Asset Pricing Models with News Sentiment

I declare that this thesis and the work presented in it are my own and have been generated by me as the result of my own original research.

I confirm that:

1. This work was done wholly or mainly while in candidature for a research degree at this University;
2. Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
3. Where I have consulted the published work of others, this is always clearly attributed;
4. Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
5. I have acknowledged all main sources of help;
6. Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
7. None of this work has been published before submission.

Signature: **Date:** 18/11/2025

Acknowledgements

I would like to express my sincere gratitude to my primary supervisor, Dr Michael Hatcher, for his guidance, encouragement, and support throughout my PhD. His advice and constructive feedback have been highly valuable in shaping the development of this thesis. I am also deeply grateful to my second supervisor, Professor José Olmo, whose continuous supervision, thoughtful insights, and rigorous academic advice have played a central role in this work. I have benefited enormously from his expertise, his generosity with his time, and his unwavering support at every stage of the research process. My thanks extend as well to my second supervisor, Dr Serhiy Stepanchuk, for his helpful discussions and constructive suggestions, which substantially improved the clarity and depth of several chapters.

Moreover, I am indebted to Professor Hersh Shefrin, whose seminal contributions to behavioural finance and personal guidance have significantly shaped the direction of my research. His encouragement and intellectual inspiration were instrumental in developing the behavioural foundations of this thesis. Finally, I wish to express my heartfelt appreciation to my parents and my sister for their unconditional love, patience, and unwavering support. Their encouragement throughout the most challenging stages of the PhD has been a constant source of strength, and I could not have completed this journey without them.

Chapter 1

Introduction

Dating back to at least Keynes (1936), behavioral economists have explored how psychological biases drive investor sentiment to push security prices away from their fundamental values. These mispricings distort capital allocation, affect investment decisions, and undermine the financial sector's intermediation role. In recent decades, this research has expanded to include anomalies such as the the equity premium–risk-free rate puzzle, stock price overreaction and underreaction, implied volatility patterns in options, and sentiment-driven price shifts linked to seasonal patterns and major events. Notable contributions come from scholars including Shiller (1981), De Bondt and Thaler (1985), De Bondt and Thaler (1987), Lee et al. (1991a), Jegadeesh and Titman (1993a), Jegadeesh and Titman (2001), Loughran and Ritter (1995), Daniel et al. (1998), Kamstra et al. (2000), Barberis et al. (2001), Poteshman (2001), Baker and Wurgler (2004), Edmans et al. (2007), and Xiong and Yan (2010).

Behavioral finance differs from neoclassical finance through two central features: psychology and limits to arbitrage (Barberis & Thaler, 2003). The psychological dimension includes heuristics, biases, and framing effects (Kahneman & Tversky, 1979a;

Kahneman et al., 1982). Limits to arbitrage, arising from short-selling constraints, transaction costs, illiquidity, capital limitations, noise trader risk, and regulatory barriers, restrict rational investors from fully correcting mispricings. Sentiment reflects psychological biases, while limits to arbitrage allow these biases to affect market prices.

Although sentiment is a key element of behavioral finance, the literature lacks a single, standardized definition. Traditionally, sentiment is typically defined contextually and proxied by factors like closed-end fund discounts, returns on new equity issues, or premiums on dividend-paying stocks. Moreover, sentiment variables have generally been incorporated into empirical asset pricing models as non-tradable pricing factors rather than as integral components of a formal SDF model. This thesis contributes by estimating a theoretically grounded measure of sentiment using consumption data, market returns, option prices and risk-free rates, where sentiment is defined as a change of measure that connects objective and subjective beliefs.

To integrate a well defined sentiment within an SDF-based asset pricing framework, the derivation of SDF from the fundamental asset pricing equation is first presented. The SDF, is a state-dependent function that discounts payoffs based on time and risk preferences, as well as a state price per unit objective probability, which encapsulates investor preferences for payoffs across various states of the world. Under a no-arbitrage framework, every asset price can be written by the expected product of the pricing kernel and its corresponding payoff. Therefore, when combined with a model of state probabilities, the SDF offers a comprehensive description of asset prices, expected returns, and the associated risk premia.

Starting from the investor's problem of maximizing intertemporal utility given an initial wealth level, consider an infinitely lived representative investor in complete markets who ranks consumption paths according to a time-additive expected utility model. The investor's objective is to maximize total expected discounted utility, expressed as

$$U = \max_{c_t,\omega_t} E_t \left\{ \sum_{j=0}^{\infty} \beta^j U(c_t + j) \right\}, \quad \beta > 0, \quad c_t \geq 0$$
$$(\omega_t R_{t+1} + (1-\omega_t) R^f_{t+1})(A_t - c_t) = A_{t+1} \tag{1.1}$$

where E_t is expectation conditional on the investor's time-t information, A_t is household wealth at the beginning of period t, R_{t+1} and R^f_{t+1} are stock return and risk-free return respectively between time t and $t+1$, c_t represents investors' consumption expenditure, ω_t represents the share of investors' savings invested in the stock market, U(.) is the period utility function which is increasing, concave, and twice differentiable, β represents the subjective discount rate that measures the degree of investors' impatience.

Consider a general asset pricing framework that accommodates heterogeneity in agents' beliefs and preferences. This flexible structure nests many structural asset pricing models as special cases and enables a systematic analysis of how sentiment-driven belief formation influences asset pricing outcomes. To price an asset, structural economic models must, at a minimum, specify the following three components: (1) the stochastic process that determines the set of marginal agents $\{\mathcal{M}_t\}_{t=0}^{\infty}$ responsible for pricing the asset across all periods and contingencies; (2) the one-step-ahead SDF $\{m^m_{t+1}\}_{t=0}^{\infty}$, which at least one marginal agent uses to discount future payouts from period $t+1$ back to period t; and (3) the subjective probability measure f_{pS} of the marginal agent, which specifies in each period t the perceived probability distribution over the asset's next-period payoff x_{t+1}.

Building on the three elements outlined above, by taking the first-order conditions, the investor balances the marginal utility loss from reducing current consumption against the discounted marginal utility gain from future asset payoffs. In the absence of arbitrage, this tradeoff can be succinctly expressed by the Lucas-Breeden asset pricing equation:

$$p_t = E_t^S \left[\beta \frac{u'(c_{t+1})}{u'(c_t)} x_{t+1} \right] = E_t^S (m_t x_{t+1}) \tag{1.2}$$

where p_t is the current asset price, x_{t+1} is the asset payoff in one period, m_t represents the SDF. The SDF is often referred to as the pricing kernel because, when the expectation E is expressed as an integral, m serves as the kernel function in integral transform. Other common terms for the SDF include the "marginal rate of substitution" (representing the ratio of utilities across states when utility is separable and additive, adjusted by the risk-neutral discount rate), a "change of measure," a "state-price deflator". Notably, E_t^S denotes the expectation under the subjective probability measure f_{pS} investors hold at time t, indicating that the pricing equation remains valid even when expectations deviate from the true probability measure assumed under rational expectations. This stands in contrast to the rational expectations framework, in which agents' beliefs are assumed to exactly match the true underlying probability measure. The subjective probability measure f_{pS} is essentially the mathematical representation of an agent's subjective beliefs, which plays a foundational role in asset pricing, particularly in models that relax rational expectations. This subjective expectation is computed over the probability distribution of future states of the world, reflecting the agent's subjective beliefs about how uncertainty will unfold over time. Nagel (2013) uses the term "sentiment" as a catch-all label to describe this type of expectation formation, where beliefs diverge from rational expectations and Bayesian learning benchmarks.

According to Shefrin (2008), sentiment arises from erroneous beliefs. Consider that the asset payoff x_{t+1} is governed by an objectively correct probability density function (pdf), yet individual investors hold only subjective beliefs about it. An investor 's subjective beliefs are correct exhibits zero sentiment, whereas any deviation from it results in nonzero sentiment. A central issue in behavioral asset pricing theory is how, in equilibrium, the market aggregates individual investors' subjective probability density functions (pdfs) into a single market-wide probability density function underling equation (1.1). Naturally, the concept of sentiment applicable to individual investors' pdfs also applies to the market-based distribution (Shefrin, 2008). The aggregation question has been studied since at least Lintner (1969). Shefrin (2008) shows that in a complete markets model with power utility, the equilibrium market-based pdf f_{p_R} is a generalized weighted Hölder average of individual investors' pdfs-where the weights reflect relative wealth or consumption and the exponents correspond to coefficients of relative risk aversion.

Within the neoclassical SDF framework, investors' beliefs hinge on an underlying state variable—such as aggregate consumption growth. In this context, an individual's sentiment is characterized by the discrepancy between the objective probability density functions and the subjective pdf. One of the prominent models is Lucas Jr (1978)'s consumption-based asset pricing model, where the pricing kernel represented as the intertemporal marginal rate of substitution (IMRS) between consumption at dates t and T, is given by $m_t = \frac{U'(C_T)}{U'(C_t)}$. Under the assumption of Constant Relative Risk Aversion (CRRA) preferences, the utility function is typically modeled as a power utility function, and consequently, the pricing kernel is expressed as

$$m_t(\theta) = \theta_0 \left(C_T/C_t\right)^{-\theta_1} \tag{1.3}$$

where θ_0 is the time preference parameter measuring the degree of impatience and θ_1

represents the coefficient of relative risk aversion, and $\theta = (\theta_0, \theta_1)$. C_t is the initial aggregate consumption at time t, C_T denotes the aggregate consumption at a future time T.

A large number of research estimates the pricing kernel using aggregate consumption data. Early work in consumption-based asset pricing models the pricing kernel directly as a function of aggregate consumption growth. Hansen and Singleton (1982) assume a power-utility function of aggregate U.S. consumption and estimate the pricing kernel via maximum likelihood and the generalized method of moments (GMM). Chapman (1997a) models the pricing kernel using current and lagged values of consumption as state variables, and represents the kernel through an orthogonal polynomial expansion. Hansen and Jagannathan (1991) establish limits on the mean and volatility of the consumption-based pricing kernel by relating them to the corresponding mean and standard deviation of excess returns on the market portfolio. A persistent empirical drawback in all consumption-based approaches is measurement and aggregation problems in available aggregate consumption proxies, which can undermine empirical findings.

In general, the pricing kernel is influenced not only by current and future consumption but also by all state variables affecting marginal utility. For example, Eichenbaum et al. (1988) incorporate leisure, Startz (1989) integrates durable goods purchases into the pricing kernel, and Bansal and Viswanathan (1993) specify the kernel as a function of equity returns, Treasury bill yields, and the term spread. Moreover, the habit formation models developed by Campbell and Cochrane (1999) emphasize the role of past consumption in shaping current marginal utility. Similarly, the long-run risks framework of Bansal and Yaron (2004), extended by Bansal and Yaron (2012), underscores how persistent shocks to consumption volatility have sizable effects on the pricing kernel.

There is ongoing debate among researchers about which state variables should be incorporated into the pricing kernel. To sidestep this issue, scholars explore to estimate the pricing kernel by projecting it onto the tradable asset payoff x_{t+1} without having to specify the debated state variables. For example, Aït-Sahalia and Lo (2000b) and Jackwerth (2000) estimate the pricing kernel by projecting it onto equity return states using equity index option prices. Rosenberg and Engle (2002) estimate the pricing kernel use index option price data to estimate the projection of the pricing kernel onto the returns to the $S\&P$ 500. According to Rosenberg and Engle (2002), the original pricing kernel is defined as $m_t = m_t(z_t, z_{t+1})$, where z_t is a vector of state variables entering the pricing kernel. The equation(1.1) can be rewritten by decomposing the joint density $f_t\left(x_{t+1}, z_{t+1}\right)$ into the product of the conditional density $f_t\left(z_{t+1} \mid x_{t+1}\right)$ and the marginal density $f_t\left(x_{t+1}\right)$. The expectation can be evaluated in two steps:

$$p_t = \mathrm{E}_t\left[m_t^*\left(x_{t+1}\right) x_{t+1}\right], \quad m_t^*\left(x_{t+1}\right) = \mathrm{E}_t\left[m_t\left(z_t, z_{t+1}\right) \mid x_{t+1}\right] \tag{1.4}$$

First, the projected pricing kernel $m_t^*\left(x_{t+1}\right)$ is derived by integrating the original pricing kernel $m_t\left(z_t, z_{t+1}\right)$ over the conditional density $f_t\left(z_{t+1} \mid x_{t+1}\right)$. Second, the asset price p_t is obtained by integrating the product of the projected pricing kernel and payoff using the marginal density $f_t\left(x_{t+1}\right)$. This nonparametric method avoids consumption data entirely and imposes minimal functional assumptions, yielding more flexible estimates of the SDF. As Cochrane (2009) explains, the projected pricing kernel-which solely depends on the asset payoff x_{t+1}, preserves the same pricing implications as the original pricing kernel that relies on the full state vector z_{t+1}. Thus, the projected pricing kernel can be succinctly expressed as a function of the asset payoff, which summarizes the influence of all state variables into a function directly tied to the asset's observable payoff.

Starting with prices and payoffs in a complete market, the $p = E(mx)$ representation can be interpreted as an inner product, providing an intuitive, visual framework for most of the theorems without specifying any utility functions. According to Arrow (1964) and Debreu (1959), when there is a complete set of state, that is, for every possible state of the world there is a corresponding security that pays off only if that state occurs, then the price of any security can be expressed as a weighted average of the prices of these state-contingent claims, which are now commonly known as Arrow-Debreu prices.

$$\begin{aligned} p(x) &= \sum_s \pi(x_i)x(x_i) = \sum_s f_p(x_i)\left(\frac{\pi(x_i)}{f_p(x_i)}\right)x(x_i) = E(mx) \\ &= \frac{1}{R^f}\sum f_q(x_i)x(x_i) = \frac{E^*(x)}{R^f} \end{aligned} \tag{1.5}$$

where, π is the price of a contingent claim, as known state price, $f_p(x_i)$ is the objective probability of state s. The pricing kernel m is defined as the ratio of state price to probability $m(x_i) = \frac{\pi(x_i)}{f_p(x_i)}$, which means that in a complete market, the pricing kernel is just a set of contingent claim prices scaled by probabilities. As shown in (1.4), the price of any security can be priced as the expected payoff discounted by the risk-free rate as the discount factor. However, the expectation must be taken with respect to $f_q(x_i)$, called risk-neutral probability, instead of the objective probability $f_p(x_i)$. This concept yields another common transformation of the pricing equation $p = E(mx)$, which derives risk-neutral probabilities from objective probabilities.

In neoclassical asset pricing theory, the risk-neutral density $f_q(x_i)$ is transformed from the objective density f_p by applying a change of measure using the normalized pricing kernel (Cochrane, 2009). This transformation of physical probabilities is used to simplify the mathematics of no-arbitrage pricing. Under the risk-neutral measure, all assets are assumed to earn the risk-free rate, effectively stripping out the risk premia associated with investors' risk aversion, thereby allowing assets to be priced as if investors were indifferent to risk. In particular, risk-neutral probabilities are defined as $f_q(x_i) \equiv R^f m(x_i) f_p(x_i) = R^f \pi(x_i)$, with $R^f \equiv 1/\sum \pi(x_i) = 1/E(m)$, which means that risk-neutral probabilities are obtained by scaling state prices by the normalized pricing kernel.

The distinction between state prices and state probabilities emerges from the fact that market prices must reflect both the likelihood of a future outcome and investors' risk-preferences. The theoretical relationship between state prices and state probabilities is well established, as outlined in C.-f. Huang and Litzenberger (1988). This concept shows discrepancies between the risk-neutral density f_q and the objective density f_p arise from investors' risk preferences, and there is a utility function capturing investors' preferences allowing to match both density functions. Thus, under the assumptions of complete, frictionless markets for a single asset, the relationship between f_q and f_p is captured by the representative investor's utility function. In equilibrium, the investor allocates all available wealth to the risky stock at every moment prior to T and ultimately consumes the stock's terminal value at T, so that $C_T = S_T$. In a representative investor CRRA-framework, the pricing kernel equation (1.2) can be rewritten as $m_t(\theta) = \theta_0 \left(S_T/S_t\right)^{-\theta_1}$, the pricing kernel can then be defined using this relationship as follows:

$$m(x_i) = \frac{f_q(x_i)}{R_f f_p(x_i)} = \theta_0 \left(S_T/S_t\right)^{-\theta_1} \tag{1.6}$$

S_t is the initial value $S\&P$ 500 index at time t, which serves as a proxy for the market portfolio, while S_T denotes the value of the $S\&P$ 500 index at a future time T. (1.4) provides a framework for identifying the projected pricing kernel. This relationship stems from the optimizing condition that equates the marginal rate of substitution (for expected utility) with relative state prices, using consumption at $t = 0$ as the numeraire, which implies that the ratio of the objective density $f_p(x_i)$ to the risk-neutral density $f_q(x_i)$ is proportional to the representative investor's marginal utility, formally, $f_p(x_i)/f_q(x_i) \propto U'(x_i)$.

In a behavioral framework, an analogous relationship exists between the investors' subjective density f_{pS} and f_q instead of between f_p and f_q. The resulting expression for m as the ratio between function f_{pS} and f_q is as follows:

$$\tilde{m}(x_i) = \frac{f_q(x_i)}{R_f f_{pS}(x_i)} = e^{\Lambda} \left(S_T/S_t\right)^{-\theta_1} \tag{1.7}$$

where $e^{\Lambda} = (\theta_0/\theta_{0,e})\,(f_{pS}/f_p)$, Λ is the sentiment function defined in Shefrin (2008), which is a scaled log-change of measure that transforms the objective density f_p into investor's subjective density f_{pS}. This sentiment function is expressed as the logarithm of a composite measure Φ:

$$\Lambda = \ln\left(\Phi\left(x_i\right)\right) = \ln\left(\frac{f_{pS}}{f_p} \cdot \frac{\theta_{0,t}}{\theta_{0,t,p}}\right) = \ln\left(\frac{f_{pR}}{f_p}\right) + \ln\left(\frac{\theta_{0,t}}{\theta_{0,t,p}}\right) \tag{1.8}$$

where, Φ summarizes two types of deviations arising from investors' errors, the log-change of measure $\ln\left(\frac{f_{pR}}{f_p}\right)$ quantifies the percentage difference between the subjective probability weight assigned to a state by investors and its objective probability, while $\ln\left(\frac{\theta_{0,t}}{\theta_{0,t,p}}\right)$ reflects the deviation in the representative investor's equilibrium time discount factor from its value under objectively correct beliefs. When all investors hold objectively correct beliefs, $\Phi = 1$, and hence $\Lambda = 0$.

For simplicity, sentiment Λ refers to $\ln\left(\frac{f_{pR}}{f_p}\right)$ throughout this thesis, then $e^{\Lambda} = \frac{f_{pR}}{f_p}$. When the distortion factor e^{Λ} remains close to one-indicating that the subjective pdf is only slightly different from the objective pdf. Under this conditions, by taking the property of the logarithmic approximation $\ln(1+\epsilon) \approx \epsilon$ for small ϵ, $\Lambda = \ln\left(e^{\Lambda}\right) = \ln\left(1 + \left(e^{\Lambda} - 1\right)\right) \approx e^{\Lambda} - 1$ and thus $e^{\Lambda} \approx 1 + \Lambda$. Therefore, the pricing kernel augmented with sentiment is given by:

$$\tilde{m}(x_i) = e^{\Lambda} \left(S_T/S_t\right)^{-\theta_1} \cong \left(S_T/S_t\right)^{-\theta_1} [1 + \Lambda] \equiv m_c + m_s \tag{1.9}$$

where $m_c = \left(S_T/S_t\right)^{-\theta_1}$ and $m_s = \left(S_T/S_t\right)^{-\theta_1} \ln \Lambda$. This behavioral pricing kernel in equation (1.9) will employed in the second chapter of this thesis.

The behavioral asset pricing approach focuses on the role of sentiment, which gauges the extent of excessive optimism or pessimism among investors. Importantly, sentiment is not just an average measure of these attitudes but reflects the entire distribution of investors' errors. In particular, heterogeneous beliefs can produce distinctive smile patterns in the sentiment function. According to Shefrin (2008), nonzero sentiment underlies the non-monotonic shape of the pricing kernel—a feature seen as a puzzle by neoclassical economists. In contrast, the behavioral approach predicts such non-monotonicity as a signature of strong sentiment, establishing a link between the character of investors' biases and the pricing kernel shape. Notably, the time series of overconfidence estimates corroborates this theoretical link, as the shape of the pricing kernel evolves in tandem with these measures over time.

Moreover, following to Polkovnichenko and Zhao (2013), the pricing kernel based on the rank-dependent family of models can be defined by multiplying the marginal utility by the derivative of the probability weighting function

$$m = \frac{q}{R_f p} = \frac{u'(S_T)}{u'\left(S_t\right)} Z(P) = \left(S_T/S_t\right)^{-\theta_1} Z(P) \tag{1.10}$$

where Z(P) represents the derivative of the probability weighting function.

This thesis comprises three core research chapters, each contributing to the understanding of how investor sentiment influences asset pricing through distinct theoretical and empirical lenses.

Chapter 2 is targeted to reconcile the inconsistence between the high cross-sectional explanatory power of the conditional asset pricing model and the fundamental assumptions by developing a behavioral pricing kernel with news sentiment. This chapter primarily contributes by formalizing the role of news sentiment within the SDF framework. Although news sentiment has frequently been introduced as a nontradable factor in the ad hoc factor models, its integration as a structural component of the SDF remains largely underexplored. Additionally, the chapter employs Chebyshev polynomial approximations to model the pricing kernel in a flexible and robust manner. This approach leverages the expressive power of Chebyshev polynomials to capture nonlinearities and oscillatory behavior induced by sentiment dynamics, while also allowing for time-varying risk aversion through the inclusion of instrumental variables.

Chapter 3 develops quantile preferences asset pricing model by integrating sentiment constructed from PCA to explain asymmetric investor preferences. This chapter contributes to the understanding of how investor sentiment explains the cross-section of extreme tail events in the distribution of stock returns. It employs the Panel Quantile Regression Model, which effectively accounts for unobserved heterogeneity across financial assets and enables the identification of common patterns in the tail behavior of returns across different securities. This approach advances traditional asset pricing models by addressing the limitations of homogeneous preferences and expected utility maximization, offering a more comprehensive framework for analyzing the influence of sentiment on extreme tail events in asset returns. Moreover, this chapter introduces a novel composite sentiment index constructed via Principal Component Analysis (PCA), integrating multiple sentiment proxies—especially those derived from machine learning-based news sentiment.

Chapter 4 examines the correlation between the tail overweighting parameters, embedded in the probability weighting function of the rank-dependent family of models, and external sentiment measures (PC1). The primary contribution of this chapter lies in addressing a gap in the literature by empirically examining the relationship between time-varying overweighting preference parameters of probability weighting functions and investor sentiment toward tail events. Drawing on insights from non-expected utility models—particularly cumulative prospect theory—it investigates how subjective probability weighting, especially the overweighting of tail probabilities, relates to investor sentiment concerning extreme events.

1.1 Theoretical Framework

1.1.1 Belief Formation Beyond Rational Expectation and Bayesian-Learning

Almost all economic decisions are based on by agents' perceptions of the current economy and their expectations about future economic outcomes. For example, the consumer Euler equation highlights how households' inflation expectations shape intertemporal consumption choices via changing households' incentives to save. The standard approach to modeling these beliefs has been full-information rational expectations. Rational Expectations (RE), introduced by Muth (1961) and later developed by economists such as Jr. (1970) and Sargent (1973), assumes that investors fully understand the model and its parameters, making unbiased and objective forecasts about the future using all available information without systematic errors. This ensures that investors' expectations accurately reflect the true underlying data-generating process of the economy. As a result, the expectation of an economic variable—such as prices or returns—aligns with its statistical mean under the true distribution. Moreover, RE stipulates that all subjective probability distributions should align with the objective distributions implied by the asset pricing model in equilibrium (Sargent, 2008). While expectations are central to macroeconomic models, the process of belief formation remains poorly understood.

In financial economics, asset prices inherently embody a forward-looking perspective, driven by investors' expectations regarding future payouts and prices. Thus, asset prices reflect a dynamic blend of expectations, discounting, and risk adjustments, continuously evolving as new information, economic events, and shifts in investor sentiment. To understand the temporal behavior of asset prices, it's essential to understand how expectations about both prices and payouts evolve over time. The conventional approach to this challenge relies on the assumption of rational expectations (RE), which implies that asset prices reflect all available information, forming the basis of the Efficient Market Hypothesis (EMH). Additionally, research in psychology and neuroscience reveals that only a small fraction of brain activity operates at the level of conscious reflection, with a substantial portion occurring unconsciously and being influenced by emotions. However, RE fails to adequately specify the process through which individuals form decisions based on these expectations. As a result, much of the asset pricing literature has focused on identifying specifications of preferences or technology that can account for sufficiently volatile risk premia Adam and Nagel, 2022.

Although rational expectations (RE) are analytically convenient, they represent a strong assumption. Whether or not this assumption is realistic is ultimately an empirical question. A growing body of recent research explores whether allowing investors' subjective beliefs to deviate from RE could help address some of the empirical challenges faced by RE-based asset pricing models. Developing models where investors price assets based on their subjective beliefs about future payouts and prices then requires making assumptions about how these subjective beliefs are formed. Without the strong link between beliefs and an underlying model of objective reality that RE requires, there are numerous possibilities Adam and Nagel, 2022. Attempting to reverse-engineer subjective beliefs from asset prices appears unappealing. It is likely that there are numerous belief formation mechanisms that are observationally equivalent when it comes to predicting asset price movements based on historical price data.

Expectations data, which measure how investors form beliefs about the future, are therefore essential in asset pricing models. The mechanisms used to form subjective beliefs should not only generate asset price behavior that aligns with empirical observations but also be plausible based on observable investor expectation data. With the growing availability of data from surveys and machine learning, research on investor expectations has become a highly active field of research. Mapping survey-based expectation data into agent expectations for asset pricing models is challenging due to measurement errors and belief heterogeneity. These challenges have sparked promising research opportunities, including the exploration of subjective risk perceptions, the aggregation of measured beliefs, and the relationship between asset market expectations and the macroeconomy. To develop asset pricing models that align with survey evidence on subjective belief dynamics, researchers have recently explored various approaches, including different ways in which agents update their beliefs and expectations about future outcomes based on new information related to payout or price dynamics, adaptive expectations, extrapolative expectations, and diagnostic expectations. Adam and Nagel (2022) argue that individual rationality, which implies agents updating their subjective beliefs using Bayes' rule and making optimal decisions based on their subjective beliefs about factors beyond their control, such as future market conditions, government policies, or technological changes, is less stringent than the RE assumption. Kamdar (2019) demonstrates, by using survey data in a model of rationally inattentive consumers facing fundamental uncertainty, that consumers' economic beliefs are primarily driven by one key factor: sentiment.

1.1.2 The Psychology of Tail Events

In traditional asset pricing models grounded in the assumption of investor rationality, investor preferences adhere to the expected utility paradigm, which comprises two key components: a utility function and a set of probability beliefs. Specifically, rational investors are assumed to utilize information efficiently, forming beliefs based on optimal statistical procedures, and their utility functions are concave in wealth, with this concavity capturing investors' aversion to risk. Specifically, expected utility is calculated by weighting the utility of different outcomes according to their respective probabilities. However, in many real-world situations, probabilities are unknown, leading to Knightian uncertainty or ambiguity, which necessitates assumptions about these probabilities. With expectations anchored in this manner, significant fluctuations in asset prices over time are primarily attributed to changes in risk premia, rather than shifts in anticipated future payouts and prices. The expected values of various decisions can be highly sensitive to these assumptions, particularly in cases involving long-tailed distributions, where expectations are heavily influenced by rare, extreme events.

Over the past decade, a large body of empirical research examined biased beliefs driven by sentiment through the lens of tail events. Much of the progress in understanding the psychology behind the tail events focused on the concept of probability weighting of cumulative prospect theory (Barberis, 2013a). The psychology of tail events can be understood through a two-step framework (Fox and Tversky 1998). In the first step, an individual assesses the probability of a tail event. In the second step, based on this probability judgment, a decision is made. The first step pertains to the formation of beliefs, while the second concerns preferences (Barberis, 2013a). On the beliefs side, a broad summary of the available evidence—though it oversimplifies many nuances—suggests that when asked to estimate the probability of a tail event, individuals generally tend to overestimate this probability. On the preferences side, the prevailing view is that if an individual is aware of a potential tail event, they will place greater weight on this outcome in their decision-making than the expected utility framework would suggest. This perspective is most closely associated with the "probability weighting function," a key element of Tversky and Kahneman (1992)'s cumulative prospect theory model of decision-making under risk. The weighting function transforms subjective probabilities into decision weights, and its main effect is to cause individuals to overweight the tails of any distribution they are considering. The original motivation for this was the simultaneous demand many people have for both lotteries and insurance.

1.1.3 Noise Trader Theory; Why Arbitrage Fails to Eliminate Mispricing

The study of sentiment-driven market behavior in finance originates from early noise trader theories, particularly those of Black (1986) and De Long et al. (1990), which challenge both the Rational Expectations Hypothesis and the Efficient Market Hypothesis (EMH). These theories contend that when some investors make trading decisions based on noisy, non-fundamental signals, asset prices can deviate persistently from their fundamental values, thereby undermining the assumption that markets fully and efficiently incorporate all available information. In traditional finance theories, there is no space for investor sentiment in the factors affecting stock returns. Friedman (1953b) and Fama (1965) argue that noise traders are not important in the process of asset price formation, as trades by rational arbitrageurs bring prices close to their fundamental values. However, persistent market anomalies, such as the underreaction and overreaction of stock prices and the closed-end mutual fund premium/discount puzzle, challenge the efficient market theory and raise questions about the extent to which arbitrage can eliminate the divergence between prices and fundamental values.

Noise traders are investors whose decisions are influenced by sentiment rather than fundamentals, leading to unpredictability to the market. The noise trader theory underscores that markets are not always rational. Behavioral finance incorporates this concept to explain real-world phenomena like the dot-com bubble, real estate bubbles and sudden market crashes, which are often fueled by irrational sentiment rather than changes in fundamentals. Since Black (1986) introduced noise as “expectations that need not follow rational rules”, noise traders are welcome in modelling financial markets as they provide liquidity and solve theoretical problems like the information-paradox formulated by Grossman and Stiglitz (1980). In financial markets, noise is essential because it provides the small, frequent fluctuations that keep trading activity alive. In a perfectly rational market, prices would stabilize quickly around intrinsic values, leaving little room for speculation or arbitrage. But because markets contain so much noise—driven by news events, investor behavior, and random economic data—prices are constantly in flux. This allows trading to thrive and markets to be observed dynamically, but it also introduces inefficiencies and makes it harder for investors to capitalize on any given inefficiency before it disappears.

The influential study by De Long et al. (1990) (DSSW) models the impact of noise trading on equilibrium prices, demonstrating that coordinated actions of noise traders based on non-fundamental signals introduce a priced, systematic risk. In their model, price deviations from fundamental values driven by investor sentiment are unpredictable. Arbitrageurs who bet against this mispricing face a risk, at least in the short term, that investor sentiment could become even more extreme, pushing prices further from fundamental values. The potential for loss and the arbitrageurs' risk aversion make them unwilling to take large positions. Consequently, arbitrage does not fully eliminate mispricing, and investor sentiment continues to influence security prices in equilibrium.

The noise trader model of De Long et al. (1990) motivated empirical research aimed at verifying the hypothesis that "noise trader" risks influence price formation. Since closed-end fund shares are primarily held by individual investors, Lee et al. (1991b) infer that fluctuations in closed-end fund discounts serve as a proxy for changes in investor sentiment. They find that changes in closed-end fund discounts are highly correlated with the returns of small-cap stocks, which are predominantly held by individual investors. Neal and Wheatley (1998) also find that larger closed-end fund discounts predict higher small firm returns and note that net redemptions reflect investor sentiment within closed-end fund discounts. Surprisingly, another commonly used investor sentiment indicator—the ratio of odd-lot sales to purchases—appears unable to predict returns for small or large companies. Lee et al. (1991b) showed that fluctuations in closed-end fund discounts closely correlate with the returns of small-cap stocks primarily held by individual investors, making it a proxy for shifts in investor sentiment. Neal and Wheatley (1998) also find that larger closed-end fund discounts predict higher small firm returns; they find that net redemptions capture sentiment shifts in closed-end fund discounts, but note that the odd-lot sales-to-purchases ratio, another popular sentiment indicator, shows no predictive power for small or large firm returns. In contrast to De Long et al. (1990), Bhushan et al. (1997) demonstrate that investor myopia is neither required nor sufficient to cause noise in asset prices.

Similarly, Bodurtha et al. (1995) reported that changes in the country fund discount reflect a previously unidentified risk factor, which they attributed to U.S. investor sentiment. Kelly (1997), using household data, showed that the likelihood of an individual becoming a noise trader decreases with rising income; that is, the high participation of low-income households (noise traders) in the stock market is associated with the low participation of high-income households (smart money or informed traders). Additionally, high participation by low-income households (noise traders) leads to negative future returns. Further examining the predictability of short- and long-term returns based on survey-based and indirect sentiment indicators, G. W. Brown (1999) found weak evidence for short-term predictability, but sentiment was highly correlated with long-term (2-3 year) returns. They also noted the existence of institutional sentiment in addition to individual sentiment and challenged the conventional view that sentiment is primarily driven by individual investors and only affects small-cap stocks.

Additionally, Barberis et al. (1998) introduced a model of investor sentiment to explain the two pervasive phenomena identified in stock market: underreaction and overreaction. Evidence of underreaction shows that, over a period of approximately 1 to 12 months, security prices respond slowly to new information, meaning prices don't fully adjust to the new information immediately. As a result, news is only gradually reflected in prices, causing prices to exhibit positive autocorrelation over this period. In other words, current positive news has predictive power for future positive returns. Evidence of overreaction, on the other hand, shows that over longer time horizons—around 3 to 5 years—security prices tend to overreact to consistent patterns of positive or negative news. This means that securities with a long history of positive news often become overvalued, resulting in lower average returns afterward. Put differently, securities with sustained strong performance, however measured, reach extremely high valuations, but these valuations generally revert to the mean.

Jegadeesh and Titman (1993b) provide compelling evidence of underreaction in stock markets through their examination of U.S. stock returns. Their study demonstrates that over a six-month horizon, stock returns exhibit positive autocorrelation, meaning that stocks with strong past performance tend to continue outperforming, while underperformers tend to lag further. This phenomenon, often referred to as momentum, is interpreted as evidence of slow information diffusion, where new information is not immediately and fully reflected in stock prices. Similar to findings on post-earnings announcement drift, this behavior highlights market inefficiencies and the gradual adjustment of prices to reflect underlying fundamentals. Carhart (1997) extend the Fama-French Three-Factor Model to the Carhart Four-Factor Model by incorporating the momentum factor, which demonstrated that the momentum factor operates independently of the existing market risk, size, and value factors and enhances the ability of multifactor models to explain the performance of mutual funds.

These evidencse pose a challenge to the efficient markets theory, as it suggests that sophisticated investors may achieve excess returns by leveraging underreaction and overreaction without taking on additional risk. Specifically, underreaction creates a pattern where momentum investors see returns by buying stocks that have recently risen in price and selling those that have fallen, while overreaction gives rise to mean-reversion strategies, where investors capitalize on the tendency of stock prices to revert to their average values after extreme movements. Fama and French (1996) notably attempt to address this issue from the perspective of the EMH framework, proposing that their three-factor model can account for the long-term overreaction pattern but could not explain the persistence of short-term returns (underreaction). Furthermore, these findings also challenge behavioral finance theories, as early models have failed to successfully explain these facts. The critical issue is how to interpret the beliefs formed by investors that lead to price underreaction and overreaction.

According to De Long et al. (1990) and Shleifer and Vishny (1997), a significant reason of limits to arbitrage is that investor sentiment is partly unpredictable, creating at least a short-term risk that sentiment may intensify, pushing prices further from fundamental values. The noise trader risk may decrease arbitrage positions due to short-term losses, which is particularly significant when arbitrageurs are risk-averse, rely on leverage, or manage capital on behalf of others and losing client funds for underperforming. As such, arbitrage cannot fully eliminate mispricing and investor sentiment impacts the stock prices in equilibrium. In equilibrium, prices reflect a mix of fundamentals and sentiment-driven deviations. These deviations persist because arbitrage, which theoretically should correct them, is constrained by practical and psychological factors. Noise traders, driven by irrational sentiment, can push prices further away from their fundamental value, since sentiment is unpredictable, creating what is termed "noise trader risk." Therefore, the sentiment-driven price deviations would persist even if the model incorporates arbitrageurs, which aligns with the theoretical framework that arbitrage is limited due to noise trader risk. These constraints explain why deviations from efficient prices can persist. Even in well-functioning markets, psychological factors create barriers to the correction of mispricings, challenging the assumptions of traditional efficient market theory. This framework is foundational in explaining why real-world markets often exhibit anomalies inconsistent with the Efficient Market Hypothesis (EMH). For instance, bubbles and crashes are extreme examples of situations where sentiment-driven mispricing remains uncorrected due to limited arbitrage.

The studies mentioned above illustrate why deviations from efficient prices can persist and arbitrage fails to eliminate mispricing (De Long et al., 1990; Shleifer & Vishny, 1997). According to these studies, one of the key reasons for limits to arbitrage is that changes in investor sentiment are, to some extent, unpredictable. As a result, arbitrageurs betting on mispricing face the risk that, at least in the short term, investor sentiment could become even more extreme, causing prices to deviate further from their fundamental value. Due to the existence of this "noise trader risk," arbitrage positions can incur losses in the short run. When arbitrageurs are risk-averse, use leverage, or manage other people's funds and face the risk of losing assets under management in the event of poor performance, the increased risk of deepening mispricing reduces the size of positions they are willing to hold. Therefore, arbitrage fails to fully eliminate mispricing, and investor sentiment continues to influence security prices in equilibrium. In the model below, investor sentiment is indeed somewhat unpredictable, so even with the introduction of arbitrageurs, arbitrage remains limited.

Building on sentiment indexes, researchers explored the relationship between sentiment and market anomalies like underreaction and overreaction. Barberis et al. (1998) proposed a model in which investors' tendency to overreact to past trends could explain long-term reversals, while Hong and Stein (1999) linked underreaction to gradual information diffusion among investors. These models highlighted that investor sentiment could explain previously puzzling patterns, offering a sentiment-based perspective on anomalies that EMH could not fully address. Studies also delved into cross-sectional effects of sentiment, examining how sentiment impacts different asset classes or market segments. For example, Stambaugh et al. (2012) demonstrated that sentiment-driven mispricing could explain the abnormal returns associated with the well-known anomalies like the value effect, momentum effect, and size effect.

1.2 Literature Review

1.2.1 Behavioral Pricing Kernel and News Sentiment

Asset pricing models have undergone significant evolution over time. The origins of modern asset pricing research can be traced to the mean-variance framework introduced by Markowitz (1952) in modern portfolio theory, which conceptualizes expected return and risk in terms of the mean and variance of asset returns. Following this, the Capital Asset Pricing Model (CAPM), developed by Sharpe (1964b), Linter (1965), and Mossin (1966), builds on the mean-variance framework and establishes a linear relationship, using a single factor—the stock market index—to explain the returns on common stocks. However, this static version CAPM has some restrictive assumptions such as normal distribution, single-period, rational investors, which is difficult to justify on both empirical and academic grounds. Merton (1973) states that the static CAPM fails to account for the intertemporal hedging component of asset demand. In order to adjust the assumption of single-period, Merton (1973) proposed intertemporal capital asset pricing model (ICAPM) with stochastic investment opportunities based on continuous-time economic framework, which is a modified CAPM with multiple betas considered in multiple periods. Although the intertemporal extension is theoretically significant, it is difficult to test and implement in financial practices, because all of the state variables are not easily identified. Another important criticism for CAPM was that the market return cannot be adequately proxied by broad stock indexes, and thus the unique role of the market in the model is dubious (Roll, 1977). Since true market portfolios are unobservable, multifactor models were introduced by adding additional variables to the CAPM. Ross (1976) developed Arbitrage Pricing Theory (APT), which does not require a market portfolio specification. Instead, APT considers a range of macroeconomic indicators as factors that influence stock returns, capturing systematic risk.

Subsequently, Lucas Jr (1978) and Breeden (2005) developed a single-beta capital asset pricing model in a multi-good, continuous-time model with uncertain consumption-goods prices and uncertain investment opportunities, which is known as consumption capital asset pricing model (CCAPM). In the CCAPM, the expected excess return on any security is proportional to its beta (or covariance) with respect to aggregate real consumption alone, where aggregate real consumption is calculated by the price deflator for consumption expenditures in the National Income and Product Accounts

(Breeden et al., 2015). More importantly, the CCAPM provides a fundamental understanding of the relationship between wealth and consumption and an investor's risk aversion, which can be used to estimate the investors' expected return and how the consumption-driven price volatility can influence the return.

Moreover, given the limitation of the Markowitz model, Black (1992) apply Bayesian approach that integrates investors' views and market equilibrium to develop Black-Litterman model, where the expected returns of a portfolio are estimated by equilibrium risk premiums caculated by CAPM and investors' views. In response to the shortcomings of the CAPM and the CCAPM in explaining the size and value premia in the cross-section of expected stock returns, Fama and French (1993) introduced the size factor (SMB) and the value factor (HML) as additional sources of systematic risk premia, leading to the development of the three-factor model. Later, Carhart (1997) expanded on this by proposing a four-factor model that included a momentum factor based on one-year stock returns, adding to the three-factor model. To date, the literature has identified over 300 factors that influence stock returns (Harvey and Liu, 2014). Recently, the sources of risk and return in a portfolio extracting from alternative data are used to predict stock price movements and capture market inefficiencies or risk premia on algorithmic trading (Birbeck & Cliff, 2018; Cremonesi et al., 2018).

Among all these factor pricing models, the consumption-based CAPM stands out for its unmatched theoretical purity. Various factor pricing models, in which observed factors serve as proxies for aggregate consumption risk, can be interpreted as special cases within the broader framework of the consumption-based CAPM (Cochrane, 2009). In this sense, most asset pricing theory is about how to model marginal utility using observable indicators, and hence amount to alternative ways of connecting the stochastic discount factor to data. In the SDF-based CAPM model, the single factor is the return on aggregate wealth portfolio that is often proxied by a broad market index, the expected utility associated with any portfolio on mean-variance efficient frontier. The wealth portfolio including all existing assets such as financial assets, human capital, real estate, durable consumption goods, non-corporate businesses etc in the economy is a claim to all future consumption (Cochrane, 2009). In the SDF-based Fama-French three-factor model, the size risk and value risk factors are added to the market risk factor in CAPM for the proxies of marginal utility. The small stocks and value stocks deliver low returns when marginal utility rises during recessions when durable consumption falls (Yogo, 2006). Wealth is only valuable because it gives people access to more consumption, thus utility functions over wealth should eventually be defended as arising from a more fundamental desire for consumption.

Although the consumption-based CAPM is theoretically appealing, empirical tests of the model yield parameter estimates that are inconsistent with stylized financial facts Hansen and Singleton (1983) and Mehra and Prescott (1985). Hansen and Singleton (1982, 1983) demonstrated that the CCAPM fails to simultaneously account for both the time-varying nature of interest rates and the cross-sectional patterns of average returns on stocks and bonds using US data. Similarly, Mehra and Prescott (1985) highlighted the equity premium puzzle, which suggests that an unrealistically high level of risk aversion would be needed to explain the historically large outperformance of stocks over government bonds, given the observed Sharpe ratio of the stock market and the low volatility in consumption. A series of empirical tests shows that the standard CAPM and consumption-based CAPM returns have less explanatory power for the cross-sectional and time variation in expected stock returns, and the risk-aversion parameter is not always negative and the SDF is not always positive, which imply that investors are at least locally risk-seeking and the existence of arbitrage opportunities.

The poor empirical performance of the consumption-based model under rational expectation theory motivates a search for alternative asset pricing models, namely alternative functions for stochastic discount factor m = f (data). One line of alternative models seeks to adjust the inflexibility of the expected utility specification, recognizing that the shortcomings of the consumption-based model may arise from the restrictive functional form of the utility function. While the mainstream economics still often relies on expected utility (EU) and rational agent assumptions, the existence of expected utility is fragile with both the assumptions of return distribution and changes in the time and risk preferences and the mean and variance of the discount factors. In the standard time-additive and state-independent expected utility framework, the two distinct aspects of preference are intertwined, namely intertemporal substitution is the reciprocal of relative risk aversion. Particularly, for the power utility function, the curvature parameter γ simultaneously controls intertemporal substitution and risk aversion. However, aversion to a consumption stream that varies over time and aversion to a consumption stream that varies across states of nature are two distinct aspects of consumers' preferences, thus they should be estimated separately. Utility functions are typically represented in three dimensions, which encompass preferences for goods within a specific time period and state of nature, preferences regarding the allocation of goods over time, and preferences concerning uncertainty. These dimensions correspond to the intra-temporal elasticity of substitution, intertemporal substitution, and the second-order risk aversion parameter. For time preference parameter β, it is used to measure the degree of investors' impatience, whose empirical estimates tend to 1. The individual's subjective discount rate can not be estimated without knowing their risk preference.

More importantly, the standard consumption-based CAPM, along with its conditional versions, also fails to explain the cross-section of asset pricing. In response, part of the literature has turned to ad-hoc factor models, which offer a parsimonious summary of the cross-section of expected returns, or to production-based approaches that suggest a link between firm characteristics and expected returns. However, without imposing restrictions on investor preferences and beliefs, neither of these approaches can answer the fundamental question of why investors price assets the way they do Nagel (2012). Within rational asset pricing theory, research that imposes such restrictions has focused on the ICAPM, habit persistence Campbell and Cochrane (1999), and recursive non-(time) separable preferences Epstein and Zin (1989), long-run risk models, and the effects of frictions and liquidity risk. Although these preference specifications based on the rational expectations framework can partially successfully explain the equity premium puzzle, these rational asset pricing models have less explanatory power for the cross-section of stock returns. Specifically, rational asset pricing models continue to struggle in matching the explanatory power of empirically driven factor models that use firm-level characteristics as risk factors, such as firm size and book-to-market ratio, as highlighted in the studies by Fama and French (1992, 1993), momentum effects documented by Jegadeesh and Titman (1993b), and idiosyncratic volatility, as shown in Ang et al. (2006, 2009) and Chabi-Yo (2011).

While there are existing studies on investors' beliefs and preferences, most empirical asset pricing research operates within the framework of rational expectations. In traditional aggregation theory, individual investors are typically modeled as a representative agent with Bayesian beliefs and an exponential discount factor. However, it is more realistic to consider that heterogeneous beliefs and preferences lead to a representative investor with non-Bayesian beliefs, non-exponential discounting, and time-varying risk aversion. The anomalies may reflect an omitted factor problem rooted in systematic investor errors—specifically, those driven by sentiment. Model specifications that exclude sentiment-related factors significantly impair the estimation of SDF parameters, often yielding results that conflict with standard assumptions of risk aversion Potì and Shefrin (2014). Hens and Reichlin (2013a) suggest that incomplete markets, risk-seeking behavior, and incorrect beliefs can produce upward-sloping segments in the pricing kernel—features that may help resolve the equity premium puzzle. Furthermore, behavioral preferences such as reference points, aspiration levels, and rank-dependent expected utility contribute to the time-varying volatility and shape of the pricing kernel. While research consistently underscores the important role of sentiment in asset pricing (Smales, 2016), sentiment alone does not fully account for asset price dynamics (Chau et al., 2016). Although some models incorporate macroeconomic sentiment indicators, the integration of news-based sentiment into factor models remains underexplored.

1.2.2 Quantile Preference and Sentiment Driven Tails

Traditional asset pricing equilibrium models, such as the Sharpe-Lintner Capital Asset Pricing Model (Sharpe (1964a); Lintner (1965)), as well as more recent factor-based alternatives, are grounded in optimal portfolio decision-making under the Expected Utility (EU) framework. Within this context, the investor's optimal portfolio choice is highly dependent on how the utility function is specified to capture individual preferences. Typically, it is assumed that agents make decisions by maximizing expected utility in accordance with their risk preferences. However, these assumptions are often criticized for being overly restrictive and may fall short in accurately capturing the complexities of real-world investor behavior. To move beyond traditional expected utility models, researchers aim to introduce heterogeneity into dynamic economic models, assuming that agents maximize their future quantile utilities. The fundamental distinction between quantile utility and expected utility lies in how individuals make decisions under risk. Expected utility focuses on an overall averaging of outcomes based on probabilities and treats the entire distribution symmetrically. In contrast, quantile utility shifts attention to specific quantiles of the outcome distribution, such as the median or extreme percentiles. This model is particularly useful when agents care more about particular parts of the distribution, such as exetreme losses (focusing on lower quantiles) or extreme gains (focusing on higher quantiles). Unlike expected utility, quantile utility allows for asymmetric risk preferences, where decision-makers may behave differently across different quantiles, emphasizing tail risks rather than the average.

Quantile preferences (QP) have been explored through a growing body of theoretical and empirical work. In early work, Mendelson (1987) introduced the concept of the quantile-preserving spread, a measure of risk aversion within the quantile utility framework that mirrors the role of mean-preserving spreads in the traditional expected utility (EU) model. Manski (1988b) developed the decision-theoretic foundations of quantile maximization and examined the risk preferences with quantile utility maximization. Subsequently, QP have been formally axiomatized by Chambers (2009) and Rostek (2010). In the context of preferences over distributions, Chambers (2009) demonstrated that quantile preferences are characterized by the axioms of monotonicity, ordinal covariance, and continuity. Rostek (2010) provided an axiomatization of quantile preferences within the Savage (1954) framework by introducing a 'typical' consequence scenario. In particular, Giovannetti (2013) extends the conventional linear average asset pricing formulations to the quantile process. More recently, De Castro and Galvão (2019) proposed a dynamic framework for decision-making under uncertainty, wherein agents seek to maximize a sequence of future quantile utilities over time. de Castro, Galvão, Noussair, and Qiao (2022) present one of the few experimental investigations into quantile preferences (QP), wherein participants make pairwise choices between risky lotteries. Their findings suggest that behavior consistent with QP is exhibited by approximately 30% to 50% of the sample population. de Castro, Galvão, Montes-Rojas, and Olmo (2022) applied the quantile preference (QP) model to portfolio selection.

In addition, quantile preferences are closely associated with Choquet expected utility, as outlined by Chambers (2007) and Bassett et al. (2004), who explore how individuals make decisions based on specific quantiles. Bassett et al. (2004), using Choquet utility functions, demonstrate that pessimistic optimization can be formulated as a linear quantile regression problem. This approach provides a valuable alternative to expected utility and serves as a plausible complement for analyzing rational behavior under uncertainty, which accommodates rational expectations by allowing agents to form expectations that align better with real-world scenarios. This asymmetry can be linked to the behavioral biases in prospect theory by Kahneman and Tversky (1979a), where losses weigh heavier than equivalent gains in the decision-making process. Moreover, the method of Value-at-Risk (VaR), widely used in finance, can be seen as an instance of quantile-based decision-making Engle and Manganelli (2004).

Furthermore, quantile preferences relates to the extensive literature on quantile regression (QR) methods in econometrics. Since the seminal work of Koenker and Jr. (1978), a growing body of research has focused on estimation and inference in QR panel data models, which offer powerful methods for examining the heterogeneous impacts of policy variables. These models are particularly valuable in program evaluation in economics and statistics, in which they enable researchers to assess how treatments or social programs influence the distribution of outcomes across different quantiles. Koenker (2004) presents a general framework for estimating quantile regression models with panel data. By incorporating fixed effects to control for individual-specific heterogeneity while allowing heterogeneous covariate effects within the QR framework, this approach provides a flexible and robust method for analyzing panel data structures.

In finance, Engle and Manganelli (2004) were pioneers in applying quantile regression to create the Conditional Autoregressive Value-at-Risk (CAViaR) model, which captures the conditional quantiles of asset returns. Baur, Dimpfl, and Jung (2012) utilize quantile autoregressions to analyze conditional return distributions, Cappiello et al. (2014) identify comovements between random variables through time-varying quantile regression. Žikeš and Baruník (2016) demonstrate that different realized measures are effective in predicting future return quantiles without relying on assumptions about the underlying conditional distributions. The resulting semi-parametric modeling approach effectively captures the conditional quantiles of financial returns within the flexible framework of quantile regression. White et al. (2015) extend the research focus toward a multivariate framework, emphasizing the interrelationships between the quantiles of multiple assets. A different branch of the multivariate quantile regression literature focuses on factor-based analysis (X. Chen et al. (2016); Ando and Bai (2017)). Moreover, QR has been extensively applied in program evaluation (Chernozhukov and Hansen (2005); Firpo (2007)), the identification of nonseparable models (Chesher (2003); Imbens and Newey (2009)), nonparametric identification and estimation of nonadditive random functions (Matzkin (2003)), and testing models with multiple equilibria (Echenique and Komunjer (2009)). More importantly, De Castro and Galvão (2019) provide the microeconomic foundations of QR by formulating a dynamic optimization decision model that yields a conditional quantile restriction (Euler equation).

Within the framework of bounded rationality Conlisk (1996), investor sentiment plays a significant role in influencing asset prices Baker and Wurgler (2006), De Long et al. (1990), and Hughen and McDonald (2005). In this context, a growing body of empirical studies investigates the effects of investor sentiment on both asset returns and market volatility Baker et al. (2012), Black (1986), Bouri et al. (2022), Corbet et al. (2020), De Long et al. (1990), Frugier (2016), Li et al. (2017), Suardi et al. (2022), and Yarovaya et al. (2022).To measure investor sentiment, the literature employs a variety of proxies, broadly categorized into three main types. Firstly, market-based indicators such as trading volume, turnover rates, and dividend premiums are frequently utilized, reflecting the underlying investor behavior in financial markets Baker and Stein (2004). Secondly, sentiment surveys offer direct insight into investor expectations and mood, with widely used examples including the American Association of Individual Investors' sentiment index Lee et al. (2002), the UBS/Gallup survey Qiu and Welch (2006), and animusX sentiment data Lux (2011). Thirdly, an emerging strand of research leverages machine learning methods to extract sentiment from textual data, such as financial news and commentary, providing more dynamic and granular sentiment indicators Calomiris and Mamaysky (2019) and A. Shapiro (2022).

Although a substantial body of literature examines the impact of sentiment on the cross-section of returns, much of the existing literature concentrates on average sentiment effects, thereby neglecting how different parts of the return distribution react to shifts in sentiment—especially under extreme market conditions. Quantile preference models are particularly effective in capturing downside risk aversion and asymmetric risk perceptions, making them well-suited for explaining sentiment-driven investor behavior, such as overreactions to extreme market events. In a multivariate framework, the static quantile approach provides a flexible framework to model tail events within the conditional distributions through sentiment, delivering a clear interpretation of how various factors—including sentiment—affect the return distribution.

1.2.3 Time-varying Probability Weighting and Investor Sentiment Toward Tails

Although expected utility theory remains the dominant framework for modeling instrumental rationality, it exhibits significant empirical limitations. In particular, its assumptions of full rationality and linear probability weighting often fail to capture actual human behavior. Economists have noted that individuals are often more responsive to changes in the probabilities of extreme outcomes than to those of more common, intermediate events. The Expected Utility (EU) model, with its linear weighting of probabilities, fails to capture this pattern. This led to the development of non-expected utility models, including prospect theory (Kahneman and Tversky, 1979a), rank-dependent expected utility (RDEU)(Quiggin, 1982, 1993), and cumulative prospect theory (Kahneman and Tversky (1992)). A key component of these models is the probability weighting function.

Kahneman and Tversky (1979a) proposed the original version of prospect theory, which captured the theory's fundamental behavioral insights. However, the original version has notable limitations: it was restricted to analyzing gambles with no more than two nonzero outcomes and, in some cases, predicts that individuals will choose dominated options. To address these limitations, Tversky and Kahneman (1992) introduced a revised version of their theory—cumulative prospect theory—building on the framework of rank-dependent expected utility (RDEU) developed in 1982. This modified version is the one typically used in economic analysis, and it is the version briefly reviewed in this thesis. Under cumulative prospect theory, people's decision making is characterized by the four elements of prospect theory: 1) reference dependence, 2) loss aversion, 3) diminishing sensitivity, and 4) probability weighting.

First, reference dependence refers to the idea that individuals derive utility from gains and losses relative to a reference point, rather than from absolute levels of wealth. Second, the value function embodies the principle of loss aversion—the observation that individuals exhibit more sensitive to losses compared to gains of the same magnitude. Informally, this is modeled by making the value function steeper for losses than for gains, reflecting the disproportionate psychological impact of losses. Third, diminishing sensitivity is reflected in the curvature of the value function. Specifically, the value function is concave for gains indicating risk aversion for moderate-probability gains, and convex for losses indicating risk-seeking for moderate-probability losses, meaning that the marginal value of both gains and losses decreases with their distance from a reference point. The fourth and final component is probability weighting, which indicates that people assign weights to outcomes based on transformed probabilities that overweight the tails of the distribution rather than objective probabilities. These features typically produce an asymmetric, inverse S-shape of weighting function.

According to Barberis (2013b), the research progress in understanding the psychology of tail events has largely revolved around the concept of "probability weighting," particularly its applications across different fields of economics. Intuitively, the probability weighting function serves as a tool to model individuals' risk preferences towards probabilities of ranked events. Notably, emphasize that these transformed probabilities should not be interpreted as erroneous beliefs; instead, they represent decision weights—a reflection of how individuals subjectively evaluate the impact of probabilities in decision-making. Moreover, **barberis2008stocksaslotteries<empty citation>** argue that the primary effect of the weighting function is to overweight the tails of the distribution it is applied to, which is not a bias in belief but a modeling tool aimed at capturing the common preference for positively skewed return distribution.

The theoretical restrictions on the probability weighting function are quite minimal, leaving its shape to be empirically determined. Numerous experimental studies, for example Camerer and Ho (1994), Wu and Gonzalez (1996), and , show that people tend to overweight events in the tails of the distribution compared to those in the middle, which can be characterized by an inverse-S-shaped probability weighting function and a corresponding U-shaped density function. For instance, Camerer and Ho (1994) review evidence from various studies supporting the non-linear nature of probability weights with a higher sensitivity to tail events. Wu and Gonzalez (1996) identify the inverse-S shape of probability weighting in experimental data using a method that does not require specific assumptions about the utility or probability weighting functions' forms. Numerous other researchers have conducted experiments, both hypothetical and with real payouts, to explore probability weighting functions, consistently finding that the inverse-S shape aligns well with observed data, as seen in works by Quiggin (1987), Lopes (1987), and Prelec (1998). Additionally, Berns et al. (2007) provide insights from non-monetary experiments. Comprehensive surveys of theoretical and empirical literature on risk-related choice are available from

Shoemaker (1982), Camerer (1995), and Starmer (2000).

According to Polkovnichenko and Zhao (2013), most estimated probability weighting functions display an inverse-S shape, characterized by the overweighting of tail probabilities. By fitting a parametric form to nonparametric weight estimates, it is possible to construct a time series of the parameters that define its shape. While the inverse-S pattern dominates—reflecting an overweighting of tail events, there are periods in which the function shifts to an S-shape, indicating a tendency to underweight tail probabilities. These temporal variations in the shape of the weighting function offer a potential lens through which to measure investor sentiment toward tail events (Polkovnichenko & Zhao, 2013). Despite the theoretical link between probability weighting and investor sentiment, few studies to date have explored this relationship. The chapter 4 seeks to address this gap by examining how time-varying parameters of probability weighting functions correlate with investor sentiment toward tail risks.

As Barberis (2013b) observes, it is somewhat surprising that, prospect theory has seen relatively few prominent and widely accepted applications in economics, despite the decades since its seminal introduction by Kahneman and Tversky in 1979. One might therefore conclude that, although it captures behavior accurately in laboratory settings, its relevance beyond the lab is limited. Chapter 4 challenges that view by providing empirical support for probability weighting using non-experimental data and independently corroborates the strong correlation between the time-varying parameters of probability weighting functions and external investor sentiment measures. In doing so, it extends the reach of prospect theory well beyond the confines of experimental setting.

This thesis covers a broad spectrum and connects with multiple strands of literature in economic theory, econometrics and machine learning, focusing on the evolution of decision-making models under uncertainty in dynamic settings.

1.3 Summary

This chapter provides a comprehensive overview of the key challenges inherent in traditional asset pricing models and examines how both theoretical frameworks and empirical research in behavioral asset pricing have sought to address them. Specifically, it explores how incorporating investor sentiment into behavioral pricing kernels, quantile preference models, and the rank-dependent family of utility models offers a promising avenue for resolving persistent anomalies and improving the understanding of asset price dynamics.

Chapter 2

Behavioral Pricing Kernel and News Sentiment

2.1 Introduction

The parameter estimates in the empirical tests of CCAPM are inconsistent with financial stylized facts (Hansen and Singleton, 1983; Mehra and Prescott, 1985). Hansen and Singleton (1982, 1983) found the CCAPM cannot simultaneously explain the time-variation of interest rates and the cross-sectional pattern of average returns on stocks and bonds based on US data. Mehra and Prescott (1985) observed the equity premium puzzle, which implies an unrealistically high level of risk aversion to explain the excessively high historical outperformance of stocks over government bonds, given the observed Sharpe ratio of the stock market and low volatility of consumption. Although the preference specifications based on the rational expectations framework—including the lifetime recursive utility Epstein and Zin (1989), the backward-looking habit model Campbell and Cochrane (1999), and the forward-looking long-run risk model Bansal and Yaron (2004)—can partially explain the equity premium puzzle, they have not yet succeeded in the task of explaining the cross-section of stock returns.

One possibility is that the unconditional CCAPM fails to take the time variation of the SDF parameters into consideration (Lettau & Ludvigson, 2001b). A large number of empirical tests show that expected excess return on broad market indexes are forecastable, which indicates that risk premia is time-varying (Campbell and Shiller, 1988; Fama and French, 2021; Lamont, 1998; Lettau and Ludvigson, 2001a). According to the Habit persistence model developped by Campbell and Cochrane (1999), investors' expected risk premia changes with the time-varying changes in the economic conditions (business cycles), the investment opportunity sets, and their time preference and risk aversion, thus the discount factor should be a state-dependent function of the pricing factors such as consumption growth, market return. Based on the time-varying risk premia, the parameter in stochastic discount factor will depend on investor expectations of future excess returns.

The conditional versions of the CCAPM with time varying parameters are most empirically successful in pricing the cross-sectional variations of stock returns. For example, Lettau and Ludvigson (2001a) and Lettau and Ludvigson (2001b) show that the conditional version of a multi-factor model has more explanatory power in the cross section of stock returns. However, Lewellen et al. (2010) warn the estimates may be unreliable because of the flexibility resulting from allowing time varying parameters scaled with conditioning variables unless the preference restrictions are imposed on the pricing kernel model. In the paper by Lettau and Ludvigson (2001a) and Lettau and Ludvigson (2001b), both the risk-aversion parameter and the SDF are negative for certain sample realizations of the conditioning variable, which imply that investors are at least locally risk-seeking and the existence of arbitrage opportunities.

Therefore, there is a puzzling contrast in the conditional asset pricing between the high cross-sectional explanatory power of model and the in-consistence between the coefficient estimates the fundamental assumptions implied by the model. There is the possibility that an omitted factor problem might be behind this puzzle. The fit of model specifications without sentiment-related factors significantly hindered the SDF parameters estimates that are inconsistent with standard risk aversion assumptions. Poti and Shefrin (2014) propose a polynomial pricing kernel that augments with the Baker–Wurgler sentiment factor. Although the model partially reduced the negative values of the SDF and risk aversion, the extent is limited. This chapter tends to reconcile the inconsistence between the high cross-sectional explanatory power of the conditional asset pricing model and the fundamental assumptions by developing a behavioral pricing kernel with news sentiment.

In the behavioral pricing-kernel framework, sentiment is encoded in the kernel's shape— a feature derived from Shefrin's decomposition theorem (Shefrin, 2008). This relationship arises because the shape of the sentiment function mirrors how both behavioral preferences and cognitive biases are distributed across investors. For instance, correlated optimism and overconfidence among investors can produce a pricing kernel with oscillatory features and locally upward-sloping segments (Barone-Adesi et al., 2017). Accordingly, the time-varying shape of the pricing kernel tracks evolving patterns of cognitive errors and behavioral preferences. A central empirical challenge in this context is the accurate measurement of investor sentiment. The literature employs a wide range of proxies, including the Baker and Wurgler (2006) investor sentiment index, as well as various measures based on survey, textual, and media data. A. H. Shapiro et al. (2022) constructed news sentiment indicator from economic and financial news paper articles and show the news sentiment indicator is predictive of movements of survey-based measures of consumer sentiment. To capture informational content distinct from Baker–Wurgler-type market-based proxies commonly used in the literature (e.g., Shefrin, 2014), this chapter adopts the Shapiro–Sudhof–Wilson (2022) news-sentiment index as the primary sentiment measure.

Moreover, the choice of conditioning variable plays a key role on the conditional models. Given that stock price change over time, which must reflect changing conditional moments of something on the dynamic asset pricing relation, asset prices should be modeled as a function of investors' conditioning information. Although the unobservability of the conditioning information may make a conditional factor model untestable since it is impossiple for researchers to observe and incorparte all the conditioning information used by investors in their models, Cochrane (2009) and Lettau and Ludvigson (2001b) argue that using a conditional variable that summarises investors' expectation for excess returns can largely circumvent this issue. Santos and Veronesi (2006) state that instruments that are informative about expected returns or the state of the economy can be natural candidates of conditional information in empirical tests. The instrumental variables should be parsimonious due to power considerations in GMM estimation(Tauchen (1986)).

There are a set of popular predictive variables including dividend-price ratio, dividend yield, consumption-wealth ratio, value spread used in the empirical test of conditional factors models. According to Lettau and Ludvigson (2001a, 2001b), in a broad range of optimization models for consumer behavior, the log consumption–aggregate wealth (including both human capital and asset holdings) ratio has predictive power for expected returns on aggregate wealth, or the market portfolio.

The main challenge in implementing this idea lies in the fact that the consumption-aggregate wealth ratio, particularly its human capital component, cannot be directly observed.Lettau and Ludvigson (2001b) address this limitation by introducing an observable counterpart-defined as the cointegrating residual among log consumption (c), log asset (nonhuman) wealth (a), and log labor income (y)-which they denote as cay inLettau and Ludvigson (2001a). Specifically, log consumption and log aggregate wealth (including asset holdings and human capital) tend to move together in the long run, despite that they may deviate from each other in the short run based on changing expectation for future asset returns. Thus, the observable version of this ratio will be a good proxy for agent's expectations on aggregate wealth or market portfolio.

Lettau and Ludvigson (2001a, 2001b) show that, scaling the conditional consumption-based CAPM by cay substantially improves its ability to account for the well-known "value premium", which implies that that an asset's risk is not governed by its unconditional correlation with the model's underlying factor, but instead by its correlation that varies with the prevailing state of the economy. The study emphasizes that the conditioning variable is central to the empirical framework, as the linear factor model depends on investors' unobservable information sets, which are inherently unobservable to the econometrician. To identify a variable that summarizes investor expectations, the study draws on a core principle of forward-looking models—that agents' observable behavior reflects their beliefs about future outcomes. Moreover, compared with other commonly used variables like the dividend-price ratio, consumption can be considered as the dividend paid from aggregate wealth, which means the deviations in the consumption–aggregate wealth ratio represent expectations about the entire market portfolio, not just the stock market part of it. Accordingly, the cay_t

supplied by Lettau and Ludvigson (2001a, 2001b), which captures investors' expectation for returns on the market portfolio, will be used as the conditioning variable that drives the variation of the SDF parameters in this chapter.

Following Poti and Shefrin (2014), this chapter will study whether the CCAPM with sentiment factors, and thus allowing for a behavioural influence on asset pricing, can improve the explanatory power of the (C)CAPM while admitting risk premia estimates consistent with the assumption of the representative investor's preferences and thus with the underlying economic theory. Consistent with Shefrin (2008), the SDF can be decomposed into two terms, one being the behavioural component represented by sentiment and the other being rational component that is represented by aggregate marginal utility growth of consumption. Research has consistently demonstrated that sentiment plays a key role in asset pricing models Smales (2015b) and Allen et al. (2019). However, sentiment alone cannot fully explain asset prices (Chau et al., 2016). The estimation approach involves estimating a series of models in which the pricing kernel is modeled as a function of both an index of sentiment and a set of pricing factors. The analysis begins by estimating the traditional model, which is then compared to the behavioral specification that incorporates sentiment. All these models are nested within the more general behavioral pricing kernel framework developed by Shefrin (2008), allowing for a systematic evaluation of sentiment's impact on asset pricing. The performance of the approximate kernels is significantly improved when news sentiment is incorporated into the kernel.

Moreover, to complement the SDF approximated using a Taylor series, the second pricing kernel specification is constructed using a linear combination of Chebyshev polynomial basis functions, evaluated over two fundamental state variables: consumption growth and sentiment. In a model economy with a well-defined structural foundation, the asset pricing kernel can be represented as a continuous function of its state variables (Chapman, 1997a). This function can be approximated using orthogonal polynomials defined over those variables. The use of orthonormal polynomials offers a key advantage: while the resulting pricing kernel remains nonlinear in the state variables—thus capable of capturing nonlinear payoffs—it is linear in the model's parameters. This linearity yields closed-form solutions for the approximating coefficients, simplifying both estimation and interpretation. Furthermore, the polynomial terms naturally serve as factors in a linear asset pricing framework, thereby connecting the model to traditional factor-based approaches while preserving the flexibility needed to capture complex, state-dependent dynamics.

This chapter conducts an empirical test of the (C)CAPM by approximating the pricing kernel as a linear function of consumption growth and news sentiment. The proposed approach retains many of the desirable features found in nonparametric analyses while while mitigating several of their key limitations. Specifically, it approximates the unknown marginal utility function within a static framework using a combination of Taylor series expansion and Chebyshev polynomials, thereby offering both local accuracy and global stability in functional approximation. In addition, departing from the common assumption of constant parameters over time, the analysis allows for time variation in the parameters by introducing a scaling variable,consumption–wealth ratio. This approach captures potential shifts in the pricing kernel driven by changes in the broader economic environment.

The primary contribution of this chapter lies in formalizing the role of sentiment within the SDF framework. While sentiment has often been introduced informally into empirical asset pricing models—typically as a non-tradable pricing factor—its integration as a structural component of the SDF remains underexplored. Although some studies have augmented SDF models with market and survey-based sentiment indicators, few have systematically incorporated news-based sentiment measures. This chapter advances the literature by embedding sentiment, derived from economic news, directly into the pricing kernel, enabling a more rigorous evaluation of its influence on asset pricing. Secondly, this chapter employs Chebyshev polynomial approximations to model the pricing kernel as a flexible function of aggregate consumption growth and sentiment. This method exploits the flexibility of Chebyshev polynomials to capture nonlinearities and more pronounced oscillations driven by sentiment, while also accommodating time variation in risk aversion through the inclusion of instrumental variables. By incorporating sentiment alongside consumption growth, the model achieves a notable improvement in performance, underscoring the informational value of sentiment in asset pricing.

The remain of this chapter is organized as follows. Section 2 summaries the evolution of asset pricing models. Section 3 introduced the framework from Shefrin (2005) for decomposing the SDF into a linear factor model, along with the identification restrictions used to separate the sentiment-driven and (C)CAPM-related components of the SDF. Section 4 present the GMM estimation. Section 5 details the construction of data. Section 6 report the empirical findings, in this part, the performance of different versions of conditional factor models are compared. Section 7 concludes.

2.2 Literature Review

While the CCAPM offers a plausible economic interpretation, the parameter estimates obtained from empirical tests do not align with the financial stylized facts (Hansen and Singleton, 1983; Mehra and Prescott, 1985). In addition to the equity premium puzzle, the CCAPM also fails to explain the cross section of expected stock returns Mankiw and Shapiro (1986) and Breeden et al. (1989). One issue is the measurement and aggregation problems of aggregate consumption proxies Breeden et al. (1989). Specifically, the consumption data is limited to the measurement problems including the durables problem, infrequent reporting, sampling error (Breeden et al., 1989). To address this issue, D. P. Brown and Gibbons (1985) proposed a conditional model based on a static framework, where consumption is assumed to be equivalent to wealth. In this setup, the stochastic discount factor can be approximated as a function of aggregate wealth.

The weak empirical performance of the consumption-based asset pricing model has prompted the development of alternative asset pricing models, most of which can ultimately be interpreted as variations in the functional form that maps observable data to the pricing kernel (Cochrane, 2009). This section provides a concise overview of several such approaches, setting the stage for more detailed discussions and analyses in later chapters. One strand of these alternative models is to adjust the inflexibility of the expected utility specification, as the shortcomings of the consumption-based model might stem from the chosen functional form for utility. In the standard CCAPM approach, utility is additively separable over time in a single aggregate consumption good that lumps together durables with nondurables. However, in practice, the marginal utility of nondurable consumption can be affected by the durable goods influences the marginal utility of nondurable goods, today's marginal utility can be affected by leisure or yesterday's consumption. To capture these richer dynamics, researchers have proposed more general, non-separable utility functions that allow for a separation between risk aversion and the elasticity of intertemporal substitution. These include the parametric class of non-expected utility preferences (Kreps and Porteus, 1978), lifetime recursive utility (Epstein and Zin, 1989), backward-looking habit formation models (Campbell and Cochrane, 1999), and forward-looking long-run risk models (Bansal & Yaron, 2004). More recently, Delikouras and Kostakis (2019) introduced a novel utility specification—a single-factor consumption-based model featuring disappointment aversion combined with second-order risk neutrality—to explain the cross-section of expected stock returns.

A more fundamental question concerns which variables determine marginal utility, and one prominent approach—adopted by factor pricing models—is to model the discount factor directly as a linear function of variables that serve as proxies for aggregate marginal utility growth, rather than relying solely on consumption data. The consumption-based CAPM, where an asset's systematic risk is measured by its covariance with the marginal utility of consumption and the SDF is the marginal rate of substitution in consumption, has a degree of theoretical purity that is unmatched by other asset pricing models. A set of asset pricing models, corresponding to particular forms of the the marginal rate of substitution, are the special cases of consumption-based CAPM. In the SDF-based CAPM model, the single factor is the return on aggregate wealth portfolio that is often proxied by a broad market index, the expected utility associated with any portfolio on mean-variance efficient frontier. The wealth portfolio including all existing assets such as financial assets, human capital, real estate, durable consumption goods, non-corporate businesses etc in the economy is a claim to all future consumption (Cochrane, 2009). In the SDF-based Fama-French three-factor model, the size risk and value risk factors are added to the market risk factor in CAPM for the proxies of marginal utility. The small stocks and value stocks deliver low returns when marginal utility rises during recessions when durable consumption falls (Yogo, 2006). Wealth is only valuable because it gives people access to more consumption, thus utility functions over wealth should eventually be defended as arising from a more fundamental desire for consumption. In these asset pricing models, asset prices are all set by an expected utility maximizing representative household endowed with rational expectations. According to Cochrane (2009), consumption and marginal utility respond to news, which imply that consumption will fluctuate with any variable whose changes today can signal income changes tomorrow, by permanent income logic. This fact opens the door to predictive indicators: any indicator that predicts asset returns ("changes in the investment opportunity set") or macroeconomic variables is a candidate pricing factor.

Although these alternative models have had some success in addressing the equity premium puzzle, they generally fall short in explaining the cross-section of stock returns. Lettau and Ludvigson (2001a, 2001b) introduced a conditional multifactor model that demonstrates greater explanatory power in capturing cross-sectional return patterns. However, a notable issue arises: for certain realizations of the conditioning variable, both the estimated risk-aversion parameter and the SDF turn negative, implying that investors are at least locally risk-seeking and the existence of arbitrage opportunities. Therefore, there is a puzzling contrast in the conditional asset pricing between the high cross-sectional explanatory power of model and the in-consistence between the coefficient estimates the fundamental assumptions implied by the model.

There is the possibility that an omitted factor problem might be behind the puzzle. The fit of model specifications without sentiment-related factors significantly hindered the SDF parameters estimates that are inconsistent with standard risk aversion assumptions (Poti & Shefrin, 2014). Hens and Reichlin (2013b) proposed that incomplete markets, risk-seeking behavior and incorrect beliefs can induce increasing parts in the pricing kernel and can be seen as potential solutions for the equity premium puzzle. Moreover, a set of behavioral preference elements such as reference point, aspiration point, and rank-dependent expected utility serve to accentuate the volatility and shape of the pricing kernel over time. Research has consistently highlighted that sentiment plays a significant role in asset pricing models (Smales, 2015a). However, sentiment alone cannot fully explain asset prices (Chau et al., 2016). While some studies have enhanced asset pricing models with macroeconomic-based sentiment indicators, there has been limited exploration of factor models that incorporate news sentiment.

Another strand is the nonparametric models, exemplified by Bansal, Coleman, and Hansen (1993), Bansal and Viswanathan (1993), and Chapman (1997b), which relax the linearity assumption altogether. Despite their superior empirical performance in explaining cross-sectional variations of average returns compared to the static CCAPM, these approaches are constrained by the reliance on ad hoc specifications—either in the selection of priced factors or the formulation of the nonlinear functional form. Given the vast range of potential factors and nonlinear specifications, researchers retain substantial discretion in specifying the model to be examined. Moreover, the form of the pricing kernel derived from nonparametric approaches lacks grounding in first principles. Specifically, the nonlinear kernels explored in the nonparametric approaches do not emerge endogenously from well-defined assumptions about investor preferences or return distributions Dittmar (2002). These limitations of nonparametric approaches present significant challenges for empirical application. First, tests based on ad hoc assumptions may lack statistical power, as they neglect preference-based theoretical restrictions implied by structural models. Second, the flexibility of these approaches raises concerns about overfitting and data dredging, as highlighted by Lo and MacKinlay (1990) and Fama (1991). In contrast, the static CCAPM endogenously determines both its pricing kernel (linear) and its factor structure (the consumption growth), thereby avoiding criticisms related to arbitrary factor selection and functional form specification.

2.3 Approximating the Pricing Kernel

According to Cochrane (2009), the household's first-order conditions (FOC) for the consumption and portfolio choice problem can be summarized by the following the Euler equation:

$$E_t(m_{t+1}R_{t+1}) = 1 \tag{2.1}$$

where $R_{i,t+1}$ is the gross return on asset i, which can be seen as a payoff with price one, m_t represents the pricing kernel, E_t denotes the mathematical expectation operator conditional on information available at time t. Harrison and Kreps (1979) demonstrate that m_{t+1} represents a pricing kernel that prices all risky payoffs and remains nonnegative under no-arbitrage conditions. By assuming a representative agent, the pricing kernel can be expressed as a function of the intertemporal marginal rate of substitution (IMRS) between aggregate consumption at dates t and $t+1$, $m_{t+1} = \frac{U'(C_{t+1})}{U'(C_t)}$. Although this specification is theoretically attractive, significant attention has been paid to the measurement and aggregation issues in available aggregate consumption proxies (e.g., Breeden et al. (1989)). The underperformance of the consumption based model has motivated the search for alternative asset pricing models-that is, alternative functional forms $m = f$ (data) (Cochrane, 2009). While one of the intuitive responses to poor empirical performance of consumption-based model is to try different utility functions, the more crucial question is which variables drive marginal utility.

2.3.1 Taylor Expansions

Given the inconsistency between parameter estimates and financial stylized facts when assuming a standard utility function, a suitable representation for the representative agent's utility function remains elusive. To mitigate this issue, Dittmar (2002) approximates the pricing kernel as a nonlinear function of the return on aggregate wealth using a Taylor series expansion.

$$\begin{aligned} m_c = 1 &+ h_{1t}\frac{U''(C_t)}{U'(C_t)}R_{ct+1} + h_{2t}\frac{U'''(C_t)}{U'(C_t)}R_{ct+1}^2 \\ &+ h_{3t}\frac{U'''(C_t)}{U'(C_t)}R_{ct+1}^3 + \cdots \end{aligned} \tag{2.2}$$

where, $R_{ct+1} \equiv \frac{\Delta c_{t+1}}{C_0}$ is the growth on end-of-period aggregate consumption, the h_i terms, $i = 1, 23\ldots$, are non-negative expansion parameters. As demonstrated in Equation (2.3), the marginal rate of substitution can be approximated by a polynomial function of aggregate consumption in a static setting. One key challenge with the polynomial expansion is choosing the maximum order of the approximating polynomial. Bansal, Hsieh, and Yaron (1993) allowed the data to determine this point, but this data-driven approach can reduce power and risk overfitting, making the resulting kernel's economic interpretation less clear. Dittmar (2002) employs a more powerful method by using the preference theory developed by Arrow (1971) to guide the truncation process.

According to Arrow (1971), investor utility functions should exhibit non-satiation (NS), risk aversion (RA), and non-increasing absolute risk aversion (NIARA)-the latter being linked to prudence (Kimball (1990, 1993)). In this framework, non-satiation implies positive marginal utility ($U' > 0$) and risk aversion implies diminishing marginal utility ($U'' < 0$). Under these restrictions, the pricing kernel is linear with a negative coefficient on aggregate wealth returns, thereby nesting the static CAPM. Moreover, NIARA requires that the decline in marginal utility does not become steeper as wealth increases, i.e., $\frac{d\left(-\frac{U''(W)}{U'(W)}\right)}{dW} \leq 0$. A necessary condition for NIARA, as demonstrated by Arditti (1967), is that the third derivative of the utility function satisfies $U''' \geq 0$, reflecting an aversion to negative skewness. By combining this condition with a truncation of the series expansion after the quadratic term, the resulting pricing kernel becomes quadratic in aggregate wealth returns, aligning with the three-moment CAPM. These conditions implies that by imposing standard risk aversion on agents' preferences, the signs of the coefficients for the first three polynomial terms in a Taylor series expansion can be determined. Since preference theory does not dictate the signs of higher-order polynomial terms, it is assumed that these additional terms have negligible impact on pricing, implying that the covariance between returns and polynomial terms of order greater than three in aggregate wealth is zero. Truncating terms of order higher than the second, the polynomial pricing kernel $m_{c,t+1}$ in equation (2.2) can be rewritten more compactly as follows:

$$m_{c,t+1} = 1 + \theta_{1t} R_{ct+1} + \theta_{2t} R_{ct+1}^2 \cdots \tag{2.3}$$

where, the intercept θ_{0t} of the kernel is set to 1, $\theta_{1t} \equiv h_{1t} \frac{U''(C_t)}{U'(C_t)}$ and $\theta_{2t} \equiv h_{2t} \frac{U'''(C_t)}{U'(C_t)}$. $U' > 0$, $U'' < 0$, and $U''' \geq 0$ imply that $\theta_{1t} < 0$ and $\theta_{2t} \geq 0$ for all t. In the case of a first-order Taylor series approximation of the marginal utility, the utility function simplifies to the familiar power utility form with constant relative risk aversion (CRRA)$-\gamma$, $U(c) = \frac{c^{1-\gamma}}{1-\gamma}$, where the marginal utility ratio is $\frac{U''(c)}{U'(c)} = -\frac{\gamma}{c}$, the coefficient θ_{1t} becomes $-\gamma$ while the coefficients for higher-order terms are zero.

As shown in Equation (2.3), one notable advantage of approximating the pricing kernel with Taylor polynomials is that, although the estimated kernel is nonlinear in the underlying state variables—thereby capable of pricing nonlinear payoffs—it remains linear in the polynomial coefficients. This linear structure enables closed-form solutions for the parameters and effectively treats each polynomial term as an independent factor within a linear pricing model. Thus, the polynomial pricing kernel is inherently nested within the general conditional factor pricing models presented in Cochrane (2009), which is develped to address the empirical challenges of the consumption-based model by modeling marginal utility directly using observable variables. In this framework, the SDF is specified as a (conditionally) linear function of a set of proxies:

$$m = \theta_{0t} + \theta_{it}' f_{t+1}, \quad i = 1, \ldots, k. \tag{2.4}$$

Following Shefrin (2008), the stochastic process for the pricing kernel can be modeled as the sum of rational component and behavioral component:

$$\tilde{m}(x_i) = e^{\Lambda} m_c \cong (1 + \Lambda)(\theta_{0t} + \theta_t' f_{t+1}) \equiv m_c + m_s \tag{2.5}$$

where $m_c = \theta_{0t} + \theta_t' f_{t+1}$ is the rational component, $m_s = \Lambda(\theta_{0t} + \theta_t' f_{t+1})$ is the product of marginal utility growth and the log of the ratio between the representative investor's subjective and the objective pdf, representing the behavioral component. Λ is sentiment, a log-change of measure defined as $\ln\left(\frac{f_{pR}}{f_p}\right)$ in equation (1.8), which measures the percentage by which, in equilibrium, the market-based pdf exceeds the objective pdf. In essence, the pricing kernel is modeled as a function of both an index of sentiment, Λ, and a set of pricing factors, f.

Substituting equation (2.3) into equation (2.5), denoting Λ by s for simplicity, dropping terms of third and higher orders, the pricing kernel can be expressed as a conditional multi-factor model:

$$\begin{aligned}\tilde{m}(x_i) &\approx \left(1+\theta_{1t}R_{c,t+1}+\theta_{2t}R_{c,t+1}^2\right)(1+s_{t+1}) \\ &\approx 1+\theta_{1t}R_{c,t+1}+\theta_{2t}R_{c,t+1}^2+s_{t+1}+\theta_{1t}s_{t+1}R_{c,t+1}+\mathcal{O}\text{ (higher order)}\end{aligned} \tag{2.6}$$

Following Cochrane (1996), the conditional factor pricing models are evaluated by by explicitly modeling the unconditional parameters θ_{0t} and θ'_{it} in (2.4) as a linear function on a time t information variable z_t, which serves as a predictive variable for excess returns. Then the time-varying parameters in equation (2.4) can be specified as a linear function of the first lag of z_t.

$$\theta_{it} = \theta_i^{'} z_t \qquad i = 1, ..., k \tag{2.7}$$

The unconditional factor model implied by the equation (2.4) can then be expressed as follows,

$$m_t = \theta_0 + \theta' F_{t+1}. \tag{2.8}$$

where $F_t = f_t - \mathbf{E}[f_t]$ represents the demeaned factors, θ' represents the vector of preference parameters. Given that the mean of the SDF is not identified when the excess returns are used as test assets payoffs, as noted in Cochrane (2009), by setting the constant term normalization $\theta_0 = 1+\theta' E(f)$ instead of simply $\theta_0 = 1$, the resulting pricing kernel takes the form$m = 1 - \theta'[f - E(f)]$ with mean $E(m) = 1$,

Consider the excess returns on asset i are used as test asset payoffs, in the case of a kth order approximating polynomial in a single state variable, the zero price for excess turns R^e imply a system of unconditional Euler equations:

$$\begin{aligned}E[\hat{m}_{t+1}R_{i,t+1}^e] &= E[\hat{m}_{t+1}]E[R_{i,t+1}^e] + Cov(\hat{m}_{t+1}, R_{i,t+1}^e) \\ &= 0\end{aligned} \tag{2.9}$$

where $R_{i,t+1}^e$ for $i = 1, \ldots, N$ denotes the excess returns on N portfolios over risk free rate.

Based on (2.8), the Equation (2.8) can then be expressed as a linear factor model

$$\mathbf{E}\left[R_{it}^{e}\right] = -\frac{\text{Cov}(\hat{m}_{t+1}, R_{i,t+1}^{e})}{E[\hat{m}_{t+1}]} = -\theta'\text{Cov}(F_{t+1}, R_{i,t+1}^{e}) \tag{2.10}$$

where, the covariance $\text{Cov}(F_{t+1}, R_{i,t+1}^{e})$ measures how deviations of each factor from its mean correlate with the excess return on asset i, thereby quantifying the asset's sensitivity to each systematic risk factor. By demeaning the factors, the analysis isolates only the variations—rather than absolute levels—which are the true drivers of risk premiums in asset pricing.

The equation (2.10) is essentially a cross-sectional regression of expected excess returns on the covariance between returns and factors, which indicates that the expected premium on asset i is determined by its covariances with the underlying factors the risk prices is encapsulated in θ'. Since the factor loadings in beta representation pricing model is defined as $\beta_i = Var(F_{t+1})^{-1}Cov(F_{t+1}, R_{t+1}^{e})$, which can be interpreted as a N $\times$ K coefficient vector in the multiple time series regression of excess returns $R_{t+1}^{e})$ on the discount factors F_{t+1}, thus the linear factor model is equivalent to the familiar beta representation pricing model

$$E(R_{i,t+1}^{e}) = \left(\frac{\text{Cov}(F_{t+1}, R_{i,t+1}^{e})}{\text{Var}(F_{t+1})}\right)(-\text{Var}(F_{t+1})\,\theta') = \beta_i^{'}\lambda \tag{2.11}$$

where, β can be interpreted as the quantity of risk, and $\lambda = -Var(F_{t+1})\theta$ is the factor risk premium, which can be interpreted as the price of risk. Note that λ depends on the volatility of the discount factor and is the same for each asset, while β varies across assets. In essence, the equation (2.10) implies the familiar multi-beta pricing model that a security's return is a linear function of a set of factors, including higher order approximations for consumption growth and interaction term between consumption growth and sentiment, the higher order and interaction term capture the effects of the nonlinearity in the pricing kernel.

Then the SDF parameter in the unconditional model can be obtained:

$$\theta = -\lambda / Var(F_{t+1}) \tag{2.12}$$

As seen above, each scaled multifactor model corresponds to a conditional factor model from which the scaled multifactor model is derived. The time-varying SDF parameters can be obtained by the linear function on conditional variable, however, the correspondingly time-varying risk premia of fundamental factors can not be estimated without the further assumptions on the conditional covariance. The equation (2.10) indicates the factor risk price is determined by the SDF parameters b. The linear function of b on conditional variables presumes that fluctuations in b are primarily driven by fluctuations in risk premia and implies a linear prediction equation for future excess returns. Although these prediction equations have good predictive power on excess returns, the excess returns often take negative value. Since the risk premia inherits the properties of these linear prediction models, the sign of risk premia may change over time. While it is reasonable to obtain the nonnegative factor risk premia on average in the conditional factor model, this condition does not imply that the individual risk premia from the unconditional model in equation (2.10) should be nonnegative.

Based on (2.22) and (2.8), and dropping terms of third and higher orders, the unconditional version of the equation (2.6) is specified as

$$\begin{aligned} m_{t+1}^{3M(C)\text{CAPMS}} &= \theta_0' z_t + \theta_1' z_t R_{ct+1} + \theta_2' z_t R_{ct+1}^2 + \theta_3' z_t S_{t+1} R_{ct+1} \\ &= \underbrace{\theta_{00} + \theta_{01} z_t}_{\theta_0} + \underbrace{\theta_{10} R_{c,t+1} + \theta_{11} z_t R_{c,t+1}}_{\theta_1 R_{c,t+1}} + \theta_{20} R_{c,t+1}^2 + \theta_{30} s_{t+1} R_{c,t+1} \end{aligned} \tag{2.13}$$

As can be shown, the conditional linear factor model can be converted into an unconditional multi-factor model in which the additional factors are simply scaled versions of the original factors.

To evaluate the performance of the kernel in Equation (2.13) relative to alternative pricing kernel specifications, four scaled multifactor models and three familiar unconditional models-the static CCAPM, CAPM, and the Fama-French three-factor model-are estimated. The goal is to assess their explanatory power for the cross-sectional variation in returns of 25 size and book-to-market sorted portfolios.

$$m_{t+1}^{(C)\text{CAPM}} = \underbrace{\theta_{00} + \theta_{01} z_t}_{\theta_0} + \underbrace{\theta_{10} R_{c,t+1} + \theta_{11} z_t R_{c,t+1}}_{\theta_1 R_{c,t+1}} \tag{2.14}$$

$$m_{t+1}^{3M(C)\text{CAPM}} \underbrace{\theta_{10} R_{c,t+1} + \theta_{11} z_t R_{c,t+1}}_{\theta_1 R_{c,t+1}} + \theta_{20} R_{c,t+1}^2 \tag{2.15}$$

$$m_{t+1}^{(C)\text{CAPMS}} = \underbrace{\theta_{10} R_{c,t+1} + \theta_{11} z_t R_{c,t+1}}_{\theta_1 R_{c,t+1}} + \theta_{20} s_{t+1} R_{c,t+1} \tag{2.16}$$

$$\begin{aligned} m_{t+1}^{3M(C)\text{CAPMS+S}} &= \underbrace{\theta_{00} + \theta_{01} z_t}_{\theta_0} + \underbrace{\theta_{10} R_{c,t+1} + \theta_{11} z_t R_{c,t+1}}_{\theta_1 R_{c,t+1}} + \theta_{20} s_{t+1} R_{c,t+1} \\ &+ \theta_{30} R_{c,t+1}^2 + \theta_{40} s_{t+1} \end{aligned} \tag{2.17}$$

$$\begin{aligned} m_{t+1}^{3M(C)\text{CAPMS}-\text{cay}} &= \underbrace{\theta_{10} R_{c,t+1} + \theta_{11} z_t R_{c,t+1}}_{\theta_1 R_{c,t+1}} + \theta_{20} s_{t+1} R_{c,t+1} \\ &+ \theta_{30} R_{c,t+1}^2 + \theta_{40} s_{t+1} \end{aligned} \tag{2.18}$$

$$m_{t+1}^{\text{CCAPM}} = \theta_0 + \theta_1 R_{c,t+1} \tag{2.19}$$

$$m_{t+1}^{\text{CAPM}} = \theta_0 + \theta_1 R_{m,t+1} \tag{2.20}$$

$$m_{t+1}^{\text{FF3}} = \theta_0 + \theta_1 R_{m,t+1} + \theta_2 R_{SMB,t+1} + \theta_3 R_{HML,t+1} \tag{2.21}$$

2.3.2 Chebyshev Polynomials

The second pricing kernel specification is approximated by a linear combination of Chebyshev polynomial basis functions, which are evaluated over two key state variables: consumption growth and sentiment. This approach leverages the flexibility of Chebyshev polynomials to capture more pronounced, sentiment-driven oscillations and it also permits time variation in risk aversion.

$$\hat{m}(x_t, \theta) = \sum_{j=0}^{k} \theta_j T_j(x_t) \equiv T_t' \theta \tag{2.22}$$

Typically, using a degree k polynomial for each state variable, the total number of parameters of the basis functions becomes $(k+1)^d$, where d is the number of state variables. The Chebyshev polynomials of the first kind are defined on the domain $[-1, 1]$ by $T_n(x) = \cos(n \cos^{-1}(x))$which immediately yields $T_0(x) = 1$ and $T_1(x) = x$. Higher-order polynomials, such as $T_3(x) = 4x^3 - 3x$ and $T_4(x) = 8x^4 - 8x^2 + 1$ exhibit oscillatory behavior and display periodic-like patterns over the domain. In order to extend these desirable approximation properties to functions defined on a different closed interval $[a, b]$, a linear transformation is applied to map the original variable into $[-1, 1]$. Specifically, the transformation $x = \frac{2x_{t+1} - (a+b)}{b-a}$ converts the state variables from the interval $[a, b]$ into the standard Chebyshev domain. The underlying cosine function brings periodic and oscillatory behavior into the polynomial structure. This approach yields generalized Chebyshev polynomials that maintain the advantageous features of orthogonality and minimal approximation error while ensuring accurate function representation over the new interval.

To capture the impact of sentiment on the pricing kernel, an additional state variable is introduced that represents a scaled log-change of measure that transforms the objective density f_p into investor's subjective density f_{pS}. In this framework, the orthonormal polynomials are constructed using a two-dimensional tensor product of one-dimensional polynomials, as presented in Judd (1992):

$$\hat{m}(x_{1t}, x_{2t}) = \sum_{j=0}^{k} \sum_{l=0}^{k} \theta_{jl} T_j(x_{1t}) T_l(x_{2t}) \tag{2.23}$$

where x_{1t} is consumption growth and x_{2t} is sentiment. where x_{1t} denotes consumption growth from time $t-1$ to t and x_{2t} denotes consumption growth from time t to

$t + 1$. This approach can be extended to incorporate more than two state variables by constructing the kernel through higher-fold tensor products of the corresponding one-dimensional polynomials.

2.4 GMM Estimation

In empirical asset pricing, all econometric techniques focus on the parameters estimation, the standard errors calculation of the estimated parameters and the pricing errors. Cochrane (2009) states all of the techniques come down to one of two basic ideas: time-series regression or cross-sectional regression. It is a natural way to use Generalized Method of Moments (GMM) for the estimation test of the model based on that the linear factor model is a set of moment restrictions on asset returns. Compared to the conventional Fama–MacBeth two-stage OLS regression, GMM estimation offers several key advantages.

First, Following Shanken (1992) and Jagannathan and Wang (1996), the two-pass Fama MacBeth (FM) procedure suffers from a generated-regressor, or errors-in-variables (EIV), bias because the second-pass regressions treat the estimated betas as known. Specifically, the Fama-MacBeth (FM) procedure treats first-stage betas as known in the cross-sectional regression. With few time-series observations, this generated-regressor (errors-in-variables) problem induces attenuation bias and noisy inference, even after applying Shanken-type adjustments; weak or unstable betas further exacerbate finite-sample distortions. In contrast, the two-step GMM approach estimates the risk premia directly from the Euler equation moments, thereby avoiding beta-estimation error, exploiting the full return covariance structure through an optimal weighting matrix, and providing an omnibus Hansen J-test of model validity.

Second, OLS relies on the factors being exogenous, a condition often violated in asset pricing, leading to biased estimates. In contrast, GMM does not strictly require exogeneity of the regressors because it can use appropriately chosen instruments and impose moment restrictions that account for any correlation with the errors. Third, GMM estimates can tackle the challenge from the presence of heteroscedasticity and autocorrelation by employing Newey-West estimator of the optimal weighting matrix , which results in more efficient and reliable inference even when error variances are time-varying or serially correlated. The two-step Fama–MacBeth estimator effectively corrects for cross-sectional correlation in asset-pricing tests (Fama and MacBeth (1973)), but it does not address time-series dependence in the residuals. To handle autocorrelation in panel regressions, Petersen (2009) recommends using Newey–West–corrected standard errors in the second-stage estimates. Finally, GMM enables more flexible model specifications by accommodating nonlinearities in the moment conditions, eliminating the need to linearize the SDF or assume log-normality, and thereby allowing for the capture of more complex dynamics. In this section, the first-stage GMM approach with an identity weighting matrix developed by Hansen and Singleton (1982) is used to estimate the preference parameters of the competing asset pricing models.

Using the Taylor series approximation with time-varying coefficients, equation (2.6), the Euler equation (2.1) can be expressed as

$$\mathrm{E}\Big[\Big(R^e_{i,t+1} \otimes z_t\Big)\hat{m}_{t+1}\Big] = 0. \tag{2.24}$$

where $R^e_{i,t+1}$ for $i = 1, \ldots, N$ denotes the excess returns on N portfolios over risk free rate, z_t is the $l \times 1$ vector of instrumental variables observed at time t and capture the relevant state information.

The covariance representation in (2.9) allows the coefficients θ to be directly interpreted as preference parameters, unlike the coefficients λ of the beta representation in (2.10), thus the moment conditions can be constructed from the equation (2.9) as follows:

$$g_T(\theta) = E_T[u(\theta)] = -E_T\left[z \otimes (\hat{m}R^e)\right] = -E_T\left[z \otimes R^e\right] + E_T\left[z \otimes R^e F'\right] \tag{2.25}$$

where $F' \equiv f - E(f)$, $E_T(\cdot)$ denotes the sample average over T time series observations and N represents the number of assets as test payoffs. the instrument variable z_t is chosen to align with the conditioning variables. By taking the sample average $E_T(\cdot) = \frac{1}{T}\sum_{t=1}^{T}(\cdot)$ of these moment conditions, the moment condition tests whether, on average, the pricing errors scaled by the instruments are zero. The Kronecker product $\otimes$ expands the pricing error across each instrument, thereby generating a set of moment restrictions, which jointly enforces that there is no systematic mispricing given the information available. The polynomial approximation's elements are scaled by the instrumental variables z_t, it is effectively turning the constant coefficients of the unconditional polynomial approximation in (2.12) into time-varying ones, which is equivalent to enabling the θ coefficients to vary over time as functions of these instruments, as shown by Shanken (1991) and Cochrane (1996).

When the moment conditions used for estimation are biased or inconsistent due to endogeneity, instrumental variables provide a way to "clean" the variation in the regressors, which helps construct moment conditions that are valid for identification. By instrumenting for the problematic state variables with z_t, one can obtain estimates of the time-varying parameters that are less likely to suffer from bias caused by the correlation between the regressors and the error term. cay from (Lettau and Ludvigson (2001a)) is used as an instrumental variable because it embodies key economic information about consumers' liquidity and wealth, correlates with expected returns, and—crucially—is assumed to be exogenous relative to the error term in the asset pricing equation. This makes it an effective tool for addressing endogeneity, thereby helping to achieve unbiased and consistent parameter estimates in the conditional asset pricing framework.

Equation (2.25) thus represents a system of $N \times K$ equations, where K is the moment conditions employed in the model. The total number of parameters, p, that need to be estimated is determined by the restrictions imposed in each model. For instance, for a linear approximation, $p = 2K$, for a quadratic approximation, $p = 3K$. As shown in Yogo (2003), the GMM estimator constructed from these covariance-based moment conditions is equivalent to an estimator of risk prices derived from a cross-sectional regression. Following Cochrane (1996), each individual element $R_{i,t+1} \cdot z_t$ of $R \otimes z$ can be viewed as the return on a "managed portfolio" where investment return in asset i is driven by the observable state variable z. For notational convenience, denote $R_z = R \otimes z$.

Exploiting the linear structure in the parameters, the first stage estimate of θ can be obtained by minimizing the quadratic objective function:

$$J(\theta) = \mathbf{g}_T(\theta)'\mathbf{W}\mathbf{g}_T(\theta) \tag{2.26}$$

where $g_T(\cdot)$ denotes the $(nk \times 1)$ vector of sample orthogonality conditions, and W is an $(nk \times nk)$ weighting matrix, which determines how much each moment condition contributes to the objective function. The overall adequacy of the approximate kernel can be measured by the J statistic. Specifically, the sample orthogonality conditions in equation (2.25) can be viewed as the pricing errors derived from the approximate kernel. When the model specifications are consistent with the data, these pricing errors will be nearly zero, leading to the minimized value in equation (2.26). Hansen and Singleton (1982) shows that $J_T = T \times J\left(\hat{\theta}_{\text{GMM}}\right)$ converges to a χ^2_{q-p} distribution (with $q = nl$ and $p = \dim(\theta)$), giving the familiar "J-test" of overidentifying restrictions. Intuitively, J_T tests the magnitude of a weighted average of all pricing errors. However, a low J_T value can arise for two very different reasons: either the pricing errors are small and precisely estimated, or they are relatively large but come with even larger estimated standard deviations.

The various GMM approaches hinge on the specification chosen for the weighting matrix, which determines both their efficiency and robustness. In the simplest one-step GMM one simply sets $W = I$ yielding a consistent but generally inefficient estimator.

$$\hat{\theta}_1 = \text{argmin}_\theta \; g_T(\theta)' W g_T(\theta) \tag{2.27}$$

Hansen and Singleton (1982) demonstrates that the optimal weighting matrix is the inverse of the long-run covariance matrix of the moment conditions, which, by design, accounts for heteroskedasticity and correlations among the moments. In the first stage, the identity matrix is commonly chosen, ($\mathbf{W} = \mathbf{I}$), which simplifies computation by giving equal weight to all moment conditions. This initial estimate, while consistent, is not necessarily efficient, however, it provides a useful starting point for constructing an optimal weighting matrix in a subsequent estimation step. The first stage estimate $\hat{\theta}_1$ can be used to form an estimate $\widehat{S}$ of the spectral density matrix (long-run covariance matrix), $S \equiv \sum_{j=-\infty}^{\infty} E\left[u_t(\theta) u_{t-j}(\theta)'\right]$, summarizes the variability and serial correlation of the moment conditions $u_t(\theta)$ over time, thus the second stage estimate of θ can be obtained by minimizing a quadratic form in the sample moments as follows:

$$\hat{\theta}_2 = \text{argmin}_\theta \; g_T(\theta)'\widehat{S}^{-1}g_T(\theta),$$
$$\text{var}\left(\hat{\theta}_2\right) = \frac{1}{T}\left(d'S^{-1}d\right)^{-1}, \quad d \equiv \frac{\partial g_T(\theta)}{\partial \theta}. \tag{2.28}$$

where the Jacobian matrix d, the derivative of the sample moment conditions with respect to the parameters, measures how sensitively the moment conditions respond to changes in the parameters, which is critical both for parameter identification and for determining the estimator's asymptotic variance. By capturing the entire covariance structure of the moment conditions, the matrix S enables the estimator to properly account for heteroskedasticity and autocorrelation, leading consistent and efficient estimates. The scaled estimation error, $\sqrt{T}\left(\hat{\theta}_1 - \theta_0\right)$, converges in distribution to a normal random variable with mean zero and asymptotic variance, $N(0, \text{var}(\hat{\theta}))$. Under standard regularity conditions, the scaled estimation error converges in distribution to a normal random variable with mean zero and asymptotic variance var $\left(\hat{\theta}_1\right)$, $\sqrt{T}\left(\hat{\theta} - \theta\right) \xrightarrow{d} N\left(0, \text{var}\left(\hat{\theta}_1\right)\right)$. Consider $d = -\frac{\partial g_T(b)}{\partial b'} = E\left(R^e f'\right)$, the second-moment matrix of returns and factors, the first-order condition to $\min g_T' W g_T$ is $-d'W\left[E_T\left(R^e\right) - db\right] = 0$, thus the first stage GMM estimates can be written as $\hat{\theta}_1 = \left(d'd\right)^{-1} d'E_T\left(R^e\right)$, and the second stage estimates can be written as $\hat{\theta}_2 = \left(d'S^{-1}d\right)^{-1} d'S^{-1}E_T\left(R^e\right)$.

2.5 Data

The data consist of observations on consumption, portfolio returns, news sentiment, and the instrumental variable cay used as conditioning information. The 25 size and book to market sorted portfolios from Fama-French are used as test assets payoffs, which is collected from Kenneth French website. The returns on the three-month Treasury bill rate is used as risk-free rate, which is obtained from CRSP. The consumption data and the consumption–wealth ratio estimates cay_{t-1} used as the conditioning variable are collected from Lettau and Ludvigson (2001a). The news sentiment indicator developed by A. H. Shapiro et al. (2022), based on economic and financial newspaper articles, is employed in this chapter as a proxy for investor sentiment, denoted s_t. To mitigate potential endogeneity arising from contemporaneous interactions between sentiment and stock returns, the first-order lag of the sentiment indicator is used in the analysis. Although stock returns are often modeled at a monthly frequency, consumption and labor income data are only reliably available on a quarterly basis. Measuring cay_t quarterly therefore ensures a more accurate estimation of the cointegrating relationship among c, a, and y, as monthly consumption figures tend to be noisy or interpolated. Moreover, the quarterly specification produces stable and economically meaningful deviations from the long-run equilibrium, capturing persistent shifts in investor expectations rather than transitory fluctuations driven by short-term noise. In essence, quarterly cay_t isolates the slow-moving component of expected returns, providing precisely the type of information relevant for scaling factors or conditioning instruments in macro-finance models. The data period starts from the first quarter of 1980 to the third quarter of 2019, including 159 observations.

2.6 Empirical findings

For each model in the cross-sectional regression of average retruns, the fitted expected excess returns for each of the 25 portfolios versus their sample portfolio premia are plotted. As illustrated in Figure 2.1, the fitted risk premia from the 3M(C)CAPMS+S specification-augmented with both the interaction term $R_c \cdot s$ and the standalone sentiment component s as in model (9)-align along the 30-degree line even more closely than those produced by the Fama-French three-factor model, indicating superior cross-sectional explanatory power.

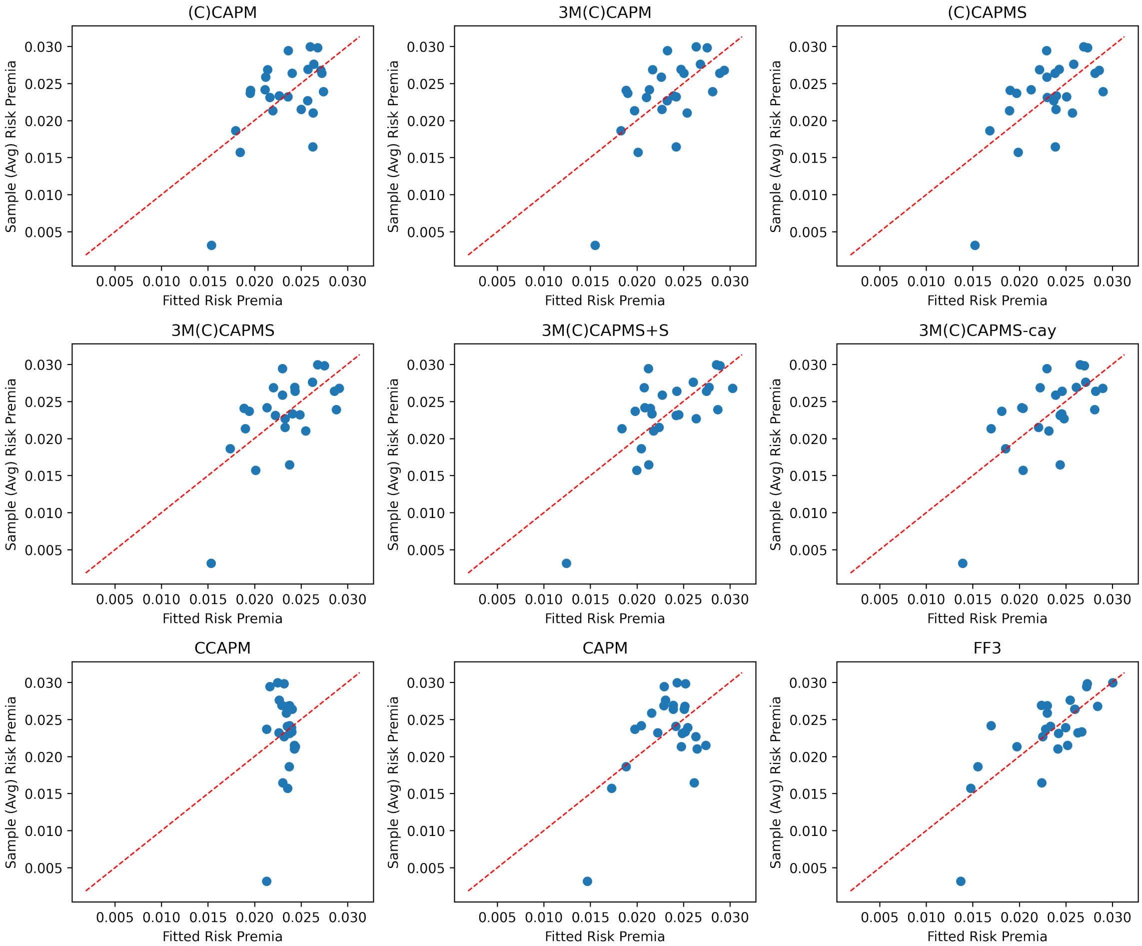

Figure 2.1: Predicted Mean Excess Return VS Mean Excess Return

Table 2.1 reports the the estimated factor risk premia (in percentage) and the corresponding Wald t-statistics for 25 size and book-to-market portfolios across a variety of pricing kernel models over the period 1980-2019. The Wald t-statistics are estimated using the optimal GMM weighting matrix. To establish a benchmark for comparison, results from the familiar unconditional models are first presented. In the static CAPM specification, the estimated beta on the market return is statistically significant but carries the incorrect sign, highlighting a fundamental empirical weakness of the model relative to the predictions of CAPM theory. In the unconditional Consumption-based CAPM (CCAPM), the t-statistic shows that the beta on the consumption growth is not statistically significant, reflecting the model's limited ability to account for cross-sectional return variation. This is further reflected in the model's poor explanatory power, with an R^2 of just 2 percent, and a negative adjusted R^2, underscoring the model's failure to account for cross-section of average returns. In contrast, the conditional version (C-CAPM), scaled by the consumption-wealth ratio (cay), though the Rc factor is insignificant, substantially improves explanatory power, accounting for approximately 36 percent of the cross-sectional variation in average returns.

Table 2.1:

The risk premia λ estimates 25 size and book-to-market portfolios beta-pricing representation (1980–2019).

	(C)CAPM	3M(C)CAPM	(C)CAPMS	3M(C)CAPMS	3M(C)CAPMS+S	3M(C)CAPMS-cay	CCAPM	CAPM	FF3	FF3S
constant	2.9836	3.2779	2.7948	3.3352	2.6766	2.2665	2.5751	4.2162	5.5078	5.3968
	(7.32)	(4.69)	(5.67)	(4.84)	(6.52)	(5.62)	(5.88)	(4.84)	(4.23)	(4.19)
cay	-0.5592				-0.3444					
	(-3.87)				(-1.39)					
Rc	-0.1824	-0.0397	-0.1853	-0.2082	-0.3929	-0.4342	-0.1313			
	(-0.91)	(-0.19)	(-0.62)	(-0.72)	(-1.35)	(-1.5)	(-0.51)			
cayRc	-0.0035	-0.0056	-0.0025	-0.0052	-0.0042	-0.0047				
	(-1.46)	(-1.89)	(-1.07)	(-1.86)	(-1.72)	(-1.71)				
R_c^2		-0.0008		-0.0004	0.0012	0.0020				
		(-0.41)		(-0.23)	(0.68)	(1.58)				
sRc			0.0247	0.0772	0.1288	0.1466				
			(0.29)	(0.9)	(1.33)	(1.6)				
s					1.8103	4.3106				-3.6954
					(0.55)	(1.08)				(-1.24)
Rm								-1.7569	-3.3609	-3.2482
								(-2.13)	(-2.63)	(-2.57)
SMB									0.0094	0.0371
									(0.05)	(0.2)
HML									0.7013	0.6777
									(3.39)	(3.25)
R^2	0.37	0.40	0.41	0.41	0.53	0.46	0.20	0.30	0.57	0.59
R^2 adj	0.25	0.28	0.29	0.26	0.34	0.27	0.16	0.24	0.49	0.48

Notes: This table presents the factor risk premia estimates in percentage and their corresponding t-statistics in brackets.

The coefficient on the scaling variable $\widehat{\text{cay}}$, which represents the time-varying component of the intercept, is positive and statistically significant across models. Notably, a comparison between the $3M(C)CAPMS+S$ and $3M(C)CAPMS-cay$ models reveals that omitting the cay term from the cross-sectional regression significantly affects both the statistical significance of other coefficients and the overall model fit. This result contrasts with the findings of Lettau and Ludvigson (2001a, 2001b), who report that excluding the scaling variable does not qualitatively affect their estimates. In addition, the interaction term cayRe is also consistently significant across the models, which suggests that, in the current framework, the scaling variable cay plays a more substantial role in shaping cross-sectional return variation, particularly through its interaction with other pricing factors. Obviously, the conditioning variable improves the fit of the CCAPM, which implies that the correlation between stock returns and consumption growth change over the economic conditions.

Table 2.2

	(C)CAPM	3M(C)CAPM	(C)CAPMS	3M(C)CAPMS	3M(C)CAPMS+S	3M(C)CAPMS-cay	CCAPM	CAPM	FF3	FF3S
cay	98				60					
Rc	53	12	54	61	115	127	38			
cayRc	9927	15697	7151	14597	11778	13336				
R_c^2		1326		680	-1903	-3172				
sRc			-95	-298	-497	-566				
s					-1	-2				1
Rm								3	5	5
SMB									-0	-0
HML									-2	-2

Notes: This table presents the parameters $\boldsymbol{\theta}$ vector of the unconditional pricing kernel, the negatives of the estimated risk prices, implied by the λ

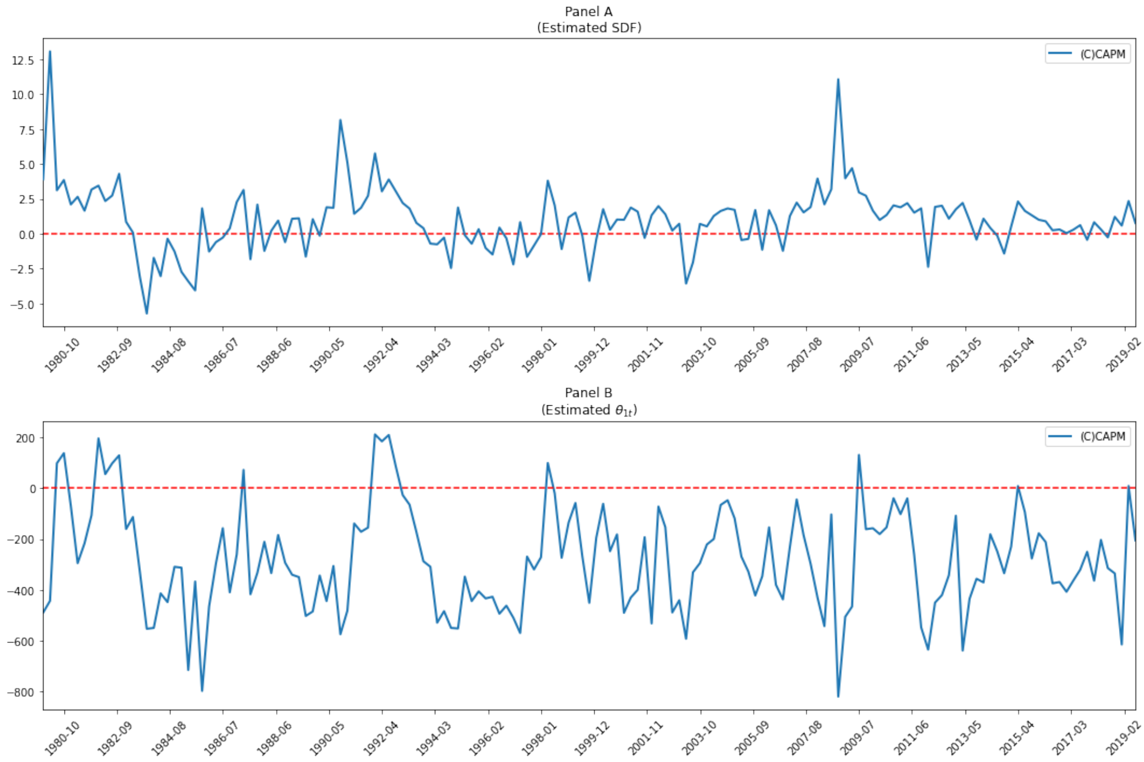

Figure 2.2

Panel A of this figure reports the SDF implied by GMM estimates of the (C)CAPM. Panel B reports the estimate of the conditional θ_{1t} parameter for the same model

Moreover, the estimated parameters θ imply that both the SDF and the time varying risk aversion θ_{1t}, frequently take negative values for prolonged portions of the sample period. As Panel A of Figure 2.2, the SDF often falls below zero, thereby assigning negative prices to payoffs that are almost surely nonnegative, which implies arbitrage opportunities and violates the non-satiation assumption imposed by preference theory. In addition, Panel B of Figure 2.2 shows that θ_{1t} occasionally takes non-negative values, thereby contradicting the fundamental assumption of risk aversion. Interestingly, the time-varying risk-aversion coefficient θ_{1t} becomes positive roughly during episodes when anecdotal evidence points to overoptimistic market forecasts of future earnings growth, such as the early 1990s, around 2000 at the onset of the tech-stock correction—and then turns positive again in the aftermath of the 2008 financial crisis, before gradually reverting below zero, roughly at the beginning of the stock market correction. The estimates of the 3M(C) CAPM exhibit similar issues, namely that the SDF is not consistently positive, and the parameter θ_{1t} does not remain uniformly negative as theoretically expected. The findings reveal strong evidence that risk aversion behaves counter cyclically across different states of the $S\&P$ 500 return. Specifically, it is observed that empirical risk aversion increases during economic downturns and declines in expansions. More importantly, compared with the SDF produced by Poti and Shefrin (2014), the SDF in this chapter takes rarely negative values, which is probably due to the choice of variables in different timeframe and the effect of intercept in the second stage regression.

When the news sentiment is incorporated to the kernel, (C)CAPMS and its three-moment extension 3M(C)CAPMS demonstrate significantly improved empirical performance relative to their non-sentiment counterparts, with most of the gain in explanatory power attributable to the inclusion of the sentiment component. Specifically, the R^2 of both the (C)CAPMS and 3M(C) CAPMS models increase to 0.41 . While these gains appear modest, they coincide with statistically significant coefficients on the interaction term between consumption growth and sentiment, captured by the cross-product $s_{t+1}R_{ct+1}$. This indicates that sentiment-adjusted marginal utility plays an economically meaningful role in explaining cross-sectional return variation. Most importantly, among the family of scaled multifactor models, the 3M(C)CAPMS+S specification augmented with both the interaction term $(R_c \cdot s)$ and the standalone sentiment component s delivers the strongest empirical performance, achieving an R^2 of 0.53. This explanatory power is comparable to that of the

Fama-French three-factor model, underscoring the value of incorporating sentiment directly into the pricing kernel.

The last column of Table 2.1 reports the results for the sentiment-augmented Fama-French three-factor model (FF3S). Relative to the standard FF3 specification, the inclusion of the sentiment factor produces a modest increase in explanatory power (R^2 rises from 0.57 to 0.59), suggesting that sentiment contributes additional information beyond the traditional size, value, and market factors. Compared with the consumption-based models augmented with sentiment, the FF3S model achieves a higher R^2, indicating that sentiment interacts more closely with risk-based return factors than with aggregate consumption dynamics. This evidence supports the view that investor sentiment captures behavioral variations in risk premia that are not fully reflected in fundamental macroeconomic variables.

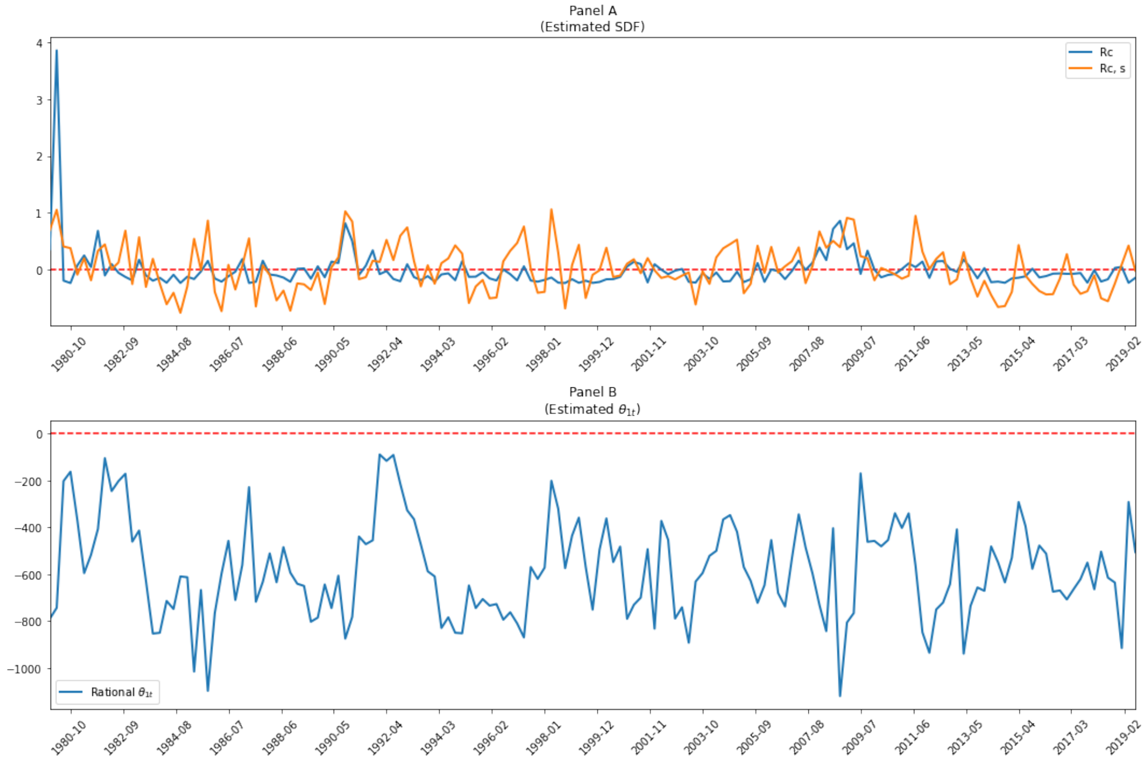

Figure 2.3

Panel A of this figure reports the kernels approximated by Chebyshev polynomials based on consumption growth and sentiment, Panel B reports the estimate of the conditional θ_{1t} parameter

More importantly, the point estimates for the elements of $\boldsymbol{\theta}$ suggest a more plausible behavior of the rational component of the SDF relative to models that omit sentiment. As illustrated in Panel A of Figure 2.3, while the full SDF still frequently assumes negative values, its rational component (depicted by the blue line) is rarely negative. Notably, as shown in Panel B of Figure 2.3 , the parameter θ_{1t} is consistently negative across all periods. This finding is crucial, as it implies that incorporating sentiment ensures the sign of the systematic conditional covariance price remains consistent with the risk-aversion assumption across all realizations of the conditioning variable. The sentiment component of the SDF plays a central role in reconciling the model with observed cross-sectional patterns in average excess returns, while also allowing the representative investor's intertemporal marginal rate of substitution (IMRS) to more closely reflect rational, utility-maximizing behavior.

Table 2.3: Testing the Statistical Significance of Pricing Errors from Consumption Growth and Sentiment-Based Kernels Approximated by Chebyshev Polynomials

	k=1	k=2	k=3	k=4
Model 1: g	46	32	43	31
	(0.006)	(0.161)	(0.013)	(0.174)
Model 2: g, s	35	23	20	16
	(0.095)	(0.578)	(0.75)	(0.904)
Model 3: cay, g, s	27	15	12	12
	(0.373)	(0.937)	(0.988)	(0.987)

The table summarizes Hansen's J-test results for pricing errors across three models and four polynomial orders.

Table 2.3 presents the parameter estimates of the unconditional SDF approximated by Chebyshev polynomial, and the Hansen J–statistic of the over-identifying restrictions evaluated at the two-step (optimal) estimate. The overidentifying restrictions tests, as measured by the $J_T(i)$ statistics, consistently fail to reject the null hypothesis that the moment conditions average to zero and thus that the extra theoretical moment conditions hold across all polynomial orders in both specifications, which implies that the model specifications are consistent with the data. Moreover, the sharp decline in the statistics at higher polynomial orders of the kernel based on both consumption growth and news sentiment indicates that increasing the model's functional flexibility leads to a much closer alignment between the model's predicted moments and the observed data.

2.7 Conclusion

This chapter investigates the role of investor sentiment—specifically, news sentiment—as a structural component of the stochastic discount factor (SDF), presenting a behavioral extension to the standard (C)CAPM framework. Motivated by persistent empirical challenges—particularly the inconsistency between the high cross-sectional explanatory power of conditional models and the violations of their underlying theoretical assumptions—this study examines whether integrating sentiment can enhance both the empirical accuracy and theoretical coherence of asset pricing models. By incorporating sentiment alongside consumption growth, the SDF becomes significantly less negative, underscoring the informational value of sentiment in asset pricing.

Chapter 3

A Quantile Asset Pricing Model with Sentiment Driven Tails

3.1 Introduction

Behavioral finance theory highlights the critical role of investor sentiment in explaining the cross-sectional relationship between risk and expected returns. While predictions of expected returns are fundamental to classical asset pricing, there remains limited understanding of how factors can accurately capture extreme tail events in the return distribution. Moreover, even less is understood about the commonalities across different assets in this regard. Standard asset pricing models typically assume that agents maximize expected utility based on their preferences over future consumption or wealth. However, the underlying preference theory may be too restrictive to adequately adapt to more realistic models of decision-making under risk. One of the key limitations of traditional asset pricing models is the assumption of homogeneous preferences, which fails to capture the diverse risk attitudes that arise under different conditions, particularly in the context of tail events. To mitigate this limitation, researchers have moved away from relying on standard expected utility models and instead introduced heterogeneity into dynamic economic frameworks by assuming that agents maximize their future stream of quantile utilities Chambers (2007), Rostek (2010), and De Castro and Galvão (2019).

The concept of quantile utility maximization provides a more flexible framework to account for these varying preferences. Unlike traditional utility maximization, which relies on the mean of the distribution (a form of risk aversion based on expected returns), quantile utility enables agents to focus on different quantiles of the return distribution, allowing for a more nuanced representation of risk preferences under different conditions. It is noteworthy, however, that the study of quantile utility maximization as an alternative model remains largely unexplored in the existing literature, despite the fact that quantile maximization is straightforward and, in certain contexts, offers a compelling approach. This chapter aims to examine how sentiment explains the cross-section of extreme tail events in the return distribution, based on the Panel Quantile Regression Model developed by Cech and Barunik (2017). This model effectively controls for unobserved heterogeneity among financial assets and allows for the identification of commonalities across different assets in this regard.

According to Fox and Tversky (1998), the psychology of tail events can be understood through a two-step framework. In the first step, an individual assesses the probability of a tail event, which reflects beliefs; in the second, given this probability judgment, the individual makes a decision, which reflects preferences. On the beliefs side, a tentative summary of the available evidence indicates that individuals tend to overestimate the probability of tail events when asked to assess their likelihood. On the preferences side, the prevailing view is that when individuals are aware of a potential tail, the prevailing view is that when individuals are aware of a potential tail event, they tend to assign it greater weight in their decision-making than it would receive under the expected utility framework. This notion is captured by the probability weighting function, a key element of Tversky and Kahneman (1992)'s cumulative prospect theory, which maps subjective probabilities into decision weights, leading individuals to overweight the tails of the distribution they consider.

In quantile preference models, the separation between tastes and beliefs is made explicit. Beliefs determine which quantile of the outcome distribution is being evaluated—that is, which part of the distribution the decision-maker focuses on—while tastes determine how outcomes at that quantile are valued. In this framework, risk attitude is fully captured by the single-dimensional parameter τ, which specifies the quantile of the return distribution that the investor evaluates (de Castro & Galvao, 2021). This formulation is considerably simpler than in expected utility models, where risk attitude is determined by an utility function. Accordingly, changes in sentiment can therefore be interpreted as shifts in the effective quantile level at which investors evaluate risky prospects. During pessimistic periods, investors place more weight on lower quantiles (low τ), while optimistic periods correspond to higher τ evaluation. Taken together, sentiment shapes both the subjective distribution and the investor's relative concern for tail outcomes, providing a natural link between quantile preferences and behavioral measures of sentiment.

Measuring investor sentiment is a non-trivial task, and current literature commonly employs three primary approaches to gauge it. The first approach uses market-based indicators as proxies for investor sentiment, including trading volume, closed-end fund discounts, IPO first-day returns, IPO volume, option-implied volatilities (VIX), or mutual fund flows Baker and Wurgler (2006). The second approach relies on survey-based indices, such as the American Association of Individual Investors' sentiment index Lee et al. (2002), the UBS/Gallup survey Qiu and Welch (2006), the animusX Investor Sentiment Index Lux (2011), and the University of Michigan Consumer Sentiment Index Barone-Adesi et al. (2017). The third approach, more recently, has seen an increasing use of machine learning techniques to analyze news sentiment, as demonstrated in works like A. Shapiro (2022), Calomiris and Mamaysky (2019), and the Thomson Reuters MarketPsych Indices (Reuters (2013)).

This chapter constructs a sentiment index using Principal Component Analysis (PCA) based on five widely used sentiment indices in academic research: the Shapiro, Sudhof, and Wilson (SSW) Sentiment Index, the AAII Investor Sentiment Survey as a proxy for individual investors' sentiment, the Baker and Wurgler (BW) Sentiment Index, the Michigan Consumer Sentiment Index (MCSI), Shiller's Crash Confidence Index (CCI) and CBOE Volatility Index (VIX). The attractiveness of the PCA-based sentiment measure becomes clearer when compared to other alternatives, as it offers several key advantages. First, it allows for the dimensionality reduction of multiple sentiment proxies into a single composite measure, thereby simplifying the analysis without losing critical information. Second, PCA captures the underlying common factors that drive investor sentiment, enhancing its robustness and accuracy by reducing noise from individual indicators. Finally, by aggregating multiple sources of sentiment data, PCA provides a more comprehensive and reliable measure, making it a valuable tool for capturing the complexity of investor sentiment compared to single-indicator approaches.

This chapter contributes to understanding how investor sentiment explains the cross-section of extreme tail events in stock returns distribution, using the Panel Quantile Regression Model, which effectively controls for unobserved heterogeneity across financial assets and allows for the identification of commonalities in the tail behavior of returns across different assets. Additionally, the chapter introduces a novel approach to measuring investor sentiment by constructing a sentiment index through Principal Component Analysis (PCA). This composite measure aggregates multiple widely used sentiment indices, with a particular focus on sentiment driven by machine learning-based news sentiment, offering a more robust and adaptable tool for sentiment analysis compared to single-indicator methods. By reducing the dimensionality of multiple sentiment proxies, PCA captures the underlying common factors driving sentiment, enhancing the model's ability to link sentiment with tail risk. This approach not only improves upon traditional models by addressing the limitations of homogeneous preferences and expected utility maximization but also provides a more comprehensive framework for analyzing the role of sentiment in explaining extreme tail events in asset returns.

In order to assess how much unobserved heterogeneity affects the relationship between sentiment and returns, the individual Univariate Quantile Regression(UQR) without controlling for unobserved heterogeneity is introduced to establish a benchmark for comparison. The comparison of panel quantile regression (PQR) estimates with individual univariate quantile regression (UQR) parameter estimates reveals that, once unobserved heterogeneity is controlled for through PQR, PC1 exerts a more pronounced negative influence on the upper quantiles of returns than most individual UQR estimates. This finding suggests that accounting for heterogeneity exposes a stronger downward impact of negative sentiment on the higher quantiles of the return distribution, which aligns with the asymmetric effect of sentiment observed in Table 2.2. Importantly, it indicates that, without controlling for unobserved heterogeneity, the influence of sentiment factors on return distributions may be systematically underestimated, especially at the tails. This comparison underscores the importance of accounting for unobserved factors in panel data to avoid biased estimates.

3.2 Literature review

The classic asset pricing models are built on the assumption that investors act rationally and that markets are efficient, leaving little room for behavioral influences. Although it's widely acknowledged that there are an abundance of irrational noise traders with erroneous stochastic beliefs in the market, early economists tend to overlook their impact on asset price formation. Friedman (1953b) and Fama (1965) argue that rational arbitrageurs counteract irrational traders by trading against them, pushing prices toward fundamental values Fama (1965) and Friedman (1953a). However, the standard finance models struggle to account for the stylized anomalies. To address these discrepancies, researchers in behavioral finance have proposed an alternative framework based on two main assumptions. The first assumption presented by De Long et al. (1990) posits that investors are influenced by sentiment. This sentiment refers to irrational beliefs about future cash flows and investment risks that lack a solid factual basis. The second assumption, highlighted by Shleifer and Vishny (1997), states that it is both costly and risky to bet against these sentimental investors, and thus there are limits to arbitrage.

A large body of literature has both theoretically and empirically examined whether and how investor sentiment influences asset prices—driven by overly optimistic or pessimistic beliefs and the behavioral biases of irrational agents. Baker and Wurgler (2006) show that sentiment has a greater impact on securities with valuations that are highly subjective and hard to arbitrage. In line with this prediction, the empirical evidence indicates that when beginning-of-period sentiment proxies are low, subsequent returns are relatively high for stocks that are younger, smaller, more volatile, unprofitable, non–dividend paying, distressed, or with extreme growth potential. In contrast, when sentiment is high, these categories of returns experience relatively low subsequent returns. Baker and Wurgler (2007) argue that the key factor making those stocks more speculative than others lies in the difficulty and subjectivity of determining their intrinsic values. For example, in the case of a young, currently unprofitable but potentially highly profitable growth firm, the combination of a lack of earnings history and a highly uncertain future enables investors to justify valuations ranging from unrealistically low to excessively high, depending on their prevailing sentiment. From a psychological perspective, such uncertainty amplifies behavioral biases such as overconfidence (Daniel et al., 1998), representativeness, and conservatism (Barberis et al., 1998).

Liang (2018) develops a consumption-based asset pricing model that integrates sentiment and presents empirical evidence that portfolios with greater sentiment loadings earn higher expected returns, although this pattern does not extend to individual stocks. Likewise, Hirshleifer et al. (2020) conceptualize investor mood as a specific form of sentiment and show that assets' sensitivity to mood fluctuations accounts for the seasonal variation in cross-sectional returns. More recently, Y. Chen et al. (2021) employ the Baker and Wurgler sentiment index to demonstrate that exposure to market sentiment helps explain the cross-sectional variation in hedge fund returns, with the observed sentiment premium largely attributable to managers' sentiment-timing skill.

In the time-series context, however, there remains an ongoing empirical debate about whether investor sentiment can reliably forecast aggregate stock returns. D.

Huang et al. (2015) construct an aligned investor sentiment index using six sentiment proxies proposed by Baker and Wurgler (2007) and demonstrate that this composite measure serves as a powerful negative predictor of aggregate market returns over the long run. Lux (2012) argues that the underlying relationship between investor sentiment and stock returns may be nonlinear. Using two forms of nonlinear causality testing, Dergiades (2012) finds that investor sentiment exhibits a significant Granger-causal effect on stock returns. However, Bekiros et al. (2016), employing a nonparametric causality approach developed by Nishiyama et al. (2011) and using the same data as D. Huang et al. (2015), find little evidence that investor sentiment has predictive power for aggregate market returns.

Although a well-known line of research on investor sentiment and aggregate stock returns has evolved following the pioneering work of Baker and Wurgler (2006), and the focus has shifted from questioning whether sentiment affects stock prices to how it can be measured and quantified, this influence is complicated by cognitive biases that shape the asymmetry of investors' risk preferences, which in turn gives rise to another debate regarding the concerns when investor sentiment exerts the greatest impact on stock returns. While the argument of Shleifer and Vishny (1997) suggests that the impact of investor sentiment on stock returns should be symmetric, variations in short-selling costs across different market conditions create a more complex and nuanced relationship. On one hand, Stambaugh et al. (2012) demonstrate that in a market with short-sale constraints, overpricing following high-sentiment periods tends to be more prevalent than sentiment-driven underpricing following low-sentiment periods, as evidenced by the returns generated from long–short strategies based on these anomalies. Grinblatt and Keloharju (2009) and Lemmon and Portniaguina (2006) show that unsophisticated investors tend to enter the stock market following good times rather than during downturns, return predictability is therefore expected to be stronger under good market conditions. Using NBER business-cycle-based market states and a Markov-switching framework, the empirical findings of Chung et al. (2012) provide evidence supporting this view. On the other hand, prospect theory posits that individuals feel greater pain from losses than pleasure from equivalent gains (Kahneman & Tversky, 1979b). Human behavior differs markedly between periods of anxiety and fear and those characterized by prosperity and calm. For instance, anxiety tends to heighten irrational decision-making (Gino et al., 2012). Moreover, several studies find that return predictability is more pronounced during recessions (Cujea & Hasler, 2016; Garcia, 2013).

Despite these findings, the majority of existing literature tends to focus on average sentiment effects, which overlooks how different segments of the return distribution respond to sentiment changes, particularly in extreme market conditions. The traditional decision making models that rely on average behaviors often fail to capture the complexities inherent in investor decision-making, particularly during periods of extreme market conditions. Moreover, much of this literature lacks a strong theoretical foundation and primarily focuses on incorporating sentiment variables into empirical asset pricing models as non-tradable pricing factors. The developments in quantile preferences provide a framework to analyze these dynamics. Quantile preference models are especially useful for understanding downside risk preferences and asymmetric risk perceptions, making them more adaptable in accounting for sentiment-driven investor behavior, where investors may overreact to the likelihood of extreme market events. This chapter introduces the appealing quantile utility approach, which enables researchers to capture the behavioral biases without assuming a specific utility function.

In the field of economics, expressing risk preferences using quantiles or distribution functions is not new, notable contributions include the works of Machina (1982), Yaari (1987), and Dybvig (1988). Quantile preferences were first proposed by Manski (1988b)and subsequently axiomatized by Chambers (2009) and Rostek (2010). Manski (1988b)develops the decision-theoretic attributes of quantile maximization and and investigates the risk preferences of the quantile utility maximize agents. Chambers (2009) demonstrates that quantile preferences over distributions are characterized by monotonicity, ordinal covariance, and continuity. de Castro and Galvao (2017), building on quantile-based preferences (Manski (1988a); Rostek (2010)), developped a dynamic model where agents maximize future quantile utilities, aligning with the preference for specific quantiles rather than just expected values. This study shows that maximizing a lower quantile of portfolio returns reflects more risk aversion than maximizing a higher quantile.

Unlike traditional regression techniques, quantile regression characterizes both the location and scale of the conditional distribution, thereby providing a comprehensive picture of how sentiment influences stock returns across the entire distribution. Traditional mean-based approaches may fail to capture the true effect due to heterogeneity—a pervasive feature of time-series variables. Moreover, quantile regression can directly capture asymmetric effects (Ding et al., 2016). It is an efficient tool for addressing nonlinearity without imposing a specific functional form on the model, such as dividing the sample into ex ante regimes (e.g., recessions versus expansions), selecting an index to represent the business cycle, or specifying a transition function that governs regime changes (Linnemann & Winkler, 2016). The possible state dependence of the sentiment–return relationship can thus be examined through different quantiles of the conditional distribution, where the upper (lower) quantiles correspond to bullish (bearish) market conditions (Baur, Dimpfl, & Jung, 2012; Ding et al., 2016).

Quantile regression provides a more nuanced and comprehensive understanding of market dynamics by capturing relationships between variables across different points of the distribution. An expanding body of research indicates that sentiment-driven tail models are particularly effective in tracing the influence of investor sentiment on the extremes of return distributions, which becomes especially relevant during periods of market stress or exuberance. More specifically, integrating sentiment measures into quantile-based frameworks can enhance predictive power in asset pricing by explicitly accounting for behavioral biases. For example, Han et al. (2016) employed both Fama–MacBeth and quantile regression techniques to examine the triangular relationship among investor sentiment, stock returns, and trading volume, uncovering intricate patterns that differ across various market states and sentiment levels. Similarly, Chang et al. (2016) applied both methodologies and found that investor sentiment Granger-causes stock returns under certain market conditions, demonstrating the predictive power of sentiment in shaping future market movements. Moreover, Li (2017) employs quantile regression to show that the impact of trading volume on stock returns becomes more pronounced under extreme market conditions, highlighting the nonlinear nature of this relationship. Abdelmalek (2023) provides evidence that the relationship between realized volatility, investor sentiment, and stock returns is asymmetric across quantiles in the quantile regression framework, in contrast to OLS regression. Moreover, this study extends behavioral theory beyond the traditional leverage hypothesis by explaining the asymmetric linkage between returns and volatility at both high and low data frequencies.

Although empirical research increasingly recognises that investor sentiment plays a disproportionately large role in shaping asset returns under extreme market conditions, existing studies mainly rely on conventional sentiment indices such as the Baker–Wurgler index. Recognising the need for reliable empirical proxies of investor sentiment, the question of how sentiment should be measured has emerged as an important topic in the literature. Following the pioneering work of Baker and Wurgler (2006), who construct a composite sentiment index by aggregating several market-based proxies, such as IPO volume, dividend premium, and the closed-end fund discount, investor sentiment has been measured through a variety of approaches. For example, Da et al. (2015), D. Huang et al. (2015), and Tetlock (2007), developed their own Investor Sentiment Indices (ISIs) using different approaches. There are two main ways to measure the ISI: direct and indirect (for additional sentiment classification methods, see Ghallab et al. (2020); and Bhardwaj et al. (2015)). In the direct measurement, primary data is typically gathered from investors through surveys. On the other hand, the indirect ISI relies on market-specific, firm-specific, and investor-specific data to serve as proxies for sentiment. Baker and Wurgler, for instance, used an indirect approach to construct a composite ISI to capture investor sentiment in the U.S. market. This study contributes to this field by constructing a quantile-based asset pricing model that integrates sentiment proxies derived from principal component analysis (PCA), thus capturing both the fundamental and behavioral factors affecting risk preferences across quantiles.

The methodological framework of this chapter is centered on the Panel Quantile Regression (PQR) Model, a robust approach suited to capturing quantile-specific sensitivities within asset pricing while accounting for individual-specific effects in panel data. This model was first conceptualized by Koenker and Bassett (1978) in the context of quantile regression, which focuses on estimating conditional quantiles of a response variable rather than the mean, and it was later extended to panel data settings by researchers such as Koenker (2004) and Galvao (2011). Koenker introduced quantile regression as a method to capture the conditional quantiles of a response variable, while Galvao and others extended this approach to panel data to account for both individual-specific effects and varying impacts across quantiles.

3.3 PCA to constructing sentiment

The sentiment indices extracted from the five sentiment indices using PCA, capturing distinct dimensions of investor sentiment, can enhance our understanding of market dynamics. Each sentiment index captures distinct dimensions of investor sentiment, reflecting similarities and differences in their data sources, time horizons, and focus on various market participants or events. Understanding these distinctions will provide a clearer picture of how investor sentiment is measured and how each index can be used in empirical research. The SSW index reflects economic sentiment derived from major U.S. economic and financial newspaper articles through state-of-the-art textual analysis techniques. In contrast, the BW Sentiment Index gains insights from market indicators such as trading volume and dividend premiums, highlighting immediate market conditions. The MCSI serves as a survey-based measure of consumer confidence, emphasizing public perceptions of economic conditions. Meanwhile, the CCI assesses investor confidence and risk perception concerning market stability and crash likelihood. The AAII Investor Sentiment Survey gauges individual investors' expectations for market performance, providing a self-reported perspective on stock market sentiment. Lastly, VIX is a forward-looking volatility indicator implied by SP 500 Index (SPX) option prices. As a measure of implied market volatility—capturing the market's expectations of risk over the next 30 days, VIX can also serves as a barometer of investor sentiment.

The sentiment index extracted from PCA has several appealing properties. PCA transforms the original correlated sentiment measures into a smaller set of uncorrelated variables (the principal components), effectively capturing the underlying patterns and dimensions of investor sentiment. Before performing PCA on these sentiment measures, ensuring stationarity and standardization are critical preprocessing steps. Ensuring stationarity is crucial to avoid spurious results that can arise from non-stationary data. Non-stationary time series may exhibit trends or changing volatility over time, leading to misleading correlations and invalid statistical inferences. Transforming the data to achieve stationarity results in more robust and interpretable outcomes. Stationarity is assessed using the Augmented Dickey–Fuller (ADF) test, whose null hypothesis is a unit root (non-stationarity). The Augmented Dickey-Fuller (ADF) test results indicate that SSW, AAII, and VIX are stationary series, while BW, MCSI, and CCI are non-stationary. To ensure that all sentiment variables used in the PCA are stationary, these series are transformed by first differencing (BW, MCSI, CCI) before proceeding. With all series rendered stationary, they are then standardized, allowing for meaningful comparisons across indices with different scales and units. This process allows for a more accurate projection of the original data onto the directions of greatest variance, ensures that the resulting principal components reflect the most significant sources of variation within the sentiment measures, thereby revealing the most significant dimensions of investor sentiment. By mitigating the influence of scale disparities among variables, standardization enhances the interpretability of the components.

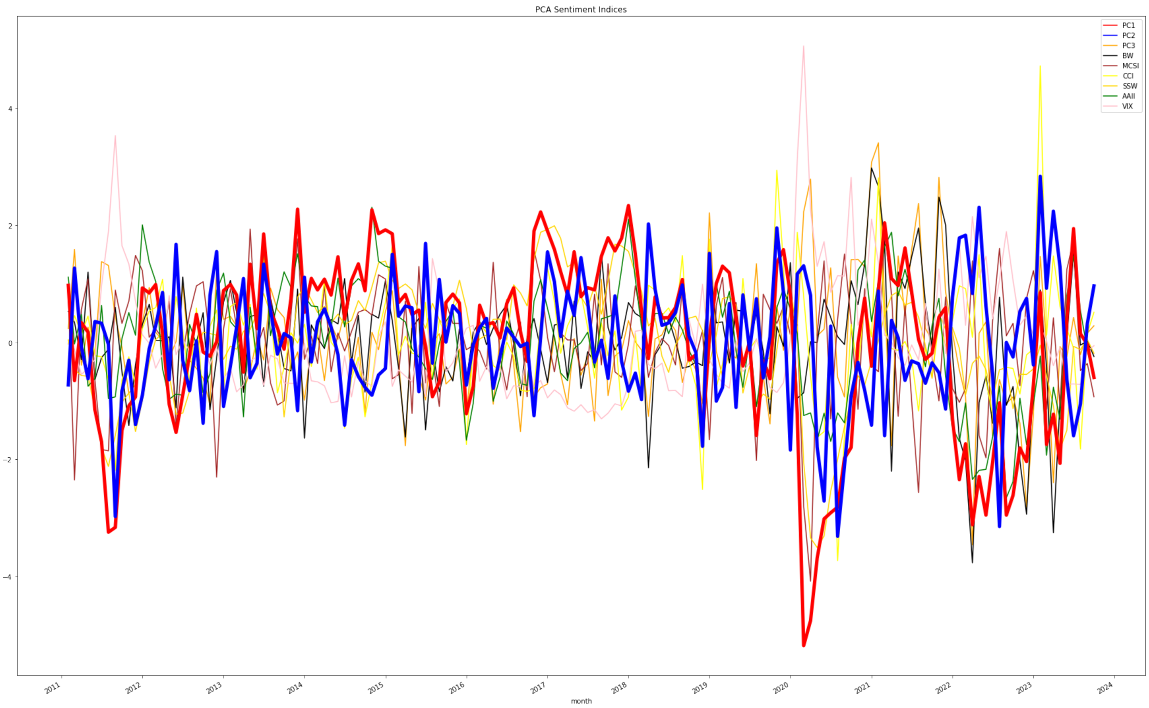

Figure 3.1: Plot of sentiment indices The figure shows the sentiment indices plotted at a monthly frequency. The sample period is February 2011 to November 2023.

The Figure 2.1 displays the plots of the first three principal components (PC1, PC2, PC3) derived from PCA, alongside the original sentiment indices. PC1, represented by the thick red line, captures the most significant dimension of variation across the five sentiment indices, while PC2, represented by the thick blue line, captures the second most significant uncorrelated dimension. PC3 adds a further orthogonal component that explains residual but systematic co-movement not captured by PC1–PC2. The PCA transformation allows a clearer representation of the overall sentiment trends by reducing noise from individual indices, reflecting a consolidated view of the sentiment landscape, highlighting how different indices relate to each other and contribute to overall market dynamics. The figure helps illustrate the relationship and co-movement between various sentiment measures and their aggregated principal components.

Table 3.1:

Loading matrix in PCA

	PC1	PC2	PC3
BW	0.25	-0.42	0.74
MCSI	0.26	-0.48	-0.51
CCI	0.17	0.74	0.22
SSW	0.53	0.14	-0.18
AAII	0.54	-0.10	0.24
VIX	-0.51	-0.16	0.24
EV Ratio	0.36	0.19	0.17

The table 3.1 reports the factors loading matrix values, indicating the contribution of each original feature to the principal components, and the explained variance ratios (EV Ratio) indicating the proportion of the total variance in the dataset that each principal component captures. As shown in Table 3.1, PC1 represents a general sentiment trend primarily influenced by SSW, AAII, and VIX, loading positively on AAII (0.54) and SSW (0.53)—capturing optimism—and negatively on VIX (0.51), an inverse indicator of market confidence. PC2 represents a contrasting sentiment dimension primarily influenced by CCI, MCSI, and BW, loading positively on CCI (0.74)—reflecting investors' crash confidence—and negatively on MCSI (0.48) and BW (0.42). This component captures a divergence between crash confidence (CCI) and broader market-based sentiment and consumer confidence (BW and MCSI). PC3 captures a residual contrast largely explained by BW (0.74) and MCSI (0.51), reflecting idiosyncratic variation specific to market-oriented sentiment that is not aligned with broader survey- or news-based indicators.

To understand why certain sentiment variables exhibit stationarity (short memory) while others are persistent (possessing unit roots), advanced econometric techniques such as Vector Error Correction Models (VECM) and Granger causality analysis can be employed to model their joint dynamics. These methods enable the classification of sentiment variables into short-term and long-term sentiments based on their dynamic interactions and relationships with market outcomes. For the non-stationary sentiment variables, applying VECM and Vector Autoregression (VAR) analyses enables the exploration of long-term equilibrium relationships through cointegration analysis and the identification of Granger causality among these variables. These techniques help determine which sentiment variables drive others, elucidating the underlying dynamics of investor sentiment. Integrating these econometric analyses complements the PCA by providing a deeper understanding of the temporal relationships and causality among the sentiment measures. This comprehensive approach enhances the robustness and interpretability of the results, ultimately leading to a more nuanced understanding of how different types of sentiment influence financial markets.

3.4 Panel Quantile Regression Model

Following the brief description of sentiment construction for model development, a simple linear model based on factors as proxies for aggregate consumption and control variables, including idiosyncratic factors, is proposed for the cross-section of quantiles of future returns.

$$Q_\tau\left(r_{i,t+1}\right) = \alpha_{i,\tau} + \beta_{1,\tau}\sigma_{i,t} + \beta_{2,\tau}\mathrm{PC1}_t + \beta_{3,\tau}\mathrm{PC2}_t + \beta_{4,\tau}\mathrm{PC3}_t + CF_t^T\beta_{CF,\tau} + \beta_{RF,\tau}RF_t \tag{3.1}$$

where $r_{i,t+1} = p_{i,t+1} - p_{i,t}$ are monthly logarithmic returns, $\alpha_{i,\tau}$ is the individual fixed effect at quantile τ, capturing stock-specific, time-invariant characteristics. $\sigma_{i,t}$ represents the stock-specific volatility, estimated using the Garman and Klass, 1980 rangebased estimator, which offers a realized, backward-looking measure of volatility based on intraday price ranges, rather than relying solely on the closing price. This range-based volatility measure is defined as: $\sigma_{i,t} = 0.511(u-d)^2 - 0.019\{c(u+d) - 2ud\} - 0.383c^2$, where u, d, and c represent the logarithmic differences between the highest and lowest prices of the day (u), the lowest price and the opening price (d), and the closing price and the opening price (c), respectively. Specifically, $u = \log\left(P_{i,t}^H\right) - \log\left(P_{i,t}^L\right), d = \log\left(P_{i,t}^L\right) - \log\left(P_{i,t}^O\right)$, and $c = \log\left(P_{i,t}^C\right) - \log\left(P_{i,t}^O\right)$, where $P_{i,t}^H, P_{i,t}^L, P_{i,t}^O$, and $P_{i,t}^C$ denote the daily highest, lowest, opening, and closing prices of stock i on day t, respectively. This estimator incorporates information from the entire trading day, capturing a more comprehensive view of volatility by accounting for intraday price fluctuations.

$\mathrm{PC1}_t$, $\mathrm{PC2}_t$ and $\mathrm{PC3}_t$ are the first, second and third principal components extracted from the five sentiment indices mentioned above via PCA, serving as time-varying common factors. They capture systematic influences, representing orthogonal dimensions of market sentiment that affect the quantiles of stock returns. $CF_t^T\beta_{CF,\tau}$ are the vector of Carhart Four Factors (market risk, size, value, and momentum) with quantile-specific sensitivities $\beta_{CF,\tau}$. These are traditional risk factors shared across stocks and represent underlying systematic risks affecting the returns distribution. The momentum factor captures the tendency of stocks that have performed well in the recent past (typically 3–12 months) to continue performing well, and vice versa for poor performers. RF_t is the risk-free rate at time t, typically represented by the return on the 1-month Treasury bill, and $\beta_{RF,\tau}$ represents the quantile-specific sensitivity of stock i 's returns to the risk-free rate.

The model combines common factors like the market sentiment and Carhart four factors with stock-specific volatility factor , allowing for a nuanced understanding of return distribution at different quantiles. Each factor contributes uniquely, helping to capture the varying impacts of idiosyncratic and systematic influences on future returns. Volatility and investor sentiment are both measures of market dynamics, but they capture different aspects of market behavior. Volatility is typically seen as a measure of risk or uncertainty. It reflects the market's fluctuations, often driven by reactions to macroeconomic events or unexpected shocks. It's more associated with market-wide risk and uncertainty.

3.5 Estimation Methodology and Data

In traditional asset pricing models, the focus is typically on the central tendency (mean) of the return distribution. However, these models often overlook the important insights provided by the entire distribution of returns, particularly in the context of tail risks and extreme market conditions. In contrast, quantile-based models offer a more comprehensive approach by analyzing behavior across different quantiles of the distribution, providing a better understanding of risks beyond the mean. The quantile objective function estimates the relationship between the independent variables and different quantiles of the dependent variable, which helps to capture the effects not just on the mean but across the entire distribution of the dependent variable. In a panel data setting, this means estimating the effects of sentiment and control variables on different points (quantiles) of the distribution of stock returns for multiple entities (stocks) over time. More importantly, by using Koenker (2004)'s penalized fixed effects estimator, the advantage of this approach is ability to account the unobserved heterogeneity among financial assets which will yield more precise quantile specific estimates. The coefficients from equation (3.1) can be obtained by minimizing the following Quantile Check Function:

$$\min \frac{1}{n} \sum_{t=1}^{n} \rho_\tau \left(r_{i,t+1} - \alpha_{i,\tau} - \beta_{1,\tau}\sigma_{i,t} - \beta_{2,\tau}\mathrm{PC1}_t - \beta_{3,\tau}\mathrm{PC2}_t - \beta_{4,\tau}\mathrm{PC3}_t - FF_t^\top \beta_{FF,\tau} \right) \tag{3.2}$$

where $\rho_\tau(u)$ is the quantile check function defined as:

$$\rho_\tau(u) = \begin{cases} \tau \cdot u & \text{if } u \geq 0 \\ (\tau - 1) \cdot u & \text{if } u < 0 \end{cases} \tag{3.3}$$

Here, u is the residual error , τ represents the quantile level focusing on. The quantile check function is a loss function tailored for quantile regression. Different from the more common Mean Squared Error (MSE) or Mean Absolute Error (MAE), the quantile check function is a type of loss function that penalizes residuals asymmetrically based on the quantile level τ, allowing it to target specific quantiles rather than just the mean.

To assess the impact of unobserved heterogeneity on the relationship between sentiment and returns, the individual Univariate Quantile Regression (UQR) is introduced without controlling for unobserved heterogeneity, serving as a benchmark for comparison with the Panel Quantile Regression (PQR) estimates. Panel quantile regression methods, such as Koenker's penalized quantile regression for panel data, utilize penalization techniques to account for unobserved individual effects. This specialized pooled regression combines the benefits of pooled analysis with the structure of panel data, enabling the analysis of both within-entity (time-based) and between-entity (cross-sectional) variations. This framework is particularly valuable for understanding heterogeneous effects across different parts of the conditional distribution in panel data settings. Individual UQR, on the other hand, typically operates on individual time series of a single asset (or stock), without incorporating cross-sectional data or controlling for asset-specific fixed effects. In summary, Panel quantile regression pools the data while accounting for individual heterogeneity, whereas individual quantile regression estimates completely separate models for each entity. Panel quantile regression offers a balance between pooling information and recognizing unique individual effects, which is generally more efficient and interpretable in panel data settings.

3.6 Data

The empirical data applied in the proposed models involves a panel of the 100 most liquid stocks based on market capitalization and liquidity from the S&P 500 spanning from February 2011 to November 2023, which was obtained from Yahoo Finance. The Volatility Index (VIX) data is obtained from the Chicago Board Options Exchange (CBOE) and the Fama-French Factors and momentum factor datasets are collected from Professor Kenneth R. French's official website. The selection criteria market capitalization means my research is focusing on large, well-known companies that tend to have higher liquidity and stable trading activity, and the Liquidity ensures that the stocks are actively traded, which minimizes the risk of abnormal price movements caused by low trading volumes. The sample period offers a substantial timeframe that captures a range of market conditions, making it highly valuable for analysis. This period includes the dramatic market crash and rapid recovery driven in large part by expansive monetary policies during the COVID-19 pandemic in 2020. These policies included historically low interest rates, quantitative easing, and other supportive measures from central banks, which aimed to stabilize financial markets and stimulate economic recovery. This long sample period allows to account for both long-term structural changes and sudden shocks, capturing a wide range of market behaviors across distinct market phases, making the empirical analysis more comprehensive and robust. The pandemic brought unprecedented volatility investor sentiment shifts, with global markets experiencing sharp declines followed by government interventions and a significant rebound in asset prices. Analyzing data from such a period allows for insights into how market sentiment and risk factors respond to extreme events and extraordinary policy measures, offering a unique perspective on the dynamics of stock returns across varying economic conditions. This timeframe also allows for studying how investor sentiment and other systematic factors interact under both high-stress and more stable market environments, enhancing the robustness of your empirical model.

3.7 Empirical results

Table 3.2 present the coefficient estimates from the panel quantile regression model in equation (3.1), which provides insights into the differential impact of market-wide and stock-specific factors on stock returns across quantiles ranging from the 10th to the 90th percentile. Each factor's effect varies in both magnitude and statistical significance, capturing significant variations in the sensitivities of stock returns to idiosyncratic and systematic factors across the return distribution. Additionally, to provide a clearer view of the dynamics, the coefficient estimates are plotted graphically in Figures 3.2, respectively.

Table 3.2:

Coefficient estimates of Panel Quantile Regressions.

quantile factor	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
Intercept	-0.0318	-0.0188	-0.0085	0.0020	0.0121	0.0216	0.0363	0.0550	0.0867
	(-41.46)	(-26.35)	(-12.03)	(2.82)	(17.01)	(29.26)	(43.86)	(57.74)	(60.19)
Volatility	-1.7429	-1.3614	-1.0688	-0.8433	-0.6200	-0.3647	-0.3129	-0.2648	-0.2324
	(-578.32)	(-382.51)	(-252.73)	(-166.45)	(-105.24)	(-51.15)	(-34.67)	(-22.31)	(-11.69)
PC1	-0.0034	-0.0036	-0.0037	-0.0033	-0.0037	-0.0036	-0.0044	-0.0063	-0.0089
	(-7.69)	(-8.73)	(-9.26)	(-8.3)	(-9.26)	(-8.8)	(-9.41)	(-11.6)	(-10.57)
PC2	0.0011	0.0012	0.0010	0.0004	0.0009	0.0010	0.0003	0.0011	0.0032
	(1.75)	(2.24)	(1.87)	(0.77)	(1.78)	(1.76)	(0.44)	(1.58)	(3.19)
PC3	0.0007	0.0002	-0.0006	-0.0002	0.0003	0.0007	0.0024	0.0035	0.0066
	(1.01)	(0.29)	(-1.08)	(-0.36)	(0.56)	(1.27)	(3.62)	(4.59)	(5.48)
ERM	0.0057	0.0069	0.0077	0.0083	0.0089	0.0093	0.0094	0.0095	0.0096
	(34.64)	(45.93)	(51.67)	(54.77)	(58.45)	(59.0)	(53.27)	(46.46)	(30.69)
SMB	-0.0006	-0.0008	-0.0011	-0.0010	-0.0011	-0.0009	-0.0008	-0.0001	0.0010
	(-2.04)	(-3.39)	(-4.4)	(-4.22)	(-4.63)	(-3.75)	(-2.84)	(-0.24)	(2.18)
HML	0.0011	0.0011	0.0014	0.0015	0.0015	0.0016	0.0017	0.0020	0.0017
	(5.59)	(6.38)	(8.17)	(8.67)	(8.59)	(8.69)	(8.02)	(7.71)	(4.12)
MOM	-0.0006	-0.0009	-0.0011	-0.0010	-0.0010	-0.0010	-0.0011	-0.0015	-0.0020
	(-3.03)	(-4.99)	(-5.97)	(-5.59)	(-5.62)	(-5.26)	(-5.34)	(-6.46)	(-5.26)
RF	-0.0100	-0.0123	-0.0094	-0.0138	-0.0136	-0.0058	0.0037	0.0060	0.0024
	(-1.72)	(-2.3)	(-1.79)	(-2.62)	(-2.63)	(-1.09)	(0.64)	(0.92)	(0.25)

Notes: This table displays coefficient estimates with t-statistics in parentheses, the intercept represents the quantile-specific common constant.

The intercept is the quantile-specific common constant after normalizing stock fixed effects ($\sum_i \alpha_{i,\tau} = 0$). Its monotonic increase across quantiles indicates an upward shift in the conditional distribution of returns, suggesting that the baseline level of expected returns rises progressively toward the upper quantiles after accounting for both observable and unobservable heterogeneity. The consistently significant t-statistics across all quantiles underscore the robustness and stability of the estimated fixed effects throughout the return distribution. The $\alpha_{i,\tau}$ are estimated but omitted from the table for brevity. The idiosyncratic volatility, measured by the Garman–Klass estimator as a proxy for realized, asset-specific uncertainty, exhibits a strongly negative and statistically significant effect across all quantiles, with the most pronounced impact observed in the left tail. This reflects that idiosyncratic risk is most penalized during periods of poor market performance (left tail), underscoring its relevance for downside risk management, which is consistent with the findings from Ang et al. (2009).

Figure 3.2: PQR Estimates for $\beta_{i,\tau}$ together with UQR box plots showing individual estimates with univariate individual $\mathrm{I} = 1, \ldots, 100$ QR estimates

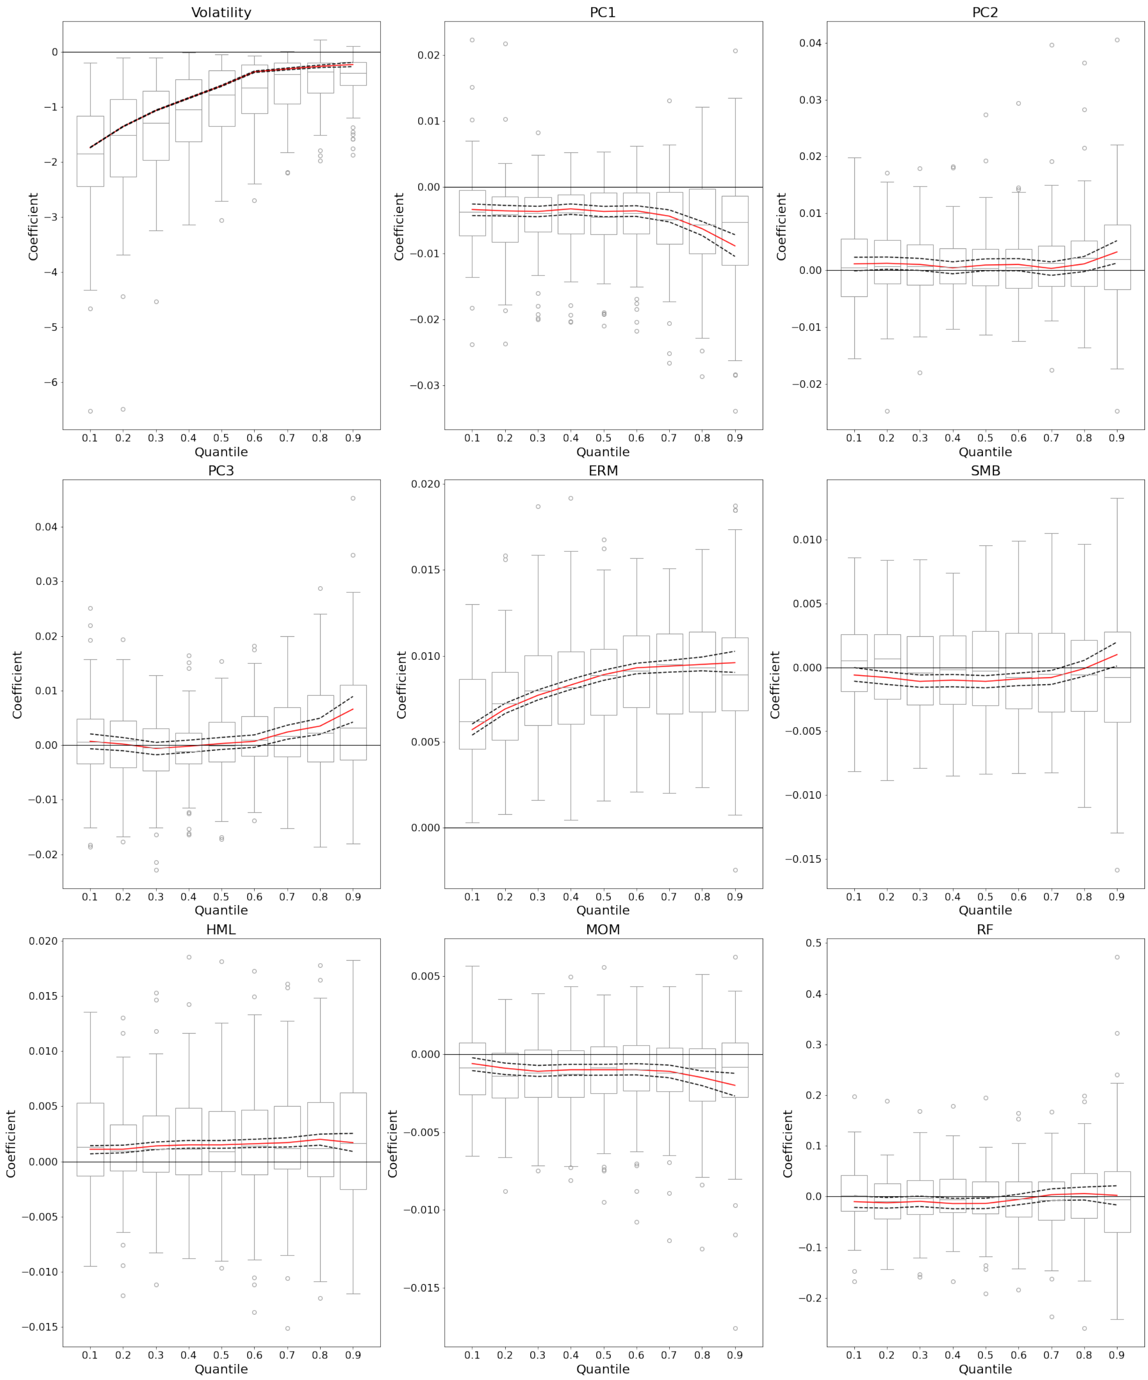

The red solid line representing the estimated coefficient, the dashed lines indicating the 95% confidence intervals and individual UQR parameter estimates plotted in boxplots

The first component, PC1, exhibits a consistently negative and statistically significant effect across all quantiles of the return distribution. While the PC1 has a relatively stable effect across lower and middle quantiles, its impact becomes more pronounced at the highest quantile, where the coefficient nearly doubles compared to the average effect. This pattern indicates that increases in sentiment disproportionately dampen the upper tail of the return distribution, reducing the likelihood of extreme positive returns. This observed pattern aligns with the finding of Baker and Wurgler (2007), who document that high sentiment correlates with lower future returns. A similar conclusion can be drawn from Figure 3.2, which presents the comparison between the PQR (Panel Quantile Regression) alongside their corresponding 95% confidence intervals and individual Univariate Quantile Regression (UQR) parameter estimates plotted in boxplots. Notably, Figure 3.2 reveals that once unobserved heterogeneity is controlled for through PQR, PC1 exerts a larger influence on the upper quantiles of returns compared to most individual UQR estimates. Specifically, the coefficient estimate from the PQR at the 0.9 quantile is significantly lower than the median of the individual UQR coefficient, indicating that the effect of sentiment is more pronounced and significant at the upper quantiles of returns. In contrast, the UQR estimates, as shown in the boxplot, exhibit less variation across quantiles. Meanwhile, the second principal component PC2, has a smaller, positive effect, reflecting orthogonal dimensions of sentiment that influence returns differently. The third component, PC3 is economically small and mostly insignificant from the 0.1–0.6 quantiles, turns positive around the upper-middle quantiles, and becomes clearly positive and statistically significant in the right tail (0.7–0.9). This pattern implies that the dimension of sentiment captured by PC3 amplifies extreme positive outcomes but has little effect on the center or left tail of the return distribution.

Among the Fama-French-Carhart factors, excess market return (ERM) is positive and highly significant throughout, with a peak at the median quantile, supporting its robustness as a cross-sectional return driver. The size factor (SMB) is negative and significant in lower quantiles but turns positive near the upper tail, suggesting a size premium conditional on market state. The value factor (HML) exerts a positive and significant effect across quantiles, slightly increasing toward the upper tail—highlighting the persistent value premium. The momentum factor (MOM) exhibits a negative and statistically significant effect across all quantiles of the return distribution, indicating a consistent pattern of return reversal. The coefficient profile is relatively stable between $\tau = 0.3$ and $\tau = 0.7$, suggesting a modest, persistent drag of momentum on mid-range returns. However, the effect intensifies substantially in the upper tail: by $\tau = 0.9$, the magnitude of the negative coefficient nearly doubles, implying that when past winners keep outperforming (high MOM), subsequent very large positive pay-offs become less likely for the average asset in the panel. This pattern suggests that, contrary to the classical momentum anomaly (where past winners continue to win), the cross-sectional return dynamics in this sample exhibit a contrarian effect: assets that have performed well in the recent past tend to underperform in the future. This dynamic underscores the strength of the panel quantile regression framework in capturing distributional heterogeneity, enabling the detection of asymmetric effects across the return distribution that traditional mean-based models might overlook. Moreover, this reversal effect can be interpreted in light of behavioral finance theories, such as investor overreaction, where extreme past performance leads to excessive optimism and subsequent corrections. The findings also suggest that momentum-based strategies may yield diminishing or even negative returns in certain market segments, especially when conditioning on return quantiles.

This reversal effect underscores the value of Panel Quantile Regression (PQR) in revealing asymmetric effects that mean-based models overlook. The findings align with behavioral finance theories, such as investor overreaction, where excessive optimism following strong past performance leads to corrections.

Table 3.3:

Corrleation matrix: Sentiment variables and Observables.

	PC1	PC2	PC3	ERM	SMB	HML	MOM	RF
PC1	1.00	-0.00	0.00	0.27	0.09	0.14	-0.14	-0.09
PC2	-0.00	1.00	0.00	-0.01	-0.10	-0.02	0.04	0.11
PC3	0.00	0.00	1.00	0.12	0.14	-0.03	-0.12	-0.15
ERM	0.27	-0.01	0.12	1.00	0.32	0.04	-0.41	-0.10
SMB	0.09	-0.10	0.14	0.32	1.00	0.04	-0.27	-0.13
HML	0.14	-0.02	-0.03	0.04	0.04	1.00	-0.28	-0.13
MOM	-0.14	0.04	-0.12	-0.41	-0.27	-0.28	1.00	-0.04
RF	-0.09	0.11	-0.15	-0.10	-0.13	-0.13	-0.04	1.00

Table 3.3 reports the correlation matrix between the PCA-derived sentiment factors and the control variables. The first principal component (PC1) exhibits generally weak correlations with the control variables, except for a moderate positive correlation with the market excess return (ERM, $r = 0.27$). This indicates that the sentiment factor is largely orthogonal to the size, value, and momentum premia. Overall, these results suggest that while the sentiment factor partially co-moves with the market

factor, it captures a distinct behavioral dimension of investor expectations beyond systematic risk exposures. The low correlations with the remaining Carhart factors and the risk-free rate further support the interpretation that the PCA-derived sentiment component represents an independent source of variation relevant to the tail behaviour of returns.

Importantly, Figure 3.2 confirms that PQR, by controlling for unobserved heterogeneity, shows a weaker momentum effect at lower quantiles and a stronger reversal effect at higher quantiles compared to UQR. Specifically, the PQR coefficients at the 0.1 quantile are less negative than the median UQR estimates, indicating a milder effect of momentum in the lower tail. However, at the 0.8 and 0.9 quantiles, the PQR coefficients are more negative than the median UQR estimates, highlighting a stronger momentum reversal at the upper quantiles once unobserved heterogeneity is accounted for. In contrast, the UQR estimates exhibit less variation across quantiles, demonstrating how PQR uncovers momentum's distributional dynamics that traditional models miss. Finally, the risk-free rate (RF) shows no consistent pattern and is statistically insignificant at most quantiles, suggesting limited explanatory power in differentiating cross-sectional returns when controlling for other factors.

In addition, as shown in both Table 3.2 and Figure 3.2, the regressors have an asymmetric impact on the quantiles of future returns, with a more pronounced effect observed in the quantiles below the median. This asymmetry is stronger in the upper quantiles compared to the lower ones. The absolute values of the PC1 parameter estimates are not symmetric around the median, which highlights the increasing importance of sentiment effects on the estimation of the upper quantiles of returns. This asymmetry suggests that negative sentiment, as captured by PC1, exerts a more pronounced downward influence on higher quantiles, indicating that sentiment has a stronger impact on the upper range of the return distribution compared to lower quantiles.

The difference between the PQR and UQR estimates underscores the importance of controlling for unobserved heterogeneity. Without accounting for stock-specific characteristics (in UQR), the influence of sentiment (PC1) appears less pronounced in the upper quantiles. The PQR model reveals that controlling for these fixed effects provides a clearer picture of how sentiment influences different parts of the return distribution, particularly highlighting its stronger impact on high quantiles.

3.8 Robustness tests

To assess the robustness of the results, several alternative PCA specifications are estimated by varying the set of sentiment indicators included in the analysis. In addition to the baseline configuration that incorporates the BW, MCSI, CCI, SSW, AAII, and VIX indices, two alternative specifications are considered: (i) excluding the VIX and (ii) using only survey-based sentiment indices. The resulting first principal components (PC1) exhibit high pairwise correlations across specifications, indicating that the extracted sentiment factor is stable with respect to the inclusion or exclusion of specific sentiment variables. This stability suggests that a common latent sentiment factor is consistently present across diverse sentiment measures. The PCA framework serves to extract this underlying factor by filtering out idiosyncratic noise from individual series and emphasizing their shared variation. The stability of the PCA results across specifications therefore confirms that the extracted component represents a genuine, pervasive sentiment dimension rather than an artifact of a particular variable choice. The corresponding quantile regression estimates show consistent signs and magnitudes across models, particularly at the upper quantiles where high sentiment continues to predict lower future returns. Moreover, correlations between PC1 and observable variables—Carhart Four Factors—reveal that the extracted sentiment factor is closely related to market volatility and equity issuance activity, highlighting its behavioral interpretation. These findings confirm that the identified sentiment component captures a robust and economically meaningful source of variation in return tail behavior.

3.9 Conclusion

To conclude, this chapter investigates how investor sentiment influences the cross-sectional distribution of stock returns, particularly in the tails, using a quantile-based asset pricing framework. Departing from the traditional expected utility approach, this analysis adopts a Panel Quantile Regression (PQR) model to capture heterogeneity in investor behavior and quantify the impact of sentiment across different points of the return distribution. Quantile preference models provide a valuable mechanism for understanding downside risk preferences and asymmetric risk perceptions, making them more adaptable in accounting for sentiment-driven investor behavior. Moreover, by constructing a sentiment index using Principal Component Analysis (PCA) on five widely recognized sentiment measures—including both survey-based and machine learning-derived indicators, this chapter develops a robust composite proxy that more effectively captures the complexity of investor sentiment than the single-indicator approach.

Chapter 4

Time-varying Probability Weighting and Investor Sentiment Toward Tails

4.1 Introduction

Over the past several decades, researchers have accumulated extensive experimental evidence on attitudes toward risk. This evidence reveals that people's risk evaluations often deviate from the predictions of expected utility theory. Empirical studies in financial markets have shown that pricing kernels estimated from index options are non-monotonic (Rosenberg and Engle (2002); Bakshi et al. (2010)) and that risk aversion can even be negative (Aït-Sahalia and Lo (2000a); Jackwerth (2000)), which contradicts the standard theory. According to traditional pricing models, the pricing kernel is expected to decrease monotonically with investor wealth and exhibit strictly positive risk aversion. This contradiction between theory and empirical evidence is known as the pricing kernel puzzle, highlighting that the standard expected utility model may be missing crucial elements needed to accurately capture investor behavior and asset pricing. To address these discrepancies, economists have developed non-expected utility models, including prospect theory (Kahneman and Tversky, 1979a), rank-dependent expected utility (RDEU) (Quiggin, 1982, 1993), and cumulative prospect theory (Kahneman and Tversky (1992)). A key component of these models is the probability weighting function.

Polkovnichenko and Zhao (2013) develop a pricing kernel based on the rank-dependent expected utility framework with a probability weighting function and show that the shape of the weighting function assigns greater emphasis to tail events. Specifically, they find that the weighting function is predominantly inverse-S shaped—indicating an overweighting of tail probabilities—but that it occasionally shifts to an S-shape during certain periods, reflecting an underweighting of tail risks. This finding is consistent with experimental evidence on probability weighting behavior. According to Barberis (2013a), the research progress in understanding the psychology of tail events has largely revolved around the concept of "probability weighting," particularly its applications across different fields of economics.

The probability weighting function from Cumulative Prospect Theory proposed by Kahneman and Tversky (1992), reflects how individuals subjectively overweight or underweight objective probabilities. Intuitively, the probability weighting function serves as a tool to model individuals' risk preferences towards probabilities of ranked events.The primary effect of the weighting function is to overweight the tails of the distribution it is applied to, which is not a bias in belief but a modeling tool aimed at capturing the common preference for positively skewed return distribution Barberis and Huang, 2008. Despite the theoretical connection between probability weighting functions and investor sentiment toward tail events, this relationship remains largely unexplored in empirical studies. Polkovnichenko and Zhao (2013) state that the time series of probability weighting parameters can be used to construct investor sentiment to tail events, but leave it to future work to investigate how this relates to expected returns and volatility. This chapter fills this gap by focusing on examining the relationship between time-varying overweighting preference parameters of probability weighting functions and investor sentiment toward tail events.

In the theoretical section of this chapter, the pricing kernel is expressed as the product of marginal utility and the derivative of the probability weighting function within the framework of non-expected utility models. Since the SDF is estimated non-parametrically, it is not feasible to separately identify both components using a non-parametric approach. Following Polkovnichenko and Zhao (2013), the time series of the parameters of the probability weighting function is estimated by fitting non-parametric estimates of probability weights to a parametric specification. To estimate the probability weighting function without imposing parametric restrictions, the power utility function is employed. The estimation of probability weighting relies on the assumption that, on average, investors' subjective density estimates shaped by behavioral biases through probability weighting functions, should align with the distribution of realized returns (Bliss & Panigirtzoglou, 2004). The subjective density function can be derived by transforming risk-neutral probabilities into real-world probabilities, the linkage provided by the pricing kernel. The discrepancies between the risk-neutral distribution (Q-measure) and the real-world distribution (P-measure), particularly in the tails, arise from investors' risk preferences. The preferences can be captured through the assumed power utility function, enabling a reconciliation between the two distributions. Mathematically, the linkage between them often involves a change of measure. Accordingly, estimating the probability weighting function parameters requires three essential quantities inputs, the risk-neutral, the physical return density, and nonparametric probability weighting function.

In the empirical section, the pricing kernel is derived from the ratio of the risk-neutral density (RND) to the physical density of returns and is subsequently used to estimate the implied probability weighting function. A non-parametric approach is employed to directly estimate the risk-neutral density from the market prices of widely traded S&P 500 (SPX) index options across a wide range of strike prices with the same maturity. This method, first introduced by Breeden and Litzenberger (1978) and later applied to equity index options by Aït-Sahalia and Lo (1998, 2000a) and Jackwerth (2000), derives the SDF based solely on the no-arbitrage principle, without relying on specific utility assumptions. The S&P 500 index is used as a proxy for market return because of its strong correlation with overall wealth in the U.S. economy, making it a reliable representation of U.S. market performance. The rationale behind this approach is that the density implied from option prices reflects both investors' beliefs about future probabilities and their risk preferences, but since it is derived from a no-arbitrage pricing model where all risk is hedged away, it represents a world where investors are indifferent to risk, hence the term "Risk Neutral Density" (RND). The key advantage of the non-parametric approach is its flexibility, as it imposes no restrictions on the option pricing model or the functional form of the pricing kernel. The physical density is derived semi-parametrically using a GJR-GARCH model proposed by Glosten et al. (1993). Additionally, to estimate probability weighting functions non-parametrically, the standard power utility function is assumed. Each of these quantities is estimated separately.

In order to assess how effectively the time-varying parameters of the probability weighting function capture investor sentiment toward tail events, this chapter select five families of probability weighting functions—the two-parameter weighted average power function from Lopes (1987), the one- and two-parameter functions from Kahneman and Tversky (1992), and the one- and two-parameter functions from Prelec (1998). To validate its effectiveness, the time-varying parameters are estimated on a monthly basis and compared against a benchmark investor sentiment proxy, represented by the first principal component (PC1) extracted via the PCA method in Chapter 3. Moreover, time-series regressions are employed to assess the explanatory power of investor sentiment, alongside a set of control variables, in capturing the time variations of the overweighting parameters of tail probabilities.

The primary contribution of this chapter is to fill a gap in the literature by empirically examining the relationship between time-varying overweighting preference parameters of probability weighting functions and investor sentiment toward tail events. By integrating insights from non-expected utility models—particularly prospect theory and cumulative prospect theory—it investigates how subjective probability weighting, especially the overweighting of tail probabilities, relates to investor sentiment toward extreme market outcomes. This chapter provides strong evidence that investor sentiment is linked to time-varying investor preferences, as proxied by the parameters of the probability weighting function. The measures derived from probability weighting functions offer a distinct, behaviorally grounded dimension to the analysis of investor sentiment. The findings complement those of Polkovnichenko et al. (2019), who use the estimated parameters α and β from the Prelec (1998) two-parameter probability weighting function as proxies for investor preferences. In their framework, α is orthogonalized to β, with a lower α indicating a stronger preference for upside potential and a lower β reflecting a stronger preference for downside protection.

In addition, the first principal component (PC1) derived via principal component analysis (PCA) is employed as a benchmark for investor sentiment to examine its relationship with the time-varying parameters of five commonly used probability weighting functions. By aggregating data from multiple sentiment indicators, PCA provides a more comprehensive and robust measure of overall investor sentiment, making it an invaluable tool for capturing the multidimensional nature of sentiment—particularly when compared to single-indicator approaches. Polkovnichenko et al. (2019) demonstrate that investor preferences proxied by the Prelec structural parameters are not a simple reflection of investor sentiment. This is evidenced by their finding that the effect of investor risk preferences on active fund flows remains robust in regressions of fund flows on the Prelec parameters α and β, even after controlling for the Baker and Wurgler (2006, 2007) sentiment index, the National Bureau of Economic Research (NBER) recession indicator, and the market volatility index (VIX). However, in this chapter, the PC1 sentiment proxy exhibits a strong correlation of 80% with the curvature parameter α, suggesting that it serves as a valuable and integrative measure for capturing diverse dimensions of investor sentiment.

The findings show that, among these functions, the Prelec weighting function curvature parameter alpha exhibits the highest 68% correlation with PC1. The findings provide empirical evidence supporting the correlation between the parameters of probability weighting and external sentiment measure in real-world, non-experimental data, contributing to the broader literature on behavioral asset pricing by offering new insights into the dynamic nature of probability weighting and its implications for the pricing kernel puzzle.

4.2 Literature review

Non-expected utility theories, including prospect theory and rank-dependent models, have been broadly employed to address issues in finance, including risk premia, the cross-section of expected returns, and portfolio diversification. For example, Epstein and Zin (1990) employ rank-dependent utility to analyze the equity risk premium. Shefrin and Statman (2000) and Polkovnichenko (2005) explore the effects of probability weighting on portfolio selection and diversification. Barberis and Huang (2008) examine cross-sectional risk premia through the lens of cumulative prospect theory. Additional contributions include studies by Levy and Levy (2004), Nicholas Barberis and Thaler (2006), as well as Chapman and Polkovnichenko (2009). Polkovnichenko and Zhao (2013) demonstrated that individual risk preferences with probability weighting can be used to explain the non-monotonic pricing kernel observed in options markets.

Under this framework, an individual's perceived utility is determined by both a value function and a probability weighting function. Earlier studies on the pricing effects of prospect theory have largely focused on the kink in the value function Barberis et al. (2001) and Benartzi and Thaler (1995). For instance, Benartzi and Thaler (1995) introduced the concept of myopic loss aversion (MLA) based on prospect theory to tackle the equity premium puzzle, while Barberis et al. (2001) illustrate how a dynamic prospect theory model can account for a high equity premium, return predictability, and excess volatility. Additionally, Baele et al. (2017) demonstrate that prospect theory helps resolve the variance premium puzzle. Baele et al. (2019) demonstrate that prospect theory helps resolve the variance premium puzzle.

However, as emphasized by Barberis et al. (2016), a substantial portion of the predictive ability of prospect theory value for stock returns stems from the "probability weighting" component of prospect theory, which leads individuals to overweight the tails of return distributions, capturing the widespread preference for lottery-like gambles. A large body of experimental research has found that probability weighting functions typically exhibit an inverse-S shape, where events in the tails of the outcome distribution are overweighted. Camerer and Ho (1994) provide a comprehensive summary of various studies that support non-linearity of probability weights with greater sensitivity of preferences to tail events. Wu and Gonzalez (1996) reveal the inverse-S shaped probability weights in experimental data through a method that does not rely on any assumptions regarding the functional form of utility or probability weighting functions. Numerous other studies, utilizing both hypothetical and real-payoff experiments, have explored probability weighting functions and found that the inverse-S shape provides an excellent fit for the data. Notable examples include the works of Quiggin (1987), Lopes (1987), Tversky and Kahneman (1992), and Prelec (1998). Berns et al. (2007) contribute further evidence from non-monetary experiments, broadening the applicability of these findings. For an overview of the theoretical and empirical literature on decision-making under risk, see the surveys by Shoemaker (1982), Camerer (1995), and Starmer (2000).

Polkovnichenko and Zhao (2013) propose that time-varying probability-weighting parameters can serve as an indicator of investor sentiment toward tail risk, but leave it to future research to examine how these parameters relate to expected returns and volatility. While subsequent studies have advanced the understanding of how the probability weighting function relates to asset-market dynamics, none have directly tested its effectiveness as a measure of investor sentiment in the way pursued in this chapter. For example, Chabi-Yo and Song (2013) employ both S&P 500 and VIX options to nonparametrically estimate the joint probability weighting function of the distributions of both market return and volatility, finding that while the joint weighting function of return and volatility remains stable over time, the weighting function of the return exhibits time variation that depends on VIX. He et al. (2017) examine how variations in inverse S-shaped probability weighting affect optimal portfolio choice within a rank-dependent utility framework. They show that a stronger inverse S-shaped probability weighting—reflecting greater overweighting of extreme outcomes—typically results in a lower allocation to risky assets, regardless of whether returns are left- or right-skewed, as long as it offers a non-negligible risk premium. Only for lottery-type stocks with low expected returns and extremely high positive skewness does an increase in inverse S-shaped weighting lead to larger portfolio allocations.

Moreover, Félix et al. (2020) focus on the relationship between implied-volatility-sentiment and the overweighting of tail probabilities. Azimi et al. (2024) develop a representative agent asset pricing model with probability weighting and show that the resulting measure of market optimism can be used to construct optimal dynamic investment strategies that consistently outperform both the buy-and-hold strategy and strategies generated by 17 leading equity premium predictors. Most recently, Jha et al. (2025) develops a unified framework that embeds behavioral distortions into rational portfolio optimization by extracting implied probability weighting functions (PWFs) from optimal portfolios. They estimate PWFs that capture nonlinear belief distortions consistent with fear and greed, showing that greater tail fatness amplifies these distortions and that shifts in the term structure of risk-free rates modify their curvature. The findings underscore the importance of jointly modeling return asymmetry and belief distortions in portfolio risk management and capital allocation under extreme-risk environments.

This chapter advances the literature by establishing a direct link between the time-varying overweighting preference parameters of probability weighting functions and investor sentiment toward tail events. Specifically, it introduces a unified sentiment benchmark—the first principal component (PC1) extracted through principal component analysis (PCA)—to examine how changes in probability weighting relate to shifts in aggregate investor sentiment. By synthesizing information from multiple sentiment indicators, the PCA-based benchmark provides a comprehensive and robust measure of overall market sentiment, offering a richer representation of investor attitudes than any single-indicator approach. This framework enables a novel empirical exploration of how sentiment-driven probability distortions shape investors' perceptions of tail risks.

The estimation methodology for the probability weighting function in most previous studies relies on the theoretical relationship between state prices and state probabilities, as established by C.-f. Huang and Litzenberger (1988). This approach typically involves separately estimating the risk-neutral density from options prices and the objective density from historical prices of the underlying asset. These two independently derived functions are then used to infer the pricing kernel, which is subsequently employed to construct the implied probability weighting function. The literature on extracting and interpreting the risk-neutral distribution from market option prices is extensive, and it expands significantly when considering research on implied volatilities and modeling the returns distribution. These methods generally fall into three categories: fitting a parametric density function to the market data, approximating the RND using a nonparametric technique, or developing a returns process model that generates the empirical RND as the density for the underlying asset's value on the option's expiration day M. Figlewski, 2010 .

Breeden and Litzenberger (1978) and Banz and Miller (1978) pioneered the concept of extracting the risk-neutral density (RND) from option prices by utilizing the second-order derivative of the option price with respect to the strike price. This approach fundamentally connects option pricing to the distribution of possible future outcomes of the underlying asset under the risk-neutral measure. Since then, numerous methods have been developed to enhance the accuracy and reliability of extracting RNDs from observed option prices. Notable examples include the works of Aït-Sahalia and Lo (2000a), Jackwerth (2000), Rosenberg and Engle (2002), and Bliss and Panigirtzoglou (2004), M. Figlewski (2010), Polkovnichenko and Zhao (2013), Félix et al. (2020). The key advantage of M. Figlewski (2010)'s method, compared to other approaches, is its ability to separately extract both the body and the tails of the distribution, allowing it to account for fat tails. This aligns with studies that suggest people tend to agree on the central tendency of a distribution, but disagree on the entire distribution.

The estimation of the physical distributions has primarily involved two key approaches. Early studies, such as those by Jackwerth (2000, 2004b) and Aït-Sahalia and Lo (2000a), employed kernel densities, which do not require any specific distributional assumptions aside from the stationarity of returns. However, in the presence of time-varying volatility and structural breaks, more recent work has shifted toward GARCH models. Notably, Rosenberg and Engle (2002), Barone-Adesi et al. (2008), and Barone-Adesi and Dall'O (2010) have applied the GJR-GARCH model (Glosten et al., 1993) to capture the conditional volatility in historical returns, thereby improving the accuracy of the estimated objective density.

4.3 Model specifications

4.3.1 Utility function with probability weights

Expected utility (EU) theory is designed to capture utility preferences through a real-valued function that is increasing (higher values correspond to stronger preferences) and concave (diminishing marginal utility), based on the utilitarian theory of choice under certainty. In expected utility theory, the decision maker maximizes the expected utility, which is the weighted sum of utilities of possible outcomes, with the weights being the probabilities of those outcomes. In expected utility theory, the decision maker seeks to maximize the expected utility, which is the weighted sum of the utilities across possible across states, with the weights corresponding to the probabilities of those states. This reflects the assumption that expected utility functionals are linear in probabilities. When the states are equiprobable, the expected utility simplifies to the sum of the utilities of all states. Formally, for any prospect $\{\boldsymbol{x};\boldsymbol{p}\}$, where $\boldsymbol{x} = \{x_1, x_2, \ldots, x_n\}$ represents possible outcomes and $\boldsymbol{p} = \{p_1, p_2, \ldots, p_n\}$ represents the probabilities associated with these outcomes, the expected utility $V(\{x;p\})$ is defined as:

$$V(\{x;p\}) = \sum_{i=1}^{n} p_i U(x_i) = \mathrm{E}[U(x)] \tag{4.1}$$

where $U : x \rightarrow \mathbb{R}$ is a utility function mapping the set of possible outcomes to real numbers, representing the satisfaction associated with that outcome.

Non-expected utility models modify the expected utility framework by transforming objective probabilities through a probability weighting function. The key distinction is that utility is weighted by a probability weighting function assigned to ranked outcomes. As a result, the utility of an outcome depends not only on its objective value but also on its rank among possible outcomes, adjusted by the probability weighting function. This chapter considers the continuous-state case of RDEU developed by Quiggin (1993), which is defined based on outcomes ranked from the worst to the best. In this context, possible outcomes are ranked by investor wealth w, which is assumed to be a random variable with a cumulative distribution function $P(w)$ and a probability density function $p(w) = P'(w)$. A probability weighting function $G(P)$ is characterized as a continuous, differentiable, non-decreasing function $G(\cdot) : [0, 1] \to [0, 1]$, with $G(0) = 0$ and $G(1) = 1$. The corresponding probability weighting density is defined as $Z(P) \equiv G'(P) \geq 0$. The RDEU can be written as:

$$U(w) = \int u(w)\, dG(P) = \int u(w)\, G'(P)\, dP = E\{u(w)Z(P)\} \tag{4.2}$$

where $u\,(w)$ is the utility of outcome w, P represents the objective probability distribution of outcomes. The probability weighting function $G(P)$ transforms the objective probability measure P to reflect the decision maker's subjective perception of risk. Since decision weights sum to 1 , it follows that $E[Z] = \int dG(P) = 1$. When $G(P) = P$ ($Z = 1$), RDEU simplifies to standard expected utility (EU).

The theoretical constraints on the weighting function are minimal, and its shape is ultimately determined empirically. Curvature and elevation are two distinct properties of probability weighting functions, and they affect how individuals perceive probabilities in different ways. Curvature refers to the concavity or convexity of the probability weighting function, which can be characterized as concave, convex, or S-shaped. Elevation refers to the vertical shift of the function, reflecting the overall level of optimism or pessimism in the decision maker's perception of probabilities. Numerous experimental studies, including those by Camerer and Ho (1994), Wu and Gonzalez (1996), and Kahneman and Tversky (1992), have shown that individuals often overweight events at the tails of the distribution relative to events in the middle. This characteristic is typically represented by an inverse-S shaped probability weighting function G, which corresponds to a U-shaped weighting density z. However, as noted in Camerer and Ho (1994), some experiments have estimated S-shaped weighting functions, where tail events are underweighted instead.

In this chapter, along with nonparametric estimation, three alternative parametric forms of weighting functions commonly used in the literature are considered. These forms can capture a wide range of shapes, including concave, convex, inverse S-shaped, and S-shaped functions, providing the flexibility necessary to accurately approximate the observed data. The simplest of these parametric forms, originally introduced by Lopes (1987), is a weighted average power function.

$$G(P) = \lambda P^{\phi} + (1-\lambda)\left(1-(1-P)^{\phi}\right), \quad \phi \geq 1, \quad \lambda \in [0,1] \tag{4.3}$$

This weighting function is highly effective in modeling various risk preferences. The power parameter ϕ controls over-weighting of the probabilities in the tails by changing the curvature of the weighting function. When ϕ takes integer values, the weighting function simplifies to a polynomial form, enabling the semi-analytically comparative statics effects of the transformation on the moments of a distribution. The parameter λ governs the balance between "optimism" and "pessimism" by changing the relative strength of the concave and convex segments. At the extremes of its range, $\lambda = 0$ corresponds to a globally concave, pessimistic weighting function, while $\lambda = 1$ represents a globally convex, optimistic weighting function. Thus, λ can be interpreted as a "sentiment" parameter.

Another commonly used weighting functions in the literature, proposed by Kahneman and Tversky (1992) and Prelec (1998), are as follows:

$$G\_KT1(P) = \frac{P^{\alpha}}{(P^{\alpha} + (1-P)^{\alpha})^{\frac{1}{\alpha}}} \tag{4.4}$$

$$W^{+}(P) = \frac{P^{\gamma}}{\left(P^{\gamma} + (1-P)^{\gamma}\right)^{1/\gamma}}, \quad W^{-}(P) = \frac{P^{\delta}}{\left(P^{\delta} + (1-P)^{\delta}\right)^{1/\delta}} \tag{4.5}$$

The tail overweight parameter γ and δ determine the curvature of the probability weighting function for gains and losses, respectively. When γ and δ parameters are close to one, the weighting function approximates neutral (unweighted) probabilities, whereas parameters near zero indicate a significant overweighting of small probabilities, leading to inverse S-shaped probability distortion functions.

A limitation of the Cumulative Prospect Theory (CPT) model is that it permits non-strictly increasing functions, which can hinder invertibility. For this reason, more recent literature on probability distortion functions tends to favor strictly monotonic alternatives, such as Prelec (1998) function, $w(p) = e^{-(-\ln(p))^{\beta}}$, for probability weighting.

$$G\_Prelec1(P) = \exp\left(-(-\log(P))^{\alpha}\right) \tag{4.6}$$

Experimental studies often find that $\alpha < 1$, indicating an inverse- S shaped weighting function that overweights probabilities in the tails. However, there is variability across studies, with some estimating $\alpha > 1$, which corresponds to underweighting the tails, as highlighted by Camerer and Ho (1994). While single-parameter weighting functions are appealing for experimental applications due to their simplicity, they are less suitable for our analysis.

To address this, Prelec (1998) introduced a two-parameter weighting function, which is used in this empirical work:

$$G\_Prelec2(P) = \exp\left(-\beta(-\log P)^{\alpha}\right) \tag{4.7}$$

where $\alpha > 0$ controls the curvature of the weighting function, thereby determining the extent to which tail probabilities are overweighted. Experimental evidence typically finds $\alpha \in [0.5, 1]$, corresponding to an inverse-S shaped curve that reflects overweighting of probabilities in the tails of the distribution. As α is lower, the inverse-S shape becomes more pronounced, and the overweighting of the tails becomes strong. When $\alpha = 1$, the function reduces to a straight 45-degree line, corresponding to the Expected Utility benchmark without probability distortion. $\beta > 0$ serves as the elevation parameter that governs the relative strength of the concave and convex segments of the weighting function. When setting $\alpha = 1$, the case of $\beta < 1$ produces a concave function, indicating risk aversion through overweighting of the left tail, whereas $\beta > 1$ yields a convex function, reflecting upside-seeking behavior through overweighting of the right tail. Thus, in the compound representation, the coefficient β can independently amplify or attenuate risk attitudes in the left tail through its interaction with α. A lower β corresponds to a uniform increase in risk aversion, whereas a lower α implies stronger risk aversion in the left tail accompanied by greater risk seeking in the right tail. The introduction of two-parameter extensions in weighting functions allows for a more flexible representation of investor preferences.

4.3.2 Pricing kernel with probability weighting function

The estimation methodology for the probability weighting function builds on the well-established theoretical relationship between state prices and state probabilities (C.-f. Huang & Litzenberger, 1988)), along with the extension of this concept to continuous distributions and the connection between risk-neutral and objective density functions.

Following Bliss and Panigirtzoglou (2004), the formulaiton can be expressed as follows:

$$\frac{f_p\left(S_T\right)}{f_q\left(S_T\right)} = \xi \frac{U'\left(S_T\right)}{U'\left(S_t\right)} \equiv \zeta\left(S_T\right) \tag{4.8}$$

where $U\left(S_T\right)$ is the representative investor's utility function, S_T denotes the value of the $S\&P$ 500 index at a future time T, $f_q\left(S_T\right)$ represents the risk-neutral density function (implied by market prices), $f_p\left(S_T\right)$ represents the objective density function (investors' true beliefs), ξ is a constant and $\zeta\left(S_T\right)$ represents the pricing kernel. A cornerstone result behind this relationship is that the ratio of the true objective density f_q to the risk-neutral density f_p is proportional to the representative investor's marginal utility (the pricing kernel). This concept shows discrepancies between the risk-neutral density f_q and the objective density f_p arise from investors' risk preferences, and there is a utility function capturing investors' preferences allowing to match both density functions under the assumptions of complete, frictionless markets and a single asset. Therefore, knowing any two of the three functions-the risk-neutral density, the objective density, or the pricing kernel (equivalently the utility function)-allows the inference of the third.

This chapter considers a simplified setting to focus on the properties of the SDF with probability weighting functions. Following Jackwerth (2000), consider a static model in which an investor's utility is a function of their terminal wealth w, where investor is endowed with an initial wealth w_0 and has access to N+1 traded securities. Suppose a complete market, there are N securities and S states of nature, the payoff of security n in state s is denoted as R_{is}, investors derive utility from their wealth and prefer more wealth to less in each state. The investor begins with an initial wealth w_0, which can be allocated across $N+1$ traded securities. Each security $i = 0, 1, \ldots, N$ offers a return R_i^s across the states $s = 0, 1, \ldots, S$, and the share of the portfolio invested in security i is denoted by θ^i. The portfolio is constrained by $\theta^0 = 1 - \sum_{i=1}^{N} \theta^i$ and the terminal wealth received in state s is given by the sum of the returns of all assets, weighted by their respective portfolio shares, as shown in the equation:

$$w_s = w_0 \left(\sum_{i=0}^{N} R_s^i \theta^i \right) \tag{4.9}$$

This expression reflects how the initial wealth is transformed based on the returns and portfolio shares. By selecting portfolios, investors are essentially deciding how to distribute their wealth across different possible states of nature, constrained by their available wealth and the security prices.

To obtain the necessary first-order optimality conditions, the results in Ai (2005) and Carlier and Dana (2003) are utilized, which demonstrate the differentiability of the RDEU functional with respect to continuously distributed random variables. Ai (2005) shows that when X is continuously distributed (with a density), the derivative of the RDEU functional with respect to a continuously distributed random variable Y is given by:

$$\frac{\partial}{\partial \alpha} U(X + \alpha Y)\bigg|_{\alpha=0} = E\left[u'(X) Z\left(P_X\right) Y\right] \tag{4.10}$$

Using the equation (4.10), the first-order condition for optimal portfolio shares θ^n in equation (4.9) is derived:

$$E\left[u'(w) Z(P)\left(R^i - R^0\right)\right] = 0, \quad \text{for } i = 1, \ldots, N \tag{4.11}$$

where $u'(w)$ is the marginal utility of terminal wealth, which reflects the investor's risk preferences and how they value wealth in different states of the world. $Z(P)$ is the weighting function density.

The Euler equation (4.10) is commonly expressed as $E\left[m\left(R^k - R^0\right)\right] = 0$ with the pricing kernel can be defined by multiplying the marginal utility by the derivative of the probability weighting function:

$$m = \frac{u'(w)}{u'(w_0)} Z(P) \tag{4.12}$$

The marginal utility of terminal wealth, $u'(w)$, reflects the investor's risk preferences and valuation of wealth across different states. The subjective probability weighting function density, $Z(P)$, represents probability distortions, capturing subjective probability weighting driven by behavioral biases or ambiguity aversion. The SDF m is positive everywhere, ensures no arbitrage, and implies a linear pricing operator, consistent with standard asset-pricing theory, regardless of the weighting function's shape.

Following Polkovnichenko and Zhao (2013), if denoting the gross return on total investor wealth as $R \equiv \frac{w}{w_0}$, by assuming a representative agent who derives utility from the market return with power utility $u(w) = \frac{w^{1-\gamma}}{1-\gamma}$ and normalizing initial wealth $w_0 = 1$, the marginal utility of wealth becomes:$u'(R) = R^{-\gamma}$, by combining with the state-price representation in equation (4.8), the SDF under RDEU can be rewritten as follows:

$$m = \frac{f_q}{R_f f_p} = u'(R)Z(P) = R^{-\gamma} Z(P) \tag{4.13}$$

This equation shows that, in a single-period model with initial wealth w_0, by normalizing initial wealth to 1, the ratio of marginal utilities simplifies to $m = u'(R)$, thus the pricing kernel is proportional to marginal utility, which is consistent with the idea that the SDF reflect the marginal utility of wealth in the future state, relative to the current state (normalized to 1). Normalization effectively removes the need to explicitly track initial wealth, reducing the complexity of the expression.

Using equation (4.13), the probability weighting density $Z(P)$ can be given:

$$Z(P) = \frac{m}{u'(R)} = \frac{f_q}{(R_f f_p)\, u'(R)} \tag{4.14}$$

where f_q and f_p denote the risk-neutral density and physical density of R.

4.4 Estimation Methodology

4.4.1 Non-parametric estimation for risk-neutral distribution

A more effective approach to extracting the risk-neutral distribution, compared to directly fitting a function of call option prices across strike prices, is to first fit a smooth function to implied volatilities across strike prices. The corresponding call option prices function are then calculated from the fitted implied volatilities. The risk-neutral probability distribution is derived by taking the second derivative of the fitted call price function with respect to strike prices, following the Breeden-Litzenberger insight. The key benefit of working with implied volatilities is that they tend to exhibit more consistent magnitudes across strike prices compared to call option prices, which can vary significantly. This chapter adopts the fast and stable curve-fitting method proposed by Jackwerth (2004a) to estimate the risk-neutral density (RND) from observed option prices. The procedure begins by transforming market option prices into Black-Scholes implied volatilities, denoted $\{\bar{\sigma}_i\}_{i=1}^{I}$, and proceeds by solving the following optimization problem:

$$\min_{\sigma_j} \underbrace{\frac{\Delta^4}{2(J+1)} \sum_{j=0}^{J} \left(\sigma_j''\right)^2}_{\text{Smoothness term}} + \underbrace{\frac{\lambda}{2I} \sum_{i=0}^{I} \left(\sigma_i - \bar{\sigma}_i\right)^2}_{\text{Fit term}} \tag{4.15}$$

where the smooth term measuring the curve curvature enforces that the fitted implied volatility curve $\{\sigma\}_{j=1}^{J}$ doesn't have large oscillations or kinks, which is achieved by minimizing the sum of the squared second derivatives of the implied volatility curve with respect to strike price. This is important because a smooth implied volatility curve is more consistent with the no-arbitrage principle and reflects a well-behaved risk-neutral distribution. The second derivative is typically approximated using a central finite difference formula: $\sigma_j'' \approx \frac{\sigma_{j+1}-2\sigma_j+\sigma_{j-1}}{\Delta^2}$. The fit term ensures that the fitted implied volatilities σ_i which are taken from the implied volatility curve σ_j at the observed strike prices K_i passes close to the market-observed implied volatilities $\bar{\sigma}_i$, which is achieved by minimizing the sum of squared errors $\sum_{i=0}^{I} \left(\sigma\left(K_i\right) - \bar{\sigma}_i\right)^2$

The observed implied volatilities $\{\bar{\sigma}_i\}_{i=1}^{I}$, derived from discrete market option prices at specific strikes K_i, are inherently discrete, inconsistent and noisy due to the limited set of traded strikes. A fine grid with J points introduces additional strike prices K_j beyond those provided by the market, with the grid spacing Δ controlling the grid's density and, by extension, the resolution of the fitted implied volatility curve, enabling to create a smooth, continuous curve from discrete, noisy market data.

Finally, the smooth implied volatility curve is converted back into a call option price curve, where the second derivative represents the compounded risk-neutral density. This method builds on the foundational insight of Breeden and Litzenberger (1978), which established that the risk-neutral density can be obtained as the scaled second derivative of the call option price with respect to the strike price: $q(K) = R_f \frac{\partial^2 \text{ Callprice } (K)}{\partial K^2}$. This formula highlights that the RND can be derived directly from option price data, provided the second derivative of the call price curve is accurately estimated. The key benefit of employing this data-driven density shape rather than a parametric distribution is its ability to capture empirical features such as fat tails and negative skewness, which conventional parametric models like the normal distribution fail to represent adequately.

4.4.2 Semi-parametric estimation for Physical distribution

Following Engle (2002), this chapter employs the asymmetric GARCH model of Glosten et al. (1993) to estimate the physical distribution of returns. While standard GARCH models perform well empirically in modeling conditional variances, they assume a symmetric response of volatility to positive and negative innovations, which can lead to systematic underestimation of risk during market downturns. The GJR-GARCH model extends the standard GARCH framework to account for volatility clustering and the leverage effect, which refers to the tendency of volatility to increase more after negative shocks than positive ones of the same magnitude. The GARCH framework allows flexible assumptions about the distribution of the innovations, which provides flexibility in capturing real-world return dynamics. Moreover, to ensure comparability with the time-varying risk-neutral distribution, the physical distribution is estimated on a monthly basis using the GJR-GARCH model. Sensitivity analysis confirms that the main results are robust to alternative volatility specifications, such as the symmetric GARCH(1,1), though the asymmetric GJR-GARCH model achieves a better fit in periods marked by pronounced negative shocks, consistent with its theoretical advantages.

The daily log return r_t is modeled as the sum of the constant mean μ of the return and an error term or shock (innovations) ε_t :

$$r_t = \ln\left(\frac{S_t}{S_{t-1}}\right) = \mu + \varepsilon_t \tag{4.16}$$

The S_t and S_{t-1} are the price of the S&P 500 index at time t and $t-1$, respectively. The $\varepsilon_t \sim N(0, \sigma_t)$ is the raw residual (actual return minus expected return), which is modeled as $\varepsilon_t = z_t\sigma_t$ using a GARCH process to account for time-varying volatility. z_t represents the innovations, and the time-varying volatility σ_t can be recursively computed by the $\mathrm{GJR}-\mathrm{GARCH}(1,1,1)$ formula:

$$\sigma_t^2 = \omega + (\alpha + \gamma I_{t-1})\epsilon_{t-1}^2 + \beta\sigma_{t-1}^2, \quad where \quad I_{t-1} = \begin{cases} 0 & \text{if } r_{t-1} \geq \mu, \\ 1 & \text{if } r_{t-1} < \mu. \end{cases} \tag{4.17}$$

where σ_t^2 is the conditional variance at time t, $\epsilon_{t-1} = r_{t-1} - \mu$ represents the lagged innovation or return shock, and I_{t-1} is an indicator function capturing the asymmetry in volatility response, ω denotes the long-run average volatility, α reflects the sensitivity of volatility to past shocks (ARCH effect), β captures volatility persistence (GARCH effect), and γ models the asymmetric response to negative shocks. Specifically, when returns fall below the mean ($I_{t-1} = 1$), the impact of shocks on volatility is amplified by γ, thus allowing the model to capture the empirical observation that bad news tends to increase volatility more than good news.

The conditional distribution of returns is determined by the time-varying variance σ_t and the distribution of the innovations (standardized residuals) z_t. The innovations (z_t) introduce randomness into the simulation process and are the primary drivers of return distributions under the GARCH framework. In the estimation step, the innovations ($z_t \sim \mathcal{N}(0,1)$) is assumed as the standardized shock with zero mean and unit variance that represents the random noise is a zero-mean white noise driving the process. In the simulation step, the actual observed (filtered) innovations $\hat{z}_t = \frac{\hat{\epsilon}_t}{\sqrt{\widehat{\sigma}_t}}$ from its empirical estimation are drawn. By sampling from this empirical distribution, the non-normal distribution characteristics of financial returns can be captured. The innovations typically show fat tails and skewness when drawn from real-world financial data, making them a better alternative to standard normal innovations.

Subsequently, the estimated conditional volatility $\widehat{\sigma}_t$ and standardized residuals $\widehat{z}_t$ are used simulate the daily paths of index log returns at each time t. To better capture the time-varying characteristics of the innovations distribution, for each simulated path, the innovations are drawn with replacement from a one-year rolling window of time t. A one-year rolling window ensures the innovations reflect recent dynamics without being overly volatile (small window) or overly smooth (large window). This rolling-window approach ensures the simulated innovations reflect the local time-varying nature of volatility and provides enough data for sampling.

A large sample of 250,000 paths is generated at each time t to minimize variance in the estimated moments and distribution function under the physical measure, thereby ensuring that the principal source of variability in the probability weighting function is the risk-neutral estimates rather than sampling variation under the physical measure. In the context of GJR-GARCH models, using historical return data $\{r_t\}$, the parameters in equation (4.17) are estimated through Maximum Likelihood Estimation (MLE), the estimation involves maximizing the conditional log-likelihood under the assumption of normally distributed random variables. When kernel density is used, the historical data is based on non-overlapping 30 day returns that match the maturity of the underlying options because the option payoff depends on returns over that time horizon. To remain consistent with this setup, the simulated daily returns using the estimated parameters from the GJR-GARCH model must also be aggregated to 30 day returns.

4.4.3 Nonparametric estimation of probability weighting functions

Given physical cumulative distribution $P(R)$, its density $p(R)$, risk-neutral cumulative distribution $Q(R)$, and its density $q(R)$, for a specific return R_0 with a corresponding probability P_0, such that the probability of realizing R_0 is given by $P(R_0) = P_0$, the probability distribution function over returns can be expressed as follows:

$$
\begin{aligned}
G(P_0) &= \int_0^{P_0} Z(P)dP \\
&= \int_0^{R_0} Z(P(R))p(R)dR \\
&= c\int_0^{R_0} \frac{q(R)}{u'(R)}dR \\
&= c\left[\frac{Q(R_0)}{u'(R_0)} - \int_0^{R_0} Q(R)\frac{u''(R)}{u'(R)^2}dR\right]
\end{aligned}
\tag{4.18}
$$

where $c = \left(\int_0^{\infty} \frac{q(R)}{u'(R)}dR\right)^{-1}$ is the normalization constant, ensuring $G(1) = 1$. When the utility function is linear, the marginal utility $u'(R)$ is constant, and $u''(R) = 0$, meaning the decision-maker is risk-neutral. In this case, the probability weighting function simplifies to $G(R_0) = Q(R_0)$, indicating that probability weighting purely reflects the change of measure driven by risk-neutral distribution $Q(R)$. The probability distribution function $Q(R)$ (or the CDF) can often be estimated more efficiently than the PDF, especially in finite samples. The ECDF is a non-parametric estimator and is much easier to compute directly from observed data without requiring complex smoothing, kernel functions, or density estimation techniques.

To avoid imposing any prior restrictions on the components of the SDF, only the combined SDF can be estimated and the marginal utility and the probability weighting density can not be identified separately. Thus, the utility function is assumed as the standard power expected utility (CRRA) commonly used in empirical and theoretical asset pricing, ensuring only that the utility is non-convex. This approach enables us to derive nonparametric estimates of the probability weights. Finally, the nonparametric estimates of the probability weights are employed to fit the parametric specification of the weighting function.

4.4.4 Parametric estimation of probability weighting functions

The parameters of the probability weighting function are estimated by minimizing the distance between the values implied by the parametric function and the nonparametric estimates of the probability weighting functions. The objective function for minimization can be expressed as follows:

$$(b^*) = \arg\min_b \sum_{i=1}^{n} \left[G_{\text{Nonpara}} \left(P_i; b \right) - \hat{G} \left(P_i \right) \right]^2 \tag{4.19}$$

where: - b^* are the optimal parameters of the parametric weighting function. - G_{Nonpara} represents the parametric probability weighting function. - $\hat{G}(P_i)$ represents the nonparametric estimate of the probability weighting function at P_i. This formulation seeks the values of the parameter vector b that minimize the sum of squared differences between the parametric function-implied values and the nonparametric estimates.

4.4.5 Tails overweight parameters and PC1 sentiment

The relationship between time-varying overweighting of small probabilities and investor sentiment is examined through multivariate OLS regressions, where each parameter of the five parametric probability weighting functions serves as the dependent variable.

$$\text{Parameter}_{i,t} = c + \beta_1 \text{PC1}_t + \beta_2 \text{VIX}_t + \beta' X_t + \epsilon_t \tag{4.20}$$

The explanatory variables include PC1, the benchmark sentiment proxy derived as the first principal component, and VIX, the market volatility measure. A group of control variables from those examined by Welch and Goyal (2008) as potential predictors of the equity market is also incorporated, including E12, the twelve-month moving sum of earnings of the S&P 5000 index, B/M, the book-to-market ratio, NTIS, the net equity expansion, TBL, the risk-free rate, INFL, the annual inflation rate, CORPR, the corporate bond spread, SVAR, the stock variance, and CSP, the cross-sectional premium. By incorporating these variables, the model isolates the incremental effect of sentiment on the overweighting of small-probability events, providing insight into probability distortions in investor behavior.

4.5 Data

The choice of options can significantly affect the accuracy of risk-neutral density (RND) estimates derived from option prices. The End-of-Month (EOM) option series, expiring on the last trading day of each calendar month, align directly with calendar-month sentiment variables and are thus selected for estimating the risk-neutral density (RND). This alignment enables a clearer reflection of market sentiment and trading behavior within standard monthly cycles. The S&P 500 index serves as the preferred asset that exhibits a strong correlation with overall wealth in the economy, as it is widely regarded as a reasonable proxy for market returns. As one of the most active index options markets globally, SPX options have been widely studied in empirical research, including by Bakshi et al. (1997), Aït-Sahalia and Duarte (2003), Rosenberg and Engle (2002), Polkovnichenko and Zhao (2013), Cuesdeanu and Jackwerth (2018), and Sichert (2018) among others.

All end-of-month S&P 500 options from January 2014 to December 2022 are obtained from OptionMetrics, using the first trading day of each month as the observation date. For each observation date, the midpoint of bid and ask prices serves as the option price, reducing pricing noise. To filter out illiquid options, option quotes with zero trading volume, a best bid price below 0.50, deep in-the-money (moneyness outside the range $0.8 \leq moneyness \leq 1.2$), and those allowing arbitrage across strikes are excluded. Deep ITM options behave more like the underlying asset and less like options, whose prices are heavily influenced by intrinsic value rather than time value or volatility. Removing them helps focus on options that are more responsive to changes in implied volatility, time decay, and other factors relevant to pricing kernel analysis.

Table 1 presents summary statistics for option quotes during the sample period, categorized by moneyness, including the total number of quotes, the Black-Scholes (BS) implied volatility (mean, minimum, and maximum), as well as the average trading volume and open interest for call and put options that are 20%, 10%, 5%, and 1% out-of-the-money (OTM). OTM puts are defined as options with a moneyness below 1, while OTM calls have a moneyness above 1. On average, there are approximately 169 options per month, with OTM puts outnumbering OTM calls. The average BlackScholes implied volatilities follow the well-documented "smirk" pattern observed in option pricing literature. Additionally, the average trading volumes indicate good liquidity, particularly for options closer to at-the-money.

Table 4.1: Summary Statistics of Out-of-the-Money (OTM) Option Quotes

		Number_of Quotes	Implied_Vol			Avg_Trading Volume	Avg_Open Interest
			Avg	Min	Max		
Option_Type	Moneyness						
OTM Put	[0.8, 0.9)	5562	0.31	0.19	0.70	138	1657
	[0.9, 0.95)	3065	0.23	0.13	0.62	398	2418
	[0.95, 0.99)	2467	0.19	0.08	0.57	401	2108
	[0.99, 1.0]	618	0.16	0.07	0.53	374	1357
OTM Call	(1.0, 1.01]	612	0.15	0.06	0.52	532	2373
	(1.01, 1.05]	2453	0.13	0.05	0.51	477	2150
	(1.05, 1.1]	2204	0.14	0.07	0.46	132	1208
	(1.1, 1.2]	1753	0.19	0.12	0.39	29	912

4.6 Empirical results

4.6.1 Risk-neutral and Physical Distribution

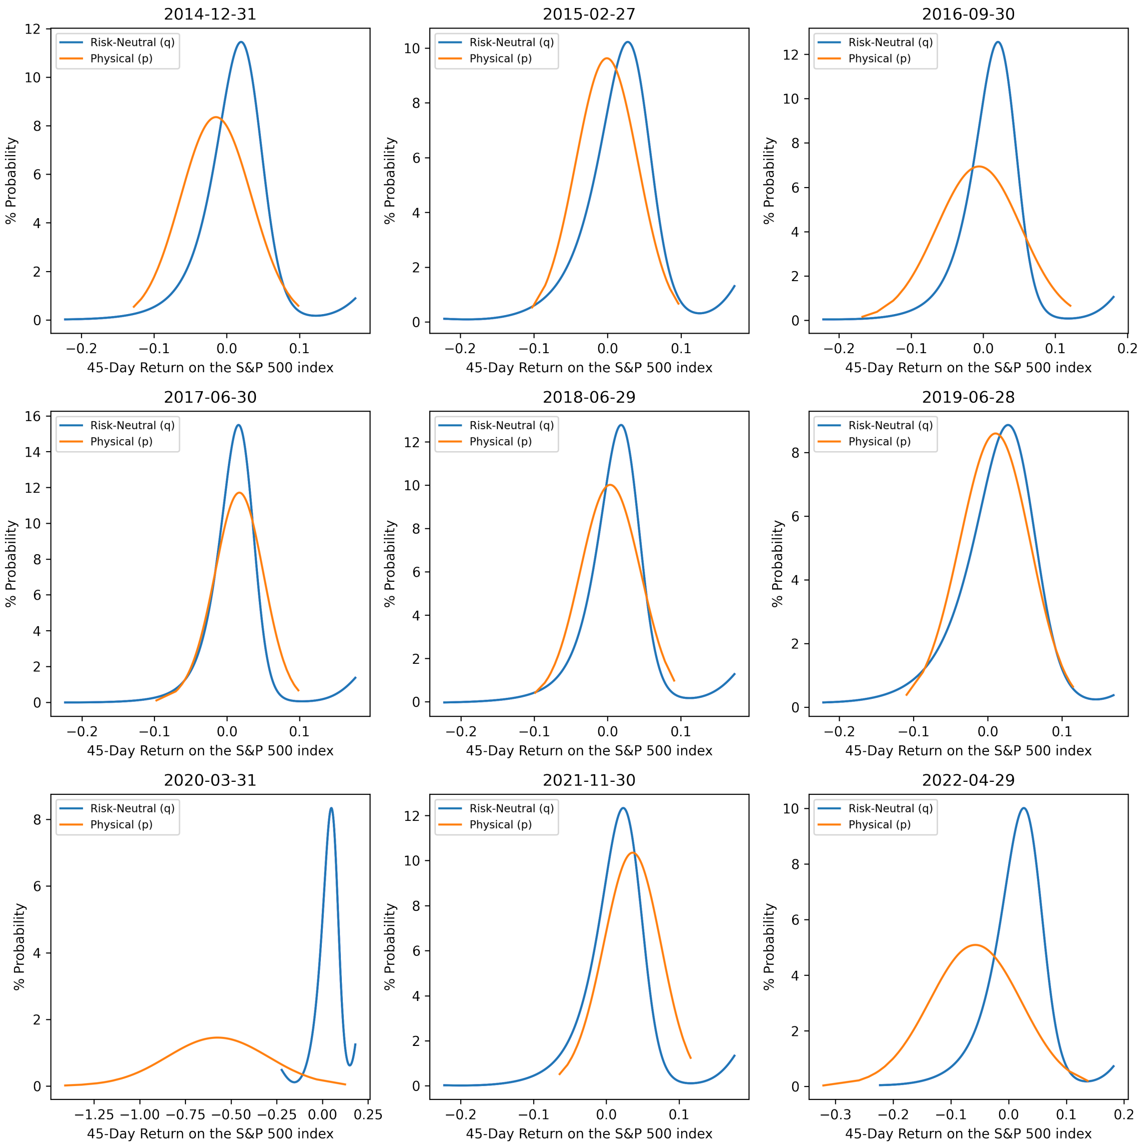

Figure 4.1: Selected representative results for Risk-neutral and Physical distributions. The Physical distributions are calculated with the same return horizon as the time-to-expiration of the options.

Figure 4.1 presents selected examples of risk-neutral (blue curves) and physical (orange curves) return distributions for the S&P 500 index, each evaluated over the same horizon as the time to expiration of the respective options. As expected, the risk-neutral densities place relatively more mass on the left tail (reflecting higher implied downside risk), while the physical densities-obtained via GJR-GARCH estimation and residual bootstrapping-exhibit shapes more consistent with historical return behavior. An important feature of these simulated physical distributions arises from the residual bootstrapping procedure. Because the algorithm randomly draws past innovations from historical residuals, large outliers (such as the extreme negative and positive residuals during March and April 2020) can be included with notable frequency. Consequently, the bootstrapped physical distributions can become inflated in their tails-particularly in scenarios where volatility spikes were exceptionally pronounced (e.g., the COVID-19 crash). While this feature captures real-world tail risks, it also means that the resulting simulated gross returns may appear higher or lower than anticipated for certain time periods. Overall, these comparisons underscore the distinct perspectives offered by risk-neutral versus physically estimated return distributions and highlight the practical implications of residualbased simulation in capturing (and potentially magnifying) historical extreme events.

Table 4.2: Estimation of Conditional Density of Daily S&P 500 Log-Returns (2012-2022)

Panel A: Parameter estimates

Parameter	Estimate	Std. Error	p_value
mu	0.006457	0.004639	0.163955
omega	0.000002	0.000000	0.000000
alpha[1]	0.050005	0.485945	0.918041
gamma[1]	0.199996	0.284569	0.482177
beta[1]	0.829997	0.113988	0.000000

Panel B: Properties of the estimated standardized innovations

Skewness	Excess Kurtosis	Normality test p-value	ARCH test p-value	Serial correlation test p-value
-0.615688	3.239504	0.000000	0.491528	0.219915

In table 4.2, Panel A reports the estimated coefficients for the GJR-GARCH model specified in equation **??**, which includes the following components: ω (long-run baseline variance), α (ARCH), γ (asymmetric term), β (GARCH). The estimation employs an GJR-GARCH$(1,1)$ model for daily S&P 500 log-returns spanning 2012 to 2022, using maximum likelihood estimation. The results indicate that the estimated mean return ($\mu = -0.0033, p < 0.001$) is negative and significant, indicating an overall slight downward drift in daily returns during the sample period. The estimated ARCH parameter, $\alpha[1] = 0.05(p = 0.0385)$, is statistically significant at the 5% level, indicating that current volatility is moderately sensitive to recent shocks. The GARCH coefficient, $\beta[1] = 0.83(p = 0.0000)$, suggests high volatility persistence, meaning volatility shocks have a strong and lasting impact on future volatility. Consequently, volatility clusters are likely to persist for extended periods. In addition, the asymmetric term ($\gamma = 0.20, p < 0.001$) is strongly significant, which means negative shocks (bad news) substantially increase future volatility more than positive shocks (good news) of similar magnitude, clearly confirming a leverage effect. This asymmetry implies a negatively skewed unconditional distribution, emphasizing the importance of modeling downside risk explicitly. Overall, the findings strongly support the necessity of incorporating asymmetric volatility dynamics when modeling stock returns.

Panel B presents the properties of the estimated standardized innovations. The normality test rejects the hypothesis that the residuals (standardized innovations) follow a standard normal distribution. The ARCH test (Engle, 1982) assesses unexplained stochastic volatility, rejecting the null hypothesis of homoskedasticity. The Ljung-Box test (Ljung & Box, 1978) evaluates unexplained serial correlation, rejecting the null hypothesis of no autocorrelation.

4.6.2 Probability weighting functions

Figure 4.2: Selected representative results for Prelec parametric estimates of the probability weighting functions

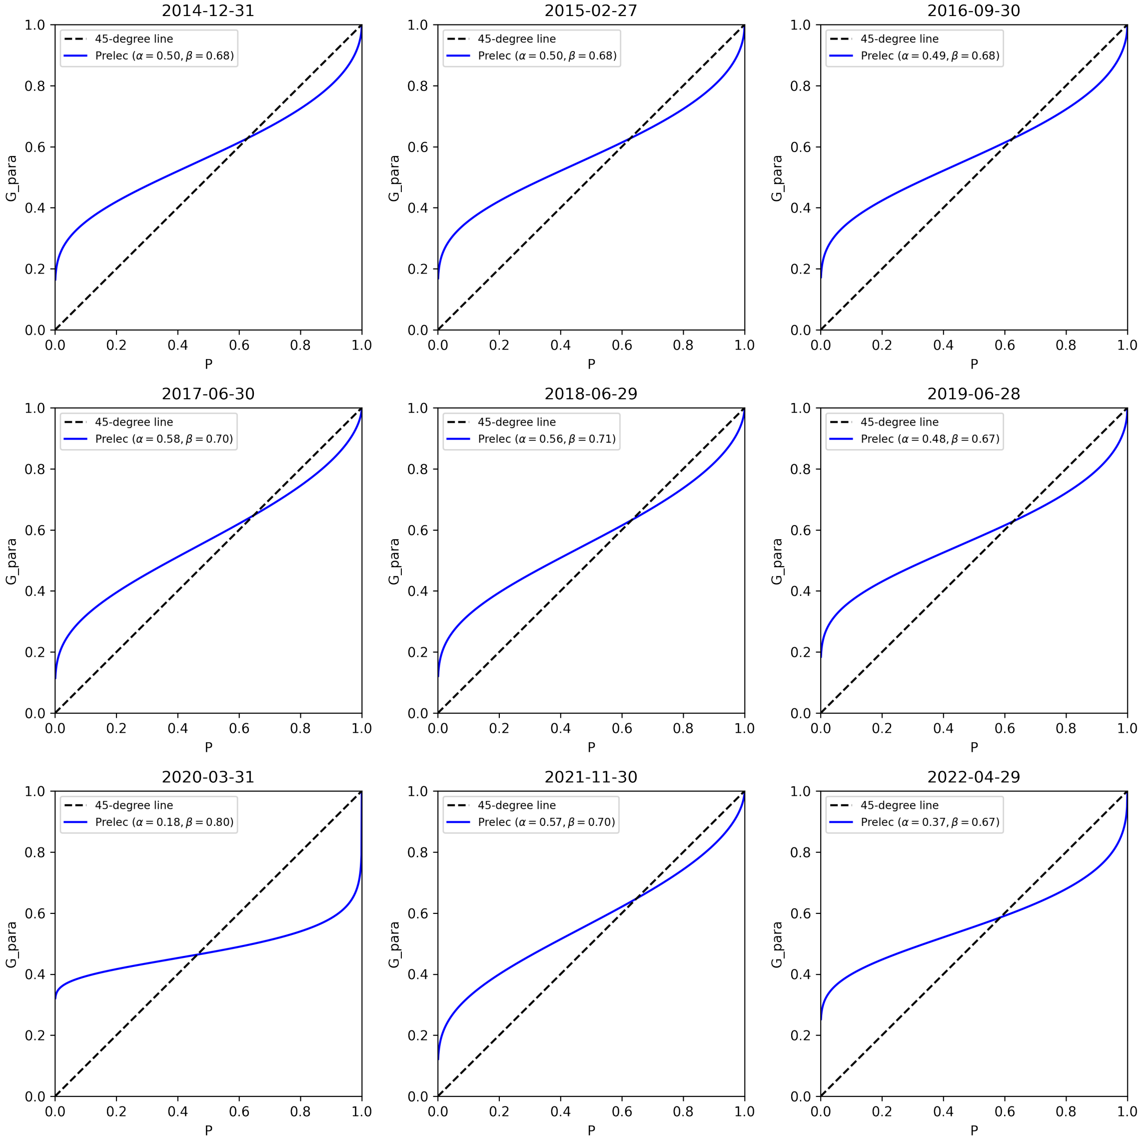

The horizontal axis represents the cumulative distribution function under physical measure and the vertical axis represents the probability weighting function.

Figure 4.2 illustrates the estimated shape of the probability weighting functions from nine representative months across different years within the sample period, using a constant relative risk aversion (CRRA) coefficient of $\gamma = 1$. The blue curves depict the estimated weighting functions, whereas the dashed diagonal represents a linear weighting of probabilities. The estimated curves typically deviate from the 45-degree line, typically forming an inverse-S shape observed in experiments in most months, suggesting investors systematically overweight tail probabilities and underweight moderate ones. However, shifts toward an S-shaped function in certain periods indicate that investor sentiment and market conditions can fluctuate, temporarily altering risk perception. These variations highlight the dynamic nature of probability weighting, reinforcing the role of behavioral biases in financial decision making.

Figure 4.3: Time-varying parameters of the five parametric weighting functions approximating the non-parametric estimates of the probability weighting functions obtained from end of month options and PC1 sentiment.

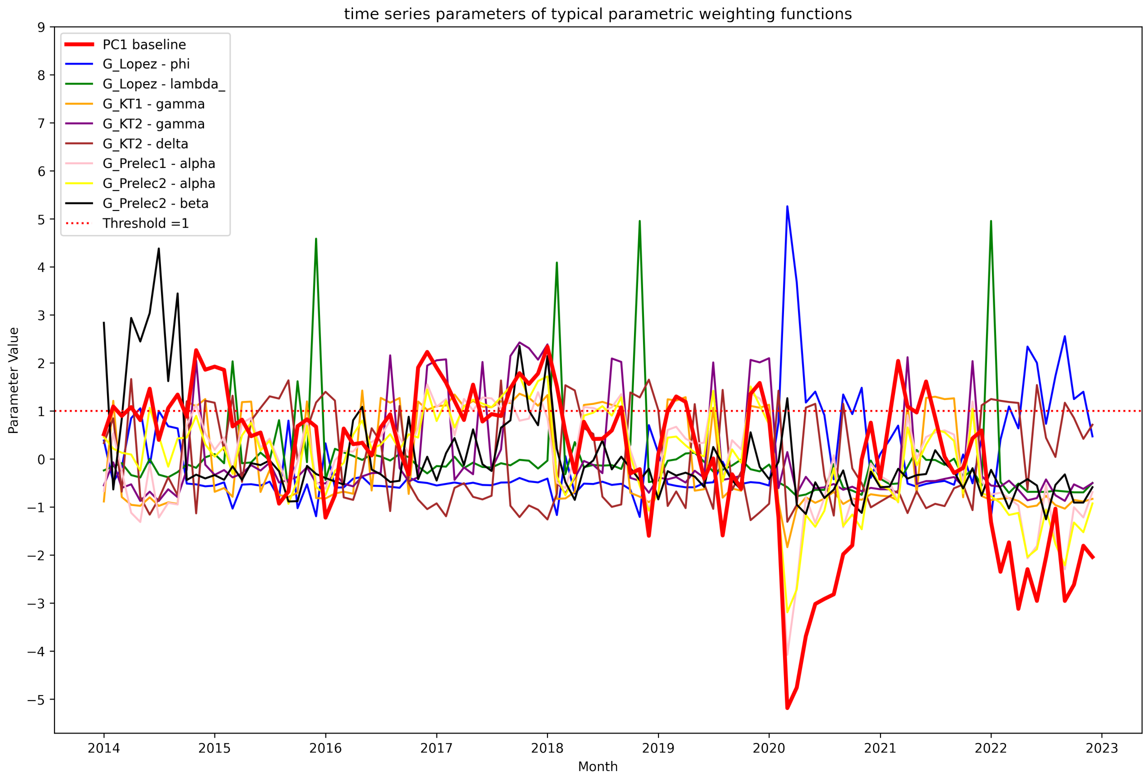

Figure 4.3 depicts the time-varying parameters of five parametric probability weighting functions alongside the benchmark sentiment proxy (PC1). The parameters are estimated by minimizing the distance between parametric function-implied values and the nonparametric estimates of the probability weighting functions derived from end-of-month options data. These parameters, including Lopes, two versions of the Tversky and Kahneman functions (KT1, KT2), and two versions of the Prelec functions (Prelec1, Prelec2), exhibit substantial variation over time, highlighting dynamic changes in investor probability weighting.

Table 4.3: Correlation Matrix: Time-varying weighting function parameters and PC1 investor sentiment

		G_Lopez		G_KT1	G_KT2		G_Prelec1	G_Prelec2						PCA				
		phi	lambda	gamma	gamma	delta	alpha	alpha	beta	PC1	PC2	PC3	BW	MCSI	CCI	SSW	AAII	VIX
G_Lopez	phi	1.00	-0.46	-0.52	-0.30	0.07	-0.84	-0.75	0.04	-0.70	-0.00	0.17	-0.15	-0.34	-0.10	-0.60	-0.46	0.68
	lambda	-0.46	1.00	-0.06	-0.05	0.14	0.07	0.06	-0.01	0.17	0.05	-0.05	0.02	0.05	0.05	0.25	0.08	-0.12
G_KT1	gamma	-0.52	-0.06	1.00	0.61	-0.40	0.81	0.77	0.03	0.58	0.04	-0.09	0.12	0.17	0.08	0.49	0.44	-0.57
G_KT2	gamma	-0.30	-0.05	0.61	1.00	-0.46	0.61	0.62	0.11	0.42	0.13	-0.02	0.05	0.07	0.17	0.36	0.38	-0.34
	delta	0.07	0.14	-0.40	-0.46	1.00	-0.32	-0.41	-0.27	-0.35	0.04	-0.31	-0.33	0.05	-0.10	-0.12	-0.47	0.24
G_Prelec1	alpha	-0.84	0.07	0.81	0.61	-0.32	1.00	0.92	-0.02	0.72	0.07	-0.12	0.16	0.27	0.16	0.56	0.52	-0.72
G_Prelec2	alpha	-0.75	0.06	0.77	0.62	-0.41	0.92	1.00	0.38	0.80	0.06	-0.10	0.19	0.26	0.15	0.64	0.59	-0.77
	beta	0.04	-0.01	0.03	0.11	-0.27	-0.02	0.38	1.00	0.32	-0.01	-0.00	0.10	0.06	0.01	0.29	0.28	-0.26
PCA	PC1	-0.70	0.17	0.58	0.42	-0.35	0.72	0.80	0.32	1.00	0.02	0.00	0.35	0.37	0.25	0.82	0.84	-0.78
	PC2	-0.00	0.05	0.04	0.13	0.04	0.07	0.06	-0.01	0.02	1.00	-0.02	-0.44	-0.48	0.79	0.18	-0.06	-0.12
	PC3	0.17	-0.05	-0.09	-0.02	-0.31	-0.12	-0.10	-0.00	0.00	-0.02	1.00	0.79	-0.54	0.25	-0.20	0.25	0.26
	BW	-0.15	0.02	0.12	0.05	-0.33	0.16	0.19	0.10	0.35	-0.44	0.79	1.00	-0.04	-0.01	0.07	0.40	-0.07
	MCSI	-0.34	0.05	0.17	0.07	0.05	0.27	0.26	0.06	0.37	-0.48	-0.54	-0.04	1.00	-0.15	0.24	0.20	-0.22
	CCI	-0.10	0.05	0.08	0.17	-0.10	0.16	0.15	0.01	0.25	0.79	0.25	-0.01	-0.15	1.00	0.19	0.21	-0.07
	SSW	-0.60	0.25	0.49	0.36	-0.12	0.56	0.64	0.29	0.82	0.18	-0.20	0.07	0.24	0.19	1.00	0.55	-0.58
	AAII	-0.46	0.08	0.44	0.38	-0.47	0.52	0.59	0.28	0.84	-0.06	0.25	0.40	0.20	0.21	0.55	1.00	-0.48
	VIX	0.68	-0.12	-0.57	-0.34	0.24	-0.72	-0.77	-0.26	-0.78	-0.12	0.26	-0.07	-0.22	-0.07	-0.58	-0.48	1.00

The correlation matrix in Table 4.3 offers insights into the relationships between the probability weighting parameters and the PC1 sentiment index. Among the analyzed parameters, the curvature parameter α of the two-parameter Prelec weighting function stands out, exhibiting the strongest positive correlation (80%) with PC1 sentiment. This significant association underscores the Prelec α parameter's sensitivity to shifts in investor sentiment, reflecting its prominent role in capturing the sentiment-driven variability in probability weighting. Other weighting function parameters also display meaningful relationships with sentiment: the curvature parameter (α) from the one-parameter Prelec weighting function closely follows, showing a 72% correlation with PC1. Furthermore, the gamma parameters from KT1 and KT2 weighting functions demonstrate moderate correlations of 58% and 42%, respectively. Conversely, the parameter ϕ of the Lopes weighting function exhibits a strong negative correlation of approximately -0.70 with PC1, indicating an inverse relationship with PC1. The strong negative correlation ($r = -0.75$) between PC1 investor sentiment and the Lopes power parameter ϕ indicates that when market sentiment becomes more optimistic, the curvature of the Lopes weighting function flattens, implying a reduction in probability distortion. In such periods, investors perceive probabilities in a way that is closer to objective or linear weighting, reflecting more rational or risk-neutral evaluations of uncertain outcomes. Conversely, during periods of low sentiment or pessimism, the value of ϕ tends to increase, steepening the curvature of the weighting function. This heightened curvature suggests that investors overweight extreme (tail) probabilities, consistent with fear-driven reactions to rare but impactful events. The lambda parameter from Lopes and the beta parameter from Prelec2 reveal relatively weaker correlations (17% and 32%), suggesting less pronounced, yet still relevant sentiment influences. Collectively, these findings confirm that investor sentiment substantially influences the shape and curvature of probability weighting functions, thus providing compelling evidence of sentiment-driven probability distortions in investor behavior under risk.

Table 4.4: Regression results: parameters of five weighting functions on PC1 and control variables

	G_Lopez		G_KT1	G_KT2		G_Prelec1	G_Prelec2	
	phi	lambda	gamma	gamma	delta	alpha	alpha	beta
const	0.49	0.45	0.70	0.69	-0.33	0.44	0.55	0.56
	(3.8)	(1.72)	(2.76)	(3.48)	(-1.51)	(7.32)	(7.39)	(7.1)
PC1	-0.05	0.03	0.11	0.06	-0.06	0.03	0.04	0.01
	(-6.86)	(1.93)	(6.9)	(5.43)	(-4.21)	(8.54)	(9.33)	(2.72)
E12	-0.00	0.00	0.00	0.00	0.00	-0.00	-0.00	-0.00
	(-0.07)	(0.86)	(0.95)	(0.04)	(1.92)	(-0.09)	(-0.76)	(-1.15)
B/M	0.51	-0.13	-1.44	-1.29	1.98	-0.45	-0.17	0.72
	(1.75)	(-0.22)	(-2.51)	(-2.9)	(4.01)	(-3.28)	(-1.02)	(4.01)
NTIS	3.43	-3.62	-6.56	-3.48	-0.78	-1.91	-0.81	2.69
	(3.3)	(-1.73)	(-3.23)	(-2.19)	(-0.44)	(-3.93)	(-1.36)	(4.21)
TBL	2.56	-3.77	-5.25	-1.73	0.62	-1.14	-0.04	2.59
	(1.68)	(-1.23)	(-1.77)	(-0.75)	(0.24)	(-1.61)	(-0.04)	(2.78)
INFL	0.84	-4.86	-7.25	-4.94	1.13	-1.52	-0.51	2.38
	(0.27)	(-0.79)	(-1.21)	(-1.06)	(0.22)	(-1.06)	(-0.29)	(1.26)
CORPR	0.15	-0.90	0.45	0.37	0.37	0.11	-0.03	-0.28
	(0.41)	(-1.2)	(0.61)	(0.65)	(0.58)	(0.61)	(-0.12)	(-1.23)
SVAR	8.96	0.58	0.32	4.90	-6.71	-2.62	-1.93	1.69
	(5.67)	(0.18)	(0.1)	(2.03)	(-2.51)	(-3.55)	(-2.13)	(1.74)
R^2	0.67	0.08	0.45	0.31	0.29	0.68	0.67	0.38
R^2 adj	0.65	0.01	0.41	0.26	0.23	0.65	0.64	0.33

Notes: Table displays coefficient estimates with t-statistics in parentheses

Table 4.4 reports the regression results for equation (4.20), examining the influence of investor sentiment (PC1) and control variables on the parameters of five probability weighting functions. The multivariate regressions exhibit substantial explanatory power, with the highest R^2 reaching 67% for the power parameter ϕ from G_Lopes. As expected, PC1 sentiment is generally positively and statistically significantly related to weighting function parameters, notably γ from G_KT1 (0.11), γ from G_KT2 (0.06), α from G_Prelec1 (0.03), and α from G_Prelec2 (0.04). This suggests that elevated sentiment exacerbates the overweighting of small probabilities as measured by these curvature parameters. However, the relationship is significantly negative for ϕ from G_Lopes (-0.05) and δ from G_KT2 (-0.06), indicating that higher sentiment may induce pessimism in certain aspects of probability weighting. In particular, the parameter δ from G_KT2, which governs the curvature of the probability weighting function in the loss domain, suggests that when sentiment rises, individuals exhibit a reduced tendency to overweight lower-probability losses.

Once examining the remaining control variables, the relationships appear less stable compared to the sentiment proxies. Notably, several variables display varying signs and inconsistent statistical significance across the different weighting function parameters. E12 is largely insignificant across all weighting functions, except for the delta parameter of the G_KT2 specification. B/M (book-to-market ratio) exhibits notable explanatory power, with significant negative coefficients for gamma parameters from G_KT1 (-1.44) and G_KT2 (-1.29), as well as the alpha parameter from G_Prelec1 (-0.45). In contrast, B/M has significant positive associations with delta from G_KT2 (1.98) and beta from G_Prelec2 (0.72), highlighting inconsistent directional effects across parameters. NTIS (net equity expansion) exhibits strong and statistically significant coefficients under most specifications, particularly in the

Lopez and KT1 weighting functions where the estimates are large (3.43 and 6.56, respectively, with t-statistics exceeding 3 in absolute value). This indicates that NTIS plays a notable role in explaining variation in the estimated parameters, suggesting that periods of net equity issuance are systematically associated with changes in the shape of the probability weighting functions. TBL (Treasury bill rate) exhibits partially significant and directionally mixed effects across the weighting functions. It attains positive significance for the β parameter of G_Prelec2 and a negative, marginally significant effect for the γ parameter of G_KT1. SVAR displays mixed signs and significance across the weighting functions, suggesting that the effect of market volatility on probability weighting is model-dependent. The other predictive factors—INFL, and CORPR exhibit statistical insignificance and inconsistent signs, highlighting their unstable relationships.

Overall, the findings indicate a weak and unstable link between fundamental control variables and the parameters of weighthing function. In contrast, the strong and stable relationship between sentiment factors and the tail overweighting parameters in multivariate regressions suggests a robust connection between sentiment and the overweighting of small probabilities.

4.7 Implication

Building on the studies from Polkovnichenko et al. (2019) and Azimi et al. (2024), the findings suggest that probability weighting parameters capture behaviorally grounded biases toward tail events that are not fully reflected in existing sentiment indicators. These parameters have direct implications for optimal portfolio allocation and dynamic trading strategies. Therefore, it becomes essential to assess the multidimensional nature of investor sentiment. Within the behavioral pricing kernel framework introduced in Chapter 2, Barone-Adesi et al. (2012) identify finer psychological dimensions of sentiment by quantifying deviations between the representative investor's subjective return distribution (Prep) and the objective distribution (Pobj). Specifically, sentiment is decomposed into its first- and second-moment components, corresponding to excessive optimism (mean distortion) and overconfidence (variance distortion). Furthermore, Barone-Adesi et al. (2012) argue that overconfidence constitutes a critical dimension of sentiment that shapes the pricing kernel but is not fully captured by the Baker–Wurgler series. This view aligns with the observation that, although Baker and Wurgler discuss several psychological heuristics and biases underlying sentiment—such as excessive optimism, overconfidence, and representativeness—they place primary emphasis on excessive optimism in their analysis (Barone-Adesi et al., 2012). Such deviations from standard utility models highlight a promising research avenue for identifying the underlying factors that drive investor preferences and influence asset prices.

4.8 Conclusion

This chapter empirically examines the relationship between investor sentiment and the overweight of tail events, proxied by the parameters of parametric probability weighting functions. The findings indicate that among the examined functions, the curvature parameter α of the Prelec weighting function exhibits the strongest correlation with PC1, reaching 68%. This provides compelling empirical evidence of a link between probability weighting parameters and external sentiment measure in real-world, non-experimental settings. Moreover, the multivariate regression analysis reveals a consistently strong and stable relationship between sentiment factors-particularly the first principal component (PC1)-and the parameters characterizing the curvature of probability weighting functions.

Overall, these results confirm a strong connection between probability weighting and investor sentiment. Utility functions incorporating probability weighting have been extensively applied to key areas in finance, including risk premia, the cross-section of expected returns, and portfolio diversification. The findings offer robust non-experimental support for the behavioral assumptions underlying this strand of research on investor decision-making. By shedding light on the dynamic nature of probability weighting, the study contributes to the broader literature on behavioral asset pricing and provides new insights into the resolution of the pricing kernel puzzle.

Chapter 5

Conclusion

5.1 Introduction

This concluding chapter summarises the roles of investor sentiment estimated in Chapters 2, 3, and 4, and discusses professional implications arising from these findings. It also evaluates the robustness of the results, emphasizing the significance of alternative model specifications in light of the persistence of irrational beliefs across stocks and over time. Recognizing that mainstream economics continues to rely heavily on expected utility theory and the rational agent paradigm, this thesis addresses several key critiques and proposed extensions to these foundational frameworks. Finally, the chapter concludes with implications for future research and situates the thesis's contributions within the broader context of behavioral finance literature.

The theoretical framework first reviews the full-information rational-expectations paradigm, noting that investors' beliefs are often shaped by adaptive, extrapolative, and sentiment-driven processes that depart from objective probability models. It then introduces the psychology of tail events, showing how probability-weighting functions cause individuals to systematically overweight extreme gains and losses, thereby distorting risk perceptions in ways that standard expected-utility theory cannot capture. Finally, it reviews the research evolution of investor sentiment on stock returns from early skepticism about noise traders to current widely recognized sentiment-driven stock returns. From the DSSW model's pioneering work that measures investor sentiment using closed-end fund discounts in the 1990s to today's sophisticated sentiment analysis using machine learning, sentiment studies continue to enhance the understanding of market anomalies. This evolution underscores the importance of integrating psychology, data analytics, and traditional financial theories to build a more comprehensive picture of markets that are not only efficient but also influenced by behavioral bias. Noise-trader theory and limits-to-arbitrage models explain why these behavioral distortions persist: risk-averse arbitrageurs confront "noise-trader risk" when betting against sentiment-driven mispricings, so deviations from fundamental values remain priced in equilibrium. Together, these insights establish a unified basis for incorporating heterogeneous beliefs and behavioral biases into modern asset-pricing models.

The literature section first traces the evolution of asset pricing theory from traditional asset pricing models to behavioral asset pricing models, highlighting how the empirical failures of consumption-based models spurred a search for more flexible SDF specifications. It then surveys the rise of quantile-preference models, and the corresponding econometric advances in quantile regression. Finally, it examines the behavioral asset pricing literature on news- and survey-based sentiment proxies and non-linear probability-weighting functions from the rank-dependent family of models. Together, these threads—alternative preference structures, richer factor models, and sentiment-driven behavioral biases—provide the theoretical and empirical foundations for the subsequent chapters' investigation of how dynamic tail beliefs and sentiment measures drive asset-pricing anomalies.

5.2 Behavioral pricing kernel and news sentiment

Chapter 2 conducts an empirical test of the conusmption-based CAPM by approximating the pricing kernel as a linear function of consumption growth and news sentiment. This approach seeks to reconcile the apparent inconsistency between the strong cross-sectional explanatory power of conditional asset-pricing models and the underlying assumptions by developing a behavioral pricing kernel that incorporates news sentiment. Building on Shefrin (2008)'s decomposition of the SDF into rational and behavioral components, the model integrates sentiment—proxied by a lagged news-based indicator—alongside consumption growth and uses the consumption–wealth ratio (cay) as a conditioning variable to allow time variation in preferences. The time series of realized values of the empirical SDF function and the corresponding risk aversion parameters are estimated using alternative specifications. In traditional RE asset pricing models, risk aversion is often assumed to be constant. However, empirical evidence suggests that risk aversion changes over time, influenced by economic cycles, wealth levels, and recent returns. This limitation can be addressed by estimating time-varying risk aversion the time-varying risk aversion, where agents become more or less risk-averse based on market conditions. The SDF is approximated using both a second-order Taylor series expansion and a flexible Chebyshev polynomial basis, allowing the model to capture nonlinearities and more nuanced dynamics in asset pricing. The proposed approach retains many of the desirable features found in nonparametric analyses while mitigating several of their key limitations. Specifically, it approximates the unknown marginal utility function within a static framework using a combination of Taylor series expansion and Chebyshev polynomials, thereby offering both local accuracy and global stability in functional approximation. In addition, departing from the common assumption of constant parameters over time, the analysis allows for time variation in the parameters by introducing a scaling variable,consumption–wealth ratio. This approach captures potential shifts in the pricing kernel driven by changes in the broader economic environment. Notably, the Chebyshev-based specification offers greater flexibility than the quadratic form derived from the Taylor expansion, enabling the pricing kernel to exhibit more pronounced oscillations driven by sentiment and better accommodate complex state-dependent behavior.

The Generalized Method of Moments (GMM) is used to estimate the parameters of various asset pricing models, leveraging its flexibility in handling endogeneity, heteroskedasticity, autocorrelation, and nonlinearity in moment conditions. A two-step GMM procedure is applied. The first stage uses an identity weighting matrix to generate consistent but possibly inefficient estimates. The second stage improves efficiency by applying the optimal weighting matrix, based on the long-run covariance matrix of the moment conditions. This yields robust standard errors and allows for hypothesis testing via the Hansen J-test for overidentifying restrictions.

The findings indicate that the sentiment-driven SDF provides a superior explanation of the cross-section of size and book-to-market sorted equity returns—both economically and statistically—compared to alternative models such as the Capital Asset Pricing Model (CAPM), and the standard Consumption CAPM (CCAPM), and performs comparably to the Fama-French three-factor model. The integration of sentiment into the behavioral pricing kernel significantly reduced the negative values that the SDF take at some time points. Overall, the findings underscore the critical role of sentiment in shaping asset prices and highlight the value of embedding behavioral elements structurally within the SDF framework.

One limitation of this chapter aligns with a critique raised by Shefrin (2008). In Behavioral asset pricing models, sentiment is commonly represented as a scalar variable, such as a bias in the mean of a given distribution, which is typically interpreted as a form of error—whether stemming from individual investors or emerging at the aggregate market level. While this approach may be suitable for simplified, ad hoc models, it is generally overly simplistic. In reality, sentiment behaves more like a stochastic process, evolving according to a distribution that interacts with underlying fundamental variables. In markets characterized by heterogeneous beliefs, sentiment is rarely uniform—some assets may reflect excessive optimism, while others exhibit excessive pessimism. This heterogeneity gives rise to the concept of an oscillating SDF, which captures the nonuniform nature of sentiment across assets. This chapter estimated the time series of the empirical SDF, which is inherently estimating the projection of the SDF onto a subspace, at each observation date. In a neoclassical world, with a stable SDF, doing so complements a cross-sectional approach. Over time, the SDF can be stitched together, or traced out as a function. In a behavioral world, sentiment varies so much that it's unclear exactly what can stitch together, because the SDF is so volatile. Consequently, within behavioral SDF theory, sentiment is more appropriately modeled as a change of measure—a function rather than a simple scalar variable. This theoretical perspective introduces complications when using news sentiment as a proxy, which are often limited in capturing the full functional nature implied by the theory. For estimating the time series of realized SDF values, the sentiment that gets measured, theoretically, should be the probability misestimation, in relative terms, associated with the underlying realization of the state variable. These are issues that are in need of exploration in future research.

Moreover, it is also appealing to compare with the alternative models. One obvious possibility is through the R2 of a cross-sectional regression but there are other alternatives such as the test proposed by Barillas and Shanken (2017). Another promising direction is to treat sentiment as a state variable, alongside the consumption–wealth ratio (cay), to assess its incremental predictive power beyond that of cay in a behavioral setting. While this approach is conceptually related to the framework presented in this chapter, it differs in that sentiment would function as a predictor of returns rather than as a pricing factor within the SDF. It is necessary to think more about the conceptual differences of using sentiment as a state variable rather than as a pricing factor. It might be the same analytically.

Another promising avenue for extension involves incorporating sentiment—following the approach taken in this paper to augment the (C)CAPM—into long-run risk models in the spirit of Bansal and Yaron (2004). While much of the empirical evaluation of these models has centered on characteristics of the conditional distribution of market portfolio returns, such as the risk premium and the volatility and persistence of excess returns, Hansen et al. (2008) suggest that long-run risk also plays a critical role in explaining average return differentials between value and growth portfolios.

5.3 Quantile Preference and Sentiment

Chapter 3 introduces a quantile preferences asset pricing model that incorporates a sentiment index constructed via Principal Component Analysis (PCA), aiming to investigate how sentiment influences the cross-section of extreme tail events in the return distribution. This approach seeks to capture and explain the asymmetric nature of investor preferences observed across different segments of the return distribution. By targeting specific return quantiles, this quantile-utility approach leverages the full distribution without imposing a parametric utility function, allowing an investor's chosen quantile to serve as a proxy for risk aversion. This contrasts with expected utility models, where risk attitude is determined by the utility function.

This chapter constructs a composite sentiment index using Principal Component Analysis (PCA), leveraging five widely recognized sentiment measures from academic literature: the Shapiro, Sudhof, and Wilson (SSW) Sentiment Index; the AAII Investor Sentiment Survey, as a proxy for individual investor sentiment; the Baker and Wurgler (BW) Sentiment Index; the Michigan Consumer Sentiment Index (MCSI); and Shiller's Crash Confidence Index (CCI). The appeal of this composite sentiment index lies in its enhanced transparency compared to traditional approaches. In contemporary literature literature, researchers have typically relied on three primary methodologies to quantify investor sentiment. The first involves market-based proxies, including metrics such as trading volume, closed-end fund discounts, IPO first-day returns, IPO issuance volume, option-implied volatility (e.g., the VIX), and mutual fund flows. These indicators infer sentiment indirectly through observed market behavior. The second approach employs sentiment surveys, which provide direct insight into investor expectations and psychological states. Common examples include Michigan Consumer Sentiment Index (MCSI), Shiller's Crash Confidence Index, AAII Investor Sentiment Survey, Conference Board's Consumer Confidence Index (CCI). More recently, a third stream of research has emerged, applying machine learning techniques to textual data—such as financial news, earnings calls, and online commentary—to derive sentiment measures that are both dynamic and highly granular. Compared to these disparate and sometimes opaque approaches, the PCA-based composite index offers a more systematic and interpretable framework for capturing the multifaceted nature of investor sentiment.

In the proposed quantile asset pricing model, in addition to the sentiment factors PC1 and PC2, a set of control variables—including idiosyncratic volatility (Garman & Klass, 1980), changes in the VIX, the risk-free rate (RF), and the Carhart Four Factors—is incorporated to ensure that the effects of sentiment are not confounded with well-known sources of well-known risk premia. These controls are essential for accurately modeling the cross-section of future return quantiles. Notably, the theoretical foundations of the momentum factor in the Carhart Four Factors span both behavioral and rational perspectives. From a behavioral standpoint, momentum is often linked to investor sentiment and cognitive biases—such as overreaction, representativeness, and herding—which can cause asset prices to deviate from fundamentals in a systematic and predictable manner. These behavioral distortions may lead to under- or overreactions that fuel momentum over intermediate horizons. Conversely, a body of research within rational finance shows that momentum can also emerge in markets characterized by information asymmetries, transaction costs, or delayed price adjustments. Even in such environments, evolving investor expectations—often shaped by sentiment—can sustain return persistence. Therefore, momentum and sentiment are conceptually and empirically intertwined. Including momentum as a control variable is thus critical to disentangle its effects from those of sentiment, particularly in the tails of the return distribution where sentiment-induced mispricing is most likely to occur.

The Panel Quantile Regression (PQR) model, as introduced by Koenker (2004), is employed to estimate the parameters defined in the proposed linear asset pricing equation. The advantage of this approach lies in its ability to effectively control for unobserved heterogeneity across financial assets while simultaneously identifying common sentiment-driven patterns, thereby yielding more precise, quantile-specific estimates. Consequently, these refined estimates are expected to enhance forecasting performance directly. Moreover, this approach can be effectively applied to obtain precise estimates of Value-at-Risk (VaR), a widely used risk measure in the financial industry. In this context, the panel data framework retains all the advantageous properties of standard time series analysis, while its cross-sectional dimension allows for the identification and control of common shocks affecting multiple assets simultaneously.

To evaluate the extent to which unobserved heterogeneity influences the relationship between sentiment and returns, the individual Univariate Quantile Regression (UQR)—which does not control for such heterogeneity—is first employed as a benchmark. Comparing the parameter estimates from the Panel Quantile Regression (PQR) with those from the individual UQR models reveals a notable pattern: once unobserved heterogeneity is accounted for through the PQR framework, the first principal component (PC1) exhibits a more pronounced negative effect on the upper quantiles of the return distribution than most UQR estimates suggest. This finding indicates that controlling for heterogeneity amplifies the observed downward influence of negative sentiment, particularly in the higher return quantiles, thereby reinforcing the asymmetric effect of sentiment documented. Crucially, it suggests that failing to account for unobserved heterogeneity can lead to systematic underestimation of sentiment impact, especially at the distributional tails. Overall, this comparison underscores the critical importance of incorporating unobserved factors in panel data models to produce more accurate and unbiased estimates.

This chapter advances the understanding of how investor sentiment influences the cross-section of extreme tail events in stock return distributions by employing a Panel Quantile Regression Model. This model effectively accounts for unobserved heterogeneity across financial assets and facilitates the identification of shared patterns in the tail behavior of returns across different stocks. A key innovation of this chapter is the development of a novel investor sentiment index, constructed using Principal Component Analysis (PCA). This composite index integrates several widely recognized sentiment proxies, with particular emphasis on machine learning-based news sentiment, offering a more robust and flexible tool compared to single-indicator approaches. By reducing the dimensionality of multiple sentiment measures, PCA isolates the underlying common factors that drive investor sentiment, thereby enhancing the model's capacity to link sentiment dynamics with tail risk. This methodology not only improves upon traditional models that rely on assumptions of homogeneous preferences and expected utility maximization but also establishes a more comprehensive framework for analyzing the impact of sentiment on extreme asset return events.

The empirical results presented in this Chapter provide compelling evidence of the asymmetric and quantile-dependent effects of investor sentiment and volatility on stock returns, thereby reinforcing the importance of modeling cross-sectional heterogeneity and tail risk. Overall, the findings reveal a nuanced, quantile-specific view of stock returns, illustrating the heterogeneous influence of systematic risk factors, market sentiment components, and idiosyncratic volatility across the distribution. The empirical findings reveal significant heterogeneity in how investor sentiment impact stock performance at different quantiles, underscoring the nuanced relationship between sentiment and return distributions. The research highlights the critical role of unobserved heterogeneity among financial assets, particularly across different quantiles of the return distribution. Using Panel Quantile Regression (PQR), the first principal component (PC1) of sentiment is found to exert a consistently negative and statistically significant effect across all quantiles. While this impact remains relatively stable in the lower and middle quantiles, it intensifies markedly at the upper tail—where the coefficient nearly doubles—indicating that elevated investor sentiment disproportionately suppresses the probability of extreme positive returns. This asymmetry aligns with the findings of Baker and Wurgler (2007), who show that high sentiment is associated with lower subsequent returns.

The values of the PC1 parameter estimates are not symmetric around the median, which highlights the increasing importance of sentiment effects on the estimation of the upper quantiles of returns. This asymmetry suggests that negative sentiment, as captured by PC1, exerts a more pronounced downward influence on higher quantiles, indicating that sentiment has a stronger impact on the upper range of the return distribution compared to lower quantiles. The results highlight the asymmetric impact of stock-specific characteristics and sentiment across the conditional return distribution, providing a deeper understanding on how sentiment indices drive performance at different levels of the return distributions, offering insights beyond traditional mean based regression models. This nuanced relationship between sentiment and returns is consistent with prospect theory of Kahneman and Tversky (1979a), which posits asymmetric perceptions of gains and losses, further emphasizing why sentiment effects concentrate in the tails. Meanwhile, the second principal component (PC2) has a smaller, positive, and more stable effect, likely capturing orthogonal dimensions of sentiment that influence returns through different channels.

In addition, a major strength of the proposed methodology lies in its ability to account for unobserved heterogeneity among financial assets, thereby enabling a clearer separation of overall market risk into its systemic and idiosyncratic components. By-comparing PQR estimates to those from individual Univariate Quantile Regressions (UQR). The PQR estimates, particularly at the highest quantile, reveal a significantly more pronounced negative effect of PC1 than the median UQR estimates, underscoring how controlling for unobserved heterogeneity unveils stronger sentiment effects on high-return stocks. In contrast, UQR estimates display less variation across quantiles, potentially masking the true distributional impact of sentiment. The findings underscores the importance to control for unobserved heterogeneity among the assets.

Beyond investor sentiment, volatility, measuring the magnitude of fluctuations in asset prices, is another significant factor shaping the tail events. Higher volatility tends to increase the frequency and intensity of these tail events, underscoring the need to understand how such risk factors affect not only the central tendency of returns but also their spread, particularly under conditions of financial stress or exuberance. By focusing on quantiles, the model can specifically capture how these risk factors—sentiment, volatility, and others affect the tails of the return distribution, revealing the degree to which both negative and positive extremes may occur. By integrating volatility factors alongside sentiment within a quantile regression model, the model provides a comprehensive perspective on how systematic risks, captured through volatility and sentiment, shape the tails of the distribution. This integration of volatility and sentiment essentially streamlines the analysis by addressing both fundamental and behavioral drivers of risk, leading to better insights into the full distribution of returns.

For future research, it is valuable to extend the static quantile utility framework into a dynamic setting by developing an intertemporal consumption model, which would enable the separation of the elasticity of intertemporal substitution (EIS) from risk preferences within the quantile framework and allow for the estimation of the implied discount factor and EIS across different quantiles. It involves an equilibrium model given by an Euler equation for the quantile preference approach, which is more appealing from a theoretical (macro) perspective but may be more difficult to implement in an asset pricing framework. This might be methodologically challenging and add value to the chapter. In the paper by de Castro and Galvao (2019), the corresponding recursive equation is defined by the sum of the current period's utility function and the discounted value of the certainty equivalent, which defines a utility function based on quantiles.

5.4 Time-varying Probability Weighting and Investor Sentiment Toward Tails

Chapter 4 investigates the relationship between tail-overweighting parameters embedded in the probability weighting functions of rank-dependent family of models, and external sentiment measures (PC1). To evaluate the extent to which time-varying parameters of probability weighting functions capture investor sentiment toward tail events, this chapter examines five families of weighting functions: the two-parameter weighted average power function introduced by Lopes (1987), the one- and twoparameter functions proposed by Tversky and Kahneman (1992), and the one- and two-parameter specifications from Prelec (1998). The effectiveness of these models is validated by estimating the parameters on a monthly basis and comparing them to a benchmark proxy for investor sentimentnamely, the first principal component (PC1) extracted through the PCA methodology detailed in Chapter 3. In addition, time-series regressions are employed to analyze the explanatory power of investor sentiment, alongside a set of control variables, in accounting for temporal variations in the overweighting of tail probabilities, proxied by the parameters of five parametric probability weighting functions. The PC1, the benchmark sentiment proxy derived as the first principal component, is used as the explanatory variables. A set of variables identified by Welch and Goyal (2008) as potential predictors of equity market returns is included as control variables. These include: E12 (the twelve-month moving sum of SP 500 earnings), B/M (the book-to-market ratio), NTIS (net equity issuance), TBL (the risk-free rate), INFL (the annual inflation rate), CORPR (the corporate bond spread), SVAR (stock return variance), and CSP (the cross-sectional premium). By integrating these controls, the model isolates the incremental effect of investor sentiment on the overweighting of small-probability events, thereby shedding light on probability distortions in investor behavior.

This chapter estimates the pricing kernel using index option data and use it to derive the corresponding implied probability weighting function. A non-parametric method is employed to extract the risk-neutral density from observed option prices. In doing so, moving beyond the calibration of parametric utility functions, this chapter explores direct estimation methods based on the no-arbitrage principle for the pricing kernel by integrating both the risk-neutral distribution (implied by asset returns) and the physical distribution (driving asset returns). The risk-neutral density is estimated using a non-parametric method, while the physical density is obtained through a semi-parametric approach, as previously outlined. The pricing kernel is computed by taking the ratio of the risk-neutral density to the physical density and subsequently dividing this ratio by the risk-free rate. A one-month horizon is selected for the benchmark analysis, as it is the most widely studied timeframe in empirical pricing kernel research and corresponds to a maturity with highly liquid option contracts. The robustness test demonstrates that the findings are consistent across other commonly analyzed horizons.

This chapter underscores the significance of probability weighting in explaining investor behavior toward tail events. The pricing kernel is adjusted by incorporating a probability weighting function that distorts cumulative probabilities. This distortion reflects a key behavioral insight: individuals tend to perceive probabilities non-linearly—overweighting small probabilities and underweighting large ones. The primary contribution of this chapter lies in addressing a gap in the literature by empirically examining the relationship between time-varying overweighting preference parameters of probability weighting functions and investor sentiment toward tail events. Drawing on insights from non-expected utility models—particularly prospect theory and cumulative prospect theory—it explores how subjective probability weighting, especially the overweighting of tail probabilities, correlates with sentiment related to extreme market outcomes. To benchmark investor sentiment, the first principal component (PC1) extracted through principal component analysis (PCA) is employed, facilitating a comparison with the time-varying parameters of five widely used probability weighting functions. By aggregating multiple sentiment indicators, PCA yields a more comprehensive and robust representation of overall investor sentiment, offering a nuanced alternative to single-indicator approaches and enhancing the analysis of sentiment's role in behavioral asset pricing. This chapter is complementary to the literature that introduces probability weighting function to compute weighted average utility, and show similar characteristics with the probability weighting functions derived from options.

The results indicate that, among the examined functions, the curvature parameter α of the Prelec weighting function exhibits the highest correlation with PC1, at

68%. These findings offer empirical support for the link between probability weighting parameters and news sentiment in real-world, nonexperimental settings. This contributes to the broader behavioral asset pricing literature by providing new insights into the dynamic nature of probability weighting and its relevance to the pricing kernel puzzle. The multivariate regression analysis reveals a consistently strong and stable relationship between sentiment factors-particularly the first principal component (PC1)-and the parameters characterizing the curvature of probability weighting functions. Specifically, higher sentiment significantly increases the overweighting of tail events, underscoring sentiment's key role in investor risk perception and pricing behavior. Compared to the instability and inconsistent significance of economic fundamentals across regressions, sentiment proxies demonstrate robust explanatory power and directional consistency, reinforcing sentiment as a key driver of investors' probability distortions. The findings highlight that investor sentiment provides a reliable and significant explanation of the overweighting of tail risks, offering valuable insights into behavioral aspects of financial decision-making under uncertainty. Importantly, the study extends these findings to real-world financial markets, validating the significance of sentiment-driven probability weighting beyond controlled laboratory settings.

One limitation of this chapter is the estimation of RND, S. Figlewski (2010) proposed an approach that allows extracting the body and tails of the distribution of returns separately.

Given the strong correlation between the time-varying parameters of probability weighting functions and investor sentiment toward tail events, a particularly promising avenue for future research is to address the pricing kernel puzzle by calibrating probability weighting functions using external sentiment measures. According to Dierkes et al. (2024), probability weighting is not only closely linked to volatile market environments but also plays a central role in explaining the pricing kernel puzzle. Their analysis reveals that, once probability weighting is incorporated, the pricing kernel becomes monotonically decreasing in wealth, and the corresponding risk aversion functions remain significantly positive. In contrast, the raw pricing kernel exhibits a pronounced U-shape, indicating episodes of negative implied risk aversion. Advancing this line of inquiry will necessitate a more nuanced exploration of the structure and dynamics underlying the mapping between sentiment indicators and probability weights. Moreover, the findings suggest that probability weighting is time-varying and, in certain periods, may lead to the underweighting of tail events. This observation opens another interesting avenue for future research: exploring the relationship between this time variation and both expected stock returns and market volatility.

There is considerable room for exploring the solutions for the pricing kernel puzzle, beginning with simple models that rely on a single state variable, and gradually progressing to more complex frameworks. Similar to the testing process, many of the proposed solutions are stand-alone models with limited empirical support. Most of these solutions focus on calibration, using stylized facts to illustrate the pricing kernel puzzle. However, there is still considerable work to be done in evaluating the alternative models and assessing their compatibility with the data. Ideally, some of these solutions could be combined into a nested model, enabling a thorough test of their individual components. For the future of the pricing kernel puzzle, one might look at the bivariate estimation of risk-neutral and physical distributions presented in Jackwerth and Vilkov (2017). Typically, bivariate risk-neutral distributions can only be derived using options written on both assets at the same time. However, Jackwerth and Vilkov (2017) were able to accomplish this by analyzing index returns and volatility, using longer-dated options on the SP 500. By dividing the two distributions, they were able to extract a bivariate pricing kernel for the first time.

5.5 Implications for professional practice

The insights derived from this thesis carry important implications for financial professionals, including portfolio managers, risk analysts, investment strategy design, and policymakers. By embedding investor sentiment into asset pricing frameworks, this research equips practitioners with enhanced tools for understanding and managing market anomalies, particularly those driven by behavioral distortions.

First, the integration of sentiment into the SDF framework in chapter 2 offers a more accurate and forward-looking measure of risk. For portfolio managers, this enhanced pricing kernel offers a refined approach to estimating expected returns under varying sentiment conditions. The model allows for more accurate risk assessment by accounting for sentiment-induced mispricing that is not captured by traditional models. This can lead to better timing in security selection and improved hedging strategies during sentiment-driven market dislocations. Moreover, financial analysts can benefit from the model's ability to decompose pricing dynamics into fundamental and sentiment-driven components. This decomposition enables more nuanced narrative-based valuation approaches, particularly useful for firms sensitive to public perception, regulatory developments, or macroeconomic shocks.

Second, the application of quantile-based methods in Chapter 3 highlights how sentiment influences different parts of the return distribution, especially the tails. This provides risk managers with a powerful framework for stress testing and scenario analysis, as it allows for a more granular understanding of how assets behave under extreme market conditions. Additionally, the use of Panel Quantile Regression methods enhances the ability to control for firm-level heterogeneity, offering a refined approach to cross-sectional stock selection. portfolio managers can implement quantile-based optimization strategies that overweight or hedge exposures depending on prevailing sentiment regimes.

Third, the findings in Chapter 4—linking probability weighting functions to sentiment—have practical relevance for the design of financial products and pricing models. Derivatives pricing, particularly for out-of-the-money (OTM) options, can be improved by accounting for the sentiment-driven overweighting of tail probabilities. Market participants involved in derivatives trading, structured product design, or volatility arbitrage can leverage these insights to anticipate and exploit pricing biases.

Lastly, the behavioral foundations of this research are increasingly applicable in financial technology and AI. As asset managers and platforms integrate machine learning to model investor behavior and preferences, the tractable, sentiment-aware models proposed here can enhance algorithmic trading strategies, robo-advisory services, and client portfolio customization—ensuring alignment with how investors actually perceive and respond to risk, rather than how they are assumed to in classical models.

In sum, this thesis bridges theoretical innovation with actionable insights, providing finance professionals with empirically grounded, sentiment-aware models to improve decision-making in increasingly complex and behaviorally influenced markets.

5.6 Implications for future research

The three research chapters in this thesis are closely interrelated, giving rise to numerous compelling questions for future exploration. Chapter 2 demonstrates that incorporating sentiment into the (C)CAPM—thereby capturing systematic investor error—significantly enhances the model's ability to explain cross-sectional variation in average returns. This finding suggests that, in a context where heterogeneous beliefs and preferences are intertwined, observed anomalies may result either from cognitive biases in investors' assessments of the joint distribution of the stock returns and aggregate consumption, or from aggregate preferences that deviate from standard risk aversion assumptions. Distinguishing between these two possibilities remains an open and important avenue for further research. The impact of beliefs—such as expected returns and perceived risk—differs from the time-varying preferences typically defined over aggregate consumption in traditional asset pricing models. Specifically, beliefs can be market- or sector-specific; for example, an investor might be optimistic about the housing market but pessimistic about the stock market. Therefore, beliefs data becomes crucial for identifying time-varying investment opportunities, particularly with regard to composition risk.

Chapter 3 contributes to our understanding of how investor sentiment influences the cross-section of extreme tail events in the distribution of stock returns, employing a Panel Quantile Regression Model that leverages the longitudinal structure of the data. A key advantage of the quantile-based preference framework used in this chapter is its ability to separate tastes from beliefs—a desirable feature in which the belief component of utility remains invariant to the specification of tastes. As a result, the framework is robust to any monotonic transformation, not just affine ones, enhancing its theoretical and empirical flexibility. Economists typically assume that a decision maker's tastes are stable, while beliefs evolve with new information. In this context, enhancing and extending Shefrin's behavioral SDF framework to the quantile utility model presents a particularly promising avenue for capturing the role of belief distortions in asset pricing.

Moreover, it is appealing to explore the root causes of investors' tendency to overweight small probabilities in chapter 4. Although much of the existing literature, for simplicity of exposition, attributes this behavior to preferences (behavioral biases), it remains unclear whether the overweighting of tail events is driven solely by such preferences or instead reflects biased beliefs, that is, distorted investor expectations. Barberis (2013b) provides a clear and insightful discussion of how these two mechanisms—preferences and beliefs—are fundamentally distinct, and how each, either independently or jointly, could help explain the persistent overpricing of out-of-the-money (OTM) options. Félix et al. (2020) demonstrates that investor sentiment is linked to time-varying preferences for lottery-like assets, even after controlling time-varying for investor optimism, suggesting that the overweighting of tail events during periods of high sentiment is more indicative of behavioral biases than of changes in beliefs alone.

Furthermore, one limitation of the quantile preference model, as noted by Giovannetti (2013), is that a purely quantile-maximizing agent likely does not accurately represent typical risk preferences. Therefore, the model developed in this study should be interpreted as a stylized and parsimonious framework within the class of models that incorporate asymmetric preferences toward gains and losses—such as those found in prospect theory and disappointment aversion. Extending the quantile utility maximization approach to practical applications—such as asset pricing, consumption behavior, and scenario-based policy analysis—offers a promising path for advancing its theoretical and empirical relevance. One particularly interesting direction for future research involves exploring the linkage between quantile utility model and the rank-dependent family of models. One approach to integrating quantile utility with RDEU is to define a utility function based on quantiles and then apply a rank-dependent weighting function to these quantiles.

Additionally, in chapter 2, sentiment within Shefrin's behavioral pricing kernel framework is represented by the ratio of an investor's subjective probability assigned to a given state relative to the objective probability of that state occurring. This formulation closely resembles the probability weighting function embedded in the SDF integrated with the rank-dependent family of models in Chapter 4, where the pricing kernel is derived from the ratio of risk-neutral and physical probability distributions, in accordance with the no-arbitrage principle. Therefore, exploring the relationship between these frameworks presents a promising avenue for future research.

Meanwhile, sentiment-based approaches have focused on developing empirical proxies for sentiment and the limits to arbitrage that allow sentiment to affect prices. The advances of machine learning with increasingly available datasets has provided macroeconomists and financial economists with a transformative machine learning-driven approach to directly measure sentiment and beliefs data from large volumes of unstructured text sources, such as news articles, financial reports, social media posts, and analyst commentary. According to Nagel (2019), beliefs data can be possibly extended backwards in time with proxies derived from textual analysis of news using machine learning methods, which is similar to what Manela and Moreira (2017) have done to extend the VIX index. This transition to machine learning has enabled a faster, more accurate pulse on market sentiment, enhancing our ability to identify sentiment-driven market dynamics more accurately than ever before. This contrasts with the reverse-engineering approach that deduces sentiment indirectly by piecing together beliefs based on observed effects in the market. Using beliefs data directly provides a more straightforward way of integrating beliefs into models, ideally yielding more precise insights into how beliefs drive decisions.

5.7 Summary

This thesis presents a cohesive exploration of how investor sentiment influences asset pricing dynamics. Across three research chapters, the work systematically demonstrates that accounting for investor sentiment provides a robust explanation for several persistent anomalies in asset pricing, particularly those associated with cross-sectional return variation and tail risks. Chapter 2 establishes a behavioral pricing kernel that embeds news sentiment directly into the SDF. This framework enhances the explanatory power of conditional asset pricing models while substantially mitigating violations of their underlying assumptions. Chapter 3 introduces a quantile-based asset pricing model that incorporates sentiment-driven asymmetries, capturing how sentiment disproportionately affects extreme returns. Chapter 4 further deepens the behavioral foundation by linking investor sentiment to the tail overweighting parameters proxied by the parameters of probability weighting functions. This thesis makes contributions to both methodological approaches and empirical testing in the literature. The findings can be applied to develop more complex models for further research.

This thesis investigates the integration of machine-learning-extracted news sentiment within several innovative asset pricing models to address longstanding anomalies in financial economics. Across three chapters, it examines how sentiment is integrated into different asset pricing frameworks, particularly emphasizing consumption-based and quantile-based models, as well as time-varying probability weighting. This research builds upon established financial theories, extending them with behavioral insights to better capture the complexities of market sentiment. The work aims to shed light on the critical role of investor sentiment in shaping asset price distributions and risk perceptions, especially tail risk, advancing both theoretical and empirical understandings in financial economics.

Bibliography

Abdelmalek, W. (2023). Investor sentiment, realized volatility and stock returns. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.XXXXXXX

Adam, K., & Nagel, S. (2022). Expectations data in asset pricing [Working Paper No. 2022-59]. In M. P. Devereux & S. J. Peeters (Eds.), *Handbook of economic expectations* (pp. xx–xx). Elsevier.

Ai, H. (2005, August). *Smooth nonexpected utility without state independence* (tech. rep. No. Working Paper 637). Federal Reserve Bank of Minneapolis Research Department. https://researchdatabase.minneapolisfed.org/downloads/sj1392036

Aït-Sahalia, Y., & Duarte, J. (2003). Nonparametric option pricing under shape restrictions. *Journal of Econometrics*, *116*(1–2), 9–47. https://doi.org/10.1016/S0304-4076(03)00102-7

Aït-Sahalia, Y., & Lo, A. W. (1998). Nonparametric estimation of state-price densities implicit in financial asset prices. *Journal of Finance*, *53*(2), 499–547. https://doi.org/10.1111/0022-1082.215228

Aït-Sahalia, Y., & Lo, A. W. (2000a). Nonparametric risk management and implied risk aversion. *Journal of Econometrics*, *94*(1-2), 9–51. https://doi.org/10.1016/S0304-4076(99)00016-0

Aït-Sahalia, Y., & Lo, A. W. (2000b). Nonparametric risk management and implied risk aversion. *Journal of Econometrics*, *94*(1–2), 9–51. https://doi.org/10.1016/S0304-4076(99)00015-0

Allen, D. E., McAleer, M. J., & Singh, A. K. (2019). Daily market news sentiment and stock prices. *Applied Economics*, *51*(30), 3212–3235. https://doi.org/10.1080/00036846.2018.1564115

Ando, T., & Bai, J.-S. (2017). Multivariate quantile regression with factor structures. *Econometric Theory*, *33*(2), 345–371.

Ang, A., Hodrick, R. J., Xing, Y., & Zhang, X. (2006). The cross-section of volatility and expected returns. *Journal of Finance*, *61*(1), 259–299.

Ang, A., Hodrick, R. J., Xing, Y., & Zhang, X. (2009). High idiosyncratic volatility and low returns: International and further u.s. evidence. *Journal of Financial Economics*, *91*(1), 1–23.

Arditti, F. D. (1967). Risk and the required return on equity. *The Journal of Finance*, *22*(1), 19–36. https://doi.org/10.1111/j.1540-6261.1967.tb01651.x

Arrow, K. J. (1964). The role of securities in the optimal allocation of risk-bearing. *Review of Economic Studies*, *31*(2), 91–96. https://doi.org/10.2307/2296188

Arrow, K. J. (1971). *Essays in the theory of risk-bearing*. North-Holland.

Azimi, M., Ghazi, S., & Schneider, M. (2024). Probability weighting and equity premium prediction: Investing with optimism [Forthcoming, DOI: 10.1111/fima.12477]. *Financial Management*.

Baele, L., Driessen, J., Ebert, S., Londono, J. M., & Spalt, O. G. (2017). Cumulative prospect theory, option returns, and the variance premium. https://doi.org/10.2139/ssrn.2932382

Baele, L., Driessen, J., Ebert, S., Londono, J. M., & Spalt, O. G. (2019). Cumulative prospect theory, option returns and the variance premium. *Review of Financial Studies*, *32*(9), 3667–3723. https://doi.org/10.1093/rfs/hhy127

Baker, M., & Stein, J. C. (2004). Market liquidity as a sentiment indicator. *Journal of Financial Markets*, *7*(4), 437–480.

Baker, M., & Wurgler, J. (2004). A catering theory of dividends. *Journal of Finance*, *59*(3), 1125–1165.

Baker, M., & Wurgler, J. (2006). Investor sentiment and the cross-section of stock returns. *Journal of Finance*, *61*(4), 1645–1680. https://doi.org/10.1111/j.1540-6261.2006.00885.x

Baker, M., & Wurgler, J. (2007). Investor sentiment in the stock market. *Journal of Economic Perspectives*, *21*(2), 129–152. https://doi.org/10.1257/jep.21.2.129

Baker, M., Wurgler, J., & Yuan, Y. (2012). Global, local, and contagious investor sentiment. *Journal of Financial Economics*, *104*(2), 272–292.

Bakshi, G., Cao, C., & Chen, Z. (1997). Empirical performance of alternative option pricing models. *Journal of Finance*, *52*(5), 2003–2049. https://doi.org/10.1111/j.1540-6261.1997.tb02749.x

Bakshi, G., Madan, D., & Panayotov, G. (2010). Returns of claims on the upside and the viability of u-shaped pricing kernels. *Journal of Financial Economics*, *97*(1), 130–154. https://doi.org/10.1016/j.jfineco.2010.03.009

Bansal, R., Coleman, W., & Hansen, L. P. (1993). Nonparametric methods in asset pricing. *Journal of Financial Economics*, *34*(1), 3–35. https://doi.org/10.1016/0304-405X(93)90002-7

Bansal, R., Hsieh, D. A., & Yaron, A. (1993). No arbitrage and arbitrage pricing: A new approach. *The Journal of Finance*, *48*(4), 1231–1262. https://doi.org/10.1111/j.1540-6261.1993.tb04747.x

Bansal, R., & Viswanathan, S. P. M. (1993). A nonparametric approach to stochastic discount factor estimation. *Journal of Finance*, *48*(4), 1749–1781. https://doi.org/10.1111/j.1540-6261.1993.tb04090.x

Bansal, R., & Viswanathan, S. (1993). No arbitrage and the term structure of interest rates with switching state variables. *Journal of Finance*, *48*(5), 1543–1571.

Bansal, R., & Yaron, A. (2004). Risks for the long run: A potential resolution of asset pricing puzzles. *Journal of Finance*, *59*(4), 1481–1509. https://doi.org/10.1111/j.1540-6261.2004.00670.x

Bansal, R., & Yaron, A. (2012). Risks for the long run: Estimation and inference. *Econometrica*, *80*(6), 2219–2277. https://doi.org/10.3982/ECTA7610

Banz, R. W., & Miller, M. H. (1978). Prices for state-contingent claims: Some estimates and applications. *Journal of Business*, *51*(4), 653–672. https://doi.org/10.1086/296026

Barberis, N. (2013a). The psychology of tail events: Progress and challenges. *The American Economic Review*, *103*(3), 611–616. https://doi.org/10.1257/aer.103.3.611

Barberis, N. (2013b). Thirty years of prospect theory in economics: A review and assessment. *Journal of Economic Perspectives*, *27*(1), 173–196.

Barberis, N., & Huang, M. (2008). "stocks as lotteries: The implications of probability weighting for security prices". *American Economic Review*, *98*(5), 2066–2100. https://doi.org/10.1257/aer.98.5.2066

Barberis, N., Huang, M., & Santos, T. (2001). Prospect theory and asset prices. *The Quarterly Journal of Economics*, *116*(1), 1–53.

Barberis, N., Mukherjee, A., & Wang, B. (2016). Prospect theory and stock returns: An empirical test. *The Review of Financial Studies*, *29*(11), 3108–3139. https://doi.org/10.1093/rfs/hhw049

Barberis, N., Shleifer, A., & Vishny, R. W. (1998). A model of investor sentiment. *Journal of Financial Economics*, *49*(3), 307–343. https://doi.org/10.1016/S0304-405X(98)00027-0

Barberis, N., & Thaler, R. (2003). A survey of behavioral finance. *Handbook of the Economics of Finance*, *1*, 1053–1128.

Barillas, F., & Shanken, J. (2017). Which alpha? *Review of Financial Studies*, *30*(4), 1316–1338. https://doi.org/10.1093/rfs/hhw101

Barone-Adesi, G., & Dall'O, H. (2010, April). *Is the price kernel monotone?* (Tech. rep. No. 10-03). Swiss Finance Institute. https://ssrn.com/abstract=1540446

Barone-Adesi, G., Engle, R. F., & Mancini, L. (2008). *A garch option pricing model with filtered historical simulation* (tech. rep.). CFA Institute Research Foundation. https://ssrn.com/abstract=603382

Barone-Adesi, G., Mancini, L., & Shefrin, H. (2012). Behavioral finance and the pricing kernel puzzle: Estimating risk aversion, optimism, and overconfidence. *Unpublished Manuscript Swiss Finance Institute*, *3*(7), 1–50.

Barone-Adesi, G., Mancini, L., & Shefrin, H. (2017). Estimating sentiment, risk aversion, and time preference from behavioral pricing kernel theory. *Swiss Finance Institute Research Paper*, (12-21).

Bassett, G. W., Koenker, R., & Nychka, D. (2004). Linear programming for pessimistic optimization: A choquet utility approach. *Journal of Economic Theory*, *119*(2), 243–264.

Baur, D. G., Dimpfl, T., & Jung, R. (2012). Quantile autoregressions for conditional return distribution analysis. *Journal of Financial Econometrics*, *10*(3), 391–417.

Baur, D. G., Dimpfl, T., & Jung, R. C. (2012). Stock return autocorrelations revisited: A quantile regression approach. *Journal of Empirical Finance*, *19*(2), 254–265. https://doi.org/10.1016/j.jempfin.2011.12.002

Bekiros, S. D., Gupta, R., & Kyei, C. (2016). On the relationship between investor sentiment and stock market dynamics: A nonlinear causality approach. *Physica A: Statistical Mechanics and its Applications*, *463*, 122–138.

Benartzi, S., & Thaler, R. H. (1995). Myopic loss aversion and the equity premium puzzle. *Quarterly Journal of Economics*, *110*(1), 73–92. https://doi.org/10.2307/2946696

Berns, G. S., Capra, C. M., Moore, S., & Noussair, C. N. (2007). A shocking experiment: New evidence on probability weighting and common ratio violations. *Judgment and Decision Making*, *2*(4), 234–242.

Bhardwaj, A., Narayan, Y., & Dutta, M. (2015). Sentiment analysis for indian stock market prediction using sensex and nifty. *Procedia Computer Science*, *70*, 85–91. https://doi.org/10.1016/j.procs.2015.10.043

Bhushan, R., Brown, D. P., & Mello, A. S. (1997). Do noise traders "create their own space?" *Journal of Financial and Quantitative Analysis*, *32*(1), 25–45. https://doi.org/10.2307/2331315

Birbeck, E., & Cliff, D. (2018). Using stock prices as ground truth in sentiment analysis to generate profitable trading signals. *2018 IEEE Symposium Series on Computational Intelligence (SSCI)*, 1–8. https://doi.org/10.1109/SSCI.2018.8628741

Black, F. (1986). Noise. *The Journal of Finance*, *41*(3), 529–543. https://doi.org/10.1111/j.1540-6261.1986.tb04513.x

Black, F. (1992). Global risk and asset pricing. *Financial Analysts Journal*, *48*(3), 28–38.

Bliss, R. R., & Panigirtzoglou, N. (2004). Option-implied risk aversion estimates. *The Journal of Finance*, *59*(1), 407–446. https://doi.org/10.1111/j.1540-6261.2004.00636.x

Bodurtha, J. N., Kim, D.-S., & Lee, C. M. C. (1995). Closed-end country funds and u.s. market sentiment. *The Review of Financial Studies*, *8*(3), 879–918.

Bouri, E., Molaverdikhani, K., & Shahzad, S. J. H. (2022). Investor sentiment and stock market volatility: Evidence from us and european markets. *Journal of International Financial Markets, Institutions and Money*, *79*, 101519.

Breeden, D. T. (2005). *An intertemporal asset pricing model with stochastic consumption and investment opportunities*. World Scientific.

Breeden, D. T., Gibbons, M. R., & Litzenberger, R. H. (1989). Empirical tests of the consumption-oriented capm. *The Journal of Finance*, *44*(2), 231–262.

Breeden, D. T., & Litzenberger, R. H. (1978). Prices of state-contingent claims implicit in option prices. *Journal of Business*, *51*(4), 621–651. https://doi.org/10.1086/296025

Breeden, D. T., Litzenberger, R. H., & Jia, T. (2015). Consumption-based asset pricing, part 1: Classic theory and tests, measurement issues, and limited participation. *Annual Review of Financial Economics*, *7*, 35–83.

Brown, D. P., & Gibbons, M. R. (1985). A simple econometric approach for utility-based asset pricing models. *The Journal of Finance*, *40*(2), 359–381.

Brown, G. W. (1999). Volatility, sentiment, and noise traders. *Financial Analysts Journal*, *55*(2), 82–90. https://doi.org/10.2469/faj.v55.n2.2263

Calomiris, C. W., & Mamaysky, H. (2019). How news and its context drive risk and returns around the world. *Journal of Financial Economics*, *133*(2), 299–336.

Camerer, C. F. (1995). Individual decision making: A survey. *Journal of Economic Behavior & Organization*, *27*(1), 1–31.

Camerer, C. F., & Ho, T.-H. (1994). Violations of the betweenness axiom and nonlinearity in probability. *Journal of Risk and Uncertainty*, *8*(2), 167–196. https://doi.org/10.1007/BF01065371

Campbell, J. Y., & Cochrane, J. H. (1999). By force of habit: A consumption-based explanation of aggregate stock market behavior. *Journal of Political Economy*, *107*(2), 205–251.

Campbell, J. Y., & Shiller, R. J. (1988). The dividend-price ratio and expectations of future dividends and discount factors. *The Review of Financial Studies*, *1*(3), 195–228.

Cappiello, L., Engle, R. F., & Sheppard, K. (2014). Time-varying comovements in quantile regression. *Journal of Econometrics*, *178*(1), 85–101.

Carhart, M. M. (1997). On persistence in mutual fund performance. *The Journal of Finance*, *52*(1), 57–82. https://doi.org/10.1111/j.1540-6261.1997.tb03808.x

Carlier, G., & Dana, R.-A. (2003). Core of convex distortions of a probability. *Journal of Economic Theory*, *113*(2), 199–222. https://doi.org/10.1016/S0022-0531(03)00122-4

Cech, F., & Barunik, J. (2017). Measurement of common risk factors: A panel quantile regression model for returns. *arXiv preprint arXiv:1708.08622*.

Chabi-Yo, F. (2011). Explaining the idiosyncratic volatility puzzle using stochastic discount factors. *Journal of Banking & Finance*, *35*(8), 1971–1983.

Chabi-Yo, F., & Song, Z. (2013). *Recovering the probability weights of tail events with volatility risk from option prices* [Working Paper, Fisher College of Business, Ohio State University].

Chambers, C. P. (2007). Choquet expected utility and quantile decisions. *Journal of Economic Theory*, *136*(1), 146–172.

Chambers, C. P. (2009). *An axiomatization of quantiles on the domain of distribution functions* (Vol. 19). Wiley Online Library.

Chang, C., Hsieh, S., & Wang, Y. (2016). Investor sentiment and the predictability of stock returns: A quantile and granger causality approach. *International Review of Financial Analysis*, *47*, 122–134.

Chapman, D. A. (1997a). A characterization of the stochastic discount factor and the asset pricing kernel. *Journal of Finance*, *52*(1), 281–308.

Chapman, D. A., & Polkovnichenko, V. (2009). "heterogeneous preferences and asset prices: Evidence from individual investor portfolios". *Journal of Finance*, *64*(4), 1365–1430. https://doi.org/10.1111/j.1540-6261.2009.01467.x

Chapman, D. A. (1997b). Kernel estimation of pricing kernels. *Review of Economics and Statistics*, *79*(3), 387–395. https://doi.org/10.1162/003465397557069

Chau, F., Deesomsak, R., & Koutmos, D. (2016). Does investor sentiment really matter? *International Review of Financial Analysis*, *48*, 221–232.

Chen, X., Hong, H., & Lu, Y. (2016). Factor models in multivariate quantile regression. *Review of Economics and Statistics*, *98*(4), 612–627.

Chen, Y., Han, B., & Pan, J. (2021). Sentiment trading and hedge fund returns [Forthcoming]. *Journal of Finance*.

Chernozhukov, V., & Hansen, C. (2005). An iv model of quantile treatment effects. *Econometrica*, *73*(1), 245–261. https://doi.org/10.1111/j.1468-0262.2005.00571.x

Chesher, A. (2003). Identification in nonseparable models. *Econometrica*, *71*(5), 1405–1441. https://doi.org/10.1111/1468-0262.00412

Chung, H., Hung, M.-C., & Yeh, C.-Y. (2012). When does investor sentiment predict stock returns? *Journal of Empirical Finance*, *19*(2), 217–240. https://doi.org/10.1016/j.jempfin.2011.11.002

Cochrane, J. H. (1996). A cross-sectional test of an investment-based asset pricing model. *Journal of Political Economy*, *104*(3), 572–621. https://doi.org/10.1086/262029

Cochrane, J. H. (2009). *Asset pricing: Revised edition*. Princeton University Press. https://www.johnhcochrane.com/asset-pricing

Conlisk, J. (1996). Why bounded rationality? *Journal of Economic Literature*, *34*(2), 669–700.

Corbet, S. W., Meegan, A., Larkin, C., Lucey, B., & Yarovaya, L. (2020). Cryptocurrency market reactions to covid-19: Fundamentals, uncertainty and sentiment. *Finance Research Letters*, *35*, 101554.

Cremonesi, P., Ceccarani, C., Curone, G., Severgnini, M., Pollera, C., et al. (2018). Milk microbiome diversity and bacterial group prevalence in a comparison between healthy holstein friesian and rendena cows. *PLOS ONE*, *13*(10), e0205054. https://doi.org/10.1371/journal.pone.0205054

Cuesdeanu, H., & Jackwerth, J. C. (2018). The pricing kernel puzzle: Survey and outlook. *Annals of Finance*, *14*(3), 289–329. https://doi.org/10.1007/s10436-017-0317-9

Cujea, S., & Hasler, M. (2016). Sentiment and stock returns: The role of media — a non-linear approach. *Finance Research Letters*, *18*, 306–311. https://doi.org/10.1016/j.frl.2016.05.001

Da, Z., Engelberg, J., & Gao, P. (2015). The sum of all fears: Investor sentiment and asset prices. *Journal of Financial Economics*, *114*(1), 1–24. https://doi.org/10.1016/j.jfineco.2014.05.005

Daniel, K., Hirshleifer, D., & Subrahmanyam, A. (1998). Investor psychology and security market under- and overreactions. *Journal of Finance*, *53*(6), 1839–1885.

De Bondt, W. F. M., & Thaler, R. H. (1985). Does the stock market overreact? *Journal of Finance*, *40*(3), 793–805.

De Bondt, W. F. M., & Thaler, R. H. (1987). Further evidence on investor overreaction and stock market seasonality. *Journal of Finance*, *42*(2), 557–581.

De Castro, L., & Galvão, A. F. (2019). Dynamic quantile models of rational behavior. *Econometrica*, *87*(6), 1893–1939.

De Long, J. B., Shleifer, A., Summers, L. H., & Waldmann, R. J. (1990). Noise trader risk in financial markets. *Journal of Political Economy*, *98*(4), 703–738. https://doi.org/10.1086/261703

Debreu, G. (1959). *Theory of value: An axiomatic analysis of economic equilibrium.* Yale University Press. https://cowles.yale.edu/sites/default/files/files/pub/mon/m17-all.pdf

de Castro, L., & Galvao, A. F. (2017, March). Dynamic quantile models of rational behavior. https://doi.org/10.2139/ssrn.2932382

de Castro, L., & Galvao, A. F. (2019). Dynamic quantile models of rational behavior. *Econometrica*, *87*(6), 1893–1939. https://doi.org/10.3982/ECTA15146

de Castro, L., & Galvao, A. F. (2021). Static and dynamic quantile preferences. *Economic Theory*, *73*, 747–779. https://doi.org/10.1007/s00199-021-01355-8

de Castro, L., Galvão, A. F., Montes-Rojas, G., & Olmo, J. (2022). Portfolio selection in quantile decision models. *Annals of Finance*, *18*(2), 133–181. https://doi.org/10.1007/s10436-021-00405-4

de Castro, L., Galvão, A. F., Noussair, C. N., & Qiao, L. (2022). Do people maximize quantiles? *Games and Economic Behavior*, *132*, 22–40. https://doi.org/10.1016/j.geb.2021.11.010

Delikouras, S., & Kostakis, A. (2019). A single-factor consumption-based asset pricing model. *Journal of Financial and Quantitative Analysis*, *54*(2), 789–827.

Dergiades, T. (2012). Do investors' sentiment dynamics affect stock returns? evidence from european markets. *Finance Research Letters*, *9*(1), 1–7.

Dierkes, M., Krupski, J., Schroen, S., & Sibbertsen, P. (2024). Volatility-dependent probability weighting and the dynamics of the pricing kernel puzzle. *Review of Derivatives Research*, *27*(1), 1–35. https://doi.org/10.1007/s11147-023-09197-3

Ding, D., Hyung, N., & Sung, S. (2016). Asymmetric effects of investor sentiment on stock returns: Evidence from quantile regression. *Applied Economics Letters*, *23*(5), 327–330. https://doi.org/10.1080/13504851.2015.1071462

Dittmar, R. F. (2002). Nonlinear pricing kernels, kurtosis preference, and evidence from the cross section of equity returns. *The Journal of Finance*, *57*(1), 369–403. https://doi.org/10.1111/1540-6261.00425

Dybvig, P. H. (1988). Inefficient dynamic portfolio strategies or how to throw away a million dollars in the stock market. *The Review of Financial Studies*, *1*(1), 67–88.

Echenique, F., & Komunjer, I. (2009). Multiple equilibria in econometric models: Testing and confidence sets. *Journal of Econometrics*, *148*(1), 17–31.

Edmans, A., García, D., & Norli, Ø. (2007). U.s. stock price crashes: What, when, and why? *Review of Financial Studies*, *20*(5), 2233–2278.

Eichenbaum, M., Hansen, L. P., & Singleton, K. J. (1988). Testing the consistency of consumption and asset returns with a representative consumer model. *Journal of Econometrics*, *39*(1-2), 295–321.

Engle, R. F. (2002). Dynamic conditional correlation: A simple class of multivariate generalized autoregressive conditional heteroskedasticity models. *Journal of Business & Economic Statistics*, *20*(3), 339–350. https://doi.org/10.1198/073500102288618487

Engle, R. F., & Manganelli, S. (2004). Caviar: Conditional autoregressive value at risk by regression quantiles. *Journal of Business & Economic Statistics*, *22*(4), 367–381. https://doi.org/10.1198/073500104000000370

Epstein, L. G., & Zin, S. E. (1989). Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework. *Econometrica*, *57*(4), 937–969.

Epstein, L. G., & Zin, S. E. (1990). "first-order risk aversion and the equity premium puzzle". *Journal of Monetary Economics*, *26*(3), 387–407. https://doi.org/10.1016/0304-3932(90)90042-T

Epstein, L. G., & Zin, S. E. V. (1989). Risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework. *V Econometrica*, *57*(4), 937.

Fama, E. F. (1965). The behavior of stock market prices. *The Journal of Business*, *38*(1), 34–105. https://doi.org/10.1086/294743

Fama, E. F. (1991). Efficient capital markets ii. *Journal of Finance*, *46*(5), 1575–1617. https://doi.org/10.1111/j.1540-6261.1991.tb04636.x

Fama, E. F., & French, K. R. (1992). The cross-section of expected stock returns. *Journal of Finance*, *47*(2), 427–465.

Fama, E. F., & French, K. R. (1993). Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics*, *33*(1), 3–56.

Fama, E. F., & French, K. R. (1996). Multifactor explanations of asset pricing anomalies. *The Journal of Finance*, *51*(1), 55–84. https://doi.org/10.1111/j.1540-6261.1996.tb05202.x

Fama, E. F., & French, K. R. (2021). *Dividend yields and expected stock returns*. University of Chicago Press.

Fama, E. F., & MacBeth, J. D. (1973). Risk, return, and equilibrium: Empirical tests. *Journal of Political Economy*, *81*(3), 607–636. https://doi.org/10.1086/260061

Félix, L. F. F., Kräussl, R., & Stork, P. (2020). Implied volatility sentiment: A tale of two tails. *Quantitative Finance*, *20*(5), 823–849. https://doi.org/10.1080/14697688.2019.1696018

Figlewski, M. (2010). Estimating the implied risk neutral density for the u.s. market portfolio. In R. F. Engle (Ed.), *Volatility and time series econometrics: Essays in honor of robert f. engle*. Oxford University Press.

Figlewski, S. (2010). Estimating the implied risk-neutral density for the u.s. market portfolio. In T. Bollerslev, J. R. Russell, & M. W. Watson (Eds.), *Volatility and time series econometrics: Essays in honor of robert f. engle* (pp. 323–353). Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199549498.003.0015

Firpo, S. (2007). Efficient semiparametric estimation of quantile treatment effects. *Econometrica*, *75*(1), 259–276. https://doi.org/10.1111/j.1468-0262.2007.00725.x

Fox, C. R., & Tversky, A. (1998). A belief-based account of decision under uncertainty. *Management Science*, *44*(7), 879–895. https://doi.org/10.1287/mnsc.44.7.879

Friedman, M. (1953a). The case for flexible exchange rates. In *Essays in positive economics*. University of Chicago Press.

Friedman, M. (1953b). *Essays in positive economics*. University of Chicago Press.

Frugier, F. (2016). Investor sentiment and equity returns: A behavioral perspective. *Journal of Economic Behavior & Organization*, *130*, 321–336.

Galvao, A. F. (2011). Quantile regression for dynamic panel data with fixed effects. *Journal of Econometrics*, *164*(1), 142–157. https://doi.org/10.1016/j.jeconom.2011.02.016

Garcia, D. (2013). Sentiment during recessions. *The Journal of Finance*, *68*(3), 1267–1300. https://doi.org/10.1111/jofi.12027

Garman, M., & Klass, M. (1980). On the estimation of the covariance matrix of stock returns. *Journal of Financial and Quantitative Analysis*, *15*(1), 1–13.

Ghallab, A., Mohsen, A., & Ali, Y. (2020). Arabic sentiment analysis: A systematic literature review. *Applied Computational Intelligence and Soft Computing*, *2020*, 7403128. https://doi.org/10.1155/2020/7403128

Gino, F., Brooks, A. W., & Schweitzer, M. E. (2012). Anxiety, advice, and the ability to discern: Feeling anxious motivates individuals to seek and use advice. *Journal of Personality and Social Psychology*, *102*(3), 497–512. https://doi.org/10.1037/a0026413

Giovannetti, B. C. (2013). Asset pricing under quantile utility maximization. *Review of Financial Economics*, *22*(4), 169–179. https://doi.org/10.1016/j.rfe.2013.05.008

Glosten, L. R., Jagannathan, R., & Runkle, D. E. (1993). On the relation between the expected value and the volatility of the nominal excess return on stocks. *Journal of Finance*, *48*(5), 1779–1801. https://doi.org/10.1111/j.1540-6261.1993.tb05128.x

Grinblatt, M., & Keloharju, M. (2009). Sensation seeking, overconfidence, and trading activity. *The Journal of Finance*, *64*(2), 549–578. https://doi.org/10.1111/j.1540-6261.2009.01443.x

Grossman, S. J., & Stiglitz, J. E. (1980). On the impossibility of informationally efficient markets. *The American Economic Review*, *70*(3), 393–408. https://www.jstor.org/stable/1805228

Han, Y., Zhou, G., & Zhu, H. (2016). Investor sentiment, stock returns, and trading volume: Evidence from quantile regression and fama–macbeth analysis. *Journal of Behavioral Finance*, *17*(3), 245–262.

Hansen, L. P., Heaton, J. C., & Li, N. (2008). Consumption strikes back?: Measuring long-run risk. *Journal of Political Economy*, *116*(2), 260–302. https://doi.org/10.1086/588200

Hansen, L. P., & Jagannathan, R. (1991). Implications of security market data for models of dynamic economies. *Journal of political economy*, *99*(2), 225–262.

Hansen, L. P., & Singleton, K. J. (1982). Generalized instrumental variables estimation of nonlinear rational expectations models. *Econometrica*, *50*(5), 1269–1286. https://doi.org/10.2307/1911873

Hansen, L. P., & Singleton, K. J. (1983). Stochastic consumption, risk aversion, and the temporal behavior of asset returns. *Journal of Political Economy*, *91*(2), 249–265. https://doi.org/10.1086/261141

Harrison, J. M., & Kreps, D. M. (1979). Martingales and arbitrage in multiperiod securities markets. *Journal of Economic Theory*, *20*(3), 381–408. https://doi.org/10.1016/0022-0531(79)90043-7

Harvey, C. R., & Liu, Y. (2014). Evaluating trading strategies. *The Journal of Portfolio Management*, *40*(5), 108–118.

He, X. D., Kouwenberg, R., & Zhou, X. Y. (2017). Inverse S-shaped probability weighting and its impact on investment [Working paper]. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=xxxxx

Hens, T., & Reichlin, C. (2013a). Three solutions to the pricing kernel puzzle. *Review of Finance*, *17*(3), 1065–1098. https://doi.org/10.1093/rof/rfs008

Hens, T., & Reichlin, C. (2013b). Three solutions to the pricing kernel puzzle. *Review of Finance*, *17*(3), 1065–1098.

Hirshleifer, D., Jiang, D., & DiGiovanni, N. (2020). Mood beta and seasonalities in stock returns. *Journal of Financial Economics*, *137*(2), 495–520.

Hong, H., & Stein, J. C. (1999). A unified theory of underreaction, momentum trading, and overreaction in asset markets. *Journal of Finance*, *54*(6), 2143–2184.

Huang, C.-f., & Litzenberger, R. H. (1988). *Foundations for financial economics*. North-Holland.

Huang, D., Jiang, F., Tu, J., & Zhou, G. (2015). Investor sentiment aligned: A powerful predictor of stock returns. *Review of Financial Studies*, *28*(3), 791–837. https://doi.org/10.1093/rfs/hhu080

Hughen, B., & McDonald, S. J. (2005). Investor sentiment and market reaction: An empirical analysis. *Journal of Behavioral Finance*, *6*(4), 210–224.

Imbens, G., & Newey, W. (2009). Identification and estimation of triangular simultaneous equations models without additivity. *Econometrica*, *77*(5), 1481–1505. https://doi.org/10.3982/ECTA7511

Jackwerth, J. C. (2000). Recovering risk aversion from option prices and realized returns. *Review of Financial Studies*, *13*(2), 433–451.

Jackwerth, J. C. (2004a). *Option-implied risk-neutral distributions and risk aversion* (tech. rep.) (Monograph). The Research Foundation of the CFA Institute. Charlottesville, VA.

Jackwerth, J. C. (2004b, March). *Option-implied risk-neutral distributions and risk aversion* (Research Foundation Monograph). CFA Institute Research Foundation. https://ssrn.com/abstract=1303779

Jackwerth, J. C., & Vilkov, G. (2017). *Asymmetric volatility risk: Evidence from option markets* (Working Paper No. 2325380). SSRN Electronic Journal. https://ssrn.com/abstract=2325380

Jagannathan, R., & Wang, Z. (1996). The conditional capm and the cross-section of expected returns. *The Journal of Finance*, *51*(1), 3–53. https://doi.org/10.1111/j.1540-6261.1996.tb05201.x

Jegadeesh, N., & Titman, S. (1993a). Returns to buying winners and selling losers: Implications for stock market efficiency. *Journal of Finance*, *48*(1), 65–91.

Jegadeesh, N., & Titman, S. (1993b). Returns to buying winners and selling losers: Implications for stock market efficiency. *The Journal of Finance*, *48*(1), 65–91. https://doi.org/10.1111/j.1540-6261.1993.tb04702.x

Jegadeesh, N., & Titman, S. (2001). Momentum. *Journal of Finance*, *55*(4), 1685–1713.

Jha, A., Shirvani, A., Jaffri, A., Rachev, S. T., & Fabozzi, F. J. (2025). Behavioral probability weighting and portfolio optimization under semi-heavy tails. *arXiv preprint arXiv:2507.04208*. https://arxiv.org/abs/2507.04208v1

Jr., R. E. L. (1970). Expectations and the neutrality of money. *Journal of Economic Theory*, *4*(2), 103–124. https://doi.org/10.1016/0022-0531(70)90002-6

Judd, K. L. (1992). Projection methods for solving aggregate growth models. *Journal of Economic Theory*, *58*(2), 410–452. https://doi.org/10.1016/0022-0531(92)90061-L

Kahneman, D., Slovic, P., & Tversky, A. (1982). *Judgment under uncertainty: Heuristics and biases*. Cambridge University Press.

Kahneman, D., & Tversky, A. (1979a). Prospect theory: An analysis of decision under risk. *Econometrica*, *47*(2), 263–291.

Kahneman, D., & Tversky, A. (1979b). Prospect theory: An analysis of decision under risk. *Econometrica*, *47*(2), 263–292. https://doi.org/10.2307/1914185

Kahneman, D., & Tversky, A. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, *5*(4), 297–323. https://doi.org/10.1007/BF00122574

Kamdar, R. (2019, June). *The inattentive consumer: Sentiment and expectations* (Meeting Paper No. 647). Society for Economic Dynamics. https://red-files-public.s3.amazonaws.com/meetpapers/2019/paper_647.pdf

Kamstra, M., Kramer, D., & Levi, M. D. (2000). Winter blues: A seasonal affective disorder–based model of stock returns. *American Economic Review*, *90*(2), 83–106.

Kelly, M. (1997). Do noise traders influence stock prices? *Journal of Money, Credit and Banking*, *29*(3), 351–363. https://doi.org/10.2307/2953699

Keynes, J. M. (1936). *The general theory of employment, interest, and money.* Macmillan.

Kimball, M. S. (1990). Precautionary saving in the small and in the large. *Econometrica*, *58*(1), 53–73. https://doi.org/10.2307/2938334

Kimball, M. S. (1993). Standard risk aversion. *Econometrica*, *61*(3), 589–611. https://doi.org/10.2307/2951719

Koenker, R. (2004). Quantile regression for longitudinal data. *Journal of Multivariate Analysis*, *91*(1), 74–89. https://doi.org/10.1016/j.jmva.2004.05.006

Koenker, R., & Bassett, G. (1978). Regression quantiles. *Econometrica*, *46*(1), 33–50. https://doi.org/10.2307/1913643

Koenker, R., & Jr., G. B. (1978). Regression quantiles. *Econometrica*, *46*(1), 33–50. https://doi.org/10.2307/1913643

Kreps, D. M., & Porteus, E. L. (1978). Temporal resolution of uncertainty and dynamic choice theory. *Econometrica: journal of the Econometric Society*, 185–200.

Lamont, O. (1998). Earnings and expected returns. *The journal of Finance*, *53*(5), 1563–1587.

Lee, C. M. C., Jiang, W., & Indro, D. C. (2002). Stock market volatility, excess returns, and investor sentiment: Evidence from the uk and the us. *Journal of Banking & Finance*, *26*(12), 2277–2299.

Lee, C. M. C., Shleifer, A., & Thaler, R. H. (1991a). Investor sentiment and asset prices. *Journal of Finance*, *46*(3), 881–906.

Lee, C. M. C., Shleifer, A., & Thaler, R. H. (1991b). Investor sentiment and the closed-end fund puzzle. *Journal of Finance*, *46*(1), 75–109.

Lemmon, M., & Portniaguina, E. (2006). Consumer confidence and asset prices: Some empirical evidence. *The Review of Financial Studies*, *19*(4), 1499–1529. https://doi.org/10.1093/rfs/hhj038

Lettau, M., & Ludvigson, S. (2001a). Consumption, aggregate wealth, and expected stock returns. *the Journal of Finance*, *56*(3), 815–849.

Lettau, M., & Ludvigson, S. (2001b). Resurrecting the (c) capm: A cross-sectional test when risk premia are time-varying. *Journal of political economy*, *109*(6), 1238–1287.

Lettau, M., & Ludvigson, S. C. (2001a). Consumption, aggregate wealth, and expected stock returns. *The Journal of Finance*, *56*(3), 815–849. https://doi.org/10.1111/0022-1082.00347

Lettau, M., & Ludvigson, S. C. (2001b). Resurrecting the (c)capm: A cross-sectional test when risk premia are time-varying. *Journal of Political Economy*, *109*(6), 1238–1287. https://doi.org/10.1086/323281

Levy, H., & Levy, M. (2004). "prospect theory and mean-variance analysis". *Review of Financial Studies*, *17*(4), 1015–1041. https://doi.org/10.1093/rfs/hhg062

Lewellen, J., Nagel, S., & Shanken, J. (2010). A skeptical appraisal of asset pricing tests. *Journal of Financial economics*, *96*(2), 175–194.

Li, X. (2017). Trading volume and stock returns in extreme market conditions: Evidence from quantile regression. *Finance Research Letters*, *21*, 163–172.

Li, X., Guo, J., & Park, J. (2017). Stock market sentiment and return predictability: Evidence from asian markets. *Pacific-Basin Finance Journal*, *44*, 156–171.

Liang, B. (2018). Investor sentiment and stock returns: Evidence from a consumption-based asset pricing model. *Journal of Financial Economics*, *128*(3), 637–658.

Linnemann, L., & Winkler, R. (2016). Estimating nonlinear effects of fiscal policy using quantile regression. *Oxford Bulletin of Economics and Statistics*, *78*(4), 575–599. https://doi.org/10.1111/obes.12122

Linter, J. (1965). The valuation of risk assets and the selection of risk investements.

Lintner, J. (1965). The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets. *The Review of Economics and Statistics*, *47*(1), 13–37. https://doi.org/10.2307/1924119

Lintner, J. (1969). Dividends, earnings, leverage, stock prices and the supply of capital to corporations. *The Review of Economics and Statistics*, *51*(3), 243–269.

Lo, A. W., & MacKinlay, A. C. (1990). Data-snooping biases in tests of financial asset pricing models. *Review of Financial Studies*, *3*(3), 431–467. https://doi.org/10.1093/rfs/3.3.431

Lopes, L. L. (1987). Between hope and fear: The psychology of risk. In *Advances in experimental social psychology* (pp. 255–295, Vol. 20). Elsevier.

Loughran, T., & Ritter, J. R. (1995). The new issues puzzle. *Journal of Finance*, *50*(1), 23–51.

Lucas Jr, R. E. (1978). *Asset prices in an exchange economy*. JSTOR.

Lux, T. (2011). Sentiment and the dynamics of speculative markets: Animusx sentiment data analysis. *Journal of Economic Dynamics and Control*, *35*(8), 1097–1114.

Lux, T. (2012). Sentiment dynamics and stock returns: The nonlinear case. *Journal of Economic Dynamics and Control*, *36*(10), 1362–1383.

Machina, M. J. (1982). "expected utility" analysis without the independence axiom. *Econometrica: Journal of the Econometric Society*, 277–323.

Manela, A., & Moreira, A. (2017). News implied volatility and disaster concerns. *Journal of Financial Economics*, *123*(1), 137–162.

Mankiw, N. G., & Shapiro, M. D. (1986). News or noise? an application of signal extraction to economic time series. *Quarterly Journal of Economics*, *101*(4), 797–815. https://doi.org/10.2307/1884377

Manski, C. F. (1988a). Decision-theoretic attributes of quantile maximization. *Journal of Econometrics*, *42*(1), 83–98.

Manski, C. F. (1988b). Ordinal utility models of decision making under uncertainty. *Theory and Decision*, *25*(1), 79–104.

Markowitz, H. M. (1952). *Portfolio selection*. Yale university press.

Matzkin, R. L. (2003). Nonparametric identification and estimation of nonadditive models. *Econometrica*, *71*(5), 1443–1472. https://doi.org/10.1111/1468-0262.00413

Mehra, R., & Prescott, E. C. (1985). The equity premium: A puzzle. *Journal of monetary Economics*, *15*(2), 145–161.

Mendelson, H. (1987). Quantile-preserving spread. *Journal of Economic Theory*, *42*(2), 334–351. https://doi.org/10.1016/0022-0531(87)90091-3

Merton, R. C. (1973). An intertemporal capital asset pricing model. *Econometrica: Journal of the Econometric Society*, 867–887.

Mossin, J. (1966). Equilibrium in a capital asset market. *Econometrica: Journal of the econometric society*, 768–783.

Muth, J. F. (1961). Rational expectations and the theory of price movements. *Econometrica*, *29*(3), 315–335. https://doi.org/10.2307/1909635

Nagel, S. (2013). Empirical cross-sectional asset pricing. *Annual Review of Financial Economics*, *5*(1), 167–199. https://doi.org/10.1146/annurev-financial-110112-121009

Nagel, S. (2019). Nber panel on the future of asset pricing: Models of beliefs.

Nagel, S. (2012). What moves stock prices? macro variables or micro shocks? *Journal of Finance*, *67*(4), 1377–1407.

Neal, R., & Wheatley, S. M. (1998). Do measures of investor sentiment predict returns? *Journal of Financial and Quantitative Analysis*, *33*(4), 523–547.

Nicholas Barberis, M. H., & Thaler, R. (2006). “individual preferences, monetary gambles, and stock market participation: A case for narrow framing”. *American Economic Review*, *96*(4), 1069–1090. https://doi.org/10.1257/aer.96.4.1069

Nishiyama, Y., Hitomi, K., Kawasaki, Y., & Jeong, K. (2011). A consistent nonparametric test for nonlinear causality—specification in time series regression. *Journal of Econometrics*, *165*(1), 112–127.

Petersen, M. A. (2009). Estimating standard errors in finance panel data sets: Comparing approaches. *The Review of Financial Studies*, *22*(1), 435–480. https://doi.org/10.1093/rfs/hhn053

Polkovnichenko, V. (2005). "household portfolio diversification: A case for rank-dependent preferences". *Review of Financial Studies*, *18*(4), 1467–1502. https://doi.org/10.1093/rfs/hhi034

Polkovnichenko, V., Wei, K. D., & Zhao, F. (2019). Cautious risk takers: Investor preferences and demand for active management. *The Journal of Finance*, *74*(2), 1025–1075. https://doi.org/10.1111/jofi.12747

Polkovnichenko, V., & Zhao, F. (2013). Probability weighting functions implied in options prices [Originally circulated as working paper, January 2012]. *Journal of Financial Economics*, *107*(3), 580–609. https://doi.org/10.1016/j.jfineco.2012.09.005

Poteshman, A. (2001). The information in option volume for future stock prices. *Review of Financial Studies*, *14*(4), 1097–1131.

Poti, V., & Shefrin, H. (2014). The signature of sentiment in conditional consumption capm estimates: A note. *Journal of Behavioral and Experimental Finance*, *2*, 1–9.

Potì, V., & Shefrin, H. (2014). The signature of sentiment in conditional consumption capm estimates: A note. *Journal of Behavioral and Experimental Finance*, *2*(100), 1–9. https://doi.org/10.1016/j.jbef.2014.02.004

Prelec, D. (1998). *The probability weighting function*. JSTOR.

Qiu, J., & Welch, I. (2006). Ub survey does something... *Some Journal*, *10*, Y–Z.

Quiggin, J. (1982). A theory of anticipated utility. *Journal of economic behavior & organization*, *3*(4), 323–343.

Quiggin, J. (1987). Decision weights in anticipated utility theory: Response to segal. *Journal of Economic Behavior & Organization*, *8*(4), 641–645.

Quiggin, J. (1993). *Generalized expected utility theory: The rank-dependent model*. Springer Science & Business Media.

Reuters, T. (2013). Marketpsych indices (trmi).

Roll, R. (1977). Of the asset pricing theory's tests part 1: On past and potential testability of the theory. *Journal of Financial Economics*, *4*(2), 129–176.

Rosenberg, J. V., & Engle, R. F. (2002). Empirical pricing kernels. *Journal of Financial Economics*, *64*(3), 341–372. https://doi.org/10.1016/S0304-405X(02)00128-9

Ross, S. A. (1976). The arbitrage theory of capital asset pricing. In *Handbook of the fundamentals of financial decision making: Part i* (pp. 11–30). World Scientific.

Rostek, M. (2010). Quantile maximization in decision theory. *The Review of Economic Studies*, *77*(1), 339–371.

Santos, T., & Veronesi, P. (2006). Labor income and predictable stock returns. *The Review of Financial Studies*, *19*(1), 1–44.

Sargent, T. J. (1973). Rational expectations, the real rate of interest, and the natural rate of unemployment. *Brookings Papers on Economic Activity*, *1973*(2), 429–472. https://doi.org/10.2307/2534360

Savage, L. J. (1954). *The foundations of statistics*. Wiley.

Shanken, J. (1991). Conditional asset pricing: A note. *Journal of Finance*, *46*(5), 1569–1575. https://doi.org/10.1111/j.1540-6261.1991.tb04627.x

Shanken, J. (1992). On the estimation of beta-pricing models. *The Review of Financial Studies*, *5*(1), 1–33. https://doi.org/10.1093/rfs/5.1.1

Shapiro, A. (2022). Machine learning-based sentiment extraction in financial news. *Journal of Behavioral Finance*, *23*(4), 369–386.

Shapiro, A. H., Sudhof, M., & Wilson, D. J. (2022). Measuring news sentiment. *Journal of Econometrics*, *228*(2), 221–243.

Sharpe, W. F. (1964a). Capital asset prices: A theory of market equilibrium under conditions of risk. *The Journal of Finance*, *19*(3), 425–442. https://doi.org/10.1111/j.1540-6261.1964.tb02865.x

Sharpe, W. F. (1964b). Capital asset prices: A theory of market equilibrium under conditions of risk. *The journal of finance*, *19*(3), 425–442.

Shefrin, H. (2005). *A behavioral approach to asset pricing*. Elsevier Academic Press.

Shefrin, H. (2008). *A behavioral approach to asset pricing*. Elsevier.

Shefrin, H., & Statman, M. (2000). “behavioral portfolio theory”. *Journal of Financial and Quantitative Analysis*, *35*(2), 127–151. https://doi.org/10.2307/2676187

Shiller, R. J. (1981). Do stock prices move too much to be justified by subsequent changes in dividends? *American Economic Review*, *71*(3), 421–436.

Shleifer, A., & Vishny, R. W. (1997). The limits of arbitrage. *The Journal of Finance*, *52*(1), 35–55. https://doi.org/10.1111/j.1540-6261.1997.tb03807.x

Shoemaker, P. J. H. (1982). Behavioral decision theory: A research survey and conceptual framework. *Acta Psychologica*, *50*(1), 13–45.

Sichert, T. (2018). The pricing kernel is u-shaped [Available at https://ssrn.com/abstract=3095551]. https://doi.org/10.2139/ssrn.3095551

Smales, L. A. (2015a). Asymmetric volatility response to news sentiment in gold futures. *Journal of International Financial Markets, Institutions and Money*, *34*, 161–172.

Smales, L. A. (2015b). Time-variation in the impact of news sentiment. *International Review of Financial Analysis*, *37*, 40–50. https://doi.org/10.1016/j.irfa.2014.11.019

Smales, L. A. (2016). The importance of fear: Investor sentiment and stock market returns [Forthcoming; working paper available at https://ssrn.com/abstract=2749518]. *Applied Economics*.

Stambaugh, R. F., Yu, J., & Yuan, Y. (2012). The short of it: Investor sentiment and anomalies. *Journal of Financial Economics*, *104*(2), 288–302.

Starmer, C. (2000). Developments in non-expected utility theory: The hunt for a descriptive theory of choice under risk. *Journal of Economic Literature*, *38*(2), 332–382. https://doi.org/10.1257/jel.38.2.332

Startz, R. (1989). Monetary policy and the consumption of durable goods. *Journal of Monetary Economics*, *23*(2), 201–222.

Suardi, S., Rasel, M. M., & Liu, X. (2022). Investor sentiment and stock market volatility: Evidence from emerging markets. *Journal of Behavioral Finance*, *23*(3), 225–241.

Tauchen, G. (1986). Finite state markov-chain approximations to univariate and vector autoregressions. *Economics Letters*, *20*(2), 177–181. https://doi.org/10.1016/0165-1765(86)90090-4

Tetlock, P. C. (2007). Giving content to investor sentiment: The role of media in the stock market. *Journal of Finance*, *62*(3), 1139–1168. https://doi.org/10.1111/j.1540-6261.2007.01232.x

Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and uncertainty*, *5*(4), 297–323.

Welch, I., & Goyal, A. (2008). A comprehensive look at the empirical performance of equity premium prediction. *The Review of Financial Studies*, *21*(4), 1455–1508. https://doi.org/10.1093/rfs/hhm014

White, H., Kim, K., & Manganelli, S. (2015). Multivariate quantile regression and asset interdependencies. *Journal of Multivariate Analysis*, *132*, 123–140.

Wu, G., & Gonzalez, R. (1996). Curvature of the probability weighting function. *Management Science*, *42*(12), 1676–1690. https://doi.org/10.1287/mnsc.42.12.1676

Xiong, W., & Yan, Z. (2010). Investor sentiment and the cross-section of stock returns. *Review of Financial Studies*, *23*(5), 2345–2370.

Yaari, M. E. (1987). The dual theory of choice under risk. *Econometrica: Journal of the Econometric Society*, 95–115.

Yarovaya, L., Matkovskyy, R., & Jalan, A. (2022). Do investors' attention and sentiment predict cryptocurrency returns? *Finance Research Letters*, *47*, 102711.

Yogo, M. (2003). Efficient tests of stock return predictability [Available at SSRN: https://ssrn.com/abstract=457548]. *NBER Working Paper No. 10026*. https://doi.org/10.2139/ssrn.457548

Yogo, M. (2006). A consumption-based explanation of expected stock returns. *The Journal of Finance*, *61*(2), 539–580.

Žikeš, F., & Baruník, J. (2016). Predicting return quantiles with realized measures. *Journal of Banking & Finance*, *69*, 235–248.