Mispricing, Momentum, and Market Timing - Essays on Stock Market Puzzles and Capital Structure Decisions

von der Wirtschaftswissenschaftlichen Fakultät der
Gottfried Wilhelm Leibniz Universität Hannover
zur Erlangung des akademischen Grades

Doktor der Wirtschaftswissenschaften
– Dr. rer. pol. –

genehmigte Dissertation
von

M.Sc. Jan Krupski
geboren am 6. November 1990 in Burgwedel

2023

Referent: Prof. Dr. Maik Dierkes
Korreferent: Prof. Dr. Marcel Prokopczuk
Tag der Promotion: 22.03.2023

Eigenständigkeitserklärung – Declaration of Original Authorship

Hiermit bestätige ich, dass ich die vorliegende Arbeit selbständig ohne Hilfe Dritter verfasst und keine anderen als die angegebenen Hilfsmittel benutzt habe. Die Stellen der Arbeit, die dem Wortlaut oder dem Sinn nach anderen Werken entnommen sind, wurden unter Angabe der Quelle kenntlich gemacht.

I hereby confirm that this is my own work and the use of all material from other sources has been properly and fully acknowledged.

Datum	Jan Krupski

Abstract

This doctoral thesis comprises one essay on the risk management of momentum strategies and three essays on the implications of skewness preferences on financial markets. Chapter 1 provides an extensive summary and links all projects within the framework of behavioral finance.

In Chapter 2 (co-authored with Maik Dierkes), we investigate momentum in stock returns and propose a novel approach to manage the downside risk of momentum strategies. Across markets, momentum is one of the most prominent anomalies and leads to high risk-adjusted returns. However, these returns come at the cost of substantial tail risk as there are short but persistent periods of highly negative returns. Momentum crashes occur in rebounding bear markets, when the momentum portfolio exhibits a negative beta and momentum volatility is high. Based on ex-ante estimates of these risk measures, we construct a crash indicator that effectively isolates momentum crashes. Subsequently, we propose an implementable trading strategy that combines both systematic and momentum-specific risk and more than doubles the Sharpe ratio of the original momentum strategy. Moreover, it outperforms existing risk management approaches over the 1928-2020 period, in sub-samples, and internationally.

In Chapter 3 (co-authored with Maik Dierkes and Sebastian Schroen), we address the effects of time-varying skewness preference, referred to as lottery demand, on first-day returns and the long-term performance of initial public offerings (IPOs). Following the identification approach of Dierkes (2013), we measure lottery demand in terms of option-implied probability weighting functions and find a significantly positive impact on first-day returns, tantamount to higher IPO underpricing and more money left on the table. Furthermore, disentangling the effects of lottery demand and cross-sectional expected skewness reveals that IPO returns are particularly driven by the interaction of market-wide lottery demand

and asset-specific lottery characteristics. In the long run, firms that went public during periods of high lottery demand perform poorly for up to five years after the IPO.

In Chapter 4 (co-authored with Maik Dierkes, Sebastian Schroen, and Philipp Sibbertsen), we perform a simulation-based approach to estimate volatility-dependent probability weighting functions and investigate the impact of probability weighting on the pricing kernel puzzle. We first obtain risk neutral and physical densities from the Pan (2002) stochastic volatility and jumps model and then estimate probability weighting functions according to the identification strategy presented in Chapter 3. Across volatilities, we find pronounced inverse S-shapes. Hence, small (large) probabilities are overweighted (underweighted), and probability weighting almost monotonically increases in volatility, suggesting higher skewness preferences in volatile markets. Moreover, by estimating probabilistic risk attitudes, equivalent to the share of risk aversion related to probability weighting, we shed further light on the pricing kernel puzzle. While pricing kernels estimated from the Pan (2002) model display the typical U-shape documented in the literature, adjusted pricing kernels are monotonically decreasing and thus in line with economic theory. As a result, risk aversion functions are positive throughout wealth levels.

Finally, in Chapter 5 (co-authored with Maik Dierkes), we employ idiosyncratic skewness as a proxy for firm-specific mispricing and investigate the impact of market timing on capital structure decisions. Consistent with the market timing theory, idiosyncratic skewness is significantly positively related to equity issues, while the impact on debt issues is negative and less important. Moreover, we find equity issues to be accompanied by debt retirement programs. Challenging the market timing theory, effects are not persistent and vanish after about three years. In line with Alti (2006), our results are therefore consistent with a modified version of the trade-off theory, including market timing as a short-term factor.

Keywords: Momentum, IPO, Skewness Preferences, Probability Weighting, Pricing Kernel Puzzle, Market Timing, Capital Structure

Zusammenfassung

Diese Dissertation umfasst einen Aufsatz zum Risikomanagement von Momentum-Strategien und drei Aufsätze über die Auswirkungen von Schiefepräferenzen auf Finanzmärkte. Kapitel 1 enthält eine ausführliche Zusammenfassung und ordnet die Forschungsprojekte in den Rahmen der verhaltensorientierten Finanztheorie ein.

In Kapitel 2 (gemeinsam mit Maik Dierkes verfasst) untersuchen wir das Momentum von Aktienrenditen und entwickeln einen neuartigen Ansatz zur Risikosteuerung von Momentum-Strategien. Die Momentum-Anomalie ist eine der bekanntesten Finanzmarkt-Anomalien und erzielt hohe risikobereinigte Renditen. Diese sind jedoch mit einem erheblichen Verlustrisiko verbunden, da sich wiederholt mehrmonatige Phasen stark negativer Renditen ereignen. Momentum-Crashes treten insbesondere in sich erholenden Bärenmärkten auf, wenn das Momentum-Portfolio zeitgleich ein negatives Beta und eine hohe Momentum-Volatilität aufweist. Auf Grundlage von ex-ante Schätzungen dieser Risikomaße konstruieren wir einen Crash-Indikator, welcher Momentum-Crashes erfolgreich isoliert. Infolgedessen stellen wir eine implementierbare Handelsstrategie vor, die systematisches und momentumspezifisches Risiko kombiniert und die Sharpe-Ratio der ursprünglichen Momentum-Strategie mehr als verdoppelt. Darüber hinaus übertrifft sie bestehende Risikomanagement-Strategien im Zeitraum von 1928-2020 sowie in Subperioden und im internationalen Kontext.

In Kapitel 3 (gemeinsam mit Maik Dierkes und Sebastian Schrön verfasst) untersuchen wir die Auswirkungen von zeitlich variierenden Schiefepräferenzen, im Folgenden als Lotterienachfrage bezeichnet, auf kurz- und langfristige Renditen nach Börsengängen (IPOs). Aufbauend auf der Identifikationsstrategie von Dierkes (2013) messen wir die Lotterienachfrage anhand optionsimplizierter Wahrscheinlichkeitsgewichtungsfunktionen und stellen einen signifikant positiven Einfluss auf die

Renditen am ersten Handelstag fest. Dieses Resultat ist gleichbedeutend mit einer stärkeren Unterbewertung der Emittenten (bezüglich des Eröffnungspreises) sowie höheren Opportunitätskosten. Darüber hinaus werden IPO-Renditen insbesondere durch die Interaktion von marktweiter Lotterienachfrage und firmenspezifischer Lotteriecharakteristika getrieben. Abschließend stellen wir fest, dass Unternehmen, deren Börsengang in Zeiten starker Lotterienachfrage erfolgt, über einen Zeitraum von bis zu fünf Jahren nach dem Börsengang schlechtere Renditen aufweisen.

In Kapitel 4 (gemeinsam mit Maik Dierkes, Sebastian Schrön und Philipp Sibbertsen verfasst) nutzen wir stattdessen einen simulationsbasierten Ansatz, um volatilitätsabhängige Wahrscheinlichkeitsgewichtungsfunktionen zu schätzen und deren Auswirkungen auf das Pricing Kernel Puzzle zu untersuchen. Zunächst elizitieren wir risikoneutrale und physische Dichtefunktionen auf Basis des stochastischen Volatilitäts- und Sprungmodells von Pan (2002) und schätzen damit Wahrscheinlichkeitsgewichtungsfunktionen gemäß der in Kapitel 3 vorgestellten Identifikationsstrategie. Über alle Volatilitätsniveaus hinweg weisen diese eine ausgeprägte inverse S-Form auf, gleichbedeutend mit der Übegewichtung (Untergewichtung) kleiner (großer) Wahrscheinlichkeiten. Bemerkenswerterweise nimmt die Wahrscheinlichkeitsgewichtung mit der Volatilität beinahe monoton zu, was auf ausgeprägtere Schiefepräferenzen in volatilen Märkten hinweist. Darüber hinaus schätzen wir die probabilistische Risikoeinstellung, also den Anteil der Risikoaversion, der durch Wahrscheinlichkeitsgewichtung hervorgerufen wird, und untersuchen damit das Pricing Kernel Puzzle. Während die mit Pan (2002) geschätzten Pricing Kernel, übereinstimmend mit der Literatur, U-förmig sind, weisen die um die probabilistische Risikoeinstellung bereinigten Kernel-Funktionen einen monoton fallenden Verlauf auf und stehen somit im Einklang mit der ökonomischen Theorie. Infolgedessen ist die Risikoaversion über alle Vermögensniveaus hinweg positiv.

Abschließend verwenden wir in Kapitel 5 (gemeinsam mit Maik Dierkes verfasst) die idiosynkratische Schiefe als einen Proxy für unternehmensspezifische Fehlbewertungen und untersuchen anhand dessen

die Auswirkungen von Market Timing auf Kapitalstrukturentscheidungen. Im Einklang mit der Market-Timing-Theorie hat die idiosynkratische Schiefe einen signifikant positiven Effekt auf die Emission von Aktien, während der Einfluss auf die Emission von Schuldtiteln negativ und von geringerer Bedeutung ist. Zudem stellen wir fest, dass Aktienemissionen in der Regel durch den Abbau von Schulden begleitet werden. Entgegen der Market-Timing-Theorie sind diese Effekte jedoch nicht von Dauer und verschwinden nach etwa drei Jahren. In Übereinstimmung mit Alti (2006) unterstützen unsere Ergebnisse daher eine modifizierte Version der Trade-Off-Theorie, welche Market Timing als kurzfristigen Faktor einbezieht.

Schlagwörter: Momentum, IPO, Schiefepräferenz, Wahrscheinlichkeitsgewichtung, Pricing Kernel Puzzle, Market Timing, Kapitalstruktur

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Motivation

To gain a better understanding of financial markets, neoclassical finance assumes that market participants act strictly rational and maximize utility. Consequently, asset prices should reflect fundamental values and markets are assumed to be in line with the efficient market hypothesis (Fama, 1970), which, in its strongest form, states that asset prices reflect all available information, both public and private. Hence, it is not possible to outperform the market on a risk-adjusted basis.

Although neoclassical models such as Sharpe (1964)'s capital asset pricing model (CAPM) are still widely accepted and taught as normative models, a growing strand of literature questions the existence of both rational market participants and market efficiency. The latter asserts that shares always trade at their fair value since mispricing is immediately corrected by rational arbitrageurs. However, De Long et al. (1990) show that prices can deviate considerably, even in the absence of fundamental risk. Hence, arbitrageurs may encounter limits to arbitrage if irrational noise traders drive up existing overvaluations. In fact, arbitrage is only conducted by a small number of professional traders who apply the capital of less sophisticated retail investors (Shleifer and Vishny, 1997). Since these are unaware of fair values and focus on past performance, arbitrageurs may have to liquidate positions without correcting the mispricing. As a result, predictable return patterns that contradict neoclassical finance – so-called anomalies – may persist over a long time period. In Chapter 2, we focus on one of the most prominent of these anomalies, the momentum anomaly, and propose a novel risk management strategy.

Contradicting another cornerstone of neoclassical finance – expected utility theory – a large and growing body of literature, both experimental and empirical, documents behavioral patterns that are hard to reconcile with rational decision making. For example, experimental evidence shows that losses are perceived more negatively than equivalent gains and small probabilities for tail events are overweighted. Moreover, decision makers tend to focus on changes in wealth rather than total wealth. Kahneman and Tversky (1979) and Tversky and Kahneman (1992) summarize these findings in their famous prospect theory, which is still considered to be one of the most influential descriptive models for decision making under risk. Notably, several studies, such as Kliger and Levy (2009), Barberis et al. (2016), and Baele et al. (2019), find probability weighting – and thus skewness preference – to be the model's key component. Consistent with this finding, Kraus and Litzenberger (1976) show that many empirical contradictions of the CAPM can be attributed to the omission of skewness as a risk factor. Hence, in Chapters 3 and 4 we elicit probability weighting functions from S&P 500 option prices and the Pan (2002) stochastic volatility and jumps model, respectively, and relate them to empirical puzzles (the IPO underpricing puzzle and the pricing kernel puzzle).

Finally, Barberis and Huang (2008) show that prospect theory investors overweight securities with highly right-skewed return distributions, causing them to be overvalued. In line with this prediction, Boyer et al. (2010) and Conrad et al. (2013) find a negative relation between idiosyncratic skewness and subsequent returns. In Chapter 5, we therefore employ idiosyncratic skewness as a proxy for firm-specific mispricing and investigate the impact of market timing on capital structure decisions.

1.2 Outline

Each chapter of this thesis provides an independent introduction and conclusion to the respective research question. The remainder of this chapter summarizes the contribution of each paper.

Chapter 2: Isolating Momentum Crashes (co-authored with Maik Dierkes) Jegadeesh and Titman (1993) show that past winners tend to outperform past losers in the near future. A subsequent zero-cost strategy that buys past winners and short sells losers therefore earns significant returns of 1.49% per month. Moreover, after controlling for Fama and French (1993) factors, risk-adjusted returns even increase to 1.69% per month, challenging neoclassical finance and, in particular, the CAPM. However, despite earning high average returns, the momentum strategy exhibits both a high kurtosis and a negative skewness, exposing it to significant tail risk. The two most prominent crashes took place in 1932 and 2009 and resulted in a draw-down of 91% and 73%, respectively. Notably, momentum crashes primarily occur in rebounding bear markets, when the momentum portfolio displays a negative beta and the volatility of momentum returns is high.

In order to control the time-varying exposure to momentum, the influential studies of Barroso and Santa-Clara (2015) and Daniel and Moskowitz (2016) both propose scaling approaches. However, while Barroso and Santa-Clara (2015) focus on momentum-specific risk, Daniel and Moskowitz (2016)'s baseline approach is restricted to systematic risk. We therefore propose a novel crash indicator strategy that accounts for both sources of risk. We first show that an ex-ante crash indicator, based on systematic risk, largely separates momentum crashes from momentum bull markets. Subsequently, we study the interaction between the crash indicator and momentum-specific risk and find that the explanatory power further improves.

Building on these insights, we propose an implementable trading strategy that scales time-varying exposure based on momentum-specific risk and reverses weights when the crash indicator predicts a crash. Empirically, we find the crash indicator strategy to outperform the approaches of Barroso and Santa-Clara (2015) and Daniel and Moskowitz (2016) both over the full sample period from 1926 to 2020 and in sub-samples. Our conclusions are robust to spanning tests, an international momentum portfolio, and the inclusion of transaction costs.

Chapter 3: Option-implied Lottery Demand and IPO Returns (co-authored with Maik Dierkes and Sebastian Schroen) IPOs not only earn anomalously high returns on their first day of trading, but also exhibit substantial variation over time. For example, the average return increased from 15% in 1990-1998 to 65% in 1999-2000, before falling back to 12% in 2001-2003 (Loughran and Ritter, 2004). In stark contrast, the performance in the long run is exceptionally poor (Ritter, 1991).

While earlier studies focus on rational explanations like information asymmetries (Beatty and Ritter, 1986) and a changing risk composition (Ritter, 1984), more recent studies suggest a behavioral perspective. For example, Loughran and Ritter (2002) propose a prospect theory (PT) approach, where firms evaluate outcomes based on a reference point and aggregate the loss from leaving money on the table and the increased valuation of retained shares. Usually, this results in a net profit and explains why firms only partially adjust offer prices to high demand during the book building period. In contrast, Barberis and Huang (2008) focus on PT's probability weighting component and find that investors, on average, overweight small probabilities for large gains, resulting in a preference for lottery-like stocks. As a consequence, more pronounced probability weighting should lead to higher first-day returns and a poor performance in the long run. By estimating the expected skewness of the IPO's industry, Green and Hwang (2012) document first empirical evidence consistent with these predictions. Following the identification approach of Dierkes (2013), we provide a cleaner test by directly estimating a time series of PT's probability weighting parameter gamma from S&P 500 option prices.

During our sample period from 1996 to 2020, we find several episodes of increased probability weighting (i.e. small gammas), most importantly in the late 1990s (during the DotCom bubble) and in the most recent past (2018-2020). We therefore employ gamma as an inverse predictor of market-wide lottery demand.

In line with the predictions of Barberis and Huang (2008), we find that IPOs issued in periods of high lottery demand earn higher first-day returns and are more likely to perform poorly over return horizons of up to five years. Moreover, we show that the explanatory power of expected skewness strongly depends on the respective lottery demand regime. IPO returns are thus particularly driven by the interaction of market-wide lottery demand and asset-specific lottery characteristics. Finally, we find that most of the market reaction takes place in the secondary market, driven by buying pressure from retail investors.

Chapter 4: Volatility-Dependent Probability Weighting and the Dynamics of the Pricing Kernel Puzzle (co-authored with Maik Dierkes, Sebastian Schroen, and Philipp Sibbertsen) Jackwerth (2000) defines risk neutral probabilities as the product of physical probabilities and a risk aversion adjustment. Accordingly, the pricing kernel, defined as the ratio of risk neutral to physical probabilities, is expected to monotonically decrease in wealth and distinctly reflects risk aversion. In contrast to this prediction, however, several studies find U-shaped pricing kernels, implying episodes of negative risk aversion.

We attribute this finding to irrational investors who overweight small probabilities for tail events and thus distort the pricing kernel. Importantly, recent studies find both the pricing kernel (Benzoni et al., 2011; Babaoğlu et al., 2018) and probability weighting functions (Kilka and Weber, 2001; Polkovnichenko and Zhao, 2013) to depend on the level of volatility (or uncertainty, respectively). We therefore refer time variation in pricing kernels and risk aversion to a volatility-dependent and thus time-varying degree of probability weighting. We first follow Ziegler (2007) and elicit risk neutral and physical density functions from the Pan (2002) stochastic volatility and jumps model for a wide range of volatilities. Subsequently, we adopt the identification strategy presented in Chapter 3 and estimate implied probability weighting functions for each of the obtained densities. However, in contrast to Chapter 3, we do not estimate a time series, but focus on the cross-section of probability weighting functions over volatilities, which enables us to counterfactually investigate the impact of volatility on the extent of probability weighting.

Although the Pan (2002) model was not designed to account for prospect theory preferences, our results are strikingly robust. Implied probability weighting functions are consistently inverse S-shaped and the curvature parameter gamma (probability weighting) almost monotonically decreases (increases) in volatility, implying more pronounced skewness preferences in volatile market environments. Furthermore, estimating probabilistic risk attitudes, equivalent to the share of risk aversion related to probability weighting, enables us to explore the pricing kernel puzzle. While pricing kernels estimated from the Pan (2002) model display the typical U-shape documented in the literature, adjusted pricing kernels are monotonically decreasing in wealth and thus in line with economic theory. As a direct result, risk aversion functions are positive throughout wealth levels. Our conclusions are robust to alternative return horizons, a nonparametric empirical setting, and several other robustness checks.

Chapter 5: Idiosyncratic Skewness and Market Timing of Capital Structure Decisions (co-authored with Maik Dierkes) There are three prevailing theories of capital structure: the pecking order theory, the trade-off theory, and the market timing theory. The pecking order theory predicts that firms primarily fund investments with internal funds. If these are not sufficient, they prefer debt over equity issues. According to the trade-off theory, firms choose a target leverage by balancing the costs and benefits of debt. In contrast, the market timing theory predicts that managers attempt to exploit temporary fluctuations in the cost of equity and therefore issue (repurchase) equity when shares are perceived to be overvalued (undervalued). According to Baker and Wurgler (2002), market timing should have a long-lasting impact on capital structure.

Again, our approach builds on the theoretical insights of Barberis and Huang (2008). Investors with prospect theory preferences demand securities with highly right-skewed payoffs, causing them to be overvalued. In line with this prediction, several empirical studies find a significantly negative relation between subsequent returns and both idiosyncratic and market-wide skewness (Boyer et al., 2010; Chang et al., 2013). Moreover, Green and Hwang (2012) find their measure of industry-specific skewness to be positively (negatively) related to first-day (long-term) IPO returns. We therefore employ a firm-specific version of their measure as a proxy for mispricing and investigate both the short-term impact and the persistence of (equity) market timing.

Our results provide further evidence for a strong market timing effect in the short run. Idiosyncratic skewness is significantly positively related to equity issues and negatively related to debt issues, with the former effect being the predominant one. Moreover, we find equity issues to be accompanied by debt retirement programs. However, in contrast to the predictions of Baker and Wurgler (2002), the market timing effect is not persistent and disappears after about three years. This key result is confirmed by both partial adjustment models and interaction effects with the firm-specific financing deficit and is robust to a wide range of robustness checks. Our findings are thus consistent with a long-run validity of the trade-off theory, including market timing as a short-term factor.

Chapter 2

Isolating Momentum Crashes

This chapter refers to the following publication:

Dierkes, Maik and Jan Krupski (2022): 'Isolating Momentum Crashes', Journal of Empirical Finance **66**: 1-22.

Available online at:
https://doi.org/10.1016/j.jempfin.2021.12.001

Abstract

Across markets, momentum is one of the most prominent anomalies and leads to high risk-adjusted returns. On the downside, momentum exhibits huge tail risk as there are short but persistent periods of highly negative returns. Crashes occur in rebounding bear markets, when momentum displays negative betas and momentum volatility is high. Based on ex-ante calculations of these risk measures, we construct a crash indicator that effectively isolates momentum crashes from momentum bull markets. An implementable trading strategy that combines both systematic and momentum-specific risk more than doubles the Sharpe ratio of original momentum and outperforms existing risk management strategies over the 1928-2020 period, in 5 and 10-year sub-samples, and an international momentum portfolio.

Keywords: Asset Pricing, Market Anomalies, Momentum, Crash Indicator

JEL: G11, G12.

2.1 Introduction

Jegadeesh and Titman (1993) show that past winners continue to outperform past losers in the near future. Their zero-cost strategy buys previous winners, short sells losers, and earns significant returns of 1.49% per month. Notably, controlling for Fama and French (1993) risk factors over the period from 1928 to 2020 results in even higher risk-adjusted returns of 1.69%.[1] These findings challenge neoclassical finance and, in particular, the capital asset pricing model (CAPM).[2] Furthermore, momentum is robust across industry portfolios, international markets, and asset classes.[3]

Average momentum returns are high, yet they display huge tail risk, i.e. a high kurtosis and negative skewness. Since 1926, there have been several momentum crashes that feature short but persistent periods of highly negative returns. For example, the momentum portfolio lost about 91% from June to August 1932, followed by a second draw-down in April to July 1933. Another prominent crash took place in 2009 when momentum lost more than 73% within a period of three months. Remarkably, crashes are driven by large gains of previous losers while winners still exhibit

[1] Controlling for Fama and French (2015) factors, risk-adjusted returns decrease to 1.38% but remain highly significant. Note that Fama and French (2015) factor data starts in 1963.

[2] The CAPM was independently proposed by Sharpe (1964), Lintner (1965) and Mossin (1966).

[3] See Moskowitz and Grinblatt (1999) for industry portfolios and Rouwenhorst (1998) and Rouwenhorst (1999) for international evidence in developed and emerging markets, respectively. Chan et al. (2000) confirm results for both markets. Notably, there is no significant momentum in the Japanese market, as shown by Asness (2011). See Okunev and White (2003) and Menkhoff et al. (2012) for currency markets as well as Erb and Harvey (2006) and Asness et al. (2013) for commodity futures and bond markets.

modestly positive returns.

Our study is most closely related to Barroso and Santa-Clara (2015) and Daniel and Moskowitz (2016) who employ scaling strategies that adjust the time-varying exposure to momentum. Barroso and Santa-Clara (2015) propose a risk management strategy solely based on momentum-specific risk. Thereby, exposure to momentum is scaled by the ratio of a pre-defined target volatility and the realized volatility of momentum returns. The strategy almost doubles the Sharpe ratio of original momentum and provides an intuitive way to adjust risk according to individual risk aversion. In contrast, Daniel and Moskowitz (2016) focus on systematic risk and adjust exposure with respect to expected returns, the conditional variance, and a time-invariant scaling parameter. This dynamic approach (DYN) significantly increases momentum returns and outperforms the constant volatility strategy (CVOL) of Barroso and Santa-Clara (2015).

We introduce a novel crash indicator strategy (CI) that considers both systematic and momentum-specific risk and improves existing risk management approaches. Our contribution to the literature is threefold. First, we show that an ex-ante crash indicator based on systematic risk measures largely separates momentum crashes from momentum bull markets. In our sample from 1928 to 2020, average returns when a crash is indicated amount to -3.63%, while the mean in non-crash periods is 1.49%. Second, we perform predictive regressions with crash indicators motivated from the literature and momentum volatility to show that a combination of systematic and momentum-specific risk measures further improves explanatory power. Third, we propose an implementable trading strategy based on our crash indicator and a constant scaling approach. In momentum bull markets, the strategy corresponds to Barroso and Santa-Clara (2015), but replaces the volatility of momentum returns by the variance. In crash periods, the weight in momentum is reversed, i.e. the strategy invests in past losers and short sells winners.[4]

With a Sharpe ratio of 1.12, CI is superior to both CVOL (0.94) and DYN (1.03).[5] Moreover, both strategies are spanned by the CI strategy, while CI's alpha remains significant at the 1% level (even after adjusting for CVOL and DYN simultaneously).[6] By dividing the full sample into decades and 5-year rolling windows, we further show that CI consistently offers superior returns even in shorter time periods. DYN is outperformed in all decades and 90% of 5-year periods, while CVOL returns are exceeded in seven out of nine decades and 75% of rolling windows. In addition to that, one of the remaining decades offers almost equivalent returns. Most importantly, CI clearly outperforms both strategies in major crash periods. Remarkably, implementing CVOL and CI requires only six and 24 months, respectively, while DYN rests on an expanding window regression that gains power by applying a long sample of data. By re-estimating the strategies in each sub-sample, we show that DYN

[4] Our study is also related to Blitz et al. (2011) and Blitz et al. (2020) who rank stocks on residual returns – adjusted for Fama and French (1993) factors – instead of raw returns. However, we focus on the interaction of systematic (beta) risk and momentum-specific risk. Furthermore, our estimation of beta requires less data (six months) since the estimation of residual stock returns relies on regressions that include the previous 36 months.

[5] As break-even round-trip costs show, CI remains superior even after including transaction costs.

[6] Our results are robust to a 2×3 style momentum portfolio that is formed on previous performance and size.

performs substantially worse when applied in markets that have not yet experienced a momentum crash. More precisely, the re-estimated strategy is clearly inferior in all decades and 99% of 5-year periods. In six out of nine decades, DYN is even outperformed by original momentum.

Moreover, by *ex-post* scaling risk-adjusted returns to have an annualized volatility of 19%, Daniel and Moskowitz (2016) circumvent a highly relevant problem in implementing their strategy: there is no hint on how an investor could intuitively adjust risk exposure *ex-ante*. Therefore, we perform a sensitivity analysis and show that downside risk strongly depends on the choice of the ex-ante unknown scaling parameter. In a worst case scenario, even a reasonable calibration can involve a loss of the full investment, rendering risk management ineffective.

Finally, we present international evidence by estimating an international momentum portfolio consisting of the most important markets outside of the United States: France, Germany, Japan, and the United Kingdom.[7] Again, the Sharpe ratio of CI (1.77) clearly exceeds both CVOL (1.48) and DYN (1.47). To prove that these results are not driven by the specific choice of countries, we re-estimate all strategies for a Global-Ex-USA and regional portfolios and still find superior performance.

[7] See Barroso and Santa-Clara (2015)'s Table 5.

2.2 Momentum in US Equity Markets

2.2.1 Data and Portfolio Construction

We determine momentum returns based on daily and monthly return-sorted decile portfolios that are kindly provided by Kenneth French.[8] We classify the 10% best (worst) performing stocks as winners (losers) and rebalance portfolios on a monthly basis. Monthly (daily) data cover the period from January 1927 (October 1926) to May 2020. Supplementary data on the Fama and French (1993) three factor model, the risk-free rate, and 2×3 portfolios formed on size and momentum are provided by Kenneth French as well. Furthermore, we employ daily and monthly country-specific momentum returns (provided by AQR Capital Management) to construct an international momentum portfolio.[9] Finally, to perform robustness checks, we deploy a Global-Ex-USA momentum portfolio and several regional portfolios (Europe, North America, and Pacific) that are also provided by AQR. We cover the entire period of available data (January 1987 to May 2020).

2.2.2 Momentum Crashes

Fig. 2.1 illustrates momentum crashes by plotting cumulative momentum returns over the full sample (Panel A) and the two most important crash decades: the 1930s (Panel B) and 2000s (Panel C). Although a $1

[8] For details, see https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html.

[9] See https://www.aqr.com/Insights/Datasets. For further information on the construction of country-specific momentum portfolios, see Asness and Frazzini (2013) and Asness et al. (2019).

Fig. 2.1: Cumulative Momentum and the Market

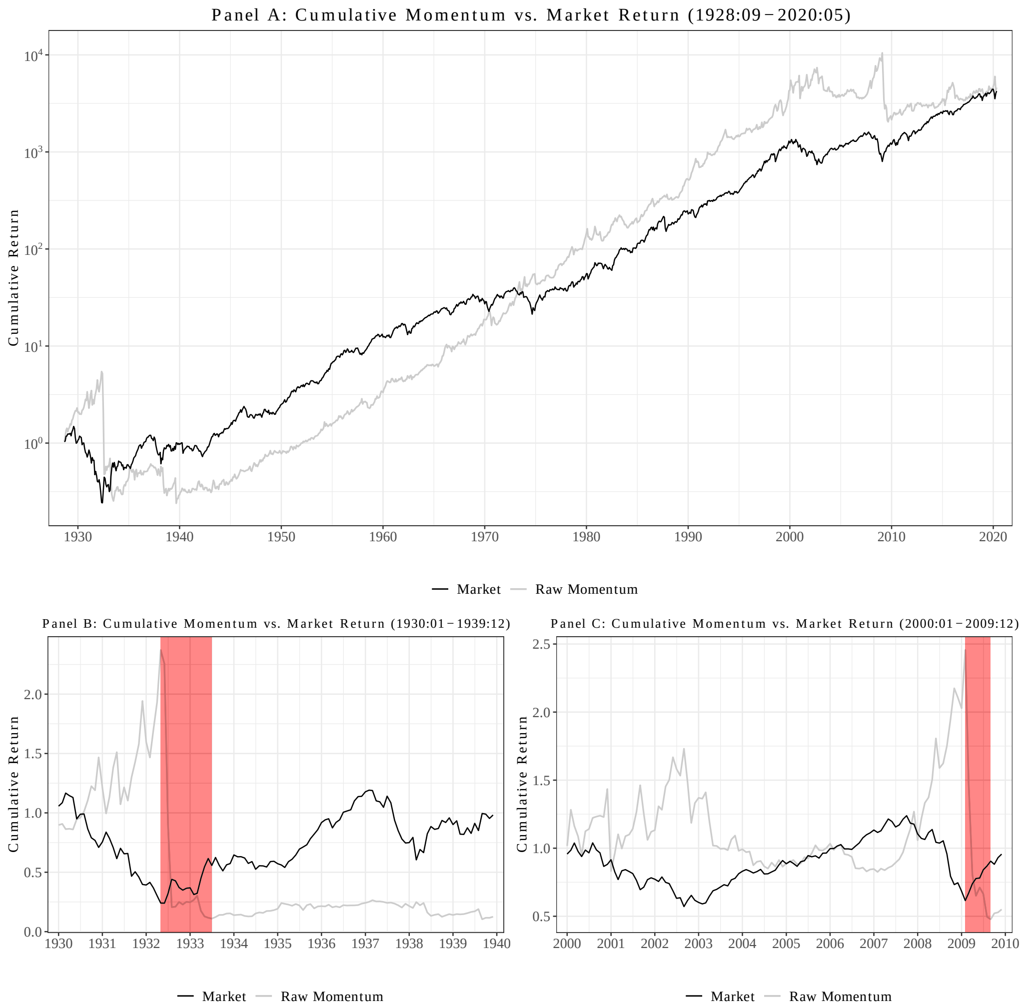

Fig. 2.1 plots cumulative momentum and markets returns over the full sample (Panel A) and the two most prominent crash periods: The 1930s (Panel B) and 2000s (Panel C). The y-axis of Panel A is logarithmized to improve visibility. Highlighted areas in Panel B and C mark crash periods.

investment in September 1928 would have led to $4,607 in May 2020, several crashes occurred that took decades to recover from. Due to short but persistent crashes in June to August 1932 and April to July 1933, the momentum strategy occasionally lost more than 95% of its initial value. Another large crash from March to May 2009 similarly involved a loss of more than 73%. In addition to that, there have been several smaller crashes in 1938/1939, 1974/1975, and 2001/2002. While being smaller in size, each involved at least one monthly loss of more than 19%. Therefore, high monthly returns of 1.15% come with a large kurtosis of 16.6 and a highly negative skewness of −2.3.[10] Barroso and Santa-Clara (2015) note that even the high average returns of the momentum strategy do not compensate investors for taking the risk of suchlike momentum crashes.

2.3 Predicting Momentum Crashes

2.3.1 Time-varying Risk of Momentum

Cooper et al. (2004) find a positive correlation between momentum returns and the state of the market, where up-markets (down-markets) are defined by positive (negative) 3-year market returns. More precisely, average momentum returns following up-markets (0.93%) are significantly higher than following down-markets (−0.37%). Daniel and Moskowitz (2016) confirm this finding for market states based on 2-year returns and show that crash periods display positive 1-month returns. Consequently, they define this turning point as 'Market Rebound'.[11] Panels B and C of Fig. 2.1 illustrate this finding as both show negative 2-year and positive 1-month market returns during momentum crashes (highlighted area). Crashes occur exactly when the market starts to rebound.

[10] Note that monthly returns reported by Jegadeesh and Titman (1993) are higher since the momentum crash of 2009 is not included in their sample.

[11] Asem and Tian (2010) have been the first to show that momentum returns particularly depend on market dynamics. They find returns to be higher when markets remain in their current state.

Grundy and Martin (2001) find a time-varying market beta of momentum returns.[12] This result is intuitive as in bull markets winners (losers) tend to be high-beta (low-beta) stocks, while in bear markets winners (losers) are those who co-vary the least (most). Based on this finding, Grundy and Martin (2001) propose a dynamically hedged portfolio that adjusts momentum returns for market and size risk. However, since momentum returns are regressed on market and size returns in month t to month $t+5$, betas have a look-ahead bias and the strategy is not tradable. Moreover, Daniel and Moskowitz (2016) show that ex-ante hedging does not improve performance.[13] To illustrate the interaction of momentum and its market beta, Panel A of Fig. 2.2 displays cumulative momentum returns and betas for the period from January 2000 to May 2020.[14] While, in fact, beta is negative in crash periods (e.g. in 2009), it is also negative prior to crashes when momentum exhibits exceptionally large returns (e.g. in 2008). Thus, despite comprising some information, beta alone does not avoid momentum crashes.

As of yet, we only considered market risk (i.e. systematic risk) of momentum. Another important source of risk is presented by Stivers and Sun (2010) and Barroso and Santa-Clara (2015) who investigate the impact of momentum-specific risk. Stivers and Sun (2010) find a negative relationship between the cross-sectional dispersion in stock returns and subsequent momentum returns, net of several macroeconomic variables.

[12] Time-variation in betas of return-sorted portfolios was first shown by Kothari and Shanken (1992).

[13] See also Barroso (2014).

[14] Appendix 2.A.1 outlines the calculation of momentum betas.

Fig. 2.2: Momentum and Beta

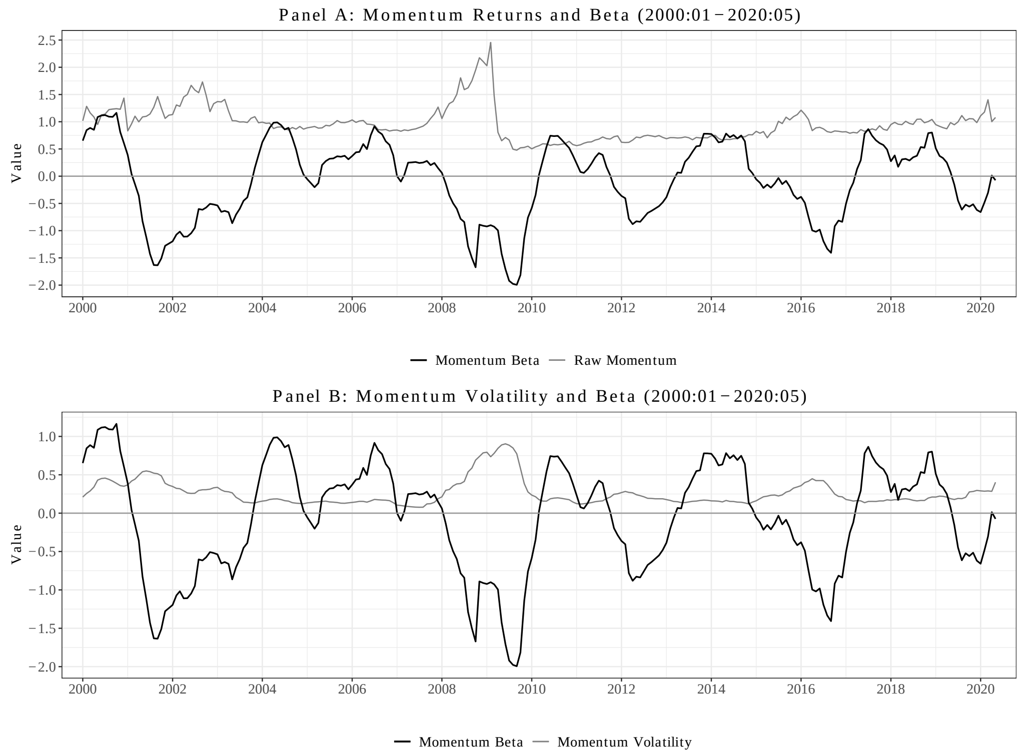

Fig. 2.2 plots cumulative momentum returns (Panel A) and momentum volatility (Panel B) compared to momentum beta. Both panels include monthly data from January 2000 to May 2020. At the beginning of each month, beta is estimated by a simple regression of the 126 preceding daily momentum returns on the CAPM. Momentum volatility is calculated as the realized volatility of the 126 daily momentum returns preceding the start of the current month.

Barroso and Santa-Clara (2015) find momentum risk to be predictable by its own realized variance and document a negative relationship between momentum volatility and subsequent returns.

Fig. 2.3 therefore compares cumulative momentum returns to momentum volatility.[15] Panels A and B display the 1930s and 2000s crash decades, respectively, where highlighted areas denote the momentum crashes of 1932/1933 and 2009. Both panels confirm Barroso and Santa-

[15] Momentum volatility is calculated according to Appendix 2.A.2.

Fig. 2.3: Risk and Return of Momentum

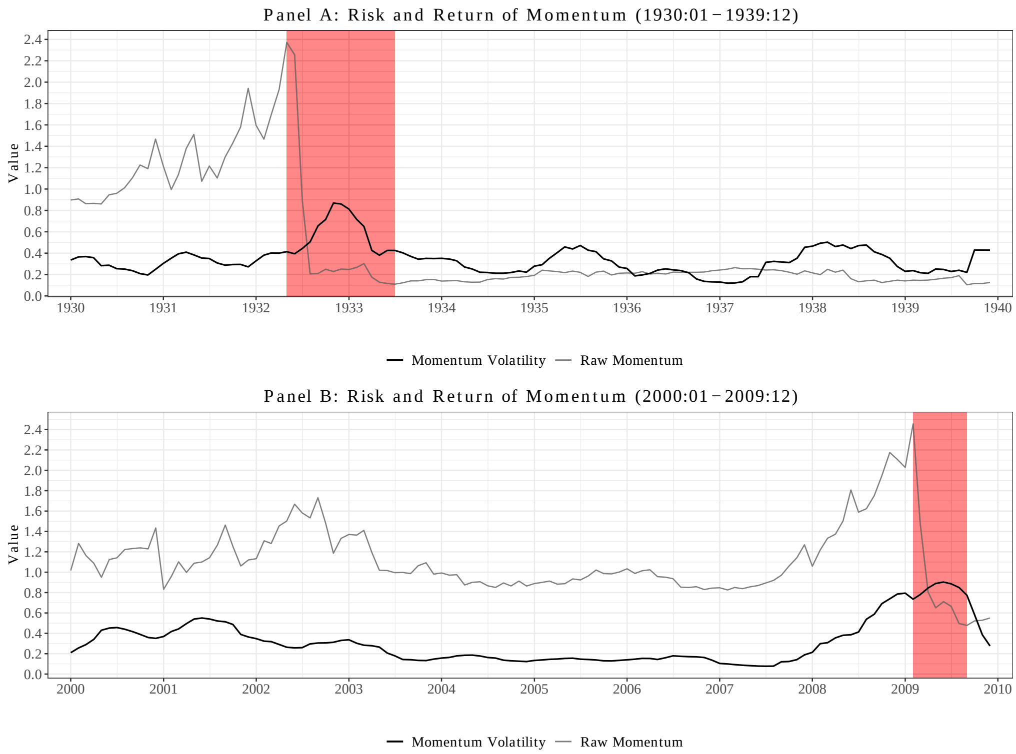

Fig. 2.3 plots cumulative momentum returns and annualized momentum volatility in the 1930s (Panel A) and 2000s (Panel B). Highlighted areas mark crash periods. Momentum volatility is calculated as the realized volatility of the 126 daily momentum returns preceding the start of the current month.

Clara (2015)'s result that momentum volatility increases during momentum crashes. Moreover, Panel B of Fig. 2.2 shows that high-volatility periods feature negative betas. This result is intuitive as periods of high momentum volatility are also periods of market distress, which in turn implies a low (high) beta of previous winners (losers).

Table 2.1 summarizes the time-varying risk of momentum by collecting the 15 worst momentum returns of our sample period and the corresponding risk measures. In each of these months, momentum incurred a loss of at least 19.7%. Notably, we find the worst months to be

Table 2.1: Worst Momentum Returns and Corresponding Risk Measures

Rank	Date	Momentum	Systematic Risk			Specific Risk
			$Market_{1M}$	$Market_{2Y}$	β_{MOM}	σ_{MOM}
1	1932 – 08	−77.0%	37.6%	−67.6%	−0.84	0.51
2	1932 – 07	−60.2%	33.9%	−74.8%	−0.79	0.44
3	2009 – 04	−45.6%	10.2%	−40.6%	−1.00	0.84
4	1939 – 09	−45.2%	16.9%	−21.6%	−0.08	0.22
5	2001 – 01	−42.0%	3.9%	10.7%	0.39	0.37
6	1933 – 04	−41.9%	39.0%	−59.0%	−0.11	0.43
7	2009 – 03	−39.8%	9.0%	−44.9%	−0.93	0.78
8	1938 – 06	−33.2%	24.0%	−27.7%	−1.29	0.44
9	1931 – 06	−29.0%	14.2%	−47.6%	−1.05	0.35
10	2020 – 04	−28.7%	13.6%	−0.8%	0.01	0.28
11	1933 – 05	−26.9%	21.6%	−36.7%	−0.13	0.38
12	2009 – 08	−25.4%	3.4%	−27.2%	−1.98	0.85
13	2002 – 11	−20.1%	6.0%	−36.2%	−0.51	0.31
14	2016 – 04	−19.8%	0.9%	10.9%	−1.00	0.45
15	1975 – 01	−19.7%	14.0%	−41.8%	−0.40	0.17

Table 2.1 presents the 15 worst monthly momentum returns over the period from September 1928 to May 2020 as well as corresponding risk measures. $Market_{2Y}$ and $Market_{1M}$ are the 2-year and contemporaneous 1-month market return. σ_{MOM} and β_{MOM} are momentum volatility and beta. At the beginning of each month, beta is estimated by a simple regression of the 126 preceding daily momentum returns on the CAPM. Momentum volatility is calculated as the realized volatility of the 126 daily momentum returns preceding the start of the current month.

clustered as 12 out of 15 either took place in the 1930s or the 2000s and the two worst months are subsequent losses from July to August 1932. The maximum loss occurred in August 1932 when momentum lost 77% of its initial value. In *all* months, the contemporaneous market return is positive and all but two months display a negative 2-year market return. Similarly, 13 out of 15 months exhibit a negative beta. With respect to momentum-specific risk, only one month shows an annualized volatility of slightly below the average (0.18), while still exceeding the median annualized volatility of 0.14.

2.3.2 Isolation of Crash Periods

Based on the results presented in Section 2.3.1, we propose three feasible crash indicators.

1. A bear market indicator, $I_{B,t-1}$, based on Daniel and Moskowitz (2016), which equals one if the 2-year market return preceding the start of month t ($Market_{2Y,t-1}$) is negative and zero otherwise.

$$I_{B,t-1} = \begin{cases} 1 & \text{if } Market_{2Y,t-1} < 0, \\ 0 & \text{otherwise} \end{cases} \tag{2.1}$$

2. As Daniel and Moskowitz (2016) find momentum crashes to occur during market rebounds, we further motivate a rebound indicator, $I_{R,\tau}$, which equals zero unless bear markets display positive 1-month returns ($Market_{1M,\tau}$). Depending on the view (ex-post or ex-ante), $Market_{1M,\tau}$ is either the contemporaneous ($\tau = t$) or lagged 1-month return ($\tau = t - 1$).

$$I_{R,\tau} = \begin{cases} 1 & \text{if } Market_{2Y,t-1} < 0 \ \& \ Market_{1M,\tau} > 0, \\ 0 & \text{otherwise} \end{cases} \tag{2.2}$$

3. Although momentum beta alone does not avoid momentum crashes (Daniel and Moskowitz, 2016), a combination with $I_{R,\tau}$ seems reasonable. Therefore, we propose a crash indicator, $I_{C,\tau}$, which is one

if a rebound takes place and momentum beta is negative.

$$I_{C,\tau} = \begin{cases} 1 & \text{if } Market_{2Y,t-1} < 0 \ \& \ Market_{1M,\tau} > 0 \ \& \ \beta_{MOM,t-1} < 0, \\ 0 & \text{otherwise} \end{cases} \quad (2.3)$$

Table 2.2 presents average momentum returns with respect to the state of the indicators. Most importantly, each indicator features significantly lower returns in periods defined as a crash. Considering ex-ante measures, $I_{C,t-1}$ displays the highest and most significant absolute difference in means (−5.12%). In months depicted as 'no crash', average returns increase from the full sample mean of 1.15% to 1.49%, while dropping to −3.63% in crash periods. For the rebound indicator, $I_{R,t-1}$, the absolute difference in means decreases (−3.84%) and significance reduces to the 5%-level. While average returns in non-crash periods are similar to those of $I_{C,t-1}$ (1.47%), returns in crash periods increase to −2.38%. For the bear market indicator, $I_{B,t-1}$, the difference remains significant at the 5%-level, although the absolute spread shrinks to 2.63%. The crash period average increases to roughly −1.04%, whereas the non-crash mean amounts to 1.59%.[16] Considering ex-post indicators, differences and absolute t-values strongly increase. However, $I_{C,t}$ still displays the highest and most significant difference (10.40%).

While these results suggest that I_C is a meaningful crash indicator, it seems reasonable to ask how an investor could have known this in 1930. Note that bear markets are often initiated by a market crash. In these

[16] The latter result is driven by the fact that $I_{R,t-1}$ and $I_{C,t-1}$ occasionally miss the first month of momentum crashes, whereas $I_{B,t-1}$ covers the whole period of market stress. As bear markets also include positive returns prior to momentum crashes, the crash period mean is increased, while the non-crash average enlarges as well.

Table 2.2: Comparison of Mean Returns

Indicator (I_j)	$I_j = 0$	$I_j = 1$	Diff.	t-value	Implementation
$I_{B,t-1}$	1.59%	−1.04%	−2.63%	−2.49**	ex-ante
$I_{R,t-1}$	1.47%	−2.38%	−3.84%	−2.53**	ex-ante
$I_{C,t-1}$	1.49%	−3.63%	−5.12%	−2.83***	ex-ante
$I_{R,t}$	1.95%	−6.66%	−8.61%	−5.65***	ex-post
$I_{C,t}$	1.94%	−8.47%	−10.40%	−5.94***	ex-post

Table 2.2 presents average momentum returns with respect to indicator j, where B denotes the bear-market indicator, R is the rebound indicator and C is the crash indicator. To calculate means, we apply the full sample period from September 1928 to May 2020. Ex-ante (ex-post) implementation indicates that lagged (contemporaneous) market returns are used. Stars indicate significance at the 10% (*), 5% (**) and 1% (***) level.

particular periods, momentum winners (losers) are those stocks with the smallest (largest) losses. In line with Shleifer and Vishny (1997), losers are more likely to be undervalued than winners and their expected returns increase.[17] By construction, losers also exhibit a higher beta than winners, resulting in a negative beta of the momentum portfolio.[18] When the market starts to recover, contemporaneous returns are positive, but the overall market condition is still considered to be a bear market.[19] Finally, as losers move back to their fair value, they display higher returns and the return of the momentum portfolio is negative.

To investigate whether a combination of the crash indicators and momentum-specific risk further improves momentum predictability, we employ predictive regressions of monthly momentum returns on the

[17] According to Shleifer and Vishny (1997), professional traders (arbitrageurs) apply the capital of less sophisticated (potentially irrational) retail investors who are unaware of fair values and focus on past performance. Therefore, larger losses in the recent past likely entail more withdrawals, resulting in an undervaluation of loser stocks.

[18] See Barroso and Santa-Clara (2015) and Daniel and Moskowitz (2016).

[19] See Daniel and Moskowitz (2016).

interaction of crash indicators and annualized momentum volatility. The regression framework is set up as follows

$$r_{MOM,t} = \alpha + \gamma \cdot I_{j,t-1} + \delta \cdot I_{j,t-1} \cdot \sigma_{MOM,t-1} + \eta \cdot \sigma_{MOM,t-1} + \lambda \cdot \vec{X}_{t-1} + \epsilon_t,$$

where $r_{MOM,t}$ and $\sigma_{MOM,t-1}$ denote monthly momentum returns and the annualized momentum volatility of the preceding 126 daily momentum returns, respectively. Ex-ante crash indicators are depicted by $I_{j,t-1}$. In addition to that, $\vec{X}_{t-1}$ and ϵ_t denote a vector of lagged Fama and French (1993) risk factors and the monthly residuals.

Table 2.3 presents results for the full sample period from September 1928 to May 2020. The first three models simply regress momentum returns on indicator dummies. By construction, coefficients correspond to the difference in means presented in Table 2.2 and t-statistics are lowest for $I_{C,t-1}$ (well below -5). In Model (4), momentum returns are solely regressed on the preceding momentum volatility. Consistent with Stivers and Sun (2010) and Barroso and Santa-Clara (2015), we find that momentum volatility has a significantly negative impact on subsequent returns. Models (5) to (7) present results for the interaction of crash indicators and momentum volatility, where we obtain four important results. First, significance increases for each of the indicators. Second, in absolute terms, coefficients, t-values, and R^2's exceed those of Models (1) to (4), suggesting that a combination of market risk and momentum-specific risk improves crash predictability. Third, $I_{C,t-1}$ remains the most significant indicator, both statistically and economically. Fourth, as explained variability ranges from 2.2% ($I_{B,t-1}$) to 3.4% ($I_{C,t-1}$), R^2's can be considered

Table 2.3: Predictive Regressions

	Dependent variable:								
	$r_{MOM,t}$								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
$I_{B,t-1}$	−0.026***								
	(−4.18)								
$I_{R,t-1}$		−0.038***							
		(−4.47)							
$I_{C,t-1}$			−0.051***						
			(−5.43)						
$\sigma_{MOM,t-1}$				−0.078***					
				(−4.14)					
$I_{B,t-1} \cdot \sigma_{MOM,t-1}$					−0.083***				
					(−5.06)				
$I_{R,t-1} \cdot \sigma_{MOM,t-1}$						−0.123***		−0.103***	
						(−5.63)		(−4.47)	
$I_{C,t-1} \cdot \sigma_{MOM,t-1}$							−0.144***		−0.125***
							(−6.30)		(−5.23)
FF3 included?	No	No	No	No	No	No	No	Yes	Yes
Adj. R^2	0.016	0.018	0.026	0.014	0.022	0.027	0.034	0.040	0.046

Table 2.3 presents OLS-regressions of monthly momentum returns on the interaction of several combinations of lagged crash indicators and lagged annualized momentum volatility over the full sample period from September 1928 to May 2020. In addition to that, model (8) and (9) include Fama and French (1993) risk-factors. The regression framework is set up as follows:

$$r_{MOM,t} = \alpha + \gamma \cdot I_{j,t-1} + \delta \cdot I_{j,t-1} \cdot \sigma_{MOM,t-1} + \eta \cdot \sigma_{MOM,t-1} + \lambda \cdot \vec{X}_{t-1} + \epsilon_t$$

$r_{MOM,t}$ and $\sigma_{MOM,t-1}$ denote monthly momentum returns and the annualized momentum volatility of the preceding 126 daily momentum returns, respectively. Ex-ante crash indicators j are depicted by $I_{j,t-1}$. In addition to that, $\vec{X}_{t-1}$ and ϵ_t state a vector of lagged Fama and French (1993) risk-factors and the monthly residuals. Stars indicate significance at the 10% (*), 5% (**) and 1% (***) level, t-values are stated in parentheses.

high.[20] To examine whether these findings are robust to common risk factors, the last two models add Fama and French (1993) factors. For both indicators, $I_{R,t-1}$ and $I_{C,t-1}$, coefficients and t-values only slightly decrease (in absolute terms) and $I_{C,t-1}$ continues to exhibit the highest predictive

[20] See Campbell and Thompson (2008). Even in-sample regressions do not exceed $R^2 = 1.35\%$ (their Table 2).

power. We thus find $I_{C,t-1}$ to effectively isolate momentum crashes from momentum bull markets.

To illustrate why beta adds information, Panel A of Fig. 2.4 plots cumulative momentum returns in contrast to 2-year market returns. Highlighted areas of both panels begin when momentum starts to recover and end when the 2-year market return (or beta) recognizes that the momentum crash is over (i.e. changes sign). A closer look at the years 2008 to 2010 reveals highly positive returns prior to momentum crashes. According to the bear-market indicator (which is solely based on 2-year market returns), the crash period would not only include the crash but also a large proportion of preceding positive returns. Moreover, the 2-year market return remains negative until September 2010, whereas momentum already starts to recover in October 2009. In contrast, the rebound indicator ($I_{R,t-1}$) additionally incorporates 1-month market returns and is thus capable of separating crashes from the preceding rise. However, $I_{R,t-1}$ still fails to isolate momentum crashes from the beginning recovery (as 2-year-market returns remain negative until late 2010). In Panel B, the market return is replaced by momentum beta, which recognizes the recovery in March 2010, six months earlier than the 2-year market return. The reasoning is simple. In contrast to the 2-year market return, beta estimates are based on the previous six months and therefore show a faster response to market changes.

To examine the information content of beta in more detail, we perform a simple regression of momentum volatility on the contemporaneous and

Fig. 2.4: Isolation of Crash Periods

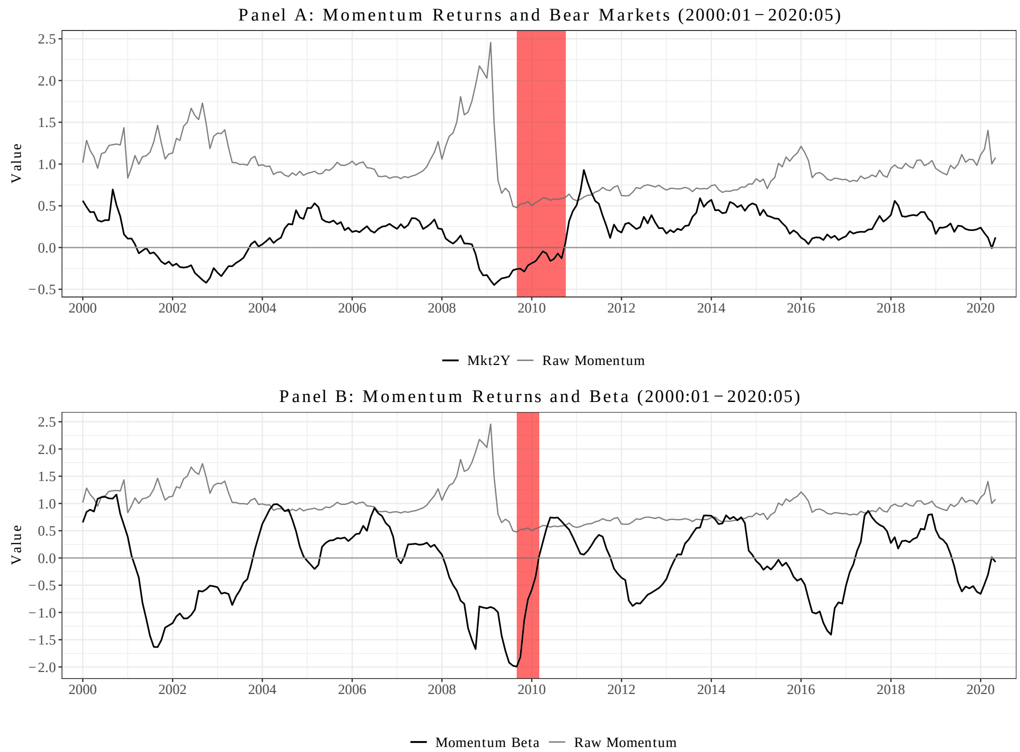

Fig. 2.4 opposes cumulative momentum returns over the period from January 2000 to May 2020 to 2-year market returns (Panel A) and momentum beta (Panel B). Highlighted areas mark the period from the beginning of momentum recovery to the point when 2-year market returns and beta change signs. At the beginning of each month, beta is estimated by a simple regression of the 126 preceding daily momentum returns on the CAPM.

lagged beta, respectively

$$\sigma_{MOM,\tau} = \alpha + \gamma \cdot \beta_{MOM,t} + \epsilon_t,$$

where τ is either t or $t+1$. Table 2.4 shows that both the contemporaneous and the lagged beta add information, suggesting that beta predicts future momentum volatility. More precisely, a decrease of beta by one unit increases contemporaneous (future) momentum volatility by economically and statistically significant 7.6 (7.2) percentage points.

Table 2.4: The Information Content of Beta

Dependent Variable	Independent Variables		
	α	$\beta_{MOM,t}$	R^2
$\sigma_{MOM,t}$	0.190	−0.076***	0.157
(t-value)	(54.14)	(−14.34)	
$\sigma_{MOM,t+1}$	0.190	−0.072***	0.143
(t-value)	(53.54)	(−13.57)	

Table 2.4 presents results for OLS-regressions of momentum volatility on the contemporaneous and lagged momentum beta over the full sample period from September 1928 to May 2020, respectively:

$$\sigma_{MOM,\tau} = \alpha + \gamma \cdot \beta_{MOM,t} + \epsilon_t,$$

where τ is either t or $t+1$. At the beginning of each month, beta is estimated by a simple rolling regression of the 126 preceding daily momentum returns on the standard market model. Momentum volatility is calculated as the realized volatility of the 126 daily momentum returns preceding the start of the current month. Stars indicate significance at the 10% (*), 5% (**) and 1% (***) level, t-values are stated in parentheses.

2.4 Risk-Managed Momentum

2.4.1 Risk Management Strategies

There is a growing literature on managing momentum crashes. Two of the most promising strategies are presented by Barroso and Santa-Clara (2015) and Daniel and Moskowitz (2016) who apply scaling strategies that take into account the current level of risk.

Barroso and Santa-Clara (2015) propose a constant volatility scaling strategy (CVOL) based solely on momentum-specific risk. Exposure to momentum (w_{t-1}^{CVOL}) is scaled by the ratio of a pre-defined target volatility (σ_{target}) and the realized volatility of daily momentum returns over the preceding 126 trading days ($\sigma_{MOM,t-1}$)

$$r_{MOM_t^{CVOL}} = \underbrace{\frac{\sigma_{target}}{\hat{\sigma}_{MOM,t-1}}}_{\widehat{=} w_{t-1}^{CVOL}} r_{MOM,t}. \tag{2.4}$$

Barroso and Santa-Clara (2015) propose a target volatility of 12%. Although the actual volatility over our sample period exceeds this target (16.5%), constant volatility scaling increases momentum returns and simultaneously reduces volatility. Moreover, including a constant target supports investors in adjusting risk exposure according to their individual risk preferences.

In contrast, the baseline strategy of Daniel and Moskowitz (2016) focuses on market risk and scales momentum exposure dynamically (DYN). More precisely, the optimal weight in their momentum strategy (w_{t-1}^{DYN}) maximizes the Sharpe ratio of an intertemporal version of Markowitz' (1952) portfolio optimization

$$r_{MOM_t^{DYN}} = \underbrace{(\frac{1}{2\lambda})\frac{\hat{\mu}_{t-1}}{\hat{\sigma}_{t-1}^2}}_{\widehat{=} w_{t-1}^{DYN}} r_{MOM,t}. \tag{2.5}$$

While expected returns ($\hat{\mu}_{t-1}$) are estimated by a full sample regression of monthly momentum returns on the interaction of the bear market indicator and lagged market variance, expected momentum variance ($\hat{\sigma}_{t-1}^2$) is calculated by a linear combination of a GJR-GARCH model and realized momentum volatility. The scaling parameter λ is a pre-defined constant that controls the unconditional risk and return of the dynamic portfolio. Notably, there is no intuitive choice of λ, potentially entailing high risk exposure. Importantly, as $\hat{\mu}_{t-1}$ is estimated over the full sample, the strategy is not implementable. To overcome this issue, Daniel and Moskowitz (2016) propose to replace full sample regressions by expanding window regressions.[21] Furthermore, *ex-ante* expected variance is based on the realized variance of momentum returns over the preceding 126 days. In their sample period from 1934 to 2013, DYN outperforms both original momentum and CVOL (in terms of Sharpe ratios).

We propose a novel crash indicator strategy (CI) that combines both market and momentum-specific risk

$$r_{MOM_t^{CI}} = \underbrace{\frac{\sigma^2_{target}}{\hat{\sigma}^2_{t-1}} \cdot (-1)^{I_{C,t-1}}}_{\widehat{=} w^{CI}_{t-1}} r_{MOM,t}. \tag{2.6}$$

First, the strategy comprises momentum-specific risk since baseline weights are determined by a scaling approach that is similar to Barroso and Santa-Clara (2015). However, our scaling approach slightly differs as returns are scaled by the momentum variance. We motivate this approach using the theoretical results of Daniel and Moskowitz (2016) who find the optimal weight to be determined by the variance.[22] The actual volatility resulting from a 12% target volatility (i.e. $\sigma^2_{target} = 0.12^2 = 0.0144$) is roughly 17.2% and thus only slightly exceeds constant volatility scaling (16.5%). Therefore, it is still possible to adjust risk exposure with respect to individual risk preferences. Second, by including the (ex-ante)

[21] We recap their ex-ante estimations in Appendix 2.A.3. For ex-post estimations see Daniel and Moskowitz (2016) and Glosten et al. (1993).

[22] This finding is in line with general portfolio theory (Campbell and Viceira, 2002). In an earlier version of this study, the crash indicator strategy was based on volatility scaling. However, conclusions with respect to our main results have been the same.

crash indicator, $I_{C,t-1}$, we also take into account systematic sources of risk, leading to enhanced momentum scaling. More precisely, when the indicator predicts a crash ($I_{C,t-1} = 1$), exposure to the momentum strategy is reversed (we are long in losers and short sell previous winners). Thus, CI not only mitigates momentum crashes but also benefits from them.

2.4.2 Risk-managed Performance

To make results comparable, we follow Daniel and Moskowitz (2016) and scale all strategies to have an annualized in-sample volatility of 19% over the sample period from September 1928 to May 2020.[23]

The time-varying exposure of all strategies is presented in Fig. 2.5.[24] Panel A presents the weights of CI (dotted line) and DYN (solid line). Except for the first years (prior to the first momentum crash), non-crash weights of CI and DYN are highly correlated (96.5%) and in non-crash periods after 1933, correlation even increases to 98.8%. Considering that far more data is needed to reasonably estimate DYN, this is an interesting observation. While both strategies exhibit large maximum weights (CI: 4.88, DYN: 4.26), CI's minimum (−1.78) is distinctly lower than that of DYN (−0.40).[25] Panel B displays weights for the constant volatility strategy. Even though in-sample volatility is scaled to 19%, weights are similar to those of Barroso and Santa-Clara (2015), because the actual

[23] This approach does not affect the Sharpe ratio since average returns and volatilities change proportionally.

[24] By definition, original momentum exhibits a weight of one.

[25] Daniel and Moskowitz (2016) report weights ranging from −0.60 to 5.37. Our results differ as Daniel and Moskowitz (2016) only report values related to their ex-post strategy.

Fig. 2.5: Weights in Momentum

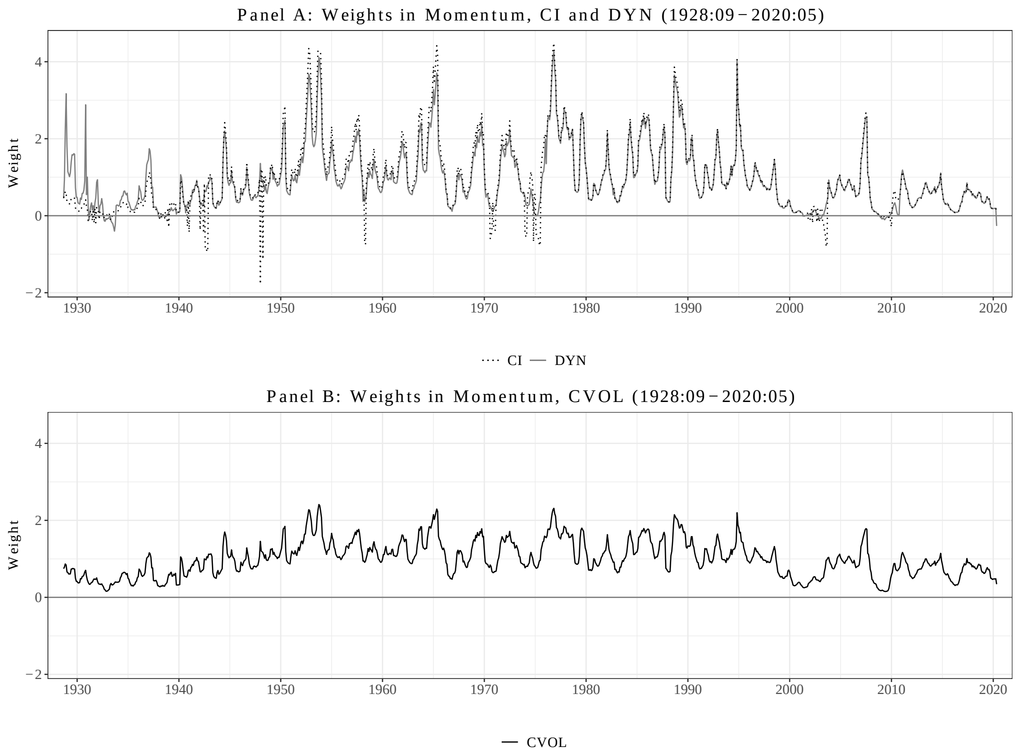

Fig. 2.5 plots scaling weights of risk management strategies over the full sample period from September 1928 to May 2020. Panel A displays the weights of the crash indicator (CI) and dynamic scaling strategy (DYN), while Panel B plots weights with respect to the constant volatility strategy (CVOL). Following Daniel and Moskowitz (2016) and to make results comparable, all risk management strategies are scaled to have an annualized in-sample volatility of 19%.

volatility resulting from a 12%-target is roughly 16.5%, and thus close to the in-sample volatility imposed here. While weights reported in Barroso and Santa-Clara (2015) range from 0.13 to 2.00, respective values in our analysis are 0.15 and 2.41. The standard deviation of weights is equal to 0.45 and thus falls below DYN (0.80) and CI (0.92).

Fig. 2.6 presents cumulative returns of all risk management strategies over the full sample period from September 1928 to May 2020 (Panel A) and the crash periods of the 1930s (Panel B) and 2000s (Panel C).

Fig. 2.6: Risk-Managed Performance: Cumulative Returns

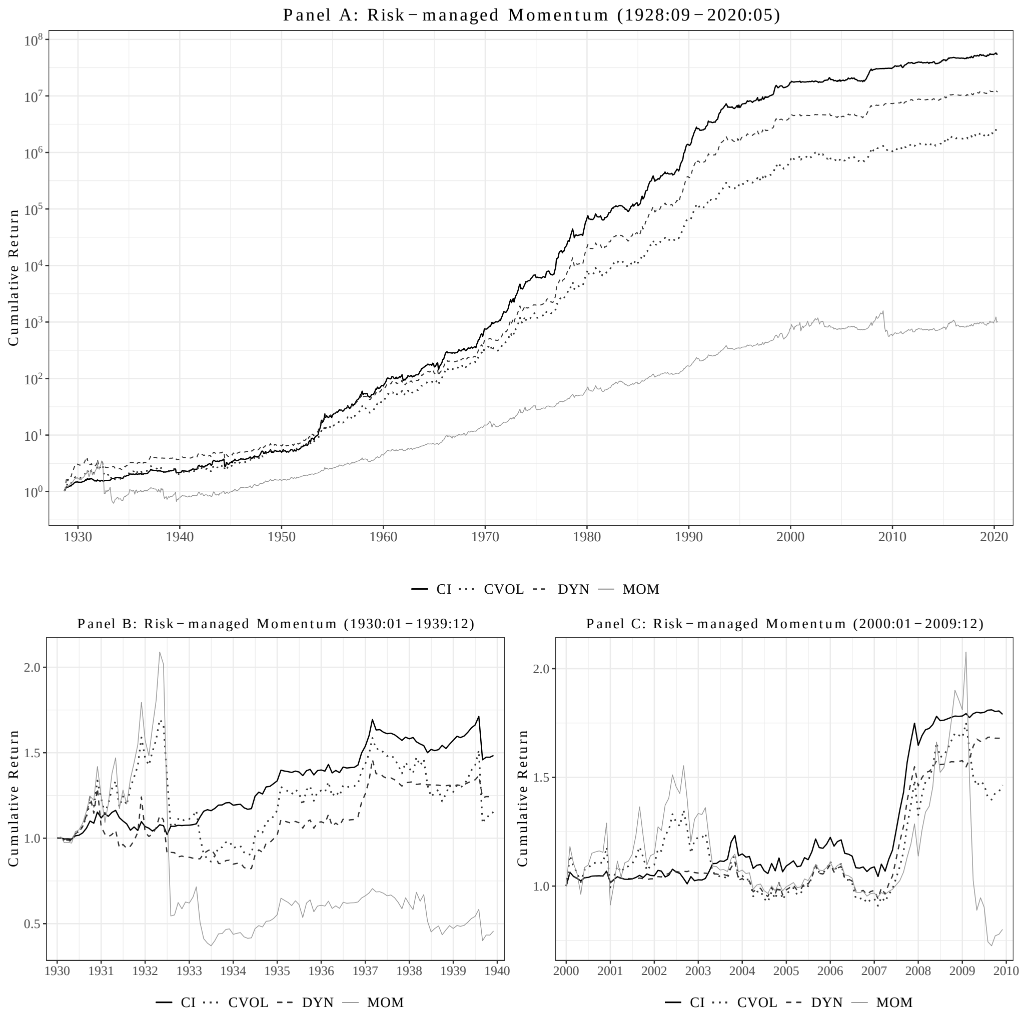

Fig. 2.6 presents cumulative returns of original momentum (MOM), the constant volatility strategy (CVOL), dynamic scaling (DYN), and the crash indicator strategy (CI). Panel A shows cumulative returns over the full sample period from September 1928 to May 2020, while Panels B (1930s) and C (2000s) display cumulative returns in crash periods. Following Daniel and Moskowitz (2016) and to make results comparable, all risk management strategies and original momentum are scaled to have an annualized in-sample volatility of 19%. Cumulative returns of Panels B and C are calculated applying full-sample returns, i.e. returns are not re-scaled within decades. The y-axis of Panel A is logarithmized to improve visibility.

Panel A shows that CI clearly outperforms original momentum, CVOL, and DYN. While a $1 investment in the CI strategy in September 1928

Table 2.5: Risk-Managed Performance: Descriptive Statistics

Statistic	Full Period			
	MOM_{raw}	CVOL	DYN	CI
Mean	13.81%	17.85%	19.62%	21.33%
Median	17.64%	18.80%	10.73%	12.46%
Minimum	−77.02%	−28.26%	−24.62%	−25.39%
Maximum	26.16%	24.99%	42.18%	44.18%
Volatility	27.32%	19.00%	19.00%	19.00%
Sharpe Ratio	0.51	0.94	1.03	1.12
Skew	−2.27	−0.32	0.85	0.75
Kurtosis	16.58	2.02	6.47	6.07

Table 2.5 presents descriptive statistics for original momentum (MOM_{raw}), the constant volatility strategy (CVOL), dynamic scaling (DYN), and the crash indicator strategy (CI). Calculations cover the full sample period from September 1928 to May 2020. Following Daniel and Moskowitz (2016) and to make results comparable, all risk management strategies are scaled to an annualized in-sample volatility of 19%. Both, mean and median returns are annualized, while minimum and maximum returns are stated as monthly returns. Furthermore, we annualize volatilities and Sharpe ratios.

would have led to almost 55 million dollars, the same investment in DYN would have generated roughly 11.8 million dollars. This is equal to an outperformance of 367%. With respect to CVOL, results are even clearer as cumulative returns of CI are almost 24 times higher. Original momentum is staggeringly outperformed by roughly 55 million dollars.

As shown in Table 2.5, all risk management strategies successfully reduce extreme losses. While original momentum exhibits monthly crashes of up to 77%, risk-managed crashes reduce to about 28.3% (CVOL), 24.6% (DYN), and 25.4% (CI), respectively. Moreover, annualized mean returns of CI (21.33%) exceed those of original momentum (13.81%) by almost

55%. CVOL and DYN are outperformed by 19.5% and 9%, respectively.[26]

At a first glance, the increase in monthly returns may appear small, yet the actual improvement is strong. First, a large chunk of these benefits can be reaped precisely when marginal utility is highest, namely during crash periods (Panels B and C of Fig. 2.6). Without rescaling to 19%, average monthly returns of the CI strategy in the 1930s (2000s) amount to roughly 0.35% (0.53%), whereas returns of DYN and CVOL decrease to 0.26% (0.48%) and 0.28% (0.43%), respectively. By rescaling strategies to 19%, results become even clearer. Average CI returns are 0.78% (1.02%), while CVOL and DYN earn only 0.25% (0.52%) and 0.30% (0.96%), respectively.

Second, Table 2.6 presents t-tests for differences in means. CI returns are significantly different from original momentum, both over the full sample and in respective states of the crash indicator. Most importantly, although not being significant over the full sample (when considering CVOL and DYN), CI returns are significantly higher in crash periods. Thus, the crash indicator strategy particularly mitigates crash risk, yet keeping momentum's upside potential alive.

Third, in Section 2.5 we perform robustness checks that confirm a superior performance of CI in sub-samples.[27] Moreover, they reveal a significantly worse performance of the dynamic strategy when applied to

[26] We also construct a strategy that adjusts the Daniel and Moskowitz (2016) approach for the crash indicator (instead of a bear market indicator). Risk adjusted returns largely exceed original momentum but trail DYN.

[27] In a first draft, only pre-Covid data until October 2019 was available. While April 2020 has been one of the 15 worst monthly momentum returns (−28%), the momentum strategy merely lost 2.8% throughout February 2020 to May 2020. Thus, the Covid induced market crash did not result in a momentum crash, which is why CI correctly did not indicate a crash. Hence, CI only lost 0.3%, whereas DYN performed worse than original momentum (−3.1%). We consider this as out-of-sample evidence.

Table 2.6: T-Tests for Differences in Average Returns

Sample	*p*-value		
	$CVOL$	DYN	MOM_{raw}
Full Sample	0.2150	0.5423	0.0000***
$I_C = 0$	0.5130	0.7141	0.0004***
$I_C = 1$	0.0051***	0.0390**	0.0146**

Table 2.6 presents p-values of t-tests for differences in means of the crash indicator strategy (CI) and the other risk management strategies. Positive numbers imply higher average returns of CI. $I_{C,t-1} = 0$ and $I_{C,t-1} = 1$ indicate the current state of the crash indicator proposed in section 2.3.2. Full Sample tests cover data from September 1928 to May 2020. Stars indicate significance at the 10% (*), 5% (**) and 1% (***) level.

relatively new markets.

Fourth, although CI requires less data than DYN, even non-crash returns ($I_{C,t-1} = 0$) are higher (1.85% vs. 1.76% per month). Our results thus suggest that DYN's outperformance with respect to CVOL is exclusively driven by applying variance scaling instead of volatility scaling. In fact, variance scaling in the spirit of Barroso and Santa-Clara (2015), i.e. without having to estimate μ, earns higher full period returns than DYN (1.68% vs. 1.63%).

Finally, by *ex-post* scaling risk-adjusted returns to 19%, Daniel and Moskowitz (2016) circumvent a highly relevant problem of implementing their strategy: there is no hint on how investors could intuitively choose the scaling parameter λ *ex-ante*.[28] To investigate the sensitivity with respect to the choice of λ, we recalculate risk-managed returns for several parameter choices. As the full sample λ to achieve 19% annualized volatility is roughly 0.49, we choose values from $\lambda = 0.1$ to $\lambda = 1.0$ to cal-

[28] See Equation (2.5).

Fig. 2.7: Sensitivity Analysis of DYN (1928:09 - 2020:05)

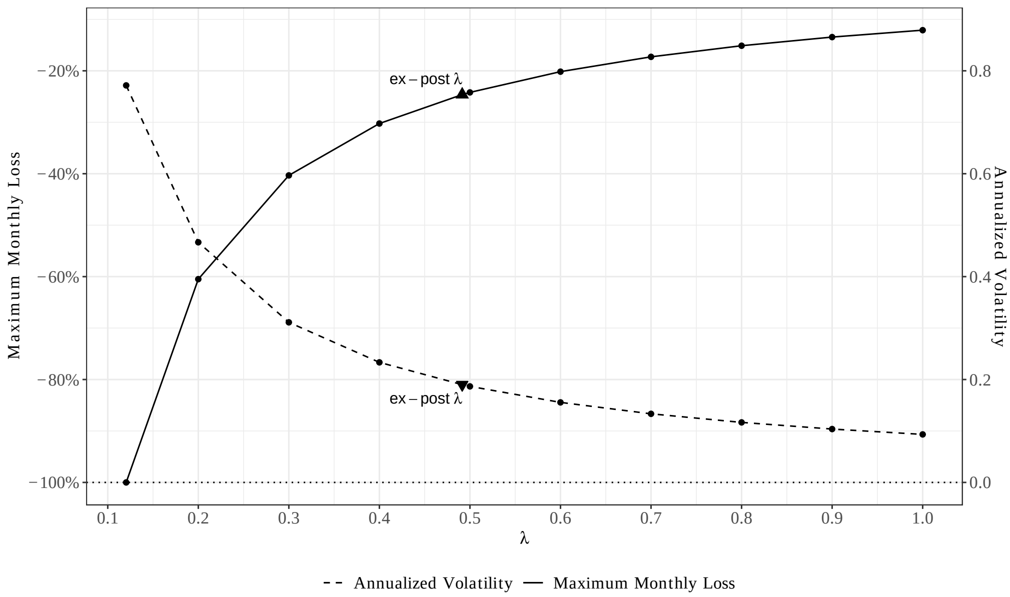

Fig. 2.7 presents maximum monthly losses and annualized in-sample volatilities of the dynamic scaling strategy (DYN) with respect to the scaling parameter λ, presented in Equation (2.5). Returns are calculated over the full sample period from September 1928 to May 2020. *Ex-post* λ (= 0.491) denotes the full sample λ that is chosen to achieve an in-sample annualized volatility of 19%.

culate both maximum monthly losses and annualized volatilities. Fig. 2.7 presents results. While $\lambda > 0.49$ slightly reduce losses and annualized volatilities, $\lambda < 0.49$ result in sharply increased draw-downs, accompanied by large volatilities. Most importantly, with $\lambda < 0.12$ investors would have lost their full investment, rendering risk management ineffective.[29] Nevertheless, it is questionable whether very small λ's are implementable at all ($\lambda = 0.12$ would produce a maximum weight of 17.3).[30] Rational $\mu-\sigma$ investors might note that Sharpe ratios are not affected by λ's greater than 0.12. However, Shleifer and Vishny (1997) show that professional

[29] This is particularly relevant for the 1930s momentum crash when no empirical data was available.

[30] We thank an anonymous referee for this suggestion.

arbitrage is usually conducted by a small number of arbitrageurs, using the capital of less sophisticated, potentially irrational, investors who are unaware of fair values and act rather myopic. Thus, although the choice of $\lambda > 0.12$ does not influence Sharpe ratios directly, it may significantly affect monthly draw-downs, potentially leading to fund withdrawals and ultimately an indirect effect on Sharpe ratios.

Apart from that, CI and DYN offer the highest maximum returns (44.2% and 42.2%), but also exhibit the lowest medians (12.5% and 10.7%) and a high kurtosis (6.1 and 6.5), suggesting that average returns are driven by a small number of particularly high returns. However, the kurtosis is still clearly reduced and both strategies show a positive skewness of 0.75 and 0.85, respectively. In contrast, CVOL displays a kurtosis of roughly 2.0 and returns are still negatively skewed (−0.32). Hence, all strategies effectively reduce tail risk. Finally, as volatilities coincide by construction, Sharpe ratios reflect average returns: CI displays the highest Sharpe ratio (1.12), followed by DYN (1.03) and CVOL (0.94).[31]

Despite offering the largest Sharpe ratio, the high variability of CI's weights makes it reasonable to ask whether our results hold after including transaction costs. Following Barroso and Santa-Clara (2015) and Hanauer and Windmueller (2021), we therefore first calculate each strategy's turnover.[32] Table 2.7 presents results. For the original momentum strategy, we obtain an average monthly turnover of roughly 81.1%, which

[31] Note that the discussion of average returns extends to Sharpe ratios. By applying an ex-post crash indicator, average returns increase to 22.8% and the Sharpe ratio improves to roughly 1.20.

[32] The calculation of the monthly turnover is outlined in Appendix 2.A.4.

is similar to the results of Barroso and Santa-Clara (2015) who report a monthly turnover of 74%. However, note that their sample is limited to the period from March 1951 to December 2010.[33] Average monthly turnovers of risk-managed strategies are 81.5% (CVOL), 78.3% (DYN), and 84.4% (CI), respectively. While we expected CI to display the highest turnover, results for the dynamic strategy are surprisingly low. Although the average absolute change of w_{t-1}^{DYN} (0.154) more than doubles the change of w_{t-1}^{CVOL} (0.073), the portfolio turnover is 3.2 percentage points lower and even falls below MOM_{raw}. We explain this finding by DYN's mean and median weight in the momentum strategy (0.94/0.76), which is lower than for both CI (0.97/0.76) and CVOL (1.00/0.96). In high turnover periods (when more stocks enter or leave one of the portfolios), smaller weights reduce monthly turnover, while in periods of low turnover, the larger variation in momentum weights becomes more relevant. In these periods, DYN indeed displays a higher portfolio turnover. However, we find the former effect to be stronger.[34] By construction, these findings

[33] Hanauer and Windmueller (2021) find a clearly lower turnover of roughly 54%. This finding is caused by constructing the momentum strategy with HML-style portfolios based on 70%/30% percentile breakpoints and double-sorts instead of momentum deciles. Moreover, they report a significantly larger increase of turnover when risk-managed strategies are considered. We explain this difference by higher scaling weights (their risk-managed strategies are scaled to have the same annualized volatility as original momentum). Daniel and Moskowitz (2016) do not report turnovers.

[34] We examine DYN's turnover relative to the turnover of CVOL ($\Delta = TO_{DYN} - TO_{CVOL}$). First, we calculate DYN's average turnover with respect to the monthly change in weights. If the difference is above the median (0.038), we find Δ to be 0.08 (i.e. TO_{DYN} exceeds TO_{CVOL} by eight percentage points), while otherwise Δ is clearly negative (-0.14). Although the variation of weights indeed influences monthly turnover, the positive effect of a lower mean exposure is prevailing. Moreover, we calculate Δ with respect to the turnover of original momentum. If TO_{MOM} is below (above) its median (0.796), Δ amounts to roughly -4.0 (-2.5) percentage points. Thus, in high turnover periods the impact of a smaller median weight increases.

Table 2.7: Turnover and Break-even Round Trip Costs

	MOM_{raw}	CVOL	DYN	CI
Turnover	81.11%	81.47%	78.27%	84.44%
Round-trip costs (5% level)	10.14%	17.14%	20.10%	20.65%
Round-trip costs (1% level)	7.97%	15.64%	18.54%	19.20%

Table 2.7 presents the average monthly portfolio turnover of original momentum (MOM_{raw}), the constant volatility strategy (CVOL), dynamic scaling (DYN), and the crash indicator strategy (CI). Furthermore, we report annualized break-even round trip transaction costs, i.e. round trip costs that would render profits of each strategy insignificant at the 5% and 1% level, respectively. Both, turnovers and break-even round trip costs are calculated according to Appendix 2.A.4. We cover the full sample period from September 1928 to May 2020.

also come into effect when considering the CI strategy. However, in this case, turnover increases by incorporating the crash indicator.

To prove that the CI strategy remains superior after including transaction costs, we follow Barroso and Santa-Clara (2015), Hanauer and Windmueller (2021), and Grundy and Martin (2001) and calculate the round trip transaction costs that would render profits of each risk management strategy insignificant at the 5% and 1% level, respectively.[35] As depicted by Table 2.7, CI's annualized break-even transaction costs at the 1% level are 19.20% and exceed both DYN (18.54%) and CVOL (15.64%). Results at the 5% level are very similar. We thus conclude that the crash indicator strategy still outperforms CVOL and DYN after including transaction costs.

[35] The calculation of break-even transaction costs is outlined in Appendix 2.A.4.

2.4.3 Spanning Tests

To investigate whether risk management strategies are spanned by common risk factors or each other, we perform regressions of monthly risk-managed returns on variations of Fama and French (1993) factors and risk-managed returns of the remaining strategies. Table 2.8 reports the corresponding regression alphas. For ease of comparison, we follow Daniel and Moskowitz (2016) and scale risk-managed returns to have the same annualized in-sample volatility. That is, we apply scaled full-sample returns, as presented in Section 2.4.2.[36]

Table 2.8: Spanning Tests

	Model								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	$FF3$	$FF3$ +MOM	$FF3$ +$CVOL$	$FF3$ +DYN	$FF3$ +CI	$FF3$ +$CVOL$ +DYN	$CVOL$ +DYN	CI +DYN	CI +$CVOL$
α_{CVOL}	1.70***	0.62***		0.23***	0.19**			−0.01	
	(10.67)	(6.81)		(3.00)	(2.26)			(−0.07)	
α_{DYN}	1.72***	0.91***	0.14*		0.03				−0.01
	(10.37)	(6.66)	(1.79)		(0.52)				(−0.15)
α_{CI}	1.82***	1.10***	0.29***	0.22***		0.18***	0.26***		
	(10.96)	(7.59)	(3.28)	(3.42)		(2.82)	(4.01)		

Table 2.8 presents full-sample (1928:09-2020:05) regression alphas of monthly risk-managed momentum returns with respect to Fama and French (1993) factors (FF3) and the other risk management strategies. CVOL and DYN denote the constant and dynamic scaling strategy, whereas CI corresponds to the crash indicator strategy. Following Daniel and Moskowitz (2016), all strategies are scaled to have the same annualized in-sample volatility of 19%. Alphas are stated in percentage points. Stars indicate significance at the 10% (*), 5% (**) and 1% (***) level, t-values are stated in parentheses.

[36] For their spanning tests, Daniel and Moskowitz (2016) change scaling to have an annualized in-sample volatility of 23%. For consistency, we keep our previously used target of 19%.

In Model (1), we simply regress risk-managed momentum returns on the Fama and French (1993) three factor model (FF3). All strategies display large and significant alphas, whereby both economic and statistical significance is highest for CI. Model (2) extends the right-hand side by original momentum, leading to the Carhart (1997) four factor model. Although t-statistics and alphas decrease, all strategies remain highly significant. More precisely, t-statistics range from 6.66 (DYN) to 7.59 (CI) and risk-adjusted returns of CI and DYN (1.10% and 0.91%) clearly exceed those of CVOL (0.62%). This finding is reasonable as CVOL only depends on momentum-specific information. Models (3) to (6) provide combinations of Fama and French (1993) factors and risk-managed momentum returns. Model (3) controls for CVOL. While $\alpha_{CI} = 0.29\%$ remains highly significant ($t = 3.28$), α_{DYN} is distinctly smaller (0.14%) and merely significant at the 10% level ($t = 1.79$). Model (4) replaces CVOL by DYN and displays highly significant alphas, both for CVOL and CI. However, α_{CI} is still the most significant. Controlling for CI (Model 5), α_{CVOL} shows a lower t-value (2.26) but is still significant at the 5% level. In contrast, α_{DYN} reduces to 0.03% and is not significant at any conventional level. To investigate the joint impact of DYN and CVOL, Model (6) combines Fama and French (1993) factors with both competing strategies. Emphasizing previous results, α_{CI} remains large (0.18%) and significant at the 1% level. Finally, Models (7) to (9) test each strategy with respect to both of the remaining strategies (jointly and without FF3). Results are impressive. While both α_{DYN} and α_{CVOL} (Models 8 and 9) show no significance at all, α_{CI} (Model 7) remains highly significant ($t = 4.01$).

In summary, all regressions suggest a superior performance of CI. Since CVOL returns show reduced significance after controlling for FF3 and CI, and there is no significance when controlling for CI and DYN simultaneously, the CVOL strategy is partially spanned. Likewise, after including CI as an explanatory variable, DYN alphas become very small and are not significant at all. We find our results to be confirmed when CVOL and DYN are solely regressed on CI (not reported). We thus conclude that DYN and CVOL returns are (partially) spanned by CI. Moreover, our findings are in contrast to Daniel and Moskowitz (2016) who find CVOL returns to be spanned by the dynamic strategy.[37]

In unreported results, we replace the value factor (HML) by Asness and Frazzini (2013)'s 'HML Devil' factor (HML_d), which depends on more recent data and is thus considered to be a better predictor of momentum. Results are similar. CI is not spanned by one of the other strategies and DYN returns are spanned by CI. However, DYN is now also spanned by CVOL and α_{CVOL} displays higher t-values.[38] Moreover, we re-estimate all strategies with a 2×3 style momentum portfolio.[39] Results correspond to earlier findings. CI offers the highest Sharpe ratio and is not spanned by one of the other strategies, while DYN continues to be spanned by CI.[40]

[37] Results differ from Daniel and Moskowitz (2016) for two reasons. First, Daniel and Moskowitz (2016) only report spanning tests for their ex-post strategy. Second, reported alphas are estimated with respect to interaction terms of Fama and French (1993) factors and a market stress indicator. By construction, average returns in these periods are lower.

[38] We thank an anonymous referee for this suggestion.

[39] Two portfolios are formed on size and three portfolios are formed on prior returns (breakpoints are the 30^{th} and 70^{th} percentile). Returns of the winner (loser) portfolio are then calculated by the average of small and big previous winners (losers).

[40] When Fama and French (1993) factors are excluded, CVOL returns are spanned as well. We thank an anonymous referee for this suggestion.

2.4.4 International Evidence

To investigate the international performance of risk management strategies, we construct an international momentum portfolio. Thereby, we follow Barroso and Santa-Clara (2015) and focus on the most important markets outside of the United States: France, Germany, Japan, and the United Kingdom. In order to construct the portfolio, we first estimate all risk management strategies in each respective market. Thereafter, we scale all strategies to have an annualized in-sample volatility of 19% and calculate monthly international portfolio returns by equally weighting each market. After deducting months that are needed to estimate the strategies, we cover a 31-year period from March 1989 to May 2020. Table 2.9 reports results.

Although strategies in each country are scaled to have an in-sample volatility of 19%, the *portfolio* volatility is lower and ranges from 11.9% (CI) to 13.8% (CVOL). Thus, independent of risk management, risk can be reduced by constructing an international portfolio. Nevertheless, except for original momentum (9.7%), average returns are similar to those in the US.[41] The highest mean returns are offered by CI (21.1%), whereas DYN and CVOL earn average returns of 18.0% and 20.3%, respectively. As a result, CI also offers the highest Sharpe ratio (1.77), followed by CVOL (1.48) and DYN (1.47).[42] All of which clearly exceed the Sharpe ratio of original momentum (0.76). Apart from that, maximum losses

[41] Low returns of original momentum are driven by Japanese momentum. This finding is consistent with Barroso and Santa-Clara (2015) and Asness (2011).

[42] Results hold when considering a value-weighted international portfolio.

Table 2.9: International Performance: Descriptive Statistics

Statistic	International Portfolio			
	MOM_{raw}	$CVOL$	DYN	CI
Mean	9.74%	20.32%	17.95%	21.07%
Median	11.09%	20.27%	13.60%	17.48%
Minimum	−24.97%	−12.20%	−11.89%	−9.05%
Maximum	15.23%	17.55%	16.38%	13.74%
Volatility	12.85%	13.75%	12.18%	11.89%
Sharpe Ratio	0.76	1.48	1.47	1.77
Skewness	−0.92	−0.08	0.30	0.60
Kurtosis	7.89	0.74	2.27	1.09

Table 2.9 presents descriptive statistics with respect to an international momentum portfolio. In order to construct the portfolio, we first estimate all risk management strategies in each of the four most important markets outside of the United States: France, Germany, Japan, and the United Kingdom. We choose countries according to Barroso and Santa-Clara (2015). Second, all strategies are scaled to have an annualized in-sample volatility of 19%. Third, we calculate monthly portfolio returns by equally weighting each of the international markets. Both, mean and median returns are annualized, while minimum and maximum returns are stated as monthly returns. Furthermore, we also annualize volatilities and Sharpe ratios. CVOL and DYN denote the constant and dynamic scaling strategy, whereas CI corresponds to the crash indicator strategy. Reported values are calculated for the period from March 1989 to May 2020.

strongly decrease. While original momentum offers a minimum return of roughly −25.0%, returns increase to about −9.1% when considering CI. Results for CVOL (−12.2%) and DYN (−11.9%) are similar. Due to smaller volatilities, we also observe decreasing maximum returns. Nonetheless, except for original momentum, maxima exceed minima by 38% (DYN) to 52% (CI). Consequently, the risk management strategies either display a positive (CI and DYN) or only slightly negative skewness (CVOL). Moreover, as kurtosis ranges from 0.7 (CVOL) to 2.3 (DYN), all strategies show a negative excess kurtosis.[43] To prove that our findings are not simply driven by the specific choice of countries, we repeat our analysis for a Global-Ex-USA portfolio and several regional momentum portfolios.[44] In all regions, CI strongly improves momentum returns. Furthermore, CI offers the highest Sharpe ratio with respect to the global portfolio (CI 1.33, CVOL 1.31, DYN 1.18), in North America (CI 1.35, CVOL 1.16, DYN 1.14), and in the Pacific region (CI 0.88, CVOL 0.74, DYN 0.67), while slightly trailing CVOL in Europe (CI 1.74, CVOL 1.76, DYN 1.66).

Finally, to study whether the risk management strategies span each other in an international environment, we also repeat the spanning tests.[45] Results are presented in Table 2.10. After adjusting for Fama and French (1993) factors and original momentum (Models 1 and 2), all strategies remain highly profitable. In Models (3) to (6), α_{CI} remains significant at the 1% level, both after adjusting for CVOL and DYN separately (Models 3 and 4) and jointly (Model 6). When adjusting for CI, α_{DYN} is barely significant at the 5% level and α_{CVOL} is not significant at any conventional level (Model 5). Without Fama and French (1993) factors (Models 7-9), results are even clearer. CI remains significant at the 1% level, while α_{CVOL} and α_{DYN} show no significance at all. We find our results to be confirmed when CVOL and DYN are regressed on CI only (not reported).

In summary, we find CI to outperform the other risk management

[43] In each respective country – including Japan – CI and CVOL clearly improve the performance. In Japan, the Sharpe ratio of CI (0.27) almost quadruples the ratio of original momentum (0.07), while DYN's improvement is small (0.08).

[44] We thank an anonymous referee for this suggestion.

[45] Country-specific Fama and French (1993) factors are provided by AQR capital. In accordance with the calculation of portfolio returns, factor returns are also weighted equally.

Table 2.10: Spanning Tests for International Momentum Strategies

	Model								
	(1) *FF3*	(2) *FF3 +MOM*	(3) *FF3 +CVOL*	(4) *FF3 +DYN*	(5) *FF3 +CI*	(6) *FF3 +CVOL +DYN*	(7) *CVOL +DYN*	(8) *CI +DYN*	(9) *CI +CVOL*
α_{CVOL}	1.86***	0.94***		0.36**	0.24			−0.01	
	(9.11)	(9.05)		(2.55)	(1.61)			(−0.04)	
α_{DYN}	1.67***	1.08***	0.38***		0.27**				0.17
	(9.33)	(7.54)	(3.13)		(1.99)				(1.52)
α_{CI}	1.87***	1.43***	0.61***	0.60***		0.46***	0.52***		
	(10.36)	(8.65)	(4.74)	(4.53)		(3.77)	(4.21)		

Table 2.10 presents regression alphas of the risk-managed international momentum portfolio with respect to Fama and French (1993) factors (FF3) and the other risk management strategies. In order to construct the portfolio, we first estimate all risk management strategies in each of the four most important markets outside of the United States: France, Germany, Japan, and the United Kingdom. We choose countries according to Barroso and Santa-Clara (2015). Second, all strategies are scaled to have an annualized in-sample volatility of 19%. Third, we calculate monthly portfolio returns by equally weighting each of the international markets. Corresponding Fama and French (1993) factors are calculated by equally weighting respective factors of each country. CVOL and DYN denote the constant and dynamic scaling strategy, whereas CI corresponds to the crash indicator strategy. Alphas are stated in percentage points. Stars indicate significance at the 10% (*), 5% (**) and 1% (***) level, t-values are stated in parentheses.

strategies in terms of Sharpe ratios, both in the United States and internationally. Moreover, the CI strategy sharply decreases maximum losses and tail risk, while still preserving high maxima. As a result, CI is not spanned by one of the other strategies, whereas DYN and CVOL are at least partially spanned, both in the United States and internationally.

2.5 Robustness Checks

2.5.1 Full-sample Scaling and Sub-sample Performance

To verify the insights of Section 2.4.2, we now introduce sub-samples. Each sub-sample contains one decade of monthly returns, starting in January 1930. To calculate cumulative returns of each strategy, we apply full sample risk-adjusted returns, i.e. we do not rescale sub-samples to have an annualized in-sample volatility of 19%. The interpretation is as follows. Assume an investor who starts at the end of our sample, i.e. May 2020, and has to live through one of the previous decades with its respective returns. Fig. 2.8 illustrates the results.

Remarkably, DYN is outperformed in all decades, including the 1960s, 1980s, and 1990s, when no crash was indicated. This finding confirms that dynamic scaling is inferior to simple variance scaling (which does not require the estimation of μ). Furthermore, CVOL is outperformed in seven out of nine decades, while another decade offers almost equivalent returns (1940s, Panel B). In the 2010s (Panel I), CI is inferior to CVOL, but still offers significantly positive returns. However, in this decade CI also displays a very low volatility of roughly 9% (CVOL: 14.7%). As a result, the Sharpe ratio of CI exceeds CVOL by more than 24%. Most importantly, original momentum, DYN, and CVOL are clearly outperformed in the crash periods of the 1930s and 2000s (Panels A and H). Notably, in both decades CI also displays a low volatility (8.5% and 10.0%, respectively), suggesting that the strategy effectively reduces crash risk. Original momentum is outperformed in all but one sub-sample (2010s, Panel I).

Fig. 2.8: Sub-sample Performance

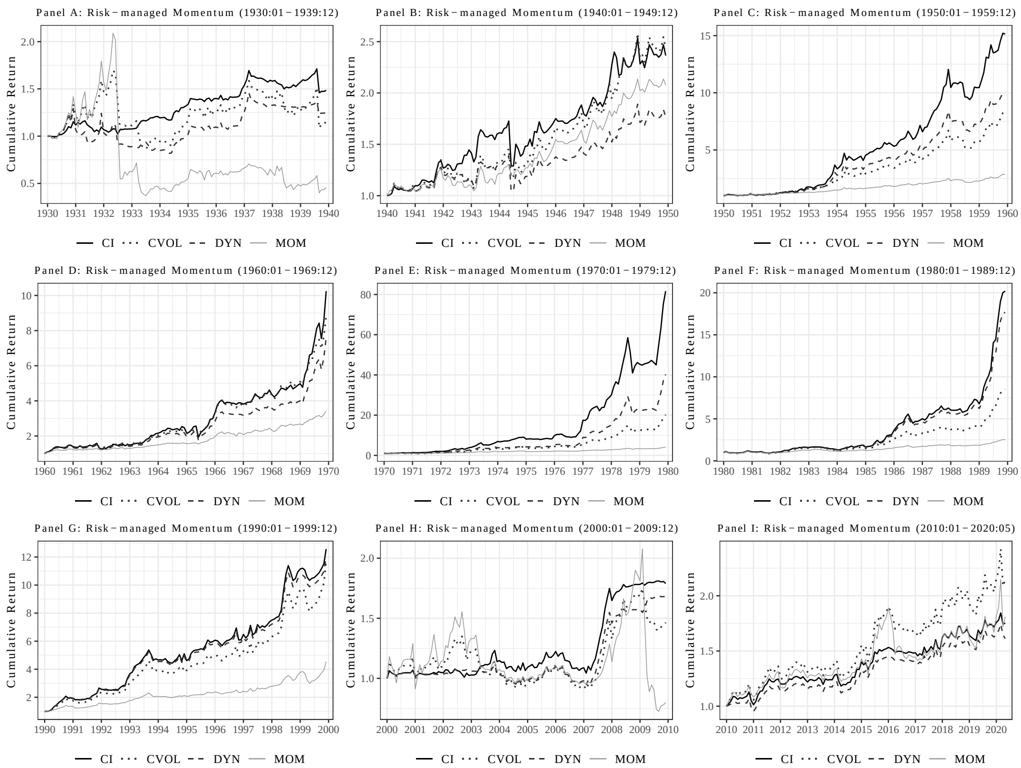

Fig. 2.8 presents sub-sample performance of original momentum (MOM), the constant volatility strategy (CVOL), dynamic scaling (DYN), and the crash indicator strategy (CI). Panels show cumulative returns for each of the covered decades. Following Daniel and Moskowitz (2016) and to make results comparable, full sample strategies are scaled to have an annualized in-sample volatility of 19%. Cumulative returns are calculated applying full-sample returns, i.e. returns are not re-scaled within decades.

However, cumulative returns in this decade are almost equal and original momentum's volatility is clearly higher (16.3%). As a consequence, CI's Sharpe ratio exceeds original momentum by more than 68%. These findings support earlier insights and, given CI's superiority, confirm the power of the crash indicator in separating crashes from momentum bull markets.

To study whether our results are driven by specific starting points, we also calculate 5-year rolling windows. Again, we apply full sam-

Fig. 2.9: 5-Year Rolling Windows

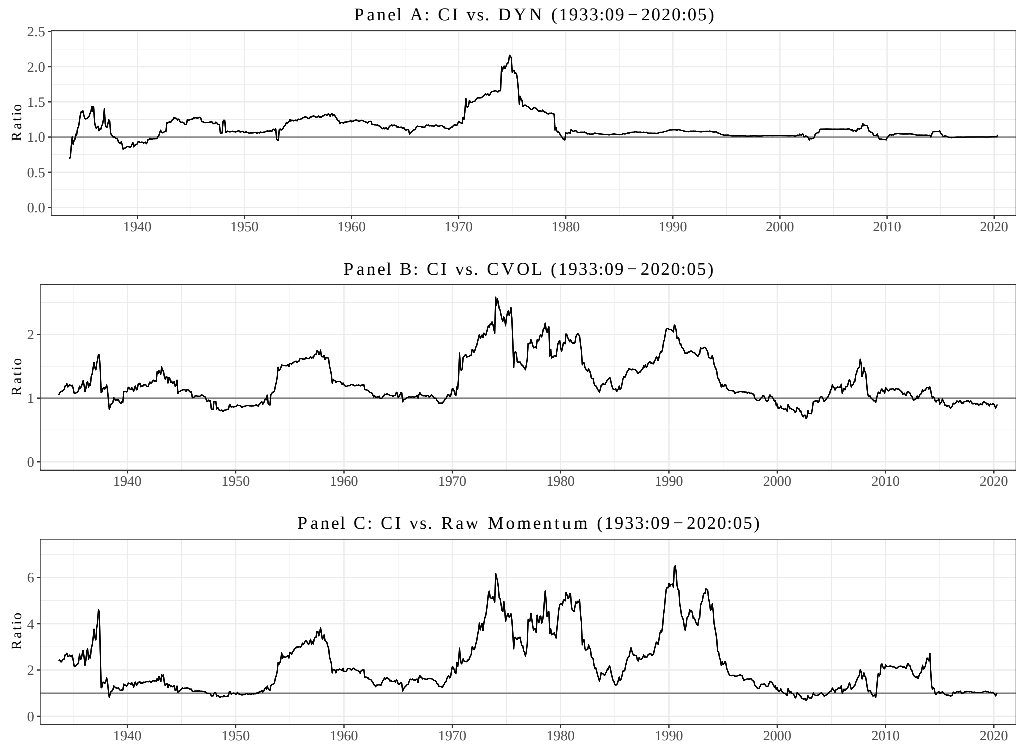

Fig. 2.9 presents ratios of rolling 5-year returns of the crash indicator strategy (CI) relative to dynamic scaling (DYN, Panel A), the constant volatility strategy (CVOL, Panel B), and unscaled original momentum (MOM, Panel C). As we calculate cumulative 5-year returns, reported values cover the period from September 1933 to May 2020. Following Daniel and Moskowitz (2016), full sample returns are scaled to have an annualized in-sample volatility of 19%. Cumulative returns are calculated applying full-sample returns, i.e. returns are not re-scaled within 5-year periods.

ple risk-adjusted returns that are not re-scaled within 5-year periods. Fig. 2.9 presents ratios of cumulative CI returns relative to returns of the DYN strategy (Panel A), the CVOL strategy (Panel B), and raw momentum (Panel C), starting in September 1933 (the first month containing five years of data). A ratio greater than one is tantamount to superior performance of CI, while values below one indicate underperformance. Cumulative returns of CI exceed those of the DYN strategy in roughly

90% of 5-year periods, including major crash periods as well as momentum bear markets of the mid-1970s. In contrast, DYN shows superior performance only in the late 1930s, while the strategy is outperformed by up to 50% in the early 1930s.[46] With respect to CVOL, we find superior CI returns in more than 75% of 5-year periods, including major crash periods of the early 1930s and late 2000s. The extent of outperformance peaks in the 1950s, 1970s, and late 1980s. However, note that in crash periods CI displays a lower volatility. Thus, in terms of Sharpe ratios, the outperformance in these periods is higher than depicted. In contrast, the extent of underperformance within the remaining 25% is of minor importance. Lastly, outperformance further increases when considering original momentum, which is clearly outperformed in more than 88% of rolling windows, and particularly peaks in the 1930s, 1970s, and early 1990s, when CI returns exceed original momentum by up to 550%. Underperformance in the remainder is negligible.

2.5.2 Re-estimated Strategies and Sub-samples

While CVOL and CI only require six and 24 months of data, respectively, DYN performs an expanding window regression that gains power by applying a long sample. To investigate the implementability in markets with shorter time-series, we repeat the robustness checks of Section 2.5.1 with a re-estimated version of the DYN strategy. Hence, we reset expand-

[46] The very first months indicate a superior performance of DYN. However, this finding is related to the fact that until September 1930 (when there was the first bear market), μ is equal to the average of previous momentum returns, resulting in a large weight and higher risk-adjusted returns.

Fig. 2.10: Sub-sample Performance (re-estimated)

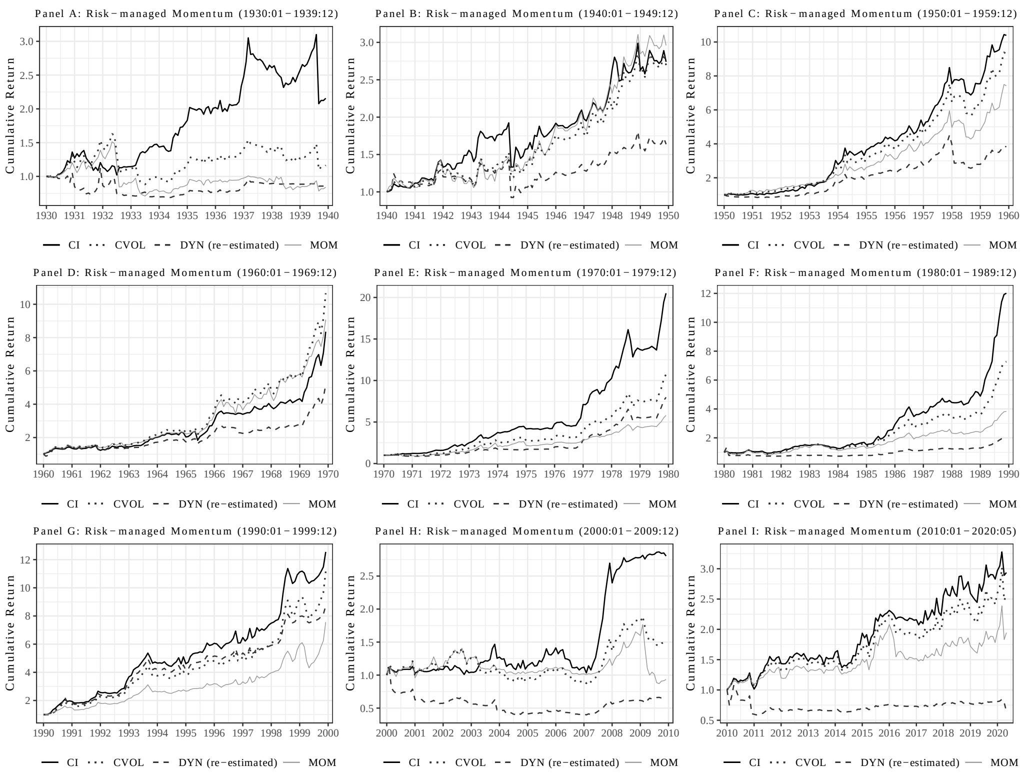

Fig. 2.10 presents the re-estimated sub-sample performance of original momentum (MOM), the constant volatility strategy (CVOL), dynamic scaling (DYN), and the crash indicator strategy (CI). This is, we estimate expanding window regressions for each of the decades. Therefore, we do not apply full sample returns, but scale strategies in each decade to have an annualized in-sample volatility of 19%.

ing windows for each decade and 5-year period. Fig. 2.10 presents the sub-sample performance of re-estimated DYN, CI, CVOL, and original momentum. To make results comparable, we do not apply full sample returns, but scale all strategies to have an annualized in-sample volatility of 19%.

Re-estimated DYN is clearly outperformed in all sub-samples. The performance particularly declines in crash periods (Panels A and H) and in the 2010s (Panel I), when the strategy incurs losses. Moreover, in six out of nine decades re-estimated DYN is even outperformed by original momentum, including the momentum bull markets of the 1950s (Panel C) and 1980s (Panel F). The reasoning is as simple as important. With respect to Equation (2.5), the weight of the dynamic strategy largely depends on expanding window regressions of the expected return $\hat{\mu}_{t-1}$. Until the first bear market, $\hat{\mu}_{t-1}$ is solely based on the intercept and equals the average of preceding momentum returns. As non-crash periods display positive returns, expected momentum returns are rather high and DYN exhibits a high exposure to the momentum strategy. Consequently, the first momentum crash has an exceptionally large impact. This process is particularly obvious in the 1930s (Panel A) and the 2000s (Panel H). Additionally, CVOL and original momentum are outperformed by CI in eight out of nine and seven out of nine periods, respectively. However, in these decades (the 1940s and 1960s), CI still earns significantly positive returns. Most importantly, CI particularly improves in crash decades. In terms of cumulative returns, competing strategies are outperformed by at least 83.5%.

Analogously to Section 2.5.1, we also re-calculate returns for rolling 5-year periods. In order to make results comparable, we now report absolute differences in Sharpe ratios of CI and the other strategies. Fig. 2.11 presents results for the re-estimated DYN strategy (Panel A), CVOL (Panel B) and original momentum (Panel C). Values above zero imply superior risk-adjusted performance of CI, while values below zero indicate underperformance.

CI staggeringly outperforms the re-estimated DYN strategy in almost

Fig. 2.11: 5-Year Rolling Windows (re-estimated)

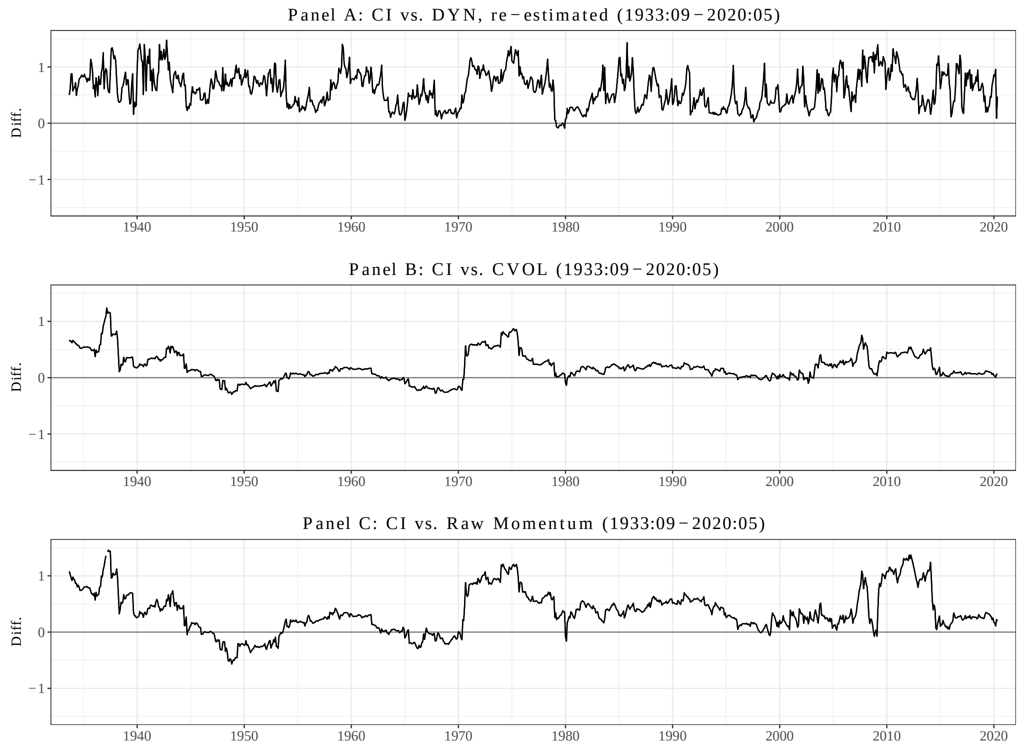

Fig. 2.11 presents the 5-year rolling window performance of original momentum (MOM), the crash indicator strategy (CI), the constant volatility strategy (CVOL), and *re-estimated* dynamic scaling (DYN, re-estimated). This is, we estimate expanding window regressions for each of the 5-year periods. As we calculate 5-year returns, reported values cover the period from September 1933 to May 2020. In order to make results comparable, we report differences in Sharpe ratios of the crash indicator strategy and DYN (Panel A), CVOL (Panel B), and unscaled original momentum (Panel C).

every 5-year period (99%). Average outperformance in terms of Sharpe ratios amounts to more than 0.61, particularly peaking in the crash periods of the 1930s, 1974/1975, and 2009. With respect to CVOL, more than 80% of 5-year periods are outperformed. Differences now peak in the 1930s, 1974/1975, and the 2000s. The same applies to original momentum since the majority of differences (84.1%) largely exceeds zero (with a mean of roughly 0.36).[47] Notably, the average Sharpe ratio of original momentum (0.75) exceeds that of re-estimated DYN (0.50), while the ratios of CI (1.11) and CVOL (0.93) are consistent with the full sample results.

In summary, our findings strongly suggest that DYN is not particularly successful in markets that have not yet experienced a momentum crash.[48] In contrast, both CVOL and CI provide successful risk management without requiring long time series.

2.6 Concluding Remarks

On average, investing in past winners and short-selling past losers provides highly significant returns that cannot be explained by common risk factors. However, the momentum strategy also displays substantial tail risk as there are short but persistent periods of highly negative returns. Crashes particularly occur in rebounding bear markets, when the momentum portfolio displays a negative market-beta and momentum volatility is high. We show that a crash indicator based on systematic risk measures isolates momentum crashes from momentum bull markets. Furthermore, we find enhanced predictive power when combining systematic (measured by the crash indicator) and momentum-specific risk (measured by momentum volatility). Based on this finding, we propose

[47] From 1948 to 1953 and 1965 to 1970, the risk-adjusted Sharpe ratio of CI is inferior to original momentum, which is caused by a rather high standard deviation of CI. This finding coincides with Fig. 2.10.

[48] Interestingly, Table 7 of Daniel and Moskowitz (2016) likely exaggerates DYN's performance because they report the performance from 1934 onward, *after* the first momentum crash, although their sample starts in July 1927.

an implementable trading strategy that outperforms the existing risk management strategies of Barroso and Santa-Clara (2015) and Daniel and Moskowitz (2016) over the period from 1928 to 2020, in sub-samples, and after accounting for transaction costs. Moreover, by constructing an international momentum portfolio, we show that our results extend to foreign markets. Importantly, we find that the dynamic scaling strategy of Daniel and Moskowitz (2016) is not successful in markets that have not yet experienced a momentum crash and does not give a hint on how investors could intuitively adjust scaling parameters *ex-ante*, potentially leading to large losses.

2.A Appendix

2.A.1 Estimation of Momentum Beta

At the beginning of each month, beta is estimated by a rolling regression of the 126 preceding daily momentum returns on the CAPM

$$r_{MOM,t} - r_{f,t} = \alpha + \beta_{MOM} \cdot (r_{Mkt,t} - r_{f,t}) + \epsilon_t, \tag{2.A.1}$$

where $r_{MOM,t}$ and $r_{Mkt,t}$ denote daily momentum and market returns, respectively. $r_{f,t}$ is the daily risk-free rate and ϵ_t are residuals at time t.

2.A.2 Estimation of Momentum Volatility

Following Barroso and Santa-Clara (2015), the volatility at the beginning of month t is estimated by the realized volatility of the previous 126 daily momentum returns

$$\sigma_{MOM,t} = \sqrt{12}\sqrt{21\sum_{j=0}^{125} r^2_{MOM,d_{t-1-j}}/126}, \tag{2.A.2}$$

where $\sigma_{MOM,t}$ states the annualized momentum volatility in month t and $r^2_{MOM,d_{t-1-j}}$ denotes the daily squared momentum return at day $t-1-j$.

2.A.3 Estimation of the Ex-Ante Dynamic Strategy

Daniel and Moskowitz (2016) deduce the optimal weight in the momentum portfolio (w_t^{DYN}) by maximizing the in-sample unconditional Sharpe ratio (see their Appendix C) which gives

$$w_t^{DYN} = (\frac{1}{2\lambda})\frac{\hat{\mu}_{t-1}}{\hat{\sigma}^2_{t-1}}. \tag{2.A.3}$$

In their ex-ante approach, the conditional expected return of the momentum portfolio ($\hat{\mu}_{t-1}$) is estimated by an expanding window regression of monthly momentum returns on the interaction between the bear market indicator ($I_{B,t-1}$) and the realized market variance ($\hat{\sigma}^2_{m,t-1}$) over the preceding 126 daily momentum returns

$$r_{MOM,t} = \gamma_0 + \gamma_{int} \cdot I_{B,t-1} \cdot \hat{\sigma}^2_{m,t-1} + \hat{\epsilon}_t. \tag{2.A.4}$$

The conditional variance of the momentum portfolio ($\hat{\sigma}^2_{t-1}$) is estimated by the realized variance of the 126 preceding daily momentum returns. In addition to that, λ is a constant that controls for unconditional risk and return of the risk-managed portfolio. Daniel and Moskowitz (2016) choose λ to achieve an ex-post annualized volatility of 19%. There is no hint on how to determine λ ex-ante.

2.A.4 Turnover Calculation and Break-even Round Trip Costs

To calculate momentum turnover, we follow Barroso and Santa-Clara (2015) and Hanauer and Windmueller (2021). For each leg (the winner and loser portfolio), we compute monthly turnover by

$$TO_{t,l} = 0.5 \cdot \sum_{t}^{N_t} |w_{i,t} - \widetilde{w}_{i,t-1}|, \tag{2.A.5}$$

where $TO_{t,l}$ denotes the turnover of leg l in month t, $w_{i,t}$ is the weight of stock i in the respective leg at time t, and N_t is the corresponding number of stocks. Moreover, $\widetilde{w}_{i,t-1}$ is the weight of stock i right before portfolio re-balancing. More precisely, we adjust weights of the previous period by including each constituent's return ($r_{i,t-1}$) during that period

$$\widetilde{w}_{i,t-1} = \frac{w_{i,t-1} \cdot (1 + r_{i,t-1})}{\sum_t^{N_t} w_{i,t-1} \cdot (1 + r_{i,t-1})}. \tag{2.A.6}$$

Monthly turnover of the momentum portfolio is then calculated by summing up the turnover of the long and the short leg. Finally, to obtain the turnover of the risk-managed strategies, we adjust Equation (2.A.5) by including the risk-managed weights (i.e. w_{t-1}^{DYN}, w_{t-1}^{CVOL}, and w_{t-1}^{CI}) at time t (denoted by $w_{scaled,t}$)

$$TO_{t,l} = 0.5 \cdot \sum_t^{N_t} |w_{scaled,t} \cdot w_{i,t} - w_{scaled,t-1} \cdot \widetilde{w}_{i,t-1}|. \tag{2.A.7}$$

Following Grundy and Martin (2001), Barroso and Santa-Clara (2015), and Hanauer and Windmueller (2021), we further calculate break-even round trip costs, i.e. round trip transaction costs that would render profits of each strategy insignificant at the 5% and 1% level, respectively

$$\text{Round-trip costs} = (1 - \frac{z_{1-\alpha/2}}{t_s}) \cdot \frac{\overline{\mu}_s}{\overline{TO}_s}, \tag{2.A.8}$$

where t_s is the t-statistic of strategy s, $\overline{\mu}_s$ is the average monthly return, and $\overline{TO}_s$ is the average monthly turnover of strategy s. Furthermore, $z_{1-\alpha/2}$ denotes the z-value corresponding to the desired level of α (1.96 for $\alpha = 5\%$ and 2.576 for $\alpha = 1\%$).

Chapter 3

Option-implied Lottery Demand and IPO returns

This chapter refers to the following publication:

Dierkes, Maik, Jan Krupski and Sebastian Schroen (2022): 'Option-implied Lottery Demand and IPO Returns', Journal of Economic Dynamics and Control **138**, 104356.

Available online at:
https://doi.org/10.1016/j.jedc.2022.104356

Abstract

We study the impact of time-varying lottery demand on first-day returns and the poor long-term performance of IPOs. Lottery demand – measured in terms of option-implied probability weighting – is associated with significantly higher first-day returns, tantamount to higher IPO underpricing and more money left on the table. Interacting the time variation in lottery demand with cross-sectional expected skewness reveals that IPO returns are particularly driven by the interaction between market-wide lottery demand and asset-specific lottery characteristics. When expected skewness meets low lottery demand, there is virtually no effect of skewness on first-day returns. In the long run, IPOs issued in high lottery demand regimes are more likely to perform poorly for up to five years after the IPO.

Keywords: IPO, Lottery Demand, Skewness Preferences

JEL: G12, G41.

3.1 Introduction

We present a novel behavioral explanation for the anomalously high first-day returns of initial public offerings (IPOs) and their underperformance in the long run. Our explanation is based on time-varying demand for lottery-like investments, measured in terms of market-wide probability weighting, and our rationale is as follows. As outlined by Barberis and Huang (2008), the positively skewed returns of IPOs are attractive to investors with a preference for lottery-like returns, also referred to as lottery demand. Due to lottery demand, IPOs earn high first-day returns but perform poorly in the long run, tantamount to a correction of the initial overpricing. We account for time variation in lottery demand and find that expected lottery demand significantly predicts high first-day returns as well as a poor long-term performance. Furthermore, we disentangle the pricing effects of expected idiosyncratic skewness and lottery demand and find that skewness significantly predicts IPO returns – as documented by Green and Hwang (2012) – but only if there is market-wide lottery demand to cater for. Our results suggest that institutional investors have a lower preference for skewed returns as they are less likely to exert buying pressure in IPOs, both in the primary and the secondary market.

We follow Dierkes (2013) and estimate Cumulative Prospect Theory's (CPT) probability weighting parameter gamma nonparametrically from S&P 500 index options. During our sample period from 1996 to 2020, we find several episodes of increased probability weighting, for example, during the run-up of the DotCom bubble in the late 1990s or the most recent surge of US stock indices in 2019-2020. Since lower values of gamma imply stronger probability weighting, we use gamma as an inverse proxy for expected lottery demand in the subsequent month. More specifically, gammas below one are tantamount to overweighting small probabilities, consistent with higher lottery demand.

We find expected lottery demand to perform well in explaining both the anomalously high first-day returns and the poor long-term performance of IPOs. With respect to first-day returns, a sample split at $\gamma_{t-1} = 1$ yields a highly significant difference of 14.28 percentage points (t-value = 12.27) between periods of high and low lottery demand (26.76% versus 12.48%). The effect is strongest for firms younger than four years. Our baseline results extend to a regression analysis which controls for various deal, firm, and market characteristics known to affect first-day IPO returns. A one unit *decrease* of gamma – and thus an *increase* in lottery demand – is associated with higher first-day returns of up to 13 percentage points (t-value = −9.42). Net of control variables and industry fixed-effects, this effect reduces to six percentage points but remains statistically significant at the 1%-level (t-value = −3.63).

Given our time series of gammas, we are able to disentangle the effects of skewness and lottery demand and find the interaction to be an important driver of IPO returns. More precisely, we revisit the findings of Green and Hwang (2012) who relate first-day returns to the expected skewness of stock returns within the IPO's industry and also document a positive impact. Compared to Green and Hwang (2012), the performance of expected skewness is not only weaker once we include control variables, but also depends on the prevailing lottery demand regime. Interacting lottery demand and expected skewness highlights that skewness significantly predicts first-day IPO returns, but only if there is lottery demand to cater for. Without lottery demand, the relation between skewness and first-day returns is virtually flat.

Moreover, expected lottery demand explains the poor long-term performance of IPOs. We compare long-term IPO returns to the returns of non-issuers by matching each IPO to the firm with the closest book-to-market ratio in the same size decile. Although matching firms outperform IPOs for several return horizons (one year, three years, and five years), the outperformance is only significant for all three time horizons if the IPO took place during a high lottery demand regime. We thus conclude that IPOs are more likely to become overpriced on the first day of trading if they go public during periods of high lottery demand. Since this overpricing tends to be corrected in the long run, respective IPOs are more likely to underperform matching firms.

Finally, a closer look at the relationship between expected lottery demand and the trading behavior of institutional and individual investors suggests that our results are driven by the differential in the skewness preference between institutional and individual investors. Institutional investors in the primary market only incorporate a fraction of 12.5%-25% of expected lottery demand in their revision of the offer price, whereas the lion's share of the market reaction occurs in the secondary market. Well in line with this finding, expected lottery demand predicts retail buying pressure on the first trading day of IPOs, while being less important for the trading behavior of institutional investors, both measured in terms of the Barber et al. (2009) herding measure. This finding likely explains why expected lottery demand particularly predicts IPO returns in the secondary market and suggests that retail investors indeed have a higher skewness preference than institutional investors. Apart from that, we perform several robustness checks and find that our baseline results hold for alternative definitions of lottery demand regimes and different sub-periods.

Our study contributes to a large and growing body of literature on the underpricing of IPOs and their subsequent long-term performance. Early studies focus on traditional explanations like information asymmetries (Beatty and Ritter, 1986; Rock, 1986), litigation and reputation risk (Tinic, 1988; Lowry and Shu, 2002), and a changing risk composition (Ritter, 1984).[1] Under the changing risk composition hypothesis, riskier IPOs are expected to be more underpriced, equivalent to higher first-day returns. Ljungqvist and Wilhelm (2003) explain the exceptionally high first-day returns in 1999-2000 by a reduced (fractional) CEO ownership and a strongly increased proportion of 'family & friends' shares.[2] Since the latter entitle to purchase shares *at* the offer price, there is an incentive to underprice. Loughran and Ritter (2004) reject this finding together with the changing risk composition hypothesis. Instead, they propose that

[1] Ritter and Welch (2002) later argue that asymmetric information is not the primary determinant of IPO underpricing.

[2] 'Family & friends' shares are, for example, distributed to family members, employees, and suppliers.

return differences trace back to a changing objective function of issuing companies. In the late 1990s, issuers were more willing to accept underpricing due to an increased emphasis on analyst coverage (resulting in an oligopoly of underwriters) and personal brokerage accounts (resulting in incentives to seek underpricing).[3]

Ritter and Welch (2002) suggest a behavioral perspective to explain IPO underpricing. In this direction, Loughran and Ritter (2002) propose a prospect theory approach to explain why issuers agree to leave 'money on the table' (which is tantamount to underpricing). According to prospect theory, issuing firms focus on the *change* in wealth rather than the *level* of wealth, which is in line with prospect theory's reference point dependent valuation (see Kahneman and Tversky, 1979; Tversky and Kahneman, 1992). Therefore, they will aggregate the loss from leaving money on the table and the large gain from an increased valuation of retained shares, resulting in a net profit.[4] As a consequence, issuers only partially adjust offer prices to news and high demand during the book building period.[5] In a more recent study, Loughran and McDonald (2013) show that IPO underpricing is also related to the tone of the S-1 form. By conducting a text mining approach, they find higher first-day returns for issuers whose S-1 form contains more words that are uncertain or negative. Loughran and McDonald (2013) conclude that their results are consistent with

[3] Other related studies focus on regulations (Loughran et al., 1994), internet IPOs (Ofek and Richardson, 2003), and industry peers (Purnanandam and Swaminathan, 2004).

[4] This behavior is likely related to mental accounting introduced by Thaler (1985).

[5] Since the model does not distinguish between public and private information, first-day returns should be predictable based on public information. The partial adjustment phenomenon was first shown by Hanley (1993).

prospect theory as well.[6]

Barberis and Huang (2008) take a closer look at the relation of IPO underpricing and probability weighting, which is one of CPT's key preference components. Investors, on average, overweight small probabilities for large gains, resulting in a preference for lottery-like stocks with highly right-skewed returns. Consequently, Barberis and Huang (2008) predict that IPOs with positively skewed returns earn higher first-day returns but perform poorly in the long run. As the study most closely related to ours, Green and Hwang (2012) provide first evidence in favor of this prediction. They find expected skewness, based on the IPO's industry, to be positively (negatively) related to first-day returns (long-term returns).[7] This result is in line with Boyer et al. (2010) who find a negative impact of idiosyncratic skewness on subsequent returns and Kumar (2009) who shows that individual investors prefer stocks with lottery-like features and underperformance is strongest for those who overweight lottery stocks the most.[8]

We provide additional evidence in favor of behavioral explanations for IPO underpricing and extend the findings of Green and Hwang (2012) in several directions. First, we provide a cleaner test of the predictions in Barberis and Huang (2008) by directly estimating the impact of probability weighting on IPO pricing. Second, the empirical literature usually studies the asset pricing implications of time-varying skewness under

[6] Other behavioral studies also consider sentiment demand (e.g. Ljungqvist et al., 2006) as an alternative mechanism.

[7] Aissia (2014) confirms results for French IPOs. For a seminal study on long-term IPO returns, see Ritter (1991).

[8] See also Bali et al. (2011), Eraker and Ready (2015), and Kumar et al. (2016).

constant preferences for skewness. However, there is ample evidence that skewness preferences (and thus lottery demand) vary over time. Kumar (2009) finds higher gambling demand during economic downturns, whereas Polkovnichenko and Zhao (2013) and Dierkes (2013) show that option-implied skewness preferences exhibit significant time variation. By accounting for this time variation in aggregate lottery demand, we shed further light on the otherwise puzzling episodes of high IPO underpricing documented in Loughran and Ritter (2004).[9] Third, by studying the interaction between expected lottery demand and cross-sectional expected skewness, we are able to emphasize the role of aggregate preferences for skewness as the main driver of the findings in Green and Hwang (2012). High expected lottery demand strongly amplifies the otherwise flat relationship between idiosyncratic skewness and first-day returns, further highlighting the interaction between skewness preferences and asset-specific skewness as an important determinant of IPO returns.

3.2 Data

3.2.1 Options Data

In order to derive a monthly time series of option-implied gammas (and thus lottery demand), we follow the approach of Dierkes (2013) and obtain S&P 500 option prices provided by OptionMetrics. Consistent with the availability of data, our sample includes monthly option expiries

[9] For example, the average IPO return increased from 7% in 1980-1989 to 15% in 1990-1998 and 65% in 1999-2000, before falling back to 12% in 2001-2003.

from February 1996 to December 2020.[10] Data preparation follows the previous literature, most importantly Dierkes (2013) and Ait-Sahalia et al. (2001). Option prices are determined by the average of bid and ask prices and we only keep options with a positive volume. Interest rates to determine physical and risk neutral distributions are linearly inter- or extrapolated and obtained from OptionMetrics as well. We outline the calculation of gamma in more detail in Section 3.3.

3.2.2 IPO Data

Our sample of IPOs is based on the Field-Ritter dataset of IPO founding dates, as used in Field and Karpoff (2002) and Loughran and Ritter (2004).[11] Among other firm characteristics, the Field-Ritter dataset includes offer and founding dates as well as internet and venture capital dummies. In accordance with the availability of options data, we choose a sample period from February 1996 to December 2020. We thus cover several relevant periods: the DotCom bubble of 1999-2000, the financial crisis of 2007-2008, and the 2020 stock market crash due to the Covid-19 pandemic. We obtain offer prices and deal characteristics from the Refinitiv Financial Securities database and keep all IPOs with a share code of 10 or 11 (common stocks) that are covered by the Center for Research in Security Prices (CRSP) within three days.[12] We calculate first-day returns in three steps. If an IPO is covered by CRSP within the first day of trading, the return is determined by the offer price and

[10] An initial month is needed to estimate the first gamma.
[11] For more details, see Appendix A of Loughran and Ritter (2004).
[12] The database was formerly known as the Thomson Financial Securities database.

the CRSP closing price of the same day. If an IPO is *not* covered within the first day, we use first-day returns provided by the Refinitiv Financial Securities database. In case first-day returns are still missing after step 2, we extend step 1 up to the third day of trading. As a result, our final sample comprises 4,673 IPOs. Long-term IPO returns are based on the split-adjusted 252/756/1260 trading days (1/3/5 years) ahead price from CRSP (excluding initial returns).[13] To compare these returns to matching non-issuers, we obtain balance sheet data from Compustat and NYSE size breakpoints from the Kenneth French Data Library.[14] To control our results for expected skewness (as proposed by Green and Hwang, 2012) and coskewness (Harvey and Siddique, 2000), we further obtain industry-sorted portfolio returns from Kenneth French and SIC codes from CRSP. In contrast to Green and Hwang (2012), we do not remove IPOs that are related to the industry 'other'. Lastly, we obtain trade and quotes data from the Trades and Quotes (TAQ) transaction database from March 1996 to December 2000. We thank all fellow researchers for sharing their data.

3.3 Option-implied Lottery Demand

3.3.1 Probability Weighting and Lottery Demand

We derive our measure of lottery demand from the probability weighting parameter gamma in the cumulative prospect theory (CPT) of Tversky and Kahneman (1992). CPT has been widely used to explain the demand of

[13] We choose these return horizons to produce results that are comparable to Green and Hwang (2012).
[14] See French (2022*a*).

individuals for lotteries, which is hard to reconcile with expected utility theory (EUT). Under CPT, the decision maker evaluates the lottery L according to the CPT value

$$CPT(L) = \sum_{i=1}^{m} w_i^- \cdot v(x_i) + \sum_{i=m+1}^{n} w_i^+ \cdot v(x_i), \tag{3.1}$$

with the value function

$$v(x) = \begin{cases} x^\alpha & if\ x \geq 0 \\ -k(-x)^\beta & if\ x < 0, \end{cases} \tag{3.2}$$

where $x := \tilde{x} - x^*$ is the vector of outcomes $\tilde{x}$ of lottery L over the reference value x^*, $\alpha \in (0,1)$ ($\beta \in (0,1)$) is the value sensitivity for gains (losses), and $k > 1$ is the loss aversion parameter. Outcomes x_i are associated with probability p_i ($i = 1,\ldots,n$) and rank-ordered such that $x_1 < \ldots < x_m \leq 0 < x_{m+1} < \ldots < x_n$. The outcome of the value function then enters the CPT value with distorted probability weights according to

$$w_i^- = w^-(\sum_{j=1}^{i} p_j) - w^-(\sum_{j=1}^{i-1} p_j), \tag{3.3}$$

$$w_i^+ = w^+(\sum_{j=i}^{n} p_j) - w^+(\sum_{j=i+1}^{n} p_j). \tag{3.4}$$

To capture the overweighting of small probabilities, Tversky and Kahneman (1992) propose the probability weighting functions

$$w^+(p) = \frac{p^\gamma}{(p^\gamma + (1-p)^\gamma)^{1/\gamma}}, \quad w^-(p) = \frac{p^\delta}{(p^\delta + (1-p)^\delta)^{1/\delta}}. \tag{3.5}$$

For $\gamma \in (0,1)$ and $\delta \in (0,1)$, the probability weighting functions w^+ and w^- imply an overweighting of low probabilities such that $w(p) > p$. Since the weighting functions are monotone, they honor first-order stochastic dominance for $\gamma, \delta \geq 0.28$ only (see Barberis and Huang, 2008). When estimating gamma from option data (see below), we thus winsorize $\gamma < 0.28$ to 0.28.[15] As outlined by Barberis and Huang (2008), who set $\gamma = \delta$, probability weighting has a distinct effect on the demand for lottery-like assets. The relation between gamma and lottery demand is inverse, meaning that a lower value of gamma is tantamount to stronger overweighting of low probabilities and thus higher lottery demand.

IPOs are a perfect example to study the asset pricing implications of aggregate probability weighting. The IPO return distribution is highly positively skewed and offers an IPO investor the small chance of a very large return (see Section 3.4). Thus, IPOs are attractive to lottery investors, become overpriced, and earn low long-term returns (see Barberis and Huang, 2008).

3.3.2 Implications of Probability Weighting for Option Pricing

Our estimation approach closely follows Dierkes (2013) who is the first to introduce this fully nonparametric estimation procedure for probability weighting from option prices. We generalize the discrete version of CPT in Section 3.3.1 to allow for continuous distributions as in, for example, Barberis and Huang (2008) and Polkovnichenko and Zhao (2013). Furthermore, we make the simplifying assumption that the reference point is set to zero. Hence, all outcomes are gains, allowing for a unique

[15] There are eight occurrences in our series of 299 monthly gamma estimates that have virtually no effect on our main results.

identification.[16] We assume a representative agent with monotonically increasing utility function u and monotonically increasing probability weighting function w. Both u and w are assumed to be twice continuously differentiable. This assumption is well in line with the previous literature, for example Barberis and Huang (2008). The representative agent derives utility from an index S that represents the return of the market portfolio. S_t (S_T) refers to the index value today (in the future). We follow the literature and normalize S_t to one (see, for example, Jackwerth, 2000).

Let f_P and f_Q denote the density functions of the data-generating process (or physical measure) and the risk neutral measure with corresponding cumulative distribution functions F_P and F_Q, respectively. Polkovnichenko and Zhao (2013) and Dierkes (2013) formally derive the pricing kernel and show that the risk neutral density with probability weighting is given by

$$f_Q(S_T) = f_P(S_T) \cdot w'\Big(1 - F_P(S_T)\Big) \cdot \beta \frac{u'(S_T)}{u'(S_t)}. \tag{3.6}$$

Equation (3.6) implies that the cumulative physical distribution, $F_P(S_T)$, gets distorted by probability weighting as $1 - w(1 - F_P(S_T))$. It follows that $-\int_0^\infty u(S_T) dw(1 - F_P(S_T))$ is the CPT value over gains (see Barberis and Huang, 2008, p. 2071). Further rearranging of Equation (3.6) yields the pricing kernel under probability weighting

$$\frac{f_Q(S_T)}{f_P(S_T)} = w'\Big(1 - F_P(S_T)\Big) \cdot \beta \frac{u'(S_T)}{u'(S_t)},$$

[16] This assumption does not affect the general estimation of over- or underweighting of small probabilities for extreme outcomes and thus lottery demand. However, it facilitates unique identification since non-zero reference points can impede normalization of the gain distribution due to mixed lotteries.

which varies with the physical distribution, $F_P(S_T)$, if w is not linear for given preferences w and u. Taking derivatives with respect to S_T and rearranging then reveals the impact of probability weighting on risk aversion as

$$\frac{f_P'(S_T)}{f_P(S_T)} - \frac{f_Q'(S_T)}{f_Q(S_T)} = \underbrace{-\frac{u''(S_T)}{u'(S_T)}}_{ARA_u(S_T)} + \underbrace{\left(\frac{w''(1-F_P(S_T))}{w'(1-F_P(S_T))} f_P(S_T)\right)}_{\text{probabilistic risk attitude}}. \tag{3.7}$$

$ARA_u(S_T)$ denotes the absolute risk aversion function across index levels S_T, associated only with the agent's utility function u. The second term $w''(1-F_P(S_T))/w'(1-F_P(S_T)) \cdot f_P(S_T)$ on the right-hand side of Equation (3.7) displays the probabilistic risk attitude. The denominator is always positive due to the strictly increasing weighting function. Equation (3.7) reveals the additional impact of probability weighting on the option-implied risk aversion. With linear w, i.e. no probability weighting, the probabilistic risk attitude part vanishes and Equation (3.7) boils down to a variant well-known from classic economic theory (see e.g. Jackwerth, 2000). With probability weighting, however, observed risk aversion contains a probabilistic risk attitude attributable to the probability weighting function w (in addition to the Arrow-Pratt measure of u). For instance, in the case of an inverse S-shaped probability weighting function w, the probabilistic risk attitude is positive for low wealth levels and steadily decreases until it is negative for high wealth levels. The intuition here is that convex parts of the probability weighting function locally overweight bad states in the economy and thereby increase the observed risk aversion. Concave parts reduce the probabilistic risk attitude due to a local overweighting of favorable states.

We now set out to identify w and u nonparametrically by starting from the equilibrium condition (3.6). For notational convenience, we drop the time index T and summarize constants $u'(S_t)$ and β to a single normalization constant β. Our identification strategy to estimate w and u makes use of the fact that the probabilistic risk attitude associated with w varies with the physical distribution, F_P, whereas the risk attitude associated with u, i.e. $ARA_u(S_T)$, stays constant. To illustrate this, we take Equation (3.6) for the two different physical distributions P_1 and P_2 and rearrange for $u'(S)$

$$\frac{f_{Q_1}(S)}{w'(1-F_{P_1}(S))f_{P_1}(S)\cdot\beta_1} = u'(S), \tag{3.8}$$

$$\frac{f_{Q_2}(S)}{w'(1-F_{P_2}(S))f_{P_2}(S)\cdot\beta_2} = u'(S). \tag{3.9}$$

Obviously, the term $w'(1-F_{P_i}(S))$ changes with P_i, $i \in \{1,2\}$, whereas $u'(S)$ remains constant. Equating the left-hand sides of Equations (3.8) and (3.9) eliminates $u'(S)$ and yields

$$w'(1-F_{P_2}(S)) = \frac{f_{Q_2}(S)}{f_{Q_1}(S)}\frac{f_{P_1}(S)}{f_{P_2}(S)}\frac{\beta_1}{\beta_2}\cdot w'(1-F_{P_1}(S)),\ \forall\ S. \tag{3.10}$$

The only unknown in this equation is the function w, since we can estimate the other quantities from market outcomes. We refer to Appendix 3.A.1 for the estimation of the risk neutral and physical densities f_{Q_i} and f_{P_i}, respectively.

To get an estimate for w, we impose the so-called single crossing assumption on the two distribution functions F_{P_1} and F_{P_2} to make sure that Equation (3.10) constitutes a delay differential equation (DDE) of neutral type.[17] Under the single crossing assumption, $1 - F_{P_2}$ is always above (below) $1 - F_{P_1}$ if S is small (large) and both actually coincide in one point. Taking a closer look at the two distribution functions illustrates how to guarantee this assumption. If P_1 has more mass in the tails than P_2, then for some value $\hat{S}$ it holds $F_{P_1}(S) \geq F_{P_2}(S)$ for all $S \leq \hat{S}$, $F_{P_1}(S) \leq F_{P_2}(S)$ for all $S \geq \hat{S}$, and $F_{P_1}(\hat{S}) = F_{P_2}(\hat{S})$. Thus, with $S < \hat{S}$, we are able to identify $w'(1 - F_{P_2}(S))$ at some 'time point' $1 - F_{P_2}(S)$ because 'time point' $1 - F_{P_1}(S)$ lies in the past and therefore $w'(1 - F_{P_1}(S))$ is already known (and vice versa). Finally, we can solve the DDE for w on the two intervals $[0, \hat{S}]$ and $[\hat{S}, \infty)$. We follow Dierkes (2013) and use different times to maturity to ensure the single crossing property. One advantage is that this allows for a time series of estimated probability weighting functions and thus lottery demand.

Importantly, DDEs require a small *range* of starting values, not a single initial value as is the case with ordinary differential equations. To get a reasonable initial condition, we reconsider the two distributions P_1 and P_2 with their risk neutral counterparts Q_1 and Q_2, respectively. Rearranging the decomposition of aggregate absolute risk aversion in Equation (3.7) yields the following set of equations:

$$\frac{\frac{f'_{P_1}(\hat{S})}{f_{P_1}(\hat{S})} - \frac{f'_{Q_1}(\hat{S})}{f_{Q_1}(\hat{S})} - ARA_u(\hat{S})}{f_{P_1}(\hat{S})} = \frac{w''(1 - F_{P_1}(\hat{S}))}{w'(1 - F_{P_1}(\hat{S}))},$$

[17] DDEs are characterized by the fact that today's derivative of the unknown function depends on the function's behavior in the past. Neutral type means that today's derivative of the unknown function depends on its derivative in the past.

$$\frac{\frac{f'_{P_2}(\hat{S})}{f_{P_2}(\hat{S})} - \frac{f'_{Q_2}(\hat{S})}{f_{Q_2}(\hat{S})} - ARA_u(\hat{S})}{f_{P_2}(\hat{S})} = \frac{w''(1 - F_{P_2}(\hat{S}))}{w'(1 - F_{P_2}(\hat{S}))}.$$

Since the single crossing assumption guarantees that $F_{P_1}(\hat{S}) = F_{P_2}(\hat{S})$ for the state $\hat{S}$, we equate the left hand sides of both equations, rearrange, and obtain

$$ARA_u(\hat{S}) = -\frac{u''(\hat{S})}{u'(\hat{S})} = \frac{\frac{ARA_{M_1}(\hat{S})}{f_{P_1}(\hat{S})} - \frac{ARA_{M_2}(\hat{S})}{f_{P_2}(\hat{S})}}{\frac{1}{f_{P_1}(\hat{S})} - \frac{1}{f_{P_2}(\hat{S})}}, \tag{3.11}$$

where $ARA_{M_i}(S) = f'_{P_i}(S)/f_{P_i}(S) - f'_{Q_i}(S)/f_{Q_i}(S)$, $i \in \{1, 2\}$, is the market's aggregate risk aversion implied by asset prices. Importantly, Equation (3.11) can be solved without knowledge about w or u.

Equation (3.11) gives us a rough estimate of u in the tiny neighborhood of $\hat{S}$. We make a parametric assumption about u for a tiny range around $\hat{S}$, say $[\hat{S} - 0.001, \hat{S} + 0.001]$, to gain an initial condition for our DDE. Then, Equation (3.6) yields an initial condition for w' in Equation (3.10) and, together with $w(0) = 0$ and $w(1) = 1$, we can identify w nonparametrically. We assume that $ARA_u(S) = ARA_u(\hat{S})$, $\forall S \in [\hat{S} - 0.001, \hat{S} + 0.001]$. Thus, absolute risk aversion associated with u is constant in this small range and u is characterized by an exponential utility function. An unreported simulation study reveals that, even if the utility function u is not given by the exponential function, this parametric choice for u on a tiny interval around $\hat{S}$ does not affect our results for the probability weighting function w. There is virtually perfect identification.

3.3.3 Estimating Gamma from Option Prices

Now that we identified the probability weighting function w implied in option prices, we make the transfer from the yet anonymous function w back to CPT and estimate the curvature (i.e. the gamma) of the probability weighting function. More precisely, we fit the linear-in-log-odds probability weighting function $w(p) = \frac{\delta p^\gamma}{\delta p^\gamma + (1-p)^\gamma}$ (see Bleichrodt and Pinto, 2000) by fitting a linear regression to

$$\log\left(\frac{w(p)}{1-w(p)}\right) = \log(\delta) + \gamma \log\left(\frac{p}{1-p}\right) + \epsilon, \tag{3.12}$$

on the interval $p \in \{0.10, 0.11, \ldots, 0.90\}$, where ϵ is the residual.[18] The curvature index gamma corresponds to the estimated slope parameter of this regression. For each month of our sample period from February 1996 to December 2020, we estimate the nonparametric probability weighting function w from Equation (3.10) according to the procedure outlined in Section 3.3.2. Then, we estimate a time series of gammas from Equation (3.12). The time series average (median) of gamma is 0.90 (0.89) with a 95% confidence interval $[0.853, 0.938]$. This estimate is slightly larger than the values reported by psychologists who document values between 0.65 and 0.84 (see Bleichrodt and Pinto, 2000). The fact that our value for gamma is closer to one is not surprising since the estimates come from one of the most liquid and competitive option markets in the world. Thus, we can expect estimates to be closer to EUT. However, the overall probability weighting function exhibits an inverse-S shape, consistent with CPT.

[18] We discard the edges of the interval $(0,1)$ to avoid distortion of the estimates due to noise.

Fig. 3.1: Option-Implied Gamma

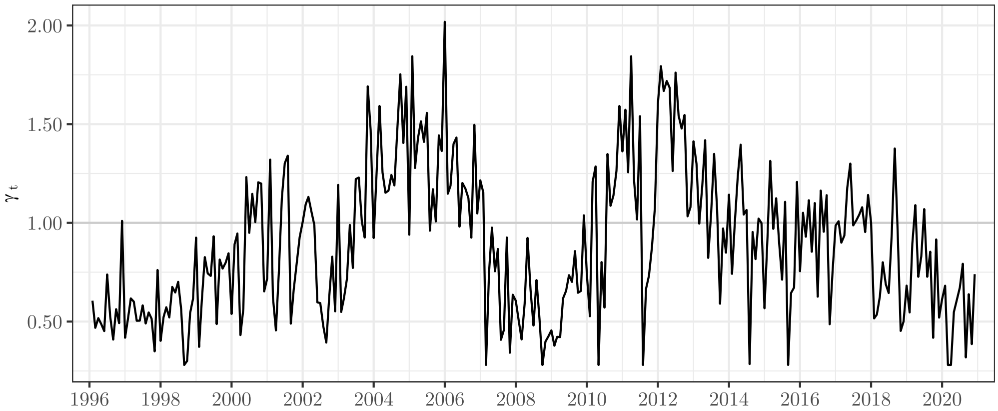

Fig. 3.1 illustrates the probability weighting parameter gamma over the period from February 1996 to December 2020. Gammas are calculated as outlined in Section 3.3.3. According to CPT, gammas below one imply an overweighting of small probabilities, i.e a higher demand for lotteries.

Fig. 3.1 depicts the option-implied probability weighting parameter gamma throughout our sample period. The time series reveals three major periods of $\gamma < 1$ and thus increased lottery demand. The first period is from February 1996 to January 2000 – with a short exception in December 1996 – and largely coincides with the run-up of the DotCom bubble. Since this episode is also referred to as irrational exuberance, our observation of increased lottery demand is to be expected. The next prolonged period of lottery demand largely covers the subprime crisis and lasts from March 2007 to December 2009. The third episode ranges from January 2018 to December 2020 with only three exceptions in September 2018, April 2019, and July 2019. Especially the recent past is characterized by increased lottery demand. The most recent minimum occurred in March 2020 when Covid-19 fully reached global stock markets.

Finally, our time series of gammas serves as a measure of expected lottery demand. More precisely, we define expected lottery demand in month t as gamma in month $t-1$.

$$E_{t-1}\left[\text{Lottery Demand}_t\right] = \gamma_{t-1}. \tag{3.13}$$

Using lagged gammas in the analysis of IPOs prevents a look-ahead bias and guarantees that lottery demand is observable from an ex-ante perspective. Recall that gamma is an inverse measure of lottery demand and lottery demand is particularly pronounced for $\gamma < 1$, as shown in Section 3.3.1.

3.4 Expected Lottery Demand and IPO Returns

3.4.1 Lottery Demand and First-Day Returns

Fig. 3.2 illustrates the distribution of first-day returns within our sample period from February 1996 to December 2020. While most IPOs offer modestly positive returns (as indicated by the median of 7.21%), five percent of IPOs display a return of more than 100% (and up to 697.5%).[19] As a result, the distribution of returns is highly right-skewed with a skewness of roughly 4.77. We thus confirm a lottery-like behavior of IPOs, as predicted by Barberis and Huang (2008), and expect lottery demand to be an important driver of first-day returns.

[19] The highest first-day return within our sample period was recorded for VA Linux in December 1999.

Fig. 3.2: Frequency Distribution of First-Day Returns

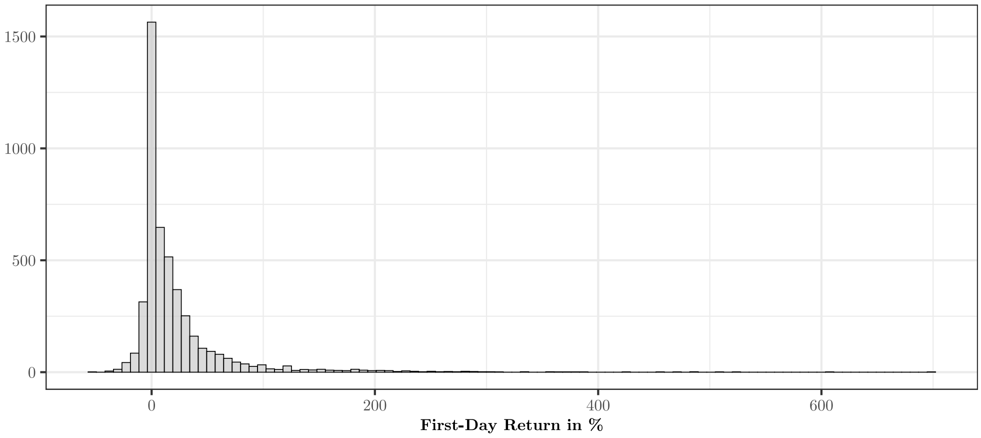

Fig. 3.2 illustrates the frequency distribution of first-day returns (stated in percent) over the period from February 1996 to December 2020.

To take a first look at this prediction, we investigate the lowess-smoothed relationship between first-day returns and expected lottery demand as defined in Equation (3.13).[20] As illustrated by Fig. 3.3, we find a strongly negative (positive) impact of γ_{t-1} (expected lottery demand).[21] Except for a *small* increase of returns between $\gamma_{t-1} = 0.28$ and roughly $\gamma_{t-1} = 0.65$, the relationship is strictly monotonically decreasing. A closer look at the small increase for $\gamma_{t-1} < 0.65$ reveals that this finding is due to the omission of important IPO characteristics in the univariate relationship depicted in Fig. 3.3. In unreported results, we find the relation to be almost exclusively driven by Internet IPOs of the years 1999 and 2000. Consequently, controlling for this effect reestablishes a strictly negative relation between γ_{t-1} and first-day returns. Furthermore, univariate regressions on a subsample with $\gamma_{t-1} \leq 0.65$ reveal that the increase is

[20] For more information on smoothing scatter plots, see Cleveland (1979).
[21] Note that the overall level of first-day returns is rather small since smoothed values are guided by the median instead of average returns.

Fig. 3.3: Expected Lottery Demand and First-Day Returns

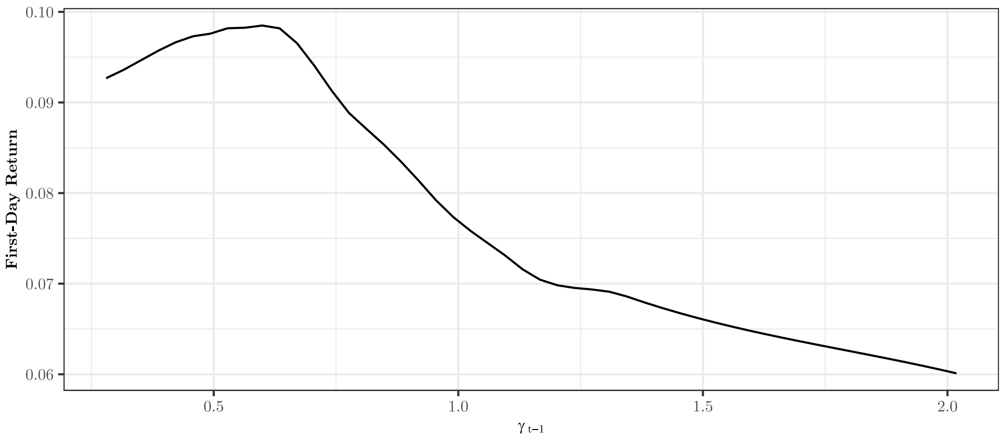

Fig. 3.3 plots the lowess-smoothed relationship between expected lottery demand (γ_{t-1}) and first-day IPO returns. We cover a sample period from March 1996 to December 2020.

not statistically significant ($t = 0.77$), while there is a highly significant negative relationship for $\gamma_{t-1} > 0.65$ ($t = -9.96$). Apart from that, we find the impact of γ_{t-1} to be particularly strong as we move closer to one and there is a salient difference between first-day returns in high ($\gamma_{t-1} \leq 1$) and low ($\gamma_{t-1} > 1$) lottery demand regimes, indicating that overweighting of small probabilities is related to higher first-day returns.

Table 3.1 presents descriptive statistics for first-day returns in low and high lottery demand regimes as well as the full sample. We choose a threshold of $\gamma_{t-1} = 1$ since gammas below one imply an inverse S-shape of the probability weighting function and thus overweighting of small probabilities. Consistent with Fig. 3.3, both mean and median returns significantly increase in lottery demand. While mean (median) returns amount to 12.48% (3.57%) in low lottery demand periods, values increase to 26.76% (9.38%) when lottery demand is high. The difference in returns (14.27 percentage points) is highly significant at the 1% level

Table 3.1: Expected Lottery Demand and IPO Returns: Descriptive Statistics

		Lottery Demand	
	Full Sample	Low	High
Mean	22.12	12.48	26.76
Median	7.21	3.57	9.38
Minimum	−54.10	−41.08	−54.10
Maximum	697.50	212.10	697.50
Volatility	166.15	83.55	191.68

Table 3.1 presents descriptive statistics of first-day IPO returns (in percent), conditional on the expected lottery demand regime. The IPO volatility is annualized. In accordance with CPT, we refer to high lottery demand if the lagged gamma in a given month is below or equal to one, and to low lottery demand otherwise. We cover a sample period from March 1996 to December 2020.

(t = 12.27). This result coincides with Green and Hwang (2012) who report significantly higher mean and median returns when their measure of expected skewness is high. Furthermore, high lottery demand regimes display substantially higher maximum returns (697.50% vs. 212.10%) and maximum losses (41.08% to 54.10%). Lastly, we report the annualized volatility of first-day returns. Consistent with the larger range of returns, high lottery demand regimes display a higher volatility (191.68% vs. 83.55%).[22]

Previous studies (e.g. Ritter, 1991; Loughran and Ritter, 2004) find the age of an issuer to be an important determinant of IPO pricing. Therefore, our next step is to investigate the impact of lottery demand on first-day returns sorted by the age of the issuer. Age classes correspond to Ritter (1991) and results are depicted in Table 3.2. Consistent with Loughran

[22] In unreported results, we find that high lottery demand regimes also exhibit a higher skewness and kurtosis.

Table 3.2: Expected Lottery Demand, IPO Returns, and Firm Age

		Lottery Demand			
Age in Years	Full Sample	Low	High	High – Low	t-value
0 – 1	14.88	3.73	21.29	17.56***	(4.35)
2 – 4	42.55	19.26	48.49	29.23***	(6.65)
5 – 9	27.21	19.21	31.14	11.93***	(4.95)
10 – 19	20.75	15.33	23.18	7.85***	(3.86)
20 – up	12.96	10.82	14.08	3.26**	(2.21)

Table 3.2 presents average first-day IPO returns (in percent), conditional on the expected lottery demand regime and age classes as defined in Ritter (1991). In accordance with CPT, we refer to high lottery demand if the lagged gamma in a given month is below or equal to one, and to low lottery demand otherwise. Stars indicate significance at the 10% (*), 5% (**) and 1% (***) level. We cover a sample period from March 1996 to December 2020.

and Ritter (2004), we find first-day returns to peak for issuers that are 2-4 years old (42.55%). Within this age class, we also find the largest and most significant difference between high (48.49%) and low (19.26%) lottery demand regimes (29.23 percentage points). Beyond that, we find large and highly significant return differences for issuers with an age of up to one year (21.29% vs. 3.73%) and 5-9 years (31.14% vs. 19.21%). However, independent of the specific category, the lottery demand premium is positive and significant for issuers of all age classes. Consistent with earlier results, full sample means are mainly driven by high lottery demand periods.

So far, we have focused on univariate results. However, our measure of expected lottery demand is likely related to other determinants of first-day returns that have been documented in the literature.

Table 3.3: Summary of Control Variables

Variable	Description and expected sign	Source
	Panel A: Firm Characteristics	
Age	We define the *Age* of an issuing firm according to the founding dates provided in the Field-Ritter dataset. Ritter (1991) and Loughran and Ritter (2004) document a negative impact on IPO returns.	Ritter (2022*a*)
Internet	*Internet* is an industry dummy that equals one if an IPO is considered to be an internet company and is zero otherwise. Ofek and Richardson (2003) show that internet firms attain higher first-day returns.	Ritter (2022*a*)
NASDAQ	*NASDAQ* is a dummy variable that takes the value one if an IPO is listed on NASDAQ and is zero otherwise. In line with Loughran and Ritter (2004) and Green and Hwang (2012), we expect a positive coefficient.	Refinitiv
CoSkew	Following Green and Hwang (2012), we calculate the coskewness of an IPO as the average firm-level coskewness over the previous 63 trading days (three months) within the IPO's industry. *CoSkew* is the parameter δ of the regression $$r_{Mkt} - r_{i,t} - r_{f,t} = \alpha + \beta(r_{Mkt,t} - r_{f,t}) + \delta(r_{Mkt,t} - r_{f,t})^2 + \epsilon_{i,t},$$ where $r_{i,t}$ is the return of the IPO's industry at time t, $r_{f,t}$ is the risk-free rate, and $r_{Mkt,t}$ is the market return. $\epsilon_{i,t}$ denotes the residual. Green and Hwang (2012) find no significant effect of *CoSkew*.	CRSP, French (2022*a*), French (2022*b*)
	Panel B: Deal Characteristics	
LN(Proceeds)	*LN(Proceeds)* is the natural logarithm of the total proceeds from an IPO in million dollars and includes over-allotments (see Ritter, 1991). We adjust values to 1997 prices according to the CPI. Proceeds should be positively related to initial returns (see Loughran and Ritter, 2004).	Refinitiv
Price Adjustment	*Price Adjustment* accounts for adjustments between the initial filing range and the offer price. We calculate adjustments as the absolute of the percentage change from the average filing price to the offer price. According to Hanley (1993), price adjustments should be positively related to first-day returns.	Refinitiv
Share Overhang	*Share Overhang* is defined as the natural logarithm of (1 + retained shares / shares offered). To calculate the ratio of retained shares, we make use of the percentage value of shares offered. According to Loughran and McDonald (2013), *Share Overhang* should have a positive impact on IPO returns.	Refinitiv
Pure Primary	*Pure Primary* is an indicator that equals one if an IPO only includes primary shares. IPOs without secondary shares are expected to offer higher first-day returns. However, as Loughran and Ritter (2004) show, the impact is largely driven by the 1999-2000 period.	Refinitiv

Continued from previous page

Variable	Description and expected sign	Source
Top Tier Underwriter	*Top Tier Underwriter* is a dummy variable that is equal to one if the Carter and Manaster (1990) rank of the lead underwriter is greater than or equal to eight and is zero otherwise. The role of underwriter reputation has been explored by several studies, for example Carter et al. (1998), Ljungqvist and Wilhelm (2005), and Loughran and McDonald (2013). We expect a positive coefficient.	Ritter (2022*b*)
Venture Capital	*Venture Capital* is a dummy variable that indicates if an IPO was venture-backed. In general, venture-backed IPOs are assumed to offer higher first-day returns.	Ritter (2022*a*)
	Panel C: Market Characteristics	
Investor Sentiment	Following Green and Hwang (2012), we include *Investor Sentiment*, measured by the Michigan State University Consumer Confidence Index	Univ. of Michigan (2022)
IPO Volatility	*IPO Volatility* is the three-month volatility of first-day returns prior to the IPO. As shown in Table 3.1, high lottery demand regimes display a distinctly higher level of IPO volatility. We thus expect IPO volatility to have a positive impact on first-day returns.	CRSP, Refinitiv
$MktRf_{t-1}$	$MktRf_{t-1}$ is the return of the market portfolio over the previous month. Loughran and Ritter (2004) document that high first-day returns often follow high market returns.	French (2022*a*)
Y(1999-2000), Y(2007-2008)	Y(1999-2000) and Y(2007-2008) are time period indicators that capture the peak of the DotCom bubble (1999-2000) and the subprime crisis (2007-2008).	Refinitiv
	Panel D: Alternative Skewness Measures	
$Skew_{i,t-1}$	Green and Hwang (2012) calculate the expected (idiosyncratic) skewness of IPO i at time t as $$Skew_{i,t-1} = \frac{(P_{99}-P_{50})-(P_{50}-P_1)}{(P_{99}-P_1)},$$ where P_j is the jth percentile of logarithmized monthly returns across *stocks within the IPO's industry* over the last three months prior to the month of the IPO. We denote expected skewness by $Skew_{i,t-1}$ since it is based on ex-ante available data. Given the findings in Green and Hwang (2012), the expected coefficient is positive.	French (2022*a*)
Market $Skew_{i,t-1}$	We calculate expected market skewness ($MarketSkew_{i,t-1}$) according to $$MarketSkew_{i,t-1} = \frac{(P_{99}-P_{50})-(P_{50}-P_1)}{(P_{99}-P_1)},$$ where P_k is the kth percentile of logarithmized monthly returns across *all stocks* over the last three months prior to the month of the IPO.	French (2022*a*)

To account for these, we follow Green and Hwang (2012) and estimate a regression framework that includes several firm, deal, and market characteristics as well as industry fixed effects.[23] Table 3.3 presents both the included control variables and their expected sign. The dependent variable of all specifications is the first-day return.

Table 3.4 presents results, where t-values in parentheses are based on White (1980) standard errors. In Model (1), we start with a univariate regression of first-day returns on γ_{t-1} (expected lottery demand) and find a highly significant, negative (positive) relationship ($t = -9.42$). An increase of gamma by one unit reduces first-day returns by economically important 13 percentage points. After including firm and deal characteristics in Model (2), statistical significance is only slightly reduced ($t = -8.46$) and the negative (positive) impact of γ_{t-1} (lottery demand) on first-day returns becomes even stronger with a coefficient estimate of −0.14. Furthermore, we find all control variables to be significant at the 1%-level and, except for *CoSkew*, signs correspond to our expectations. However, the negative sign of *CoSkew* coincides with the results of Carter et al. (2011).[24] In Model (3), we add market characteristics and industry fixed effects. With a t-value of −3.63 and a coefficient estimate of −0.06, expected lottery demand continues to be economically and statistically significant at the 1%-level. Not surprisingly and consistent with Table 3.1, we find the IPO volatility ($t = 5.31$) to be a major cause of the reduced statistical signif-

[23] As our measure of lottery demand is not based on the IPO's industry, we exclude industry characteristics. Regressions that include industry fixed effects are estimated without a constant.

[24] Note that the coskewness measure of Carter et al. (2011) is based on the returns of previous IPOs instead of industry returns.

Table 3.4: Expected Lottery Demand and IPO Returns: Regression Approach

	Dependent variable:				
	First-Day Return				
	(1)	(2)	(3)	(4)	(5)
γ_{t-1}	-0.13^{***}	-0.14^{***}	-0.06^{***}		
	(−9.42)	(−8.46)	(−3.63)		
$Skew_{t-1}$				0.10^{**}	
				(2.47)	
$MarketSkew_{t-1}$					0.01
					(0.14)
Age		-0.001^{***}	-0.001^{***}	-0.001^{***}	-0.001^{***}
		(−5.99)	(−4.83)	(−5.06)	(−5.01)
CoSkew		-0.003^{***}	-0.002^{***}	-0.001^{**}	-0.002^{***}
		(−4.14)	(−3.02)	(−2.39)	(−2.70)
Internet		0.37^{***}	0.22^{***}	0.22^{***}	0.22^{***}
		(10.40)	(6.41)	(6.25)	(6.29)
NASDAQ		0.07^{***}	0.05^{***}	0.04^{***}	0.05^{***}
		(5.73)	(3.44)	(3.28)	(3.43)
LN(Proceeds)		0.06^{***}	0.06^{***}	0.05^{***}	0.05^{***}
		(7.41)	(7.33)	(6.99)	(6.95)
Price Adjustment		0.37^{***}	0.27^{***}	0.27^{***}	0.27^{***}
		(4.13)	(3.21)	(3.16)	(3.15)
Share Overhang		0.45^{***}	0.35^{***}	0.36^{***}	0.35^{***}
		(9.81)	(8.22)	(8.24)	(8.18)
Pure Primary		0.07^{***}	0.05^{***}	0.05^{***}	0.05^{***}
		(5.13)	(3.70)	(3.79)	(3.75)
Venture Capital		0.07^{***}	0.06^{***}	0.07^{***}	0.06^{***}
		(4.29)	(3.98)	(4.00)	(3.94)
Top Tier Underwriter		0.05^{***}	0.03^{**}	0.04^{**}	0.04^{**}
		(3.07)	(2.16)	(2.41)	(2.38)
$MktRf_{t-1}$			0.74^{***}	0.75^{***}	0.74^{***}
			(3.61)	(3.66)	(3.57)
IPO Volatility			0.42^{***}	0.43^{***}	0.45^{***}
			(5.31)	(5.46)	(5.50)
Investor Sentiment			0.0002	0.001^{**}	0.001
			(0.40)	(1.97)	(1.48)
Y(1999-2000)			0.02	−0.01	−0.01
			(0.27)	(−0.09)	(−0.16)
Y(2007-2008)			-0.04^{*}	−0.01	−0.02
			(−1.80)	(−0.68)	(−1.04)
Industry fixed effects	No	No	Yes	Yes	Yes
Adjusted R^2	0.01	0.19	0.41	0.41	0.41

Table 3.4 presents results for OLS regressions of first-day IPO returns on our measure of expected lottery demand (γ_{t-1}), expected skewness ($Skew_{t-1}$), expected market skewness ($MarketSkew_{t-1}$), and several control variables motivated by the literature. Stars indicate significance at the 10% (*), 5% (**) and 1% (***) level and t-values (in parentheses) are based on White (1980) heteroscedasticity-consistent standard errors. We cover a sample period from March 1996 to December 2020 and do not report the intercept which is included in Models (1) and (2).

icance.[25] Interestingly, we do not find a significant impact of investor sentiment. This is, in part, related to the Covid-19 pandemic. While sentiment was very low, average IPO returns in 2020 have been quite high (38.4%). In contrast, the same period displayed gammas around 0.5 and thus high lottery demand.

To compare the explanatory power of expected lottery demand with other skewness measures, we replace gamma by the expected (idiosyncratic) skewness, $Skew_{t-1}$, and the expected market skewness, $MarketSkew_{t-1}$, respectively. In contrast to Green and Hwang (2012), our estimate for $Skew_{t-1}$ (Model 4) is only significant at the 5%-level. This finding is attributable to the considerable correlation between $Skew_{t-1}$ and the IPO volatility ($\rho = 23.8\%$) such that *IPO volatility* captures at least some of *Skew*'s explanatory power. After excluding *IPO volatility* from Model (4), $Skew_{t-1}$ turns out to be significant at the 1% level ($t = 4.48$), as expected from the findings in Green and Hwang (2012). Moreover, without control variables (not reported), *Skew*'s explanatory power is comparable to that of γ_{t-1}. With respect to $MarketSkew_{t-1}$ (Model 5), we do not find any explanatory power. Again, we attribute this finding to the (even higher) correlation with the IPO volatility (33.3%).

In unreported results, we further investigate the impact of expected lottery demand on money left on the table.[26] In a regression framework similar to that of Table 3.4, we find lottery demand to be significant at the 1% level. A one unit increase of γ_{t-1} is associated with a 27.1 million

[25] Excluding *IPO volatility* from Model (3) increases the significance of γ_{t-1} to $t = -6.15$.

[26] Money left on the table is defined as the difference between the first closing price and the offer price, multiplied by the number of shares sold.

dollar reduction of money left on the table. To put that in perspective: the average amount of money left on the table in our sample period is roughly 25 million dollars. However, this result is not surprising as the amount of money left on the table largely depends on first-day returns.

In summary, we find expected lottery demand to be an important determinant of IPO returns. First-day returns strongly increase in lottery demand (even after adjusting for control variables) and return differences between high and low lottery demand regimes are large and highly significant.

3.4.2 Disentangling Lottery Demand and Skewness

Now we turn to our key contribution. Considering that the rationale behind our measure of expected lottery demand (i.e. investors' lottery preferences) and expected skewness (i.e. the IPO-specific lottery characteristics) is very similar, it is important to disentangle their impact on first-day returns. Therefore, we first perform a two-way sort based on skewness terciles and lottery demand regimes.[27] Results are depicted in Table 3.5.

With respect to the full sample, the return difference between high and low expected skewness amounts to 13.91 percentage points and is very similar to the value reported in Green and Hwang (2012). Moreover, with a t-value of 7.71, the difference is significant at the 1%-level. In high lottery demand regimes, the return difference substantially increases to

[27] Following Green and Hwang (2012), we define tercile membership based on data that was available at the time of the IPO. Therefore, in each month IPOs are ranked in ascending order based on expected skewness and assigned into terciles.

Table 3.5: Expected Lottery Demand, Expected Skewness and IPO Returns: Two-Way Sort

		Lottery Demand			
	Full Sample	Low	High	High – Low	t-value
Low Skew	15.84	12.35	17.60	5.25***	(3.53)
Med Skew	21.79	13.00	26.59	13.58***	(6.80)
High Skew	29.75	12.00	36.69	24.69***	(9.94)
High – Low	13.91***	−0.35	19.09***		
t-value	(7.71)	(−0.22)	(7.87)		

Table 3.5 presents first-day IPO returns (in percent), conditional on expected lottery demand and expected skewness as defined in Green and Hwang (2012). In accordance with CPT, we refer to high lottery demand if the lagged gamma in a given month is below or equal to one, and to low lottery demand otherwise. Following Green and Hwang (2012), in each month IPOs are ranked in ascending order based on expected skewness and assigned into terciles. Classification of expected skewness is based on full sample terciles. Stars indicate significance at the 10% (*), 5% (**) and 1% (***) level. We cover a sample period from March 1996 to December 2020.

19.09 percentage points with a t-value of 7.87. During low lottery demand periods, however, there is no skewness premium at all. In contrast, the return difference between high and low lottery demand regimes is positive and significant at the 1%-level for all skewness terciles. While the return difference is 5.25 percentage points ($t = 3.53$) across low skewness IPOs, the premium increases to 24.69 percentage points ($t = 9.94$) when skewness is high.

As of yet, our results indicate that the explanatory power of expected skewness largely depends on the respective lottery demand regime. We further investigate this finding in a regression framework that incorporates interaction terms. To simplify interpretations, we replace γ_{t-1} by an indicator variable, $Lottery\ Demand_{t-1}$, which takes on a value of one

Table 3.6: Expected Lottery Demand, Expected Skewness and IPO Returns: Interaction Regressions

	Dependent variable:				
	First-Day Return				
	(1)	(2)	(3)	(4)	(5)
$Lottery\ Demand_{t-1}$	0.14***	0.14***	0.15***	0.11***	0.06***
	(12.27)	(12.02)	(12.43)	(8.95)	(4.81)
$Skew_{t-1}$		0.37***	0.01	0.09**	0.08*
		(8.75)	(0.28)	(2.06)	(1.96)
$Lottery\ Demand_{t-1} \times Skew_{t-1}$			0.53***	0.24***	0.14*
			(7.72)	(3.20)	(1.90)
Firm characteristics	No	No	No	Yes	Yes
Deal characteristics	No	No	No	Yes	Yes
Market characteristics	No	No	No	No	Yes
Industry fixed effects	No	No	No	Yes	Yes
Adjusted R^2	0.02	0.04	0.04	0.19	0.40

Table 3.6 presents results for OLS regressions of first-day IPO returns on an expected lottery demand indicator ($Lottery\ Demand_{t-1}$), expected skewness ($Skew_{t-1}$), and an interaction term. The lottery demand indicator is equal to one if the lagged gamma in a given month is below or equal to one and zero otherwise. Control variables and industry fixed effects correspond to Table 3.4. However, in Model (5), we drop *IPO Volatility* due to the large correlation with $Skew_{t-1}$, as shown in Section 3.4. Stars indicate significance at the 10% (*), 5% (**) and 1% (***) level and t-values (in parentheses) are based on White (1980) heteroscedasticity-consistent standard errors. We cover a sample period from March 1996 to December 2020 and do not report the intercept which is included in Models (1) to (3).

if $\gamma_{t-1} \leq 1$ and is zero otherwise. $Skew_{t-1}$ corresponds to the expected skewness as described above. Results are depicted in Table 3.6. In Model (1), we start with a univariate regression of first-day returns on the lottery demand dummy. By construction, the coefficient estimate and t-value of $Lottery\ Demand_{t-1}$ correspond to the average first-day returns described in Table 3.1. In Model (2), we further include expected skewness (but do not include an interaction term). Both measures have a positive and highly significant impact on first-day returns. While the coefficient and t-value of $Lottery\ Demand_{t-1}$ are very similar to Model (1), $Skew_{t-1}$ displays a coefficient estimate of 0.37 with a t-value of 8.75. Thus, increasing $Skew_{t-1}$ by one unit increases first-day returns by 37 percentage points. However, note that the difference between *Skew*'s 10% and 90% percentile is only 0.43, compared to 0.95 for gamma. In Model (3), we additionally include an interaction term. In line with Table 3.5, we find expected skewness to be insignificant, while the interaction term ($Lottery\ Demand_{t-1} \times Skew_{t-1}$) is highly significant, both statistically (t = 7.72) and economically. In Model (4), we control for firm and deal characteristics as well as industry fixed effects. Now, $Skew_{t-1}$ is significant at the 5%-level (t = 2.06), while $Lottery\ Demand_{t-1}$ (t = 8.95) and the interaction term remain significant at the 1%-level (t = 3.20). In Model (5), we also add market controls.[28] Most importantly, even though the significance of the interaction term is reduced, we still find it to be almost significant at the 5%-level (t = 1.90), while $Lottery\ Demand_{t-1}$ remains significant at the 1%-level (t = 4.81).

We conclude that the impact of expected lottery demand is different from expected skewness. Moreover, we find the *interaction* between investors' lottery demand and assets' lottery characteristics to be an important determinant of IPO returns.[29]

[28] We drop *IPO volatility* due to the large correlation with $Skew_{t-1}$.

[29] In unreported results, we replace $Skew_{t-1}$ by the expected market skewness. Except for a reduced explanatory power of lottery demand in low *MarketSkew* periods, our conclusions remain the same.

3.4.3 Lottery Demand and Long-Term Performance

In stark contrast to the high first-day returns documented in the previous section, IPOs often perform poorly in the long run, as documented by, for example, Ritter (1991) and Ritter and Welch (2002). In line with the predictions of Barberis and Huang (2008) and the results in Green and Hwang (2012), this suggests that IPOs are initially overpriced due to their lottery-like payoffs. If high first-day returns, attributable to high lottery demand regimes, represent an overpriced first-day valuation, we expect a stronger underperformance of IPOs that took place during a high lottery demand regime.

We start our long-run analysis by estimating the lowess-smoothed relationship between expected lottery demand and long-term returns. Fig. 3.4 illustrates results for 1-year, 3-year, and 5-year returns, where one year is defined by 252 trading days and initial returns are excluded. Mirroring the relationship between lottery demand and first-day returns,

Fig. 3.4: Expected Lottery Demand and Long-Term Returns

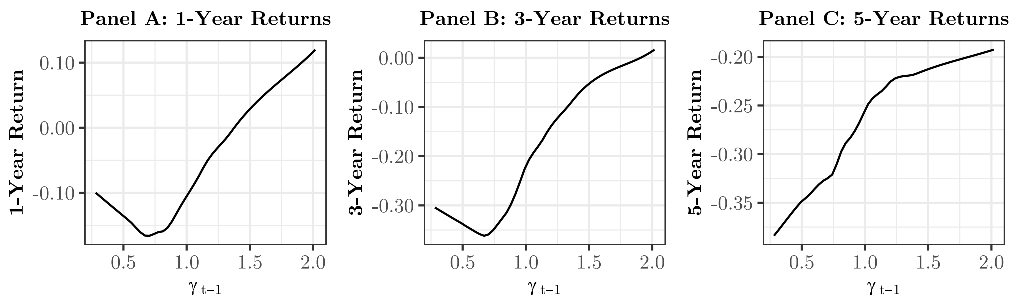

Fig. 3.4 plots the lowess-smoothed relationship between expected lottery demand (γ_{t-1}) and long-term IPO returns. Following Green and Hwang (2012), we report 1-year (252 tradings days), 3-year (756 trading days), and 5-year (1260 trading days) returns, excluding the first-day return. We cover a sample period from March 1996 to December 2020.

1-year (Panel A) and 3-year returns (Panel B) display a small decrease between $\gamma_{t-1} = 0.28$ and $\gamma_{t-1} = 0.70$, while there is a strictly positive and distinct relationship for gammas above 0.70. Again, we relate this finding to Internet IPOs of the years 1999 and 2000, which not only displayed extraordinary high first-day returns (90%), but also very low long-term returns. For example, the average (median) 3-year return over the full sample amounts to 12% (−25%), whereas DotCom IPOs offered an average (median) return of only −83% (−94%). Moreover, compared to other IPOs of the years 1999 and 2000, Internet IPOs less often survived for five years or longer (42.6% vs. 70.8%). As a result, their impact on the smoothed relationship declines and 5-year returns (Panel C) are almost monotonically increasing in gamma, except for a very small interval around $\gamma_{t-1} = 0.70$. Noteworthy, the return difference between high and low gammas (low and high lottery demand) is substantial for each of the three holding periods.

To confirm this result, we follow Ritter (1991), Ritter and Welch (2002), and Green and Hwang (2012) by comparing long-term IPO returns to the returns of matching non-issuers. Thereby, we adopt the matching procedure of Green and Hwang (2012). More precisely, we match each IPO to a non-issuing firm (with share code 10 or 11) with the closest book-to-market ratio in the same size decile that has been listed on AMEX, NASDAQ, or NYSE for at least five years. To account for delistings, we drop non-issuers as soon as less than five more years of data are available and only keep returns that are covered by both the matching firm and

Table 3.7: Expected Lottery Demand and the Long-Run Performance of IPOs

Lottery Demand	1 Year	3 Years	5 Years
Low			
IPO	9.74	20.54	29.32
Matching	12.25	33.18	43.21
Difference	−2.51	−12.64**	−13.89*
t-value	(−1.08)	(−2.34)	(−1.79)
High			
IPO	2.63	11.65	14.93
Matching	14.63	32.19	43.36
Difference	−12.00***	−20.54***	−28.42***
t-value	(−4.48)	(−3.16)	(−4.60)

Table 3.7 compares long-term IPO returns (in percent) to returns of non-issuing matching firms. We follow Green and Hwang (2012) and define a matching firm as the stock with the closest book-to-market ratio in the same size decile. Size breakpoints are based on NYSE breakpoints. The size of an IPO is defined as the market value after the first trading day. Book-to-market ratios are based on book values provided by Compustat. Non-issuing firms are defined as firms that have been listed on NYSE, AMEX, or NASDAQ (with share code 10 or 11) for more than 5 years (1260 trading days). We remove all non-issuers that have been delisted within 5 years after the IPO and report results for 1-year (252 trading days), 3-year (756 tradings days), and 5-year (1260 trading days) returns, excluding the first-day return. In accordance with CPT, we refer to high lottery demand if the lagged gamma in a given month is below or equal to one, and to low lottery demand otherwise. Stars indicate significance at the 10% (*), 5% (**) and 1% (***) level and *t*-values are based on White (1980) heteroscedasticity-consistent standard errors.

the IPO.[30] As a result, our new sample covers IPOs from 1996 to 2015. In accordance with the previous literature, size breakpoints are based on NYSE breakpoints and first-day returns are excluded. We calculate the market value of an IPO according to the first closing price, while market

[30] If an IPO firm is delisted during a year, we follow Green and Hwang (2012) and calculate the return for the remainder of the respective year by compounding the CRSP value-weighted market return.

values of matching firms are based on the previous month. Depending on data availability, book-to-market ratios of IPOs are either based on the current or the following year.[31] For matching firms, we employ the most recent available book value. After the matching procedure, 3,505 IPOs (75% of the initial sample) remain.[32] Table 3.7 presents results.

Most importantly, IPOs underperform matching firms throughout all holding periods and differences strongly increase when lottery demand at the time of the IPO was high. For low lottery demand regimes, differences vary from −2.51 (1 year) to −13.89 percentage points (5 years) and are only significant at the 5%-level when considering 3-year holding periods. In contrast, underperformance increases to 12.00 (1 year), 20.54 (3 years), and 28.42 percentage points (5 years) when lottery demand at the time of the IPO was high. Importantly, all differences are substantially larger and highly significant at the 1%-level.[33]

In summary, our results provide further evidence that market-wide lottery demand causes IPOs to become overpriced, resulting in a poor long-term performance.

[31] We drop IPOs for which no book value is available by the end of the following year.

[32] 777 IPOs (or 16.6%) of the initial IPO sample could not be matched due to missing book values.

[33] We derive similar results when considering 2-year or 4-year return horizons. In unreported results, we also compare median differences and find them to strongly increase when lottery demand at the time of the IPO was high (while increments for low lottery demand periods are rather moderate).

3.5 Dissecting the Asset Pricing Implications of Lottery Demand

3.5.1 Primary versus Secondary Markets

The high first-day returns of IPOs suggest that the offer price, which is determined in the primary market, does not fully incorporate the asset pricing implications of skewness. Green and Hwang (2012) point out two alternative explanations. Either institutional investors in the primary market have a lower skewness preference, or institutional investors don't recognize that retail investors are willing to pay higher prices for highly skewed securities.[34] In this section, we use our proxy for expected lottery demand to distinguish between these possible explanations.

As a first step in this regard, we quantify the extent to which each market segment incorporates lottery demand in the IPO process. To measure the market reaction in the primary market, we use the revision of the offer price (OP) relative to the midpoint (Mid) of the filing range, $(OP - Mid)/Mid$, as the dependent variable. The combined market reaction (primary and secondary) is measured as $(Close - Mid)/Mid$, where $Close$ is the first-day closing price as used for the first-day returns. Replacing the first-day returns with these two measures and comparing the coefficient estimates then allows us to quantify the differential impact of expected lottery demand in secondary markets.[35] Table 3.8 presents

[34] We thank an anonymous referee for highlighting this point.
[35] Loughran and Ritter (2002) use a similar setup to quantify the incorporation of market returns 15 days prior to the offer date.

Table 3.8: Primary versus Secondary Market Adjustment to Lottery Demand

	Dependent variable:			
	$\frac{OP-Mid}{Mid}$		$\frac{Close-Mid}{Mid}$	
	(1)	(2)	(3)	(4)
γ_{t-1}	-0.02^{***}	-0.02^{**}	-0.16^{***}	-0.08^{***}
	(−4.20)	(−2.47)	(−9.13)	(−3.70)
Firm characteristics	No	Yes	No	Yes
Deal characteristics	No	Yes	No	Yes
Market characteristics	No	Yes	No	Yes
Industry fixed effects	No	Yes	No	Yes
Adjusted R^2	0.003	0.16	0.01	0.40

Table 3.8 presents OLS regressions of the primary market return (defined as the percentage change from the average filing price to the offer price) on our measure of expected lottery demand (γ_{t-1}) and control variables introduced in Table 3.4 (except for *Price Adjustment*). We match lottery demand in two ways: γ_{t-1} at the time of the IPO (Models 1 and 2) and γ_{t-1} at the filing date (Models 3 and 4). Stars indicate significance at the 10% (*), 5% (**) and 1% (***) level and t-values (in parentheses) are based on White (1980) heteroscedasticity-consistent standard errors. We cover a sample period from March 1996 to December 2020 and do not report the intercept.

results. Except for the new dependent variables and the exclusion of *Price Adjustment*, the setup is identical to Table 3.4.

Interestingly, expected lottery demand has an impact in both market segments as the coefficient estimates on γ_{t-1} are highly significant in statistical terms throughout Models (1) to (4). However, the impact of lottery demand is stronger in the secondary market. Comparing the coefficient estimates (−0.02 vs. −0.16) shows that the primary market incorporates only roughly 12.5% of expected lottery demand into the offer price. This share increases to roughly 25% (−0.02 vs. −0.08) after accounting for control variables. Thus, the lion's share of the market reaction occurs in the secondary market, in line with the conclusions of Green and Hwang (2012).

3.5.2 Institutional versus Retail Investors

There are two potential explanations for the stronger effect of lottery demand in the secondary market (as documented in the previous section). First, there are countervailing forces in the bargaining process in the primary market. As the profits due to high first-day returns in the secondary market come at the expense of the issuing company and its pre-issue shareholders, issuers prefer high offer prices and reject IPO underpricing. From an underwriter's perspective, IPO underpricing is an indirect source of compensation, because it makes it easier to find buyers and thereby reduces marketing costs. This might explain why underwriters do not completely adjust the offer price to public information, e.g. indications of high demand (see Loughran and Ritter, 2002). However, there are also several incentives for underpricing from an issuer's perspective, e.g. side payments via personal brokerage accounts (Loughran and Ritter, 2004), 'family and friends shares' (Ljungqvist and Wilhelm, 2003), and an aggregation of money left on the table with gains from an increased valuation of retained shares (Loughran and Ritter, 2002).

Second, the institutional investors who determine the offer price are simply less likely to account for skewness in their valuation of IPOs or refrain from investing due to risk considerations. As the IPO market is subject to the aforementioned institutional features which might mute or at least mitigate institutional lottery demand's impact on the offer price, we study the trading behavior of both investor groups in the secondary market during the IPO. Similar to Green and Hwang (2012), we use buyer-initiated trades from the Trades and Quotes (TAQ) transaction database and extend the scope of our analysis to the trading behavior of institutional investors. For this explanation to hold, we expect that lottery demand predicts retail buying pressure in IPOs, while being less important for institutional trading.

We adopt the Barber et al. (2009) herding measure, which relies on the proportion of buyer-initiated trades in the TAQ database. To sign a trade as buyer-initiated, we follow Barber et al. (2009) and apply the Lee and Ready (1991) algorithm to NYSE stocks and the Ellis et al. (2000) classification rule to NASDAQ stocks. Next, we match the TAQ data with our IPO sample and keep all matches with up to three days between the IPO date and the first trading date in the TAQ database. However, due to the widespread introduction of decimalization in 2001 and the growth in algorithmic trading, we cover a subset from March 1996 to December 2000.[36] We are able to match 1,921 out of 2,085 IPOs.

After matching the data and inferring the trade direction, we follow Green and Hwang (2012) and use trades with a size below \$10,000 as a proxy for retail investor trades. In line with Barber et al. (2009), trades with a total size above \$50,000 serve as a measure for institutional trading. For stock i on day t, we then compute $p_{i,t}$ as the proportion of signed trades which is buyer-initiated and use $p_{i,t}$ to compute the herding measure $HM_{i,t}$ as

$$HM_{i,t} = \left|p_{i,t} - E[p_{i,t}]\right| - E\left[\left|p_{i,t} - E[p_{i,t}]\right|\right], \tag{3.14}$$

[36] Both mechanisms undermine the distinction between individual and institutional trading based on trade size (see Barber et al., 2009; Cao et al., 2020). Our restriction of the sample period is standard in this strand of the literature.

where $E[p_{i,t}]$ is the proportion of all purchases on day t and $E\left[\left|p_{i,t} - E[p_{i,t}]\right|\right]$ accounts for the higher variation in the proportion of purchases in stocks with fewer trades.[37] We compute $HM_{i,t}$ separately for small and large trades, referred to as HM_{Retail} and $HM_{Inst.}$, respectively. The correlation between both measures is $\rho = 34\%$.

To analyze the relationship between trading behavior and lottery demand, we regress HM_{Retail} and $HM_{Inst.}$ on expected lottery demand (γ_{t-1}), expected skewness ($Skew_{t-1}$), and control variables motivated by Green and Hwang (2012)'s Table 6: *CoSkew*, *Internet*, *Age*, *NASDAQ*, *LN(Proceeds)*, $MktRf_{t-1}$, and *Investor Sentiment*. Table 3.9 presents results.

Models (1) and (2) analyze the relationship between expected lottery demand (γ_{t-1}) and retail buying pressure (HM_{Retail}). The results are well in line with our hypothesis as lottery demand significantly increases retail trading in IPOs at the 1% level. In Model (1), the coefficient estimate of γ_{t-1} is −0.04 with a t-value of −4.52. After including control variables in Model (2), the coefficient estimate of γ_{t-1} increases to −0.03, but remains economically and statistically significant ($t = -2.77$). Models (3) and (4) extend the analysis to institutional trading. While Model (3) reveals a similar relationship between γ_{t-1} and trading activity, the effect is completely captured by control variables in Model (4). Thus, after accounting for deal and market characteristics, expected lottery demand predicts retail buying pressure on the first day of trading, but not buying pressure exerted by institutional investors. A closer look at the control variables

[37] In line with Barber et al. (2009), we compute proportions in terms of the value of the trade.

Table 3.9: Impact of Lottery Demand on IPO Trades

	Dependent variable:			
	HM_{Retail}		$HM_{Inst.}$	
	(1)	(2)	(3)	(4)
γ_{t-1}	−0.04***	−0.03***	−0.06***	0.005
	(−4.52)	(−2.77)	(−4.30)	(0.38)
$Skew_{t-1}$		−0.01		−0.01
		(−0.95)		(−0.47)
CoSkew		0.0001		0.001**
		(0.31)		(2.10)
Internet		0.01		−0.03***
		(1.21)		(−5.12)
Age		0.0002		0.001***
		(1.16)		(3.71)
NASDAQ		−0.07***		−0.10***
		(−10.21)		(−11.01)
LN(Proceeds)		0.01**		−0.02***
		(2.47)		(−4.47)
$MktRf_{t-1}$		0.09*		−0.12*
		(1.67)		(−1.91)
Investor Sentiment		−0.0001		−0.001
		(−0.20)		(−1.02)
Adjusted R^2	0.01	0.11	0.01	0.17

Table 3.9 presents OLS regressions of the Barber et al. (2009) herding measure, HM_i (i = Retail, Inst.), on expected lottery demand (γ_{t-1}), expected skewness ($Skew_{t-1}$), and several control variables that are motivated by Green and Hwang (2012)'s Table 6. We use trades with a size below (above) \$10,000 (\$50,000) as a proxy for retail (institutional) investor trades. Stars indicate significance at the 10% (*), 5% (**) and 1% (***) level and t-values (in parentheses) are based on White (1980) heteroscedasticity-consistent standard errors. We cover a sample period from March 1996 to December 2000 and do not report the intercept.

facilitates the understanding of this result. The most striking differences are the changing signs on $Internet$, Age and $LN(Proceeds)$ from Model (2) to Model (4). Buying pressure from institutional investors is significantly lower for Internet IPOs and increases in the age of the firm, whereas both variables have no explanatory power for retail trading.[38] Moreover,

[38] Adding the controls sequentially reveals that $Internet$ and Age have the largest impact.

HM_{Retail} is significantly larger for IPOs with higher proceeds, whereas the opposite holds true for $HM_{Inst.}$. In stark contrast to Green and Hwang (2012), $Skew_{t-1}$ is insignificant in both cases, which is likely explained by a different definition of retail trading activity. Our results illustrate that the secondary market of IPOs exhibits a stronger pricing reaction to expected lottery demand and that this reaction is primarily driven by buying pressure from retail investors.[39]

3.6 Robustness Checks

3.6.1 Alternative Sample Splits

With respect to CPT, splitting the sample at $\gamma_{t-1} = 1$ is an intuitive way to separate high and low lottery demand regimes. Nevertheless, there might be a small area around $\gamma_{t-1} = 1$ where probability weighting is moderate and investors still act rather neutral.[40] We thus introduce an alternative sample split that assigns gammas between 0.9 and 1.1 to medium lottery demand. Gammas below 0.9 and above 1.1 are still considered as high and low lottery demand, respectively. In Appendix 3.A.2, we report results. For the sake of brevity, we focus on the most important findings.

Table 3.A.1 reports descriptive statistics of first-day returns. All measures behave monotonically, and when lottery demand his high, mean

[39] In unreported results, we extend our scope to the overall US stock market (as covered by CRSP) and evaluate the change of institutional ownership in lottery-like stocks, i.e. stocks which are identified along the lines of, for example, Bali et al. (2011) and Kumar et al. (2016). We find that institutional investors tend to relatively reduce positions in lottery stocks significantly during periods of high lottery demand. This is consistent with gamma being closely related to retail trading behavior in the secondary market.

[40] Furthermore, gamma is an estimate which is subject to estimation error.

(median) returns exceed respective values from the two-way split. With respect to the age-sort (Table 3.A.2), we derive very similar conclusions. For all age classes, returns monotonically increase in lottery demand and return differences between high and low lottery demand regimes slightly exceed those from the two-way split. Table 3.A.3 presents results for the alternative sort based on expected skewness and lottery demand regimes. Except for low skewness IPOs, first-day returns monotonically increase in lottery demand. Moreover, the deviation is very small and high–low differences increase for both expected skewness and expected lottery demand. A potential concern with respect to our alternative sample split is that thresholds are chosen arbitrarily. We therefore repeat the above analysis with alternative thresholds (for example 0.80/1.20 and 0.95/1.05, not reported) and derive very similar conclusions.[41]

We conclude that it is reasonable to define a medium lottery demand regime in which investors act rather neutral. First-day returns monotonically increase in lottery demand and results hold for alternative thresholds.

3.6.2 Sub-periods

Finally, to make sure that our results are not exclusively driven by the exceptionally high first-day returns in 1999-2000, we perform a simple sub-period test by cutting our sample period in half. The first sub-sample comprises IPOs from 1996 to 2008 and includes the DotCom bubble,

[41] For some specifications, first-day returns in medium and low lottery demand regimes converge. However, we still find a significant lottery demand premium.

Fig. 3.5: Lottery Demand Regimes and Sub-period Returns

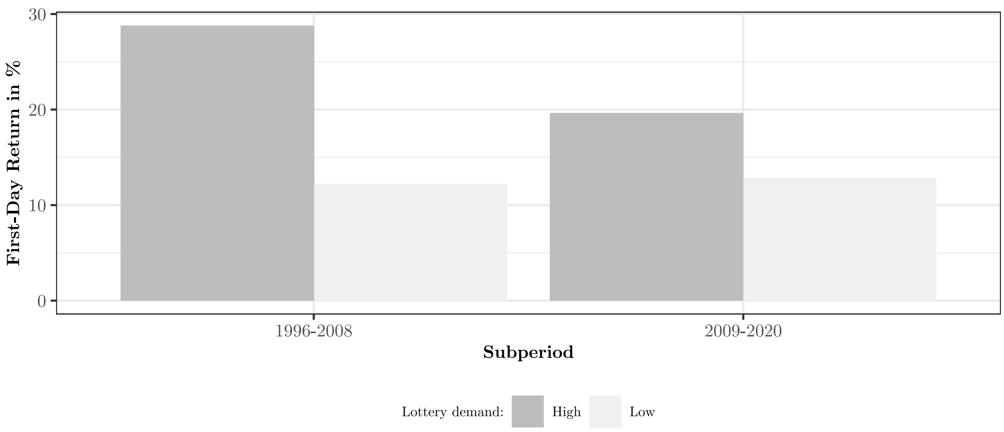

Fig. 3.5 plots average first-day IPO returns in sub-periods, conditional on the expected lottery demand regime in the month of the IPO. Thereby, we simply split our sample period in half. In accordance with Cumulative Prospect Theory, we refer to high lottery demand if the lagged gamma in a given month is below or equal to one and to low lottery demand otherwise.

the burst of the DotCom bubble, and the subprime crisis. The second sub-sample contains IPOs from 2009 to 2020 and covers the long-running bull market from 2009 to 2019 as well as the Covid-19 crash in 2020. Furthermore, by choosing this threshold, we ensure that sub-samples contain both high and low lottery demand regimes.[42] Fig. 3.5 illustrates results by comparing average first-day returns in low and high lottery demand regimes.

In the first sub-sample, average returns amount to 28.80% when expected lottery demand is high and reduce to 12.20% when lottery demand is low. We thus yield a highly significant return difference of 16.60 percentage points ($t = 11.29$). In 2009-2020, low lottery demand returns are comparable to the first sub-sample (12.83%), while returns in high lottery

[42] See Fig. 3.1 and Section 3.3.3.

demand regimes amount to 19.63%. The return difference of 6.79 percentage points is significant at the 1%-level ($t = 4.19$). We therefore conclude that our results are not particularly driven by the DotCom bubble and extent to more recent IPOs.

3.7 Concluding Remarks

Our results provide further evidence for the asset pricing implications of skewness and lottery demand. More precisely, lottery-like assets may become overpriced and earn low subsequent returns. We are the first to account for time variation in market-wide lottery demand in the context of IPO pricing. IPOs issued in periods of high lottery demand earn higher first-day returns and are more likely to perform poorly over return horizons of up to five years after the IPO. Most importantly, we show that the explanatory power of expected skewness strongly depends on lottery demand regimes. Thus, first-day IPO returns are particularly driven by the interaction of market-wide lottery demand and asset-specific lottery characteristics. Moreover, we find institutional investors to have a lower propensity to exert buying pressure on the IPO day, which likely explains why expected lottery demand particularly predicts returns in the secondary market. Our findings complement the existing literature, most importantly Barberis and Huang (2008) and Green and Hwang (2012), and suggest that buying pressure due to high lottery demand is an important determinant of IPO pricing.

3.A Appendix

3.A.1 Estimation of Risk Neutral and Physical Densities

The estimation technique for both the risk neutral and the physical distributions follows Dierkes (2013). We briefly summarize the approach here. For preparation of the OptionMetrics options data, we refer to Section 3.2.1. To abstract from dividend payments $D_{t,\tau}$ between time t and the expiry date $t+\tau$, we base our results on the option-implied future. More specifically, we infer the futures price $F_{t,\tau} = S_t e^{r_{t,\tau} - D_{t,\tau}}$ from the put-call parity. After calculating the future value, we exclude in-the-money options and use the more liquid at-the-money and out-of-the-money options for further calculations.

In general, the estimation of risk neutral densities follows the standard in the literature, most importantly Ait-Sahalia and Lo (1998). We fix times to maturity to 15 and 30 calendar dates in order to obtain two different risk neutral distributions for each expiry date (which is necessary for the identification strategy outlined in 3.3.2). Our calculation of risk neutral densities relies on Breeden and Litzenberger (1978). Since this requires a fine grid of implied volatilities $\hat{\sigma}(K/F_{t,\tau})$, we start with the estimation of nonparametric implied volatilities. For a given time to maturity t and expiry date $t+\tau$, we collect all S&P 500 options that have the same expiration date and whose time to maturity deviates by at most seven calendar days. For each option i out of these n options, we collect the Black-Scholes implied volatility σ_i, the strike price K_i, the time to maturity τ_i, and the implied future value $F_{t,\tau}$. We then apply the Nadaraya-Watson kernel estimator to obtain

$$\hat{\sigma}(K/F_{t,\tau}) = \frac{\sum_{i=1}^{n} k_{K/F_{t,\tau}}\left(\frac{K/F_{t,\tau}-K_i/F_{t_i,\tau_i}}{h_{K/F_{t,\tau}}}\right)k_\tau\left(\frac{\tau-\tau_i}{h_\tau}\right)\sigma_i}{\sum_{i=1}^{n} k_{K/F_{t,\tau}}\left(\frac{K/F_{t,\tau}-K_i/F_{t_i,\tau_i}}{h_{K/F_{t,\tau}}}\right)k_\tau\left(\frac{\tau-\tau_i}{h_\tau}\right)}, \tag{3.A.1}$$

with the kernel functions

$$k_{K/F_{t,\tau}}(x) = \frac{1}{2\pi}exp\Big(-\frac{x^2}{2}\Big), \tag{3.A.2}$$

$$k_\tau(x) == \frac{1}{2\pi}exp\Big(-\frac{x^2}{2}\Big). \tag{3.A.3}$$

The bandwidths $h_{K/F_{t,\tau}}$ and h_τ are chosen according to Ait-Sahalia and Lo (1998). The implied volatilities from Equation (3.A.1) are then used to compute call prices along a grid of equidistant strike prices $K = K_1, ..., K_n$. The resulting call prices $C_{t,\tau}(S_t, K, \tau, r_{t,\tau}, D_{t,\tau})$ enter the formula of Breeden and Litzenberger (1978) for the risk neutral density f_Q

$$f_Q(S_T) = e^{-r_{t,\tau}} \frac{\delta^2 C_{t,\tau}(S_t, K, \tau, r_{t,\tau}, D_{t,\tau})}{\delta K^2}\Bigg|_{K=S_T}. \tag{3.A.4}$$

The estimation of the physical counterpart of the return distribution follows the bootstrap approach of Kliger and Levy (2009). We collect daily returns of the S&P 500 from four years prior to each expiry, until the expiry, as motivated from Jackwerth (2000). Then, we sample returns with replacement and compound returns to match the trading days until expiry of the options. This approach contains a tiny proportion of forward looking data to account for the notion of rational expectations.[43] We set the mean of the compounded returns to the sample average of returns over the future values.[44] Finally, we estimate the physical distribution

[43] Our results are robust to the exclusion of forward looking returns.

[44] Recall that we abstract from dividend payments by using returns over the futures value, similar to Polkovnichenko and Zhao (2013).

from a kernel density estimate with Gaussian kernel and the Silverman (1986) rule for the bandwidths over 10,000 of such bootstrapped samples.

3.A.2 Alternative Sample Splits

Table 3.A.1: Expected Lottery Demand and IPO Returns: Descriptive Statistics (Alternative Sort)

		Lottery Demand		
	Full Sample	Low	Medium	High
Mean	22.12	12.60	16.50	28.00
Median	7.21	3.13	5.03	10.00
Minimum	−54.10	−33.08	−41.08	−54.10
Maximum	697.50	212.10	292.19	697.50
Volatility	166.15	86.52	121.40	198.31

Table 3.A.1 presents descriptive statistics of first-day IPO returns (in percent), conditional on the expected lottery demand regime. The IPO volatility is annualized. In contrast to Table 3.1, we now consider three lottery demand regimes. We refer to high (low) lottery demand if the lagged gamma in a given month is below or equal to 0.9 (above 1.1). Medium lottery demand regimes are defined by gammas between 0.9 and 1.1. We cover a sample period from March 1996 to December 2020.

Table 3.A.2: Expected Lottery Demand, IPO Returns, and Firm Age (Alternative Sort)

		Lottery Demand				
Age in Years	Full Sample	Low	Medium	High	High – Low	t-value
0 – 1	14.88	3.34	8.24	22.71	19.37***	(4.34)
2 – 4	42.55	20.96	24.10	50.83	29.87***	(5.95)
5 – 9	27.21	19.94	22.97	31.88	11.94***	(4.45)
10 – 19	20.75	15.51	18.64	23.50	7.99***	(3.50)
20 – up	12.96	10.21	12.69	14.38	4.17**	(2.47)

Table 3.A.2 presents average first-day IPO returns (in percent), conditional on the expected lottery demand regime and age classes as defined in Ritter (1991). In contrast to Table 3.2, we now consider three lottery demand regimes. We refer to high (low) lottery demand if the lagged gamma in a given month is below or equal to 0.9 (above 1.1). Medium lottery demand regimes are defined by gammas between 0.9 and 1.1. Stars indicate significance on the 10% (*), 5% (**) and 1% (***) level. We cover a sample period from March 1996 to December 2020.

Table 3.A.3: Expected Lottery Demand, Expected Skewness and IPO Returns: Two-Way Sort (Alternative Sort)

		Lottery Demand				
	Full Sample	Low	Med	High	High – Low	t-value
Low Skew	15.84	12.08	11.90	18.40	6.33***	(3.75)
Med Skew	21.79	13.33	16.01	27.95	14.62***	(6.43)
High Skew	29.75	12.40	21.62	39.45	27.05***	(9.43)
High – Low	13.91***	0.32	9.71***	21.04***		
t-value	(7.71)	(0.18)	(3.11)	(7.53)		

Table 3.A.3 presents first-day IPO returns (in percent), conditional on expected lottery demand and expected skewness as defined in Green and Hwang (2012). In contrast to Table 3.5, we now consider three lottery demand regimes. We refer to high (low) lottery demand if the lagged gamma in a given month is below or equal to 0.9 (above 1.1). Medium lottery demand regimes are defined by gammas between 0.9 and 1.1. Following Green and Hwang (2012), in each month, IPOs are ranked in ascending order based on expected skewness and assigned into terciles. Stars indicate significance on the 10% (*), 5% (**) and 1% (***) level. We cover a sample period from March 1996 to December 2020.

Chapter 4

Volatility-Dependent Probability Weighting and the Dynamics of the Pricing Kernel Puzzle

This chapter refers to the working paper:

Dierkes, Maik, Jan Krupski, Sebastian Schroen and Philipp Sibbertsen (2022): 'Volatility-Dependent Probability Weighting and the Dynamics of the Pricing Kernel Puzzle', Working Paper, Leibniz Universität Hannover.

Abstract

We obtain risk neutral and physical densities from the Pan (2002) stochastic volatility and jumps model to estimate volatility-dependent probability weighting functions. Across volatility levels, we find pronounced inverse S-shapes and probability weighting almost monotonically increases in volatility, indicating higher skewness preferences in volatile market environments. Moreover, by estimating probabilistic risk attitudes, we shed further light on the pricing kernel puzzle. In line with economic theory, we find pricing kernels, net of probability weighting, to be monotonically decreasing. Equivalently, we find risk aversion to be positive across wealth levels. Our results are robust to a variety of adjustments.

Keywords: Stochastic Volatility, Probability Weighting, Pricing Kernel Puzzle

JEL: G11, G14, G41.

4.1 Introduction

According to Jackwerth (2000), risk neutral probabilities are tantamount to the product of physical probabilities and a risk aversion adjustment. The pricing kernel, defined as the ratio of risk neutral and physical probabilities, is expected to monotonically decrease in wealth and distinctly reflects risk aversion. However, several studies find U-shaped pricing kernels (the pricing kernel puzzle) or, equivalently, negative episodes of risk aversion functions (the risk aversion puzzle). We attribute this finding to investors who overweight small probabilities for tail events and therefore distort the pricing kernel. Moreover, it has long been suggested that time-varying risk aversion or, put differently, a time-varying price of risk, is key to understand asset prices. For example, Fama (2014) notes that both risk and investors' risk aversion are likely to change over time, resulting in a time-varying equity premium. In line with this, we refer time variation in pricing kernels and risk aversion to a volatility-dependent and hence time-varying degree of probability weighting.

Our study thus contributes to two strands of literature: *time-varying risk preferences* and the *pricing kernel puzzle*. First, we obtain risk neutral and physical densities from the Pan (2002) stochastic volatility and jumps model and find a strikingly robust relationship between volatility and cumulative prospect theory (CPT)'s probability weighting parameter gamma (see Tversky and Kahneman, 1992). Across volatility levels, the probability weighting function exhibits an inverse S-shape, i.e. small (large) probabilities are overweighted (underweighted) and gamma (probability weighting) almost monotonically decreases (increases) in volatility, suggesting that skewness preferences are more pronounced in volatile markets. Second, by adjusting model-implied pricing kernels and risk aversion functions for probability weighting, we shed further light on the pricing kernel puzzle. More specifically, we estimate probabilistic risk attitudes, equivalent to the share of risk aversion related to probability weighting. While pricing kernels estimated from the Pan (2002) model exhibit the typical U-shape documented in the literature, pricing kernels net of probability weighting are strictly monotonically decreasing and therefore in line with economic theory. As a direct result, risk aversion functions are positive throughout wealth levels.

In a seminal study, Campbell and Cochrane (1999) propose a habit-formation model with slowly moving external habits and find both risk aversion and marginal utility to countercyclically depend on the business cycle. Moreover, the model explains several asset pricing phenomena, including the procyclical (countercyclical) variation of stock prices (volatility). Brandt and Wang (2003) extend the habit-formation model by including a process for aggregate risk aversion and also find risk preferences to vary. They conclude their results to be consistent with both an agent irrationally fearing unexpected inflation and an economy with heterogeneous preferences where risk aversion varies with the cross-sectional distribution of wealth. In a more recent study, Guiso et al. (2018) analyze portfolio data and repeated surveys of Italian bank clients in 2007 and 2009. After the financial crisis, they find both qualitative and quantitative measures of risk aversion to increase substantially. As potential mechanisms behind their findings, the authors suggest fear of losses and overweighting of salient payoffs. Hence, it seems reasonable to explain time-varying risk preferences from a behavioral perspective. In this sense, Barberis et al. (2001) propose a prospect theory framework in which, in contrast to Campbell and Cochrane (1999), changes in risk aversion are caused by changes in the level of the stock market. They find that, after recent run-ups, agents are less risk averse because prior gains cushion subsequent losses. In comparison to consumption-based models, the level of risk aversion is smaller but still explains several market characteristics. While Barberis et al. (2001) restrict their model to reference-point dependent valuation and loss aversion, several recent studies highlight the importance of CPT's probability weighting component. Kliger and Levy (2009) assess the performance of expected utility (EUT), rank-dependent expected utility (RDEU), and CPT models and find that probability weighting functions exhibit a pronounced inverse S-shape.[1] Moreover, when including probability weighting, the model fit improves substantially. Polkovnichenko and Zhao (2013) and Dierkes et al. (2022) estimate option-implied probability weighting functions and find them to substantially vary over time. Notably, variation is not erratic, but systematic. For example, Kilka and Weber (2001) find in the lab that probability weighting is more pronounced when agents are less confident in assessing the uncertainty of a decision situation – much like in a high volatility regime, especially when volatility drives up jump

[1] See also Camerer and Ho (1994), Tversky and Fox (1995), Wu and Gonzalez (1996), Gonzalez and Wu (1999), Abdellaoui (2000), and Bleichrodt and Pinto (2000).

intensity. In line with this finding, Liu et al. (2005) propose that risk and rare events should have an impact on risk preferences and Chabi-Yo et al. (2008) relate changes in preferences and beliefs to regime shifts in state variables.[2] Moreover, Gao et al. (2021) show that investors dislike high-skewness securities in low volatility regimes, while Polkovnichenko and Zhao (2013) note that periods with less inverse S-shaped probability weighting functions tend to coincide with these regimes.

We capture these findings by estimating probability weighting functions from the Pan (2002) stochastic volatility and jumps model. We choose this model as it offers the advantage that, in addition to the wealth level, it includes the volatility as an additional state variable which we can change counterfactually.[3] Following Ziegler (2007), we first obtain risk neutral and physical densities for a wide range of volatilities. Thereafter, we follow Dierkes et al. (2022) and employ these densities to estimate the probability weighting parameter gamma for any given volatility. Even though the Pan (2002) model was never designed to match CPT preferences, our results are strikingly robust and correspond to earlier studies. In our main specification, we normalize the return horizon to one year and find gamma (probability weighting) to almost monotonically decrease (increase) in volatility. For example, with the two-parameter specification of Prelec (1998), gammas vary from roughly 0.99 for very low volatilities to 0.70 for high volatilities. Results for the two-parameter linear-in-log-odds (0.90 to 0.68) and the one-parameter Tversky and Kahneman (1992)

[2] See also Bliss and Panigirtzoglou (2004) and Brown and Jackwerth (2012).
[3] Nevertheless, we find our model-based results to hold in a fully nonparametric setting.

function (0.95 to 0.82) are similar. Most importantly, we find the average probability weighting function over volatilities to display a pronounced inverse S-shape. These findings are robust to alternative return horizons (three months and six months) and a nonparametric empirical setting.

As risk aversion is closely connected to the pricing kernel, we can directly transfer our estimation approach to the pricing kernel puzzle. In a seminal study, Jackwerth (2000) recovers risk aversion from risk neutral and physical probabilities, estimated via S&P 500 options and stock returns, respectively.[4] While he finds risk aversion to be positive and decreasing in wealth prior to the 1987 stock market crash, risk aversion is partially negative and increasing in the post-crash era. Among others, the puzzle has been confirmed by Ait-Sahalia and Lo (2000) and Rosenberg and Engle (2002). Moreover, Beare and Schmidt (2016) and Golubev et al. (2014) perform statistical tests and reject pricing kernel monotonicity for the S&P 500 and the German DAX, respectively.[5] In the recent past, several studies proposed possible solutions to the pricing kernel puzzle. For example, Bakshi et al. (2010) assume heterogeneity among investors, with pessimists short selling the market portfolio and thus driving increases in the pricing kernel. Ziegler (2007) estimates risk aversion functions from the Pan (2002) model and finds them to be monotonically decreasing

[4] To obtain densities, Jackwerth (2000) employs a variation of Jackwerth and Rubinstein (1996)'s approach.

[5] Bliss and Panigirtzoglou (2004) find risk aversion estimates to be positive. However, they restrict the pricing kernel by assuming power or exponential utility functions. Linn et al. (2018) argue that prior pricing kernel estimates are inconsistent because they compare forward looking risk neutral densities to backward looking physical densities. While they find a monotonically decreasing pricing kernel, Cuesdeanu and Jackwerth (2018) attribute this result to their specific estimation procedure.

but negative for gains (implying an increasing pricing kernel). Although assuming heterogeneous investors might solve the problem, the degree of heterogeneity would need to be implausibly large. Further possible solutions include state-dependence in fundamentals (Chabi-Yo et al., 2008) and the inclusion of higher moment preferences (Chabi-Yo, 2012; Cuesdeanu and Jackwerth, 2018). Hens and Reichlin (2013) show that if at least one of the three standard assumptions (market completeness, risk aversion, correct beliefs) is violated, the pricing kernel may have increasing parts. Most importantly, they find the combination of distorted beliefs (i.e. probability weighting) and misestimation of probabilities to be a possible solution.[6] Hence, it appears reasonable that applying behavioral insights may solve the pricing kernel puzzle.

In this sense, Baele et al. (2019) develop an asset pricing model with CPT preferences (based on Barberis et al., 2001) and find the implied CPT pricing kernel to display a pronounced U-shape (implying partially negative risk aversion functions). In line with Barberis et al. (2016), they conclude that the key driver of their results is the probability weighting component. Polkovnichenko and Zhao (2013) and Dierkes et al. (2022) estimate pricing kernels to study the time variation in probability weighting functions. While their results are generally consistent with a U-shaped pricing kernel, they rather focus on the time variation in probability weighting and its asset pricing implications.

By adjusting model-implied pricing kernels and risk aversion functions for probability weighting, we shed further light on the role of proba-

[6] However, they need to assume a slightly negative expected mean return.

bility weighting as an important driver of the pricing kernel puzzle. Since Benzoni et al. (2011) and Babaoğlu et al. (2018) find pricing kernels to be variance-dependent, we again employ the Pan (2002) stochastic volatility and jumps model and estimate probabilistic risk attitudes, equivalent to the share of risk aversion related to probability weighting.[7] We add to the results of previous studies in two ways. First, we provide direct measures of pricing kernels and risk aversion functions, and second, we explicitly relate the pricing kernel puzzle to the probabilistic risk attitude. Before accounting for probability weighting, we find the average pricing kernel to exhibit a strong U-shape, implying episodes of negative risk aversion (consistent with Ziegler, 2007). However, since the probabilistic risk attitude is strikingly close to the absolute risk aversion estimated from Pan (2002), the adjusted risk aversion function is consistently positive and the corresponding pricing kernel is monotonically decreasing in wealth. Our results are robust to alternative return horizons (three and six months), wealth percentiles, an alternative functional assumption, a numerical approach to estimate the probabilistic risk attitude, and variations of the Pan (2002) coefficient estimates. We therefore conclude that probability weighting intensifies in volatile market environments and plays an important role in explaining the pricing kernel puzzle.

[7] Further studies that relate the pricing kernel and risk aversion to volatility are Christoffersen et al. (2013), Song and Xiu (2016), and Linn et al. (2018).

4.2 Methodology

4.2.1 Estimation of Probability Weights

Our framework closely follows Dierkes et al. (2022) who introduce a fully nonparametric estimation procedure to derive time-varying probability weighting functions. We assume a representative agent with monotonically increasing and twice continuously differentiable utility function u and probability weighting function w.[8] Moreover, we assume the agent to derive utility from the market return, denoted by S. Then, under EUT (i.e. with linear w) it holds

$$f_Q(S_T) = f_P(S_T) \cdot \beta \frac{u'(S_T)}{u'(S_t)}, \tag{4.1}$$

where $f_Q(S_T)$ and $f_P(S_T)$ denote the risk neutral and physical densities with corresponding distribution functions F_Q and F_P.[9] Furthermore, $u'(S_T)$ and $u'(S_t)$ are marginal utilities with respect to the future and current stock price. β is a normalizing constant. By solving Equation (4.1) for $u'(S_T)$ and calculating $u''(S_T)$, we yield the absolute risk aversion associated with the index level S_T

$$ARA(S_T) = -\frac{u''(S_T)}{u'(S_T)} = \frac{f_P'(S_T)}{f_P(S_T)} - \frac{f_Q'(S_T)}{f_Q(S_T)}, \tag{4.2}$$

which corresponds to Jackwerth (2000). However, with probability weighting, and thus non-linear w, the physical distribution function adjusts to $F_{\tilde{P}}(S_T) = 1 - w(1 - F_P(S_T))$ and the corresponding density function is

[8] This assumption is in line with Barberis and Huang (2008). See Tversky and Kahneman (1992) for more details on probability weighting.

[9] See, for example, Chapman and Polkovnichenko (2009) and Ross (2015). In accordance with Jackwerth (2000), we normalize S_t to one.

$f_{\tilde{P}}(S_T) = f_P(S_T) \cdot w'(1 - F_P(S_T))$. As a result, the risk neutral density and the pricing kernel, net of probability weighting, are given by

$$f_Q(S_T) = \underbrace{f_P(S_T) \cdot w'(1 - F_P(S_T))}_{f_{\tilde{P}}(S_T)} \cdot \beta \frac{u'(S_T)}{u'(S_t)}, \tag{4.3}$$

and

$$\frac{f_Q(S_T)}{f_P(S_T) \cdot w'(1 - F_P(S_T))} = \beta \frac{u'(S_T)}{u'(S_t)}, \tag{4.4}$$

respectively. Note that the adjusted pricing kernel varies with the physical distribution, $F_P(S_T)$, if w is not linear. In case of linear w, however, Equation (4.4) collapses to the standard pricing kernel. In analogy with the steps from Equation (4.1) to Equation (4.2), we then yield the absolute risk aversion with probability weighting as

$$\underbrace{\frac{f_P'(S_T)}{f_P(S_T)} - \frac{f_Q'(S_T)}{f_Q(S_T)}}_{ARA(S_T)} = \underbrace{-\frac{u''(S_T)}{u'(S_T)}}_{ARA_u(S_T)} + \underbrace{\frac{w''(1 - F_P(S_T))}{w'(1 - F_P(S_T))} f_P(S_T)}_{ARA_w(S_T)}, \tag{4.5}$$

where $ARA_u(S_T)$ denotes the absolute risk aversion after accounting for probability weighting, i.e. the level of risk aversion only associated with the representative agent's utility function u. In contrast, $ARA_w(S_T)$ describes the *probabilistic risk attitude* and reflects the level of risk aversion originating from the probability weighting function w.[10] Equation (4.5) also reveals that without probability weighting, the probabilistic risk attitude becomes zero and absolute risk aversion boils down to Equation (4.2). However, in case of the typically observed inverse S-shaped probability weighting function, the probabilistic risk attitude is positive (and

[10] By assuming an increasing probability weighting function, the denominator of the probabilistic risk attitude is always positive. See also Polkovnichenko and Zhao (2013) and Quiggin (1993).

decreasing) for low wealth levels and becomes negative for high wealth levels. Thus, convex episodes of the probability weighting function locally overweight bad states in the economy and increase the observed risk aversion (probabilistic risk attitude > 0), while concave parts reduce the risk aversion (probabilistic risk attitude < 0).

In order to identify u and w, we utilize the fact that the probabilistic risk attitude varies with the physical distribution, $F_P(S_T)$, while the risk aversion associated with u, $ARA_u(S_T)$, remains constant. More precisely, we assume two different physical distributions, $F_{P_1}(S_T)$ and $F_{P_2}(S_T)$, solve Equation (4.3) for $u'(S_T)$ and merge β and $u'(S_t)$ to a single normalizing constant β. After dropping the time index T for notional convenience, we yield

$$\frac{f_{Q_1}(S)}{w'(1-F_{P_1}(S))f_{P_1}(S)\cdot\beta_1} = u'(S), \tag{4.6}$$

$$\frac{f_{Q_2}(S)}{w'(1-F_{P_2}(S))f_{P_2}(S)\cdot\beta_2} = u'(S). \tag{4.7}$$

By equating the left hand sides of Equations (4.6) and (4.7) , we eliminate $u'(S)$ and obtain

$$w'(1-F_{P_2}(S)) = \frac{f_{Q_2}(S)}{f_{Q_1}(S)}\frac{f_{P_1}(S)}{f_{P_2}(S)}\frac{\beta_1}{\beta_2}\cdot w'(1-F_{P_1}(S)),\ \forall\ S. \tag{4.8}$$

Since we are able to estimate risk neutral and physical densities from the Pan (2002) stochastic volatility and jumps model (see Section 4.2.2), Equation (4.8) leaves w' as the only unknown. We impose the following assumption (single-crossing assumption) which allows the estimation of w', or equivalently w. Suppose that P_1 has more mass in the tails than P_2, such that for some value $\hat{S}$ it holds

$$F_{P_1}(S) \geq F_{P_2}(S) \; \forall \; S \leq \hat{S},$$
$$F_{P_1}(S) \leq F_{P_2}(S) \; \forall \; S \geq \hat{S},$$
$$F_{P_1}(\hat{S}) = F_{P_2}(\hat{S}).$$

Then, Equation (4.8) constitutes two Delay Differential Equations (DDE) of neutral type, one DDE for all $S \leq \hat{S}$ and one DDE for all $S \geq \hat{S}$. More precisely, on both subsets, today's derivative of the yet unknown function w depends on its derivative in the past. For example, we are able to identify $w'(1 - F_{P_2}(S))$ at 'time point' $1 - F_{P_2}(S)$ when 'time point' $1 - F_{P_1}(S)$ lies in the past and $w'(1 - F_{P_1}(S))$ is already known (and vice versa). Consequently, we can solve the DDE for w' on the two intervals $[0, \hat{S}]$ and $[\hat{S}, \infty)$. And finally, with $w(0) = 0$ and $w(1) = 1$, we identify w. However, while Dierkes et al. (2022) estimate risk neutral and physical densities from empirically observed option prices, we analytically calculate densities in the Pan (2002) model (see Section 4.2.2). In Section 4.2.3, we provide more details on the estimation of the DDE and outline differences with respect to Dierkes et al. (2022).

4.2.2 The Pan (2002) Model

We obtain risk neutral and physical densities from the Pan (2002) stochastic volatility and jumps model as it offers the advantage that, besides the wealth level, it includes the volatility as an additional state variable, which we can change counterfactually. Note that our analyses are difficult to execute in a nonparametric setup, since – for any given cross-section of option prices – we would not be able to change the volatility state counterfactually without the model framework. At the same time, the Pan (2002) model is rich enough to explain relevant characteristics of S&P 500 index returns and options written on them. Moreover, it provides closed-form expressions for the transforms of both $f_P(S_T)$ and $f_Q(S_T)$, making it appealing for our estimation technique. Given that Polkovnichenko and Zhao (2013) casually observe S-shaped probability weighting functions during times of low volatility and Kilka and Weber (2001) find probability weighting to be more pronounced when agents are less confident in assessing a decision situation, we expect volatility to be an important determinant of probability weighting.[11]

Pan (2002) fits an elaborate model, based on Bates (2000), to S&P 500 option prices and time series of the underlying. More specifically, she proposes a model with stochastic volatility and jumps in the underlying's price process, where jump intensity is correlated with the current level of volatility. The model determines three risk premia: a diffusive (Brownian) risk premium, a volatility risk premium, and a state-dependent jump risk premium. Under the physical measure, Pan (2002) proposes the following process for the underlying index price S_t and variance V_t

$$dS_t = [r_t - q_t + \eta^S V_t + \lambda V_t(\mu - \mu^*)]S_t dt + \sqrt{V_t} S_t dB_t^1 + dZ_t - \mu S_t \lambda V_t dt, \tag{4.9}$$

$$dV_t = \kappa_v(\bar{v} - V_t)dt + \sigma_v \sqrt{V_t}\Big(\rho dB_t^1 + \sqrt{1-\rho^2} dB_t^2\Big), \tag{4.10}$$

[11] In addition to that, several studies relate volatility to the pricing kernel puzzle or risk aversion, e.g. Bliss and Panigirtzoglou (2004), Ziegler (2007), and Linn et al. (2018).

where the riskless rate r_t and the dividend yield q_t both follow a square-root process with long-run means $\bar{r}$ and $\bar{q}$, mean reversion rates κ_r and κ_q, and volatility coefficients σ_r and σ_q, respectively.[12] Random innovations are introduced by two independent standard Brownian motions, dB_t^1 and dB_t^2, and a poisson (pure-jump) process, Z_t, whose jump intensity is λV_t and which is thus perfectly correlated with the instantaneous variance V_t. The logarithm of the relative jump size, conditional on a jump occurring, is normally distributed with mean $\mu_J = \ln(1+\mu) - \sigma_J^2/2$ and variance σ_J^2. Thus, the last term of Equation (4.9), $\mu S_t \lambda V_t dt$, compensates for the instantaneous change in expected index returns introduced by the pure-jump process Z_t. The premia for Brownian return risks and jump risks are estimated by $\eta^S V_t$ and $\lambda V_t(\mu - \mu^*)$, respectively. The variance process is modeled by Equation (4.10) and follows a square-root process with long-run mean $\bar{v}$, mean reversion rate κ_v, and volatility σ_v. The Brownian shocks to price S_t and variance V_t are correlated with constant coefficient ρ. Under the risk neutral measure, the dynamics of S_t and V_t evolve according to

$$dS_t = (r_t - q_t)S_t dt + \sqrt{V_t} S_t dB_t^1(Q) + dZ_t^Q - \mu^* S_t \lambda V_t dt, \tag{4.11}$$

$$dV_t = [\kappa_v(\bar{v} - V_t) + \eta^v V_t]dt + \sigma_v \sqrt{V_t}\Big(\rho dB_t^1(Q) + \sqrt{1-\rho^2} dB_t^2(Q)\Big), \tag{4.12}$$

where r_t and q_t are assumed to behave as under the physical measure. $dB_t^1(Q)$, $dB_t^2(Q)$, and Z_t^Q are two independent standard Brownian motions and the poisson (pure-jump) process under the risk neutral measure, respectively. Again, jump intensity is defined by λV_t and the logarithm of

[12] See Pan (2002)'s Equation (2.3).

Table 4.1: Pan (2002) Parameters

Panel A: Pan (2002), Table 3									
κ_v	$\bar{v}$	σ_v	ρ	η^s	η^v	$\lambda^* = \lambda$	μ (%)	σ_J	μ^* (%)
6.4	0.0153	0.30	−0.53	3.6	3.1	12.3	−0.8	0.0387	−19.2
(1.8)	(0.0029)	(0.04)	(0.07)	(2.4)	(2.2)	(1.9)	(2.4)	(0.0072)	(1.8)
Panel B: Pan (2002), Table 6									
		κ_r	$\bar{r}$	σ_r	κ_q	$\bar{q}$	σ_q		
		0.20	0.058	0.0415	0.24	0.025	0.0269		
		(0.15)	(0.016)	(0.0009)	(0.33)	(0.011)	(0.0004)		

Table 4.1 presents parameter estimates from the Pan (2002) stochastic volatility and jumps model (see her Tables 3 and 6). Pan (2002) estimates these parameters with an implied-state generalized method of moments (IS-GMM) approach and joint spot and option data from the Berkeley Options Data Base for the period from 1989 to 1996.

the jump size, conditional on a jump occurring, is normally distributed with mean $\mu_J^* = \ln(1+\mu^*) - \sigma_J^2/2$ and variance σ_J^2. The variance process in Equation (4.12) is defined by the mean reversion rate $\kappa_v^* = \kappa_v - \eta_v$, the long-run mean $\bar{v}^* = \kappa_v \bar{v}/\kappa_v^*$, and the volatility coefficient σ_v. The volatility risk premium is estimated by $\eta^v V_t$.

Pan (2002) estimates parameters with an 'implied-state' generalized method of moments (IS-GMM) approach and joint spot and option data from the Berkeley Options Data Base.[13] We provide an overview of the coefficient estimates in Table 4.1 and employ these to obtain risk neutral and physical densities via Fourier inversion. Thereby, we closely follow Ziegler (2007) and refer to Appendix 4.A.1 for more details.

An obvious concern of our approach is that Pan (2002) applies option data from 1989 to 1996 and the market environment has changed thereafter. In Section 4.4.1, we therefore provide an out-of-sample test

[13] We refer to Pan (2002) for more details.

by implementing a nonparametric empirical setting for the period from 1996 to 2020. Moreover, in Section 4.4.4 we re-run our simulation with alternative parameters. In both cases, we find our results to hold.

4.2.3 Differentiation from Earlier Studies

Although the economic theory underlying our analysis is very similar to that of Dierkes et al. (2022), the actual implementation differs significantly. While they conduct a fully nonparametric approach and obtain risk neutral and physical densities from option prices, we rely on the Pan (2002) model. Our analysis is thus entirely simulation-based.

To solve the DDE introduced in Section 4.2.1, Dierkes et al. (2022) employ different S&P 500 maturities.[14] By repeating this approach each month, they estimate a *time series* of probability weighting functions. We adjust their approach by assuming different levels of volatilities. More specifically, we solve the DDE for volatilities, $v_t = \sqrt{V_t}$, from 0.01 to 0.60 by choosing two adjacent volatility levels, e.g. $v_t = 0.10$ and $v_t + 0.01 = 0.11$. We thus provide a *cross-section* of probability weighting functions, which enables us to counterfactually investigate the relationship between probability weighting and volatilities.

4.2.4 Fitting Probability Weighting Functions

Given that we have estimated nonparametric probability weights according to Sections 4.2.1 through 4.2.3, we now have to fit these weights to

[14] We refer to their study for more details. See also Dierkes (2013) who was the first to introduce the elicitation procedure described above.

parametric functions. To do so, we make use of three different functional forms: the two-parameter Prelec (1998) function, the two-parameter linear-in-log-odds function (Tversky and Fox, 1995; Bleichrodt and Pinto, 2000), and the one-parameter Tversky and Kahneman (1992) function, as defined by Equations (4.13), (4.14), and (4.15), respectively.

$$w(p) = e^{-\delta(-\log(p))^{\gamma}}, \tag{4.13}$$

$$w(p) = \frac{\delta p^{\gamma}}{\delta p^{\gamma} + (1-p)^{\gamma}}, \tag{4.14}$$

$$w(p) = \frac{p^{\gamma}}{(p^{\gamma} + (1-p)^{\gamma})^{1/\gamma}}, \tag{4.15}$$

where $\gamma < 1$ implies overweighting of small probabilities and the typical inverse S-shape. We estimate δ and γ by fitting Equations (4.13) and (4.14) with linear regressions according to

$$\log(-\log(w(p)) = \log(\delta) + \gamma \log(-\log(p)) + \epsilon, \tag{4.16}$$

$$\log\left(\frac{w(p)}{1-w(p)}\right) = \log(\delta) + \gamma \log\left(\frac{p}{1-p}\right) + \epsilon, \tag{4.17}$$

on the interval $p \in \{0.01, 0.02, \ldots, 0.99\}$, while Equation (4.15) is fitted with non-linear least squares. ϵ denotes the residuals.

4.3 Results

4.3.1 Implied Probability Weighting Functions

In this section, we report results for the main specification of our study, characterized by a return horizon of one year and stochastic volatilities ranging from 0.01 to 0.60. However, in Section 5.4, we show that our results extend to return horizons of three and six months.

In order to estimate implied probability weighting functions, we first obtain physical and risk neutral probabilities according to Section 4.2.2. Fig. 4.1 illustrates the corresponding densities (Panel A) and distribution functions (Panel B) across wealth levels (averaged over volatilities). Dashed lines correspond to 95% point-wise confidence intervals. By construction, our results are very similar to those of Ziegler (2007). This is, the average physical density is located to the right of the risk neutral density and exhibits a more pronounced peak (at a wealth level of 1.13). However, in contrast to Ziegler (2007), we find both densities to be more dispersed and attribute this finding to a different choice of volatilities. While we average over a large set of volatilities, Ziegler (2007)'s results are based on five rather low volatilities, ranging from roughly 0.097 to 0.145.[15]

In Fig. 4.2, we illustrate the estimated probability weighting parameters for each volatility level and each of the three functional forms outlined in Section 4.2.4. While we also report the elevation parameter delta (Panel B), we follow Polkovnichenko and Zhao (2013) and focus our analysis on the curvature parameter gamma (Panel A).

Although the Pan (2002) model was never designed to match CPT preferences, the relationship between volatilities and probability weighting is strikingly clear. For all weighting functions, probability weighting is present across volatility levels and gammas almost monotonically decrease in volatility. This result is well in line with Polkovnichenko and

[15] Ziegler (2007) bases his choice on the average volatility reported in Pan (2002), i.e. he applies $\bar{v} = \sqrt{0.0153} \pm$ one and two standard errors.

Fig. 4.1: Physical and Risk Neutral Distributions, 1 Year Horizon

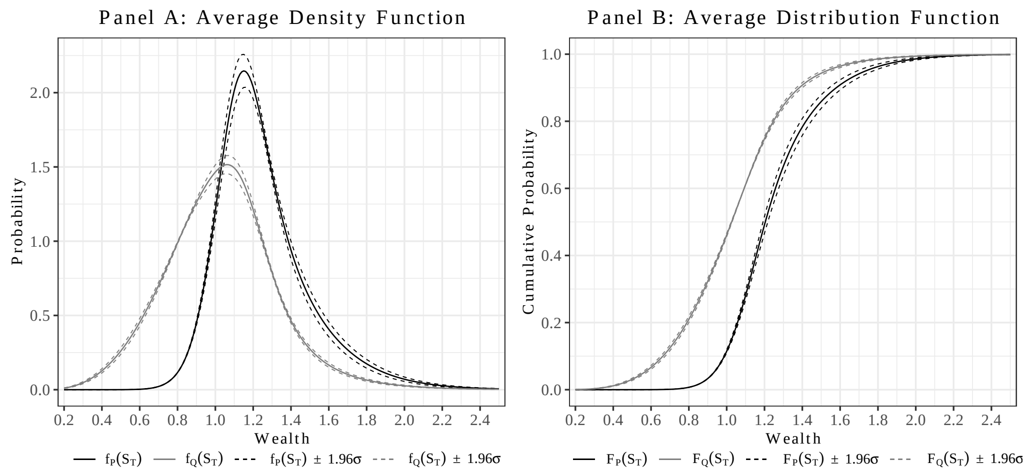

Fig. 4.1 plots physical and risk neutral densities (Panel A) and distribution functions (Panel B), estimated from the Pan (2002) stochastic volatility and jumps model and averaged over volatilities from 0.01 to 0.60. Physical (risk neutral) densities are denoted by $f_P(S_T)$ ($f_Q(S_T)$), whereas physical (risk neutral) distribution functions are denoted by $F_P(S_T)$ ($F_Q(S_T)$). We assume a return horizon of one year.

Zhao (2013) and Kilka and Weber (2001). Moreover, it corresponds to Gao et al. (2021) who find that investors dislike high-skewness securities when market volatility is low. With the two-parameter Prelec (1998) specification (γ^{Prelec}), gamma is about 0.99 for very low volatilities and 0.70 for high volatilities. With respect to the linear-in-log-odds function ($\gamma^{Log.Odds}$), our results are very similar as gammas vary from 0.90 to 0.68. In contrast, variation of the Tversky and Kahneman (1992) gamma (γ^{TK92}) is slightly reduced to 0.95 and 0.82, which is likely explained by the fact that the Tversky and Kahneman (1992) function does not include the elevation parameter δ.

In Table 4.2, we summarize our results and compare them to parameter estimates from previous studies.[16] Given that Pan (2002) estimates model

[16] We refer to Stott (2006) for an extensive overview.

Fig. 4.2: Implied Probability Weighting, 1 Year Horizon

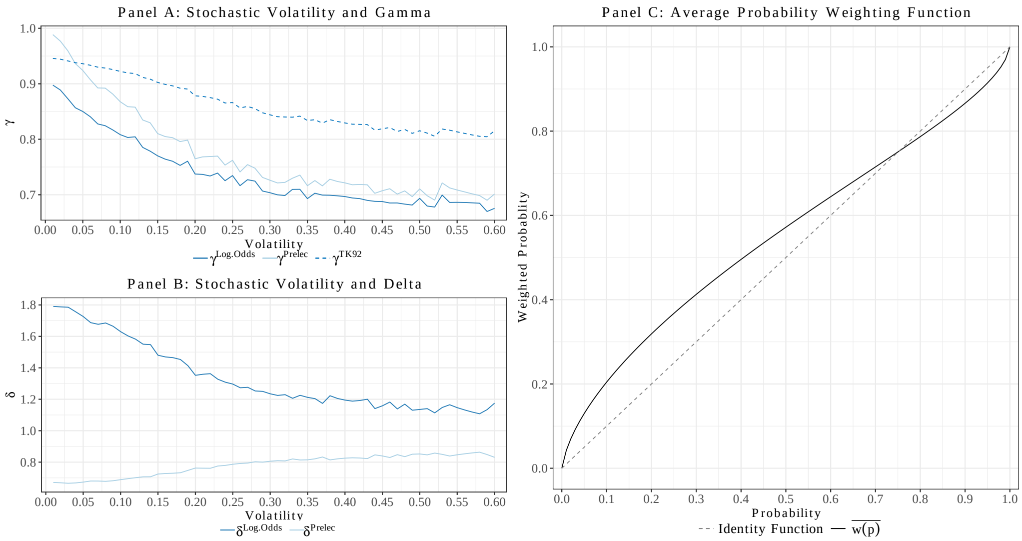

Fig. 4.2 plots results for probability weighting functions estimated from the Pan (2002) stochastic volatility and jumps model. We identify probability weights nonparametrically and estimate parameter values for three well known probability weighting functions, namely the two-parameter weighting function of Prelec (1998), denoted by *Prelec*, the two-parameter linear-in-log-odds function (as used in Tversky and Fox, 1995; Bleichrodt and Pinto, 2000), denoted by *Log.Odds*, and the one-parameter function of Tversky and Kahneman (1992), denoted by *TK92*. While Panel A and Panel B display the curvature parameter γ and the elevation parameter δ, respectively, Panel C shows the probability weighting function averaged over volatilities. We assume a return horizon of one year.

parameters from the S&P 500, i.e. one of the most liquid and competitive option markets in the world, it is not surprising that our gamma estimates are closer to one compared to, for example, Bleichrodt and Pinto (2000) and Kliger and Levy (2009). However, in line with these studies, we find persistent overweighting of small probabilities, indicating a high demand for lottery-like assets and a large potential impact of probability weighting on the pricing kernel puzzle. Moreover, note that Polkovnichenko and Zhao (2013)'s median estimates (0.90-0.95, depending on the assumed level of risk aversion) are even closer to one. As a result, the average

Table 4.2: Typical Parameters of Probability Weighting Functions

Functional Form	Study	γ	δ
$Prelec$	Kliger and Levy (2009)	0.60	0.79
$Prelec$	Polkovnichenko and Zhao (2013)	0.90 – 0.95	
$Log.Odds$	Wu and Gonzalez (1996)	0.68	0.84
$Log.Odds$	Tversky and Fox (1995)	0.69	0.77
$Log.Odds$	Bleichrodt and Pinto (2000)	0.55	0.81
$Log.Odds$	Dierkes et al. (2022)	0.89	
$TK92$	Tversky and Kahneman (1992)	0.61	
$TK92$	Wu and Gonzalez (1996)	0.71	
$TK92$	Zeisberger et al. (2012)	0.86 – 0.87	
$Prelec$	This study	0.70 – 0.99	1.19 – 1.79
$Log.Odds$	This study	0.68 – 0.90	0.68 – 0.83
$TK92$	This study	0.82 – 0.95	

Table 4.2 lists parameter estimates of previous studies for different probability weighting functions. We report results for the two-parameter weighting function of Prelec (1998), denoted by $Prelec$, the one- and two-parameter linear-in-log-odds function (as used in Tversky and Fox, 1995; Bleichrodt and Pinto, 2000), denoted by $Log.Odds$, and the one-parameter function of Tversky and Kahneman (1992), denoted by $TK92$. We outline these functions in Equations (4.13) to (4.15), where δ denotes the elevation parameter and γ defines the curvature. The last three rows correspond to our results illustrated in Fig. 4.2. We report our estimates for the lowest and highest volatility (0.01 and 0.60), respectively. Note that Polkovnichenko and Zhao (2013) apply the two-parameter Prelec (1998) function, but only report median values for the curvature parameter γ. Dierkes et al. (2022) employ the one-parameter linear-in-log-odds function. We only report parameters for gains.

probability weighting function over volatilities (Panel C) is characterized by a pronounced inverse S-shape. According to Polkovnichenko and Zhao (2013), inverse S-shaped probability weighting functions (including a convex segment) are consistent with non-monotonicity in pricing kernels and negative risk aversion functions. As we find pronounced probability weighting, we expect a strong impact of the probabilistic risk attitude on pricing kernels and risk aversion functions.

Apart from that, understanding how probability weighting varies with volatility might help us to understand how negative premia on lottery stocks such as IPOs (Green and Hwang, 2012), SEOs (Chen et al., 2019), and OTC stocks (Eraker and Ready, 2015) change with aggregate volatility. Moreover, M&A activity (Schneider and Spalt, 2017) and the equity share in new issues (Baker and Wurgler, 2000) might also depend on volatility. To quantify the relationship between probability weighting and volatility, we fit linear regressions of gamma on volatilities ($v = 0.01, 0.02, \ldots, 0.60$) and variances ($v^2 = 0.01^2, \ldots, 0.60^2$). We report the regression estimates below.

$$\gamma^{Prelec} = 0.977 - 1.181v + 1.253v^2, \qquad \text{adj. } R^2 = 0.978$$

$$\gamma^{Log.Odds} = 0.889 - 0.861v + 0.892v^2, \qquad \text{adj. } R^2 = 0.980$$

$$\gamma^{TK92} = 0.962 - 0.482v + 0.374v^2, \qquad \text{adj. } R^2 = 0.989$$

We conclude that there is a distinct and close relationship between volatility and probability weighting. In times of market distress (when volatility is high), investors overweight small probabilities and the demand for lottery-like assets increases, while during low volatility regimes weighted probabilities are close to their actual counterparts.

4.3.2 The Pricing Kernel Puzzle

According to economic theory, pricing kernels are defined as the ratio of risk neutral to physical probabilities and should monotonically decrease in wealth. Moreover, it is well established that the pricing kernel and risk aversion are two sides of the same coin. A locally decreasing (increasing) pricing kernel directly implies a locally positive (negative) risk aversion and vice versa. Thus, we can make a statement on the pricing kernel either by estimating the pricing kernel itself or by retracing it from risk aversion functions.

However, in contrast to economic theory, several recent studies have captured (locally) U-shaped pricing kernels or negative episodes of the risk aversion function, implying the pricing kernel and risk aversion puzzle, respectively.[17] We tackle these puzzles by adjusting both pricing kernels and risk aversion functions for probability weighting. Note that, although most of our results relate to risk aversion functions, we refer to both puzzles by the term 'pricing kernel puzzle' as this term is more frequently used in the literature.

In a first step, we investigate the pricing kernel puzzle by calculating both the raw pricing kernel, i.e. $f_Q(S_T)/f_P(S_T)$, and the pricing kernel net of probability weighting (as outlined in Equation 4.4). Recall that the adjusted pricing kernel varies with the physical distribution, $F_P(S_T)$, if the probability weighting function is not linear (i.e. $\gamma \neq 1$). Thus, convex parts of the probability weighting function ($w'(1-F_P(S_T)) > 1$) reduce the pricing kernel, whereas concave parts ($w'(1-F_P(S_T)) < 1$) increase it. Given our finding of pronounced probability weighting across volatilities, we expect the adjusted pricing kernel to monotonically decrease in wealth.

[17] See, for example, Jackwerth (2000), Ait-Sahalia and Lo (2000), and Rosenberg and Engle (2002).

Fig. 4.3: Average Pricing Kernel, 1 Year Horizon

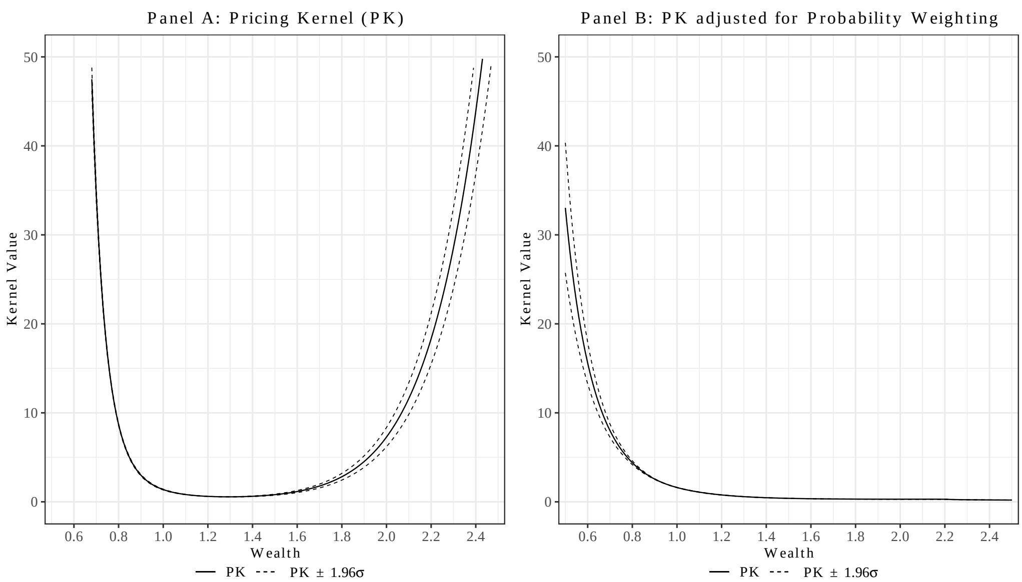

Fig. 4.3 plots pricing kernels estimated from the Pan (2002) stochastic volatility and jumps model. Following the literature (e.g. Jackwerth, 2000; Baele et al., 2019), we estimate the pricing kernel (PK, Panel A) as the ratio of the risk neutral to the physical probabilities, i.e $PK = f_Q/f_P$. In Panel B, we follow Equation (4.4) and calculate the pricing kernel, net of probability weighting, as $PK = f_Q(S_T)/f_P(S_T) \cdot w'(1-F_P(S_T))$. We assume a return horizon of one year.

Fig. 4.3 illustrates both the raw and the adjusted pricing kernel, estimated from the Pan (2002) model and averaged over volatilities. The return horizon is one year and dashed lines correspond to 95% point-wise confidence intervals. Note that, in order to obtain a smooth probabilistic risk attitude, we derive $w'(1-F_P(S_T))$ and $w''(1-F_P(S_T))$ analytically by fitting the nonparametric probability weights to the two-parameter weighting function of Prelec (1998). We obtain an almost perfect fit.[18] Consistent with the literature, we find the raw pricing kernel (Panel A) to exhibit a pronounced global U-Shape, implying a decreasing and par-

[18] Depending on the volatility level, R^2's vary from 99.74% to 99.99%. The median estimate is 99.96%.

tially negative risk aversion. Hence, as inferred by Ziegler (2007), the Pan (2002) model alone does not lead to well-behaved preferences. However, by providing closed-form expressions for the transforms of both $f_P(S_T)$ and $f_Q(S_T)$, the model is well-suited to adjust the pricing kernel and risk aversion functions for probability weighting, as outlined by Equations (4.4) and (4.5). In Panel B, we report results for the adjusted pricing kernel, which is monotonically decreasing in wealth. Thus, after accounting for probability weighting, the pricing kernel is well in line with economic theory and corresponds to Baele et al. (2019). To have a closer look at the dynamics driving this result, it is natural to investigate risk aversion functions. Fortunately, Equation (4.5) enables us to separate risk aversion related to the utility function u (denoted by ARA_u) and risk aversion originating from the probability weighting function w (the probabilistic risk attitude ARA_w). Fig. 4.4 presents results for a return horizon of one year, where risk aversion functions are averaged over volatilities and dashed lines correspond to 95% point-wise confidence intervals.

In Panel A, we report the absolute risk aversion (ARA) over wealth levels. While the Pan (2002) model indeed solves the problem of a U-shape, risk aversion still becomes negative for wealth levels greater than 1.35. By construction, this finding is consistent with Ziegler (2007) and the raw pricing kernel reported in Fig. 4.3. Moreover, it confirms Campbell and Cochrane (1999) and Brandt and Wang (2003): when the business cycle reaches the trough, wealth levels are low and the corresponding risk aversion is high. Panels B and C illustrate the adjusted risk aversion, ARA_u, and the probabilistic risk attitude, ARA_w, respectively. Most im-

Fig. 4.4: Implied Absolute Risk Aversion, 1 Year Horizon

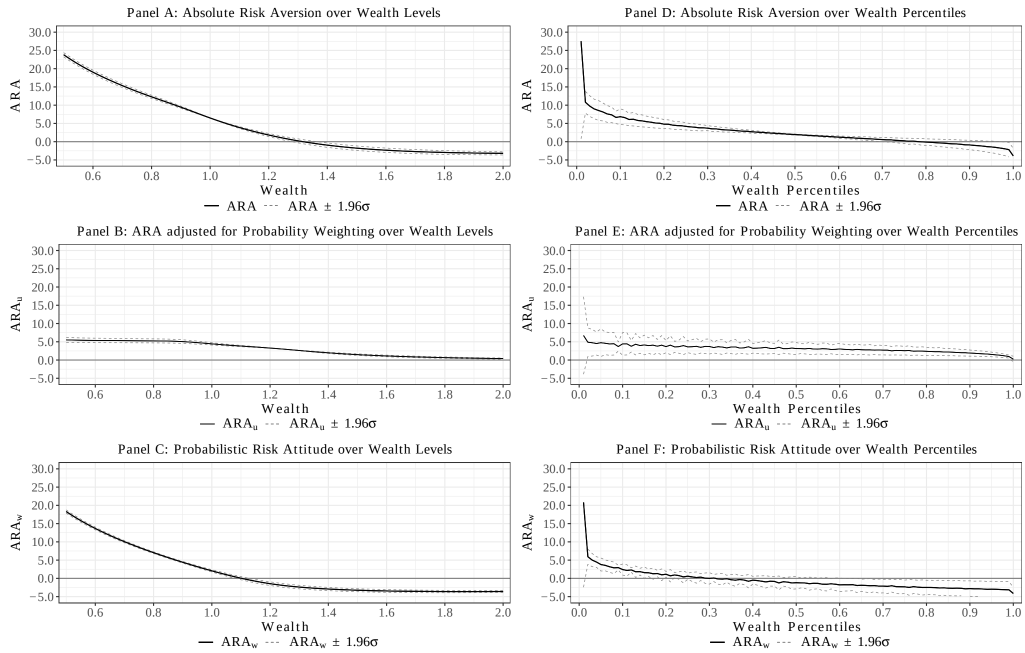

Fig. 4.4 plots implied risk aversion functions estimated from the Pan (2002) stochastic volatility and jumps model. In Panel A, we report the absolute risk aversion (ARA) over wealth levels, averaged across volatilities from 0.01 to 0.60 and calculated following Jackwerth (2000) and Equation (4.2): $ARA = F_P'(S_T)/F_P(S_T) - F_Q'(S_T)/F_Q(S_T)$. In Panel B, we report ARA functions adjusted for the probabilistic risk attitude (as outlined by Equation 4.5), i.e. $ARA_u = ARA - \frac{w''(1-F_P(S_T))}{w'(1-F_P(S_T))} \cdot f_P(S_T)$. We derive $w''(1-F_P(S_T))$ and $w'(1-F_P(S_T))$ analytically by fitting the nonparametrically estimated probability weights to the two-parameter probability weighting function of Prelec (1998). In Panels D to F, we repeat all estimations for wealth percentiles instead of wealth levels. We assume a return horizon of one year.

portantly, as ARA_w closely resembles ARA, the adjusted risk aversion is significantly positive and almost constant over wealth levels. Moreover, in accordance with the dynamics outlined in Section 4.2.1, ARA_w is positive and decreasing for low wealth levels, while it becomes negative for wealth levels greater than 1.10. To prove that our findings do not depend on the specific choice of volatilities, we repeat our calculations for wealth *percentiles* instead of *levels* and present results in Panels D to F.[19] In fact, the high risk aversion for wealth levels smaller than 0.80 appears to be driven by only a few wealth percentiles. However, the probabilistic risk attitude, ARA_w, still resembles this behavior very closely, resulting in a positive and almost constant adjusted risk aversion, ARA_u.[20] Again, this is a surprisingly clear result given that the Pan (2002) model does not account for probability weighting. However, a reasonable concern of our approach is that we calculate $w'(1 - F_P(S_T))$ and $w''(1 - F_P(S_T))$ analytically by fitting the estimated probability weights to the weighting function of Prelec (1998). In Section 4.4.3, we accommodate this concern by providing results for both an alternative functional assumption and a numerical solution.

Our results shed further light on the dynamics driving the pricing kernel puzzle. By accounting for probability weighting, we obtain a monotonically decreasing pricing kernel and a decreasing but consistently positive risk aversion. Importantly, we show that negative episodes of the risk aversion function arise due to the probabilistic risk attitude being negative for high wealth levels. We therefore conclude that the probabilistic risk attitude is a promising explanation for the pricing kernel puzzle.

[19] For example, low volatilities correspond to almost no probability mass for wealth levels greater than 1.40.

[20] Note that risk aversion estimates become insignificant for the lowest wealth percentiles as there is a strongly increased standard deviation.

4.4 Robustness

4.4.1 Empirical Relationship between Probability Weighting and Volatility

While our main results are based on the Pan (2002) model and suggest a strongly negative relation between gamma and stochastic volatility (i.e. probability weighting increases in volatility), it is reasonable to ask whether this relationship extends to a nonparametric empirical setting. We therefore follow Dierkes et al. (2022) and utilize a time series of monthly gammas from nonparametric estimates of the physical density function f_P (via S&P 500 returns) and the risk neutral density f_Q (via S&P 500 option prices).[21] To measure volatility, we employ the option-implied volatility index VIX which is provided by the Chicago Board Options Exchange (CBOE) on a daily basis. To reconcile both time series, we calculate the monthly average of daily VIX closing prices. Data on gammas is provided by Dierkes et al. (2022). Due to data availability, we cover a sample period from February 1996 to December 2020. As this period directly follows the sample period used in Pan (2002), our robustness check also serves as an out-of-sample test.

Most importantly, we once more find a strongly negative relationship. For example, a univariate regression of gamma on the VIX yields a negative and highly significant coefficient estimate ($t = -10.91$) and

[21] The authors illustrate that their measure closely reflects several stock market episodes like the DotCom bubble, the subprime crisis, and the recent surge in lottery demand in 2020 and 2021 (see their Fig. 1). Parameters are fitted based on the two-parameter log-odds function.

Fig. 4.5: Empirical Relationship between Probability Weighting and Volatility

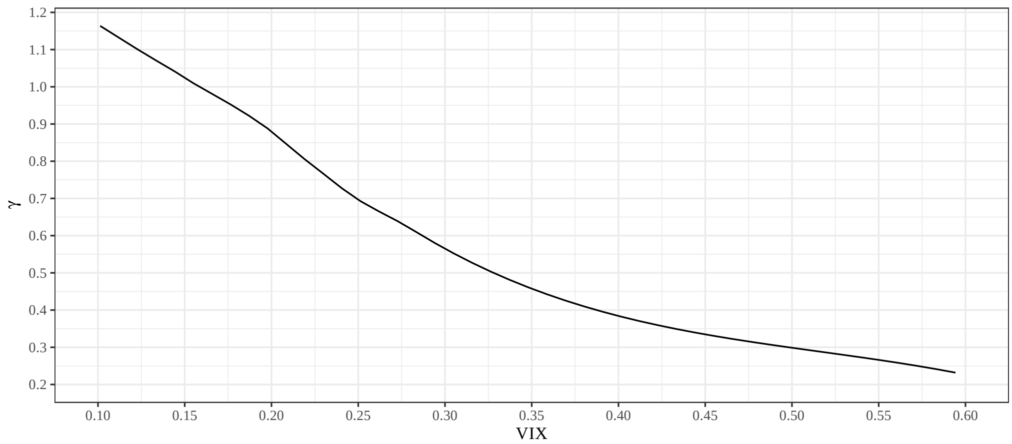

Fig. 4.5 plots the lowess-smoothed empirical relationship between the probability weighting parameter gamma and volatility. Gammas are estimated by applying nonparametric estimates of the physical density function f_P (via S&P 500 returns) and the risk neutral density f_Q (via S&P 500 option prices), while volatilities are proxied by the option-implied volatility index VIX. As data on the VIX is provided on a daily basis, we employ the monthly average of daily VIX closing prices. Due to data availability, we cover a sample period from February 1996 to December 2020.

$R^2 = 28.4\%$. By including an additional variance term (as in Section 4.3.1), R^2 even increases to 30.1%. Moreover, the difference between average gammas in high (0.71) and low volatility regimes (1.06), according to a median split of the VIX, is economically important and statistically significant at the 1%-level ($t = -8.05$).

In Fig. 4.5, we illustrate the empirical link between probability weighting and volatilities by estimating the lowess-smoothed relationship between gamma and the VIX. The range of gammas increases, yet the shape of the smoothed relationship is surprisingly close to that reported in Panel A of Fig 4.2. First, gamma monotonically decreases in volatility, indicating more pronounced probability weighting in volatile market environments.

Fig. 4.6: Empirical and Model-Implied Probability Weighting

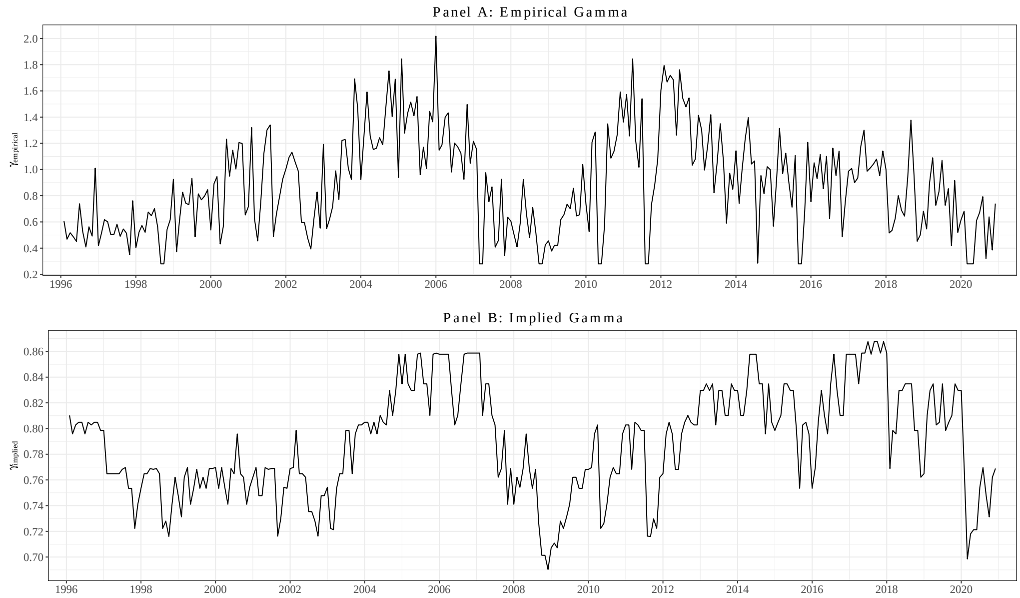

Panel A of Fig. 4.6 plots the time series of empirical gammas as described in Fig. 4.5. Panel B illustrates a time series of model-implied gammas. In order to estimate this time series, we employ the monthly average of daily VIX closing prices and match each value according to the relationship presented in Fig. 4.2. Due to data availability, we cover a sample period from February 1996 to December 2020.

Second, gamma strongly decreases for low VIX levels, while the slope is less steep for volatilities greater than 0.35, suggesting that the differential impact is stronger in low volatility periods.[22] The empirical relationship is thus well in line with Kilka and Weber (2001) and Gao et al. (2021) and confirms our simulation-based conclusions.

In Fig. 4.6, we compare the monthly empirical gammas (Panel A) to a time series of gammas implied by our simulation approach (Panel B). In order to estimate the model-implied gammas, we employ the monthly average of daily VIX closing prices and match each value according to

[22] Replacing the average VIX by the maximum VIX per month leads to very similar results. However, while maintaining its shape, the lowess-smoothed relationship is, by construction, slightly shifted upwards.

the relationship presented in Fig. 4.2. While model-implied gammas are, by construction, more condensed, the overall shape of both time series is remarkably close. For example, both time series display lower gammas (i.e. increased probability weighting) during the run-up of the DotCom bubble in 1998-2000 and the subprime crisis in 2007-2009. Moreover, both estimates nicely reflect increased probability weighting after Covid-19 reached global stock markets in March 2020. In line with this, we also find a large correlation between the two time series (54.5%). We therefore consider our results as further out-of-sample evidence. Moreover, they are well in line with the literature on time-varying risk preferences, e.g. Brandt and Wang (2003), Guiso et al. (2018), and Polkovnichenko and Zhao (2013).

While it is reassuring that the empirical results confirm our simulation-based findings, note that in such an exercise it is not possible to counterfactually change the volatility level with all else being equal. That is, an analysis using several months with varying volatility could have been confounded by additional time-varying economic state variables. This is why, in our baseline analysis, we opted for model-based results with volatility as the only additional state variable.

4.4.2 Alternative Maturities

To prove that our results hold for alternative assumptions, we now repeat our estimations for return horizons of six and three months and focus on the most important components: volatility-dependent proba-

Fig. 4.7: Implied Probability Weighting, 6 Months and 3 Months Horizon

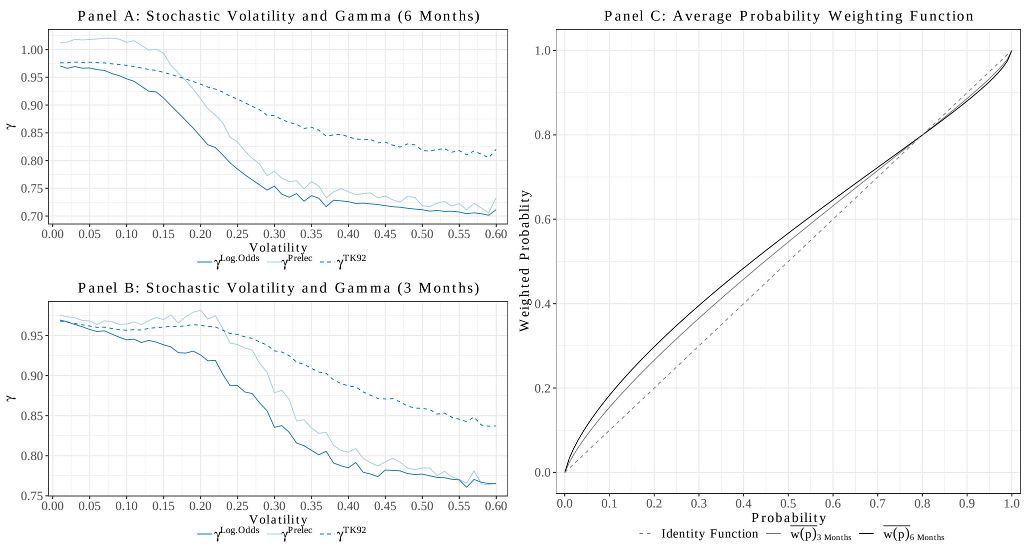

Fig. 4.7 plots results for probability weighting functions estimated from the Pan (2002) stochastic volatility and jumps model. We identify probability weights nonparametrically and estimate parameter values for three well known probability weighting functions, namely the two-parameter weighting function of Prelec (1998), denoted by *Prelec*, the two-parameter linear-in-log-odds function (as used in Tversky and Fox, 1995; Bleichrodt and Pinto, 2000), denoted by *Log.Odds*, and the one-parameter function of Tversky and Kahneman (1992), denoted by *TK92*. While Panel A and Panel B display the curvature parameter γ for a return horizon of six and three months, respectively, Panel C shows the probability weighting function averaged over volatilities. We assume a return horizon of six months.

bility weighting and the impact of the probabilistic risk attitude on risk aversion functions. However, we report corresponding risk neutral and physical densities as well as pricing kernels in Appendices 4.A.2 and 4.A.3, respectively.

Fig. 4.7 illustrates the curvature parameter gamma as well as the average probability weighting function for return horizons of six and three months, respectively. With respect to a return horizon of six months (Panel A), the variation in gamma is very similar to our main specification. While γ^{Prelec} varies from 1.02 for low volatilities to 0.71 for high volatilities, $\gamma^{Log.Odds}$ decreases from roughly 0.97 to 0.70. Notably, γ^{Prelec} is rather constant for volatilities between 0.01 and 0.11 and then sharply decreases for volatilities between 0.12 and 0.30. Again, the variation in γ^{TK92} is somewhat smaller (0.98 to 0.80) but still reasonable. Most importantly, even though gammas seem to be shifted upwards, we still find a strongly negative relationship with volatilities. As a result, the average probability weighting function (black solid line in Panel C) displays a distinct, but slightly less pronounced, inverse S-shape. Panel B reports gammas for a return horizon of three months. γ^{Prelec} ($\gamma^{Log.Odds}$) now varies from roughly 0.98 to 0.77 (0.97 to 0.78), whereas γ^{TK92} ranges from 0.97 to 0.84. Again, γ^{Prelec} is almost constant for small volatilities and then sharply decreases. Although gammas are below one, the overall level is further shifted upwards. As a consequence, the average probability weighting function (grey solid line in Panel C) is closer to the identity function (dashed line), but still preserves an inverse S-shape. We thus conclude that the estimation of probability weights is robust to alternative return horizons.

To further investigate the pricing kernel puzzle, we focus on the adjusted risk aversion (ARA_u). Fig. 4.8 illustrates risk aversion functions for a return horizon of six months. In Panel A, we report the average risk aversion over wealth levels. While the overall shape is close to our main specification, ARA is slightly shifted upwards and becomes negative for wealth levels greater than 1.20 (compared to 1.35 for a return horizon of one year). In contrast to Fig. 4.4, the adjusted risk aversion, ARA_u, is somewhat bumpier and slightly increasing for wealth levels greater

Fig. 4.8: Implied Absolute Risk Aversion, 6 Months Horizon

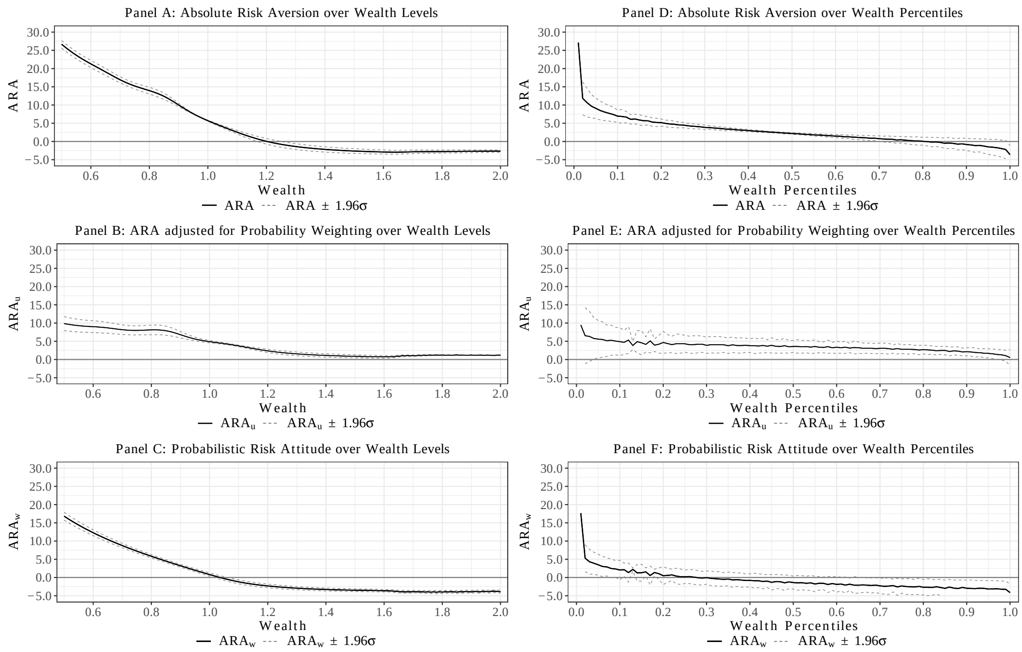

Fig. 4.8 plots implied risk aversion functions estimated from the Pan (2002) stochastic volatility and jumps model. In Panel A, we report the absolute risk aversion (ARA) over wealth levels, averaged across volatilities from 0.01 to 0.60 and calculated following Jackwerth (2000) and Equation (4.2): $ARA = F_P'(S_T)/F_P(S_T) - F_Q'(S_T)/F_Q(S_T)$. In Panel B, we report ARA functions adjusted for the probabilistic risk attitude (as outlined by Equation 4.5), i.e. $ARA_u = ARA - \frac{w''(1-F_P(S_T))}{w'(1-F_P(S_T))} \cdot f_P(S_T)$. We derive $w''(1 - F_P(S_T))$ and $w'(1 - F_P(S_T))$ analytically by fitting the derive estimated probability weights to the two-parameter probability weighting function of Prelec (1998). In Panels D to F, we repeat all estimations for wealth percentiles instead of wealth levels. We assume a return horizon of six months.

than 1.60 (with very little probability mass, Panel B). Most importantly, ARA_u remains significantly positive for all wealth levels and thus implies a monotonically decreasing pricing kernel. The probabilistic risk attitude is almost unchanged, i.e. ARA_w is positive and decreasing for low wealth levels, and negative for wealth levels greater than 1.05 (Panel C). With respect to wealth percentiles (Panels D to F), results correspond to our main specification. Notably, except for some noise around the 15% quan-

Fig. 4.9: Implied Absolute Risk Aversion, 3 Months Horizon

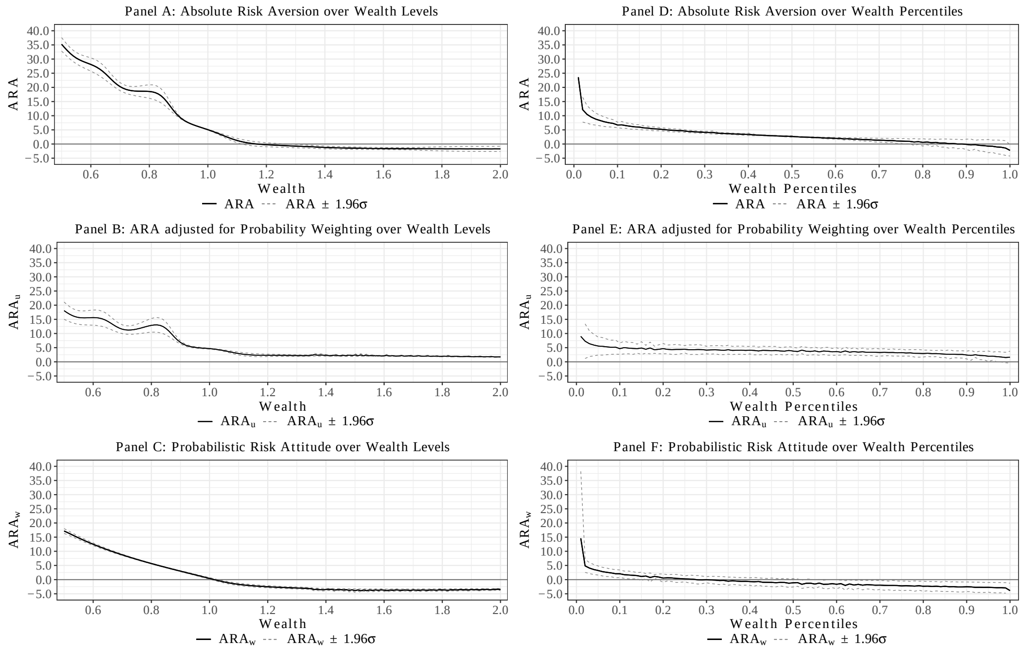

Fig. 4.9 plots implied risk aversion functions estimated from the Pan (2002) stochastic volatility and jumps model. In Panel A, we report the absolute risk aversion (ARA) over wealth levels, averaged across volatilities from 0.01 to 0.60 and calculated following Jackwerth (2000) and Equation (4.2): $ARA = F_P'(S_T)/F_P(S_T) - F_Q'(S_T)/F_Q(S_T)$. In Panel B, we report ARA functions adjusted for the probabilistic risk attitude (as outlined by Equation 4.5), i.e. $ARA_u = ARA - \frac{w''(1-F_P(S_T))}{w'(1-F_P(S_T))} \cdot f_P(S_T)$. We derive $w''(1-F_P(S_T))$ and $w'(1-F_P(S_T))$ analytically by fitting the nonparametrically estimated probability weights to the two-parameter probability weighting function of Prelec (1998). In Panels D to F, we repeat all estimations for wealth percentiles instead of wealth levels. We assume a return horizon of three months.

tile, there are no episodes of increasing ARA_u. We therefore argue that increasing segments in Panel B are merely an artifact of averaging over volatilities.

Fig. 4.9 illustrates results for a return horizon of three months. In contrast to our main specification, ARA is shifted upwards and appears to be more bumpy, but still monotonically decreases in wealth. Since the probabilistic risk attitude is only slightly affected, the bumpy shape of ARA directly transfers to the adjusted risk aversion. Hence, ARA_u exhibits increasing parts around a wealth level of 0.80 (with a physical density of almost zero). Most importantly, ARA_u is consistently positive and Panel E confirms that increasing episodes are, again, an artifact of averaging over volatilities.

In summary, we find our results to be robust to alternative maturities. Although risk aversion functions are less smooth, we find the risk aversion – net of probability weighting – to remain significantly positive over both wealth levels and wealth percentiles, implying a monotonically decreasing pricing kernel.

4.4.3 Alternative Estimation of the Probabilistic Risk Attitude

A natural concern of our approach is that we derive $w'(1-F_P(S_T))$ and $w''(1-F_P(S_T))$ analytically by fitting nonparametric probability weights to the two-parameter function of Prelec (1998). Thus, our results might reflect the specific functional assumption. To accommodate this concern, we first replace our functional assumption by the linear-in-log-odds probability weighting function and then provide an entirely numerical solution. Results for the linear-in-log-odds function and a return horizon of one year are presented in Fig. 4.10.[23]

By construction, absolute risk aversion functions in Panels A and D are not affected by a change of the functional assumption since ARA (see

[23] Again, we find an almost perfect fit.

Fig. 4.10: Implied Absolute Risk Aversion, Linear-in-Log-Odds, 1 Year Horizon

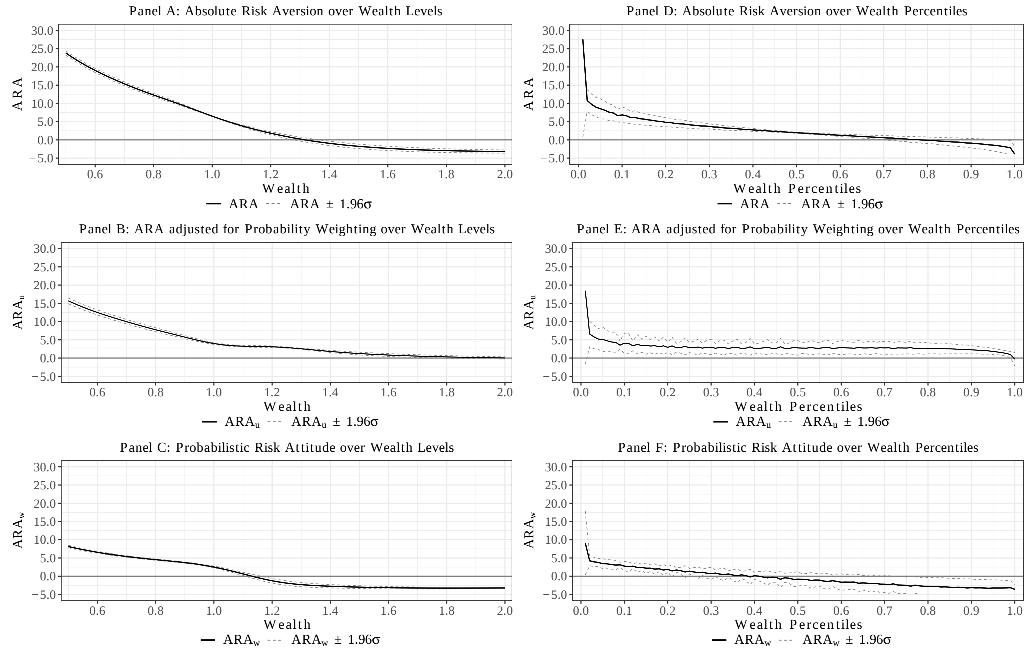

Fig. 4.10 plots implied risk aversion functions estimated from the Pan (2002) stochastic volatility and jumps model. In Panel A, we report the absolute risk aversion (ARA) over wealth levels, averaged across volatilities from 0.01 to 0.60 and calculated following Jackwerth (2000) and Equation (4.2): $ARA = F_P'(S_T)/F_P(S_T) - F_Q'(S_T)/F_Q(S_T)$. In Panel B, we report ARA functions adjusted for the probabilistic risk attitude (as outlined by Equation 4.5), i.e. $ARA_u = ARA - \frac{w''(1-F_P(S_T))}{w'(1-F_P(S_T))} \cdot f_P(S_T)$. We derive $w''(1-F_P(S_T))$ and $w'(1-F_P(S_T))$ analytically by fitting the nonparametrically estimated probability weights to the two-parameter linear-in-log-odds probability weighting function (Tversky and Fox, 1995; Bleichrodt and Pinto, 2000). In Panels D to F, we repeat all estimations for wealth percentiles instead of wealth levels. We assume a return horizon of one year.

Equation 4.2) does not depend on $w'(1-F_P(S_T))$ or $w''(1-F_P(S_T))$. Hence, changes in ARA_u (Panels B and E) solely depend on ARA_w (Panels C and F). In contrast to our main specification, we find the probabilistic risk attitude to be less pronounced for low wealth levels.[24] Thus, for a wealth level of 0.50, we find $ARA_u \approx 15$, whereas in our main specification it

[24] See Dierkes and Sejdiu (2019) for differences in the probabilistic risk attitude of various probability weighting functions for probabilities near zero.

holds $ARA_u \approx 5$. Most importantly, ARA_u remains positive throughout wealth levels and all but the highest wealth percentile, again implying a monotonically decreasing pricing kernel. Moreover, except for some noise, ARA_u is monotonically decreasing. Even though ARA_u is not significant for wealth levels greater than 1.77, it should again be noted that there is only little probability mass for these wealth levels (see Fig. 4.1). We therefore conclude that our findings are robust to a different functional assumption.

In Fig. 4.11, we present the resulting risk aversion functions when $w'(1 - F_P(S_T))$ and $w''(1 - F_P(S_T))$ are derived numerically, i.e. without fitting probability weights to a parametric weighting function. A word of caution is in order. As mentioned in Section 4.2.4, we estimate probability weights on the interval $p \in \{0.01, 0.02, \ldots, 0.99\}$. Hence, we are not able to estimate probability weights unless it holds for at least one volatility that $F_P(S_T) \geq 0.01$ or $F_P(S_T) \leq 0.99$. Moreover, to derive $w'(1 - F_P(S_T))$ and $w''(1 - F_P(S_T))$ numerically, we lose two more observations. We therefore propose a fine grid for the two probabilities at the extremes and a regular grid in between, i.e. $p \in \{0.001, 0.002, 0.01, 0.02, \ldots, 0.99, 0.998, 0.999\}$. By this means, we obtain the probabilistic risk attitude for at least one volatility and wealth levels between 0.72 and 2.00. Consequently, Fig. 4.11 is also limited to this range.

In Panel A, we report the absolute risk aversion which is slightly increasing for wealth levels below 0.78 and greater than 1.47. Note that, even though ARA does not depend on $w'(1 - F_P(S_T))$ or $w''(1 - F_P(S_T))$, the function differs from Fig. 4.4 and 4.10. The rationale is given by the

Fig. 4.11: Implied Absolute Risk Aversion, Numerical Solution, 1 Year Horizon

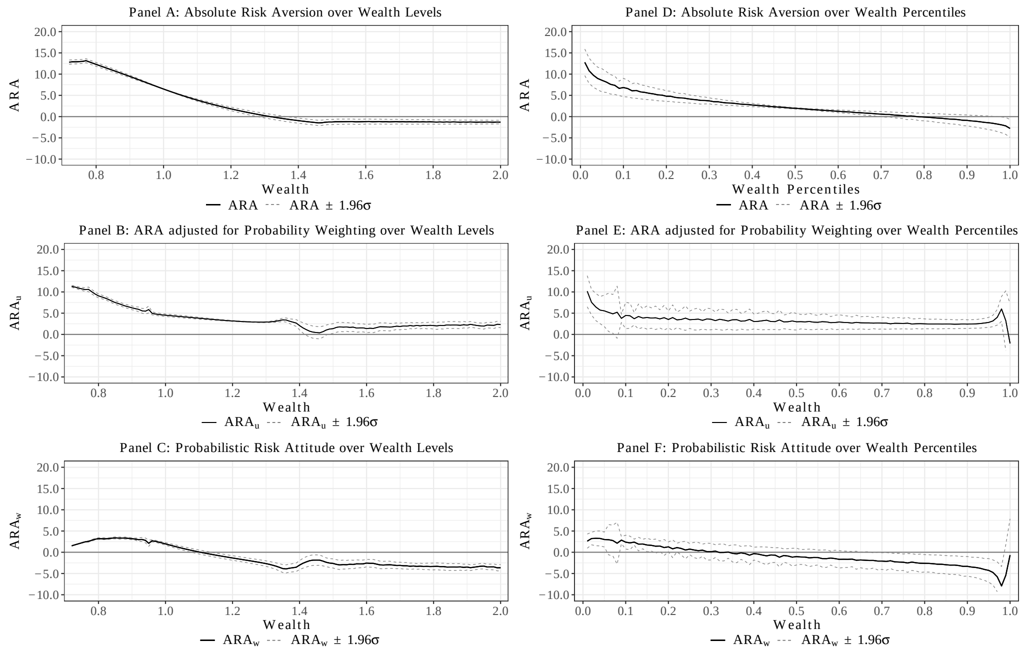

Fig. 4.11 plots implied risk aversion functions estimated from the Pan (2002) stochastic volatility and jumps model. In Panel A, we report the absolute risk aversion (ARA) over wealth levels, averaged across volatilities from 0.01 to 0.60 and calculated following Jackwerth (2000) and Equation (4.2): $ARA = F_P'(S_T)/F_P(S_T) - F_Q'(S_T)/F_Q(S_T)$. In Panel B, we report ARA functions adjusted for the probabilistic risk attitude (as outlined by Equation 4.5), i.e. $ARA_u = ARA - \frac{w''(1-F_P(S_T))}{w'(1-F_P(S_T))} \cdot f_P(S_T)$. In contrast to our main specification, we derive $w''(1-F_P(S_T))$ and $w'(1-F_P(S_T))$ numerically. In Panels D to F, we repeat all estimations for wealth percentiles instead of wealth levels. We assume a return horizon of one year.

fact that for wealth levels below 0.78 and above 1.47, $w'(1-F_P(S_T))$ and $w''(1-F_P(S_T))$ are only known for *some* of the examined volatilities. For wealth levels *between* 0.78 and 1.47, ARA is equal to Fig. 4.4. Resulting from the estimation instabilities described above, the probabilistic risk attitude is less smooth and locally increasing (Panel C). Nevertheless, we find the global shape, i.e. $ARA_w > 0$ for small wealth levels and $ARA_w < 0$ for high wealth levels, to be preserved. As a consequence, ARA_u remains significantly positive for the vast majority of wealth levels (Panel B). Panels D to F present results for wealth percentiles. Except for the highest wealth percentile, ARA_u is consistently positive and significant throughout the vast majority of wealth percentiles. Even though the estimation procedure is less stable, our results are thus robust to a numerical solution. In unreported results, we repeat our calculations for return horizons of three and six months and find our conclusions to hold.[25]

4.4.4 Sensitivity Analysis

Our simulation approach is not only based on the Pan (2002) model, but also relies on the corresponding parameter estimates. Although results in Section 4.4.1 hint that our findings related to gamma extend to the 1996-2020 period, a reasonable concern of our approach is that Pan (2002) employed options data from 1989 to 1996 and the market environment has changed thereafter. Below, however, we show that this concern is not warranted as we re-run our simulation with alternative parameters. More precisely, we adjust each parameter for ± one standard error ($\hat{\sigma}_{SE}$) and summarize results in Table 4.3.[26]

We find our results to be remarkably robust. For example, variation in parameters related to the interest rate r and the dividend yield q does not involve any considerable impact on probability weighting and risk aversion functions. Moreover, our conclusions remain unaffected by

[25] Confidence intervals in the numerical solution tend to increase. However, ARA_u remains significantly positive for the vast majority of wealth levels and percentiles.

[26] We thereby assume the remaining parameters to be constant. Strictly speaking, a changed market environment will also result in different standard errors. However, given our results in section 4.4.1, we assume sufficient accuracy.

Table 4.3: Sensitivity Analysis

Parameter	Estimate (Std. Error)	Comments
		Panel A: Pan (2002), Table 3
κ_v	6.4 (1.8)	In case of $\hat{\kappa}_v - 1\hat{\sigma}_{SE}$, gamma increases for $\sqrt{V} > 0.50$. Probability weighting remains strongly inverse S-shaped.
$\bar{v}$	0.0153 (0.0029)	In case of $\hat{\bar{v}} - 1\hat{\sigma}_{SE}$, our code becomes numerically unstable for $\sqrt{V} > 0.45$. $ARA_u > 0$, but not significant for wealth levels > 1.75 (probability mass is below 1%).
σ_v	0.30 (0.04)	
ρ	−0.53 (0.07)	
η^s	3.6 (2.4)	In case of $\hat{\eta}^s + 1\hat{\sigma}_{SE}$, our code becomes numerically unstable for $\sqrt{V} > 0.53$ and gamma increases for $\sqrt{V} > 0.43$. Probability weighting remains strongly inverse S-shaped.
η^v	3.1 (2.2)	In case of $\hat{\eta}^v - 1\hat{\sigma}_{SE}$, even ARA is always positive. In case of $\hat{\eta}^v + 1\hat{\sigma}_{SE}$, ARA (and also ARA_w) are strongly negative. ARA_u is slightly negative for wealth levels > 1.50. ARA_u over wealth percentiles, however, remains positive for all but the highest percentile.
$\lambda^* = \lambda$	12.3 (1.9)	
μ (%)	−0.8 (2.4)	
σ_J	0.0387 (0.0072)	In case of $\hat{\sigma}_J + 1\hat{\sigma}_{SE}$, our code becomes numerically unstable for $\sqrt{V} > 0.48$. Numerically derived ARA_u is < 0 for wealth levels > 1.92 (corresponding probability mass is almost zero).
μ^* (%)	−19.2 (1.8)	
		Panel B: Pan (2002), Table 6
κ_r	0.20 (0.15)	
$\bar{r}$	0.058 (0.016)	
σ_r	0.0415 (0.0009)	
κ_q	0.24 (0.33)	
$\bar{q}$	0.025 (0.011)	
σ_q	0.0269 (0.0004)	

Table 4.3 presents parameter estimates from the Pan (2002) stochastic volatility and jumps model (see her Tables 3 and 6). To check whether our results are robust to variations of these parameters, we re-run our calculations with $\pm 1\hat{\sigma}_{SE}$. While $\sqrt{V}$ indicates volatilities, ARA, ARA_u, and ARA_w denote the absolute risk aversion, the adjusted risk aversion (related to the utility function), and the probabilistic risk attitude, respectively. We assume a return horizon of one year and derive $w''(1 - F_P(S_T))$ and $w'(1 - F_P(S_T))$ analytically by fitting nonparametrically estimated probability weights to the two-parameter probability weighting function of Prelec (1998).

variation in the jump risk premium $\lambda V_t(\mu - \mu^*)$, the volatility coefficient σ_v, and the correlation between Brownian shocks ρ. Nevertheless, there are some minor instabilities which we want to outline hereafter.

Considering the mean reversion rate of the variance process, $\hat{\kappa}_v - 1\hat{\sigma}_{SE}$ results in an increasing gamma for volatility levels greater than 0.50. The average probability weighting function, however, is still strongly inverse S-shaped and ARA_u is not affected. With respect to the average volatility, we find that $\hat{\bar{v}} - 1\hat{\sigma}_{SE}$ results in an unstable estimation of the DDE for $v > 0.45$. Moreover, ARA_u is positive but not significant for wealth levels greater than 1.75 (with a corresponding probability mass of below 1%). In case of $\hat{\eta}^S + 1\hat{\sigma}_{SE}$ (related to the premium for Brownian risk), estimating the DDE becomes numerically unstable for $v > 0.53$ and gamma increases for $v > 0.43$. Nevertheless, probability weighting remains strongly inverse S-shaped. The largest impact on our results is given by a variation of the variance risk premium $\hat{\eta}^v$. In case of $\hat{\eta}^v - 1\hat{\sigma}_{SE}$, even ARA is consistently positive, while for $\hat{\eta}^v + 1\hat{\sigma}_{SE}$ both ARA and the probabilistic risk attitude are strongly negative. As a result, ARA_u is slightly negative for wealth levels greater than 1.50. Importantly, ARA_u over wealth *percentiles* remains positive for all but the highest percentile. Finally, variation in the jump size volatility ($\hat{\sigma}_J + 1\hat{\sigma}_{SE}$) causes the DDE to become unstable for $v > 0.48$. Despite that, our conclusions with respect to ARA_u and probability weighting functions remain unaffected.

In summary, we find some specifications for which solving the DDE becomes numerically unstable in case of large volatilities. Our conclusions, however, remain largely unaffected. We still find a strongly negative (positive) relationship between volatility and gamma (probability weighting), and after accounting for probability weighting, risk aversion functions are mostly significantly positive. Lastly, note that volatilities greater than 0.50 only occur with a probability of less than one percent.[27]

4.5 Concluding Remarks

We contribute to a large body of literature on time-varying risk preferences and the pricing kernel puzzle. Following Ziegler (2007), we obtain risk neutral and physical densities from the Pan (2002) stochastic volatility and jumps model for a large set of volatilities. Thereafter, we employ these densities to estimate nonparametric probability weights, which we fit to three well-known probability weighting functions: the two-parameter Prelec (1998) function, the two-parameter linear-in-log-odds function, and the one-parameter Tversky and Kahneman (1992) function. Even though the Pan (2002) model was not designed to account for CPT preferences, our results are strikingly clear. Implied probability weighting functions are strongly inverse S-shaped and the curvature parameter gamma almost monotonically decreases in volatility, suggesting that skewness preferences are more pronounced in volatile market environments. Moreover, we estimate probabilistic risk attitudes, equivalent to the share of risk aversion related to probability weighting. In doing so, we fit the estimated probability weights to the functional assumption of Prelec (1998) and calculate derivatives analytically. This enables us to ad-

[27] Given daily VIX closing prices from January 1990 to August 2022, the probability for volatilities greater than 0.50 (0.45) is 0.90% (1.36%).

just pricing kernel and risk aversion functions for probability weighting and to shed further light on the pricing kernel puzzle. We find the raw pricing kernel, implied by the Pan (2002) model, to display a pronounced U-shape, implying episodes of negative risk aversion. After taking into account probability weighting, however, the pricing kernel is monotonically decreasing in wealth and risk aversion functions remain significantly positive. Our results are robust to alternative return horizons, wealth percentiles, an alternative functional assumption and both a numerical approach to estimate the probabilistic risk attitude and variations of the Pan (2002) coefficient estimates. Moreover, we provide an out-of-sample test by implementing a nonparametric empirical setting for the period from 1996 to 2020, confirming that Pan (2002)'s parameter estimates are still appropriate. We therefore conclude that probability weighting is not only closely related to volatile market environments, but is also a key driver of the pricing kernel puzzle.

4.A Appendix

4.A.1 Estimation of Physical and Risk Neutral Densities

To estimate physical and risk neutral densities, we closely follow Ziegler (2007) who provides time-t conditional Fourier transforms of $\ln(S_T)$.[28] Given the notation outlined in Section 4.2.2 and initial values for the interest rate r, the dividend yield q, the volatility v, and the return horizon $\tau = T - t$, Ziegler (2007) provides the time-t conditional transform under the physical measure as

$$\psi(S; v, r, q, \tau) = \exp\Big(\alpha_r(S) + \alpha_q(S) + \alpha_v(S) + \beta_r(S)r + \beta_q(S)q + \beta_v(S)v\Big), \tag{4.A.1}$$

where α_i and β_i $(i = r, q, v)$ are defined as

$$\alpha_r = -\frac{\kappa_r \bar{r}}{\sigma_r^2}\left((\gamma_r - \kappa_r)\tau + 2\ln\left(1 - (\gamma_r - \kappa_r)\frac{1-\exp(-\gamma_r\tau)}{2\gamma_r}\right)\right), \tag{4.A.2}$$

$$\alpha_q = -\frac{\kappa_q \bar{q}}{\sigma_q^2}\left((\gamma_q - \kappa_q)\tau + 2\ln\left(1 - (\gamma_q - \kappa_q)\frac{1-\exp(-\gamma_q\tau)}{2\gamma_q}\right)\right), \tag{4.A.3}$$

$$\alpha_v = -\frac{\kappa_v \bar{v}}{\sigma_v^2}\left((\gamma_v + b)\tau + 2\ln\left(1 - (\gamma_v + b)\frac{1-\exp(-\gamma_v\tau)}{2\gamma_v}\right)\right), \tag{4.A.4}$$

$$\beta_r = -\frac{2(1-S)(1-\exp(-\gamma_r\tau))}{2\gamma_r - (\gamma_r - \kappa_r)(1-\exp(-\gamma_r\tau))}, \tag{4.A.5}$$

$$\beta_q = -\frac{2S(1-\exp(-\gamma_q\tau))}{2\gamma_q - (\gamma_q - \kappa_q)(1-\exp(-\gamma_q\tau))}, \tag{4.A.6}$$

$$\beta_v = -\frac{a(1-\exp(-\gamma_v\tau))}{2\gamma_v - (\gamma_v + b)(1-\exp(-\gamma_v\tau))}, \tag{4.A.7}$$

[28] See his Appendix B, where transforms are based on Pan (2002)'s Appendix B and Duffie et al. (2000).

with

$$\gamma_r = \sqrt{\kappa_r^2 + 2(1-S)\sigma_r^2}, \tag{4.A.8}$$

$$\gamma_q = \kappa_q^2 + 2S\sigma_q^2, \tag{4.A.9}$$

$$\gamma_v = \sqrt{b^2 + a\sigma_v^2}, \tag{4.A.10}$$

$$a = S(1-S) - 2\lambda\Big(\exp(S\mu_J + S^2\frac{\sigma_J^2}{2}) - 1 - S\mu^*\Big) - 2S\eta^S, \tag{4.A.11}$$

$$b = \sigma_v \rho S - \kappa_v. \tag{4.A.12}$$

Given the parameter estimates reported in Table 4.1 and the level of the underlying S, the physical density f_P can be obtained via numerical integration of

$$f_P(S;v,r,q,\tau) = \frac{1}{2\pi}\int_{-\infty}^{\infty} \psi(iz;v,r,q,\tau)\exp(-izS)dz. \tag{4.A.13}$$

To obtain the risk neutral density f_Q, some of the parameters have to be replaced by their risk neutral counterparts. The time-t conditional transform under the risk neutral measure is then given by

$$\psi^*(S;v,r,q,\tau) = \exp\Big(\alpha_r(S) + \alpha_q(S) + \alpha_v^*(S) + \beta_r(S)r + \beta_q(S)q + \beta_v^*(S)v\Big), \tag{4.A.14}$$

where α_r, α_q, β_r, and β_q are defined as in Equation (4.A.1) and

$$\alpha_v^* = -\frac{\kappa_v^*\bar{v}^*}{\sigma_v^2}\Big((\gamma_v^* + b^*)\tau + 2\ln\Big(1 - (\gamma_v^* + b^*)\frac{1-\exp(-\gamma_v^*\tau)}{2\gamma_v^*}\Big)\Big), \tag{4.A.15}$$

$$\beta_v^* = -\frac{a^*(1-\exp(-\gamma_v^*\tau))}{2\gamma_v^* - (\gamma_v^* + b^*)(1-\exp(-\gamma_v^*\tau))}, \tag{4.A.16}$$

with

$$\kappa_v^* = \kappa_v - \eta^v, \tag{4.A.17}$$

$$\bar{v}^* = \kappa_v \bar{v} / \kappa_v^*, \tag{4.A.18}$$

$$a^* = S(1-S) - 2\lambda \Big(\exp(S\mu_J^* + S^2\sigma_J^2/2) - 1 - S\mu^* \Big), \tag{4.A.19}$$

$$b^* = \sigma_v \rho S - \kappa_v^*, \tag{4.A.20}$$

$$\gamma_v^* = \sqrt{(b^*)^2 + a^* \sigma_v^2}. \tag{4.A.21}$$

Given the parameter estimates reported in Table 4.1 and the level of the underlying S, f_Q can again be obtained via numerical integration:

$$f_Q(S; v, r, q, \tau) = \frac{1}{2\pi} \int_{-\infty}^{\infty} \psi^*(iz; v, r, q, \tau) \exp(-izS) dz. \tag{4.A.22}$$

4.A.2 Alternative Maturities: Distributions

Fig. 4.A.1: Physical and Risk Neutral Distributions, 6 Months Horizon

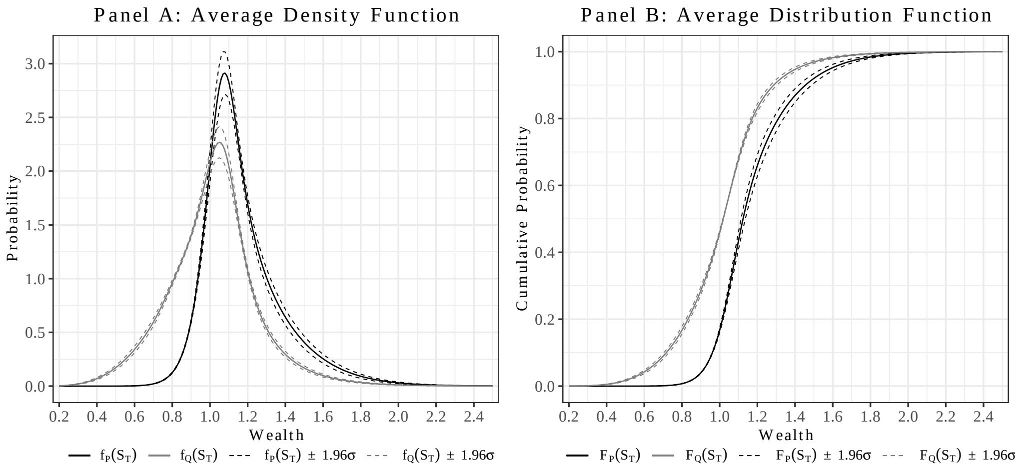

Fig. 4.A.1 plots physical and risk neutral densities (Panel A) and distribution functions (Panel B), estimated from the Pan (2002) stochastic volatility and jumps model and averaged over volatilities from 0.01 to 0.60. Physical (risk neutral) densities are denoted by $f_P(S_T)$ ($f_Q(S_T)$), whereas physical (risk neutral) distribution functions are denoted by $F_P(S_T)$ ($F_Q(S_T)$). We assume a return horizon of six months.

Fig. 4.A.2: Physical and Risk Neutral Distributions, 3 Months Horizon

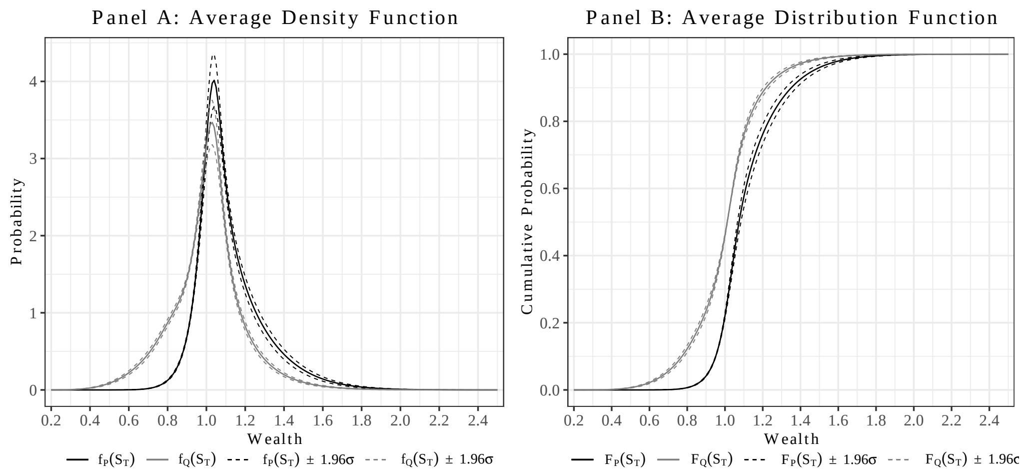

Fig. 4.A.2 plots physical and risk neutral densities (Panel A) and distribution functions (Panel B), estimated from the Pan (2002) stochastic volatility and jumps model and averaged over volatilities from 0.01 to 0.60. Physical (risk neutral) densities are denoted by $f_P(S_T)$ ($f_Q(S_T)$), whereas physical (risk neutral) distribution functions are denoted by $F_P(S_T)$ ($F_Q(S_T)$). We assume a return horizon of three months.

4.A.3 Alternative Maturities: Pricing Kernels

Fig. 4.A.3: Average Pricing Kernel, 6 Months Horizon

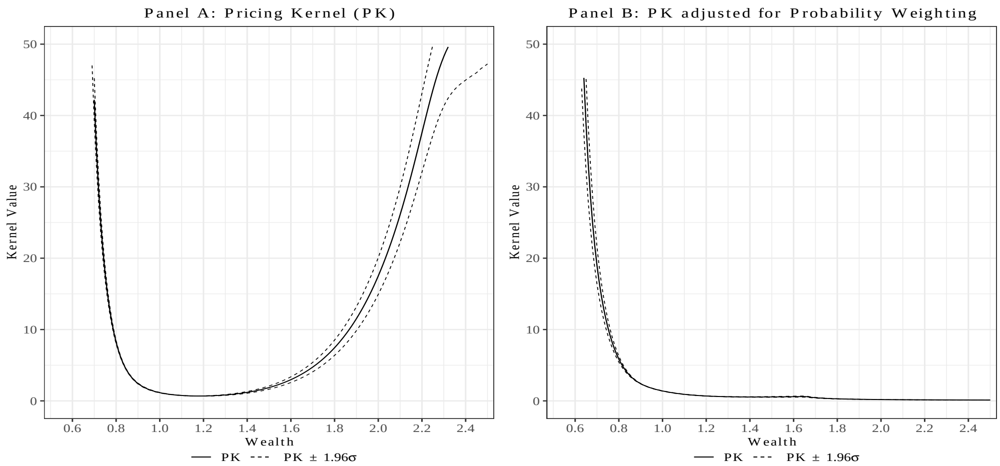

Fig. 4.A.3 plots pricing kernels estimated from the Pan (2002) stochastic volatility and jumps model. Following the literature (e.g. Jackwerth, 2000; Baele et al., 2019), we estimate the pricing kernel (PK, Panel A) as the ratio of the risk neutral to the physical probabilities, i.e $PK = f_Q/f_P$. In Panel B, we follow Equation (4.4) and calculate the pricing kernel, net of probability weighting, as $PK = f_Q(S_T)/f_P(S_T) \cdot w'(1 - F_P(S_T))$. We assume a return horizon of six months.

Fig. 4.A.4: Average Pricing Kernel, 3 Months Horizon

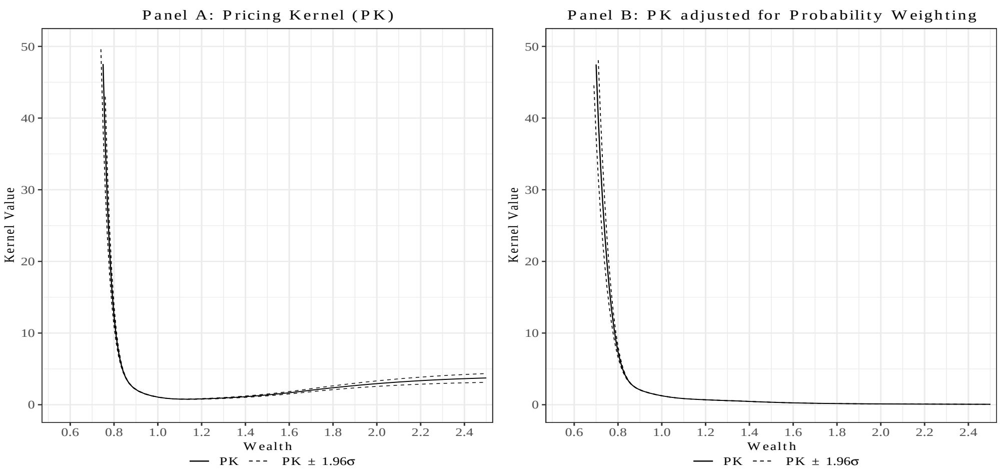

Fig. 4.A.4 plots pricing kernels estimated from the Pan (2002) stochastic volatility and jumps model. Following the literature (e.g. Jackwerth, 2000; Baele et al., 2019), we estimate the pricing kernel (PK, Panel A) as the ratio of the risk neutral to the physical probabilities, i.e $PK = f_Q/f_P$. In Panel B, we follow Equation (4.4) and calculate the pricing kernel, net of probability weighting, as $PK = f_Q(S_T)/f_P(S_T) \cdot w'(1 - F_P(S_T))$. We assume a return horizon of three months.

Chapter 5

Idiosyncratic Skewness and Market Timing of Capital Structure Decisions

This chapter refers to the working paper:

Dierkes, Maik and Jan Krupski (2022): 'Idiosyncratic Skewness and Market Timing of Capital Structure Decisions', Working Paper, Leibniz Universität Hannover.

Abstract

We investigate the impact of market timing on capital structure by employing idiosyncratic skewness as a proxy for firm-specific mispricing. Consistent with the market timing theory, idiosyncratic skewness is significantly positively related to equity issues, while the impact on debt issues is negative and of less importance. Moreover, we find equity issues to be accompanied by debt retirement programs. Contrasting the market timing theory, these effects are not persistent and vanish after about three years. Both the short-term and the long-term results are robust to a wide range of robustness checks. In line with Alti (2006), our results are therefore consistent with a modified version of the trade-off theory, including market timing as a short-term factor.

Keywords: Idiosyncratic Skewness, Market Timing, Persistence, Capital Structure

JEL Classification: G14, G32, G41.

5.1 Introduction

Building on the seminal study of Modigliani and Miller (1958) and market imperfections, there are three prevalent theories of capital structure. The pecking order theory of Myers and Majluf (1984) predicts that, due to asymmetric information, managers may refuse to issue equity because the costs to old shareholders outweigh a project's positive net present value. Firms therefore primarily fund investments with internal funds. If these are not sufficient, they prefer debt over equity issues. According to the trade-off theory, firms choose their optimal mix of debt and equity by balancing the costs (e.g. bankruptcy costs) and benefits (e.g. tax benefits) of debt.[1] At the optimum, the marginal benefit of one additional dollar of debt equals its marginal costs.[2] The market timing theory predicts that managers attempt to exploit temporary fluctuations in the cost of equity and therefore issue (repurchase) equity when shares are perceived to be overvalued (undervalued). Consequently, proxies for misvaluation should be positively related to equity issues.[3] According to Baker and Wurgler (2002), market timing should have a long-lasting impact.

We provide new evidence on the implications of firm-specific mispricing on financing decisions and capital structure. We employ idiosyncratic skewness as a proxy for mispricing and find a strong effect of market timing in the short run. More precisely, idiosyncratic skewness is significantly positively related to equity issues and negatively related to

[1] See Lemmon and Zender (2010).
[2] See Fama and French (2002).
[3] See Baker and Wurgler (2002) and Elliott et al. (2008).

debt issues, with the former effect being the predominant one.[4] Notably, equity issuance is accompanied by debt retirement programs, resulting in a negative impact on both the leverage level and changes in leverage. However, in contrast to the predictions of the market timing theory, we find that the effect is not persistent and vanishes after about three years. Consequently, our results are consistent with a long-term validity of the trade-off theory, including market timing as a short-term factor.[5]

Our study relates to the vast amount of literature on capital structure theory. For example, Lemmon and Zender (2010) show that, after accounting for debt capacity, the pecking order theory appropriately describes financing behavior. Moreover, Shyam-Sunder and Myers (1999) conclude that the pecking order is of first-order importance. Contrary to the pecking order theory, Frank and Goyal (2003) find that there are more equity issues than debt issues and that these track the financing deficit more closely (especially in the 1990s).[6]

Supporting the trade-off theory, Hovakimian et al. (2001) find that firms adjust their capital structure based on a target debt ratio and Fama and French (2002) show that leverage is mean reverting. These results align well with survey evidence from Graham and Harvey (2001) and Graham (2022) who find that 81% and 72% of firms follow some type of target leverage, respectively. Flannery and Rangan (2006) and Warr et al. (2012) find that adjustment to these targets proceeds at a fast rate of around 35% per year, implying a half-life of deviations from target

[4] This finding is consistent with Dong et al. (2012).
[5] See Alti (2006).
[6] See also Huang and Ritter (2009), Welch (2004), and Byoun (2008).

of only 1.6 years. In contrast, Fama and French (2002) document a rate of 10-18% (for book leverage), while Huang and Ritter (2009) report an annual adjustment of 17%.[7] Further studies find the speed of adjustment to depend on the relative position to the target (Byoun, 2008), adjustment costs (Leary and Roberts, 2005; Korteweg et al., 2022), and equity mispricing (Warr et al., 2012).

Contrasting the trade-off theory, Rajan and Zingales (1995) and Hovakimian et al. (2001) find that firms typically issue stock when their valuation is high and thus the perceived cost of equity is low.[8] In line with this, Graham and Harvey (2001) find that overvaluation is one of the main reasons to issue equity, while managers are reluctant to issue shares when they consider their stock to be undervalued. Baker and Wurgler (2002) conclude that capital structure is the cumulative outcome of past attempts to time the equity market. Firms issue (repurchase) equity instead of debt when they perceive their market value - in terms of the market-to-book ratio - to be high (low).[9] Long-term returns of equity issuers are usually negative, suggesting that market timing, on average, is successful. Baker and Wurgler (2002) propose two versions of market timing. First, a dynamic version of Myers and Majluf (1984) with rational managers and investors, where the extent of adverse selection

[7] Hovakimian and Li (2011) find a rate of only 5-8% and Kayhan and Titman (2007) conclude that a firm's history influences observed debt ratios for at least ten years. Chang and Dasgupta (2009) argue that the results of prior studies can be reproduced by random simulations, but Huang and Ritter (2009) and DeAngelo and Roll (2015) refute this finding.

[8] See also Kayhan and Titman (2007) and Huang and Ritter (2009).

[9] See also Hovakimian et al. (2001), Elliott et al. (2008), and Warusawitharana and Whited (2016).

varies across firms and time. Second, irrational investors (or managers) and time-varying (perceived) mispricing. Importantly, for this version to hold, the market does not necessarily need to be inefficient. Instead, it is sufficient that managers *believe* they can time the market. Baker and Wurgler (2002) consider this explanation to be more likely. Subsequently, there has been a substantial amount of studies to study the impact of equity mispricing on financing decisions and capital structure. Elliott et al. (2008) confirm that mispricing is driven by investors' irrationality rather than information asymmetries and Alti and Sulaeman (2012) find that equity issuance depends on both a high valuation and high institutional investor demand. Dong et al. (2012) show that overvaluation is more important than undervaluation and Kim and Weisbach (2008) and Bolton et al. (2013) conclude that overvaluation implies equity issues even if firms do not have financial needs.[10]

Contrasting the market timing theory, Flannery and Rangan (2006) and Alti (2006) show that firms subsequently rebalance away from the influence of market timing decisions by issuing more debt and less equity. As a result, the impact of market timing disappears after only a few years. Hovakimian (2006) and Mahajan and Tartaroglu (2008) confirm this finding for the U.S. and international equity markets, respectively.

Our study contributes to both the extensive literature on mispricing proxies and the persistence of market timing. To the best of our knowledge, we are the first to employ idiosyncratic skewness – as a proxy for

[10] Moreover, DeAngelo et al. (2010), Khan et al. (2012), and Dittmar et al. (2020) find that high valuations have a significantly positive impact on the probability for seasoned equity offerings (SEOs).

mispricing – in a broad study of financing decisions and capital structure. We motivate using this measure for several reasons. Most importantly, Barberis and Huang (2008) show that investors with prospect theory preferences demand securities with highly right-skewed payoffs (such as IPOs), which induces them to hold more concentrated portfolios and leads to an overvaluation of securities with a skewed return distribution.[11] In line with this prediction, Boyer et al. (2010), Conrad et al. (2013), and Chang et al. (2013) find a negative relation between skewness and subsequent average returns.[12] Moreover, Autore and DeLisle (2016) find idiosyncratic skewness to significantly predict returns following an SEO and Green and Hwang (2012) document that IPOs with high expected idiosyncratic skewness (based on their industry) earn high first-day (low long-term) returns.[13] Notably, they also find that expected idiosyncratic skewness predicts future return skewness. Consequently, firms recognizing the mispricing will exploit it by issuing more equity, resulting in a lower leverage. To investigate the impact of firm-specific mispricing on financing decisions and capital structure, we therefore employ the skewness measure of Green and Hwang (2012) and replace industry returns by firm-specific returns.

Our results provide further evidence for the implications of market timing. After accounting for control variables motivated by the literature,

[11] See also Mitton and Vorkink (2007), Brunnermeier and Parker (2005), and Brunnermeier et al. (2007).

[12] These studies employ the expected idiosyncratic skewness, the risk neutral idiosyncratic skewness, and the risk neutral market skewness, respectively. See also Kumar et al. (2022) who find that skewness significantly predicts the mispricing component of a combined measure of eleven prominent anomalies.

[13] See also Dierkes et al. (2022).

we find idiosyncratic skewness to be significantly positively related to equity issues, while the impact on debt issues is negative. However, in line with Dong et al. (2012), the effect on equity issues is more important, both economically and statistically ($t = 8.08$ vs. $t = -3.01$).[14] Additionally, we confirm Hovakimian et al. (2001) as we find equity issues to be accompanied by debt retirement programs. As a result, idiosyncratic skewness negatively predicts both the leverage level and changes in leverage. Furthermore, there is a significantly positive impact on the change in cash holdings, suggesting that firms do not issue equity to meet financial needs.[15] While we base our main analysis on book leverage, our results also apply to market leverage and several robustness checks motivated by the literature (e.g. subperiods, size splits, Fama and MacBeth (1973) regressions, an alternative skewness measure, and alternative leverage definitions). We therefore conclude that our results are consistent with a strong impact of market timing in the short run.

However, in stark contrast to the market timing theory, we find that the impact on leverage (leverage changes) completely vanishes after only three (four) years.[16] In line with this finding, we find an annual speed of adjustment of 32%, implying a half-life of deviations from target leverage of only 1.8 years.[17] Notably, rebalancing is not only driven by more debt issues but also by reduced debt retirements. Again, our conclusions are

[14] We derive the same conclusion when considering logit regressions ($t = 6.86$ vs. $t = -3.76$).

[15] This finding is in line with Kim and Weisbach (2008) and Bolton et al. (2013).

[16] As there is a positive drift in leverage, we adjust changes in leverage for firm fixed effects.

[17] Following Flannery and Rangan (2006), we include firm fixed effects.

robust to several alternative specifications. In line with Alti (2006), we therefore conclude that the impact of firm-specific mispricing – and thus market timing – is not persistent. Instead, our results support the idea of a modified version of the trade-off theory, including market timing as a short-term factor.

5.2 Data and Methodology

5.2.1 Idiosyncratic Skewness

Following Green and Hwang (2012), we measure idiosyncratic skewness (and thus firm-specific mispricing) as

$$Skew_{i,t} = \frac{(P_{99} - P_{50}) - (P_{50} - P_1)}{(P_{99} - P_1)}, \tag{5.1}$$

where P_j denotes the j'th percentile of the daily log return distribution of firm i over the last 250 trading days preceding the end of its fiscal year.[18] Green and Hwang (2012) favor this measure over the traditional third moment of skewness because it is solely based on the tail of the distribution and thus better captures the idea that skewness-seeking investors primarily care about tail events. Realizations of $Skew_{i,t}$ are positive if the right tail of the distribution (P_{99}) is further away from the median (P_{50}) than the left tail (P_1) and negative if the opposite is true. The denominator controls for the dispersion of the firm-specific return distribution and normalizes estimates to values between -1 and 1.

[18] We obtain returns from the Center for Research in Security Prices (CRSP) and remove stocks with a share code other than 10 or 11.

However, note that our estimation approach differs from Green and Hwang (2012) in two important ways. First, instead of cross-sectional returns within firm i's industry, we employ the return distribution of the firm itself. Even though Green and Hwang (2012) argue that their measure captures *idiosyncratic* skewness, it rather represents the skewness of stocks that are comparable to firm i. As we aim to capture *firm-specific* mispricing, it is therefore natural to make use of firm-specific returns. Second, instead of monthly log returns over the three months preceding the month of the IPO, we employ daily log returns over the 250 trading days preceding the end of the fiscal year. The reasoning for this adjustment is straightforward. On the one hand, it allows us to capture mispricing throughout the firm's fiscal year and, on the other hand, using monthly returns would lead to estimates that are not meaningful. In the following, we drop the firm index i.

5.2.2 Sample Construction

Our initial sample is based on all firms that are covered by Compustat between 1971 to 2020.[19] In line with previous studies, we remove financial firms with SIC codes between 6000 and 6999 and firms with a book value of assets below ten million dollars. Following Hovakimian et al. (2001), we define book debt as short-term debt (item 34) plus long-term debt (item 9).[20] Book equity is defined as total assets (item 6) minus total liabilities (item 181), preferred stock (item 10), and deferred taxes (item

[19] We choose this sample period because cash flow data becomes available in 1971.

[20] See also Hovakimian (2006), Flannery and Rangan (2006), Rajan and Zingales (1995), Byoun (2008), Hovakimian and Li (2011), and Warr et al. (2012).

35).[21] Market equity equals the share price (item 24) times the number of outstanding shares (item 25). Book leverage (Lev_t) is defined as book debt divided by total assets and market leverage ($Lev_{Mkt,t}$) is book debt divided by total assets minus book equity plus market equity. Finally, ΔLev_t ($\Delta Lev_{Mkt,t}$) denotes the change in book leverage (market leverage) from the end of the previous fiscal year ($t-1$) to the end of the current fiscal year (t). We choose this approach because it excludes non-debt liabilities which, according to Hovakimian (2006), are not a good indicator of whether a firm is at risk of default. Moreover, Kayhan and Titman (2007) note that a broader definition of leverage (as in Baker and Wurgler, 2002) likely overstates the financial leverage.

Throughout most of our empirical analysis, we follow Baker and Wurgler (2002) and Alti (2006) and focus on book leverage. We do this for two reasons. First, the calculation of idiosyncratic skewness is return-based and the relationship with market leverage is thus, in part, mechanical. Second, book leverage reflects active rebalancing (e.g. through issues and repurchases), whereas market leverage includes factors that are not under the control of the firm (e.g. stock returns and option exercising).[22] However, to verify our results, we also investigate the impact of idiosyncratic skewness on market leverage. Additionally, our robustness analysis includes the leverage approach of Baker and Wurgler (2002).[23] We drop observations for which either of the two leverage ratios is below zero or

[21] If missing, preferred stock is replaced by the redemption value of preferred stock (item 56).
[22] See Chang and Dasgupta (2009).
[23] We refer to their study for more details on the estimation procedure.

above one and we require at least two observations per firm.[24] Our final sample comprises 110,782 firm-year observations and displays an average leverage of 24.3% (with a standard deviation of 19.2%).

Raw equity issues correspond to the change in book equity minus the change in retained earnings (item 36), divided by total assets. Accordingly, raw debt issues are defined as the change in book debt divided by total assets. Following Hovakimian et al. (2001), we define substantial equity (e/A_t) and debt issues (d/A_t) as those exceeding 5% of total assets.[25] Likewise, substantial repurchases are defined as *negative* equity ($-e/A_t$) and debt issues ($-d/A_t$) exceeding 1.25% and 5% of total book assets, respectively.[26] The equity share in new issues (ES_t) is the ratio of equity issues to total issues. However, to reduce the number of missing values, we base the equity share on raw issues.[27] Finally, the change in cash holdings ($\Delta Cash_t/A_t$) corresponds to the change in cash and short-term investments (item 1) and investments (Inv_t/A_t) are equivalent to capital expenditures (item 128). Both measures are scaled by total assets.

5.2.3 Control Variables

The choice of control variables follows the literature, most importantly Rajan and Zingales (1995), Fama and French (2002), and Alti (2006). Firm size ($SIZE_t$) is the natural logarithm of sales in million dollars

[24] See Baker and Wurgler (2002).
[25] See Leary and Roberts (2005) and Hovakimian (2006).
[26] See Leary and Roberts (2005). Our conclusions remain unchanged if equity repurchases are defined according to a threshold of 5%.
[27] We obtain 85,593 firm-year observations. If we included the threshold, this number would shrink to 38,437 observations. Our conclusions are not affected by this choice.

(item 12) and should be positively related to leverage since larger firms are more diversified and have a better reputation in debt markets.[28] Asset tangibility (TNG_t) is defined as net property, plant, and equipment (item 8) divided by total assets and should positively predict leverage as tangible assets may serve as collateral. Profitability (PRF_t) is measured as EBITDA (item 13) divided by total assets. According to the pecking order theory, there should be a negative impact on leverage because firms prefer internal funds. With respect to the trade-off theory, however, the impact should be positive as suppliers are more willing to lend to profitable firms. The market-to-book ratio (MB_t) is defined as market value (total assets minus book equity plus market equity) divided by book assets. MB_t may either serve as a proxy for mispricing or growth opportunities. In both cases, the impact on leverage should be negative. Following Baker and Wurgler (2002) and Alti (2006), we drop market-to-book ratios larger than ten. Finally, R&D expenses (RD_t) are defined as research and development expense (item 46, replaced by zero when missing) divided by total assets and RDD_t is a dummy variable that takes the value one if R&D expenses are missing.[29] According to Hovakimian et al. (2001), RD_t also corresponds to future growth opportunities and should therefore be negatively related to leverage.

[28] If sales are below one million dollar, we assign a value of zero.
[29] See Alti (2006).

5.3 Results

5.3.1 Short-Term Impact of Market Timing

Based on the previous literature on market timing and the pricing implications of skewness, we expect a strong impact of idiosyncratic skewness on financing decisions, and thus capital structure, in the short run. To have a first look at this prediction, Table 5.1 presents summary statistics for issuance decisions and leverage ratios, where firms are sorted into quintiles based on the idiosyncratic skewness of their returns in fiscal year t. Since idiosyncratic skewness acts as a proxy for overvaluation, Quintile 1 comprises the most *undervalued* firms and Quintile 5 contains the most *overvalued* firms. We report the equal-weighted average in a given quintile as well as full sample means. Although some firm-year observations have a skewness of close to -1 or 1 (not reported), the majority of firms displays values between -0.14 (Quintile 1) and 0.22 (Quintile 5).

Most importantly, we find a strong pattern for all measures under consideration. Average equity issues (e/A_t, Row 6) increase from 3.60% of total book assets for the most undervalued firms to 5.31% for the most overvalued ones.[30] Notably, the difference of 1.71 percentage points is highly significant ($t = 11.05$) and mostly driven by overvaluation. Moreover, in line with Dong et al. (2012), we find that equity issues are more sensitive to mispricing than debt issues (d/A_t, Row 7). Nevertheless, market timing in debt issues is still significant at the 1%-level ($t = -3.18$) and mostly driven by undervaluation. As a result, the equity share in

[30] Average equity issues are below 5% since non-substantial issues are treated as zeros.

Table 5.1: Summary Statistics

	Skewness Quintile							
	Full Sample	1	2	3	4	5	High – Low	t–value
(1) $Skew_t$	0.04	−0.14	−0.02	0.04	0.11	0.22	0.35	
(2) Lev_t	24.68	25.30	24.78	24.48	24.44	24.40	−0.89***	(−4.72)
(3) $Lev_{Mkt,t}$	21.53	22.01	21.84	21.25	21.20	21.33	−0.68***	(−3.60)
(4) ΔLev_t	0.44	1.20	0.63	0.49	0.22	−0.38	−1.58***	(−16.75)
(5) $\Delta Lev_{Mkt,t}$	0.61	2.72	1.35	0.71	−0.09	−1.70	−4.42***	(−49.96)
(6) e/A_t	4.18	3.60	3.63	4.10	4.29	5.31	1.71***	(11.05)
(7) d/A_t	3.62	3.80	3.57	3.60	3.58	3.53	−0.26***	(−3.18)
(8) ES_t	57.61	53.31	56.32	57.87	58.76	61.75	8.44***	(18.01)
(9) $-e/A_t$	1.01	1.31	0.98	1.03	0.89	0.86	−0.45***	(−5.54)
(10) $-d/A_t$	1.89	1.60	1.76	1.75	1.99	2.35	0.75***	(9.10)

Table 5.1 presents summary statistics for financing decisions and leverage ratios, where firms are sorted into quintiles based on their idiosyncratic skewness in fiscal year t. We report the equal-weighted average in a given quintile as well as full sample means. The notation and construction of leverage ratios and issuance decisions follows Section 5.2.2. Except for $Skew_t$, values are stated in percent and – in the High–Low column – stars indicate significance at the 10% (*), 5% (**), and 1% (***) level. We cover a sample period from 1971 to 2020.

new issues (ES_t, Row 8) substantially increases in idiosyncratic skewness (t = 18.01). Repurchase decisions mirror these findings. Equity repurchases ($-e/A_t$, Row 9) significantly decrease in skewness ($t = -5.54$) and the effect is mostly driven by undervaluation. This finding is well in line with Baker and Wurgler (2013) who note that overvaluation is a motive for equity issues, while undervaluation facilitates repurchases.[31] In contrast, the impact on debt retirements ($-d/A_t$, Row 10) is positive and largely driven by overvaluation (t = 9.10), suggesting that equity

[31] See also Warusawitharana and Whited (2016).

issues are accompanied by debt retirement programs.[32] Consequently, we find a significantly negative impact on both book leverage (Lev_t, Row 2) and market leverage ($Lev_{Mkt,t}$, Row 3). At a first glance, it is somewhat surprising that the effect on the change in leverage (ΔLev_t, Row 4) is even stronger. However, we attribute this finding to the positive drift in leverage of around 0.44 percentage points per year.

So far, our results provide evidence for a strong market timing effect in the short run. To study this finding in a multivariate setting, Table 5.2 adds several control variables motivated by the literature (as outlined in Section 5.2.3) and industry fixed effects.[33] We adjust standard errors for both firm and year clusters. Following Alti (2006), among others, we lag control variables by one year as contemporaneous controls may be noisy. Moreover, for the sake of brevity, we focus our analysis on the impact of idiosyncratic skewness.[34]

As documented in Table 5.1, the market timing effect is driven by substantial over- and undervaluation, which is why we evaluate the economic impact of moving from the lowest ($Skew_t = -0.14$) to the highest skewness quintile ($Skew_t = 0.22$). Our results confirm that companies time their issuance and repurchase decisions based on idiosyncratic skewness (and thus equity mispricing). Most importantly, the impact on leverage (Models 1 and 2), the change in leverage (Models 3 and 4), and debt issues

[32] This finding is in line with Hovakimian et al. (2001).

[33] In line with Alti (2006), Models (3) to (7) additionally control for the lagged leverage.

[34] In Model (1), which constitutes the most common leverage regression, all controls are highly significant and signs are in line with the expectations outlined in Section 5.2.3. The negative sign of PRF_{t-1} contrasts the trade-off theory, but is a common empirical finding.

Table 5.2: Short-Term Impact of Market Timing on Capital Structure

	Dependent variable:						
	Lev_t	$Lev_{Mkt,t}$	ΔLev_t	$\Delta Lev_{Mkt,t}$	e/A_t	d/A_t	ES_t
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$Skew_t$	−0.04***	−0.08***	−0.04***	−0.12***	0.04***	−0.01***	0.20***
	(−3.97)	(−5.97)	(−11.47)	(−13.55)	(8.08)	(−3.01)	(7.76)
$SIZE_{t-1}$	0.01***	0.0002	−0.0001	−0.003***	−0.01***	−0.004***	−0.01***
	(5.53)	(0.09)	(−0.26)	(−3.15)	(−10.23)	(−8.63)	(−3.09)
TNG_{t-1}	0.20***	0.19***	0.02***	0.02***	0.02***	0.005	−0.07***
	(18.30)	(17.97)	(5.35)	(5.36)	(5.24)	(1.63)	(−3.35)
PRF_{t-1}	−0.21***	−0.19***	−0.02***	−0.01	−0.28***	0.01	−0.15***
	(−8.50)	(−7.46)	(−3.21)	(−1.23)	(−14.75)	(1.57)	(−5.46)
MB_{t-1}	−0.02***	−0.04***	0.0004	0.0001	0.03***	0.01***	0.02***
	(−12.63)	(−17.25)	(0.75)	(0.17)	(17.33)	(9.91)	(6.57)
RD_{t-1}	−0.32***	−0.33***	−0.06***	−0.07***	0.33***	−0.03***	0.36***
	(−8.35)	(−8.77)	(−6.88)	(−7.03)	(11.91)	(−4.22)	(8.81)
RDD_{t-1}	0.03***	0.03***	0.003***	0.005***	0.01***	0.01***	−0.02***
	(6.04)	(5.70)	(3.39)	(4.20)	(6.94)	(5.44)	(−3.05)
Lev_{t-1}			−0.13***		0.03***	0.03***	−0.09***
			(−28.12)		(5.53)	(8.60)	(−4.18)
$Lev_{Mkt,t-1}$				−0.12***			
				(−12.43)			
SIC fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Cluster adj.	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Adjusted R^2	0.24	0.34	0.06	0.08	0.36	0.03	0.07

Table 5.2 presents results for OLS regressions of leverage ratios and financing decisions on idiosyncratic skewness ($Skew_t$) and several control variables motivated by the literature. In line with Baker and Wurgler (2002) and Alti (2006), we lag control variables by one year as contemporaneous controls may be noisy. The notation and construction of dependent variables and controls follows Sections 5.2.2 and 5.2.3, respectively. Stars indicate significance at the 10% (*), 5% (**), and 1% (***) level and t-values (in parentheses) are based on cluster-adjusted standard errors. We cover a sample period from 1971 to 2020 and do not report the intercept.

(Model 6) remains significantly negative (at the 1%-level), while equity issues (Model 5, $t = 8.08$) and the equity share (Model 7, $t = 7.76$) are significantly positively affected. Note that the impact on book leverage (Models 1 and 3) is not driven by returns and thus only depends on active financing decisions caused by mispricing (and retained earnings). Importantly, our results are also economically powerful. In Model (1), the coefficient estimate of -0.04 implies that moving from an undervalued firm (in the first skewness quintile) to an overvalued firm (in the fifth skewness quintile) reduces book leverage by $0.04 \times 0.35 = 1.40$ percentage points (which exceeds the univariate impact). With the same coefficient estimate of -0.04, the economic impact on ΔLev_t is comparable (Model 3). Due to the mechanical relationship between idiosyncratic skewness and market leverage, the economic impact in Models (2) and (4) increases to 2.80 and 4.20 percentage points, respectively. Apart from that, overvaluation increases equity issues (Model 5) by 1.40 percentage points, whereas the impact on debt issues (Model 6) is economically negligible (-0.35 percentage points). Finally, the equity share in new issues (Model 7) increases by economically important 7.00 percentage points. In unreported results, we repeat this framework for repurchases. In line with Table 5.1, we find that $Skew_t$ is significantly negatively (positively) related to equity repurchases (debt retirements). However, economically, the issuance decision is more important.[35]

Finally, we investigate the impact of idiosyncratic skewness on the probability for issues and repurchases by performing independent logit regressions (i.e. the debt-equity choice is evaluated separately). In line with Section 5.2.2, we define a firm as issuing equity (debt) if equity issues (debt issues) exceed 5% of total assets. A firm is defined as repurchasing equity (retiring debt) if repurchases exceed 1.25% (5%) of total assets.

[35] The corresponding t-statistics are -3.71 and 3.13, respectively.

Table 5.3: Short-Term Impact of Market Timing on the Probability for Issues and Repurchases

	Dependent variable:			
	Issues		Repurchases	
	$P(e_t)$	$P(d_t)$	$P(-e_t)$	$P(-d_t)$
	(1)	(2)	(3)	(4)
$Skew_t$	0.99***	−0.48***	−0.96***	0.58***
	(6.86)	(−3.76)	(−6.83)	(4.62)
$SIZE_{t-1}$	−0.13***	−0.07***	0.23***	0.005
	(−9.88)	(−4.77)	(11.85)	(0.31)
TNG_{t-1}	0.18*	0.59***	−0.93***	−0.25**
	(1.82)	(7.20)	(−7.21)	(−2.55)
PRF_{t-1}	−2.53***	0.33**	1.01***	−1.55***
	(−11.19)	(2.05)	(2.64)	(−6.02)
MB_{t-1}	0.47***	0.06***	−0.03	−0.30***
	(22.01)	(4.15)	(−1.10)	(−10.64)
RD_{t-1}	3.43***	−1.49***	−1.98***	−2.01***
	(7.94)	(−4.51)	(−4.05)	(−5.51)
RDD_{t-1}	0.10***	0.13***	−0.04	0.08***
	(2.99)	(4.41)	(−1.02)	(2.95)
SIC fixed effects	Yes	Yes	Yes	Yes
Cluster adj.	Yes	Yes	Yes	Yes
Pseudo R^2	0.18	0.02	0.05	0.24

Table 5.3 presents results for logit regressions of issuance and repurchase decisions on idiosyncratic skewness ($Skew_t$) and several control variables motivated by the literature. In line with Baker and Wurgler (2002) and Alti (2006), we lag control variables by one year as contemporaneous controls may be noisy. The notation and construction of controls follows Section 5.2.3. We define substantial equity (e/A_t) and debt issues (d/A_t) as those exceeding 5% of total book assets. Substantial repurchases are defined as *negative* equity ($-e/A_t$) and debt issues ($-d/A_t$) exceeding 1.25% and 5% of total book assets, respectively (Leary and Roberts, 2005). Stars indicate significance at the 10% (*), 5% (**), and 1% (***) level and t-values (in parentheses) are based on cluster-adjusted standard errors. We cover a sample period from 1971 to 2020 and do not report the intercept.

Table 5.3 presents results. In line with our previous results, the probability for equity issues ($t = 6.86$) and debt retirements ($t = 4.62$) increases in idiosyncratic skewness, while the impact on debt issues ($t = -3.76$)

and equity repurchases ($t = -6.83$) is significantly negative.[36] In terms of marginal effects (not reported), the probability to issue equity (debt) increases (decreases) by economically important $0.35 \times 0.118 = 4.13$ (3.08) percentage points when comparing an overvalued firm (in the fifth skewness quintile) to an undervalued firm (in the first skewness quintile). In contrast, the probability for equity repurchases (debt retirement) decreases (increases) by 3.96 (2.28) percentage points.

Taken together, we find that idiosyncratic skewness – and thus mispricing – plays an important role in both issuance and repurchase decisions. When stocks are overvalued, firms issue significantly more equity (less debt) and retire more debt (repurchase less equity). As a result, idiosyncratic skewness is significantly negatively related to both the leverage level and changes in leverage. Our results thus confirm a strong market timing effect in the short run.

5.3.2 Persistence of Market Timing

In contrast to the literature on short-term implications of market timing, findings on the persistence of market timing are less conclusive. While, for example, Baker and Wurgler (2002) document a long-lasting effect, Alti (2006) and Hovakimian (2006) conclude that market timing is not persistent. In this section, we therefore aim to shed some light on this yet open research question. In Table 5.4, we provide first evidence by investigating the impact of idiosyncratic skewness on the level of book leverage in fiscal years $t+1$ to $t+4$.

[36] Our findings are also in line with Hovakimian et al. (2001) and Elliott et al. (2008).

Table 5.4: Persistence of Market Timing - Leverage Level

	Dependent variable:			
	Lev_{t+1}	Lev_{t+2}	Lev_{t+3}	Lev_{t+4}
	(1)	(2)	(3)	(4)
$Skew_t$	−0.03***	−0.02**	−0.01	−0.004
	(−3.90)	(−2.20)	(−1.30)	(−0.46)
$SIZE_{t-1}$	0.01***	0.01***	0.01***	0.01***
	(5.13)	(5.00)	(5.12)	(5.15)
TNG_{t-1}	0.18***	0.17***	0.16***	0.15***
	(15.65)	(14.16)	(12.96)	(12.02)
PRF_{t-1}	−0.19***	−0.17***	−0.16***	−0.16***
	(−7.65)	(−6.80)	(−6.24)	(−6.18)
MB_{t-1}	−0.02***	−0.01***	−0.01***	−0.01***
	(−10.03)	(−9.33)	(−8.73)	(−7.76)
RD_{t-1}	−0.31***	−0.28***	−0.26***	−0.27***
	(−7.53)	(−6.74)	(−5.75)	(−5.93)
RDD_{t-1}	0.03***	0.03***	0.02***	0.02***
	(5.79)	(5.45)	(4.83)	(4.58)
SIC fixed effects	Yes	Yes	Yes	Yes
Cluster adj.	Yes	Yes	Yes	Yes
Adjusted R^2	0.22	0.21	0.20	0.20

Table 5.4 presents results for OLS regressions of future leverage ratios on idiosyncratic skewness ($Skew_t$) and several control variables motivated by the literature. In line with Baker and Wurgler (2002) and Alti (2006), we lag control variables by one year as contemporaneous controls may be noisy. The notation and construction of leverage ratios and controls follows Sections 5.2.2 and 5.2.3, respectively. Stars indicate significance at the 10% (*), 5% (**), and 1% (***) level and t-values (in parentheses) are based on cluster-adjusted standard errors. We cover a sample period from 1971 to 2020 and do not report the intercept.

While both the coefficient estimates and the statistical significance of control variables remain largely unchanged, the impact of idiosyncratic skewness gradually decreases from Lev_{t+1} to Lev_{t+4}. Although the impact on Lev_{t+1} (Model 1) remains highly significant and is comparable to the short-term effect ($t = -3.90$), statistical significance with respect to Lev_{t+2} (Model 2) already reduces to the 5%-level ($t = -2.20$) and the coefficient estimate is cut in half. After only three years (Model 3), significance disappears completely ($t = -1.30$). Finally, in Model (4), both the economic and the statistical effect are close to zero. Our results are thus in line with Alti (2006), Hovakimian (2006), and Mahajan and Tartaroglu (2008) who also find that the impact of market timing is not persistent.[37]

In Table 5.5, we follow Alti (2006) and replace leverage levels by the cumulative change in leverage from fiscal year $t-1$ to $t+\tau$.[38] In Panel A, we employ the control variables outlined in Section 5.2.3 as well as industry fixed effects. Again, the impact of idiosyncratic skewness gradually decreases from Model (1) to Model (4). However, contrasting the results in Table 5.4, the economic significance in $t+1$ (Model 1) even slightly increases. Moreover, the impact on the cumulative change in leverage from $t-1$ to $t+4$ (Model 4) remains significant at the 5% level ($t = -1.97$) and finally turns insignificant in $t+6$ (not reported). At a first glance, this finding implies a higher persistence of market timing effects than previously documented. However, as outlined in Section 5.3.1, there is a positive drift in Lev_t, which is why a long-lasting impact on $\Delta Lev_{t+\tau}$ does not necessarily imply that market timing effects persist. In Table 5.A.1 (reported in the Appendix), we provide summary statistics for the cumulative change in leverage. Most importantly, we find that the drift increases in τ and exceeds the impact of idiosyncratic skewness after only two years. As a result, the change in leverage remains positive even if the highest skewness quintile is considered. In Panel B of Table 5.5, we

[37] In unreported results, we compute the external-finance-weighted average idiosyncratic skewness in the spirit of Baker and Wurgler (2002) and do not find a significant impact.
[38] τ denotes the number of fiscal years ahead.

Table 5.5: Persistence of Market Timing - Change in Leverage

Panel A: Industry Fixed Effects				
	Dependent variable:			
	ΔLev_{t+1}	ΔLev_{t+2}	ΔLev_{t+3}	ΔLev_{t+4}
	(1)	(2)	(3)	(4)
$Skew_t$	-0.05^{***}	-0.03^{***}	-0.02^{***}	-0.02^{**}
	(-9.39)	(-3.56)	(-3.25)	(-1.97)
Controls	Yes	Yes	Yes	Yes
SIC fixed effects	Yes	Yes	Yes	Yes
Firm fixed effects	No	No	No	No
Cluster adj.	Yes	Yes	Yes	Yes
Adjusted R^2	0.06	0.07	0.11	0.13

Panel B: Firm Fixed Effects				
	Dependent variable:			
	ΔLev_{t+1}	ΔLev_{t+2}	ΔLev_{t+3}	ΔLev_{t+4}
	(1)	(2)	(3)	(4)
$Skew_t$	-0.04^{***}	-0.02^{***}	-0.01^{**}	-0.01
	(-7.82)	(-3.18)	(-2.06)	(-1.00)
Controls	Yes	Yes	Yes	Yes
SIC fixed effects	No	No	No	No
Firm fixed effects	Yes	Yes	Yes	Yes
Cluster adj.	Yes	Yes	Yes	Yes
Adjusted R^2	0.30	0.35	0.38	0.44

Table 5.5 presents results for OLS regressions of the change in leverage (from $t-1$ to $t+\tau$) on idiosyncratic skewness ($Skew_t$) and several control variables motivated by the literature. In line with Baker and Wurgler (2002) and Alti (2006), we lag control variables by one year as contemporaneous controls may be noisy. The notation and construction of leverage ratios and controls follows Sections 5.2.2 and 5.2.3, respectively. For brevity, we only report results for $Skew_t$. In Panel A, we account for industry fixed effects, while Panel B controls for firm fixed effects. Stars indicate significance at the 10% (*), 5% (**), and 1% (***) level and t-values (in parentheses) are based on cluster-adjusted standard errors. We cover a sample period from 1971 to 2020 and do not report the intercept.

therefore control for firm fixed effects in $\Delta Lev_{t+\tau}$ Not surprisingly, the persistence of idiosyncratic skewness is clearly reduced and comparable to Table 5.4. More precisely, both the economic and the statistical impact in $t+1$ (Model 1) are in line with the short-term results and $Skew_t$ turns insignificant in $t+4$ (Model 4, $t=-1.00$).

Table 5.6: Long-Term Impact of Market Timing on Issues and Repurchases

Panel A: Issues and the Equity Share

	Dependent variable:					
	e_{t+1}/A_t	e_{t+2}/A_t	d_{t+1}/A_t	d_{t+2}/A_t	ES_{t+1}	ES_{t+2}
	(1)	(2)	(3)	(4)	(5)	(6)
$Skew_t$	0.04***	0.01*	0.01**	0.01***	0.06*	−0.05*
	(5.34)	(1.83)	(2.45)	(3.89)	(1.78)	(−1.90)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
SIC fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Cluster adj.	Yes	Yes	Yes	Yes	Yes	Yes
Adjusted R^2	0.27	0.24	0.02	0.02	0.06	0.05

Panel B: Repurchases

	Dependent variable:			
	$-e_{t+1}/A_t$	$-e_{t+2}/A_t$	$-d_{t+1}/A_t$	$-d_{t+2}/A_t$
	(1)	(2)	(3)	(4)
$Skew_t$	−0.0003	0.002	−0.01**	−0.01***
	(−0.09)	(0.90)	(−2.10)	(−3.44)
Controls	Yes	Yes	Yes	Yes
SIC fixed effects	Yes	Yes	Yes	Yes
Cluster adj.	Yes	Yes	Yes	Yes
Adjusted R^2	0.002	0.004	0.03	0.02

Panel A of Table 5.6 presents results for OLS regressions of future equity (Models 1 and 2) and debt issues (Models 3 and 4), as well as the equity share in new issues (Models 5 and 6), on idiosyncratic skewness ($Skew_t$) and several control variables motivated by the literature. In line with Baker and Wurgler (2002) and Alti (2006), we lag control variables by one year as contemporaneous controls may be noisy. The notation and construction of controls follows Section 5.2.3. For brevity, we only report results for $Skew_t$. In Panel B, we replace the dependent variables by equity repurchases (Models 1 and 2) and debt retirements (Models 3 and 4). We define substantial equity (e/A_t) and debt issues (d/A_t) as those exceeding 5% of total book assets. Substantial equity repurchases ($-e/A_t$) and debt retirements ($-d/A_t$) are characterized by *negative* equity and debt issues exceeding 1.25% and 5% of total book assets, respectively (Leary and Roberts, 2005). Stars indicate significance at the 10% (*), 5% (**), and 1% (***) level and t-values (in parentheses) are based on cluster-adjusted standard errors. We cover a sample period from 1971 to 2020 and do not report the intercept.

In order to identify whether the lack of persistence is driven by actual financing decisions, we now investigate the impact of idiosyncratic skewness on issues and repurchases in subsequent years. Table 5.6 presents results. In Panel A, we first focus on the issuance decision. In line with

Alti (2006), firms issue significantly more debt in both $t+1$ (Model 3, $t = 2.45$) and $t+2$ (Model 4, $t = 3.89$). However, in $t+1$, firms also continue to issue more equity (Model 1, $t = 5.34$). While this result may seem counterintuitive, it could simply reflect the fact that planning for an SEO can take time until the next fiscal year (Alti and Sulaeman, 2012). Moreover, as shown by Green and Hwang (2012), idiosyncratic skewness predicts future return skewness. Consequently, firms may anticipate the demand for lottery-like stocks in $t+1$ and issue more equity. As a result, the impact on the equity share remains positive in $t+1$, but turns negative in $t+2$. In both cases, however, statistical significance is rather low. In Panel B, we focus on repurchases. While idiosyncratic skewness is unrelated to future equity repurchases (Models 1 and 2), there is a negative impact on debt retirement in both $t+1$ (Model 3, $t = -2.10$) and $t+2$ (Model 4, $t = -3.44$). Taken together, we find that firms immediately take action in debt markets to rebalance away from the effects of market timing (Flannery and Rangan, 2006). However, firms also continue to issue equity, which likely explains why market timing effects in $t+1$ (Table 5.4) remain similar to those reported in Table 5.2. In unreported results, we repeat this analysis based on logit regressions and find our conclusions to hold.[39]

So far, our results point in the direction of a long-term validity of the trade-off theory, but do not provide a direct test. We therefore follow Flannery and Rangan (2006) and Huang and Ritter (2009) and run a partial adjustment analysis in order to estimate the speed of adjustment

[39] The statistical significance for equity issues increases, but is still exceeded by debt issues in $t+2$. Moreover, the impact on equity repurchases remains statistically negative in $t+1$, but is offset by a more negative impact on debt repurchases.

(SOA) to firm-specific leverage targets. For brevity, we limit our analysis to two approaches. First, we estimate the standard partial adjustment regression

$$Lev_{i,t} = (1-\lambda)Lev_{i,t-1} + \lambda\beta X_{i,t-1} + \epsilon_{i,t}, \tag{5.2}$$

where either book or market leverage in fiscal year t ($Lev_{i,t}$) is regressed on leverage in $t-1$ ($Lev_{i,t-1}$), idiosyncratic skewness, and the standard set of control variables ($X_{i,t-1}$).[40] $\epsilon_{i,t}$ denotes residuals. In the following, we drop the firm index i and define the speed of adjustment as one minus the coefficient estimate of Lev_{t-1}. Second, we follow Flannery and Rangan (2006) and control for firm-fixed effects (α_i)

$$Lev_{i,t} = (1-\lambda)Lev_{i,t-1} + \lambda\alpha_i + \lambda\beta X_{i,t-1} + \epsilon_{i,t}. \tag{5.3}$$

We include these unobserved characteristics to capture potential effects on the firm-specific target leverage that are intertemporally constant but cannot be measured directly. Moreover, according to Flannery and Rangan (2006) and Byoun (2008), firm fixed effects explain a large proportion of the cross-sectional variation in target leverage. We therefore base our conclusions on Equation (5.3). Table 5.7 presents results.

In Model (1), we estimate Equation (5.2) based on the level of book leverage. While the economic impact of $Skew_t$ is similar to Table 5.2, statistical significance strongly increases ($t = -11.10$). The coefficient estimate for Lev_{t-1} is 0.88, implying an annual adjustment of 12% and a half-life of deviations from target of about 5.4 years (log(0.5)/log(0.88)).

[40] These variables have been shown to predict firm-specific leverage ratios. Excluding idiosyncratic skewness does not affect our results.

Table 5.7: Speed of Adjustment to Leverage Targets

	Dependent variable:			
	Lev_t		$Lev_{Mkt,t}$	
	(1)	(2)	(3)	(4)
$Skew_t$	−0.04***	−0.04***	−0.12***	−0.11***
	(−11.10)	(−8.58)	(−13.21)	(−10.58)
$SIZE_{t-1}$	−0.0000	0.004***	−0.003***	−0.001
	(−0.09)	(3.02)	(−3.09)	(−0.28)
TNG_{t-1}	0.02***	0.03***	0.02***	0.04***
	(5.61)	(3.92)	(5.62)	(4.70)
PRF_{t-1}	−0.03***	−0.03***	−0.01	−0.01
	(−4.33)	(−3.81)	(−1.18)	(−1.57)
MB_{t-1}	0.001	−0.0000	0.0001	−0.001
	(1.34)	(−0.02)	(0.16)	(−0.91)
RD_{t-1}	−0.05***	−0.05***	−0.07***	−0.04***
	(−6.97)	(−3.99)	(−6.87)	(−4.74)
RDD_{t-1}	0.004***	0.001	0.01***	−0.001
	(4.83)	(0.70)	(4.46)	(−0.56)
Lev_{t-1}	0.88***	0.68***		
	(188.63)	(62.59)		
$Lev_{Mkt,t-1}$			0.88***	0.67***
			(94.64)	(54.09)
SIC fixed effects	Yes	No	Yes	No
Firm fixed effects	No	Yes	No	Yes
Cluster adj.	Yes	Yes	Yes	Yes
Adjusted R^2	0.79	0.82	0.80	0.82

Table 5.7 presents results for partial adjustment regressions. In Models (1) and (3), we estimate

$$Lev_{i,t} = (1-\lambda)Lev_{i,t-1} + \lambda\beta X_{i,t-1} + \epsilon_{i,t},$$

where either book or market leverage in fiscal year t ($Lev_{i,t}$) is regressed on the leverage ratio in $t-1$ ($Lev_{i,t-1}$), idiosyncratic skewness, and several control variables motivated by the literature ($X_{i,t-1}$). $\epsilon_{i,t}$ denotes residuals. In Models (2) and (4), we account for firm fixed effects (α_i) to capture potential effects on the firm-specific target leverage

$$Lev_{i,t} = (1-\lambda)Lev_{i,t-1} + \lambda\alpha_i + \lambda\beta X_{i,t-1} + \epsilon_{i,t}.$$

Following Flannery and Rangan (2006) and Huang and Ritter (2009), we interpret $\Lambda = 1-\lambda$ as the speed of adjustment toward target leverage. The notation and construction of leverage ratios and controls follows Sections 5.2.2 and 5.2.3, respectively. Stars indicate significance at the 10% (*), 5% (**), and 1% (***) level and t-values (in parentheses) are based on cluster-adjusted standard errors. We cover a sample period from 1971 to 2020 and do not report the intercept.

Since the coefficient estimate for market leverage (Model 3) remains the same, the speed of adjustment is not affected by the choice of leverage. As of yet, our results correspond to those of Fama and French (2002) and DeAngelo and Roll (2015). However, as the standard partial adjustment framework does not capture unobserved effects on the firm-specific target leverage, we include firm fixed effects in Model (2). Most importantly, we find that the estimated speed of adjustment strongly increases. The coefficient estimate of 0.68 implies an annual speed of adjustment of 32% and a half-life of deviations from target leverage of only 1.8 years. According to Warr et al. (2012), it would thus take 1/0.32 = 3.125 years to reach the target. With a coefficient estimate of 0.67, results for the market leverage (Model 4) are very similar. The estimated SOAs are therefore well in line with both Table 5.4 and the previous literature (Flannery and Rangan, 2006; Byoun, 2008; Lemmon et al., 2008; Warr et al., 2012). However, it should be noted that Huang and Ritter (2009) find both estimators to be biased.[41] While the standard partial adjustment approach (Equation 5.2) produces downward biased estimates of the speed of adjustment (upward biased estimates of λ), the SOA for the firm fixed approach (Equation 5.3) is upward biased, especially when the number of observations per firm is low. We justify using the firm fixed approach for several reasons. First, we find that the median firm is listed in Compustat for eight years, which is two years longer than reported by Huang and Ritter (2009). Second, Flannery and Rangan (2006) recognize

[41] Instead, they propose a long-differencing estimator.

this problem and find the SOA to be barely affected.[42] Third, if we account for firm fixed effects, the estimated speed of adjustment closely resembles our previous findings. To capture the effect of market timing on future leverage ratios, it therefore seems reasonable to include firm fixed effects.

Finally, we follow Huang and Ritter (2009) and investigate the interaction effect of market timing (measured by idiosyncratic skewness) and financing conditions (measured by positive financing deficits) on current and future leverage ratios. To some extent, this allows us to discriminate between the three prevalent theories of capital structure. Following Frank and Goyal (2003), the financing deficit (FD) is defined as dividend payments plus investments and the change in working capital minus internal cash flows.[43] We measure the firm-specific financing deficit as a proportion of total assets and define positive financing deficits (FD^+) as FD if $FD > 0$ and zero otherwise.[44] We then repeat our standard regression framework with FD^+ and $Skew_t \times FD^+$ as additional regressors. According to the pecking order theory, financing decisions should not be affected by mispricing. Leverage ratios should therefore only depend on

[42] Flannery and Rangan (2006) focus on market leverage and use the lagged book leverage as an instrumental variable.

[43] Cash dividends correspond to Compustat item 127. Investments equal item 128 + item 113 + item 129 + item 219 - item 107 - item 109 if the format code is 1 to 3 and item 128 + item 113 + item 129 - item 107 - item 109 - item 309 - item 310 if the format code is 7. The change in working capital is defined as item 236 + item 274 + item 301 if the format code is 1 and -item 236 + item 274 + item 301 if the format code is 2 or 3. If the format code is 7, the change in working capital equals - item 302 - item 303 - item 304 - item 305 - item 307 + item 274 - item 312 - item 301. Internal cash flows are defined as item 123 + item 124 + item 125 + item 126 + item 106 + item 213 + item 217 + item 218 if the format code is 1 to 3 and item 128 + item 113 + item 129 - item 107 - item 109 - item 309 - item 310 if the format code is 7. Since many firm-years suffer from at least one missing, we replace missing values by zeros.

[44] This approach follows Huang and Ritter (2009).

the financing deficit, while the interaction term is supposed to be insignificant. In contrast, if the market timing theory holds, the interaction term should display a significantly negative impact that persists for several years. Lastly, according to the static trade-off theory, firms should only issue equity to rebalance toward their target leverage. Thus, after including control variables, the interaction term should be insignificant.[45] In contrast, if the trade-off theory holds in the long-term and market timing is included as a short-term factor, we should see a significantly negative interaction effect in the short run, but no significant impact in the long run. Table 5.8 presents results.

The interaction term ($Skew_t \ x \ FD_t^+$) displays a significantly negative impact on the current book leverage (Model 1), again supporting a short-term market timing effect. To evaluate the economic importance of this effect, we follow Huang and Ritter (2009) and compare the effect for a firm that is severely undervalued and thus displays very high cost of equity ($Skew_t = -1$) and a firm that is severely overvalued and therefore exhibits very low cost of equity ($Skew_t = 1$).[46] Moreover, we also assume a financing deficit of ten percent. In general, the impact on leverage is defined as $(\beta_{Skew_t \times FD_t^+} \times Skew_t + \beta_{FD_t^+}) \times FD_t^+$, where $\beta_{Skew_t \times FD_t^+}$ and $\beta_{FD_t^+}$ denote the coefficient estimates for the interaction term and FD_t^+, respectively. Hence, the increase in leverage for a severely undervalued firm is $(-0.36 \times -1 + 0.38) \times 0.10 = 7.4$ percentage points, whereas the same financial deficit for a severely overvalued firm would result in an increase

[45] The *dynamic* trade-off theory allows for short-term deviations from target leverage.

[46] Huang and Ritter (2009) compare the impact of the highest and lowest equity risk premium, respectively. We assume control variables to remain constant.

Table 5.8: Interaction Effect of Market Timing and Financing Conditions

	Dependent variable:				
	Lev_t	Lev_{t+1}	Lev_{t+2}	Lev_{t+3}	Lev_{t+4}
	(1)	(2)	(3)	(4)	(5)
FD_t^+	0.38***	0.38***	0.37***	0.35***	0.32***
	(15.77)	(16.08)	(15.13)	(17.39)	(14.41)
$Skew_t$	−0.02**	−0.03***	−0.02*	−0.01	−0.005
	(−2.10)	(−2.74)	(−1.66)	(−0.95)	(−0.49)
$Skew_t$ x FD_t^+	−0.36***	−0.25***	−0.16*	−0.18**	−0.10
	(−5.12)	(−3.58)	(−1.87)	(−2.37)	(−1.25)
$SIZE_{t-1}$	0.01***	0.01***	0.01***	0.01***	0.01***
	(7.41)	(7.10)	(6.90)	(6.77)	(6.52)
TNG_{t-1}	0.18***	0.17***	0.16***	0.15***	0.14***
	(17.68)	(14.99)	(13.47)	(12.28)	(11.38)
PRF_{t-1}	−0.14***	−0.12***	−0.11***	−0.10***	−0.11***
	(−5.78)	(−5.03)	(−4.40)	(−4.10)	(−4.20)
MB_{t-1}	−0.03***	−0.02***	−0.02***	−0.02***	−0.02***
	(−16.73)	(−15.03)	(−13.58)	(−12.77)	(−11.38)
RD_{t-1}	−0.42***	−0.40***	−0.38***	−0.35***	−0.35***
	(−10.30)	(−9.36)	(−8.59)	(−7.53)	(−7.62)
RDD_{t-1}	0.02***	0.02***	0.02***	0.02***	0.02***
	(5.21)	(4.95)	(4.71)	(4.23)	(4.05)
SIC fixed effects	Yes	Yes	Yes	Yes	Yes
Cluster adj.	Yes	Yes	Yes	Yes	Yes
Adjusted R^2	0.28	0.26	0.24	0.23	0.22

Table 5.8 presents results for OLS regressions of current and future leverage ratios on idiosyncratic skewness ($Skew_t$), positive financing deficits (FD_t^+), the interaction effect ($Skew_t$ x FD_t^+), and several control variables motivated by the literature. Following Frank and Goyal (2003), the financing deficit is defined as the sum of dividend payments, investments, and the change in working capital minus internal cash flows (scaled by total book assets). For more details, see Footnote 43. In line with Baker and Wurgler (2002) and Alti (2006), we lag control variables by one year as contemporaneous controls may be noisy. The notation and construction of leverage ratios and controls follows Sections 5.2.2 and 5.2.3, respectively. Stars indicate significance at the 10% (*), 5% (**), and 1% (***) level and t-values (in parentheses) are based on cluster-adjusted standard errors. We cover a sample period from 1971 to 2020 and do not report the intercept.

of only 0.2 percentage points. Given that the residual impact of $Skew_t$ is significantly negative ($t = -2.10$), the total effect is even higher. Altogether, the magnitude of the short-term market timing effect is economically important and comparable to Huang and Ritter (2009).[47] In Model (2), we investigate the impact on book leverage in $t+1$. In line with our previous results, the impact of both the interaction effect ($t = -3.58$) and $Skew_t$ ($t = -2.74$) remains highly significant and economically important.[48] However, in subsequent periods (Models 3 to 5), the impact of $Skew_t \times FD_t^+$ (and also $Skew_t$) gradually decreases and finally becomes insignificant in $t+4$. This finding contrasts Huang and Ritter (2009) who find the interaction term to become insignificant in $t+8$. Table 5.8 therefore confirms a strong short-term market timing effect that disappears after four years.

5.3.3 Further Tests

We conclude our main analysis by investigating the impact of idiosyncratic skewness on the change in cash holdings, the change in retained earnings, and investments. The calculation of these measures follows Alti (2006) and is outlined in Section 5.2.2. Table 5.9 presents results.

Idiosyncratic skewness – and thus mispricing – significantly positively predicts the current change in cash holdings (Model 1, $t = 10.79$), suggesting that firms issue equity even if the capital is not needed immediately. This finding is well in line with Alti (2006), Kim and Weisbach (2008), and Dittmar et al. (2020).[49] While being smaller in size, the effect extends

[47] Huang and Ritter (2009) report increments of 6.62 and 0.83 percentage points, respectively.
[48] A financial deficit of ten percent now results in a leverage increase of 6.30 and 1.30 percentage points, respectively.
[49] In contrast, DeAngelo et al. (2010) document that most firms would have run out of cash if they did not issue equity.

Table 5.9: Impact of Market Timing on Liquidity and Investments

	Dependent variable:					
	$\Delta Cash_t/A_t$	$\Delta Cash_{t+1}/A_t$	$\Delta RE_t/A_t$	$\Delta RE_{t+1}/A_t$	Inv_t/A_t	Inv_{t+1}/A_t
	(1)	(2)	(3)	(4)	(5)	(6)
$Skew_t$	0.06***	0.01***	0.06***	0.05***	−0.003	0.02***
	(10.79)	(2.99)	(6.22)	(3.72)	(−1.53)	(5.97)
$SIZE_{t-1}$	0.003***	0.001*	0.004***	0.01***	−0.005***	−0.004***
	(4.37)	(1.87)	(3.27)	(4.81)	(−15.41)	(−14.01)
TNG_{t-1}	0.03***	0.02***	−0.01	0.01	0.15***	0.14***
	(6.53)	(3.61)	(−1.02)	(1.23)	(33.62)	(32.86)
PRF_{t-1}	0.10***	0.07***	0.62***	0.43***	0.08***	0.08***
	(4.51)	(4.85)	(10.97)	(14.90)	(10.47)	(10.36)
MB_{t-1}	0.01***	0.002	0.003	−0.004	0.01***	0.01***
	(7.69)	(1.28)	(1.32)	(−1.22)	(12.01)	(9.29)
RD_{t-1}	0.15***	0.09***	−0.12*	−0.19***	0.02***	0.03***
	(3.82)	(3.21)	(−1.73)	(−3.83)	(2.88)	(3.76)
RDD_{t-1}	0.01***	0.003*	0.005	−0.0001	−0.001	−0.001
	(4.01)	(1.85)	(1.06)	(−0.03)	(−1.05)	(−0.79)
SIC fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Cluster adj.	Yes	Yes	Yes	Yes	Yes	Yes
Adjusted R^2	0.03	0.01	0.19	0.11	0.37	0.36

Table 5.9 presents results for OLS regressions of the change in cash ($\Delta Cash_{t+\tau}/A_t$, Models 1 and 2), the change in retained earnings ($\Delta RE_{t+\tau}/A_t$, Models 3 and 4), and investments ($Inv_{t+\tau}/A_t$, Models 5 and 6) on idiosyncratic skewness ($Skew_t$) and several control variables motivated by the literature. $\Delta Cash_{t+\tau}/A_t$ is defined as the change in cash and short-term investments (item 1) scaled by total book assets (item 6). $\Delta RE_{t+\tau}/A_t$ and $Inv_{t+\tau}/A_t$ correspond to the change in retained earnings (item 36) and capital expenditures (item 128), both scaled by total book assets. τ is either zero or one. In line with Baker and Wurgler (2002) and Alti (2006), we lag control variables by one year as contemporaneous controls may be noisy. The notation and construction of controls follows Section 5.2.3. Stars indicate significance at the 10% (*), 5% (**), and 1% (***) level and t-values (in parentheses) are based on cluster-adjusted standard errors. We cover a sample period from 1971 to 2020 and do not report the intercept.

to $t+1$ (Model 2, $t = 2.99$) and is likely explained by the positive impact of idiosyncratic skewness on subsequent equity issues (as reported in Table 5.6).[50] Besides that, idiosyncratic skewness is positively related to the change in retained earnings (Models 3 and 4). At a first glance, this implies that our results are – at least partially – driven by $Skew_t$'s

[50] Note that $\Delta Cash_{t+1}$ involves the change in cash from t to $t+1$, not from $t-1$ to $t+1$.

impact on retained earnings. However, as with the change in leverage, this conclusion is misleading since the average change in retained earnings is roughly −1%. If we replace $\Delta RE_{t+\tau}/A_t$ by $RE_{t+\tau}/A_t$, there is no impact at all.[51] In Models (5) and (6), we investigate the impact of idiosyncratic skewness on investments. In line with Kim and Weisbach (2008) and DeAngelo et al. (2010), there is a significantly positive impact on future investments (Model 6, $t = 5.97$), while the relationship with contemporaneous investments is flat (Model 5, $t = -1.53$).[52]

Finally, Table 5.A.2 (in the Appendix) investigates the impact of skewness-induced mispricing on current and future stock returns (cumulative from t to $t+\tau$), adjusted for Fama and French (1993) risk factors.[53] In line with the literature on the pricing implications of skewness (e.g. Boyer et al., 2010; Conrad et al., 2013) and the predictions of the market timing theory (Baker and Wurgler, 2002), the impact of $Skew_t$ on current returns (Model 1) is positive, whereas the impact on long-term returns (Models 2 to 4) is significantly negative. This suggests that, on average, market timing is successful.[54]

5.4 Robustness

Our main analysis provides strong evidence for the long-run validity of the trade-off theory, while short-run decisions are determined by market timing. However, our findings could be driven by, for example, the sam-

[51] The corresponding t-statistics are −1.19 and 0.08, respectively.
[52] In unreported results, we find no impact of idiosyncratic skewness on dividends.
[53] To limit the impact of outliers, we trim returns at the 1% and 99% level.
[54] See also Loughran and Ritter (1995) and Lewis and Tan (2016).

ple period, specific firm characteristics, or our measure of idiosyncratic skewness. We therefore repeat our baseline analysis (Tables 5.2 and 5.4) with several adjustments motivated by the literature.

5.4.1 Robustness of Short-Term Results

We first investigate the robustness of the short-term effect and present results in Table 5.10. For brevity, we only report coefficient estimates and t-statistics for $Skew_t$.

First of all, we perform a simple subperiod test by splitting our sample period in half. The first subsample thus covers the period from 1971 to 1995 (Row 1a), whereas the second subsample ranges from 1996 to 2020 (Row 1b). Frank and Goyal (2003) and Huang and Ritter (2009) find that, starting in the 1990s, equity issues track the financing deficit more closely than debt issues, suggesting that the market timing effect might be driven by the second subsample. While this indeed is the case for equity issues (Model 5), we find a similar – if not smaller – effect on both the leverage level (Models 1 and 2) and changes in leverage (Models 3 and 4).[55]

Next, we perform a robustness check based on firm size. Thereby, we follow Alti (2006) and split our sample at annual sales of 50 million dollars (Rows 2a and 2b). While Alti (2006) finds the market timing effect to be comparable across firm size, Dong et al. (2012) report that mispricing is more pronounced for small firms.[56] Since the economic

[55] In unreported results, we investigate the impact of SEC rule 415, which was introduced in October 1983 and permits firms to register a certain amount of securities for SEOs by filing a shelf registration statement. While we indeed we find a stronger effect on equity issues in the post-1983 era, the impact on leverage is less pronounced.

[56] Dong et al. (2012) derive their conclusion from sorts based on total assets.

Table 5.10: Robustness of the Short-Term Impact

	Dependent variable:						
	Lev_t	$Lev_{Mkt,t}$	ΔLev_t	$\Delta Lev_{Mkt,t}$	e/A_t	d/A_t	ES_t
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
(1a) $\text{Subperiod}_{1971-1995}$	−0.04***	−0.10***	−0.04***	−0.14***	0.02***	−0.02***	0.26***
	(−3.54)	(−4.83)	(−8.94)	(−11.43)	(4.06)	(−3.88)	(6.97)
(1b) $\text{Subperiod}_{1996-2020}$	−0.03***	−0.08***	−0.04***	−0.10***	0.05***	0.00	0.16***
	(−3.00)	(−7.64)	(−7.20)	(−10.27)	(6.33)	(−0.55)	(6.64)
(2a) $\text{Size}_{Sales \leq \$50\ mio.}$	−0.06***	−0.13***	−0.05***	−0.13***	0.09***	−0.01	0.18***
	(−4.34)	(−7.27)	(−5.80)	(−13.03)	(6.23)	(−1.28)	(5.90)
(2b) $\text{Size}_{Sales > \$50\ mio.}$	−0.03***	−0.07***	−0.04***	−0.12***	0.02***	−0.01***	0.21***
	(−3.39)	(−5.49)	(−9.90)	(−13.41)	(5.72)	(−3.28)	(7.16)
(3a) Adjustment Costs_{Low}	−0.02**	−0.06***	−0.03***	−0.09***	0.03***	−0.005	0.17***
	(−2.02)	(−6.67)	(−6.58)	(−11.86)	(6.22)	(−1.03)	(6.47)
(3b) Adjustment Costs_{High}	−0.06***	−0.11***	−0.05***	−0.14***	0.05***	−0.02***	0.24***
	(−5.45)	(−6.12)	(−11.26)	(−12.52)	(6.18)	(−3.48)	(6.84)
(4) Fama and MacBeth (1973)	−0.02***	−0.08***	−0.03***	−0.10***	0.03***	−0.01***	0.16***
	(−3.92)	(−11.34)	(−12.13)	(−23.93)	(5.40)	(−2.91)	(10.77)
(5) No Gaps	−0.04***	−0.09***	−0.05***	−0.13***	0.04***	−0.01***	0.21***
	(−3.92)	(−5.92)	(−9.38)	(−12.42)	(7.24)	(−2.62)	(6.92)
(6) Alternative Skewness	−0.003***	−0.006***	−0.003***	-0.009***	0.004***	−0.001**	0.012***
	(−3.60)	(−7.00)	(−7.36)	(−13.84)	(5.18)	(−2.48)	(6.39)
(7) Baker & Wurgler (2002)	0.01	−0.10***	−0.04***	−0.20***	0.04***	−0.04***	0.24***
	(1.17)	(−4.90)	(−8.25)	(−14.19)	(7.81)	(−3.71)	(10.02)
(8) Alternative Leverage	−0.92***		−1.04***				
	(−6.91)		(−10.19)				

Table 5.10 presents robustness checks for the short-term impact of market timing (see Table 5.2). More specifically, we employ subperiods (Rows 1a and 1b), a size split (Rows 2a and 2b), and a split based on adjustment costs. Thereby, firms with low adjustment costs are characterized by having rated debt (Row 3a), whereas firms with high adjustment costs are not rated (Row 3b). Moreover, we re-estimate our regression framework based on the Fama and MacBeth (1973) methodology (Row 4) and without gaps in any of the relevant variables (Row 5). To ensure that our results do not depend on the specific measure of idiosyncratic skewness, we repeat our analysis with the traditional third moment of skewness (Row 6). Finally, we adapt the Baker and Wurgler (2002) methodology (Row 7) and replace our measure of leverage by Debt-to-EBITDA (Row 8), as suggested by Graham (2022). In Row (8), we only report results for Models (1) and (3) since, by construction, Debt-to-EBITDA is a measure of book leverage and the impact on equity and debt issues complies with those in Table 5.2. The notation and construction of dependent variables and controls follows Sections 5.2.2 and 5.2.3, respectively. For brevity, we only report results for $Skew_t$. Stars indicate significance at the 10% (*), 5% (**), and 1% (***) level and t-values (in parentheses) are based on cluster-adjusted standard errors. We cover a sample period from 1971 to 2020.

effect of idiosyncratic skewness on leverage, changes in leverage, and equity issues (Models 1 to 5) is stronger for small firms, our results point in the direction of the latter. However, for large firms, the market timing effect is still important and statistically comparable.[57]

To investigate the impact of adjustment costs, we split our sample based on whether a firm has rated debt or not. According to Byoun (2008), the existence of bond ratings proxies for financial constraints and Elliott et al. (2008) note that rated firms have access to greater quantities of debt at lower costs.[58] Lemmon and Zender (2010) find that the ability to issue public (rated) debt indicates a larger debt capacity. As a result, rated firms (Row 3a) should be less prone to issue equity and the market timing effect should be weaker than for non-rated firms (Row 3b). Economically, we find strong evidence for this hypothesis as the impact on issuance activities and both the leverage level and changes in leverage is strongly increased. Statistically, however, the market timing effect is very similar (except for Models 1 and 6).

Additionally, we re-estimate our regression framework based on the Fama and MacBeth (1973) methodology, since according to Fama and French (2002) standard errors would otherwise be understated.[59] Results are presented in Row (4). While, in economic terms, the market timing

[57] In fact, the negative impact on debt issues is rather driven by large firms.

[58] See also Faulkender and Petersen (2006).

[59] Fama and French (2002) note that year-by-year variation in slopes, determining the standard errors, includes estimation errors due to cross-sectional correlation of residuals. In contrast, Fama and MacBeth (1973) standard errors are heteroscedasticity consistent. Coefficients are based on the time series average of annual cross-sectional regressions and standard errors are calculated based on the annual coefficients. See Fama and MacBeth (1973) and Fama and French (2002) for more details.

effect is somewhat reduced, statistical significance is comparable to our main results. We attribute this finding to the fact that we have already accounted for both firm and year clusters in our OLS regressions.

Frank and Goyal (2003) show that not allowing for gaps in relevant variables has a negative impact on the validity of the pecking order theory, suggesting that the effects of market timing may amplify. In Row (5), we therefore omit all firms with discontinuous data on variables that are needed to calculate leverage ratios and find our conclusions to hold.[60]

To study whether our conclusions depend on the specific measure of idiosyncratic skewness, we repeat our analysis with the traditional third moment of skewness. Results are presented in Row (6). Again, statistical significance is comparable to our main analysis and conclusions remain unchanged.

Next, we recalculate leverage measures and issuance decisions based on Baker and Wurgler (2002) and report results in Row (7). Except for Model (1), conclusions closely resemble our main analysis. Idiosyncratic skewness significantly positively (negatively) affects equity issues (debt issues) and the impact on leverage measures in Models (2) to (4) is significantly negative. In fact, the economic significance even increases.[61] However, in Model (1), the impact of idiosyncratic skewness is not significant (and even positive). We attribute this finding to the large correlation between the Baker and Wurgler (2002) leverage and $SIZE_{t-1}$ (30.7%).

[60] These variables are total assets, short-term debt, long-term debt, total liabilities, preferred stock, deferred taxes, stock prices, and shares outstanding.

[61] Note that the average level of the Baker and Wurgler (2002) leverage is higher (46.5% vs. 24.3%) and therefore leaves more room for impact.

After omitting $SIZE_{t-1}$ (not reported), idiosyncratic skewness turns out to be highly significant, both statistically ($t = -4.40$) and economically.[62] Similarly, we find a significantly negative impact if $SIZE_{t-1}$ is replaced by an adjusted measure of firm size (not reported) that simply scales annual sales by total book assets ($t = -5.04$, with a coefficient estimate of -0.05).

Finally, Graham (2022) documents that – in contrast to the conventional wisdom – CFOs prefer the ratio of debt to EBITDA to evaluate leverage. In Row (8), we therefore report results for this alternative measure.[63] Notably, the impact of idiosyncratic skewness is even stronger. We thus conclude that the short-run effects of market timing, as measured by idiosyncratic skewness, are robust to a wide range of adjustments.

5.4.2 Robustness of Long-Term Results

Table 5.11 provides robustness checks for the persistence analysis. Again, we only report coefficient estimates and t-statistics for $Skew_t$.

In Rows (1a) and (1b), we investigate the persistence of market timing with respect to the two subperiods proposed in Section 5.4.1. Although the economic impact is similar to our main results, persistence in the first subperiod is reduced to only one year. In Rows (2a) and (2b), we repeat the size split. While the short-term impact was stronger for small firms, the persistence of market timing effects is actually lower. This tendency is in line with Alti (2006).[64] With respect to adjustment costs,

[62] The coefficient estimate is -0.05.

[63] We only report results for Models (1) and (3) since, by construction, Debt-to-EBITDA is a measure of book leverage and the impact on equity and debt issues corresponds to Table 5.2.

[64] However, in his study both subsamples become insignificant in $t+1$ (his Table 9).

Table 5.11: Robustness of the Long-Term Impact

	Dependent variable:			
	Lev_{t+1}	Lev_{t+2}	Lev_{t+3}	Lev_{t+4}
	(1)	(2)	(3)	(4)
(1a) Subperiod$_{1971-1995}$	−0.04***	−0.02	−0.01	0.01
	(−3.75)	(−1.60)	(−1.11)	(0.58)
(1b) Subperiod$_{1996-2020}$	−0.03***	−0.02**	−0.02	−0.02
	(−2.94)	(−1.97)	(−1.05)	(−1.33)
(2a) Size$_{Sales \leq \$50\ mio.}$	−0.05***	−0.03	−0.02	−0.01
	(−3.45)	(−1.64)	(−0.98)	(−0.38)
(2b) Size$_{Sales > \$50\ mio.}$	−0.03***	−0.02**	−0.01	0.00
	(−3.46)	(−2.07)	(−1.25)	(−0.37)
(3a) Adjustment Costs$_{Low}$	−0.02**	−0.02	−0.01	-0.01
	(−2.13)	(−1.36)	(−0.60)	(−0.43)
(3b) Adjustment Costs$_{High}$	−0.06***	−0.03***	−0.03**	-0.01
	(−5.48)	(−2.99)	(−2.45)	(−1.33)
(4) Market Leverage	−0.06***	−0.04***	−0.03**	−0.02*
	(−5.05)	(−3.25)	(−2.36)	(−1.84)
(5) Fama and MacBeth (1973)	−0.02***	−0.02***	−0.01	0.00
	(−4.92)	(−2.96)	(−1.57)	(−0.01)
(6) No Gaps	−0.04***	−0.02**	−0.02*	−0.01
	(−4.18)	(−2.45)	(−1.69)	(−1.29)
(7) Alternative Skewness	−0.003***	−0.002***	−0.002**	−0.001
	(−4.43)	(−3.12)	(−2.05)	(−0.75)
(8a) Baker & Wurgler (2002)$_{Book}$	−0.01	0.00	0.00	0.00
	(−0.70)	(0.00)	(0.07)	(−0.15)
(8b) Baker & Wurgler (2002)$_{Mkt}$	−0.08***	−0.05**	−0.04**	−0.04**
	(−4.02)	(−2.43)	(−2.02)	(−2.01)
(9) Alternative Leverage	−0.72***	−0.42***	−0.27**	−0.11
	(−5.79)	(−2.96)	(−2.08)	(−0.76)

Table 5.11 presents robustness checks for the persistence of market timing (see Table 5.4). More specifically, we employ subperiods (Rows 1a and 1b), a size split (Rows 2a and 2b), and a split based on adjustment costs. Thereby, firms with low adjustment costs are characterized by having rated debt (Row 3a), whereas firms with high adjustment costs are not rated (Row 3b). In Row (4), we replace book leverage by market leverage. Moreover, we re-estimate our regression framework based on the Fama and MacBeth (1973) methodology (Row 5) and without gaps in any of the relevant variables (Row 6). To ensure that our results do not depend on the specific measure of idiosyncratic skewness, we repeat our analysis with the traditional third moment of skewness (Row 7). Finally, we adapt the Baker and Wurgler (2002) methodology for both book leverage (Row 8a) and market leverage (Row 8b) and replace our measure of leverage by Debt-to-EBITDA (Row 9), as suggested by Graham (2022). In Row (9), we only report results for Models (1) and (3) since, by construction, Debt-to-EBITDA is a measure of book leverage and the impact on equity and debt issues complies with those in Table 5.2. The notation and construction of leverage ratios and controls follows Sections 5.2.2 and 5.2.3, respectively. For brevity, we only report results for $Skew_t$. Stars indicate significance at the 10% (*), 5% (**), and 1% (***) level and t-values (in parentheses) are based on cluster-adjusted standard errors. We cover a sample period from 1971 to 2020.

the expected effect is straightforward. If a firm is rated (Row 3a), it can issue debt at a lower cost and market timing effects should be rebalanced more quickly than for non-rated firms (Row 3b). Our results provide clear evidence in this direction. For firms with low adjustment costs, the impact of idiosyncratic skewness on leverage in $t+1$ is only significant at the 5%-level and turns insignificant in $t+2$. In contrast, for firms with high adjustment costs, statistical significance extends to $t+3$. In Row (4), we replace book leverage by market leverage. As mentioned in Section 5.2.2, market leverage includes factors that are not under the control of the firm, whereas book leverage reflects active rebalancing. We therefore expect an increased persistence of market timing effects. In line with Welch (2004), we indeed find this to be the case. However, in $t+4$, the coefficient estimate is only significant at the 10%-level and finally turns insignificant in $t+5$ (not reported). Next, we re-estimate the regression framework based on the Fama and MacBeth (1973) methodology. Results (reported in Row 5) correspond to both our main analysis and Table 5.10. While the economic impact is slightly reduced, statistical significance is even increased. However, the impact of idiosyncratic skewness still turns insignificant in $t+3$. When firms are restricted to have continuous time series (Row 6), the persistence of market timing effects extends to $t+3$ (but only at the 10%-level). This result is in line with our expectations in Section 5.4.1. A similar finding is achieved when idiosyncratic skewness is replaced by the traditional third moment of skewness (Row 7). In this case, however, the statistical significance in $t+3$ improves to the 5%-level. In Rows (8a) and (8b), we extend our analysis to the Baker and Wurgler

(2002) methodology. While, not surprisingly, there is no impact on book leverage (Row 8a), results for the market leverage (Row 8b) correspond to those reported in Row (4).[65] Finally, we replace book leverage by Debt-to-EBITDA (Row 9). In line with the amplified short-term impact reported in Table 5.10, the persistence of the market timing effect extends to $t+3$. In conclusion, our evidence on the persistence of market timing is robust to a wide range of robustness checks.[66]

5.5 Concluding Remarks

We provide new evidence on the broad implications of firm-specific mispricing on financing decisions and capital structure. By employing idiosyncratic skewness as a proxy for mispricing, we find a strong short-term effect of market timing. More specifically, idiosyncratic skewness is significantly positively related to equity issues and negatively related to debt issues, with the former effect being the predominant one. Moreover, in line with Hovakimian et al. (2001), we find equity issues to be accompanied by debt retirement programs.

However, in contrast to the predictions of Baker and Wurgler (2002), these effects are not persistent and disappear after about three years. Both partial adjustment models (to estimate the speed of adjustment to target leverage) and interaction effects of idiosyncratic skewness and the firm-specific financing deficit confirm this finding. Finally, our re-

[65] Again, the impact of idiosyncratic skewness turns insignificant in $t+5$ (not reported).

[66] In unreported results, we winsorize idiosyncratic skewness at the 5%-level and 10%-level, respectively, and find our conclusions to hold.

sults are robust to a wide range of adjustments, including sample splits, adjustment costs, Fama and MacBeth (1973) regressions, and both an alternative measure of skewness and alternative definitions of leverage. Our results are thus consistent with a long-run validity of the trade-off theory, including market timing as a short-term factor.

5.A Appendix

Table 5.A.1: Summary Statistics - Change in Leverage

		Skewness Quintile						
	Full Sample	1	2	3	4	5	High − Low	t−value
$Skew_t$	0.04	−0.14	−0.02	0.04	0.11	0.22	0.35	
ΔLev_{t+1}	1.15	2.11	1.30	1.29	0.81	0.24	−1.87***	(−13.12)
ΔLev_{t+2}	1.48	2.23	1.67	1.43	1.27	0.81	−1.41***	(−6.64)
ΔLev_{t+3}	1.62	2.26	1.78	1.64	1.46	0.95	−1.31***	(−6.56)
ΔLev_{t+4}	1.82	2.37	1.95	1.93	1.61	1.28	−1.09***	(−4.81)

Table 5.A.1 presents summary statistics for the change in leverage from fiscal year $t-1$ to $t+\tau$. We sort firms into quintiles based on their idiosyncratic skewness in a given fiscal year and report the equal-weighted average within each quintile. The construction of leverage ratios follows Section 5.2.2. Except for $Skew_t$, values are stated in percent and stars indicate significance at the 10% (*), 5% (**), and 1% (***) level. We cover a sample period from 1971 to 2020.

Table 5.A.2: The Impact of Idiosyncratic Skewness on Stock Returns

	Dependent variable:			
	RET_t	RET_{t+1}	RET_{t+2}	RET_{t+3}
	(1)	(2)	(3)	(4)
$Skew_t$	0.88***	−0.17***	−0.30***	−0.33***
	(14.55)	(−3.89)	(−5.10)	(−3.64)
$MktRf_t$	0.78***	−0.31*	−0.80***	−0.76***
	(9.87)	(−1.68)	(−3.49)	(−3.77)
SMB_t	0.71***	0.28	0.39*	0.49**
	(5.29)	(1.16)	(1.82)	(2.11)
HML_t	0.03	−0.05	−0.41**	−0.30
	(0.30)	(−0.34)	(−2.02)	(−1.12)
SIC fixed effects	Yes	Yes	Yes	Yes
Cluster adj.	Yes	Yes	Yes	Yes
Adjusted R^2	0.13	0.01	0.01	0.001

Table 5.A.2 presents results for OLS regressions of current (Model 1) and future stock returns (cumulative returns from t to $t+\tau$, Models 2 to 4) on idiosyncratic skewness ($Skew_t$) and Fama and French (1993) risk factors. To limit the impact of outliers, we trim returns at the 1% and 99% level. Stars indicate significance at the 10% (*), 5% (**), and 1% (***) level and t-values (in parentheses) are based on cluster-adjusted standard errors. We cover a sample period from 1971 to 2020 and do not report the intercept.

Bibliography

Abdellaoui, Mohammed (2000), 'Parameter-Free Elicitation of Utility and Probability Weighting Functions', *Management Science* **46**(11), 1497–1512.

Aissia, Dorsaf Ben (2014), 'IPO first-day returns: Skewness preference, investor sentiment and uncertainty underlying factors', *Review of Financial Economics* **23**(3), 148–154.

Ait-Sahalia, Yacine and Andrew W. Lo (1998), 'Nonparametric Estimation of State-Price Densities Implicit in Financial Asset Prices', *The Journal of Finance* **53**(2), 499–547.

Ait-Sahalia, Yacine and Andrew W. Lo (2000), 'Nonparametric risk management and implied risk aversion', *Journal of Econometrics* **94**(1), 9–51.

Ait-Sahalia, Yacine, Yubo Wang and Francis Yared (2001), 'Do option markets correctly price the probabilities of movement of the underlying asset?', *Journal of Econometrics* **102**(1), 67–110.

Alti, Aydoğan (2006), 'How Persistent Is the Impact of Market Timing on Capital Structure?', *The Journal of Finance* **61**(4), 1681–1710.

Alti, Aydoğan and Johan Sulaeman (2012), 'When do high stock returns trigger equity issues?', *Journal of Financial Economics* **103**(1), 61–87.

Asem, Ebenezer and Gloria Y. Tian (2010), 'Market Dynamics and Momentum Profits', *The Journal of Financial and Quantitative Analysis* **45**(6), 1549–1562.

Asness, Clifford (2011), 'Momentum in Japan: The Exception that Proves the Rule', *The Journal of Portfolio Management* **37**(4), 67–75.

Asness, Clifford and Andrea Frazzini (2013), 'The Devil in HML's Details', *The Journal of Portfolio Management* **39**(4), 49–68.

Asness, Clifford S., Andrea Frazzini and Lasse Heje Pedersen (2019), 'Quality minus junk', *Review of Accounting Studies* **24**(1), 34–112.

Asness, Clifford S., Tobias J. Moskowitz and Lasse Heje Pedersen (2013), 'Value and Momentum Everywhere', *The Journal of Finance* **68**(3), 929–985.

Autore, Don M. and Jared R. DeLisle (2016), 'Skewness Preference and Seasoned Equity Offers', *The Review of Corporate Finance Studies* **5**(2), 200–238.

Babaoğlu, Kadir, Peter Christoffersen, Steven Heston and Kris Jacobs (2018), 'Option Valuation with Volatility Components, Fat Tails, and Nonmonotonic Pricing Kernels', *The Review of Asset Pricing Studies* **8**(2), 183–231.

Baele, Lieven, Joost Driessen, Sebastian Ebert, Juan M Londono and Oliver G Spalt (2019), 'Cumulative Prospect Theory, Option Returns, and the Variance Premium', *The Review of Financial Studies* **32**(9), 3667–3723.

Baker, Malcolm and Jeffrey Wurgler (2000), 'The Equity Share in New Issues and Aggregate Stock Returns', *The Journal of Finance* **55**(5), 2219–2257.

Baker, Malcolm and Jeffrey Wurgler (2002), 'Market Timing and Capital Structure', *The Journal of Finance* **57**(1), 1–32.

Baker, Malcolm and Jeffrey Wurgler (2013), Chapter 5 - Behavioral Corporate Finance: An Updated Survey, *in* 'Handbook of the Economics of Finance', Vol. 2, Elsevier, pp. 357–424.

Bakshi, Gurdip, Dilip Madan and George Panayotov (2010), 'Returns of claims on the upside and the viability of U-shaped pricing kernels', *Journal of Financial Economics* **97**(1), 130–154.

Bali, Turan G., Nusret Cakici and Robert F. Whitelaw (2011), 'Maxing out: Stocks as lotteries and the cross-section of expected returns', *Journal of Financial Economics* **99**(2), 427–446.

Barber, Brad M., Terrance Odean and Ning Zhu (2009), 'Do Retail Trades Move Markets?', *Review of Financial Studies* **22**(1), 151–186.

Barberis, Nicholas, Abhiroop Mukherjee and Baolian Wang (2016), 'Prospect Theory and Stock Returns: An Empirical Test', *The Review of Financial Studies* **29**(11), 3068–3107.

Barberis, Nicholas and Ming Huang (2008), 'Stocks as Lotteries: The Implications of Probability Weighting for Security Prices', *American Economic Review* **98**(5), 2066–2100.

Barberis, Nicholas, Ming Huang and Tano Santos (2001), 'Prospect Theory and Asset Prices*', *The Quarterly Journal of Economics* **116**(1), 1–53.

Barroso, Pedro (2014), The Bottom-Up Beta of Momentum, SSRN Scholarly Paper ID 2144204, Social Science Research Network, Rochester, NY.

Barroso, Pedro and Pedro Santa-Clara (2015), 'Momentum has its moments', *Journal of Financial Economics* **116**(1), 111–120.

Bates, David S. (2000), 'Post-'87 crash fears in the S&P 500 futures option market', *Journal of Econometrics* **94**(1), 181–238.

Beare, Brendan K. and Lawrence D. W. Schmidt (2016), 'An Empirical Test of Pricing Kernel Monotonicity', *Journal of Applied Econometrics* **31**(2), 338–356.

Beatty, Randolph P. and Jay R. Ritter (1986), 'Investment banking, reputation, and the underpricing of initial public offerings', *Journal of Financial Economics* **15**(1), 213–232.

Benzoni, Luca, Pierre Collin-Dufresne and Robert S. Goldstein (2011), 'Explaining asset pricing puzzles associated with the 1987 market crash', *Journal of Financial Economics* **101**(3), 552–573.

Bleichrodt, Han and Jose Luis Pinto (2000), 'A Parameter-Free Elicitation of the Probability Weighting Function in Medical Decision Analysis', *Management Science* **46**(11), 1485–1496.

Bliss, Robert R. and Nikolaos Panigirtzoglou (2004), 'Option-Implied Risk Aversion Estimates', *The Journal of Finance* **59**(1), 407–446.

Blitz, David, Joop Huij and Martin Martens (2011), 'Residual momentum', *Journal of Empirical Finance* **18**(3), 506–521.

Blitz, David, Matthias X. Hanauer and Milan Vidojevic (2020), 'The idiosyncratic momentum anomaly', *International Review of Economics & Finance* **69**, 932–957.

Bolton, Patrick, Hui Chen and Neng Wang (2013), 'Market timing, investment, and risk management', *Journal of Financial Economics* **109**(1), 40–62.

Boyer, Brian, Todd Mitton and Keith Vorkink (2010), 'Expected Idiosyncratic Skewness', *Review of Financial Studies* **23**(1), 169–202.

Brandt, Michael W. and Kevin Q. Wang (2003), 'Time-varying risk aversion and unexpected inflation', *Journal of Monetary Economics* **50**(7), 1457–1498.

Breeden, Douglas T. and Robert H. Litzenberger (1978), 'Prices of State-Contingent Claims Implicit in Option Prices', *The Journal of Business* **51**(4), 621–651.

Brown, David P. and Jens Carsten Jackwerth (2012), The Pricing Kernel Puzzle: Reconciling Index Option Data and Economic Theory, *in* J.A. Batten and N.Wagner, eds, 'Derivative Securities Pricing and Modelling', Vol. 94 of *Contemporary Studies in Economic and Financial Analysis*, Emerald Group Publishing Limited, pp. 155–183.

Brunnermeier, Markus K., Christian Gollier and Jonathan A. Parker (2007), 'Optimal Beliefs, Asset Prices, and the Preference for Skewed Returns', *American Economic Review* **92**(2), 159–165.

Brunnermeier, Markus K. and Jonathan A. Parker (2005), 'Optimal Expectations', *American Economic Review* **95**(4), 1092–1118.

Byoun, Soku (2008), 'How and When Do Firms Adjust Their Capital Structures toward Targets?', *The Journal of Finance* **63**(6), 3069–3096.

Camerer, Colin F. and Teck-Hua Ho (1994), 'Violations of the betweenness axiom and nonlinearity in probability', *Journal of Risk and Uncertainty* **8**(2), 167–196.

Campbell, John Y. and Luis M. Viceira (2002), *Strategic Asset Allocation: Portfolio Choice for Long-term Investors*, Oxford University Press.

Campbell, John Y. and Samuel B. Thompson (2008), 'Predicting Excess Stock Returns Out of Sample: Can Anything Beat the Historical Average?', *The Review of Financial Studies* **21**(4), 1509–1531.

Campbell, John Y. and John H. Cochrane (1999), 'By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior', *Journal of Political Economy* **107**(2), 205–251.

Cao, Jie, Tarun Chordia and Xintong Zhan (2020), 'The Calendar Effects of the Idiosyncratic Volatility Puzzle: A Tale of Two Days?', *Management Science* **67**(12), 7291–7950.

Carhart, Mark M. (1997), 'On Persistence in Mutual Fund Performance', *The Journal of Finance* **52**(1), 57–82.

Carter, Richard B., Frederick H. Dark and Ajai K. Singh (1998), 'Underwriter Reputation, Initial Returns, and the Long-Run Performance of IPO Stocks', *The Journal of Finance* **53**(1), 285–311.

Carter, Richard B., Frederick H. Dark, Ioannis V. Floros and Travis R. A. Sapp (2011), 'Characterizing the Risk of IPO Long-Run Returns: The Impact of Momentum, Liquidity, Skewness, and Investment', *Financial Management* **40**(4), 1067–1086.

Carter, Richard and Steven Manaster (1990), 'Initial Public Offerings and Underwriter Reputation', *The Journal of Finance* **45**(4), 1045–1067.

Chabi-Yo, Fousseni (2012), 'Pricing Kernels with Stochastic Skewness and Volatility Risk', *Management Science* **58**(3), 624–640.

Chabi-Yo, Fousseni, René Garcia and Eric Renault (2008), 'State Dependence Can Explain the Risk Aversion Puzzle', *The Review of Financial Studies* **21**(2), 973–1011.

Chan, Kalok, Allaudeen Hameed and Wilson Tong (2000), 'Profitability of Momentum Strategies in the International Equity Markets', *The Journal of Financial and Quantitative Analysis* **35**(2), 153.

Chang, Bo Young, Peter Christoffersen and Kris Jacobs (2013), 'Market skewness risk and the cross section of stock returns', *Journal of Financial Economics* **107**(1), 46–68.

Chang, Xin and Sudipto Dasgupta (2009), 'Target Behavior and Financing: How Conclusive Is the Evidence?', *The Journal of Finance* **64**(4), 1767–1796.

Chapman, David A. and Valery Polkovnichenko (2009), 'First-Order Risk Aversion, Heterogeneity, and Asset Market Outcomes', *The Journal of Finance* **64**(4), 1863–1887.

Chen, Yi-Wen, Robin K. Chou and Chu-Bin Lin (2019), 'Investor sentiment, SEO market timing, and stock price performance', *Journal of Empirical Finance* **51**, 28–43.

Christoffersen, Peter, Steven Heston and Kris Jacobs (2013), 'Capturing Option Anomalies with a Variance-Dependent Pricing Kernel', *The Review of Financial Studies* **26**(8), 1963–2006.

Cleveland, William S. (1979), 'Robust Locally Weighted Regression and Smoothing Scatterplots', *Journal of the American Statistical Association* **74**(368), 829–836.

Conrad, Jennifer, Robert F. Dittmar and Eric Ghysels (2013), 'Ex Ante Skewness and Expected Stock Returns', *The Journal of Finance* **68**(1), 85–124.

Cooper, Michael J., Roberto C. Gutierrez and Allaudeen Hameed (2004), 'Market States and Momentum', *The Journal of Finance* **59**(3), 1345–1365.

Cuesdeanu, Horatio and Jens Carsten Jackwerth (2018), 'The pricing kernel puzzle in forward looking data', *Review of Derivatives Research* **21**(3), 253–276.

Daniel, Kent and Tobias J. Moskowitz (2016), 'Momentum crashes', *Journal of Financial Economics* **122**(2), 221–247.

De Long, J. Bradford, Andrei Shleifer, Lawrence H. Summers and Robert J. Waldmann (1990), 'Noise Trader Risk in Financial Markets', *Journal of Political Economy* **98**(4), 703–738.

DeAngelo, Harry, Linda DeAngelo and René M. Stulz (2010), 'Seasoned equity offerings, market timing, and the corporate lifecycle', *Journal of Financial Economics* **95**(3), 275–295.

DeAngelo, Harry and Richard Roll (2015), 'How Stable Are Corporate Capital Structures?', *The Journal of Finance* **70**(1), 373–418.

Dierkes, Maik (2013), Probability weighting and asset prices, SSRN Scholarly Paper ID 2253817, Rochester, NY.

Dierkes, Maik, Jan Krupski and Sebastian Schroen (2022), 'Option-implied lottery demand and IPO returns', *Journal of Economic Dynamics and Control* **138**, 104356.

Dierkes, Maik and Vulnet Sejdiu (2019), 'Indistinguishability of small probabilities, subproportionality, and the common ratio effect', *Journal of Mathematical Psychology* **93**, 102283.

Dittmar, Amy, Ran Duchin and Shuran Zhang (2020), 'The timing and consequences of seasoned equity offerings: A regression discontinuity approach', *Journal of Financial Economics* **138**(1), 254–276.

Dong, Ming, David Hirshleifer and Siew Hong Teoh (2012), 'Overvalued Equity and Financing Decisions', *The Review of Financial Studies* **25**(12), 3645–3683.

Duffie, Darrell, Jun Pan and Kenneth Singleton (2000), 'Transform Analysis and Asset Pricing for Affine Jump-diffusions', *Econometrica* **68**(6), 1343–1376.

Elliott, William B., Johanna Koëter-Kant and Richard S. Warr (2008), 'Market timing and the debt–equity choice', *Journal of Financial Intermediation* **17**(2), 175–197.

Ellis, Katrina, Roni Michaely and Maureen O'Hara (2000), 'The Accuracy of Trade Classification Rules: Evidence from Nasdaq', *The Journal of Financial and Quantitative Analysis* **35**(4), 529.

Eraker, Bjørn and Mark Ready (2015), 'Do investors overpay for stocks with lottery-like payoffs? An examination of the returns of OTC stocks', *Journal of Financial Economics* **115**(3), 486–504.

Erb, Claude B. and Campbell R. Harvey (2006), 'The Strategic and Tactical Value of Commodity Futures', *Financial Analysts Journal* **62**(2), 69–97.

Fama, Eugene F. (1970), 'Efficient Capital Markets: A Review of Theory and Empirical Work', *The Journal of Finance* **25**(2), 383–417.

Fama, Eugene F. (2014), 'Two Pillars of Asset Pricing', *American Economic Review* **104**(6), 1467–1485.

Fama, Eugene F. and James D. MacBeth (1973), 'Risk, Return, and Equilibrium: Empirical Tests', *Journal of Political Economy* **81**(3), 607–636.

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

Fama, Eugene F. and Kenneth R. French (2002), 'Testing Trade-Off and Pecking Order Predictions About Dividends and Debt', *The Review of Financial Studies* **15**(1), 1–33.

Fama, Eugene F. and Kenneth R. French (2015), 'A five-factor asset pricing model', *Journal of Financial Economics* **116**(1), 1–22.

Faulkender, Michael and Mitchell A. Petersen (2006), 'Does the Source of Capital Affect Capital Structure?', *The Review of Financial Studies* **19**(1), 45–79.

Field, Laura Casares and Jonathan M. Karpoff (2002), 'Takeover Defenses of IPO Firms', *The Journal of Finance* **57**(5), 1857–1889.

Flannery, Mark J. and Kasturi P. Rangan (2006), 'Partial adjustment toward target capital structures', *Journal of Financial Economics* **79**(3), 469–506.

Frank, Murray Z and Vidhan K Goyal (2003), 'Testing the pecking order theory of capital structure', *Journal of Financial Economics* **67**(2), 217–248.

French, Kenneth R. (2022*a*), 'Fama/French 3 Factors, Kenneth French Data Library'.
URL: https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html

French, Kenneth R. (2022*b*), 'Industry Returns, Kenneth French Data Library'.
URL: https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/ftp/30_Industry_Portfolios_daily_CSV.zip

Gao, Xiang, Kees G. Koedijk and Zhan Wang (2021), 'Volatility-Dependent Skewness Preference', *The Journal of Portfolio Management* **48**(1), 43–58.

Glosten, Lawrence R., Ravi Jagannathan and David E. Runkle (1993), 'On the Relation between the Expected Value and the Volatility of the Nominal Excess Return on Stocks', *The Journal of Finance* **48**(5), 1779–1801.

Golubev, Yuri, Wolfgang K. Härdle and Roman Timofeev (2014), 'Testing monotonicity of pricing kernels', *AStA Advances in Statistical Analysis* **98**(4), 305–326.

Gonzalez, Richard and George Wu (1999), 'On the Shape of the Probability Weighting Function', *Cognitive Psychology* **38**(1), 129–166.

Graham, John R. (2022), 'Presidential Address: Corporate Finance and Reality', *The Journal of Finance* **77**(4), 1975–2049.

Graham, John R and Campbell R Harvey (2001), 'The theory and practice of corporate finance: evidence from the field', *Journal of Financial Economics* **60**(2), 187–243.

Green, T. Clifton and Byoung-Hyoun Hwang (2012), 'Initial Public Offerings as Lotteries: Skewness Preference and First-Day Returns', *Management Science* **58**(2), 432–444.

Grundy, Bruce D. and J. Spencer Martin Martin (2001), 'Understanding the Nature of the Risks and the Source of the Rewards to Momentum Investing', *The Review of Financial Studies* **14**(1), 29–78.

Guiso, Luigi, Paola Sapienza and Luigi Zingales (2018), 'Time varying risk aversion', *Journal of Financial Economics* **128**(3), 403–421.

Hanauer, Matthias X. and Steffen Windmueller (2021), Enhanced Momentum Strategies, SSRN Scholarly Paper ID 3437919, Social Science Research Network, Rochester, NY.

Hanley, Kathleen Weiss (1993), 'The underpricing of initial public offerings and the partial adjustment phenomenon', *Journal of Financial Economics* **34**(2), 231–250.

Harvey, Campbell R. and Akhtar Siddique (2000), 'Conditional Skewness in Asset Pricing Tests', *The Journal of Finance* **55**(3), 1263–1295.

Hens, Thorsten and Christian Reichlin (2013), 'Three Solutions to the Pricing Kernel Puzzle', *Review of Finance* **17**(3), 1065–1098.

Hovakimian, Armen (2006), 'Are Observed Capital Structures Determined by Equity Market Timing?', *Journal of Financial and Quantitative Analysis* **41**(1), 221–243.

Hovakimian, Armen and Guangzhong Li (2011), 'In search of conclusive evidence: How to test for adjustment to target capital structure', *Journal of Corporate Finance* **17**(1), 33–44.

Hovakimian, Armen, Tim Opler and Sheridan Titman (2001), 'The Debt-Equity Choice', *Journal of Financial and Quantitative Analysis* **36**(1), 1–24.

Huang, Rongbing and Jay R. Ritter (2009), 'Testing Theories of Capital Structure and Estimating the Speed of Adjustment', *Journal of Financial and Quantitative Analysis* **44**(2), 237–271.

Jackwerth, Jens C. (2000), 'Recovering Risk Aversion from Option Prices and Realized Returns', *The Review of Financial Studies* **13**(2), 19.

Jackwerth, Jens Carsten and Mark Rubinstein (1996), 'Recovering Probability Distributions from Option Prices', *The Journal of Finance* **51**(5), 1611–1631.

Jegadeesh, Narasimhan and Sheridan Titman (1993), 'Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency', *The Journal of Finance* **48**(1), 65–91.

Kahneman, Daniel and Amos Tversky (1979), 'Prospect Theory: An Analysis of Decision under Risk', *Econometrica* **47**(2), 263–291.

Kayhan, Ayla and Sheridan Titman (2007), 'Firms' histories and their capital structures', *Journal of Financial Economics* **83**(1), 1–32.

Khan, Mozaffar, Leonid Kogan and George Serafeim (2012), 'Mutual Fund Trading Pressure: Firm-Level Stock Price Impact and Timing of SEOs', *The Journal of Finance* **67**(4), 1371–1395.

Kilka, Michael and Martin Weber (2001), 'What Determines the Shape of the Probability Weighting Function Under Uncertainty?', *Management Science* **47**(12), 1712–1726.

Kim, Woojin and Michael S. Weisbach (2008), 'Motivations for public equity offers: An international perspective', *Journal of Financial Economics* **87**(2), 281–307.

Kliger, Doron and Ori Levy (2009), 'Theories of choice under risk: Insights from financial markets', *Journal of Economic Behavior & Organization* **71**, 330–346.

Korteweg, Arthur, Michael Schwert and Ilya A. Strebulaev (2022), 'Proactive Capital Structure Adjustments: Evidence from Corporate Filings', *Journal of Financial and Quantitative Analysis* **57**(1), 31–66. Publisher: Cambridge University Press.

Kothari, S. P. and Jay Shanken (1992), 'Stock return variation and expected dividends: A time-series and cross-sectional analysis', *Journal of Financial Economics* **31**(2), 177–210.

Kraus, Alan and Robert H. Litzenberger (1976), 'Skewness Preference and the Valuation of Risk Assets', *The Journal of Finance* **31**(4), 1085–1100.

Kumar, Alok (2009), 'Who Gambles in the Stock Market?', *The Journal of Finance* **64**(4), 1889–1933.

Kumar, Alok, Jeremy K. Page and Oliver G. Spalt (2016), 'Gambling and Comovement', *Journal of Financial and Quantitative Analysis* **51**(01), 85–111.

Kumar, Alok, Mehrshad Motahari and Richard Taffler (2022), 'Skewness Preference and Market Anomalies'.

Leary, Mark T. and Michael R. Roberts (2005), 'Do Firms Rebalance Their Capital Structures?', *The Journal of Finance* **60**(6), 2575–2619.

Lee, Charles M. C. and Mark J. Ready (1991), 'Inferring Trade Direction from Intraday Data', *The Journal of Finance* **46**(2), 733–746.

Lemmon, Michael L. and Jaime F. Zender (2010), 'Debt Capacity and Tests of Capital Structure Theories', *Journal of Financial and Quantitative Analysis* **45**(5), 1161–1187.

Lemmon, Michael L., Michael R. Roberts and Jaime F. Zender (2008), 'Back to the Beginning: Persistence and the Cross-Section of Corporate Capital Structure', *The Journal of Finance* **63**(4), 1575–1608.

Lewis, Craig M. and Yongxian Tan (2016), 'Debt-equity choices, R&D investment and market timing', *Journal of Financial Economics* **119**(3), 599–610.

Linn, Matthew, Sophie Shive and Tyler Shumway (2018), 'Pricing Kernel Monotonicity and Conditional Information', *The Review of Financial Studies* **31**(2), 493–531.

Lintner, John (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.

Liu, Jun, Jun Pan and Tan Wang (2005), 'An Equilibrium Model of Rare-Event Premia and Its Implication for Option Smirks', *The Review of Financial Studies* **18**(1), 131–164.

Ljungqvist, Alexander, Vikram Nanda and Rajdeep Singh (2006), 'Hot Markets, Investor Sentiment, and IPO Pricing', *The Journal of Business* **79**(4), 1667–1702.

Ljungqvist, Alexander and William J. Wilhelm (2003), 'IPO Pricing in the Dot-com Bubble', *The Journal of Finance* **58**(2), 723–752.

Ljungqvist, Alexander and William J. Wilhelm (2005), 'Does Prospect Theory Explain IPO Market Behavior?', *The Journal of Finance* **60**(4), 1759–1790.

Loughran, Tim and Bill McDonald (2013), 'IPO first-day returns, offer price revisions, volatility, and form S-1 language', *Journal of Financial Economics* **109**(2), 307–326.

Loughran, Tim and Jay R. Ritter (1995), 'The New Issues Puzzle', *The Journal of Finance* **50**(1), 23–51.

Loughran, Tim and Jay R. Ritter (2002), 'Why Don't Issuers Get Upset About Leaving Money on the Table in IPOs?', *The Review of Financial Studies* **15**(2), 413–444.

Loughran, Tim, Jay R. Ritter and Kristian Rydqvist (1994), 'Initial public offerings: International insights', *Pacific-Basin Finance Journal* **2**(2), 165–199.

Loughran, Tim and Jay Ritter (2004), 'Why Has IPO Underpricing Changed over Time?', *Financial Management* **33**(3), 5–37.

Lowry, Michelle and Susan Shu (2002), 'Litigation risk and IPO underpricing', *Journal of Financial Economics* **65**(3), 309–335.

Mahajan, Arvind and Semih Tartaroglu (2008), 'Equity market timing and capital structure: International evidence', *Journal of Banking & Finance* **32**(5), 754–766.

Markowitz, Harry M. (1952), 'Portfolio selection', *Journal of Finance* **7**, 77–91.

Menkhoff, Lukas, Lucio Sarno, Maik Schmeling and Andreas Schrimpf (2012), 'Currency momentum strategies', *Journal of Financial Economics* **106**(3), 660–684.

Mitton, Todd and Keith Vorkink (2007), 'Equilibrium underdiversification and the preference for skewness', *Review of Financial Studies* **20**(4), 1255–1288.

Modigliani, Franco and Merton H. Miller (1958), 'The Cost of Capital, Corporation Finance and the Theory of Investment', *The American Economic Review* **48**(3), 261–297.

Moskowitz, Tobias J. and Mark Grinblatt (1999), 'Do Industries Explain Momentum?', *The Journal of Finance* **54**(4), 1249–1290.

Mossin, Jan (1966), 'Equilibrium in a Capital Asset Market', *Econometrica* **34**(4), 768–783.

Myers, Stewart C. and Nicholas S. Majluf (1984), 'Corporate financing and investment decisions when firms have information that investors do not have', *Journal of Financial Economics* **13**(2), 187–221.

Ofek, Eli and Matthew Richardson (2003), 'DotCom Mania: The Rise and Fall of Internet Stock Prices', *The Journal of Finance* **58**(3), 1113–1137.

Okunev, John and Derek White (2003), 'Do Momentum-Based Strategies Still Work in Foreign Currency Markets?', *The Journal of Financial and Quantitative Analysis* **38**(2), 425.

Pan, Jun (2002), 'The jump-risk premia implicit in options: evidence from an integrated time-series study', *Journal of Financial Economics* **63**(1), 3–50.

Polkovnichenko, Valery and Feng Zhao (2013), 'Probability weighting functions implied in options prices', *Journal of Financial Economics* **107**(3), 580–609.

Prelec, Drazen (1998), 'The Probability Weighting Function', *Econometrica* **66**(3), 497–527.

Purnanandam, Amiyatosh K. and Bhaskaran Swaminathan (2004), 'Are IPOs Really Underpriced?', *The Review of Financial Studies* **17**(3), 811–848.

Quiggin, John (1993), *Generalized Expected Utility Theory: The Rank Dependent Model*, Springer Science & Business Media.

Rajan, Raghuram G. and Luigi Zingales (1995), 'What Do We Know about Capital Structure? Some Evidence from International Data', *The Journal of Finance* **50**(5), 1421–1460.

Ritter, Jay R. (1984), 'The 'Hot Issue' Market of 1980', *The Journal of Business* **57**(2), 215–240.

Ritter, Jay R. (1991), 'The Long-Run Performance of initial Public Offerings', *The Journal of Finance* **46**(1), 3–27.

Ritter, Jay R. (2022*a*), 'Field-Ritter Dataset of IPO Founding Dates'.
URL: https://site.warrington.ufl.edu/ritter/files/founding-dates.pdf

Ritter, Jay R. (2022*b*), 'IPO Underwriter Reputation Rankings'.
URL: https://site.warrington.ufl.edu/ritter/files/Underwriter-Rank.xls

Ritter, Jay R. and Ivo Welch (2002), 'A Review of IPO Activity, Pricing, and Allocations', *The Journal of Finance* **57**(4), 1795–1828.

Rock, Kevin (1986), 'Why new issues are underpriced', *Journal of Financial Economics* **15**(1), 187–212.

Rosenberg, Joshua V. and Robert F. Engle (2002), 'Empirical pricing kernels', *Journal of Financial Economics* **64**(3), 341–372.

Ross, Steve (2015), 'The Recovery Theorem', *The Journal of Finance* **70**(2), 615–648.

Rouwenhorst, K. Geert (1998), 'International Momentum Strategies', *The Journal of Finance* **53**(1), 267–284.

Rouwenhorst, K. Geert (1999), 'Local Return Factors and Turnover in Emerging Stock Markets', *The Journal of Finance* **54**(4), 1439–1464.

Schneider, Christoph and Oliver Spalt (2017), 'Acquisitions as Lotteries? The Selection of Target-Firm Risk and its Impact on Merger Outcomes', *Critical Finance Review* **6**(1), 77–132.

Sharpe, William F. (1964), 'Capital Asset Prices: A Theory of Market Equilibrium Under Conditions of Risk*', *The Journal of Finance* **19**(3), 425–442.

Shleifer, Andrei and Robert W. Vishny (1997), 'The Limits of Arbitrage', *The Journal of Finance* **52**(1), 35–55.

Shyam-Sunder, Lakshmi and Stewart C. Myers (1999), 'Testing static tradeoff against pecking order models of capital structure', *Journal of Financial Economics* **51**(2), 219–244.

Silverman, B. W. (1986), *Density estimation for statistics and data analysis*, Chapman and Hall, London.

Song, Zhaogang and Dacheng Xiu (2016), 'A tale of two option markets: Pricing kernels and volatility risk', *Journal of Econometrics* **190**(1), 176–196.

Stivers, Chris and Licheng Sun (2010), 'Cross-Sectional Return Dispersion and Time Variation in Value and Momentum Premiums', *Journal of Financial and Quantitative Analysis* **45**(4), 987–1014.

Stott, Henry P. (2006), 'Cumulative prospect theory's functional menagerie', *Journal of Risk and Uncertainty* **32**(2), 101–130.

Thaler, Richard (1985), 'Mental Accounting and Consumer Choice', *Marketing Science* **4**(3), 199–214.

Tinic, Seha M. (1988), 'Anatomy of Initial Public Offerings of Common Stock', *The Journal of Finance* **43**(4), 789–822.

Tversky, Amos and Craig R. Fox (1995), 'Weighing risk and uncertainty', *Psychological Review* **102**(2), 269–283.

Tversky, Amos and Daniel Kahneman (1992), 'Advances in prospect theory: Cumulative representation of uncertainty', *Journal of Risk and Uncertainty* **5**(4), 297–323.

Warr, Richard S., William B. Elliott, Johanna Koëter-Kant and Özde Öztekin (2012), 'Equity Mispricing and Leverage Adjustment Costs', *Journal of Financial and Quantitative Analysis* **47**(3), 589–616.

Warusawitharana, Missaka and Toni M. Whited (2016), 'Equity Market Misvaluation, Financing, and Investment', *The Review of Financial Studies* **29**(3), 603–654.

Welch, Ivo (2004), 'Capital Structure and Stock Returns', *Journal of Political Economy* **112**(1), 106–132.

White, Halbert (1980), 'A Heteroskedasticity-Consistent Covariance Matrix Estimator and a Direct Test for Heteroskedasticity', *Econometrica* **48**(4), 817–838.

Wu, George and Richard Gonzalez (1996), 'Curvature of the Probability Weighting Function', *Management Science* **42**(12), 1676–1690.

Zeisberger, Stefan, Dennis Vrecko and Thomas Langer (2012), 'Measuring the time stability of Prospect Theory preferences', *Theory and Decision* **72**(3), 359–386.

Ziegler, Alexandre (2007), 'Why Does Implied Risk Aversion Smile?', *The Review of Financial Studies* **20**(3), 859–904.