

Advanced REIT Portfolio Optimization

Innovative Tools for Risk Management

W. Brent Lindquist · Svetlozar T. Rachev · Yuan Hu · Abootaleb Shirvani
with contributions from
Stephen T. Crosson · Jimmy H. Jackson

Advanced Analytics for the Real Estate Market

W. Brent Lindquist
Department of Mathematics and Statistics
Texas Tech University
Lubbock, Texas, USA

Yuan Hu
Department of Mathematics and Statistics
Texas Tech University
Lubbock, Texas, USA

Stephen T. Crosson
Jantz Analytics
Dallas, Texas, USA

Svetlozar T. Rachev
Department of Mathematics and Statistics
Texas Tech University
Lubbock, Texas, USA

Abootaleb Shirvani
Department of Applied and Computational
Mathematics
Notre Dame University
Notre Dame, Indiana, USA

Jimmy H. Jackson
Jantz Analytics
Plano, Texas, USA

Foreword

This book offers a scholarly presentation of the premises and applications of the mathematics of the proprietary software of Jantz Analytics, LLP. The Jantz models focus on publicly traded stocks of real estate investment trusts (REITs). Specifically, the models provide highly advanced analytics for REIT portfolio optimization. The Jantz models provide the following toolset:

- Portfolio optimization strategies incorporating tail-risk assessment, day-ahead return forecasting, turnover and performance attribute constraints, benchmark tracking and investor view input;
- Backtesting;
- A spectrum of risk-assessment tools augmenting traditional risk measures with early warning systems and risk budgeting;
- Option valuation; and
- Inclusion of environmental, social and governance ratings in portfolio optimization and option pricing.

These tools are employed within a unified framework consistent with dynamic asset pricing (rational finance). The Jantz software is unique among existing portfolio-optimization platforms. Many such platforms are based on historical performance, but the Jantz software is predictive. All Jantz forecasting and risk models are consistent with the regulatory requirements in Basel II and Basel III. *In short, we believe the Jantz software reflects the “state of the science” in portfolio optimization, risk analysis and option valuation.*

The expertise of the Jantz team is multidimensional. It represents a unique combination of business knowledge, real estate experience, and world-class mathematical and statistical credentials. The cofounders of Jantz have over 80 years of combined valuation and consulting experience in real estate throughout the United States and Puerto Rico. This experience includes virtually all property types. Their clients have included commercial banks, pension funds, various public-sector entities, and investment banks.

Dr. Svetlozar (Zari) Rachev is one of the world’s foremost authorities on the application of heavy-tailed distributions in finance. He is a cofounder and formerly the president of Bravo Risk Management Group, which originated the Cognition methodology. Bravo was acquired by FinAnalytica, where Zari served as chief scientist. Dr. W. Brent Lindquist is a computational mathematician with 40 years of experience in developing numerical methods. He is a cofounder of the company that marketed the Frontier package used in oil reservoir simulation and has commercially licensed his 3DMA-Rock code for studying flow at pore scales. Dr. Abootaleb Shirvani is an expert in Lévy subordinated processes applied to finance. Dr. Yuan Hu’s expertise is in option pricing in complete markets with non-Gaussian returns.

We anticipate that the readership of this book will be not only mathematical and statistical experts but also management-level professionals without deep knowledge of the mathematics presented. To assist in the latter’s comprehension, a summary of content in nonmathematical terms is presented as an abstract preceding each chapter.

Stephen T. Crosson, MAI, SRA

Jimmy H. Jackson, MAI

Cofounders, Jantz Analytics, LLC

Chapter 1	The Real Estate Investment Market: The Current State and Why Advances Are Needed	1
Chapter 2	The Data	11
2.1	REIT Asset Descriptions	11
2.1.1	Domestic REITs	11
2.1.2	International REITs	16
2.2	Real Estate Stock Descriptions	18
2.3	Benchmarks	20
2.3.1	<i>Indices</i>	20
2.3.2	<i>Exchange Traded Funds</i>	21
2.3.3	<i>Mutual Funds</i>	21
2.4	Additional Assets and Indices	22
2.5	Data Observations	24
	References	
Chapter 3	Modern Portfolio Theory	25
3.1	Return Time Series	25
3.2	MPT-Based Portfolios	26
3.2.1	<i>Markowitz Mean-Variance Portfolio</i>	27
3.2.2	<i>Markowitz Mean-Variance Tangent Portfolio</i>	30
3.2.3	<i>CVaR-Minimizing Portfolios</i>	31
3.2.4	<i>CVaR Tangent Portfolio</i>	34
3.2.5	<i>Criticisms of Mean-Variance Optimization</i>	35
3.3	Black–Litterman Model	35
3.4	Historical Optimization	38
	References	40
Chapter 4	Historical Portfolio Optimization – Domestic REITs	41
4.1	Basic Strategies, Price, and Return Performance	42
4.1.1	<i>Long-Only Strategy</i>	42
4.1.2	<i>Jacobs et al. Long–Short Strategy</i>	43
4.1.3	<i>Lo–Patel Long–Short Strategy</i>	45
4.1.4	<i>Long–Short Momentum Strategy</i>	46
4.2	Performance under Turnover Constraints	47
4.3	Performance–Risk Measures	54
4.4	Observations	61
	References	62
Chapter 5	Diversification with International REITs	63
5.1	International Portfolio Performance	63
5.1.1	<i>Long-Only International Portfolios</i>	63
5.1.2	<i>Jacobs et al. Long–Short International Portfolios</i>	67
5.1.3	<i>Lo–Patel Long–Short International Portfolios</i>	69

5.2	Global Portfolio Performance	69
5.2.1	<i>Long-Only Global Portfolios</i>	70
5.2.2	<i>Jacobs et al. Long–Short Global Portfolios</i>	73
	References	75
Chapter 6	Black–Litterman Optimization Results	76
6.1	<i>Domestic Portfolios</i>	76
6.2	<i>Global Portfolios</i>	79
Chapter 7	Dynamic Portfolio Optimization: Beyond MPT	82
7.1	Dynamic Optimization	82
7.1.1	<i>ARMA(1,1)–GARCH(1,1) with Student’s t-Distribution</i>	83
7.1.2	<i>Multivariate t-Distribution and t-Copulas</i>	84
7.1.3	<i>Generation of Dynamic Returns</i>	84
7.1.4	<i>Combining the Dynamic Approach with the Black–Litterman Model</i>	86
7.2	Portfolio Optimization Using Dynamic Returns	87
7.2.1	<i>Dynamic Long-Only Portfolios</i>	91
7.2.2	<i>Dynamic Jacobs et al. Long–Short Portfolios</i>	93
7.2.3	<i>Dynamic Lo–Patel Long–Short Portfolios</i>	95
7.3	Dynamic Optimization with the Black–Litterman Model	98
	References	99
Chapter 8	Backtesting	101
8.1	VaR Tests	103
8.1.1	<i>Binomial Test</i>	103
8.1.2	<i>Traffic Light Test</i>	102
8.1.3	<i>Kupiec’s Tests</i>	104
8.1.4	<i>Christoffersen’s Tests</i>	105
8.1.5	<i>Haas’s Tests</i>	107
8.3	Backtest Results	108
8.3.1	<i>Historical Optimization</i>	108
8.3.2	<i>Dynamic Optimization</i>	113
	References	115
Chapter 9	Diversification with Real Estate Stocks	116
Chapter 10	Risk Information and Management	122
10.1	Early Warning Systems	122
10.1.1	<i>Chow Test for a Structural Break</i>	123
10.1.2	<i>Early Warning Based on Tail-Loss Ratio</i>	124
10.1.3	<i>Early Warning Based on Mahalanobis Distance</i>	129
10.2	Asset Weighting	133
10.3	Risk Budgets: Incremental and Component Risk	138
10.3.1	<i>Incremental, Marginal, and Component VaR</i>	139
10.3.2	<i>Computing VaR, IVaR, MVaR, and ciVaR</i>	140

	10.3.3 Portfolio Results	143
10.4	Factor Analysis	151
	References	159
Chapter 11	Optimization with Performance-Attribution Constraints	162
11.1	Performance-Attribute Constraints	163
11.2	Application to the Domestic REIT Portfolio	166
	References	175
Chapter 12	Option Pricing	178
12.1	Double Subordinated Pricing Models	182
12.2	Option Pricing under the Double Subordinated IG Model	184
12.3	Empirical Example	186
	12.3.1 Choice of a and v_{\max}	189
	12.3.2 Option Price and Implied Volatility Surfaces	194
12.4	Volatility Measures	195
	Appendix A	199
	Appendix B	200
	References	200
Chapter 13	Inclusion of ESG Ratings in Optimization	203
Chapter 14	Inclusion of ESG Ratings in Option Pricing	205

Abbreviations

This list of abbreviations does not include asset and benchmark abbreviations. All such abbreviations are defined in Chapter 2 and are assigned by the trading exchange.

AA	Asset allocation effect
ADR	American depository receipts
AIC	Akaike information criterion
ARMA	autoregressive moving-average
BIC	Bayesian information criterion
BL	Black–Litterman model
CAPM	capital asset pricing model
CC	Christoffersen’s conditional coverage test
CCI	Christoffersen’s conditional coverage independence test (Chapter 8)
CDF	cumulative distribution function
CF	characteristic function
<i>ciStd</i>	component standard deviation
<i>ciVaR</i>	component value-at-risk
CML	capital market line
CVaR	conditional value-at-risk
D	dynamic optimization
EQW	equal-weighted portfolio
ESG	environment, sustainability and governance
ETF	exchange-traded fund
FFT	fast Fourier transform
GARCH	generalized autoregressive conditional heteroscedasticity
H	historical optimization
I	interaction term (Chapter 11)
IG	inverse Gaussian distribution
IQR	inter-quartile range
IVaR	incremental value at risk
M95	minimum CVaR ₉₅ portfolio
M99	minimum CVaR ₉₉ portfolio
MDD	maximum drawdown
MPT	modern portfolio theory
MStd	marginal standard deviation
MTLR	modified tail-loss ratio
MTLS	modified tail-loss spread
MVaR	marginal value-at-risk
MVP	Markowitz minimum mean-variance portfolio
NDIG	normal double inverse Gaussian
PDF	probability density function
PoF	Kupiec’s proportion of failures test

R	REITs-only portfolio
RR	Rachev ratio
R+S	REITs + stock portfolio
REIT	real estate investment trust
S&P	Standard and Poors
SE	selection effect
SR	Sharpe ratio
SS	Sortino–Satchell ratio
Std	standard deviation
T95	tangent CVaR ₉₅ portfolio
T99	tangent CVaR ₉₉ portfolio
TBF	Haas's time between failures test
TBFI	Haas's time between failures independence test
TL	traffic light test (Basel accord)
TLR	tail-loss ratio
TLS	tail-loss spread
TO	turnover
TVP	Markowitz tangent mean-variance portfolio
TUFF	Kupiec's time until first failure test
VaR	value-at-risk

Chapter 1

The Real Estate Investment Market: The Current State and Why Advances Are Needed

To prepare for the most challenging investment periods, an investment approach must comprise two critical components: (1) diversified portfolios managed under strategies that maximize return while minimizing risk (based on one or more chosen measures), supplemented with additional risk-assessment and risk-management tools, and (2) the purchase of derivatives (e.g., options) that provide insurance against “foreseeably expected” adverse conditions. *The central goal of this book is to elucidate such an investment approach, developed in the context of model portfolios of REIT-based assets.*

Risk is the greatest factor in investing. Modern investment theories (such as modern portfolio theory) and strategies provide techniques for imposing some degree of “control” over the risk–return spectrum. However, it is the appearance of outlier events (tail events, in the language of probability) that occasion the greatest investment stress and can even lead to deep changes in financial and social systems. For example, the trigger events of the Great Recession of 2008–2009 – the bursting of the U.S. housing bubble in 2005–2006 and the resulting subprime mortgage crisis in 2007–2008 – revealed significant vulnerabilities in the financial system and led to the collapse or bailout of major investment banks. The recession produced a serious disruption of normal international relations. Unlike the Great Recession, and perhaps more akin to the disruptive force of internet technology, the COVID-19 pandemic has severely tested the assumptions on which many real estate models are based. In many ways, the experiences of the real estate market over the course of the pandemic, ranging from shock and denial to reconstruction, resemble stages of human grieving. There are strong indications that the pandemic has accelerated, and perhaps made irreversible, developments that were already underway in online shopping, online streaming of entertainment content, online ordering of food, remote work, and online business meetings – all of which have repercussions for existing business and residential real estate models. Consider the following news articles, released during January, February, and May 2021 as the world began to emerge from the pandemic. We offer a brief commentary on each article.

2020 Was the Manhattan Office Market's Worst Year This Century (Bisnow, 1/7/2021¹) “The last quarter of 2020 capped a disruptive year for the world's largest office market as asking rents fell, availability increased and leasing activity dropped. Manhattan saw its slowest year for leasing since the start of the 21st century.”

This article deals with the contraction in demand for office space in Manhattan – a development that has affected office-space demand in major markets throughout the United States. Further, the pandemic’s abatement is unlikely to result in pre-COVID demand. Both employers and employees are realizing the benefits of working from home, at least part-time. The magnitude and potential duration of this phenomenon are currently unknowable. Investors will compensate by requiring greater returns and heightened scrutiny of tenants’ creditworthiness and in-office needs. (S.T.C. 6/8/2021)

¹ https://www.bisnow.com/new-york/news/office/overall-manhattan-office-leasing-volume-hit-its-lowest-point-this-century-in-2020-107283?utm_source=push.

City Apartment REITs Await the Return to the Office (*The Wall Street Journal*, 1/5/2021²).

“Covid-19 hurt Equity Residential and UDR shares last year, with rents falling as some urban employees moved to rural and suburban areas.”

The risk inherent in urban-apartment investments has not been adequately analyzed. Several factors are clearly in need of consideration, including employees’ preference for working from home, employers’ desire to reduce office-rent expenses by permitting working from home, and the desire of some employees to relocate to areas with lower tax rates, lower housing costs, less regulation, and less civil disorder, among other considerations. (S.T.C. 2/28/2021)

The Metrics You Should Be Watching in 2021 (GlobeSt.com, 1/12/2021³). “Job cuts and the Texas ratio will matter.”

This article points out some important metrics for property demand. However, sophisticated analytics involve a far deeper dive into value-influencing factors for real estate. (S.T.C. 2/28/2021)

Covid-19 Forces Co-Working Firms to Recast Their Business Model (*The Wall Street Journal*, 2/9/2021⁴). “Pandemic is accelerating the end of the industry’s lease-and-sublet model, some suggest.”

This article reflects the reality of risk in the office sector as well as firms’ creativity is adapting to diminished demand. The new coworking paradigm requires ownership to share in master lessees’ risk of demand contraction. (S.T.C. 2/28/2021)

Zillow Stock Soars to New Highs as Residential Real Estate Moves Online (*Barron’s*, 2/11/2021⁵). “Zillow Group shares have touched a new all-time high after the online real estate services company posted better-than-expected fourth-quarter results, as the pandemic accelerates consumer adoption of virtual home sales.”

The Zillow model, made possible by the use of internet technology, illustrates the consolidation of elements of the single-family residential segment. Similar models will surely follow soon. (S.T.C. 2/28/2021)

The Close: Stocks Rise as Yellen Backs More Stimulus, BlackBerry Continues to Surge (*The Globe and Mail*, 1/19/2021⁶). “Wall Street’s main indexes rose on Tuesday as U.S. Treasury Secretary nominee Janet Yellen advocated for a hefty fiscal relief package before lawmakers to help the world’s largest economy ride out a pandemic-driven slump.”

² https://www.wsj.com/articles/city-apartment-reits-await-the-return-to-the-office-11609851601?st=6ckq1zhhl479mn7&reflink=article_gmail_share.

³ <https://www.globest.com/2021/01/12/the-metrics-you-should-be-watching-in-2021/?slreturn=20210128174522>.

⁴ <https://www.wsj.com/articles/covid-19-forces-co-working-firms-to-recast-their-business-model-11612875600>.

⁵ <https://www.barrons.com/articles/zillow-stock-soars-to-new-highs-as-residential-real-estate-moves-online-51613065333?st=17x6d4eako3tw0t>.

⁶ <https://www.theglobeandmail.com/investing/markets/inside-the-market/market-news/article-premarket-world-shares-climb-ahead-of-yellen-speech-earnings-in-focus/>.

1. Recent Developments

The additional federal stimulus will likely result in countervailing impacts on commercial real estate. There is a widespread belief that higher inflation is probable in the near-intermediate term. In general, inflation has a positive impact on commercial real estate pricing. However, the pandemic has exerted a profound influence on demand for the major commercial property types, excluding the warehouse sector. The timing of a return to prepandemic demand remains unclear. Indeed, some industry observers believe the work-from-home phenomenon has reduced office demand for the foreseeable future. Lastly, most commercial lenders have exercised forbearance in dealing with nonperforming loans. The resulting uncertainty increases the risk inherent in this asset class. (S.T.C. 1/21/2021)

Becoming a Landlord Looks Good on Paper. In Reality, It's Much Harder. What to Consider before Turning Your Old Home into a Rental Property (*The Wall Street Journal*, 1/21/2021⁷).

The illustration this article presents, in which a homeowner in California could receive rental payments of \$6,000 per month on their \$1.6-million residence, reflects a nominal 4.50% return on investment (ROI). This return would be reduced by other expenses, such as taxes, insurance, and overall major maintenance. The only certainty or guaranteed allurement of this kind of investment would be the property depreciation that could be charged on the rental property. However, the recapture of the depreciation at the end of the investment creates a bit of a tax surprise or burden. For the first time in many decades, the long-standing continued home-appreciation expectation in California has become uncertain. If appreciation rates go flat, as some experts have predicted, the California homeowner's ROI after all expenses might be 1%–2%, or even lower. This nominal rate of return does not appear to be worth the risk relative to the downside of owning the \$1.6-million residence. Another real source of potential risk is rent control, which limits the magnitude of annual increases in rent. Due to these cumulative risk factors, it appears to be advisable to convert this large amount of equity to a higher-yielding class of real estate, such as a proper lower-risk allocation of REIT shares in which the investment pays a healthy dividend, has appreciation possibilities spread across many assets, and is much more passive for the investor. (J.H.J. and S.T.C. 1/22/2021)

How Top Real Estate Fund Managers Are Preparing for a Post-Covid World (*Barron's*, 2/22/2021⁸). “‘We think three to five years from now 20% of the workforce is gone on a daily basis from the existing standing inventory of office space,’ says Burl East, manager of Altegris/AACA Opportunistic Real Estate.”

The pandemic has greatly increased risk in the office, retail, and hospitality sectors of commercial real estate. The benefactors of shifting investor demand are warehouse/distribution centers as well as properties that house technology-based businesses. The emergence of “prop tech” firms (tech firms focused on property) is also a recent example of investors' expanding appetites for both real estate and superior analytics. Prop tech firms are proving to be highly-sought-after targets for special purpose acquisition companies. (S.T.C. 2/28/2021)

⁷ <https://www.wsj.com/articles/should-i-become-a-landlord-11611237640>.

⁸ <https://www.barrons.com/articles/how-top-real-estate-fund-managers-are-preparing-for-a-post-covid-world-51613993400>.

Hotel Sellers Use Car-Sales Model to Finance Lodging Deals (*The Wall Street Journal*, 2/23/2021⁹). “Banks, owners offer cheap rates and sometimes require little money down to unload unprofitable hotels.”

This article compares current seller financing to that of the savings and loan crisis of the late 1980s and early 90s. However, they have significant differences. The current situation has resulted from a sharp contraction in demand caused by the pandemic, but the savings and loan crisis was caused by a vast oversupply of commercial real estate. Moreover, the former is likely to be resolved in relatively short order, whereas it took years to restore equilibrium after the latter. (S.T.C. 2/28/2021)

The DOJ Says It Will Appeal after a Trump-Appointed Judge Struck Down a Federal Eviction Moratorium (Business Insider, 2/28/2021¹⁰).

The nationwide eviction moratorium from the Centers for Disease Control and Prevention has been ruled unconstitutional and is on appeal to the U.S. Fifth Circuit of Appeals. If this ruling is sustained on appeal, it will almost certainly lead to litigation against the federal government for lost rents caused by the moratorium. (S.T.C. 2/28/2021)

Homebuyers Are Getting Slammed by Record-High Prices. Here's When Economists Say They'll Finally Ease Up (Business Insider, 5/13/2021¹¹). “Limited inventory and low mortgage rates have sparked rising homebuyer demand, which has resulted in astronomical home-price growth over the past year.”

This article does not address the critical factor of affordability. Without concomitant growth in potential buyers’ income, home-price increases must abate. Homebuilders will be squeezed by continuing increases in material costs, on the one hand, and diminishing buyer affordability, on the other. (S.T.C. 5/19/2021)

During the COVID-19 pandemic, lingering forbearance risks have contributed to more stringent underwriting criteria. Even after the reopening of the economy, underwriting standards have not softened for the mortgage-applicant sector. These stringent lending standards will eventually slow or stall sale transactions for this large sector of buyers who utilize mortgage loans. (J.H.J. 5/19/2021)

Investors Bet on Commercial Real Estate, Undeterred by Empty Offices and Hotel Rooms (*The Wall Street Journal*, 5/18/2021¹²). “Government support and help from banks kept landlords from suffering steep losses.”

Over the next several years, the commercial real estate market will most definitely see the repositioning of now-suffering property classes (hotels, malls, and office properties) as other property classes due to the lack of market/consumer demand. For instance, the following property-classification shifts are likely to occur simultaneously in the short term:

⁹ <https://www.wsj.com/articles/hotel-sellers-use-car-sales-model-to-finance-lodging-deals-11614076202>.

¹⁰ <https://www.businessinsider.com/doj-attorneys-to-appeal-cdc-eviction-moratorium-ruling-2021-2>.

¹¹ <https://www.businessinsider.com/when-will-home-prices-drop-increase-predictions-forecast-2021-5>.

¹² https://www.wsj.com/articles/investors-bet-on-commercial-real-estate-undeterred-by-empty-offices-and-hotel-rooms-11621330204?st=9qguj7qtf2f8qv4&reflink=article_gmail_share.

1. Recent Developments

- hotels to condo/multifamily residential uses
- office properties to multifamily residential properties
- large-scale regional malls to Google-type last-mile consumer fulfillment centers, etc.

It appears that this musical-chairs-style repositioning/redevelopment is already priced into the sentiment of the commercial real estate market as well as into the institutional lender community. (J.H.J. 5/25/2021)

Blackstone REIT Plans to Name Katie Keenan as CEO (*The Wall Street Journal*, 5/19/2021¹³). “Blackstone Mortgage Trust Inc., a large commercial real-estate lender, is expected to appoint Katie Keenan as chief executive Wednesday, partly crediting the role she played in helping navigate the company through one of the most challenging markets in decades.”

The declines in REIT shares early in the pandemic were predictable and rational, given the many unknowns in play. Chief among them was the extent and timing of postvirus demand, particularly for office property. Blackstone’s recovery clearly resulted in part from the durability of demand from many of its biotech and entertainment-based tenants. This reflects the obvious premise that commercial real estate is highly differentiated; for example, not all office buildings have the same risk characteristics. (S.T.C. 5/23/2021)

House Prices Are Soaring. Rents Are Flat. What Does It Mean for Inflation? (*Barron’s*, 5/20/2021¹⁴). “House prices are rising at their fastest rate on record—and the pace of increases is accelerating. Monthly data from Zillow and from the S&P CoreLogic Case-Shiller Index both indicate that the current surge in U.S. home prices is unlike anything seen before, including the bubble of the 2000s.”

I vividly remember the housing bubble created by the subprime mortgage crisis of 2008. For several years afterwards, homebuilders could not move or sell their inventory of newly constructed homes in the Dallas–Fort Worth metroplex, one of the most active homebuilding markets in the entire United States. Many homebuilders resorted to the lease market to attempt to ride out the housing downturn and, basically, survive the corrective disruption. The concept of large-scale single-family rental communities emerged from this corrective disruption. This real estate concept resulted not from rental demand but from public and non-publicly-traded homebuilders’ need to ride out the storm. The central factor of sustaining single-family rental demand in 2021 is unemployment. A tenant simply cannot pay rent without a job. When economic stimulus payments end at some point in the future, and if unemployment is still excessively high, the reality of this situation will set in, and the single-family rental bubble will likely burst. (J.H.J. 1/8/2021)

More variables are impacting home prices than ever before and traditional metrics for the relative state of pricing are inadequate. For instance, the large institutional appetite for buying single-family homes in order to rent them is a recent phenomenon. This appears to be the equivalent to momentum investing in equities. (S.T.C. 5/23/2021)

This article provides relevant perspectives on the interplay between home prices and rental rates. However, the analysis of returns is simplistic. It is highly likely that repairs, maintenance,

¹³ https://www.wsj.com/articles/blackstone-reit-plans-to-appoint-katie-keen-an-as-ceo-11621424701?st=iaz6zfqdhi07wny&reflink=article_gmail_share.

¹⁴ <https://www.barrons.com/articles/house-prices-are-soaring-rents-are-flat-what-does-it-mean-for-inflation-51621525998?st=8ih1bt5c6uwis7h>.

and replacements affecting rented housing are not deducted from nominal returns. If so, the schism between rising home prices and achievable rents is even wider than has been assumed. (S.T.C. 5/25/2021)

U.S. Home Prices Push to Record High, Slowing Pace of Purchases (*The Wall Street Journal*, 5/21/2021¹⁵). “Cost for existing properties up 19% in April from year earlier; supply of homes remains limited.”

Declines in the pace of home sales are directly correlated with limited supply. Constraints on the supply of new homes include delays in the availability of materials and subcontractors. Furthermore, the rapid increase in home prices in many markets has made them unaffordable for many buyers. In short, the home market is in disequilibrium. (S.T.C. 5/23/2021)

Bet on a Rental Resurgence as Housing Boom Ages. Here's How (*Barron's*, 5/21/2021¹⁶). “Home builders are slowing construction, a development—or really a lack thereof—with import for the path of inflation, economic growth, and the stock market.”

The metric this article describes (the relationship between renting and the imputed rent of owned homes) is a useful tool. However, several other factors influence demand for both housing choices. Although the author mentions “a trend toward reurbanization,” there is scant evidence of such a trend. Indeed, many people have left, and continue to leave, large urban areas for reasons other than housing affordability, such as lower costs of living, personal safety, and preferable lifestyles. The work-from-home phenomenon and its hybrids are unlikely to be a passing phenomenon. (S.T.C. 5/25/2021)

U.S. Bankruptcy Tracker: Real Estate Breaks Chapter 11 Lull (Bloomberg, 5/25/2021¹⁷). “The real estate sector last week broke a short-lived lull in U.S. bankruptcy court filings, with two companies seeking protection from creditors.”

This article fails to discuss highly significant elements in the lower numbers of Chapter 11 filings. One such element is the federal government's vast provision of dollars to companies and individuals, the intended purpose of which was to offset, at least in part, the economic impacts of COVID-19. A second important factor is the extremely low interest rates. Such rates have had the effect of reducing debt service levels as well as inflating asset prices (e.g., equities). The expiration of federal assistance will place renewed stress on highly leveraged corporations and will likely result in increased levels of Chapter 11 filings. (S.T.C. 5/26/2021)

In Tight Housing Market, Thousands of Homes Are Reserved for Certain Buyers (*The Wall Street Journal*, 5/26/2021¹⁸). “‘Whisper listings,’ made directly to select customers, are growing at a time when housing inventory is near record lows.”

¹⁵ https://www.wsj.com/articles/u-s-home-prices-push-to-record-high-slowing-pace-of-purchases-11621605953?st=j4ffe5ful1nc4ke&reflink=article_gmail_share.

¹⁶ <https://www.barrons.com/articles/bet-on-a-rental-resurgence-as-housing-boom-ages-heres-how-51621624502?st=m8bog0utm3p5krk>.

¹⁷ <https://www.bloomberg.com/news/articles/2021-05-25/u-s-bankruptcy-tracker-real-estate-breaks-chapter-11-lull>.

¹⁸ <https://www.wsj.com/articles/in-tight-housing-market-thousands-of-homes-are-reserved-for-certain-buyers-11622>

1. Recent Developments

Pocket or whisper listings have always played a role in the real estate market, but these niche listings seem to be most beneficial to sellers in the luxury-home segment, primarily because luxury homeowners' extreme privacy concerns generally make them reluctant to let multiple potential buyers preview their homes via open-house functions. Whisper listings do not serve the nonluxury market very well as they dramatically reduce competitive bidding by substantially reducing exposure and, consequently, produce a smaller buyer audience. In a nutshell, although whisper listings are a beneficial selling tool in the luxury-home segment, they do not appear to benefit the typical-home segment. (J.H.J. 5/26/2021)

These news articles, released as the world began to recover from the COVID-19 pandemic, address the potentially irreversible changes in financial and social systems that impact real estate investing. In this book, we examine the components of an investment approach necessary for optimal adaptation to changing market conditions. We do so in the context of the real estate investment market and use model portfolios of REIT-based assets.

A real estate investment trust (REIT) is a company that uses pooled investor capital to purchase and, typically, operate income-producing property or finance real estate. REITs have a defined investment policy, and most of them specialize in a property type: residential, retail, lodging, healthcare, industrial, office, etc. REITs may be formed as partnerships, trusts, or corporations and must adhere to organizational and income requirements to maintain their status.¹⁹ The three categories of REITs are based on income source: Equity REITs own and rent property, deriving income from rents and capital gains from sales; mortgage REITs make loans, so their income is tied to interest rates; and hybrid REITs combine equity and mortgage activity. REITs can be either private or offered as a publicly traded security. Participation in a privately held REIT generally requires a higher buy-in threshold and tends to be an investment that is less liquid. For the general investor, publicly traded REITs are bought and sold on the open market, giving shareholders access to a percentage of the gain or loss from the income. As with any risky asset, REIT income is volatile and represents only one asset class in any well-diversified portfolio. Within this asset class, diversification and liquidity can be further achieved through purchase of a REIT exchange-traded fund (ETF), a portfolio of REIT assets that trade on an exchange.

Choosing a REIT ETF is of central importance to the general investor. A REIT ETF is designed on the basis of a number of factors, both qualitative and quantitative. The qualitative components generally come from the experience (and philosophy) of the team actively acquiring and managing the portfolio. The quantitative components center on the method for assigning weights to the individual assets in the portfolio. The details of the construction of a REIT ETF and its management are usually available to the public only in general qualitative terms. As a general benchmark, a REIT ETF is designed with the goal of performing as well as, and preferably better than, a particular REIT market index.

A market index is a time series that measures the performance of (a subset of) a financial market. It compares current prices (equivalently, returns) with past prices (returns) to quantify market performance over time. It is computed from the prices of select assets using a method for weighting prices from different assets. To generate user confidence, the index must be both

[021400.](#)

¹⁹ IRC sections 856(a), 856(c), and 856(h).

transparent and investable (Lo, 2016). With respect to transparency, the index-construction method must be specified, typically through a public “methodology” document. As noted above, investability is achieved through the existence of index-tracking ETFs. Current REIT-index-methodology documents exhibit little to no use of modern quantitative risk-management techniques that evaluate the selection and weighting of the assets comprising an index.

Of importance to the real estate investment community is the development of ETFs (portfolios) that have differing risk–return profiles. This is the subject of Chapters 2–7, in which we explore model REIT-based portfolios. *In its most essential terms, portfolio optimization consists of three components: a method of selecting assets for the portfolio, a method of choosing the weight of (percent investment in) each asset in the portfolio, and constraints placed on the optimization, which will limit the assignment of weights.* The first two methods, asset selection and weight determination, may be time varying or fixed (this is determined at the portfolio’s initiation). Generally, a fixed choice of constraints is imposed on the portfolio for finite time periods, though constraints may be changed or updated infrequently.

An exploration of all three of these components is beyond the scope of this study. For the purposes of this book, we consider only fixed asset selection. Chapter 2 provides a description of the REIT assets comprising the model portfolios examined throughout the book. We consider the 26 largest (by market capitalization as of August 2017) domestic U.S. REIT ETFs and the seven largest (by market capitalization as of August 2017) foreign REITs traded over the counter as American depository receipts (ADRs) in the United States. Similarly, Chapter 9 considers the five real estate stocks we use to examine portfolio diversification. We require benchmarks against which to compare the performance of our prototype funds. For these, also listed in Chapter 2, we consider three REIT market indices, two REIT-index-based ETFs, and one broad-based market index. Chapter 2 describes an additional 13 assets, representing major stock market classes, which are used in the factor-analysis section of Chapter 10.

We concentrate on methods for optimizing asset weights. Chapter 3 presents the basic mathematical background of modern portfolio theory (MPT), which is used for the daily optimization of the prototype funds whose performance we explore. MPT concentrates on the maximization of expected portfolio return while minimizing risk. We consider two standard measures of portfolio risk: mean variance and conditional value-at-risk (CVaR). Mean variance captures the “central risk” inherent in a distribution of returns, whereas CVaR concentrates on the risk inherent in the “tails” of a distribution. Because market disruptions (crashes) and market overexuberance (bubbles) correlate strongly with either tail of the distribution of returns, tail risk is a critical measure.²⁰ For each of these two measures, we consider either global minimization of the risk or maximization of the return–risk ratio: specifically, the Sharpe ratio (Sharpe, 1994) in the case of mean variance and the Sortino ratio (Rollinger and Hoffman, undated) in the case of CVaR. We also introduce the general formalism for applying constraints to the optimization. Because this book cannot address the entire range of constraints that may be imposed under this general formalism, we concentrate on a subset of constraints.

Portfolios optimized with respect to mean variance are subject to three well-documented criticisms, which we briefly discuss. In Chapter 3, we also explore the Black–Litterman model

²⁰ See, for example, Guerard et al. (2013) for a comparison of optimization under mean variance and optimization under CVaR in portfolios that utilize a general stock-selection model.

1. Recent Developments

(Black and Litterman, 1991), which mitigates these issues by using a Bayesian approach to incorporate market equilibrium returns and subjective views based on investment analyst estimates.

We perform our initial exploration of these methods in the context of optimization, using a moving window of historical returns to generate optimum portfolio weights for each successive trading day. In Chapter 4, we present results for optimizations performed on the portfolio of domestic REIT ETFs. We consider long-only investment performance as well as several long-short strategies. Virtually every change introduced into the portfolio, whether it be time-varying asset selection or changes in asset weights, will require the purchase or sale of asset shares, incurring transaction costs. Though transaction costs require up-front liquidity, the net effect is the degradation of overall portfolio return. We introduce turnover constraints as a proxy for controlling transaction costs.

A plethora of measures has been introduced to compare the performance of different portfolios (Cogneau and Hübner, 2009). In addition to the portfolio cumulative price and return over time, in Chapter 4 we introduce and utilize four of these performance measures to compare the performance of our model portfolios. These measures are chosen to represent different performance-measure categories introduced by Cogneau and Hübner.

In Chapter 5, we explore portfolio diversification by adding the international REIT assets into the portfolios discussed in Chapter 4.

Chapter 6 explores optimizations under the Black–Litterman model. Because subjective views are specific to the market day and analyst, our generic exploration of this model is restricted to incorporating market equilibrium returns. We use the Fidelity Real Estate Investment Portfolio (FRESX) as the equilibrium benchmark.

Optimization based on a moving window of historical return data provides a restricted set of independent observations that is used to predict the expected performance of the market during the subsequent trading session. In particular, past market behavior may not provide adequate samples of extreme market behavior that could significantly affect the expectation of future performance. In Chapter 7, we introduce dynamic portfolio optimization, which provides a statistically accurate large sample. Specifically, it uses statistical analysis of return data in a historical window to produce a dynamic prediction of returns for the subsequent trading session, which provides a basis for the optimization of asset weights to be used during the trading session after that. We explore the performance of dynamic optimization using the domestic REIT portfolio under long-only, long-short, and Black–Litterman strategies.

Under regulatory guidelines, banks with substantial trading activity are required to set aside capital (the market-risk capital requirement) to insure against extreme portfolio loss. The market-risk capital requirement is set on the basis of estimates of the VaR of the institution's investments. The estimation of this VaR is known as *backtesting*. Chapter 8 presents backtest results for the selected portfolios under both historical and dynamic optimization using long-only investment strategies.

In Chapter 9, we examine diversification by adding real estate stocks into the REIT portfolios discussed in Chapter 4.

It is well-known that no single optimization strategy can provide adequate investment protection. Optimization must be accompanied by a suite of appropriate risk-management tools.

Chapter 10 explores additional risk-information-providing and risk-management tools, namely early warning systems that attempt to discern early signals of market disruption; incremental and component risk measures that guide the reallocation of portfolio assets; and analyses that identify a set of common factors responsible for a significant component of the observed behavior of a portfolio's time-varying return.

In Chapter 11, we extend MPT to consider an approach to portfolio optimization that includes constraints on the values of two performance measures (so-called performance attributes). We consider these in terms of the minimization of CVaR, though the methods are easily adapted to optimizations using other risk measures.

Chapter 12 moves into the realm of insurance instruments by developing option-pricing methods based on our REIT model funds as the underlying asset.

There is a growing push for companies and investment funds to adopt environmental, social, and governance (ESG) criteria. Companies and funds are ranked according to their performance relative to (various) ESG criteria.²¹ Investment practices that prioritize the ESG record or “score” of a company or fund are referred to collectively as socially responsible investing. We anticipate that this trend will engender legislative support. Devoted to an extension of MPT, Chapter 13 moves from the analysis of a risk–return efficient frontier to that of an ESG–risk–return efficient frontier. The result leads to investment portfolios (funds) that maximize both return and ESG measures while minimizing risk.

References

- Black, F. & Litterman, R. (1991). Asset allocation combining investor views with market equilibrium. *Journal of Fixed Income*, 1(2), 7–18.
- Cogneau, P. & Hübner, G. (2009). The 101 ways to measure portfolio performance. *SSRN Electronic Journal*. DOI: 10.2139/ssrn.1326076
- Guerard, J. B., Rachev, S. T. & Shao, B. P. Efficient global portfolios: Big data and investment universes. *IBM Research & Development*, 57(5), paper 11.
- Lo, A. W. (2016). What is an index? *Journal of Portfolio Management*, 42(2), 21–36.
- Rollinger, T.N. & Hoffman, S. T. (undated). Sortino: A “Sharper” Ratio. Red Rock Capital. http://www.redrockcapital.com/Sortino__A__Sharper__Ratio_Red_Rock_Capital.pdf
- Sharpe, W. F. (1994). The Sharpe ratio. *Journal of Portfolio Management*, 21, 49–58.

²¹ See, for example, https://etfdb.com/esg-investing/environmental-issues/green-building/#etfs_esg&sort_name=sather_esg_fields.value&sort_order=desc&page=1.

Chapter 2

The Data

This chapter summarizes all the asset data used in the book. Section 2.1 describes the 26 domestic REIT ETFs and the seven international REITs traded over the counter as ADRs in the United States that are used to develop prototype portfolio-based funds. Section 2.2 describes the five real estate stocks used to diversify the prototype portfolios. Section 2.3 summarizes the benchmark data used to gauge the performance of our prototype portfolios. These benchmarks consist of REIT market indices, REIT-based ETFs, a REIT-based mutual fund, and a general market ETF. In section 2.4, we summarize additional assets representing major stock market classes used in the factor-analysis section of Chapter 10. Price data for all the assets was obtained from Bloomberg Professional Services, with the notable exception of data for two of the REIT market indices, which came from the Federal Reserve Bank of St. Louis. Cumulative price and return performance for the asset data are displayed in sections 2.1–2.3.

2.1 REIT Asset Descriptions

2.1.1 Domestic REITs

Because the United States is the largest public equity market for real estate, it is important to base ETF construction on portfolios of domestic assets. Specifically, our domestic portfolios comprise the 26 largest (by market capitalization as of August 2017) U.S.-traded REITs. These are listed by market capitalization in Table 2.1, which also indicates their *Wall Street Journal* (WSJ) symbol, trust name, inception date, and real estate sector of specialization.²²

Table 2.1 The 26 largest domestic U.S. REIT ETFs by market capitalization (August 2017).

WSJ	Name	Inception Date	Sector	Market Cap (\$bn) ^a
AMT	American Tower Corp.	02/26/1998	Specialty	57.90
SPG	Simon Property Group Inc.	12/13/1993	Retail	50.13
PSA	Public Storage	11/17/1980	Specialty	36.32
CCI	Crown Castle International Corp.	08/17/1998	Specialty	35.61
PLD	Prologis Inc.	11/20/1997	Industrial/Office	31.28
AVB	AvalonBay Communities Inc.	03/10/1994	Residential	26.34
WY	Weyerhaeuser Co.	05/02/1973	Specialty	25.87
EQR	Equity Residential	08/11/1993	Residential	24.66
VTR	Ventas Inc.	05/04/1997	Healthcare	24.36
BXP	Boston Properties Inc.	06/17/1997	Industrial/Office	18.30
UDR	UDR Inc.	03/16/1980	Residential	17.75
ESS	Essex Property Trust Inc.	06/06/1994	Residential	17.37
SBAC	SBA Communications Corp.	06/15/1999	Specialty	16.57
O	Realty Income Corp.	10/17/1994	Retail	15.65
VNO	Vornado Realty Trust	03/16/1980	Diversified	14.90

²² As specified by GFM Asset Management; <https://gfmasset.com/2017/08/top-200-us-listed-real-estate-investment-trusts-reits-by-market-cap-as-of-2017q3/>.

HCP ^a	Healthpeak Properties	05/22/1985	Healthcare	14.80
HST	Host Hotels & Resorts Inc.	03/16/1980	Hotel/Lodging	13.29
NLY	Annaly Capital Management Inc.	10/07/1997	Mortgage	12.81
MAA	Mid-America Apartment Communities Inc.	01/27/1994	Residential	11.49
ARE	Alexandria Real Estate Equities Inc.	05/27/1997	Industrial/Office	11.12
REG	Regency Centers Corp.	10/28/1993	Retail	10.92
SLG	SL Green Realty Corp.	08/14/1997	Industrial/Office	10.38
DRE	Duke Realty Corp.	02/05/1986	Industrial/Office	9.95
FRT	Federal Realty Investment Trust	05/02/1973	Retail	9.39
IRM	Iron Mountain Inc.	01/31/1996	Specialty	9.16
MAC	Macerich Co.	03/09/1994	Retail	8.34

^a Now trading under the symbol PEAK.

What follows is a brief business description²³ of each of the 26 domestic REITs, categorized by real estate investment sector.

Diversified

VNO: Vornado Realty Trust is a fully integrated real estate investment trust that owns, manages, and leases office properties in New York City, Chicago, and San Francisco.

Healthcare

HCP: Healthpeak Properties Inc. is a real estate investment trust that invests in healthcare-related properties throughout the United States. Properties include long-term-care facilities and acute care and medical office buildings.

VTR: Ventas Inc. is a real estate investment trust that owns senior-housing communities, skilled nursing facilities, hospitals, and medical office buildings in the United States and Canada.

Hotel/Lodging

HST: Host Hotels & Resorts Inc. is a real estate trust that owns or holds controlling interests in upscale and luxury full-service hotel lodging properties in many areas, including Washington, DC; Toronto and Calgary, Canada; Mexico City, Mexico; and Santiago, Chile – as well as in Italy, Spain, Poland, Belgium, the Netherlands, and the United States.

Industrial/Office

ARE: Alexandria Real Estate Equities Inc. acquires, manages, expands, and develops office and laboratory properties. The company leases its properties to pharmaceutical, biotechnology, diagnostic, and personal-care-products companies, research institutions, and related government agencies. It serves customers in the state of California.

BXP: Boston Properties Inc. operates as a real estate investment trust. The company owns, manages, and develops office properties in the United States, with a significant presence in Boston, Washington, DC, Midtown Manhattan, and San Francisco.

DRE: Duke Realty Corp. owns interests in industrial, office, and medical office properties across the Southeastern, Midwestern, and Southern United States. The company provides leasing,

²³ See <https://www.bloomberg.com/>.

2. Real Estate Investment Trusts

property and asset management, acquisition, development, construction, build to suit, and other related services.

PLD: Prologis Inc. is an owner, operator, and developer of industrial real estate, focused on global and regional markets across the Americas, Europe, and Asia. The company also leases modern distribution facilities to customers, including manufacturers, retailers, transportation companies, third-party logistics providers, and other enterprises.

SLG: SL Green Realty Corp. is a fully integrated, self-administered, and self-managed real estate investment trust focused exclusively on owning and operating office buildings in Manhattan.

Mortgage

NLY: Annaly Capital Management Inc. is a capital manager that invests in and finances residential and commercial assets. The company's principal business objective is to generate net income for distribution to its stockholders through capital preservation, prudent selection of investments, and continuous management of its portfolio.

Residential

AVB: AvalonBay Communities Inc. is a real estate investment trust that develops, redevelops, acquires, owns, and operates multifamily communities in the United States.

EQR: Equity Residential is a real estate investment trust that acquires, develops, and manages apartment complexes in the United States.

MAA: Mid-America Apartment Communities Inc. is a self-administered and self-managed real estate investment trust that owns, develops, acquires, and operates multifamily apartment communities in the Southeast and Midwest United States and Texas. In addition, the company conducts third-party property management, development, and construction activities through its service corporation.

UDR: UDR Inc. is a self-administered real estate investment trust that owns, operates, and develops apartment communities located nationwide.

Retail

FRT: Federal Realty Investment Trust is a self-administered real estate investment trust that specializes in the ownership, management, development, and redevelopment of prime community and neighborhood shopping centers. Federal Realty Investment Trust serves customers in the United States.

MAC: Macerich Co. is a fully integrated self-managed and self-administered real estate investment trust that focuses on the acquisition, leasing, management, development, and redevelopment of regional shopping malls throughout the United States.

O: Realty Income Corp. owns and manages a portfolio of commercial properties located across the United States. The company focuses on acquiring single-tenant retail locations, which leases to regional and national chains under long-term net-lease agreements.

REG: Regency Centers Corp. is a self-administered and self-managed real estate investment trust that owns and operates grocery-anchored neighborhood retail centers. The company currently

owns and operates properties in various states located throughout the United States.

SPG: Simon Property Group Inc. is a self-administered and self-managed real estate investment trust that owns, develops, and manages retail real estate properties, including regional shopping malls, outlet centers, community/lifestyle centers, and international properties. Simon Property Group serves customers in the state of Indiana.

Specialty

AMT: American Tower Corp. is a real estate investment trust that owns, operates, and develops wireless communications and broadcast towers in the United States. The company leases antenna sites on multitenant towers for a diverse range of wireless communications industries, including personal communications services, paging, and cellular.

CCI: Crown Castle International Corp. operates as a real estate investment trust. The company owns, operates, and leases towers and other infrastructure for wireless communications. Crown Castle manages and offers wireless communication coverage and infrastructure sites in the United States and Australia.

ESS: Essex Property Trust Inc. is a self-administered and self-managed real estate investment trust company specializing in acquiring, developing, and managing multifamily residential properties. Essex has ownership interests in residential properties and commercial properties located in the states of California and Washington.

IRM: Iron Mountain Inc. is a storage and information management company providing records management, data-management solutions, and information-destruction services.

PSA: Public Storage is a real estate investment trust whose principal business activities are the acquisition, development, ownership, and operation of self-storage facilities in the United States. Public Storage also own an equity interest in an owner and operator of self-storage facilities in Europe.

SBAC: SBA Communications Corp. owns and operates wireless communications infrastructure in the United States. The company offers site leasing and development, construction, and consulting services. SBA Communications leases antenna space on its multitenant towers to a variety of wireless service providers under long-term lease contracts.

WY: Weyerhaeuser Co. is an integrated forest-products company with offices and operations worldwide. The company primarily grows and harvests trees, develops and constructs real estate, and makes a range of forest products. Weyerhaeuser is also classified as a REIT.

The daily price data for each REIT was obtained from Bloomberg Professional Services, and it covers the common time period 12/13/1999 through 12/18/2020. Daily returns for each REIT were computed from the price data. To compare the performance of the REITs, we computed a cumulative investment price for each REIT assuming a \$100 (long-only) investment in each on 12/13/1999. The cumulative investment for each of the 26 domestic REITs is displayed in Fig. 2.1, which, for graphical clarity, groups the price plots by real estate investment sector. Fig. 2.1 also shows the resultant cumulative log-return time series for each REIT investment, grouped by the same sectors.

2. Real Estate Investment Trusts

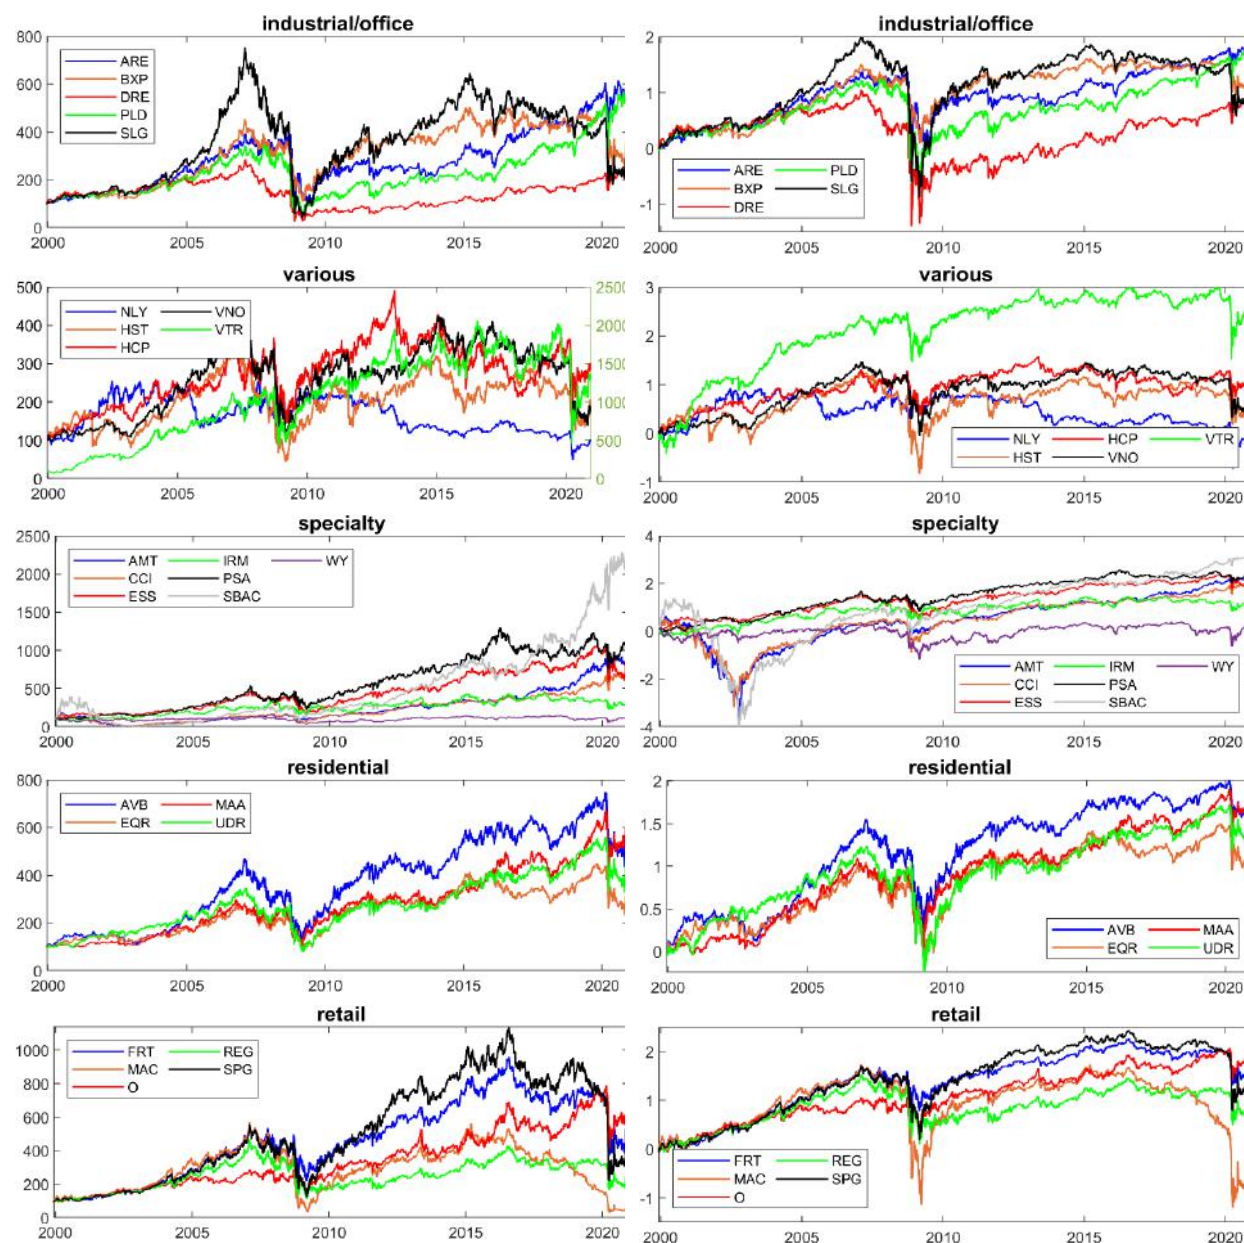


Figure 2.1 Cumulative price (left) and log-return (right) for the U.S.-traded REITs. Each time series assumes a \$100 investment on 12/13/1999. Note that the plot of the VTR time series uses the right-hand axis.

We make the following observations about the performance of these domestic REITs during this time period.

1. VTR was the outstanding performer over this time period. It was strongly affected (maximum drawdown of approximately 50%) by the 2007–2008 financial crisis but recovered strongly, although with noticeable volatility since 2013.
2. In general, the industrial/office-sector REITs were the lowest performers over this time period. These REITs experienced drawdowns exceeding 50% during the 2007–2008 financial crisis. Relative to each other, their postcrisis-recovery performance appears to mimic their precrisis

performance.

3. The performance of the few REITs representing the mortgage, hotel/lodging, and diversified sectors were on par with the performance of the industrial/office sector.
4. The performance of the REITs in the specialty group varied widely, with strong performance by VTR and ESS and weak performance by WY.
5. With the exception of EQR, the residential REITs performed on par with one another. EQR was noticeably relatively unaffected by the 2007–2008 financial crisis.
6. Although all the retail REITs were affected by the 2007–2008 crisis, their recovery after the crisis varied considerably. MAC and REG had difficulty recovering, whereas O recovered well.

2.1.2 International REITs

Studies have validated that international real estate investments can provide significant diversification benefits, including the opportunity to create cash flow and to form long-term wealth (e.g., Worzala and Sirmans, 2003). To address such diversification, we constructed international portfolios comprising the seven largest (by market capitalization as of August 2017) foreign REITs traded over the counter as ADRs in the United States. Table 2.2 lists these seven REITs and includes their WSJ symbol, trust name, inception date, country, and real estate sector of specialization.

Table 2.2 The seven largest international REIT ETFs by market capitalization (August 2017).

Ticker	Name	Inception Date	Sector	Country	Market Cap (\$bn) ^a
LKREF	Link Real Estate Invest. Trust	06/24/2008	Retail	Hong Kong	17.34
BTLCY	British Land Co. PLC ADR	01/30/2008	Diversified	United Kingdom	8.36
JNRFY	Japan Retail Fund Invest. Corp. ADR	09/29/2005	Retail	Japan	5.09
CWYUF	Smart Real Estate Invest. Trust	06/24/2008	Retail	Canada	3.95
CDPYF	Canadian Apartment Properties Real Estate Invest. Trust	11/15/2005	Residential	Canada	3.54
BOWFF	Boardwalk Real Estate Invest. Trust	08/25/2006	Residential	Canada	1.89
CIO	City Office REIT Inc.	04/15/2014	Industry/Office	Canada	0.38

^a <https://gfmasset.com/2017/08/top-200-us-listed-real-estate-investment-trusts-reits-by-market-cap-as-of-2017q3/>

Below, we present brief business summaries of each international REIT, categorized by sector.

2. Real Estate Investment Trusts

Diversified

BTLCY: British Land Co. has a Level 1 ADR program. One British Land ADR represents one existing share of the company. The Bank of New York is the depositary bank.

Industrial/Office

CIO: City Office REIT Inc. is a real estate investment trust that targets office properties in cities in the Southern and Western United States. The Canada-based firm seeks to acquire, own, and operate office properties priced between \$20 million and \$50 million in a dozen target markets with growing populations and above average employment-growth forecasts. City Office REIT, which was spun off in late 2013 by the real-estate-focused private equity fund Second City Group, went public in the United States in 2014.

Residential

BOWFF: Boardwalk Real Estate Investment Trust is a real estate company that acquires and manages multifamily residential projects throughout Western Canada.

CDPYF: Canadian Apartment Properties Real Estate Investment Trust owns and operates a portfolio of multiunit residential rental properties, including apartments, townhomes, and manufactured home communities located in and near major urban centers across Canada. Its portfolio includes fee-simple-interest apartments and townhomes, operating leasehold interests, land leasehold interests, and fee-simple-interest manufactured home community land lease sites. The company was founded in 1997, and is headquartered in Toronto, Canada.

Retail

CWYUF: Smart Centers Real Estate Investment Trust owns and operates shopping centers and focuses on acquisition, asset management, planning, development, leasing, operations, property management, and construction businesses. The trust serves clients in Canada.

JNRFY: Japan Retail Fund Investment Corp. has been listed on the REIT section of the Tokyo Stock Exchange (securities code: 8953) since March 2002. It was the first investment corporation in Japan to specifically target retail property assets. As the largest Japanese REIT that specializes in retail properties, the company strives to secure stable distributions for its unitholders and steady increases in the value of its property portfolio through selective acquisitions of prime retail properties. It operates a Level 1 ADR program.

LKREF: Link Real Estate Investment Trust is a real estate investment trust in Hong Kong. The trust owns and operates shopping centers, parking facilities, and real estate retail space.

Price data on each international REIT was obtained from Bloomberg Professional Services, and it covers the common time period 04/13/2014 through 12/18/2020. To compare performance, a price time series was computed from daily returns assuming a (long-only) investment of \$100 in each REIT on 04/13/2014. The cumulative price and log-return time series are displayed in Fig. 2.2, where the time series are color coded by real estate sector. Only CDPYF and LKREF show positive cumulative return over this time period.




Figure 2.2 Cumulative price (left) and log-return (right) for the international REITs. Each time series assumes a \$100 investment on 4/13/2014.

2.2 Real Estate Stocks

In Chapter 9, we consider diversification of the domestic REIT portfolio by the addition of real estate company (non-REIT) stocks. Table 2.3 lists the five stocks added to the portfolio and includes their New York Stock Exchange ticker symbol, company name, inception date, country, and real estate sector of specialization. The choice of companies was based on a requirement that daily stock prices from 12/13/1999 through 12/18/2020 were available from Bloomberg Professional Services. Brief business summaries of each company follow the table.

Table 2.3 The five real estate stocks used in Chapter 9.

Ticker	Name	Inception Date	Country
BVH	Bluegreen Vacations Holding Corp.	11/10/1997	United States
JLL	Jones Lang LaSalle Inc.	07/17/1997	United States
MLP	Maui Land & Pineapple Company Inc.	04/25/1994	United States
NTP	Nam Tai Property Inc.	08/06/1991	China
TRC	Tejon Ranch Co.	07/28/1980	United States

BVH: Bluegreen Vacations Holding Corp. is a Florida-based holding company whose sole investment is 100% ownership of Bluegreen Vacations Corp. Bluegreen Vacations markets and sells vacation ownership interests and manages resorts in popular leisure and urban destinations. The Bluegreen Vacation Club is a flexible, points-based, deeded vacation ownership plan with 68 Club and Club Associate Resorts and access to nearly 11,300 other hotels and resorts through partnerships and exchange networks. Bluegreen Vacations also offers a portfolio of comprehensive, fee-based resort-management, financial, and sales and marketing services to, or on behalf of, third parties.

JLL: Jones Lang LaSalle Inc., founded in the United Kingdom, is a global commercial real estate services company with offices in 80 countries. The company buys, builds, occupies, and invests in a variety of assets, including industrial, commercial, retail, residential, and hotel real estate. The company also provides investment management services worldwide, including services to institutional and retail investors and to high-net-worth individuals. Its services include investment management, asset management, sales and leasing, property management, project management, and development. The company is ranked 179th on the Fortune 500.

2. Real Estate Investment Trusts

MLP: Maui Land & Pineapple Company Inc. is a land holding and operating company founded in 1909 and based in Kapalua, Hawaii. It owns approximately 24,300 acres on the island of Maui. It develops, sells, and manages residential, resort, commercial, and industrial real estate, and it operates retail, golf, and utility operations at the Kapalua Resort. MLP also owns and manages the 8,304-acre Pu‘u Kukui Watershed Preserve, one of the largest private nature preserves in Hawaii. It formerly grew pineapples.

NTP: Nam Tai Property Inc. is a China-based holding company that conducts business primarily in Mainland China. Through its subsidiaries, the company is engaged in the development and operation of technology parks. Its main land resources are located in Shenzhen and Wuxi, China. The company provides industrial offices, industrial service spaces, and supporting dormitories to the park tenants, and it provides comprehensive industrial services to corporate tenants through a full-chain industrial model. Its main projects and properties include Nam Tai Inno Park, Nam Tai Technology Center, Nam Tai Inno Valley, Wuxi Facilities, Nam Tai Tang Xi Technology Park, and Nam Tai U-Creative Space (Lujiazui).

TRC: Tejon Ranch Co., based in Lebec, California, is one of the largest private landowners in California. The company was incorporated in 1936 to organize the ownership of a large tract of land that was consolidated from four Mexican land grants acquired in the 1850s and 1860s by ranch founder General Edward Fitzgerald Beale. The company owns over 270,000 acres in the Southern San Joaquin Valley, Tehachapi Mountains, and Antelope Valley. Tejon Ranch’s agricultural operation primarily grows almonds, pistachios, and wine grapes, along with some alfalfa and the occasional row crop. Cattle leases cover about 250,000 acres, providing grazing for as many as 12,000 head of cattle.




Figure 2.3 Cumulative price (left) and log-return (right) for the real estate stocks. Each time series assumes a \$100 investment on 12/13/1999.

Fig. 2.3 displays cumulative price and log-return time series for each stock, assuming a \$100 investment on 12/13/1999. MLP has struggled to recover from the 2008 recession; as of 12/18/2021, it is still showing negative cumulative return. Although TRC was impacted less by the recession, as of 12/18/2021, its cumulative return is also negative. The remaining three real estate stocks show positive cumulative return, with JLL being the strongest performer, though experiencing a high degree of price volatility.

2.3 Benchmarks

2.3.1. Indices

We utilize the following three REIT market indices as benchmarks against which to compare the performance of the prototype ETFs we develop in later chapters.

WILLREITIND: Wilshire U.S. REIT Total Market Index

WILLREITPR: Wilshire U.S. REIT Price Index

Values for the indices WILLREITIND²⁴ and WILLREITPR²⁵ are provided by the Federal Reserve Bank of St. Louis. They are based on the Wilshire U.S. Real Estate Investment Trust IndexSM (Wilshire U.S. REIT), which measures U.S. publicly traded real estate investment trusts and is itself a subset of the Wilshire U.S. Real Estate Securities IndexSM (Wilshire U.S. REIT). The Wilshire U.S. REIT index is designed to offer a market-based index that is more reflective of real estate held by pension funds, because indexes are unencumbered by the limitations of other appraisal-based indexes. They can serve as proxies for direct real estate investing by excluding securities whose value is not always tied to the value of the underlying real estate. Exclusions include mortgage REITs, net-lease REITs, real estate finance companies, mortgage brokers and bankers, commercial and residential real estate brokers, home builders, large landowners and subdividers of unimproved land, hybrid REITs, and timber REITs. The rationale for the exclusions listed is that factors other than real estate supply and demand, such as interest rates, influence the market value of these companies.²⁶

FNRE: FTSE Nareit Equity REITs Index

FNRE is part of the FTSE Nareit U.S. Real Estate Index Series.²⁷ In particular, it is a sector-specific subindex of the FTSE Nareit Composite Index. This investment sector includes all Equity REITs not designated as Timber REITs or Infrastructure REITs. Timber REITs invest in timber assets, including timberland and timber-related products and activities. Infrastructure REITs invest in infrastructure assets, including transportation assets (roads, bridges, tunnels, airports, etc.), energy and utilities assets (power generation, fuels, pipelines, etc.), water and waste management assets and communication assets (line-based networks, air-based networks, etc.). Infrastructure REITs do not include data center REITs. Details of the FTSE Nareit U.S. Real Estate Index Series can be found in the series rules.²⁸

For notational brevity, the REIT indices WILLRIETIND and WILLREITPR are referred to as WD and WP in the remainder of this book.

2.3.2 Exchange Traded Funds

²⁴ <https://fred.stlouisfed.org/series/WILLREITIND>

²⁵ <https://fred.stlouisfed.org/series/WILLREITPR>

²⁶ <https://wilshire.com/indexes/wilshire-real-estate-family/wilshire-us-reit>. The fact sheet and methodology for the WILLREITIND index can be downloaded from this web page.

²⁷ See <https://www.ftserussell.com/products/indices/nareit>.

²⁸ https://research.ftserussell.com/products/downloads/FTSE_Nareit_US_Real_Estate_Index_Series.pdf.

2. Real Estate Investment Trusts

Because indices are not traded, it is also important to benchmark against appropriate index-tracking ETFs. We consider two REIT-index-based ETFs and one broad-based stock market index. The stock market index is used as an out-of-class benchmark to represent the general U.S. (and to some extent, the broader worldwide) stock market.

VNQ: Vanguard Real Estate ETF

VNQ²⁹ is the most investor-popular REIT ETF, with over \$50 billion in assets. Employing over 180 holdings, the fund currently seeks to track the MSCI U.S. IMI Real Estate 25/50 Index, providing broad exposure to the U.S. real estate market. Note that the MSCI U.S. IMI Real Estate 25/50 Index was launched on 09/01/2016; due to its limited period of existence, our study does not consider it a benchmark index.

USRT: iShares Core U.S. REIT ETF

USRT³⁰ seeks to track the FT index. As of this writing, it employs 155 holdings. The fund generally invests at least 90% of its assets in the component securities of the underlying index and may invest up to 10% of its assets in certain futures, options and swap contracts, cash, and cash equivalents.

SPY: SPDR S&P 500 TRUST ETF

SPY³¹ is an ETF designed to track the S&P 500 Index.

2.3.3 Mutual Funds

Our implementation of the Black–Litterman model (section 3.6) requires a market index. In our example computations, we use the following mutual fund for this purpose.

FRESX: Fidelity Real Estate Investment Portfolio

This fund seeks above-average income and long-term capital growth, consistent with reasonable investment risk. The fund seeks to exceed the composite yield of the S&P 500 Index. At least 80% of the portfolio assets are in securities (primarily common stock) of companies principally engaged in the real estate industry and other real-estate-related investments.

Price data on the benchmarks was obtained from Bloomberg Professional Services, and it covers the time period 12/18/2007 through 12/18/2020. To compare performance, a price time series was computed from daily returns assuming a (long-only) investment of \$100 in each index or fund on 12/18/2007. Fig. 2.4 displays the cumulative price and log-return time series. SPY and the (untraded) WD index are competitive performers over this time period, with the WD indices closely tracking each other. Over this period, VNQ, USRT, and the (untraded) WP and FNRE indices underperform. The performance of VNQ closely tracks the indices WP and FNRE. The performance of the mutual fund FRESX is intermediate between that of the low and high performers over this period.

²⁹ <https://investor.vanguard.com/etf/profile/VNQ>.

³⁰ <https://www.ishares.com/us/products/239544/ishares-core-us-real-estate-etf>.

³¹ See <https://www.etf.com/SPY>.




Figure 2.4 Cumulative price (left) and log-return (right) for the benchmarks. Each time series assumes a \$100 investment on 12/18/2007.

2.4 Additional Assets and Indices

In the factor-analysis section of Chapter 10, we utilize 13 additional assets representing major market classes:³² equity, fixed income, and commodity. A brief business summary of each asset is presented here, ordered by the assets' symbols as referenced by Bloomberg Professional Services.

AGG: iShares Core U.S. Aggregate Bond ETF

An ETF incorporated in the United States that seeks to track the Bloomberg Barclays U.S. Aggregate Bond Index. The fund invests in securities within the total U.S. investment-grade bond market. This includes treasuries, government-related and corporate securities, mortgage backed securities, asset backed securities, and commercial mortgage backed securities.

EEM: iShares MSCI Emerging Market

An ETF incorporated in the United States that tracks the performance of the MSCI TR Emerging Markets Index. It holds emerging market stocks that can be classified predominantly as large- and mid-cap. It weights the holdings using a market capitalization methodology and rebalances quarterly.

EFA: iShares MSCI EAFE ETF

An ETF incorporated in the United States that tracks the performance of the MSCI EAFE Index. It holds large- and mid-cap stocks. Its investments are focused on developed countries across the world, excluding the United States and Canada. The ETF weights the holdings using a market capitalization methodology.

HYG: iShares iBoxx \$ High Yield Corporate Bond ETF

An ETF incorporated in the United States that seeks to track the investment results of an index composed of U.S.-dollar-denominated, high-yield corporate bonds.

IXIC: NASDAQ Composite

Developed with a base level of 100 as of February 5, 1971, this index is a broad-based capitalization-weighted index of stocks in all three NASDAQ tiers: Global Select, Global Market, and Capital Market.

MSCI: MSCI Inc.

MSCI Inc., formerly Morgan Stanley Capital International, provides investment-decision-

³² See <https://seekingalpha.com/article/4342477-major-asset-classes-april-2020-performance-review>.

2. Real Estate Investment Trusts

support tools to investment institutions worldwide. The company produces indices and risk-and-return portfolio analytics for use in managing investment portfolios. MSCI manages more than 145,000 daily equity, fixed income, and hedge fund indices for use by large asset management firms. MSCI is organized through two business segments: (i) Performance and Risk, which provides equity indices, portfolio-risk and performance analytics, credit analytics, and ESG products, and (ii) Governance, which provides corporate governance and specialized financial research and analysis. The company's indices act as benchmarks that measure the performance of global funds. Institutional investors use the indices as research tools and as the basis of their various investment vehicles. MSCI's Performance and Risk segment is by far its largest, accounting for 87% of the company's revenue in 2012. MSCI makes the majority of its revenues (more than 75%) from annual recurring subscriptions to its products, with nearly half of the company's revenues coming from outside the Americas.

RUA: The Russell 3000

This index is composed of 3,000 large U.S. companies, as determined by market capitalization. This portfolio of securities represents approximately 98% of the investable U.S. equity market. The index was developed with a base value of 140.00 as of 12/31/1986.

SPBDUB3T: S&P U.S. Treasury Bill 0–3 Month Index Total Return

This index, a subindex of the S&P/BGCantor U.S. Treasury Bill Index, comprises the constituents that have a maturity of zero to three months.

SPGI: S&P Global Inc.

This company provides clients with financial information services. It offers information regarding ratings, benchmarks, and analytics in the global capital and commodity markets. It operates worldwide.

TIP: iShares TIPS Bond ETF

An ETF incorporated in the United States designed to track the Bloomberg Barclays Capital U.S. Treasury Inflation Notes Index. This index measures the performance of inflation-protected public obligations of the U.S. Treasury, also known as TIPS.

USO: United States Oil Fund

An ETF that attempts to track the price of West Texas Intermediate Light Sweet Crude Oil. USO invests in oil future contracts that are traded on regulated futures exchanges.

VBINX: Vanguard Balanced Index Fund Investor Shares

The fund employs an indexing investment approach designed to track the performance of the CRSP U.S. Total Market Index and Bloomberg Barclays U.S. Aggregate Float Adjusted Index.

XAU/USD: Gold spot price in U.S. dollars/oz.

2.5 Data Observations

We close this chapter with two observations.

- Our use of market capitalization to choose REIT assets does bias the ETF prototypes developed in this book; that is, we know, a posteriori, that all the REITs in our sample survived the 2008 financial crisis. An examination of the effects of survivor bias on portfolios can be found, for

example, in Bloch et al. (1993).

- A major benefit of using a sample of REITs with high liquidity and predominant ownership by institutional investors (average ownership of 88%)³³ is that our results are unlikely to be the product of micromarket trading frictions and/or retail investor noise trading and should thus provide a cleaner test of the viability of the various risk–return portfolio-optimization techniques we examine in this book.

References

- Bloch, M., Guerard, J. B., Markowitz, H., Todd, P. & Xu, G. (2013). A comparison of some aspects of the U.S. and Japanese equity markets. *Japan and the World Economy*, 5, 3–26.
- Worzala, E. & Sirmans, C. (2003). Investing in international real estate stocks: A review of the literature. *Urban Studies*, 40(5–6), 1115–1149.

³³ Institutional ownership (IO) for 25 of the 26 domestic REITs is based on data obtained from www.marketbeat.com as of February 15, 2020. Macerich Company (MAC) is excluded from the average due to anomalous data. (Data from <https://money.cnn.com> and www.nasdaq.com list IO exceeds 100% for MAC.)

Chapter 3

Modern Portfolio Theory

This chapter introduces the basic elements of MPT that lie behind the development of our REIT-based prototype ETFs. In section 3.1, we discuss stationarity and the need to consider returns rather than prices. Both discrete and log-returns are introduced. Section 3.2 introduces MPT optimization strategies. Markowitz optimization, which minimizes the mean variance of the return, is described in section 3.2.1. That section also introduces the concept of the efficient frontier, which is central to MPT. Section 3.2.2 covers optimization designed to maximize the Sharpe ratio and demonstrates the relationship between the efficient frontier and the capital market line. Sections 3.2.3 and 3.2.4 extend the MPT of sections 3.2.1 and 3.2.2, respectively, to the case in which the risk measure is CVaR rather than mean variance. Critiques of mean-variance optimization are summarized in section 3.2.5, and the Black–Litterman model, created to address these problems, is covered in section 3.3. The implementation of these optimization techniques using moving windows of historical asset-return data is outlined in section 3.4.

3.1 Return Time Series

A stochastic process is stationary if its unconditional joint-probability distribution and all its distribution measures do not change over time. A process is weakly stationary if, at minimum, the mean and the covariance of the random variable are time independent. A great deal of time-series analysis is founded on the assumption of weak stationarity. Nonstationary behavior is determined mathematically by the presence of a unit root (a root of value 1) in the characteristic equation governing the process. Nonstationary processes are often (though not always) indicated by the presence of trends in the value of the random variable that change over time. Each price time series for the individual REITs listed in Chapter 2 is tested for the presence of a unit root using the augmented Dickey–Fuller test.³⁴ At a 95% confidence level, these tests are unable to reject the null hypothesis of the presence of a unit root for any of the price series, indicating elements of nonstationarity. Visual examination reveals clear deterministic trends as well as a market bubble collapse (the 2007–2009 Great Recession) and the 2020 COVID-19 pandemic.

Using nonstationary time-series data in financial models can produce unreliable or spurious results, which lead to poor forecasting. The solution is to transform the time-series data so that it becomes stationary. The technique of “differencing” is used to correct for nonstationarity. Applying first differences (see (3.1)) to a price time series results in the well-known return time series. Specifically, two types of return result from first differences. For asset i , the discrete (or simple) return, $R_i(t)$, and the log-return, $r_i(t)$, are defined as³⁵

$$\begin{aligned} R_i(t) &= (P_i(t) - P_i(t-1))/P_i(t-1) = P_i(t)/P_i(t-1) - 1, \\ r_i(t) &= \ln P_i(t) - \ln P_i(t-1) = \ln(P_i(t)/P_i(t-1)), \end{aligned} \quad (3.1)$$

where $P_i(t)$ is the closing price for asset i on day t . These two returns are related:

³⁴ See, for example, https://en.wikipedia.org/wiki/Augmented_Dickey%E2%80%93Fuller_test for a brief synopsis.

³⁵ See, for example, Chapter 2 of Tsay (2010).

$$R_i(t) + 1 = P_i(t)/P_i(t-1) = \exp(r_i(t)), \quad (3.2)$$

or

$$r_i(T) = \ln(1 + R_i(t)) \approx R_i(t). \quad (3.3)$$

The approximation $r_i(T) \approx R_i(t)$ holds quite well when daily asset returns are less than 1%.

At close of trading, the daily discrete return, $R_p(t)$, of a portfolio is compactly expressed in terms of the discrete returns of its assets:

$$R_p(t) = \sum_{i=1}^n w_i(t) r_i(t) = \mathbf{w}^T(t) \mathbf{r}(t), \quad (3.4)$$

where $w_i(t)$ are the asset weights applied to the portfolio during trading day t , whereas $r_i(t)$ are the asset returns computed at the close of trading day t . The second equality in (3.4) expresses the portfolio simple return using vector notation.³⁶ The relation between the discrete returns and log-returns for a portfolio is the same as that given in (3.2):

$$R_p(t) + 1 = P_p(t)/P_p(t-1) = \exp(r_p(t)), \quad (3.5)$$

where $P_p(t)$ is the value (price) of the portfolio at the close of trading day t . Note that a portfolio's cumulative log-return $r_{p,c}(T)$ satisfies

$$r_{p,c}(T) = \sum_{t=1}^T r_p(t) = \ln(P_p(T)/P_p(0)), \quad (3.6)$$

where $P_p(0)$ represents the initial investment in the portfolio. In contrast, a portfolio's cumulative discrete return $R_{p,c}(T)$ is given by

$$1 + R_{p,c}(T) = \prod_{t=1}^T (1 + R_p(t)) = P_p(T)/P_p(0). \quad (3.7)$$

As we show below, Markowitz portfolio optimizations are written in terms of discrete returns. Although the approximation (3.3) holds most of the time, such that log-returns can replace discrete returns in the optimization, for cases in which daily returns become large (e.g., $> 1\%$), the errors introduced by using log-returns in the Markowitz optimizations can become significant. For this reason, we perform optimizations using discrete asset returns. Log-returns for individual assets or for the portfolio can then be computed from (3.2) or (3.5).

In (3.1)–(3.7), we are careful to distinguish discrete returns and log-returns with notation. For the rest of the book, we drop that distinction and use $r_i(t)$ (or the vector notation $\mathbf{r}(t)$) to refer to return. Unless otherwise noted, the return is always the discrete return.

3.2 MPT-Based Portfolios

³⁶ By default, a vector \mathbf{v} is assumed to be a column vector, whereas its transpose, written \mathbf{v}^T , is a row vector. Thus, the column vector, \mathbf{v} , having elements v_1, \dots, v_n can be written $\mathbf{v} = (v_1, \dots, v_n)^T$, and the row vector \mathbf{v}^T can be written $\mathbf{v}^T = (v_1, \dots, v_n)$.

Consider a portfolio consisting of n risky assets whose vector of daily³⁷ return values³⁸ $\mathbf{r}(t) = (r_1(t), \dots, r_n(t))^T$ is considered a multivariate random variable having mean value $\bar{\mathbf{r}} = (\bar{r}_1, \dots, \bar{r}_n)^T$ and standard deviation $\boldsymbol{\sigma} = (\sigma_1, \dots, \sigma_n)^T$. The returns between different assets are generally correlated, leading to an $n \times n$, nondiagonal, covariance matrix $\boldsymbol{\Sigma}$ whose elements are $\Sigma_{ij} \equiv \sigma_{ij} = E[(r_i - \bar{r}_i)(r_j - \bar{r}_j)]$, where $E[\cdot]$ stands for expected value. The diagonal elements of $\boldsymbol{\Sigma}$ are the variances of the individual assets, $\Sigma_{ii} = \sigma_i^2$. (When $i \neq j$, σ_{ij} is the covariance between assets i and j .) If $\mathbf{w}(t) = (w_1(t), \dots, w_n(t))^T$ are the weights applied³⁹ to the assets during day t , then the return of the portfolio at the close of day t will be given by (3.4), $r_p(t) = \mathbf{r}^T(t)\mathbf{w}(t)$. Because it is assumed that the portfolio is fully invested in the n risky assets, we require the condition $\sum_{i=1}^n w_i(t) = 1$ or, in vector notation, $\mathbf{e}_n^T \mathbf{w}(t) = 1$, where \mathbf{e}_n is the n -dimensional unit vector.⁴⁰ Because the asset returns are random variables, the portfolio return r_p is also a random variable having mean value $\bar{r}_p(t) = \bar{\mathbf{r}}^T \mathbf{w}(t) = \mathbf{w}^T(t) \bar{\mathbf{r}}$ and variance $\sigma_p^2 = \mathbf{w}^T(t) \boldsymbol{\Sigma} \mathbf{w}(t)$. For notational brevity, we drop the time notation for the remainder of this chapter.

3.2.1 Markowitz Mean-Variance Portfolio

The objective of portfolio optimization under the Markowitz (1952) method is to determine the daily set of weights that minimize the return risk of the portfolio (for that day) subject to a desired expected return \bar{r}_p . The desired value of \bar{r}_p reflects the risk aversion of the investor; the larger the value, the greater the risk the investor is willing to take on. Using the variance σ_p^2 of the portfolio as the proxy for risk, we can express the Markowitz mean (return)-variance optimization as the desire to minimize the portfolio variance subject to desired expected return and full investment in the assets, that is,

$$\text{minimize } \mathbf{w}^T \boldsymbol{\Sigma} \mathbf{w} \quad \text{subject to } \bar{\mathbf{r}}^T \mathbf{w} = \bar{r}_p \quad \text{and} \quad \mathbf{e}_n^T \mathbf{w} = 1. \quad (3.8)$$

Because the desired return \bar{r}_p is changed, the optimum solution $(\sigma_p(\mathbf{w}(\bar{r}_p)), \bar{r}_p)$ traces out a curve known as the portfolio frontier. Fig 3.1 shows an example of a portfolio frontier.

Equation (3.8) is solved via standard Lagrange multiplier methods:

$$\min_{\mathbf{w}, q, \theta_0} L(\mathbf{w}, q, \theta_0) = \min_{\mathbf{w}, q, \theta_0} (\mathbf{w}^T \boldsymbol{\Sigma} \mathbf{w} / 2 + q(\bar{r}_p - \bar{\mathbf{r}}^T \mathbf{w}) + \theta_0(1 - \mathbf{e}_n^T \mathbf{w})). \quad (3.9)$$

In the Lagrangian function $L(\mathbf{w}, q, \theta_0)$, q and θ_0 are Lagrange multipliers representing, respectively, the constraints on $\bar{\mathbf{r}}^T \mathbf{w}$ and $\mathbf{e}_n^T \mathbf{w}$ in (3.8). Applying the optimizing conditions⁴¹ $\partial L(\mathbf{w}, q, \theta_0) / \partial w_i = 0$, $i = 1, \dots, n$; $\partial L(\mathbf{w}, q, \theta_0) / \partial q = 0$; and $\partial L(\mathbf{w}, q, \theta_0) / \partial \theta_0 = 0$, (3.9) becomes the following system of linear equations:

³⁷ “Daily” refers only to trading days.

³⁸ These are assumed to be simple return values, $r_i(t) = (P_i(t) - P_i(t-1)) / P_i(t-1)$, where $P_i(t)$ is the daily closing prices of asset i .

³⁹ The weights are applied at the beginning of trading day t and are assumed not to change during the trading day.

⁴⁰ $\mathbf{e}_n = (1, \dots, 1)^T$.

⁴¹ Because $L(\mathbf{w}, q, \theta_0)$ is quadratic in \mathbf{w} with a positive coefficient on the quadratic term, the optimizing condition is equivalent to a minimizing one.

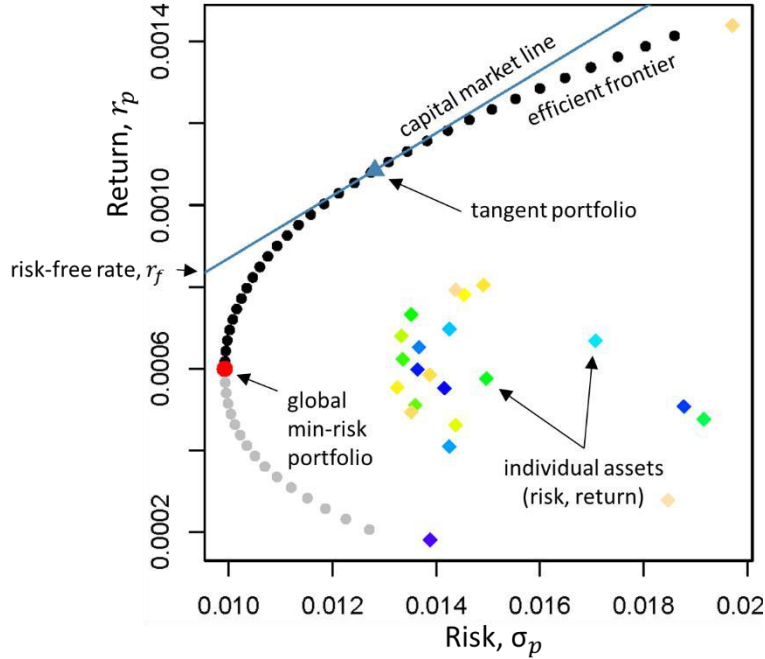


Figure 3.1. Computational example showing the fundamental elements of modern portfolio theory optimization: efficient frontier, global minimum-risk and tangent portfolios, capital market line, risk-free rate, and the (risk, mean-return) coordinates of individual assets. The portfolio frontier consists of the union of the efficient frontier (black points) with the curve segment described by the gray points. The computation shown was performed using the standard deviation of the portfolio, σ_p , as the risk measure. The picture would look similar if the risk measure were portfolio conditional value-at-risk, CVaR_α .

$$\begin{aligned}\Sigma \mathbf{w} - q\bar{\mathbf{r}} - \theta_0 \mathbf{e}_n &= 0, \\ \bar{\mathbf{r}}^T \mathbf{w} &= \bar{r}_p, \\ \mathbf{e}_n^T \mathbf{w} &= 1.\end{aligned}\tag{3.10}$$

Equation (3.10) can be solved analytically, leading to the solution

$$\mathbf{w}^* = \bar{r}_p \mathbf{w}_1 + \mathbf{w}_2,\tag{3.11}$$

where

$$\begin{aligned}\mathbf{w}_1 &= \frac{1}{\Delta} (B\Sigma^{-1}\bar{\mathbf{r}} - C\Sigma^{-1}\mathbf{e}_n), & \mathbf{w}_2 &= \frac{1}{\Delta} (A\Sigma^{-1}\mathbf{e}_n - C\Sigma^{-1}\bar{\mathbf{r}}), \\ A &= \bar{\mathbf{r}}^T \Sigma^{-1} \bar{\mathbf{r}}, & B &= \mathbf{e}_n^T \Sigma^{-1} \mathbf{e}_n, \\ C &= \bar{\mathbf{r}}^T \Sigma^{-1} \mathbf{e}_n, & \Delta &= AB - C^2.\end{aligned}\tag{3.12}$$

Solving for $\sigma_p = \sqrt{\mathbf{w}^{*T} \Sigma \mathbf{w}^*}$ gives

$$\sigma_p = \sqrt{\frac{B\bar{r}_p^2 - 2C\bar{r}_p + A}{\Delta}}.\tag{3.12a}$$

3. Modern Portfolio Theory

Thus, the portfolio frontier points $(\sigma_p(\mathbf{w}^*), \bar{r}_p)$ trace out a hyperbola in standard deviation–return space.

Equation (3.9) has a global minimum at the value $\mathbf{w}^*_{\min} = \mathbf{\Sigma}^{-1}\mathbf{e}_n/B$. This portfolio has expected return $\bar{r}_{p,\min} = C/B$ and standard deviation $\sigma_{p,\min} = \sqrt{1/B}$, and it is referred to as the *Markowitz minimum-variance portfolio*. Points on the portfolio frontier that have expected mean returns greater than $\bar{r}_{p,\min}$ are said to lie on the *efficient frontier*. (A portfolio – i.e., a solution (3.11) – is on the efficient frontier if no other portfolio having the same standard deviation but higher expected return can exist.) Fig. 3.1 identifies the efficient frontier. Because each individual asset has a (historical) mean-return value, \bar{r}_i , and a (historical) standard deviation, σ_i , each asset comprising the portfolio can also be represented (with a weight of 1) on the same plot as the efficient frontier, as indicated in Fig. 3.1.

Additional constraints can be added to the minimization problem. For example, the minimization

$$\min_{\mathbf{w}, q, \theta_0, \theta_1, \theta_2} L(\mathbf{w}, q, \theta_0, \theta_1, \theta_2) = \min_{\mathbf{w}, q, \theta_0, \theta_1, \theta_2} [\mathbf{w}^T \mathbf{\Sigma} \mathbf{w} / 2 + q(\bar{r}_p - \bar{\mathbf{r}}^T \mathbf{w}) + \theta_0(1 - \mathbf{e}_n^T \mathbf{w}) + \theta_1^T (\mathbf{w}_{\text{lb}} - \mathbf{w}) + \theta_2^T (\mathbf{w} - \mathbf{w}_{\text{ub}})] \quad (3.13)$$

imposes two additional constraints: that the weights exceed some lower bound of values, \mathbf{w}_{lb} , and do not exceed an upper bound of values, \mathbf{w}_{ub} .⁴² Again, the elements of the vectors θ_1 and θ_2 are Lagrange multipliers.⁴³ The full set of equations

$$\begin{aligned} \partial L(\mathbf{w}, q, \theta_0, \theta_1, \theta_2) / \partial w_i &= 0, \quad i = 1, \dots, n; & \partial L(\mathbf{w}, q, \theta_0, \theta_1, \theta_2) / \partial q &= 0; \\ \partial L(\mathbf{w}, q, \theta_0, \theta_1, \theta_2) / \partial \theta_0 &= 0; & \partial L(\mathbf{w}, q, \theta_0, \theta_1, \theta_2) / \partial \theta_{1,i} &= 0, \quad i = 1, \dots, n; \\ \partial L(\mathbf{w}, q, \theta_0, \theta_1, \theta_2) / \partial \theta_{2,i} &= 0, \quad i = 1, \dots, n \end{aligned}$$

provides sufficient conditions to solve for the values $\mathbf{w}, \theta_0, \theta_1$, and θ_2 .⁴⁴ Equation (3.13) can specify a long-only strategy using $w_{\text{lb}} = 0, \theta_2 = 0$ ⁴⁵ or a long–short strategy by setting appropriate entries of w_{lb} to be negative and $\theta_2 = 0$. Similarly, it is possible to constrain a subset, s , of the portfolio assets to account for a specified fraction, f_s , of the portfolio weight by employing a term $\theta_s(f_s - \mathbf{e}_s^T \mathbf{w})$ in the minimizing Lagrangian function and by adding the conditions $\partial L(\mathbf{w}, q, \theta_0, \theta_1, \theta_2, \theta_s) / \partial \theta_{s,i} = 0, i = 1, \dots, n$ to provide the additional equations necessary to solve for θ_s in the resulting linear system. Here, entries in the vector \mathbf{e}_s are given by

$$e_{s,i} = \begin{cases} 1, & \text{if asset } i \text{ is in subset } s, \\ 0, & \text{otherwise.} \end{cases} \quad (3.14)$$

The appropriate use of \mathbf{e}_s^T and f_s provides the basis for implementing the Lo–Patel 130/30 type investment strategy discussed in section 4.1.3.

As the examples (3.9), (3.12), and (3.13) indicate, the exact shape traced out by the efficient frontier depends not only on the risky assets (specifically, their number, their mean return $\bar{\mathbf{r}}$, and

⁴² Specifically, $w_{\text{lb},i} \leq w_i \leq w_{\text{ub},i}, i = 1, \dots, n$.

⁴³ They are the vectors $\theta_1^T = (\theta_{1,1}, \dots, \theta_{1,n})$, $\theta_2^T = (\theta_{2,1}, \dots, \theta_{2,n})$.

⁴⁴ Some constraints may not, in fact, affect the minimizing solution. Standard techniques recognize when this occurs and eliminate such constraints from consideration.

⁴⁵ For a long-only strategy, it is sufficient to require $0 \leq w_i, i = 1, \dots, n$ and $\mathbf{e}_n^T \mathbf{w} = 1$.

covariance structure Σ) but also on the constraints imposed. Mean-variance efficient frontiers numerically computed under limits on the percentage holding of any one asset can be found, for example, in Bloch et al. (1993).

3.2.2 Capital Market Line and the Markowitz Mean-Variance Tangent Portfolio

We now consider portfolios that include a single riskless asset having expected return r_f along with the n risky assets. Assume the weight of the riskless asset in the portfolio is w_f . Now, $\mathbf{e}_n^T \mathbf{w} + w_f = 1$. We formulate the optimization problem

$$\begin{aligned} & \min_{\mathbf{w}} \mathbf{w}^T \Sigma \mathbf{w} / 2, \\ & \text{such that } (1 - \mathbf{e}_n^T \mathbf{w}) r_f + \bar{\mathbf{r}}^T \mathbf{w} = \bar{r}_p. \end{aligned} \quad (3.15)$$

Note that the second equation combines the restrictions $\mathbf{e}_n^T \mathbf{w} + w_f = 1$ and $w_f r_f + \bar{\mathbf{r}}^T \mathbf{w} = \bar{r}_p$ into a single restriction. As a Lagrange multiplier minimization problem, this is

$$\min_{\mathbf{w}, q} L(\mathbf{w}, q) = \min_{\mathbf{w}, q} \left(\mathbf{w}^T \Sigma \mathbf{w} / 2 + q(\bar{r}_p - (1 - \mathbf{e}_n^T \mathbf{w}) r_f - \bar{\mathbf{r}}^T \mathbf{w}) \right). \quad (3.16)$$

Again, \bar{r}_p acts as a parameter (desired expected return). The analytic solution to the system $\partial L(\mathbf{w}, q) / \partial w_i = 0$; $i = 1, \dots, n$, $\partial L(\mathbf{w}, q) / \partial q = 0$, is

$$\begin{aligned} \mathbf{w} &= q(\bar{r}_p) \Sigma^{-1} (\bar{\mathbf{r}} - r_f \mathbf{e}_n), \\ q(\bar{r}_p) &= \frac{\bar{r}_p - r_f}{B r_f^2 - 2C r_f + A}, \end{aligned} \quad (3.17)$$

where A, B , and C are given by (3.12). Evaluating the portfolio standard deviation $\sigma_p = \sqrt{\mathbf{w}^T \Sigma \mathbf{w}}$ reveals

$$\sigma_p(\bar{r}_p) = (\bar{r}_p - r_f) / \sqrt{B r_f^2 - 2C r_f + A} \equiv a \bar{r}_p - b, \quad (3.18)$$

where the last form emphasizes that σ_p varies linearly with \bar{r}_p . Thus, the risk–return profile $(\sigma_p(\bar{r}_p), \bar{r}_p)$ for the portfolio follows a straight line, the capital market line (CML), as \bar{r}_p (equivalently, w_f) is varied.

- When $w_f = 1$, the CML intersects the return axis at $(0, r_f)$.
- When $w_f = 0$, the CML intersects the efficient frontier tangentially⁴⁶ at the “market” portfolio value (σ_m, \bar{r}_m) .

The condition $w_f = 0$ is equivalent to the statement that $\mathbf{e}_n^T \mathbf{w} = 1$. From (3.17), this gives the weights \mathbf{w}_m for the market portfolio:

$$\mathbf{w}_m = \frac{\Sigma^{-1} (\bar{\mathbf{r}} - r_f \mathbf{e}_n)}{C - r_f B}. \quad (3.19)$$

⁴⁶ It takes some computation to show that (3.20) satisfies (3.12a) – that is, that (σ_m, \bar{r}_m) is on the efficient frontier and that the slope of the efficient frontier determined by (3.12a) at (σ_m, \bar{r}_m) is indeed the slope of the CAPM line.

3. Modern Portfolio Theory

The mean return, \bar{r}_m , of the market portfolio is given by $\bar{r}_m = \bar{\mathbf{r}}^T \mathbf{w}_m$, and its standard deviation, σ_m , is given by (3.18) when $\bar{r}_p = \bar{r}_m$. Evaluating these quantities gives the following analytic results:

$$\begin{aligned}\bar{r}_m &= r_f + \frac{Br_f^2 - 2Cr_f + A}{C - r_f B}, \\ \sigma_m &= \frac{\sqrt{Br_f^2 - 2Cr_f + A}}{C - r_f B}.\end{aligned}\tag{3.20}$$

From (3.20), we have the identity $(\bar{r}_m - r_f)/\sigma_m = \sqrt{Br_f^2 - 2Cr_f + A}$. This identity can be used to rewrite equation (3.18) for the CML in the familiar form

$$\bar{r}_p = r_f + \frac{\bar{r}_m - r_f}{\sigma_m} \sigma_p,\tag{3.20a}$$

which is the statement that all the portfolios (σ_p, \bar{r}_p) on the CML have the same Sharpe ratio (Sharpe, 1994) as the market portfolio:

$$\frac{\bar{r}_p - r_f}{\sigma_p} = \frac{\bar{r}_m - r_f}{\sigma_m}.$$

The CML and market (tangent) portfolio are also illustrated in Fig. 3.1. Equation (3.20a) can be written in the form of the capital asset pricing model (CAPM), $\bar{r}_p = r_f + \beta_p(\bar{r}_m - r_f)$, where $\beta_p = \sigma_p / \sigma_m$ is the CML portfolio sensitivity (its *beta*) of its expected excess return to the expected excess return of the market portfolio (3.20). (The value of β_p goes from 0 when $\sigma_p = 0$ [i.e., $\bar{r}_p = r_f$], goes to 1 when $\sigma_p = \sigma_m$ [i.e., $\bar{r}_p = \bar{r}_m$], and exceeds 1 when $\bar{r}_p > \bar{r}_m$.)

The portfolio defined by (3.19) and (3.20) – and more generally, by the point on the efficient frontier to which the CML is the tangent – defines the Markowitz portfolio that maximizes the Sharpe ratio. It is also referred to as the *tangent mean-variance portfolio*.

3.2.3 CVaR-Minimizing Portfolios

The Markowitz mean-variance model establishes a framework for building similar strategies that use other measures as proxies for the portfolio risk. We now discuss portfolio optimization when CVaR is chosen as the portfolio risk measure.⁴⁷ As implied by its name, CVaR is related to the concept of VaR, which is defined as follows. Let $F(x) = \Pr\{\bar{r}_p \leq x\}$ denote the cumulative distribution function of the daily return \bar{r}_p of a portfolio. Then

$$\text{VaR}_\alpha(\bar{r}_p) = -\inf\{x \in \mathbb{R} \mid F(x) \geq \alpha\}.\tag{3.21}$$

The negative sign in (3.21) is a convention under which VaR values reported for losses are positive. We refer to α as the tail probability. Conceptually, (3.21) states that if the portfolio return, \bar{r}_p , is a random variable having $\text{VaR}_\alpha(\bar{r}_p) = x_\alpha$, then for α -percent of the time, the portfolio return will

⁴⁷ The Basel III regulatory framework for banks requires CVaR as the risk measure.

be a value that is less than (more negative than) $-x_\alpha$;⁴⁸ that is, the loss (in positive dollars) will exceed x_α . Thus, an equivalent statement is that x_α is the value corresponding to the α -th percentile of $F(x)$. $\text{CVaR}_\alpha(\bar{r}_p)$ is defined as an average over all values $\text{VaR}_\gamma(\bar{r}_p)$, where $0 \leq \gamma \leq \alpha$:

$$\text{CVaR}_\alpha(\bar{r}_p) = \frac{1}{\alpha} \int_0^\alpha \text{VaR}_\gamma(x) d\gamma. \quad (3.22)$$

A daily $\text{VaR}_{0.05}$ value of \$1 million means there is a 5% chance that a portfolio could lose \$1 million or more during a one-day period (assuming no changes in the portfolio over the day). In contrast, a daily $\text{CVaR}_{0.05}$ value is an average of VaR values (for that portfolio) corresponding to all probabilities between 0% and 5%.⁴⁹ Thus, a daily value of $\text{VaR}_{0.05} = \$1$ million with a $\text{CVaR}_{0.05}$ value of \$5 million indicates a 5% chance not only that the loss will exceed \$1 million but also that the average loss expected will in fact be \$5 million. For any given value of α , $\text{CVaR}_\alpha(\bar{r}_p) \geq \text{VaR}_\alpha(\bar{r}_p)$.

The objective of the CVaR portfolio-optimization method (Rockafellar and Uryasev, 2000, 2002; Krokmal et al., 2002; Tütüncü et al., 2003) is to maximize return while minimizing CVaR, with CVaR most commonly measured at either the 95% or 99% quantile level (i.e., $1 - \alpha = 0.95$ or 0.99). Conceptually, the minimum CVaR_α portfolio-optimization problem is formulated as

$$\min_{\mathbf{w}} \text{CVaR}_\alpha(\mathbf{w}) \quad \text{subject to } \bar{\mathbf{r}}^T \mathbf{w} = \bar{r}_p \quad \text{and} \quad \mathbf{e}_n^T \mathbf{w} = 1. \quad (3.23)$$

Unlike minimization of the variance, $\mathbf{w}^T \Sigma \mathbf{w}$, which is naturally expressed as quadratic in \mathbf{w} and leads automatically to a quadratic minimization problem, casting (3.23) into a form convenient for solution requires some work. In particular, the dependence of CVaR_α on the asset weights \mathbf{w} must be elucidated. To do so, the random portfolio return \bar{r}_p can be expressed as samples from a joint distribution function⁵⁰ $f(\mathbf{w}, \mathbf{r})$ of the asset weights and mean returns. Let $p(\mathbf{r})$ denote the probability density function that determines the daily asset returns \mathbf{r} . For any fixed value of \mathbf{w} , the cumulative distribution of the daily portfolio returns is given by

$$F(\mathbf{w}, x) = \int_{f(\mathbf{w}, \mathbf{r}) \leq x} p(\mathbf{r}) d\mathbf{r}. \quad (3.24)$$

Then

$$\text{VaR}_\alpha(\mathbf{w}) = -\inf \{x \in \mathbb{R} \mid F(\mathbf{w}, x) \geq \alpha\}, \quad (3.25)$$

and $\text{CVaR}_\alpha(\mathbf{w})$ can be expressed as

$$\text{CVaR}_\alpha(\mathbf{w}) = \frac{1}{\alpha} \int_{-f(\mathbf{w}, \mathbf{r}) \geq \text{VaR}_\alpha(\mathbf{w})} (-f(\mathbf{w}, \mathbf{r})) p(\mathbf{r}) d\mathbf{r}. \quad (3.26)$$

As required in (3.23), (3.26) achieves the objective of expressing CVaR_α in terms of asset weights, \mathbf{w} . However, (3.26) is not in a convenient form. Note that (3.25) and (3.24) give the relation

⁴⁸ Recall that VaR is defined as positive for losses.

⁴⁹ In addition, CVaR satisfies the desirable attributes of a coherent risk measure and is consistent with performance relations of risk-averse investors (see Pflug, 2000).

⁵⁰ It is common in the literature to define $f(\mathbf{w}, \mathbf{r})$ alternatively as the *loss* probability function, whose samples provide negative return values.

$$\alpha = F(\mathbf{w}, \text{VaR}_\alpha(\mathbf{w})) = \int_{-f(\mathbf{w}, \mathbf{r}) \geq \text{VaR}_\alpha(\mathbf{w})} p(\mathbf{r}) d\mathbf{r}. \quad (3.27)$$

If we add the term $\text{VaR}_\alpha(\mathbf{w})$ to and subtract the equivalent term $\alpha^{-1}\alpha\text{VaR}_\alpha(\mathbf{w})$ from (3.26) and use (3.27) to express α , we produce the following alternate form for $\text{CVaR}_\alpha(\mathbf{w})$ (Rockafellar and Uryasev, 2000):

$$\text{CVaR}_\alpha(\mathbf{w}) = \text{VaR}_\alpha(\mathbf{w}) + \frac{1}{\alpha} \int_{R^n} (-f(\mathbf{w}, \mathbf{r}) - \text{VaR}_\alpha(\mathbf{w}))^+ p(\mathbf{r}) d\mathbf{r}. \quad (3.28)$$

Here, $y^+ \equiv \max(0, y)$. If $f(\mathbf{w}, \mathbf{r})$ is a convex function of \mathbf{w} , then $\text{CVaR}_\alpha(\mathbf{w})$ is a convex function of \mathbf{w} , guaranteeing the existence of a minimum value.

In practice, (3.28) is evaluated for a portfolio using a sample of its asset returns \mathbf{r}_t , $t = 1, \dots, T$, from a finite historical time period T . The discrete form of (3.28) is therefore

$$\text{CVaR}_\alpha(\mathbf{w}) = \text{VaR}_\alpha(\mathbf{w}) + \frac{1}{\alpha T} \sum_{t=1}^T \max(0, -\mathbf{r}_t^T \mathbf{w} - \text{VaR}_\alpha(\mathbf{w})). \quad (3.29)$$

Using (3.29), we can approximate the minimization problem (3.23) with the discrete form

$$\begin{aligned} \min_{\mathbf{w}} \left\{ \text{VaR}_\alpha(\mathbf{w}) + \frac{1}{\alpha T} \sum_{t=1}^T y_t \right\}, \\ \text{such that } y_t \geq \max(-\mathbf{r}_t^T \mathbf{w} - \text{VaR}_\alpha(\mathbf{w}), 0), \\ \bar{\mathbf{r}}^T \mathbf{w} = \bar{r}_p, \quad \mathbf{e}_n^T \mathbf{w} = 1. \end{aligned} \quad (3.30)$$

In (3.30), $\bar{\mathbf{r}} = T^{-1} \sum_{t=1}^T \mathbf{r}_t$ is the vector of asset mean returns computed from the time period $[0, T]$.

The relationship between $\text{VaR}_\alpha(\mathbf{w})$ and $\text{CVaR}_\alpha(\mathbf{w})$ is such that the value of \mathbf{w} that minimizes $\text{CVaR}_\alpha(\mathbf{w})$ also minimizes $\text{VaR}_\alpha(\mathbf{w})$. This leads to the following approach (Rockafellar and Uryasev, 2000; Tütüncü et al., 2003):

$$\begin{aligned} \min_{\mathbf{w}, \gamma} F_\alpha(\mathbf{w}, \gamma) \quad \text{where} \\ F_\alpha(\mathbf{w}, \gamma) = \gamma + \frac{1}{\alpha} \int_{R^n} (-f(\mathbf{w}, \mathbf{r}) - \gamma)^+ p(\mathbf{r}) d\mathbf{r}. \end{aligned} \quad (3.31)$$

The optimization function $F_\alpha(\mathbf{w}, \gamma)$ is convex in γ , and it is convex with respect to (\mathbf{w}, γ) if $f(\mathbf{w}, \mathbf{r})$ is convex in \mathbf{w} (Rockafellar and Uryasev, 2000). In this approach, the discrete minimization problem (3.30) becomes

$$\begin{aligned} \min_{\mathbf{w}, \gamma} \left\{ \gamma + \frac{1}{\alpha T} \sum_{t=1}^T y_t \right\}, \\ \text{such that } y_t \geq \max(-\mathbf{r}_t^T \mathbf{w} - \gamma, 0), \\ \bar{\mathbf{r}}^T \mathbf{w} = \bar{r}_p, \quad \mathbf{e}_n^T \mathbf{w} = 1. \end{aligned} \quad (3.32)$$

Under this approach, the optimization function and constraints are manifestly linear in both \mathbf{w} and γ , making the optimization problem a linear programming problem solvable

by, for example, the simplex method.

As Tütüncü et al. (2003) note, the constraint on y_t guarantees only that $y_t \geq \max(0, -\mathbf{r}_t^T \mathbf{w} - \text{VaR}_\alpha(\mathbf{w}))$. However, the minimization of the objective function guarantees that at the minimizing value \mathbf{w}^* , $y_t = \max(0, -\mathbf{r}_t^T \mathbf{w}^* - \text{VaR}_\alpha(\mathbf{w}^*))$.

As the parameter \bar{r}_p varies, the minimizing solution $(\text{CVaR}_\alpha(\mathbf{w}^*), \bar{r}_p(\mathbf{w}^*))$ produces the portfolio frontier curve in risk (CVaR_α) – return (r_p) space. The portfolio frontier will have a convex (though not hyperbolic) form, with a single overall minimum value $\text{CVaR}_{\alpha, \min}$ such that $\text{CVaR}_\alpha(\mathbf{w}^*) \geq \text{CVaR}_{\alpha, \min}$ for every minimizing solution \mathbf{w}^* . The portfolio \mathbf{w}_{\min}^* satisfying $\text{CVaR}_\alpha(\mathbf{w}_{\min}^*) = \text{CVaR}_{\alpha, \min}$ is known as the *minimum CVaR $_\alpha$ portfolio*. All the points on the portfolio frontier having $\bar{r}_p(\mathbf{w}^*) > \bar{r}_p(\mathbf{w}_{\min}^*)$ comprise the *efficient frontier*.

3.2.4 Capital Market Line and the CVaR $_\alpha$ Tangent Portfolio

We again consider portfolios that include a single riskless asset having expected return r_f along with the n risky assets. Following the notation in sections 3.2.2 and 3.2.3, we formulate the optimization problem

$$\begin{aligned} & \min_{\mathbf{w}} \text{CVaR}_\alpha(\mathbf{w}), \\ \text{such that } & (1 - \mathbf{e}_n^T \mathbf{w})r_f + \bar{\mathbf{r}}^T \mathbf{w} = \bar{r}_p. \end{aligned} \quad (3.33)$$

Using the discrete form (3.32), this can be expressed as

$$\begin{aligned} & \min_{\mathbf{w}, \gamma} \left\{ \gamma + \frac{1}{\alpha T} \sum_{t=1}^T y_t \right\}, \\ \text{such that } & y_t \geq \max(-\mathbf{w}^T \mathbf{r}_t - (1 - \mathbf{e}_n^T \mathbf{w})r_f - \gamma, 0), \\ & (1 - \mathbf{e}_n^T \mathbf{w})r_f + \bar{\mathbf{r}}^T \mathbf{w} = \bar{r}_p. \end{aligned} \quad (3.34)$$

At the minimizing values $\mathbf{w}^*, \gamma^* = \text{VaR}_\alpha(\mathbf{w}^*)$, either $y_t = 0$ or $y_t = -\mathbf{w}^{*T} \mathbf{r}_t - (1 - \mathbf{e}_n^T \mathbf{w}^*)r_f - \gamma^* > 0$. Let $S = \{t \mid y_t > 0\}$ and $|S|$ denote the number of entries in S . Then

$$\text{CVaR}_\alpha(\mathbf{w}^*) = \gamma^* + \frac{1}{\alpha T} \sum_{t \in S} (-r_{p,t}(\mathbf{w}^*) - \gamma^*) = \left(1 - \frac{|S|}{\alpha T}\right) \gamma^* - \frac{f_s}{\alpha} \bar{r}_p, \quad (3.35)$$

$$\text{where } f_s \equiv \frac{1}{T \bar{r}_p} \sum_{t \in S} r_{p,t}(\mathbf{w}^*).$$

Because $\bar{r}_p = (1 - \mathbf{e}_n^T \mathbf{w}^*)r_f + \bar{\mathbf{r}}^T \mathbf{w}^* = T^{-1} \sum_{t=1}^T r_{p,t}(\mathbf{w}^*)$, f_s is the fraction of \bar{r}_p contributed by the returns in the set S .⁵¹ Equation (3.35) shows that the curve $(\text{CVaR}_\alpha(\mathbf{w}^*), \bar{r}_p)$ is a straight line, the CML in (CVaR, return) coordinates. Because $\bar{r}_p - r_f = (\bar{\mathbf{r}} - r_f \mathbf{e}_n)^T \mathbf{w}^*$, it is evident that \mathbf{w}^* scales linearly as the difference $\bar{r}_p - r_f$, with $\mathbf{w}^* \rightarrow 0$ as $\bar{r}_p \rightarrow r_f$. Thus, for $\mathbf{w}^* = 0$, we have no investment in risky assets and $\text{CVaR}_\alpha(\mathbf{w}^* \rightarrow 0) \rightarrow 0$. The slope of the straight line is the Sortino

⁵¹ The returns in the set S are negative.

ratio, $(\bar{r}_p - r_f)/\text{CVaR}_\alpha(\mathbf{w}^*)$ (Rollinger and Hoffman, undated). This line touches the efficient frontier tangentially at the point $(\text{CVaR}_\alpha(\mathbf{w}_s^*), \bar{\mathbf{r}}^T \mathbf{w}_s^*)$, where \mathbf{w}_s^* is the scaled value satisfying $\mathbf{e}_n^T \mathbf{w}_s^* = 1$. The portfolio having weights \mathbf{w}_s^* is known as the *tangent CVaR $_\alpha$ portfolio*. It maximizes the value of the Sortino ratio of the portfolio for a given value of $\text{CVaR}_\alpha(\mathbf{w}_s^*)$.

To conclude our discussion of CVaR-based portfolios in sections 3.2.3 and 3.2.4, we note that it is conventional to quote CVaR (and VaR) values based on the quantile level, that is, $1 - \alpha$. Thus, a tail probability, $\alpha = 0.05$, can be referenced as a 0.95 or a 95% quantile level. Thus, the notation $\text{CVaR}_{0.05}(r)$, referring to the fifth percentile, is often written as $\text{CVaR}_{95}(r)$, referring equivalently to the 95% quantile level. The use of $\text{CVaR}_{0.05}(\bar{r}_p)$ results in the statement that there is a 5% probability that returns will be worse than $x_{0.05}$, whereas the use of $\text{CVaR}_{95}(\bar{r}_p)$ results in the more positive statement that there is a 95% probability that returns will be better than $x_{0.05}$. We follow the convention of using percent levels to label CVaR and VaR values.

3.2.5 Criticisms of Mean-Variance Optimization

Simply stated, mean-variance optimization takes as input the expected mean returns and the covariance of the assets and produces optimized weights for each asset. Mean-variance optimization methods have three well-documented issues. To some extent, these issues also arise when the variance risk measure is replaced by that of CVaR.

- Input sensitivity. He and Litterman (1999) demonstrate how a small change in the expected return of a few assets can give rise to dramatic changes in optimized weights, not only for those few assets but also, unexpectedly, for assets whose expected returns were not changed.
- Unintuitive, highly concentrated portfolios. Best and Grauer (1991) examine the sensitivity of portfolios that consist of 10–100 assets to an increase in the expected return of a single asset. Specifically, they seek to determine the size of a single-asset increase needed to drive out of the portfolio at least one-half of the assets. As the number of assets in the portfolio grows, the increase in a single asset required to remove a fixed percentage of the assets from the portfolio decreases rapidly.
- Estimation error maximization. As noted by Michaud (1989, pp. 33-34), “Mean-variance optimization significantly overweights (underweights) those [assets] that have large (small) estimated returns, negative (positive) correlations and small (large) variances. These [assets] are, of course, the ones most likely to have large estimation errors.”

3.3 Black–Litterman Model

Black and Litterman (1991) created a portfolio-construction method to address problems with mean-variance optimization. The Black–Litterman model overcomes the input-sensitivity problem by using a Bayesian approach to incorporate subjective views based on investment-analyst estimates and market-equilibrium returns (Kolm and Ritter, 2017). It also “largely mitigates” (Lee, 2000) the problem of estimation-error maximization. By combining analyst views and equilibrium returns instead of relying only on historical asset returns, the Black–Litterman model provides a systematic way to evaluate the mean and covariance of asset returns. The model uses a Bayesian

approach in which prior (historical) knowledge of the distribution of the returns of the assets in a portfolio are combined with quantifiable views of the management team regarding future returns to estimate an (improved) posterior distribution of asset returns. One apparent restriction of the Black–Litterman method is that it relies on the additivity property of Gaussian (normal) distributions. This property will not hold if individual-asset returns are governed by non-Gaussian distributions.

Assume the n assets in the portfolio have a return vector, \mathbf{r} , that follows a multivariate normal distribution:

$$\mathbf{r} \sim N(\bar{\mathbf{r}}, \Sigma). \quad (3.36)$$

The mean return $\bar{\mathbf{r}}$ is an unknown parameter that the Black–Litterman model estimates by combining a management team's views with prior knowledge of $\bar{\mathbf{r}}$. This prior knowledge assumes that $\bar{\mathbf{r}}$ is also a normally distributed random variable:

$$\bar{\mathbf{r}} \sim N(\boldsymbol{\pi}, \mathbf{C}). \quad (3.37)$$

In the absence of a management team's views, this prior mean return $\boldsymbol{\pi}$ is most likely to be in equilibrium with the market. (This could also be interpreted in other ways. For example, $\boldsymbol{\pi}$ could be the returns for a target optimal performance of the portfolio, such as a benchmark or index, or even just for the current portfolio.) The Black–Litterman model assumes that the covariance matrix \mathbf{C} of the prior distribution of $\bar{\mathbf{r}}$ is

$$\mathbf{C} = \tau \Sigma, \quad (3.38)$$

where τ is a small constant (Attilio, 2006).

The experience of a management team provides an additional set of k views regarding the expected returns $\bar{\mathbf{r}}$. Each view consists of a statement of the expected performance of one asset or a combination of two or more assets in the portfolio. For example, a single view might be a statement about

- the absolute return of an asset – for example, SPG will underperform the equilibrium return by 1.25% (confidence of view = 25%);
- the relative performance between two assets – for example, HST will outperform NLY by 0.25% (confidence of view = 50%); or
- the relative performance between asset classes – for example, healthcare REITs will cumulatively outperform specialty REITs by 2% (confidence of view = 65%).

Each statement of a view contains (i) a quantitative measure (the expected performance), (ii) a specification of the asset(s) involved in the view, and (iii) an uncertainty (or level of confidence) of the view. Thus, a single view can be represented as

$$q_i = E[\mathbf{p}_i^T \mathbf{r} | \bar{\mathbf{r}}] + \varepsilon_i, \quad i = 1, \dots, k.$$

Here, q_i is the quantitative value assigned to view i , which is described in terms of an expected value $E[\cdot | \bar{\mathbf{r}}]$ predicated on $\bar{\mathbf{r}}$, as well as an uncertainty, ε_i . Here \mathbf{p}_i is a vector (of 0s and 1s) that identifies which asset returns are involved. The full set of views can be written as the single vector statement

3. Modern Portfolio Theory

$$\mathbf{q} = \mathbb{E}[\mathbf{P} \mathbf{r} \mid \bar{\mathbf{r}}] + \boldsymbol{\varepsilon}, \quad (3.39)$$

where \mathbf{P} is the matrix whose i th row identifies the assets involved in view i , $i = 1, \dots, k$. The uncertainties in the views, ε_i , are assumed to be normal random variables:

$$\boldsymbol{\varepsilon} \sim N(0, \boldsymbol{\Omega}). \quad (3.40)$$

If the ε_i are independent of each other, then $\boldsymbol{\Omega}$ is a diagonal matrix, $\text{diag}(w_i, \dots, w_k)$. Under assumption (3.36), (3.39) can be written

$$\mathbf{q} = \mathbf{P} \bar{\mathbf{r}} + \boldsymbol{\varepsilon}. \quad (3.41)$$

As noted, the Black–Litterman model uses the Bayesian framework to combine prior knowledge about $\bar{\mathbf{r}}$ with the likelihood that management views will produce an improved posterior estimate, $\text{posterior} \propto \text{likelihood} \times \text{prior}$, of the distribution of asset returns:

$$f(\bar{\mathbf{r}} \mid \mathbf{q}) \propto f(\mathbf{q} \mid \bar{\mathbf{r}}) f(\bar{\mathbf{r}}). \quad (3.42)$$

If we continue to assume normal (Gaussian) statistics, the likelihood is given by

$$f(\mathbf{q} \mid \bar{\mathbf{r}}) \propto \exp \left[-\frac{1}{2} \boldsymbol{\varepsilon}^T \boldsymbol{\Omega}^{-1} \boldsymbol{\varepsilon} \right] = \exp \left[-\frac{1}{2} (\mathbf{P} \bar{\mathbf{r}} - \mathbf{q})^T \boldsymbol{\Omega}^{-1} (\mathbf{P} \bar{\mathbf{r}} - \mathbf{q}) \right] \quad (3.43)$$

and the prior is expressed as

$$f(\bar{\mathbf{r}}) \propto \exp \left[-\frac{1}{2} (\bar{\mathbf{r}} - \boldsymbol{\pi})^T \mathbf{C}^{-1} (\bar{\mathbf{r}} - \boldsymbol{\pi}) \right]. \quad (3.44)$$

The posterior distribution is therefore

$$f(\bar{\mathbf{r}} \mid \mathbf{q}) \propto \exp \left\{ -\frac{1}{2} [(\bar{\mathbf{r}} - \boldsymbol{\pi})^T \mathbf{C}^{-1} (\bar{\mathbf{r}} - \boldsymbol{\pi}) + (\mathbf{P} \bar{\mathbf{r}} - \mathbf{q})^T \boldsymbol{\Omega}^{-1} (\mathbf{P} \bar{\mathbf{r}} - \mathbf{q})] \right\}, \quad (3.45)$$

which is again a normal distribution. Completing the squares on the argument of the exponential in (3.45) to cast it in the explicit normal form $(\bar{\mathbf{r}} - \mathbf{r}_{\text{BL}})^T \boldsymbol{\Sigma}_{\text{BL}}^{-1} (\bar{\mathbf{r}} - \mathbf{r}_{\text{BL}})/2$ leads to the following identifications:

$$\begin{aligned} \mathbf{r}_{\text{BL}} &= \boldsymbol{\Sigma}_{\text{BL}} (\mathbf{C}^{-1} \boldsymbol{\pi} + \mathbf{P}^T \boldsymbol{\Omega}^{-1} \mathbf{q}), \\ \boldsymbol{\Sigma}_{\text{BL}} &= (\mathbf{C}^{-1} + \mathbf{P}^T \boldsymbol{\Omega}^{-1} \mathbf{P})^{-1}. \end{aligned} \quad (3.46)$$

Equation (3.46) explicitly shows how the Black–Litterman model’s prediction of portfolio returns, \mathbf{r}_{BL} , is expressed as a combination of “equilibrium” returns $\boldsymbol{\pi}$ and management-team views, \mathbf{P}, \mathbf{q} . As (3.38) indicates, the value of τ^{-1} sets the relative weight given to $\boldsymbol{\pi}$. The work of He and Litterman (1999) uses a value of $\tau = 0.025$. Other authors have suggested using $\tau = 1/n$ (Atillio, 2006).

Combining the prediction for the posterior distribution governing $\bar{\mathbf{r}}$ with the prior distribution (3.36) for \mathbf{r} leads to the posterior distribution

$$\mathbf{r} \sim N(\mathbf{r}_{\text{BL}}, \boldsymbol{\Sigma} + \boldsymbol{\Sigma}_{\text{BL}}). \quad (3.47)$$

The primary products of the Black–Litterman method, the mean-return values \mathbf{r}_{BL} and covariance matrix $\boldsymbol{\Sigma} + \boldsymbol{\Sigma}_{\text{BL}}$, are natural inputs to a Markowitz mean-variance or tangent portfolio optimization. As noted above, these inputs overcome the input-sensitivity problem and “largely mitigate” the

problem of estimation error maximization.

By specifying management-team views and determining the returns, $\boldsymbol{\pi}$, this section provides a complete specification of the Black–Litterman method. Clearly, management views are portfolio specific and time specific. However, it is possible to provide general approaches to determining $\boldsymbol{\pi}$, including the use of *reverse optimization*. With reference to (3.9), consider the optimization problem

$$\min_{\mathbf{w}} E(\mathbf{w}) = \min_{\mathbf{w}} \left(\mathbf{w}^T \boldsymbol{\Sigma} \mathbf{w} / 2 + \gamma^{-1} (\bar{r}_p - \boldsymbol{\pi}^T \mathbf{w}) \right). \quad (3.48)$$

(Note that (3.48) does not impose any constraints on the asset weights.) The minimizing solution satisfies

$$\mathbf{w}^* = \gamma^{-1} \boldsymbol{\Sigma}^{-1} \boldsymbol{\pi}, \quad \text{i.e.} \quad \boldsymbol{\pi} = \gamma \boldsymbol{\Sigma} \mathbf{w}^*. \quad (3.49)$$

Equation (3.49) can be used to relate the returns $\boldsymbol{\pi}$ to the (properly constrained) weights of an “equilibrium” portfolio. One such approach is to use linear regression to determine the weights $\mathbf{w}^* = \mathbf{w}_{\text{mkt}}$ for the portfolio such that its mean return (linearly) tracks the return of a specific (market) benchmark over a specified historical time period. Constraints can be added to the regression in order to require, for example, that the weights be nonnegative (long-only investment) and that they sum to unity.

In (3.49), γ can be viewed as a risk-aversion coefficient; the smaller the value of γ , the greater the aversion to risk. An extant quantitative method of assigning a value to γ is to express risk aversion as $\gamma = (E(r) - r_f) / \sigma^2$ and realize this as

$$\gamma = S_{\text{bm}} / \sigma_{\text{mkt}}, \quad (3.50)$$

where $S_{\text{bm}} = \bar{r}_{\text{bm}} / \sigma_{\text{bm}}$ is the Sharpe ratio of the benchmark (mean benchmark return/standard deviation of benchmark returns), which is computed assuming a risk-free rate of zero, and σ_{mkt} is the standard deviation of the market portfolio, $\sigma_{\text{mkt}} = (\mathbf{w}_{\text{mkt}}^T \boldsymbol{\Sigma} \mathbf{w}_{\text{mkt}})^{1/2}$.

3.4 Historical Optimization

In Chapters 4–6, we apply these optimization procedures to portfolios of REIT assets. We do so in the context of optimization based solely on historical return data. Fig. 3.2 illustrates the procedure. A rolling window of length T days selects successive historical data periods. (Fig. 3.2 illustrates the data period $[t - T + 1, t]$.) Each time period provides the daily returns of the assets (based on close-of-trading-day prices). Information about the asset returns is fed into the optimizer, which produces a predicted set of optimal portfolio weights for day $t + 1$. The application of these weights to the portfolio is considered to take place at the beginning of day $t + 1$, and the weights remain unchanged throughout the day. Thus, the predicted return of the optimized portfolio at the end of day $t + 1$ would have been $r_{p,t+1} = \mathbf{w}_{t+1}^T \mathbf{r}_{t+1}$. The rolling window is then advanced to cover the period $[t - T + 2, t + 1]$, and optimized weights are produced for day $t + 2$. Thus, for example, if a data set of returns covers 11 years of returns (with 252 daily returns per year), say for the trading days falling into the period 1/1/2010 through 12/31/2020, then with a rolling window of length for years ($T = 4 * 252$), optimized portfolio weights would be computed for the

3. Modern Portfolio Theory

trading days lying within the seven-year period 1/1/2014 through 12/31/2020. The daily optimized portfolio returns, $r_{p,t} = \mathbf{w}_t^T \mathbf{r}_t$, computed for this seven-year period could be compared against actual returns of a market benchmark to determine how well the historical portfolio might have performed. This routine provides some measure of confidence in determining how well the optimization would work in practice.




Figure 3.2. Schematic of optimization based on historical returns.

Fig. 3.3 summarizes the historical information required by each of the optimization routines we have addressed. The time-varying measures $\sigma_i(t)$ and $\text{CVaR}_{\alpha,i}(t)$ for any asset are measures of the volatility of the asset returns. In the moving-window method depicted in Fig. 3.2, these measures are encapsulated as averages developed over the asset-return values in the window. The average value is assumed not to change over the time period covered by the window. For this

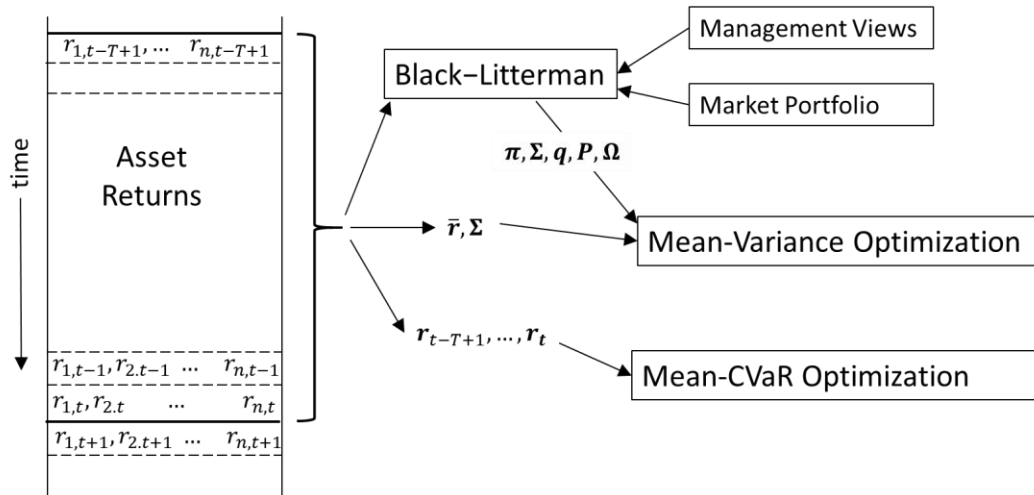


Figure 3.3. Schematic summary of required historical inputs into the optimization procedures.

reason, the historical optimization method is often referred to as a constant-volatility method.

It is critical to recognize that when utilized in practice, the performance of optimization methods based solely on historical data will be optimistic and subject to harsh realities. The most obvious of such realities is that events that take place during a given historical period do not predict (all) future events. Thus, a sample of historical returns may not contain low-probability tail events that may emerge on the next trading day. In addition, each change in daily weights must be accompanied by buy and sell orders to change the portfolio allocation accordingly. Each buy/sell order incurs transaction costs, generally a combination of broker's fees and hard-to-precisely-predict buy/sell prices because bid-ask spreads change relatively continuously over time for liquid assets and can be (relatively) large spreads in the case of low-liquidity assets. Weight changes that require unusually large market orders will therefore require time to process and will potentially occasion market awareness of the size of a position shift.

References

- Attilio, M. (2006). Beyond Black-Litterman in practice: A five-step recipe to input views on non-normal markets. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=872577
- Best, M. & Grauer, R. (1991). On the sensitivity of mean-variance-efficient portfolios to changes in asset means: Some analytical and computational results. *Review of Financial Studies*, 4(2), 315–342.
- Black, F. & Litterman, R. (1991). Asset allocation combining investor views with market equilibrium. *Journal of Fixed Income*, 1(2), 7–18.
- Bloch, M., Guerard, J. B., Markowitz, H., Todd, P. & Xu, G. (2013). A comparison of some aspects of the U.S. and Japanese equity markets. *Japan and the World Economy*, 5, 3–26.
- He, G. & Litterman, R. (1999). The intuition behind Black-Litterman model portfolios. Investment Management Research, Goldman, Sachs & Company.
- Kolm, P. N. & Ritter, G. (2017). On the Bayesian interpretation of Black-Litterman. *European Journal of Operational Research*, 258(2), 564–572.
- Krokhmal, P., Palmquist, J. & Uryasev, S. (2002). Portfolio optimization with conditional value-at-risk objective and constraints. *Journal of Risk*, 4(2), 11–27.
- Lee, W. (2000). *Advanced theory and methodology of tactical asset allocation*. John Wiley & Sons, New York.
- Markowitz, H. (1952). Portfolio selection*. *Journal of Finance*, 7(1), 77–91.
- Michaud, R. (1989). The Markowitz optimization enigma: Is optimized optimal? *Financial Analysts Journal*, 45(1), 31–42.
- Rockafellar, R. T. & Uryasev, S. (2000). Optimization of conditional value-at-risk. *Journal of Risk*, 2(3), 21–41.
- Rockafellar, R. T. & Uryasev, S. (2002). Conditional value-at-risk for general loss distributions. *Journal of Banking & Finance*, 26, 1443–1471.
- Rollinger, T.N. & Hoffman, S. T. (undated). Sortino: A “Sharper” Ratio. Red Rock Capital. http://www.redrockcapital.com/Sortino_A_Sharper_Ratio_Red_Rock_Capital.pdf
- Sharpe, W. F. (1994). The Sharpe ratio. *Journal of Portfolio Management*, 21, 49–58.

3. Modern Portfolio Theory

Tütüncü, R. H., Toh, K. C. & Todd, M. J. (2003). Solving semidefinite-quadratic-linear programs using SDPT3. *Mathematical Programming*, 95(2), 189–217.

Chapter 4

Historical Portfolio Optimization: Domestic REITs

This chapter introduces a suite of optimized REIT-based portfolios to be considered as models for either REIT-based indices or ETFs. They serve as representative prototypes of strategies implemented by institutional investment managers of actively managed portfolios. The different risk–return profiles presented by the prototype portfolios serve as asset-allocation tools for accommodating various market environments and risk tolerances. We use six portfolio optimizations based on sections 3.2.1–3.2.4. Where appropriate, we also consider a traditional, nonoptimized, equal-weighted portfolio for comparison. The equal-weighted-portfolio strategy enjoys widespread use in practice because it does not require information about risk or return and maintains the initial asset diversity of the portfolio. These seven optimized portfolios are referred to throughout the book as follows:

- **MVP** – the Markowitz minimum-variance portfolio of section 3.2.1
- **TVP** – the Markowitz tangent mean-variance portfolio of section 3.2.2
- **M95** – the minimum CVaR₉₅ portfolio of section 3.2.3
- **T95** – the tangent CVaR₉₅ portfolio of section 3.2.4
- **M99** – the minimum CVaR₉₉ portfolio of section 3.2.3
- **T99** – the tangent CVaR₉₉ portfolio of section 3.2.4
- **EQW** – the equal-weighted portfolio

CVaR optimizations at both the 95% and 99% quantile levels are included to examine the effects of the amount of asset-return information available for the tail risk in these portfolios. For example, consider a 2,000-trading-day sample of returns for an asset. Information about the 1% CVaR tail-loss risk (99% quantile level) over this period is derived from only 20 (1%) of these return values, whereas information about the 5% CVaR is derived from a larger set of 100 return values. Thus, the behavior of a portfolio optimized under the CVaR₉₅ risk measure and that of the same portfolio optimized under the CVaR₉₉ risk measure can be significantly different.

This chapter considers portfolios that consist of the domestic REIT assets introduced in section 2.1. (Chapter 5 considers diversification of these portfolios by the addition of international REITs.) Optimal portfolio weights are determined using asset returns developed solely from historical daily data; consequently, the computations implicitly assume constant volatility in each historical window. This assumption will be relaxed in Chapter 7. Constraints are imposed on the optimizations to model controls on asset allocation and transaction costs.

This chapter is structured as follows. The framework developed by Markowitz (1952) and its extensions outlined in Chapter 3 provide the basis for various strategies for obtaining increased returns under specified levels of risk. In section 4.1, we consider the performance of the seven portfolio optimizations in terms of cumulative price and return under a long-only investment strategy, two variations of long–short strategies, and a momentum strategy. With the exception of the momentum strategy, portfolio optimization (rebalancing) is performed daily. Because transaction costs associated with daily rebalancing can be expensive, to provide some control over these costs, in section 4.2 we introduce turnover as a cost proxy and consider performance under

increasing turnover constraint. In section 4.3, we introduce risk measures designed to quantify the relative risks of different portfolios. We conclude the chapter and offer some observations in section 4.4.

4.1 Basic Strategies, Price, and Return Performance

In addition to long-only investment strategies (section 4.1.1), we consider the performance of the optimizations of the portfolio of 26 domestic REITs under a general long–short strategy (section 4.1.2) and under a restricted 130/30-type strategy (section 4.1.3). In section 4.1.4, we consider the consequences of a momentum-type strategy in which rebalancing is performed less frequently.

Weights for the individual REITs in each portfolio are determined based on returns from a rolling window of 2,016 trading days (eight trading years). For example, optimized weights used for portfolio-return computations on 12/20/2007 are determined using data from the *previous* 2,016 trading days, that is, the period 12/14/1999 through 12/19/2007. The time window of 2,016 trading days ensures a sample size large enough to create a meaningful value for the weights. Iterating this recursive scheme for each portfolio generates a daily sequence of optimized portfolio weights covering the period 12/20/2007 through 12/18/2020. Once an entire time sequence of optimal portfolio weights is constructed, performance measures are computed. Note that the methodology for the EQW strategy is vastly simplified. There is no 2,016-day rolling window and no weight optimization. Weights computed for day t are computed on the basis of an equal weighting of prices on day $t - 1$ for the assets in the portfolio.

In performing portfolio optimization, a variety of additional constraints should be added to these basic formulations. The most relevant constraints include available budget, return, holding period, risk factors, transaction size, cardinality (number of assets), volatility, turnover (asset purchase/sale), and tracking error. Here, we include holding period and turnover constraints. Except for limits imposed by turnover constraints, we ignore transaction costs and, in the case of long–short strategies, margin costs and asset availability through brokers.

4.1.1 Long-Only Strategy

In this section, we consider the basic long-only strategy, in which all investment is in the assets and asset weights are restricted to the range $[0,1]$:

$$\sum_{i=1}^n w_i(t) = 1, \quad 0 \leq w_i(t) \leq 1. \quad (4.1)$$

For now, we assume transaction costs are negligible. The performance of the cumulative price of each portfolio from 12/19/2007 through 12/18/2020, assuming a \$100 investment in the portfolio on 12/18/2007, is shown in Fig. 4.1. The cumulative log-return of each portfolio is also plotted. With seven portfolios to consider, the plots shown in Fig. 4.1 become crowded if plots of the performance of all the benchmarks in Fig. 2.3 are also added. We therefore provide a low–high benchmark bracket against which to compare the performance of our portfolios by showing the performance of the better (VNQ) of the two lowest-performing ETFs and the highest-performing index (WD).

The three tangent portfolios, TVP, T95, and T99, outperform all others. Moreover, TVP and T95 outperform T99.⁵² MVP and M95 strongly track each other, and they generally equal or outperform the ETF VNQ. M99 outperforms VNQ. Interestingly, the nonoptimized EQW portfolio performs rather well, most noticeably in the postcrisis period 2010–2012. However, in the long term, it underperforms the tangent portfolios while outperforming the global risk-




Figure 4.1 Cumulative price (left) and log-return (right) for the long-only domestic portfolio optimizations compared to those for the benchmarks.

minimizing portfolios. Note the subtle differences between the performance of the tangent portfolios and the others in response to the 2020 pandemic. Although all the portfolios show a strong decline at the beginning of 2020, the tangent portfolios recover value very rapidly, whereas the others recover at a slower rate. However, the tangent portfolios begin to decline in the last quarter of 2020, a decline that is not evident in the risk-minimizing portfolios and benchmarks.

4.1.2 Jacobs *et al.* Long–Short Strategy

Intuitively, a long–short optimized strategy could improve the expected return of a portfolio for a given volatility, though short selling brings its own set of risks outside of volatility. Jacobs *et al.* (1999) provide an overview of long–short portfolio management and adjust the MVP model to allow for long–short portfolios in which the long and short portfolio optimizations are performed simultaneously. (Optimizing separately leads to a suboptimal portfolio.) We therefore consider indices that use the six MPT-based strategies⁵³ in combination with a Jacobs *et al.* (1999) long–short strategy in which each asset weight is restricted as follows:

$$\sum_{i=1}^n w_i(t) = 1, \quad -s \leq w_i(t) \leq 1 + s. \quad (4.2)$$

Thus, any asset can be shorted up to 100 s % of the total portfolio weight, or a long position can be taken in an asset up to 100 $(1 + s)$ %. Because asset weights are recomputed daily, the long and short positions are rebalanced daily with no consideration given to transaction or margin costs (maintenance margins, interest payments) or the availability of assets (e.g., through brokers).

⁵² That T95 outperforms T99 could be related to the difference in the amount of historical data contributing information to these tail-risk measures. However, see the results for the dynamic simulations in Chapter 6.

⁵³ Because the EQW portfolio is long-only, it is not included in this comparison.

4. Historical Portfolio Optimization – Domestic REITs

Fig. 4.2 displays the results of the cumulative price and log-return achieved by these portfolios under the aggressive choice $s = 0.3$. No comparison to the benchmarks is provided, because the indices do not, and index-following ETFs generally do not, include shorting of assets. Overall, the performance of this long–short strategy is worse than that of the long-only portfolios. Now, after the 2008 market crash, all portfolios have difficulty rebounding. The MVP, M95, and M99 portfolios show “relatively consistent” rebound trajectories, although the return of the M99 portfolio




Figure 4.2 Cumulative price (left) and log-return (right) for the Jacobs et al. long–short domestic portfolio optimizations constrained by (4.2) with $s = 0.3$.

falls much more than that of MVP and M95. The tangent portfolios experience “unstable recovery cycles,” rebounding slowly from lows around 2010 to highs around 2016, but then falling again and recovering to initial investment values only around 2019. Note the difference in behavior between the risk-minimizing and tangent portfolios in response to the 2020 pandemic. The risk-minimizing portfolios show a decline in early 2020, remaining essentially flat for the remainder of the year. In contrast, the tangent portfolios avoid the initial decline and undergo a period of strong growth (doubling or tripling in value). This growth is not sustained: Prices collapse back by 50% (TVP, T95) to 100% (T99), stabilize for a while, and then show a sharp drop at the end of 2020.




Figure 4.3 Cumulative price (left) and log-return (right) for the Jacobs et al. long–short domestic portfolio optimizations constrained by (4.2) with $s = 0.1$.

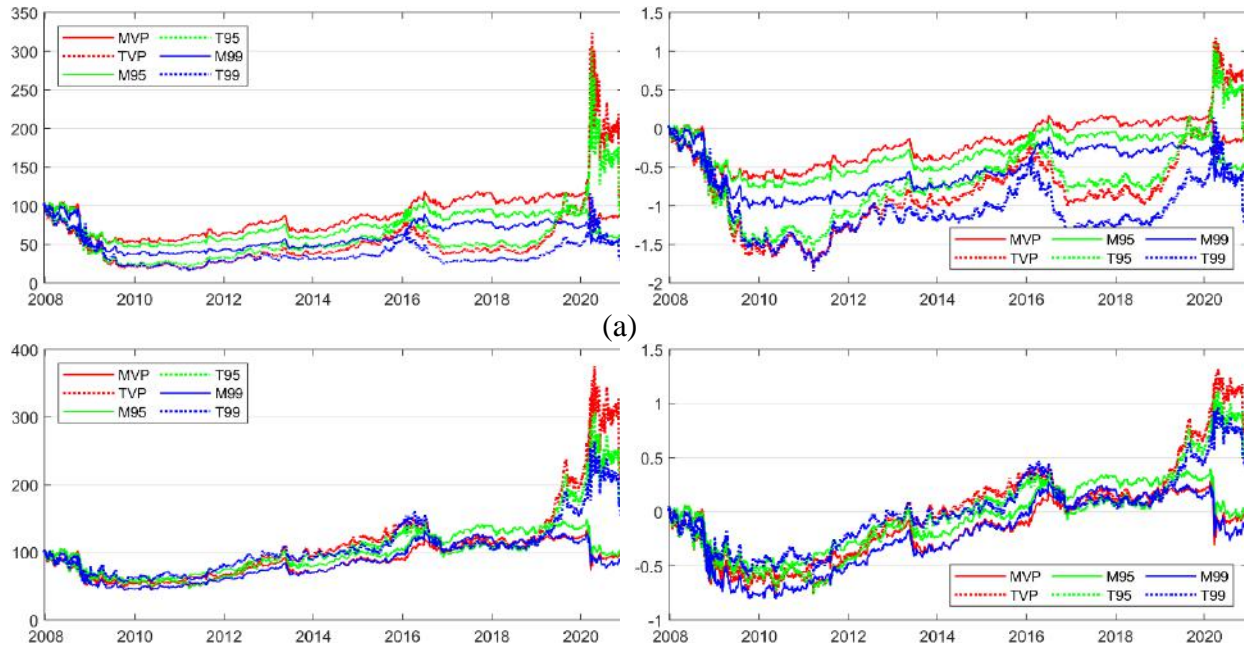
The poor performance of the Jacobs et al. approach clearly indicates that a successful long–short strategy requires careful implementation. Our prototype portfolio, consisting of 26 assets, is small. Allowing *one or more assets* to be shorted up to 30% is quite drastic. Fig. 4.3 displays the

cumulative log-return and price achieved by the optimized portfolios when a value of $s = 0.1$ is used in (4.2). Compared with Fig. 4.2, Fig. 4.3 reveals notable improvements in all the optimizations, particularly the tangent portfolios. It is interesting to contrast the long–short behavior of the tangent portfolios, illustrated in Fig. 4.3, with their long-only behavior, depicted in Fig. 4.1, over the two-year period 2019–2020. In both cases, during 2019, TVP and T95 prices rose by 160% before falling back somewhat. Whereas these long-only portfolios exhibit a sharp price decrease at the beginning of 2020 and subsequent volatile recovery, the long–short portfolios show another 160% price increase followed by a volatile period. In both the long-only and long–short cases, the tangent portfolios show price decreases at the end of 2020, and these decreases are larger for the long–short portfolios.

4.1.3 Lo–Patel Long–Short Strategy

We next consider a 130/30-inspired long–short strategy in which 30% of the starting capital comes from shorting and 130% of the starting capital is allocated to long positions. Lo and Patel (2008) construct such a 130/30 equity portfolio using the S&P 500 universe of stocks and a standard portfolio optimizer. Their strategy, employing monthly rebalancing, yields a benchmark time series of returns that they view as a 130/30 index. A 130/30 strategy is the leverage strategy

$$\begin{aligned} \sum_{i=1}^n w_i(t) &= 1, \quad -lev \leq w_i(t) \leq 1 + lev, \\ \sum_{i=1}^n \max(0, w_i(t) - w_i(t-1)) &\leq 1 + lev, \\ \sum_{i=1}^n \max(0, w_i(t-1) - w_i(t)) &\leq lev, \end{aligned} \quad (4.3)$$



(b)

Figure 4.4 Cumulative price (left) and log-return (right) for the Lo–Patel long–short domestic portfolio optimizations constrained by (4.3) with (a) $lev = 0.30$ and (b) $lev = 0.10$.

where $lev = 0.30$. It differs from the Jacobs et al. strategy in that the sum of the positive weight changes – line 2 in (4.3) – and the sum of the negative weight changes – line 3 in (4.3) – are restricted. Fig. 4.4 compares the cumulative price and log-return plots for optimizations performed under the leverage values $lev = 0.30$ and $lev = 0.10$. A leverage value of 0.30 is too aggressive for this small portfolio. Guerard et al. (2010), among others, have compared the performance of 130/30 models to long-only models under Markowitz mean-variance optimization.

4.1.4 Long–Short Momentum Strategy

The momentum strategy we consider recognizes that investor sentiment favors positively performing assets (Soros, 1987; Tanous, 1997). We implement this via a holding period in which the portfolio weights remain unchanged for a specified period of time. Fig. 4.5 compares the cumulative price and log-returns for the Jacobs et al. long–short portfolios of Fig. 4.3 ($s = 0.1$) with the restriction that rebalancing occurs only every τ trading days, where $\tau \in \{10, 20, 60, 252\}$ days. These plots reveal several interesting features. Significant cumulative price increases occur for the tangent portfolios when τ




Figure 4.5 Cumulative price for the Jacobs et al. long–short domestic portfolio optimizations subject to $s = 0.1 \leq 0.1$ and a holding period (rebalance time) varying from one day to one year.

changes from 1 to 10 days. Further increasing τ to 60 days results in a continued increase in T95 but not in TVP or T99. Changing the rebalance period to 252 days (one year) produces decreases

in TVP and T95 and a significant decrease in T99. In contrast, the minimum-risk portfolios show little response to the change in τ until 252 days, when M99 shows strong improvement.

Fig. 4.6 parallels the plots in Fig. 4.5 for the Lo–Patel strategy with $lev = 0.10$. Again, there is a strong price increase in TVP and T95, whereas the remaining optimizations show little change when τ changes from 1 (Fig. 4.4(b)) to 10. The price performance for all the portfolios remains relatively unchanged when τ changes from 10 to 20. Then, appreciable price improvements appear in all the portfolios when τ changes from 20 to 60. Rebalancing at 252 days produces continued price improvements for the minimum-risk portfolios and for T99 but some decrease for TVP and T95. Determining whether these details are specific to asset type (i.e., REITs) or a more general characteristic of the markets over this time period is beyond the scope of this book. However, what is clear is that a momentum strategy can improve the performance of strongly performing long–short portfolios.

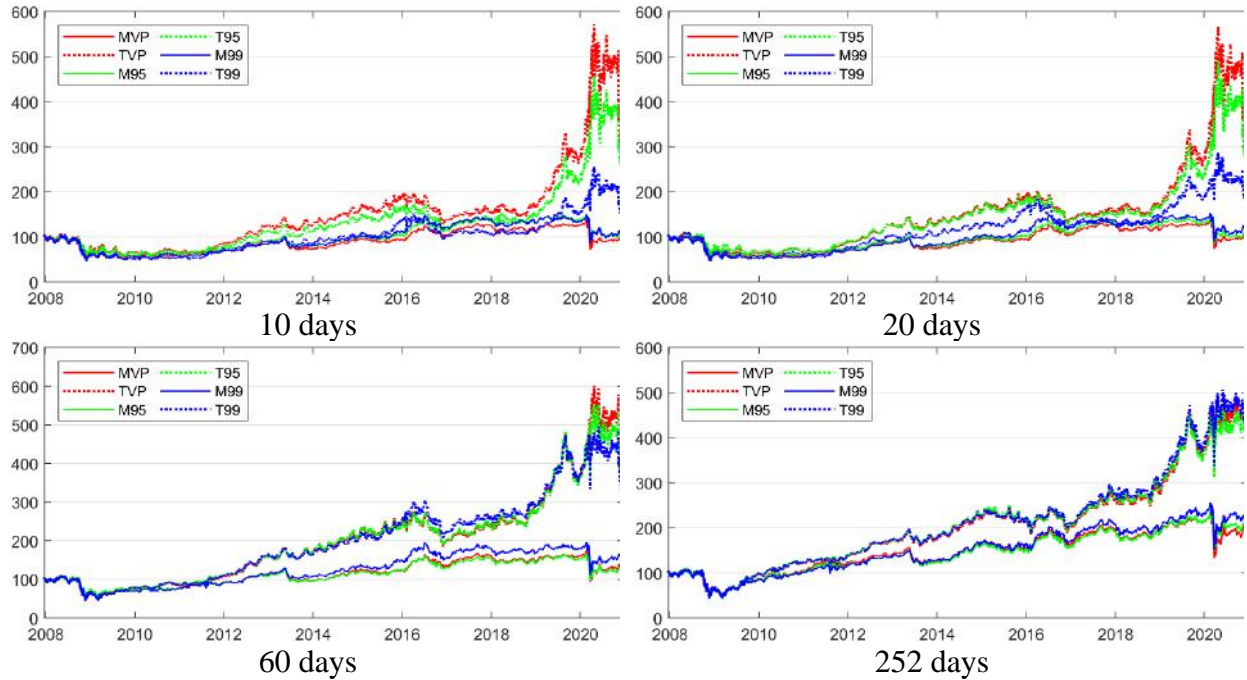


Figure 4.6 Cumulative price for the Lo–Patel long–short domestic portfolio optimizations subject to $lev = 0.1$ and a holding period (rebalance time) varying from one day to one year.

4.2 Performance under Turnover Constraints

Here, we consider the issue of transaction costs induced by the redistribution of asset weights when rebalancing. Because transaction costs are asset specific, we use *turnover* as a quantitative measure of the relative expense of transaction costs of various optimizations. The turnover can be defined in terms of the L_1 -norm of the change in asset weights from day $t - 1$ to t :

$$\|\Delta w_p(t)\|_1 = \sum_{i=1}^n |w_i(t) - w_i(t-1)|. \quad (4.4)$$

The turnover value is defined as 1/2 of the L_1 -norm:

4. Historical Portfolio Optimization – Domestic REITs

$$TO(t) = 0.5 \|\Delta w_p(t)\|_1. \quad (4.5)$$

Fig. 4.7 displays histograms of the turnover values for the six MPT long-only optimizations of Fig. 4.1. The historical period of asset returns used in these portfolio optimizations consists of 3,274 trading days. Thus, each histogram contains 3,273 samples of TO. For MVP, 92% of the turnover values are less than 0.9% (0.009). However, a few values as large as 18% do occur. For TVP, in contrast, turnover values are more strongly distributed up to 10%, with values as large as 21%. Notably, for T95, M99, and T99, turnover values as large as 45% to 60% can occur. Thus, although the tangent portfolios offer the best price-return performance, they come with high turnover values (and hence high transaction costs).




Figure 4.7 Histograms of turnover values (4.5) for the long-only domestic portfolio optimizations of Fig. 4.1.

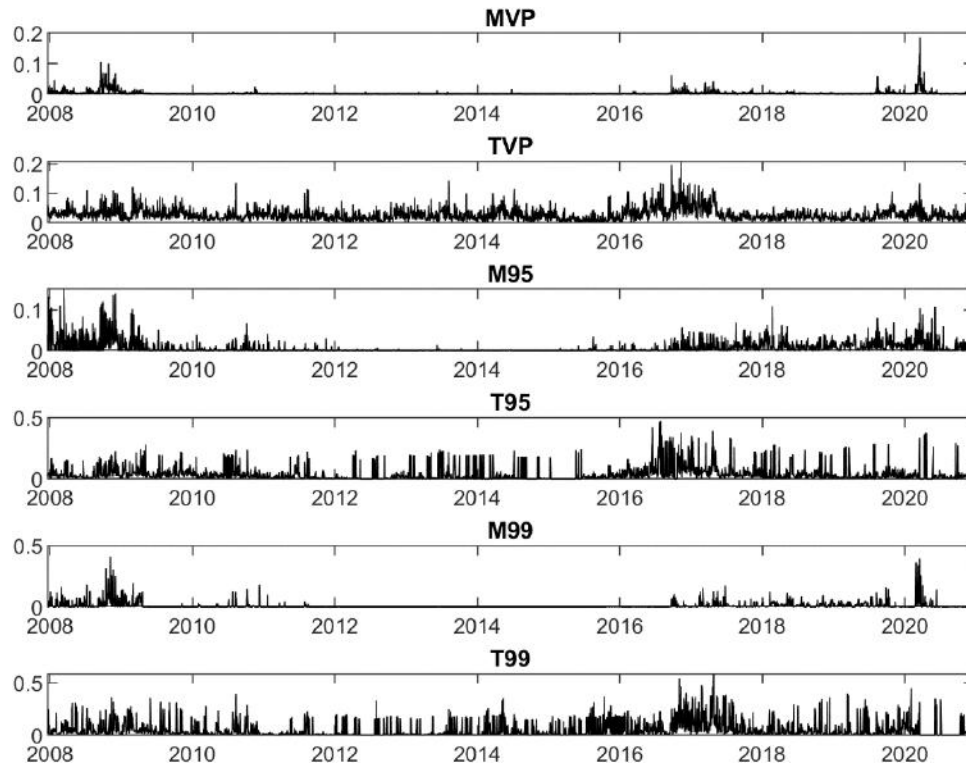


Figure 4.8 Time series of turnover values (4.5) for the long-only domestic portfolio optimizations of Fig. 4.1.

Fig. 4.8 shows the turnover values distributed over time. For the minimum-risk portfolios, the

timing of large turnover values correlates with the 2008–2009 Great Recession and the start of the 2020 pandemic. Large turnover is also evident during the 2017–2020 period. For MVP, which has the smallest turnover profile, this occurs largely during the 2017–2018 period. For M95 and M99, the turnover is stronger, extending essentially from 2017 to 2020. In the tangent portfolios, turnover rates are much more constant over time; the largest turnover values consistently occur in years 2016 and 2017.

The challenge is therefore to maintain return performance while reducing costs. Imposing a turnover constraint, such as

$$TO(t) \leq C_{TO}, \quad (4.6)$$

is one method of achieving this. For example, setting $C_{TO} = 0.04$ restricts the total of the magnitudes of asset-weight changes during a rebalance to values below 0.04 (i.e., 4%). As Figs. 4.7 and 4.8 imply, there are days on which no weight solution that satisfies (4.6) can be found. For those days, we adopt a stepped approach, increasing C_{TO} by the sequence of values $\{C_{TO}, C_{TO} + 0.01, C_{TO} + 0.02, C_{TO} + 0.04, C_{TO} + 0.08\}$ until a weight solution $\mathbf{w}(t)$ is found. If no solution is found by the end of the sequence, then we set $\mathbf{w}(t) = \mathbf{w}(t - 1)$.

Fig. 4.9 plots the cumulative price and log-return performance for these six portfolios run with a 4% turnover constraint. The price performance is relatively unchanged between Figs. 4.9 and 4.1. (In fact, T95 and T99 improve slightly, the latter over the 2012–2016 time period.)




Figure 4.9 Cumulative price (left) and log-return (right) for the long-only domestic portfolios optimized subject to a 4% turnover constraint.




Figure 4.10 Histograms of turnover values (4.5) for the long-only domestic portfolio optimizations subject to a 4% turnover constraint.

Fig. 4.10 plots the turnover-value histograms for these six portfolios run with the 4% turnover constraint. Comparing it to Fig. 4.7 demonstrates that the 4% turnover constraint plays a strong role in restricting transaction costs for TVP, T95, and T99. Fig. 4.11 shows the effect on the turnover time series. MVP, M95, and M99 have constrained turnover during the 2008–2009 and 2016–2020 time periods, and MVP requires little application of the constraint. In contrast, the

4. Historical Portfolio Optimization – Domestic REITs

tangent portfolios are constrained throughout the 2008–2020 time period.




Figure 4.11 Time series of turnover values (4.5) for the long-only domestic portfolio optimizations of Fig. 4.1 subject to a 4% turnover constraint.

Closely akin⁵⁴ to the L_1 -norm (4.4) is the L_2 -norm of the change in asset weights from day $t - 1$ to t :

$$\|\Delta w_p(t)\|_2 = \left(\sum_{i=1}^n |w_i(t) - w_i(t-1)|^2 \right)^{1/2}. \quad (4.7)$$

Rather than imposing a hard turnover constraint such as (4.6), recent scholars (e.g., Yen, 2016; Brodie et al., 2009; DeMiguel et al., 2009) have imposed an L_1 - or L_2 -norm constraint on portfolio-weight changes utilizing penalty terms in the optimization function. This method results in a more stable portfolio by facilitating sparsity (more zero changes) among the portfolio weights (i.e., producing fewer active weights) and by alleviating overfitting (overfitting requires the optimization method to work too hard to find an optimal fit for financial data that clearly has some random noise as a component). In Chapter 11, we demonstrate the use of penalty functions to impose turnover and other constraints. The calculation of these norms sheds light on the sparsity and stability of our portfolios under the various optimization strategies.

Rather than considering the histograms of the L_2 -norm data sets computed from each optimized portfolio time series, we display the TO and L_2 -norm data in box-whisker plots. Fig. 4.12 compares

⁵⁴ Consider the two vertices v_1 and v_2 joined by the hypotenuse of a right triangle. Intuitively, the L_2 -norm (L_2 -distance), $\|v_1 - v_2\|_2$, between v_1 and v_2 is the length of the hypotenuse, whereas the L_1 -norm (L_1 -distance), $\|v_1 - v_2\|_1$, is the sum of the lengths of the other two sides of the triangle. Thus, $\|v_1 - v_2\|_2 \leq \|v_1 - v_2\|_1$.

these box-whisker plots for the six MPT-based long-only optimizations using no turnover constraint and a 4% turnover constraint. Recall that the “box” establishes the interquartile range (IQR), which is the distance between first-quartile, Q_1 , and third-quartile, Q_3 , values of the data distribution. The top whisker represents either (i) the distance from the top of the box to the maximum value, if that distance is less than $1.5 \times \text{IQR}$, or (ii) a distance that is $1.5 \times \text{IQR}$ from the top of the box. The bottom whisker similarly represents either (iii) the distance from the bottom of the box to the minimum value, if that distance is less than $1.5 \times \text{IQR}$, or (iv) a distance that is $1.5 \times \text{IQR}$ from the bottom of the box. Any data point lying outside the range defined by the whiskers is considered an outlier and represented as an “x.” The middle line in each box establishes the median (Q_2) value.⁵⁵




Figure 4.12 Box-whisker summary statistics of TO (4.5) and the L_2 -norm (4.7) for the long-only domestic portfolio optimizations subject to no turnover constraint (left) and a 4% turnover constraint (right).

As the figure illustrates, for the long-only optimization with no turnover constraint, the minimum-risk portfolios, MVP, M95, and M99, have very narrow IQRs and, consequently, very short whiskers. Outliers appear exclusively as high values (which is clearly evident in the skewed histograms of Figs. 4.7 and 4.10). The number of outliers increases from MVP to M95 to M99. The tangent portfolio variant of each optimization technique has a larger IQR and more (as well as more highly valued) outliers. The results for the long-only optimizations with a 4% turnover constraint indicate a corresponding decrease in the maximum value of TO and the L_2 -norm. A comparison of Figs. 4.7 and 4.10 reveals that the MVP and M95 turnover histograms remain relatively unchanged in moving from no turnover constraint to a 4% constraint.⁵⁶ Thus, the box-whisker plots for MVP and M95 are only slightly modified in moving to a 4% turnover constraint. The box-whisker plots for the remaining optimizations exhibit strong changes. As a result of the upper limit imposed by the turnover constraint, outliers are now absent and the median values of the distribution have shifted to larger values.

Fig. 4.13 displays box-whisker plots for TO and the L_2 -norm for the six MPT portfolios optimized under (a) the long-short constraints in equation (4.2) with $s = 0.1$ and $C_{TO} = 0.04$; (b) the constraints in equation (4.3) with $lev = 0.1$; and (c) the constraints in (a) plus rebalancing only every 10 trading days. There is now a difference in behavior between TO and the L_2 -norm. For long-short investing under the constraints in (4.2) with $s = 0.1$ and $C_{TO} = 0.04$, the 4% turnover

⁵⁵ See, for example, https://en.wikipedia.org/wiki/Box_plot.

⁵⁶ After accounting for the differences in bin widths and x-axis scales.

4. Historical Portfolio Optimization – Domestic REITs

constraint places a strong upper bound on the TO distribution for the tangent portfolios, TVP, T95, and T99 – to the extent that the only outliers are low values. For each of the M95 and M99 TO distributions, all the values lie within the box-whisker range. The TO turnover constraint affects the L_2 -norm distributions less strongly.

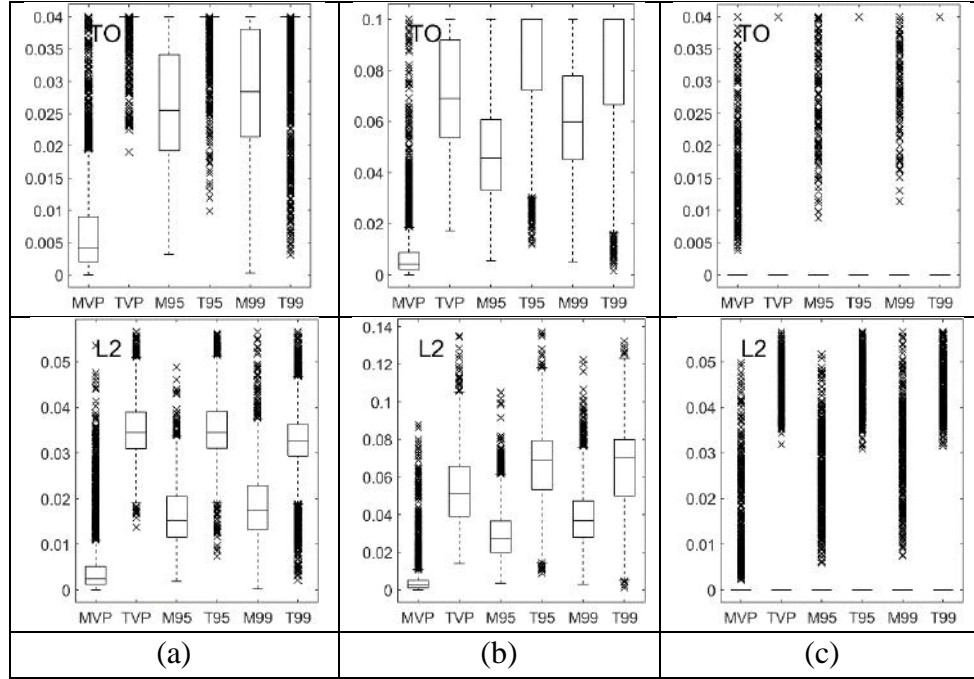


Figure 4.13 Box-whisker summary statistics of TO (4.5) and the L_2 -norm (4.7) for the domestic portfolio optimizations computed using (a) the Jacobs et al. long-short strategy with $s = 0.1$ and $C_{TO} = 0.04$; (b) the Lo-Patel long-short strategy with $lev = 0.1$; and (c) the constrained strategy in (a) with a rebalancing period of 10 trading days.

The constraints in the second and third lines of (4.3) impose TO constraints on two subsets of the assets, those undergoing positive weight change and those undergoing negative weight change from $t - 1$ to t . The asset composition of these two subsets will change over time. Thus, the box-whisker plots in Fig. 4.13(b) are less affected than those in 4.13(a). However, a comparison of the y-axis scales of Figs. 4.13(a) and (b) demonstrates that transaction costs will be larger for long-short portfolios constrained by equation (4.3).⁵⁷

By imposing a restriction on how frequently the portfolio weights are rebalanced, so that $\|\Delta w_p(t)\|_1 = \|\Delta w_p(t)\|_2 = 0$ for the majority of days, Fig. 4.13(c) explicitly indicates how the box-whisker plots are modified. By using a holding period of 10 trading days, the relation $\|\Delta w_p(t)\|_1 = \|\Delta w_p(t)\|_2 = 0$ is guaranteed to hold 90% of the time. For the tangent portfolios, TVP, T95, and T99, TO = 0 for all but a very small set of days. For the minimum-risk portfolios, MVP, M95, and M99, the TO is nonzero for virtually the entirety of the remaining 10% of the

⁵⁷ Transaction costs can of course be controlled by further reducing the value of lev .

time. In contrast, the results in Fig. 4.13(c) show that the L_2 -norm is nonzero for virtually the entirety of the remaining 10% of the time for all the long–short portfolios.




Figure 4.14 Time series of turnover values (4.5) for the Jacobs et al. long–short domestic portfolio optimizations of Fig. 4.13(a).

Fig. 4.14 shows the time series of turnover values for the portfolios in Fig. 4.13(a). As the MVP box-whisker plot in Fig. 4.13(a) reveals, 50% of the turnover values occur in the range $(0.0025, 0.009)$, with extreme values occurring in the range $(0.02, 0.04)$. In contrast, for the tangent portfolios, any turnover value less than 0.04 is an outlier. For M95 and M99, 50% of the turnover values lie between 0.02 and 0.035 and between 0.021 and 0.038, respectively. Except for MVP, no correlation in time between turnover value and any major stock market incident is evident.

Fig. 4.15(a) illustrates the behavior of the cumulative log-return and price for these optimized portfolios constrained by (4.2) with $s = 0.1$ and $C_{TO} = 0.04$. The behavior of the risk-minimizing portfolios (MVP, M95, M99) remains relatively unaffected by these constraint changes. The performance of TVP, T95, and M99 depicted in Fig. 4.15 is enhanced relative to the performance depicted in Fig. 4.4. Fig. 4.15(b) compares the cumulative price and log-return of the portfolios of Fig. 4.15(a) with the restriction that rebalancing occurs only every 10 trading days. Comparing Fig. 4.15 to Fig. 4.4 shows that the performance of MVP and M99 remains unchanged, that the performance of M95 worsens, and that the performance of all the tangent portfolios improves.

4. Historical Portfolio Optimization – Domestic REITs

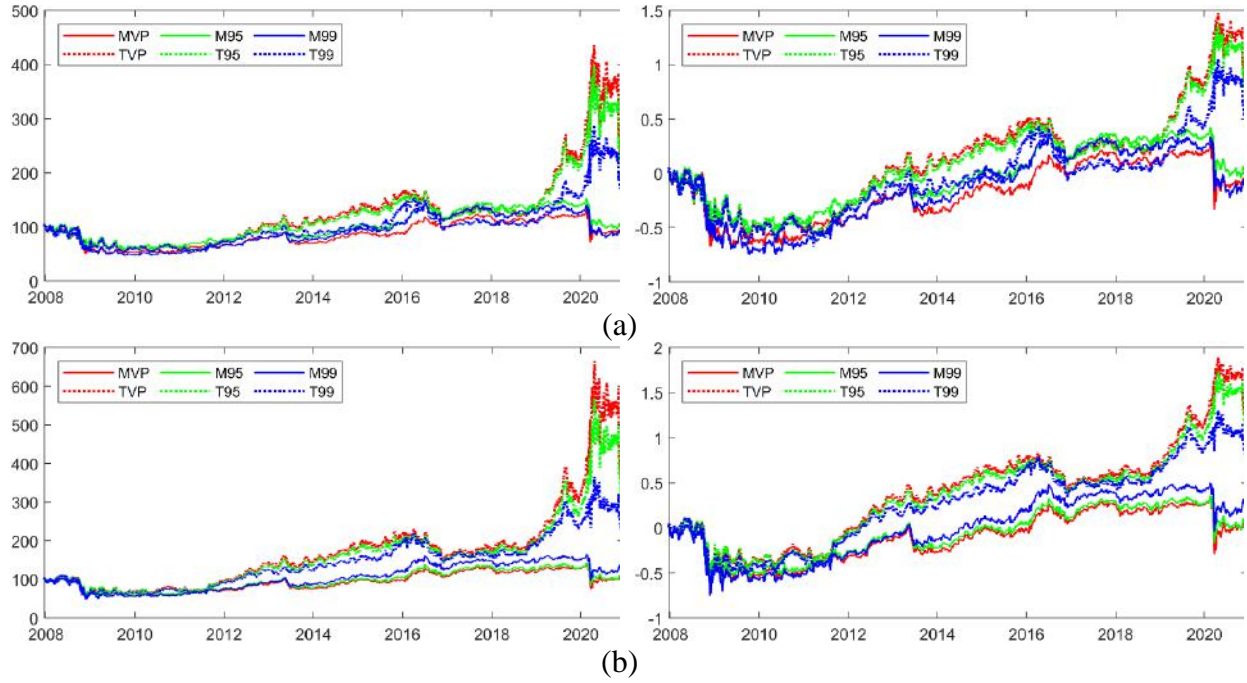


Figure 4.15 Cumulative price (left) and log-return (right) for the Jacobs et al. long-short domestic portfolio optimizations under the constraint (4.2) with $s = 0.1$ and $C_{TO} = 0.04$ with (a) daily rebalancing and (b) a rebalance period of 10 trading days.

4.3 Performance–Risk Measures

In section 4.1, we consider portfolio performance in terms of cumulative price and return. These measures are important, but they must be balanced against competing measures of performance, of which there are a multitude. Cogneau and Hübner (2009) present a classification scheme for 101 performance measures. The scheme consists of four classes, based on performance–risk ratio, incremental return, investor risk preference (via a utility function), and market timing. The largest class of measures is that of performance–risk ratios, which Cogneau and Hübner subdivide into three subclasses, depending on whether the risk component of the measure is absolute, systematic, or nonsystematic (i.e., capable of being eliminated by diversification). Factor models (see section 10.4), such as Jensen’s alpha (Jensen 1968) and the Fama–French three- and five-factor models (Fama and French 1993, 2015), are classified under incremental return. We consider four performance measures from the performance–risk class, as follows.

1. Maximum drawdown (MDD):

$$\text{MDD}(T) = \max_{t \in (0, T)} \left[\max_{s \in (0, t)} (P_p(s) - P_p(t)) \right], \quad (4.8)$$

where $P(S)$ is the cumulative price of the portfolio up through time S and $\text{MDD}(T)$ is the largest peak-to-trough decline⁵⁸ of the portfolio price during the time period $[0, T]$.

⁵⁸ Measured before a new peak is attained.

2. The Sharpe ratio (Sharpe, 1994):

$$SR(T) = \frac{\mathbb{E}[r_p(t) - r_f(t)]_{[0,T]}}{\sqrt{\text{var}[r_p(t) - r_f(t)]_{[0,T]}}} = \frac{\mu_{p[0,T]}}{\sigma_{p[0,T]}}, \quad (4.9)$$

where $r_p(t)$ is the portfolio return, $r_f(t)$ is the risk-free rate for day t , and μ_p and σ_p are, respectively, the expected mean and standard deviation of the portfolio's excess return, $r_p(t) - r_f(t)$.

3. The Sortino–Satchell ratio (Sortino and Satchell, 2001):

$$SS_2(T) = \frac{\mathbb{E} \left[\left(r_p(t) - r_f(t) \right)^+ \right]_{[0,T]}}{\left\| \left(r_f(t) - r_p(t) \right)^+ \right\|_{2[0,T]}}, \quad (4.10)$$

where $y^+ \equiv \max(0, y)$. The Sortino–Satchell ratio is defined for a general p -norm in the denominator; we choose the particular case $p = 2$.

4. The Rachev ratio (Rachev et al., 2008):

$$RR_{\beta,\gamma}(T) = \frac{\text{CVaR}_{\beta} \left(r_f(t) - r_p(t) \right)_{[0,T]}}{\text{CVaR}_{\gamma} \left(r_p(t) - r_f(t) \right)_{[0,T]}}, \quad (4.11)$$

which represents the reward potential for positive returns compared to the risk potential for negative returns at quantile levels defined by the user. In our analysis, we set $\beta = \gamma = 0.95$.

The smaller the value of MDD, the better the performance. For the ratios, the higher the value, the better the performance, with the caveat regarding the Sharpe ratio noted in the next paragraph.

The choice of these performance–risk ratios was influenced by their role as various subclasses in the Cogneau–Hübner classification and by the work of Cherdito and Kromer (2013), who consider the properties of reward–risk measures relative to the desirable qualities of monotonicity, quasi-concavity, scale invariance, and whether the measure is distribution-based. They maintain that every performance measure should be at least monotonic (a measure of “more” is better than a measure of “less”) and quasi-concave (the measure prefers averages to extremes and encourages diversification of risk rather than concentration). The Sortino–Satchell ratio (SS) satisfies all four properties. The Sharpe ratio (SR), which is the most commonly used performance measure, satisfies three of the four properties; it does not guarantee monotonicity, perhaps the most critical property of a risk measure. The Rachev ratio (RR), used by hedge funds that seek excessive returns and insure against big losses, also satisfies three of the four properties; it does not guarantee quasi-concavity.

Because the time period of our data set contains both the Great Recession and the COVID-19 pandemic, computing a single value of a performance measure for the entire time period for each portfolio will reflect performance during only one of those two events. We therefore also compute values for each performance measure for a moving window of length T , giving a time-series view of an evolving measure.

4. Historical Portfolio Optimization – Domestic REITs




Figure 4.16 MDD for the benchmarks using (a) $T = 1$ -year and (b) $T = 4$ -year moving windows. Note that each $MDD(t)$ value represents the largest drawdown seen over the period $[t - T + 1, t]$.

We first consider the MDD of the benchmarks. Fig. 4.16 shows the $T = 1$ and $T = 4$ year MDD plots for the benchmark indices (WD, WP, FT), ETFs (SPY, VNQ), and equal-weighted portfolio (EQW). The time periods covered are 12/17/2008 through 12/18/2020, in Fig. 4.16(a), and 12/16/2011 through 12/18/2020, in Fig. 4.16(b). The one-year plots reveal expectedly high MDD values during the Great Recession (60%–70%) and the pandemic (35%–45%). Between these two periods, yearly MDD values vary from roughly 10% to 25%. The stock market ETF, SPY, typically has the best MDD values. During the Great Recession period, maximum SPY drawdown is 20% better than that of the other benchmarks. Around 2014, from 2017 to 2018, and during the 2020 pandemic, the SPY one-year MDD remains relatively low compared to that of the real estate benchmarks. Only during 2019 does the MDD for SPY start to worsen relative to that for the real estate benchmarks. The real estate benchmarks generally track together. The performance of the EQW portfolio tends to fall between that of SPY and the real estate benchmarks, though with values closer to those of the real estate benchmarks.

When computed over a four-year window, MDD values retain a longer memory of the largest drawdown values. Thus, the 2012 MDD values shown in Fig. 4.16(b) reflect MDDs during the 2008–2009 period. Consequently, the behavior of four-year MDD during the period between the Great Recession and the pandemic reflects less change. The superior behavior of SPY relative to the real estate benchmarks during the period between the recession and pandemic is more apparent in the four-year plot. Because the four-year moving window averages out too much detail, we proceed with further comparison of the performance measures using only the one-year moving window.

Fig. 4.17 presents the MDD, SR, SS, and RR time series for SPY, the real estate index WD, the real estate ETF VNQ, and EQW. The MDD results for these have been presented in Fig. 4.16(a); they are presented again for convenient contrast to the behavior of the performance ratios. Low values of SR and SS during the recession and the pandemic are apparent. Again, SPY generally exhibits the best performance in these two measures, except during the Great Recession. Compared to SR and SS, the recession-related drop in RR is smaller and the pandemic-related drop in RR is of shorter duration. With respect to RR, SPY frequently lags behind or, at best, is on par with the others.




Figure 4.17 MDD, SR, SS, and RR time series (one-year moving window) for SPY, WD, VNQ, and EQW.




Figure 4.18 MDD, SR, SS, and RR time series (one-year moving window) for the long-only domestic portfolio optimizations with $C_{TO} = 0.04$.

Fig. 4.18 presents the performance-measure time series for long-only domestic portfolios with $C_{TO} = 0.04$, whose price performances are shown in Fig. 4.9. The optimized portfolios exhibit qualitative time behavior similar to that of the benchmarks of Fig. 4.17, but the quantitative details differ. For example, during the Great Recession, MDD is smaller (by $\leq 10\%$) for the minimum-risk portfolios (MVP, M95, and M99) than for the real-estate-based benchmarks, though not quite as small as for SPY, whereas the more aggressive tangent portfolios have MDD values on par with those of the EQW portfolio. During the pandemic, MDD values for the tangent portfolios are on par with SPY. The tangent portfolios generally outperform the minimum-risk portfolios in the risk measures. There are time periods of exception, such as in MDD during the Great Recession and in RR from late 2014 to early 2015. The tangent portfolios retain better performance in all the risk measures prior to and after the pandemic.

4. Historical Portfolio Optimization – Domestic REITs




Figure 4.19 Box-whisker summaries of the statistics of the performance measures for the benchmarks identified in section 2.3 and the long-only domestic portfolio optimizations with $C_{TO} = 0.04$.

Although plotting the box-whisker summary statistics results in a loss of temporal detail, it offers a better view of the performance odds for each method. Such plots also allow many more benchmarks and portfolios to be compared at the same time. This is demonstrated in Fig. 4.19, which compares MDD, SR, SS, and RR summary statistics for all seven benchmarks of section 2.3 as well as for the six domestic long-only portfolios with $C_{TO} = 0.04$. A comparison of the IQRs indicates the performance of the benchmark/portfolio over 50% of the historical period, whereas examining the whisker and outlier events gives data on the lower 25% and upper 25% occurrences. For example, with respect to MDD, we see that SPY has the overall best performance and that M95 is a close competitor. With respect to SR, SPY outperforms the other benchmarks but is outperformed by the tangent portfolios. With respect to SS, SPY outperforms the other benchmarks only in the sense that it has fewer low-value outliers, but SPY is again outperformed by the tangent portfolios. With respect to RR, it is less clear that SPY outperforms the benchmarks, and it is outperformed by all the optimized portfolios. In SR, SS, and RR measures, the tangent portfolios are clearly the top performers. The behaviors of the real-estate-based benchmarks and those of EQW are comparable to each other across all the measures.

For comparison with Fig. 4.19, we consider the computation of a single value for each performance metric over the entire 12-year period 12/17/2008 through 12/18/2020. The results are summarized in Fig. 4.20. Over 12 years, M95, MVP, and SPY have the smallest MDD. All the optimized portfolios, as well as EQW, outperform the remainder of the benchmarks. With respect to SR and SS, the tangent portfolios are the top performers, and SPY is the best-performing benchmark. EQW and the real-estate-based benchmarks are the top performers relative to RR, with the tangent portfolios being reasonably competitive. SPY's overall RR value is noticeably the worst.

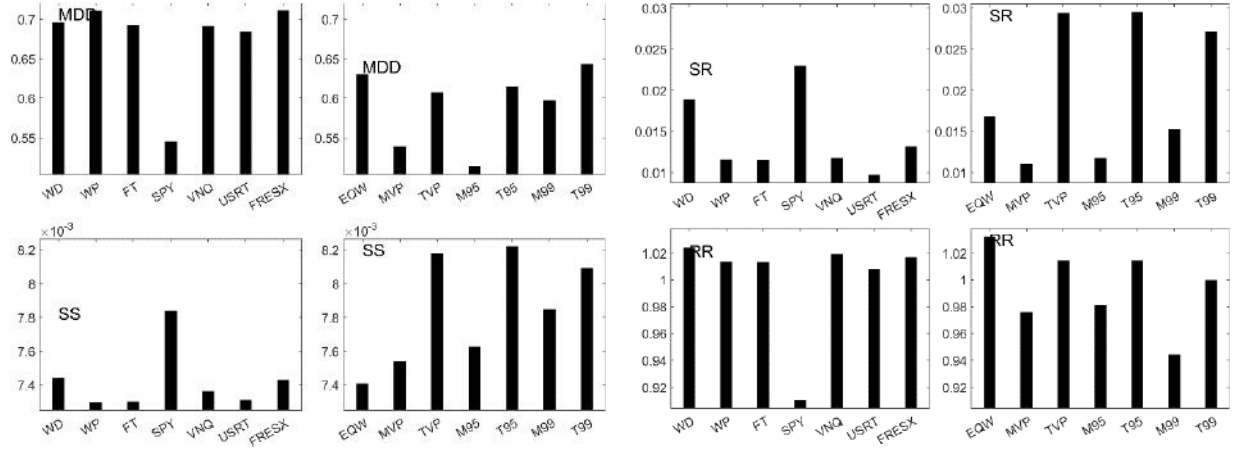


Figure 4.20 Values for each performance measure computed for the time period 12/17/2008 through 12/18/2020 for the benchmarks identified in section 2.3 and the long-only domestic portfolio optimizations with $C_{T0} = 0.04$.

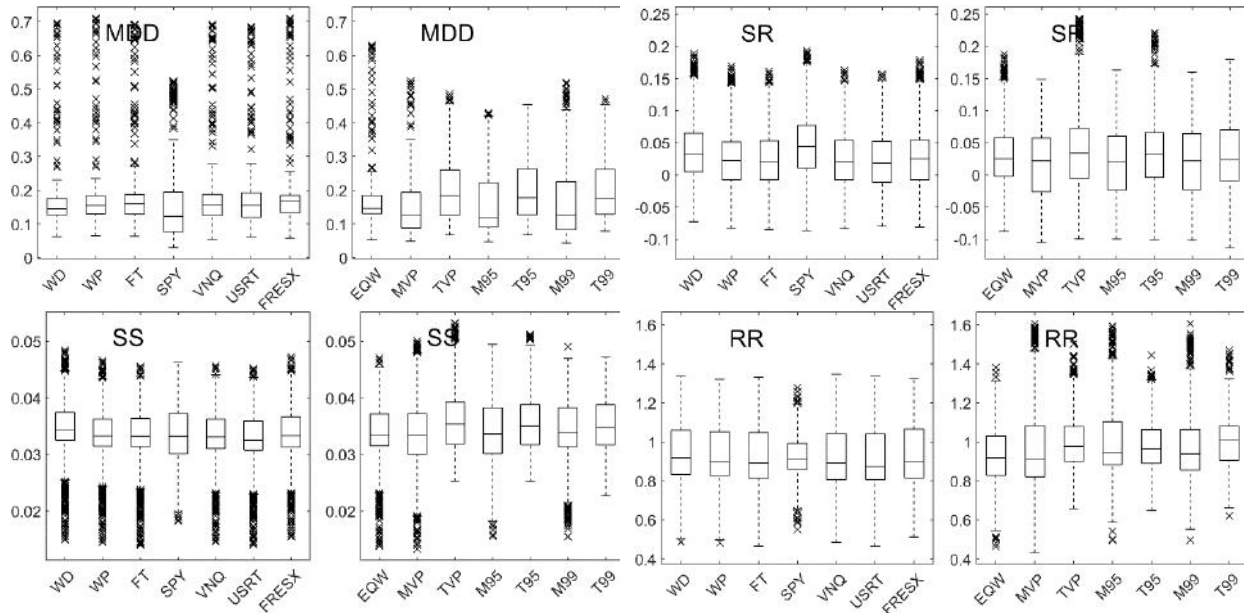


Figure 4.21 Box-whisker summaries of the statistics of the performance measures for the benchmarks identified in section 2.3 and the Jacobs et al. long-short domestic portfolio optimizations constrained by (4.2) with $s = 0.1$ and $C_{T0} = 0.04$.

Using the box-whisker summaries, we compare the performance of the long-short strategies with that of the long-only strategies. Fig. 4.21 presents a summary of the Jacobs et al. long-short domestic portfolio constrained under (4.2) with $s = 0.1$ and $C_{T0} = 0.04$, whose price performance is shown in Fig. 4.4. The performance measures for the benchmarks and EQW portfolios are identical to those illustrated in Fig. 4.19; they are repeated here to facilitate comparison. At first sight, the long-short portfolios have better MDD outlier behavior than the long-only portfolios, but note that the whiskers are longer. In addition, the IQR ranges tend to move up in response to the long-short strategy. The long-short SR values are generally worse than those for the

4. Historical Portfolio Optimization – Domestic REITs

corresponding long-only portfolios. The SS and RR values for the long–short minimum-risk portfolios are slightly better than those for the long-only portfolios.

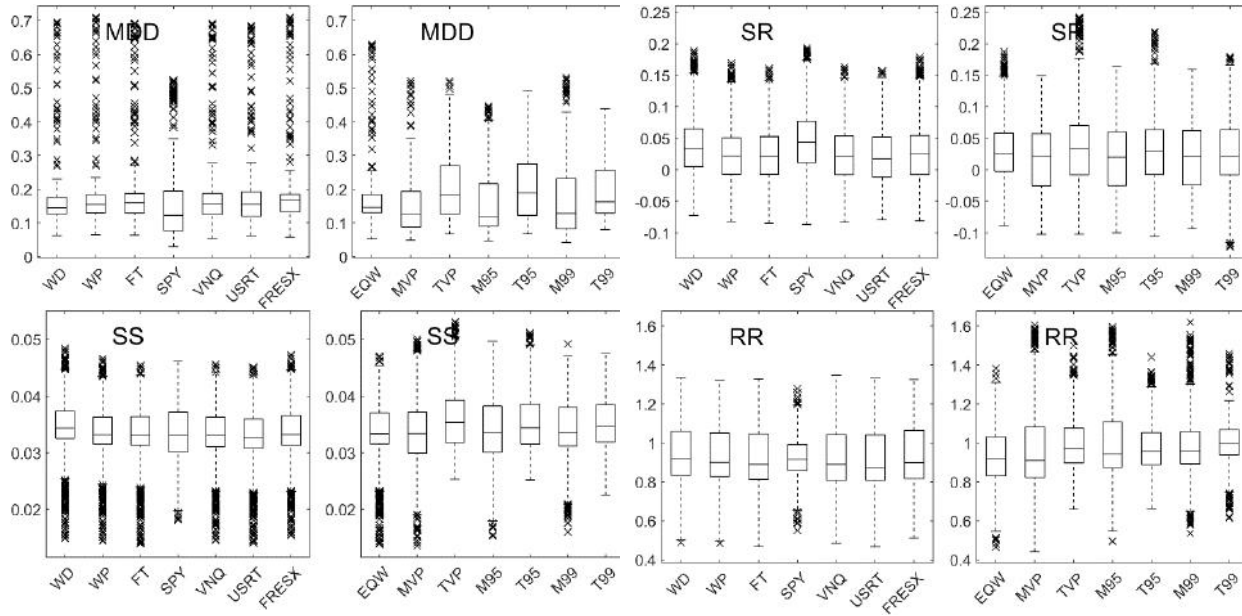


Figure 4.22 Box-whisker summaries of the statistics of the performance measures for the benchmarks identified in section 2.3 and the Lo–Patel long–short domestic portfolio optimizations constrained by (4.3) with $lev = 0.10$.

Fig. 4.22 summarizes the Lo–Patel long–short domestic portfolios constrained by (4.3) with $lev = 0.10$, whose price performance is shown in Fig. 4.4(b). Again, the values for the benchmarks and EQW are repeated to facilitate comparison. For consistency, we compare our results with the long-only portfolios of Fig. 4.19. For the Lo–Patel optimization, the range of extreme MDD outlier values is generally reduced, although the IQR ranges generally expand to larger MDD values. The SR and RR IQR ranges for the tangent portfolios contract to smaller values, while SS values improve slightly for the minimum-risk portfolios.

Without displaying the graphs, we note that for the same portfolio, changing from daily rebalancing to rebalancing every two weeks does not noticeably change the box-whisker summaries of these four performance measures. Fig. 4.23 illustrates the impact of changing the turnover constraint. We compare the performance-measure box-whisker summaries for the long-only portfolio with no turnover constraint and a turnover constraint of $C_{TO} = 0.004$ (which corresponds to, at most, a 100% change in the portfolio weight composition each year). With the exception of TVP, the increased turnover constraint serves to even the performance among all the optimizations. TVP tends to retain the best performance measures under increasing turnover constraint.

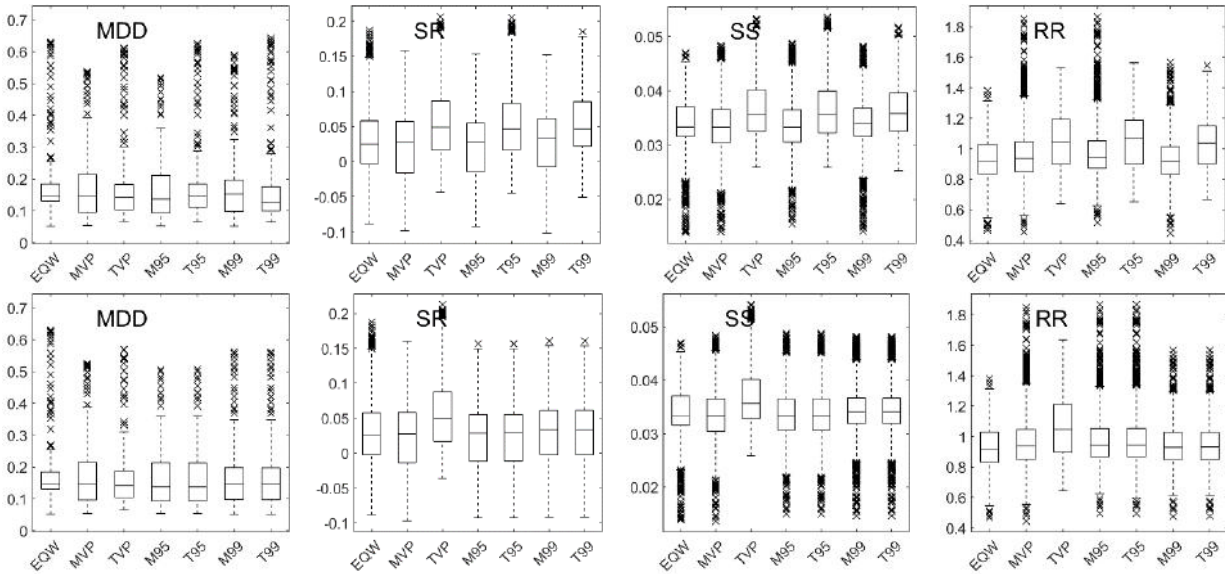


Figure 4.23 Box-whisker summaries of the performance measures for the long-only domestic portfolio optimizations subject to no turnover constraint (top) and a 4% turnover constraint (bottom).

4.4 Observations

The prototype ETFs developed in this chapter are preliminary in significant ways. They have not been subjected to market forces. They have been conditioned upon a narrow subset of institutionally held REITs. Although we have included holding period and turnover constraints, we have not included other significant investing constraints: available budget, target return, risk factors, transaction size, cardinality (number of assets), volatility, tracking error, transaction costs, and in the case of short strategies, margin costs and asset availability through brokers.

Given these critical caveats, the collective evidence, from the long-only portfolios and from the Jacobs et al. and Lo–Patel variations that allow for negative asset-allocation weights (short-selling), shows encouraging performance, in terms of returns and risk measures, relative to common industry benchmarks for the overall stock market (S&P 500) and the REIT market. In particular, the asset allocations for the top-performing indices that maximize the SR (TVP, T95, and T99) not only outperform the benchmarks but also reflect recent market-cycle trends in various property sectors over the past few years, including

- outperformance by select infrastructure and specialty REITs (AMT and SBA);
- higher uncertainty (volatility) for healthcare and retail REITs (HCP and VTR) relative to other property sectors due to disruptions in insurance markets; and
- a significant paradigm shift in retail-anchored, big-box, and mall REIT sectors (REG, MAC, O, and SPG) and in the industrial REIT sector (PLD and DRE) due to the rapid growth of the online retail market (Amazon).

Advanced techniques are required to form a better understanding of the dynamic causal relationships between complex underlying market forces and the aforementioned REIT market

trends and their manifestation as optimal portfolio weight allocations and subsequent returns. These techniques are discussed in subsequent chapters.

References

- Brodie, J., Daubechies, I., De Mol, C., Giannone, D. & Loris, I. (2009). Sparse and stable Markowitz portfolios. *Proceedings of the National Academy of Sciences*, 106(30), 12267–12272.
- Cheredito, P. & Kromer, E. (2013). Reward–risk ratios. *Journal of Investment Strategies*, 3, 1–16.
- Cogneau, P. & Hübner, G. (2009). The 101 ways to measure portfolio performance. *SSRN Electronic Journal*. DOI: 10.2139/ssrn.1326076
- DeMiguel, V., Garlappi, L., Nogales, F. J. & Uppal, R. (2009). A generalized approach to portfolio optimization: Improving performance by constraining portfolio norms. *Management Science*, 55(5), 798–812.
- Fama, E. F. & French, K. R. (1993). Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics*, 33, 3–56.
- Fama, E. F. & French, K. R. (2015). A five-factor asset pricing model. *Journal of Financial Economics*, 116, 1–22.
- Guerard, J. B., Chettiappan, S. & Xu, G. (2010). Stock-selection modeling and data mining corrections: Long-only versus 130/30 models. In J. B. Guerard (Ed.), *Handbook of portfolio construction*, pp. 621–648. Springer, Boston.
- Jacobs, B. I., Levy, K. N. & Starer, D. (1999). Long-short portfolio management. *Journal of Portfolio Management*, 25(2), 23–32.
- Jensen, M. C. (1968). The performance of mutual funds in the period 1945–64, *Journal of Finance*, 23(2), 389–416.
- Lo, A. W. & Patel, P. N. (2008). 130/30: The new long-only. *Journal of Portfolio Management*, 34(2), 12–38.
- Markowitz, H. (1952). Portfolio selection*. *Journal of Finance*, 7(1), 77–91.
- Rachev, S. T., Stoyanov, S. & Fabozzi, F. J. (2008). *Advanced stochastic models, risk assessment, and portfolio optimization*. John Wiley & Sons, Hoboken, NJ.
- Sharpe, W. F. (1994). The Sharpe ratio. *Journal of Portfolio Management*, 21, 49–58.
- Soros, G. (1987). *The alchemy of finance*. Simon & Schuster, New York.
- Sortino, F. A. & Satchell, S. (2001). *Managing downside risk in financial markets*. Butterworth-Heinemann, Oxford.
- Tanous, P. J. (1997). *Investment gurus*. New York Institute of Finance, New York.
- Yen, Y.-M. (2016). Sparse weighted-norm minimum variance portfolios. *Review of Finance*, 20(3), 1259–1287.

Chapter 5

Diversification with International REITs

Domestic investors can benefit from the addition of foreign securities to their investment portfolios when such an addition is beneficial from a risk-diversification standpoint. Eun and Resnick (1994) and Solnik (1995) find that risk reduction is possible when foreign securities are added to a diversified U.S. equity portfolio. Li et al. (2003) confirm that diversification benefits exist even in the presence of short-selling constraints. Eun et al. (2010) show that an augmented optimal portfolio that includes local-factor funds substantially outperforms a benchmark optimal portfolio that consists of only country-market indices.

We explore the impact of diversification by adding the international REITs described in section 2.2 into the domestic portfolio, producing a “global” portfolio. In section 5.1, we consider the performance of the international REITs alone, in their own portfolio (the “international” portfolio). In section 5.2, we consider the performance of the global portfolio. We explore performance in the context of the strategies covered in section 4.1. Because the available-price data period is shorter for the international REITs, in computing optimized weights for the international and global portfolios, we employ a rolling window of 1,008 trading days. Therefore, optimized portfolios are computed for the time period 04/30/2018 through 12/18/2020. A further consequence of the shortened period is that risk-measure computations in this chapter are performed using a 63-day (one financial quarter) rolling window. To compare the performance of the international and global portfolios with that of the domestic portfolio, we rerun the domestic portfolio optimizations of section 4.1 for this restricted time period. The performance of the rerun domestic portfolios is presented in section 5.2 and compared with that of the global portfolios.

5.1 International Portfolio Performance

5.1.1 Long-Only International Portfolios

We first consider the turnover data for the long-only international portfolio optimizations run with no constraint on daily turnover. Fig. 5.1(a) shows the TO time series for each optimized portfolio, and Fig 5.1(b) provides the box-whisker summary of the TO and L_2 -norm distributions. Turnover is smallest for MVP, for which the largest turnover is correlated with the pandemic onset. For M95, TO history is similar, though turnover is slightly larger. For M99, the scale of turnover outliers roughly doubles (compared to MVP and M95), though the IQR for M99 is less than that for M95. It is striking how relatively small M99 turnover is for the final three quarters of 2020. For the tangent portfolios, turnover rates are larger than those for the corresponding risk-minimizing portfolio and exhibit a much weaker correlation with the onset of the pandemic. Although T99 has larger outlier turnover values than T95, its IQR is virtually identical to that of T95.

This data shows that the imposition of a turnover constraint $C_{TO} = 0.04$ will affect only large outlier values for the minimum-risk portfolios and some of the upper-whisker range and outlier values for the tangent portfolios. Consequently (plots not shown), there are no significant differences between the performance of the long-only international portfolios optimized with no

5. Historic Portfolio Optimization – Diversification with International REITs

turnover constraint and those with $C_{TO} = 0.04$. Thus, we consider only the long-only optimizations with a 4% daily turnover constraint. Fig. 5.2 shows the box-whisker TO and L_2 -norm summary statistics when the optimizations are run with $C_{TO} = 0.04$. Although the y-axis scales are different from those in Fig. 5.1(b), close examination confirms that only outlier and upper-whisker values are affected.

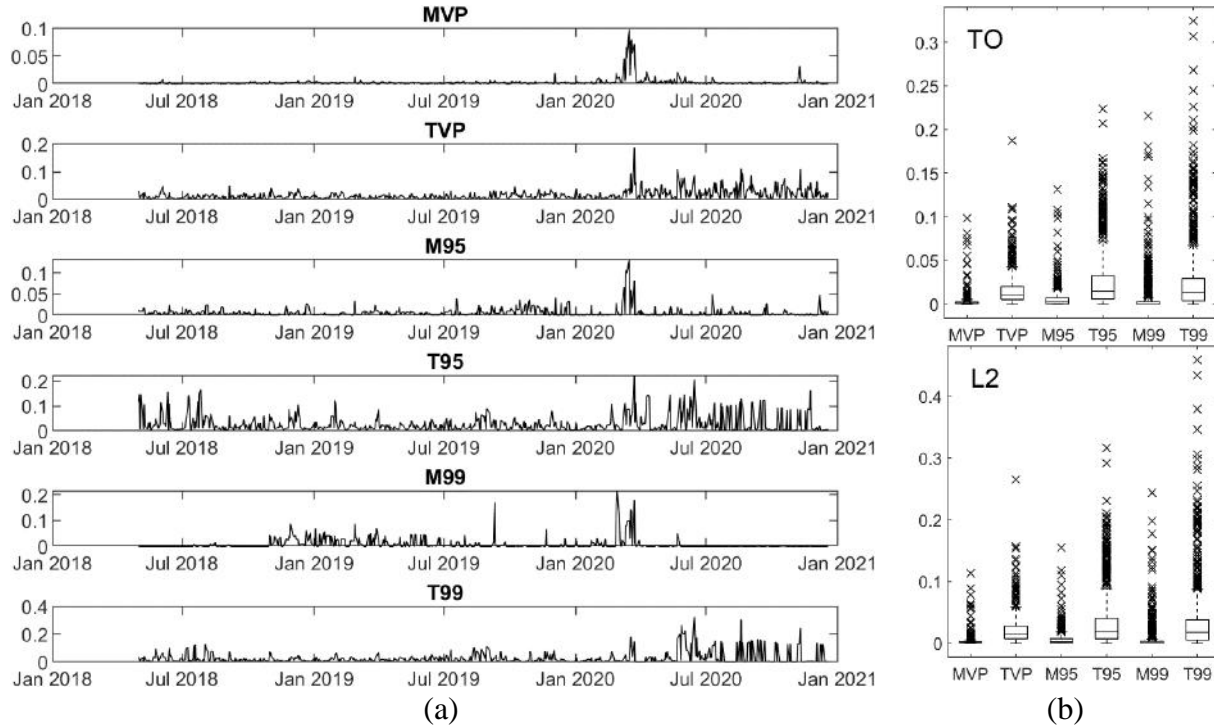


Figure 5.1 (a) Time series of turnover values (4.5) and (b) box-whisker summary statistics of TO and the L_2 -norm for the long-only international portfolio optimizations subject to no turnover constraint.

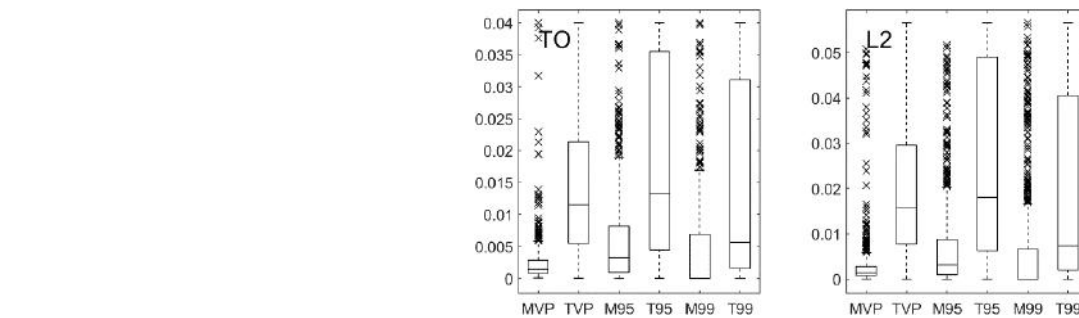


Figure 5.2 Box-whisker summary statistics of TO and the L_2 -norm for the long-only international portfolio optimizations subject to a 4% turnover constraint.

Fig. 5.3 displays the cumulative price and log-return performance for these portfolios. The downturn at the start of the pandemic divides the performance of these portfolios. Prior to the pandemic, all the optimized portfolios generally move in concert with the benchmark ETF VNQ. In addition, the tangent portfolios equal or outperform WD, the (higher-performing) benchmark

index. Interestingly, the equal-weighted portfolio of international REITs underperforms all others. This contrasts strongly with the performance of the equal-weighted portfolio of domestic REITs, which outperforms the minimum-risk domestic portfolio optimizations (see Fig. 4.9). Subsequent to the pandemic downturn, the benchmarks are the better price/return performers.




Figure 5.3 Cumulative price (left) and log-return (right) for the long-only international portfolio optimizations subject to a 4% turnover constraint compared to those for the benchmarks.

Fig. 5.4 presents the MDD, SR, SS, and RR time series for the benchmarks WD and VNQ and for the equal-weighted international portfolio EQW. Because these are calculated over a shorter moving window, they provide a more detailed picture of the changes of these risk measures over the time period 07/30/2018 through 12/18/2020 than do the series shown in Fig. 4.17. Losing 50% of its value, the EQW portfolio has the largest drawdown as a result of the pandemic and displays the largest decrease in SR, SS, and RR values at the onset of the pandemic. EQW exhibits two periods of very strong RR performance. All the risk-measure profiles for the benchmarks WD and VNQ are very similar; they incurred ~42% drawdown as a result of the pandemic. Both benchmarks show a consistent spike in SR, SS, and RR values around the end of the first quarter of 2019.




Figure 5.4 MDD, SR, SS, and RR time series (quarter-year moving window) for WD, VNQ, and EQW.

Fig. 5.5 presents the risk-measure time series for the long-only optimized international portfolios. Compared to the benchmarks and minimum-risk optimizations, the tangent optimizations suffer smaller MDD values as a result of the pandemic. Generally, the tangent portfolios outperform the minimum-risk portfolios in SR, SS, and RR prior to the start of the

5. Historic Portfolio Optimization – Diversification with International REITs

pandemic, and this situation is reversed during the pandemic. For comparison with the performance of the benchmarks, we consider the box-whisker summaries in Fig. 5.6. Among the benchmarks and the optimized portfolios, the best MDD statistics belong to SPY and M95, respectively. SPY has the best SR and SS statistics, and all the optimized portfolios are very comparable with each other in terms of these two variables. However, with respect to RR, the tangent portfolios perform the best.




Figure 5.5 MDD, SR, SS, and RR time series (quarter-year moving window) for the long-only international portfolio optimizations.

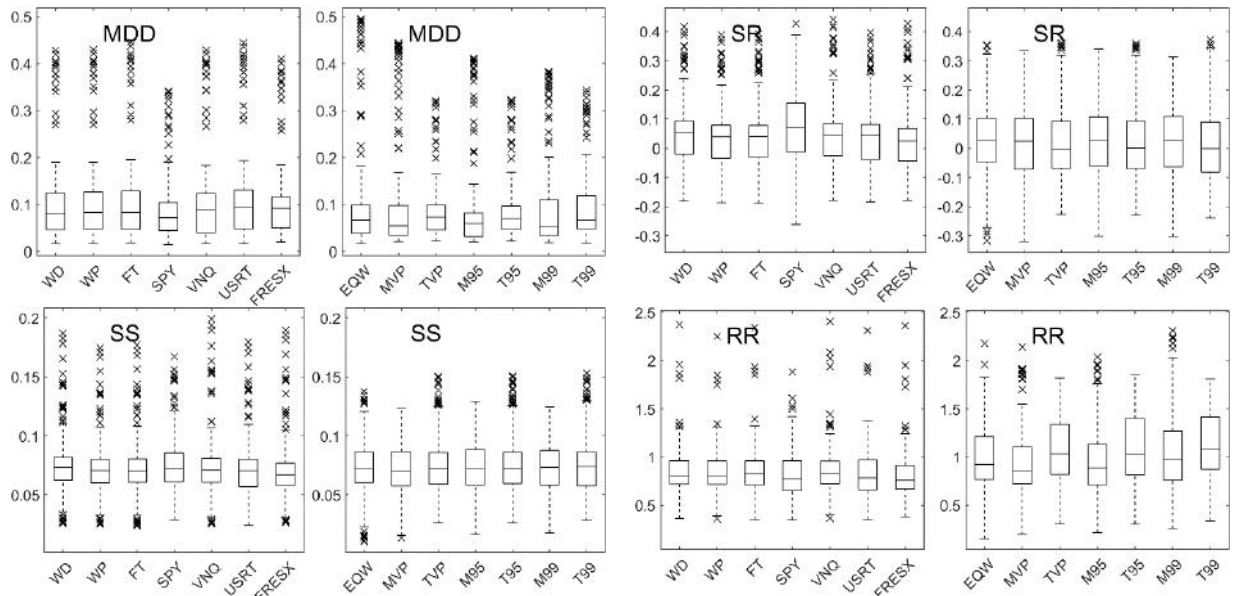


Figure 5.6 Box-whisker summaries of the statistics of the performance measures for the benchmarks identified in section 2.3 and the long-only international portfolio optimizations subject to a 4% turnover constraint.

Fig. 5.7 depicts the overall risk measures computed for the period 07/30/2018 through 12/18/2020. In terms of overall values, SPY, TVP, and T95 have the smallest drawdown, whereas SPY and the tangent portfolios have the best SS values. SPY has the best overall SR. Overall SR values are negative for the minimum-risk portfolios; only TVP and T95 have SR values

competitive with those of the benchmarks. EQW has the best RR, and only M99 and T99 have overall RR values strongly competitive with those of the benchmarks.

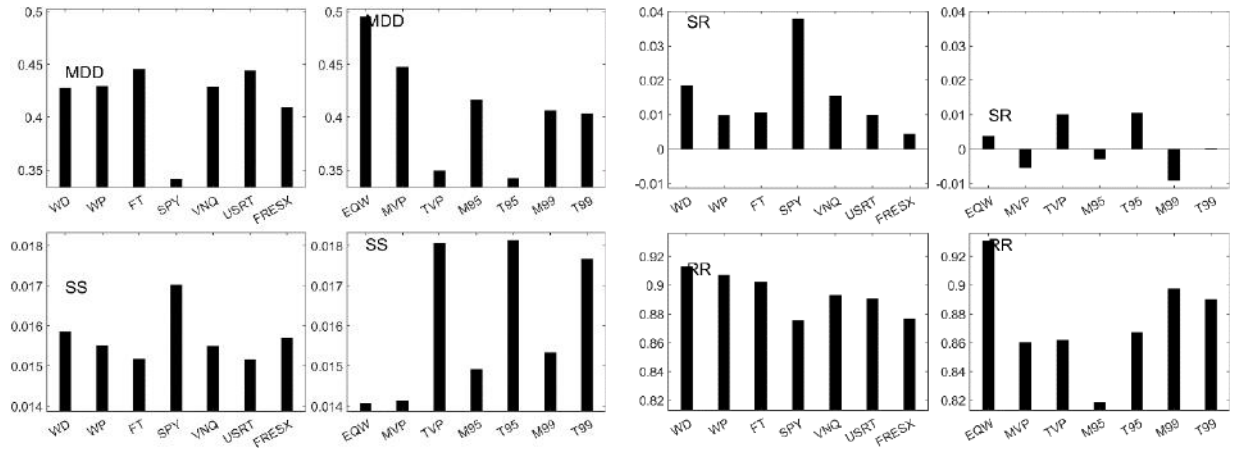


Figure 5.7 Performance measure values computed for the time period 07/30/2018 through 12/18/2020 for the benchmarks identified in section 2.3 and the long-only international portfolio optimizations subject to a 4% turnover constraint.

5.1.2 Jacobs et al. Long–Short International Portfolios

To illustrate the performance of the international portfolio under Jacobs et al. long–short investing, we consider optimization with $s = 0.1$ in (4.2) and a turnover constraint of $C_{TO} = 0.04$. Fig. 5.8 reveals that unlike long-only optimization (Fig. 5.3), these long–short optimizations exhibit a very strong separation of the cumulative price and log-return performance between the tangent and minimum-risk portfolios. In terms of MDD, the tangent portfolios are less impacted by the pandemic.




Figure 5.8 Cumulative price (left) and log-return (right) for the Jacobs et al. long–short international portfolio optimizations constrained by (4.2) with $s = 0.1$ and $C_{TO} = 0.04$.

Fig. 5.9 summarizes the TO and L_2 -norm statistics for these long–short portfolios. With the exception of MVP and T95, the IQR ranges are wider than for the long-only case (Fig. 5.2). The risk-measure time series are presented in Fig. 5.10; their box-whisker summaries are shown in Fig. 5.11; and the overall time-period risk measures are illustrated in Fig. 5.12. MDD at the start of the pandemic is better for the long–short optimizations, particularly the tangent optimizations, than

5. Historic Portfolio Optimization – Diversification with International REITs

for the long-only optimizations (Fig. 5.5). Differences between the long–short time series and box-whisker plots and those between the long-only time series and box-whisker plots (Figs. 5.5 and 5.6) are generally subtle. Clear differences are revealed by the overall time period values. Overall SRs (Fig. 11) are better for the long–short tangent portfolios than for the long-only ones (Fig. 5.7), but SR values for the minimum-risk portfolios, particularly M95, are worse. Overall, SS and RR values tend to be better for the long–short portfolios than for the corresponding long-only ones.




Figure 5.9 Box-whisker summary statistics of TO and the L_2 -norm for the Jacobs et al. long–short international portfolio optimizations constrained by (4.2) with $s = 0.1$ and $C_{T0} = 0.04$.


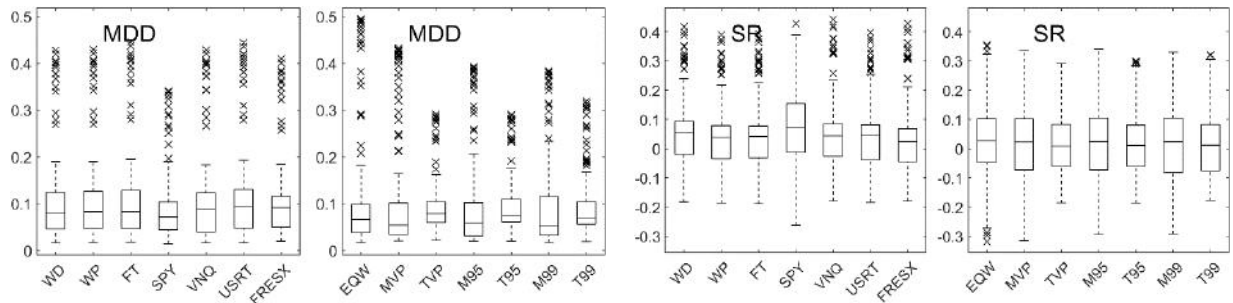


Figure 5.10 Risk-measure time series for the Jacobs et al. long–short international portfolio optimizations constrained by (4.2) with $s = 0.1$ and $C_{T0} = 0.04$.



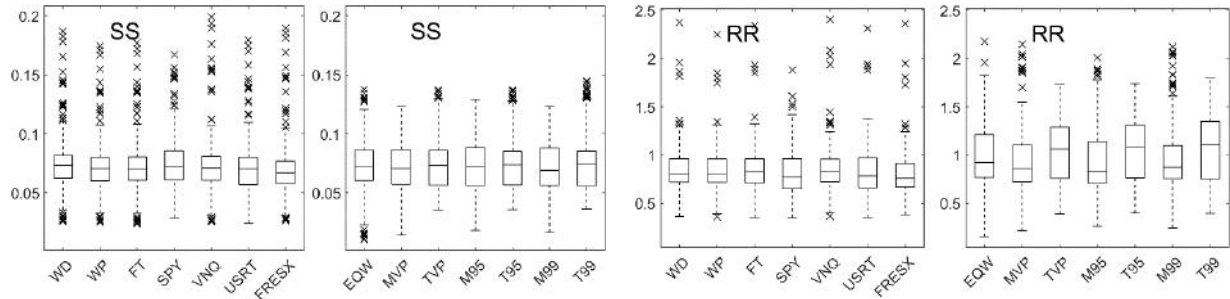


Figure 5.11 Box-whisker summaries of the statistics of the performance measures for the benchmarks identified in section 2.3 and the Jacobs et al. long-short international portfolio optimizations constrained by (4.2) with $s = 0.1$ and $C_{T0} = 0.04$.

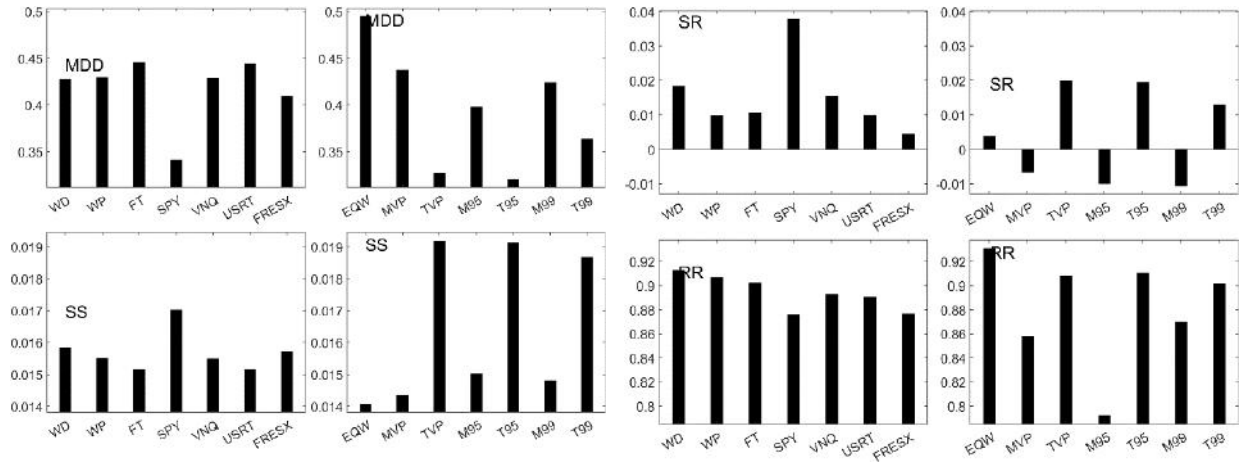


Figure 5.12 Performance measure values computed for the time period 07/30/2018 through 12/18/2020 for the benchmarks identified in section 2.3 and the Jacobs et al. long-short international portfolio optimizations constrained by (4.2) with $s = 0.1$ and $C_{T0} = 0.04$.

5.1.3 Lo–Patel Long–Short International Portfolios

Comparing Fig. 5.13 with Fig. 5.8 shows that Lo–Patel long-short optimization using $lev = 0.1$ produces results very similar to those shown in section 5.1.2. (The tangent portfolio results for Lo–Patel are slightly better.) Thus, graphs of TO and the risk-measure values for these optimized portfolios are not shown.




Figure 5.13 Cumulative price (left) and log-return (right) for the Lo–Patel long-short international

5. Historic Portfolio Optimization – Diversification with International REITs

portfolio optimizations constrained by (4.2) with $lev = 0.1$ and $C_{TO} = 0.04$.

5.2 Global Portfolio Performance

The global portfolio consists of the 26 domestic REITs described in section 2.1 and the seven international REITs of section 2.2. As mentioned, the international REIT data is uniformly available over the time period 04/13/2014 through 12/18/2020, so we rerun computations on the domestic portfolio for this shorter time period to enable comparison of the effects of this diversification.

5.2.1 Long-Only Global Portfolios

Fig. 5.14 compares the cumulative price and log-return of the global and domestic portfolios under long-only optimization with a daily turnover constraint of 4%. The global portfolio exhibits better results prior to the onset of the pandemic. However, the domestic tangent portfolios TVP, T95, and T99 suffer less drawdown during the pandemic and outperform the global tangent portfolios after the pandemic. During the pandemic, the behavior of the minimum-risk global and domestic portfolios are roughly comparable, with evidence that the domestic portfolios are recovering better by January 2020.




Figure 5.14 Cumulative price (left) and log-return (right) for the long-only (a) global and (b) domestic portfolio optimizations subject to a 4% turnover constraint.




Figure 5.15 Box-whisker summary statistics of TO values for the long-only global (left) and domestic (right) portfolio optimizations subject to a 4% turnover constraint.

Fig. 5.15 compares the turnover statistics for the global and domestic portfolios over this time period. The domestic portfolios have larger IQRs, with generally higher median values (implying generally greater fee expenses).

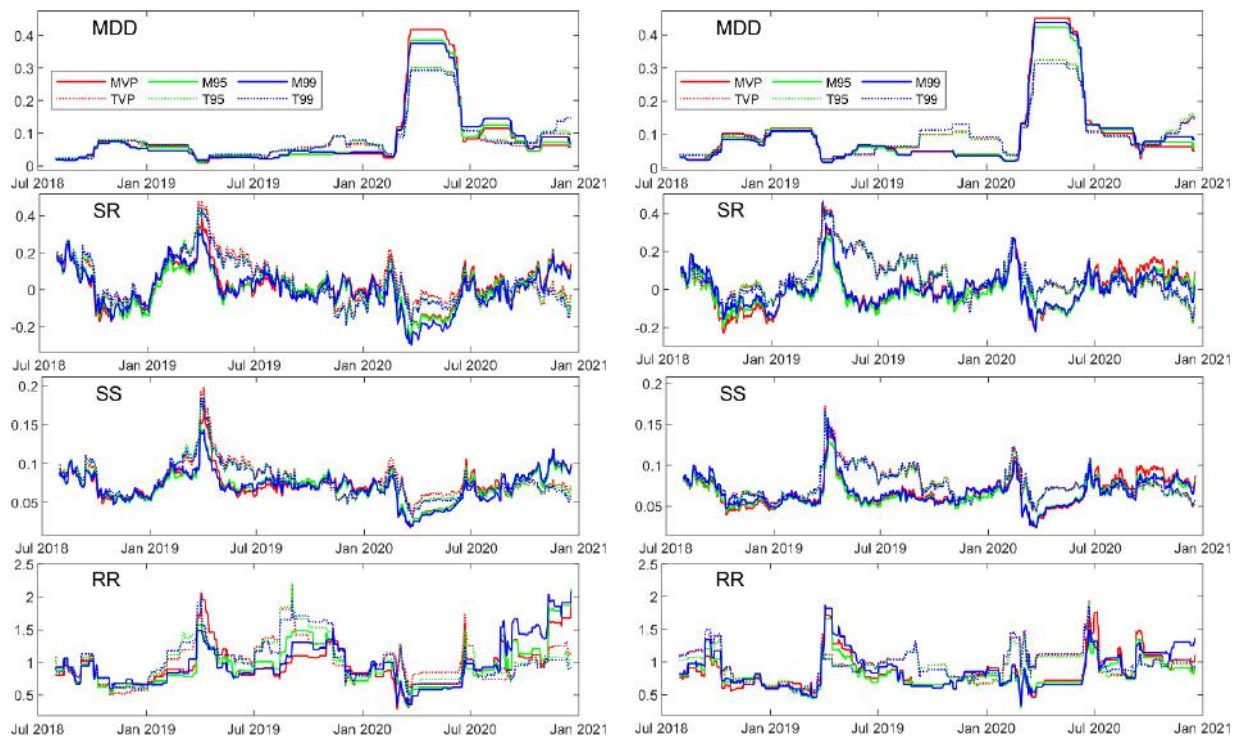


Figure 5.16 Risk-measure time series for the long-only global (left) and domestic (right) portfolio optimizations subject to a 4% turnover constraint.

Fig. 5.16 compares the risk-measure time series for the global and domestic portfolios over the time period 07/30/2018 through 12/18/2020. The global portfolios generally have better MDD values. At any fixed time point, the SR, SS, and RR values for the domestic tangent portfolios are closer together than those for the global portfolios. Figs. 5.17–5.19 provide a more extensive comparison by presenting box-whisker summary statistics for the risk-measure values. Fig 5.17

5. Historic Portfolio Optimization – Diversification with International REITs

shows the MDD statistics for the benchmarks and the global and domestic portfolios. Note that the EQW global portfolio (33 REITs) does not have the same asset composition as the EQW domestic portfolio (26 REITs). The MDD statistics for both the global and domestic optimized portfolios are better than their EQW counterparts and, with the exception of SPY, better than those for the benchmarks. MDD Q25, Q50, and Q7 values and whisker lengths are better for the global portfolios than for the domestic ones.

Fig. 5.18 compares SR statistics for the benchmarks and the global and domestic portfolios. The Q25, Q50, and Q75 values are better for the benchmarks than for the global and domestic minimum-risk portfolios, but they are not competitive with those for the optimized tangent portfolios. The Q25, Q50, and Q75 values are uniformly higher for the global minimum-risk portfolios than for their domestic counterparts. The Q25 and Q50 quantile values for the domestic tangent portfolios are higher than for their global counterparts, but the situation is reversed for the Q75 values. Because the relative ranking of the portfolios with respect to SS values is the same as for SR values, the SS box-whisker summary plots are not shown. Fig. 5.19 shows the statistics for the RR values. The global portfolios outperform the benchmarks and the domestic portfolios.




Figure 5.17 Box-whisker summary statistics of MDD values for the benchmarks and the long-only global and domestic portfolio optimizations subject to a 4% turnover constraint.




Figure 5.18 Box-whisker summary statistics of SR values for the benchmarks and the long-only global and domestic portfolio optimizations subject to a 4% turnover constraint.

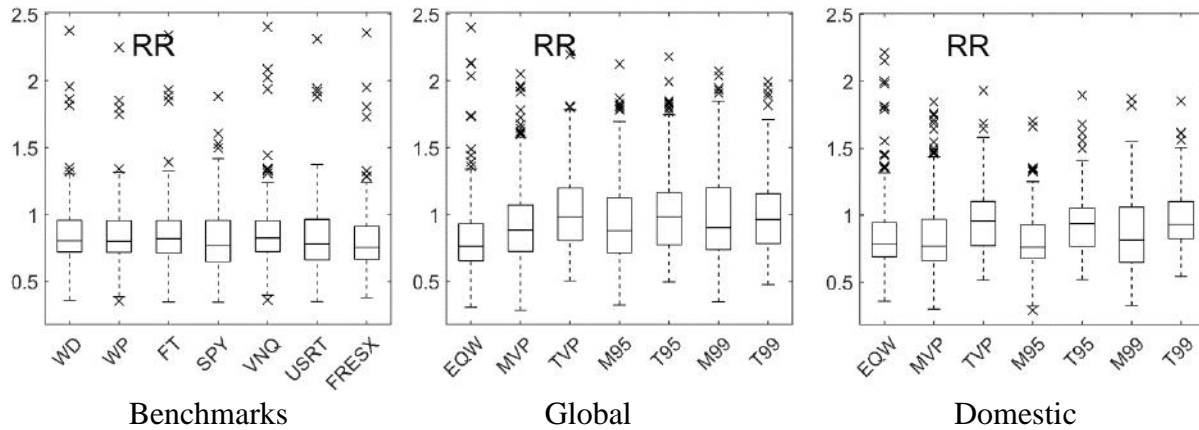


Figure 5.19 Box-whisker summary statistics of RR values for the benchmarks and the long-only global and domestic portfolio optimizations subject to a 4% turnover constraint.

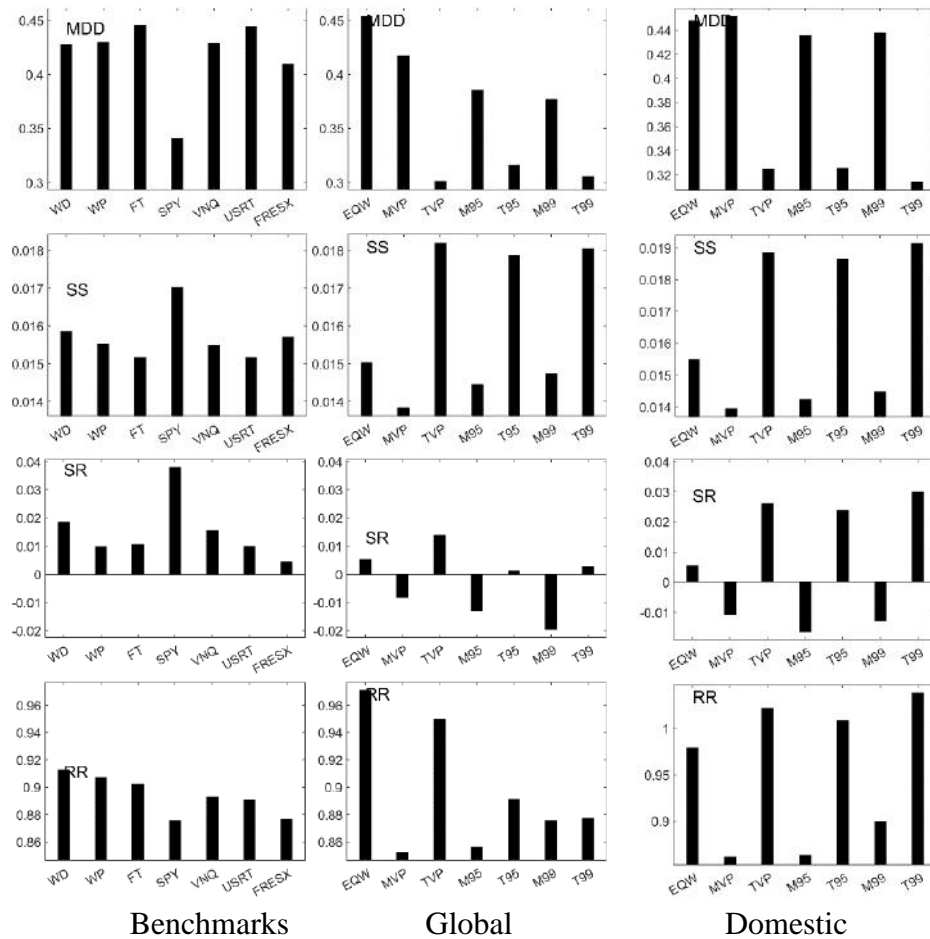


Figure 5.20 Performance measure values computed measured for the time period 07/30/2018 through 12/18/2020 for the benchmarks identified in section 2.3 and the long-only global and domestic portfolio optimizations subject to a 4% turnover constraint.

Fig. 5.20 compares the overall values of the risk measures for the time period. In terms of total MDD, which is determined by the pandemic-related decline, the global tangent portfolios have the

5. Historic Portfolio Optimization – Diversification with International REITs

smallest values, falling by 30%–31%, whereas the domestic tangent portfolio drops by 31%–32%. SPY declines by 34%, whereas all the minimum-risk portfolios and benchmarks experience 40%–45% drops. All the minimum-risk portfolios have a negative total SR over this period, a consequence of adding the international REITs (see Figs. 4.20 and 5.7). SPY shows the best total SR value (0.038), followed by the tangent domestic portfolios (0.024–0.030). The tangent portfolios, both global and domestic, have the best total SS values, and the domestic tangent values (0.019) are better than the global ones (0.018). The domestic portfolios outperform their respective global portfolios in terms of total RR. In addition, the domestic tangent portfolios outperform all the benchmarks in terms of total RR.

5.2.2 Jacobs et al. Long–Short Global Portfolios

The effect of diversification on long–short portfolios calls for a brief discussion. We consider the Jacobs et al. optimization with $s = 0.1$ under the turnover constraint $C_{TO} = 0.04$. Fig. 5.21 compares the cumulative price and log-return of the global and domestic portfolios. Paralleling the long-only case, the results are better for the global portfolios prior to the onset of the pandemic and continue to be better for the minimum-risk global portfolios during the pandemic. However, the domestic tangent portfolios perform best during the pandemic.




Figure 5.21 Cumulative price (left) and log-return (right) for the Jacobs et al. long– (a) global and (b) domestic portfolio optimizations subject to the constraints $s = 0.1$ and $C_{TO} = 0.04$.

Fig. 5.22 compares the total values of the risk measures. The global portfolios have better MDD values (by 6%–10%) than their domestic counterparts, and the global T99 portfolio shows the smallest MDD. The SRs continue to be negative for the minimum-risk portfolios; the tangent

global portfolios have the best SR values. SS values for the domestic and global portfolios are virtually identical, and all the tangent portfolios provide better SS values than the benchmarks. RR values for the domestic portfolios are better than their global counterparts, and the domestic tangent portfolios produce the best overall values.

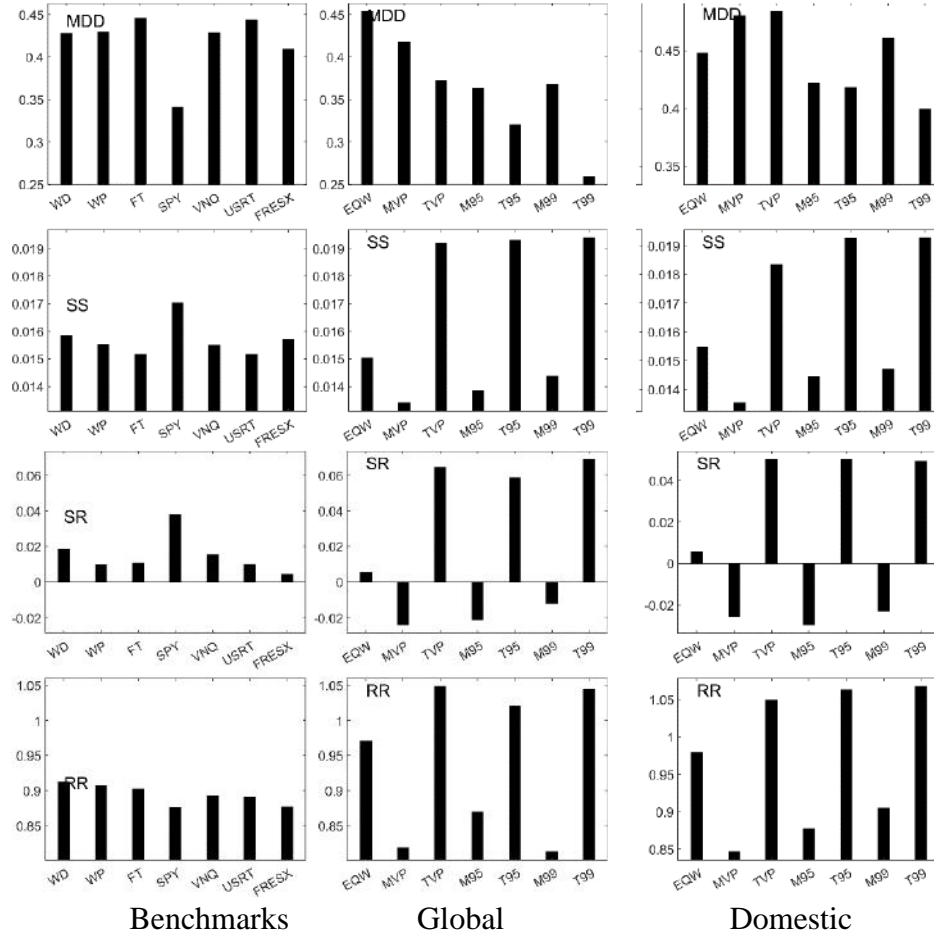


Figure 5.22 Performance measure values computed for the time period 07/30/2018 through 12/18/2020 for the benchmarks identified in section 2.3 and the Jacobs et al. long-short global and domestic portfolio optimizations subject to the constraints $s = 0.1$ and $C_{TO} = 0.04$.

References

- Eun, C. S., Lai, S., Roon, F. A. D. & Zhang, Z. (2010). International diversification with factor funds. *Management Science*, 56(9), 1500–1518.
- Eun, C. S. & Resnick, B. G. (1994). International diversification of investment portfolios: U.S. and Japanese perspectives. *Management Science*, 40(1), 140–161.
- Li, K., Sarkar, A. & Wang, Z. (2003). Diversification benefits of emerging markets subject to portfolio constraints. *Journal of Empirical Finance*, 10(1), 57–80.
- Solnik, B. H. (1995). Why not diversify internationally rather than domestically? *Financial Analysts Journal*, 51(1), 89–94.

Chapter 6

Black-Litterman Optimization Results

In this chapter, we apply the Black-Litterman optimization model, described in section 3.3, to our REIT portfolios. Because subjective views are specific to the market day and analyst, we run our Black-Litterman model on the portfolios without any specification of analyst views.⁵⁹ Without analysis views, the view-related quantities \mathbf{P} , $\mathbf{\Omega}$, and \mathbf{q} are zero; from (3.46), we have $\mathbf{r}_{BL} = \boldsymbol{\pi}$ with $\boldsymbol{\Sigma}_{BL} = \mathbf{C} = \tau\boldsymbol{\Sigma}$; and the Black-Litterman model captures the added influence of the market portfolio. Thus, the primary outputs⁶⁰ of the Black-Litterman module (equation (3.47)) are $\boldsymbol{\pi}$ and $(1 + \tau)\boldsymbol{\Sigma}$. We determine $\boldsymbol{\pi}$ using (3.49), where the weights $\mathbf{w}^* = \mathbf{w}_{mkt}$ are determined by a least-squares linear regression of our historical⁶¹ asset returns to those of the market benchmark. We employ FRESX⁶² as the market benchmark. The SR value in the risk-aversion coefficient, γ , in (3.50) is determined from the mean and standard deviation of the FRESX benchmark's returns.

6.1 Domestic Portfolios

Because FRESX employs a long-only investment strategy, the Black-Litterman values of $\boldsymbol{\pi}$ and $(1 + \tau)\boldsymbol{\Sigma}$ are used as inputs into the MVP and TVP optimization routines employing a long-only investment strategy. The goal of the Black-Litterman method is to track a benchmark subject to analyst corrections, so we compare the results of the Black-Litterman optimized portfolio against the performance of FRESX. The MVP and TVP optimizations subjected to the Black-Litterman method are labeled “BL MVP” and “BL TVP.” Those without the Black-Litterman method are labeled simply “MVP” and “TVP,” as in previous chapters. Fig. 6.1 shows the turnover time series for BL MVP, BL TVP, MVP, and TVP run with no turnover constraint. The time series for BL MVP and MVP are remarkably similar. Strong turnover differences between BL TVP and TVP reflect the impact the Black-Litterman scheme has on TVP optimization.




Figure 6.1 Turnover time series for the BL MVP and BL TVP (left) and MVP and TVP (right) long-only domestic portfolio optimizations. No turnover constraint is imposed on these optimizations.

Fig. 6.2 shows the influence of changing turnover constraints on the cumulative price of BL MVP and BL TVP run on the domestic portfolio. These are compared to the results for running

⁵⁹ Once expert views of the market are specified, they can be added into the simulation.

⁶⁰ In each rolling window; see section 3.4.

⁶¹ The regression is performed in each rolling window; see section 3.4.

⁶² The fund seeks above-average income and long-term capital growth, and it targets yield that exceeds the composite yield of the S&P 500. See <https://finance.yahoo.com/quote/FRESX/>.

the long-only optimizations MVP and TVP (without the Black–Litterman method). The performance of the minimum-risk portfolios, MVP and BL MVP, is relatively unchanged both with and without the Black–Litterman method and under changing TO constraints. It tends to underperform FRESX through most of the period but recovers to match FRESX during the pandemic. The tangent optimization shows much greater sensitivity. TVP greatly outperforms FRESX and shows strong price improvement when C_{TO} is decreased to 0.004 (100% possible turnover in the portfolio weights in one year). The Black–Litterman method strongly controls the performance of BL TVP. With no turnover constraint, BL TVP outperforms FRESX, though not to the extent that TVP does. Under $C_{TO} = 0.04$, BL TVP tracks FRESX in terms of price performance. Reducing to $C_{TO} = 0.004$, BL TVP again outperforms FRESX, though not as strongly as it does in the absence of a turnover constraint.

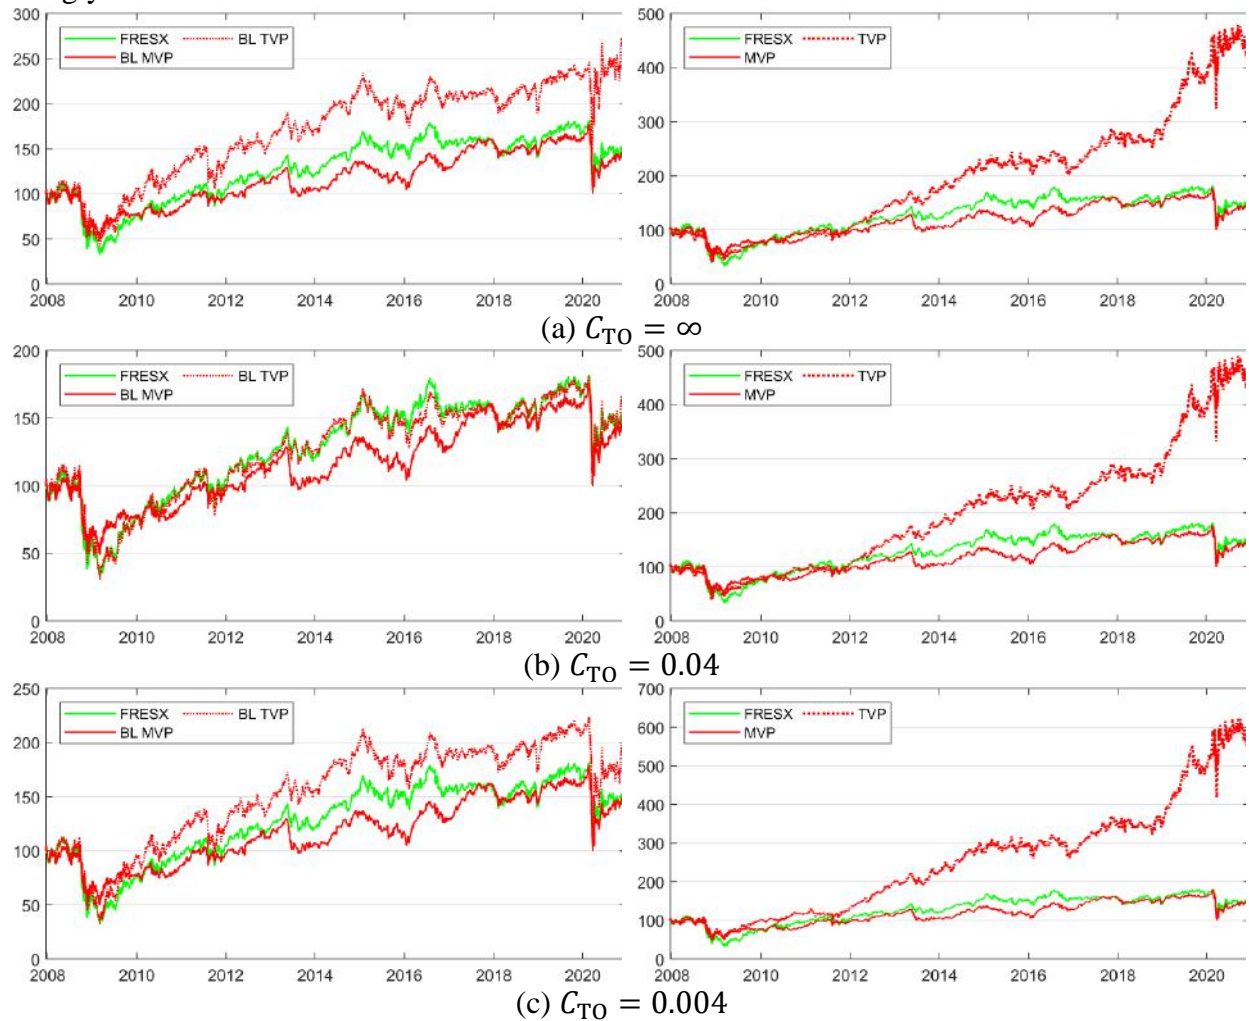


Figure 6.2 Cumulative price for the MVP and TVP long-only domestic portfolio optimizations with (left) and without (right) the Black–Litterman scheme under different levels of turnover constraint: (a) $C_{TO} = \infty$ (no constraint), (b) $C_{TO} = 0.04$, and (c) $C_{TO} = 0.004$. Price performance is compared to that of the market benchmark FRESX.

Fig. 6.3 compares risk-measure time series for BL MVP, BL TVP, MVP, and TVP when $C_{TO} =$

6. Black-Litterman Optimization Results

0.004. BL MVP and MVP time series of all four risk measures are essentially the same, as can be seen in Fig. 6.4, which summarizes the box-whisker statistics of the risk measures. In contrast, BL TVP tracks FRESX in SR, SS, and RR much better than BL MVP and clearly much better than TVP. Fig. 6.4 shows how the risk-measure box-whisker summaries for SR, SS, and RR are brought into closer alignment with those for FRESX. Interestingly, the same is not true for MDD.

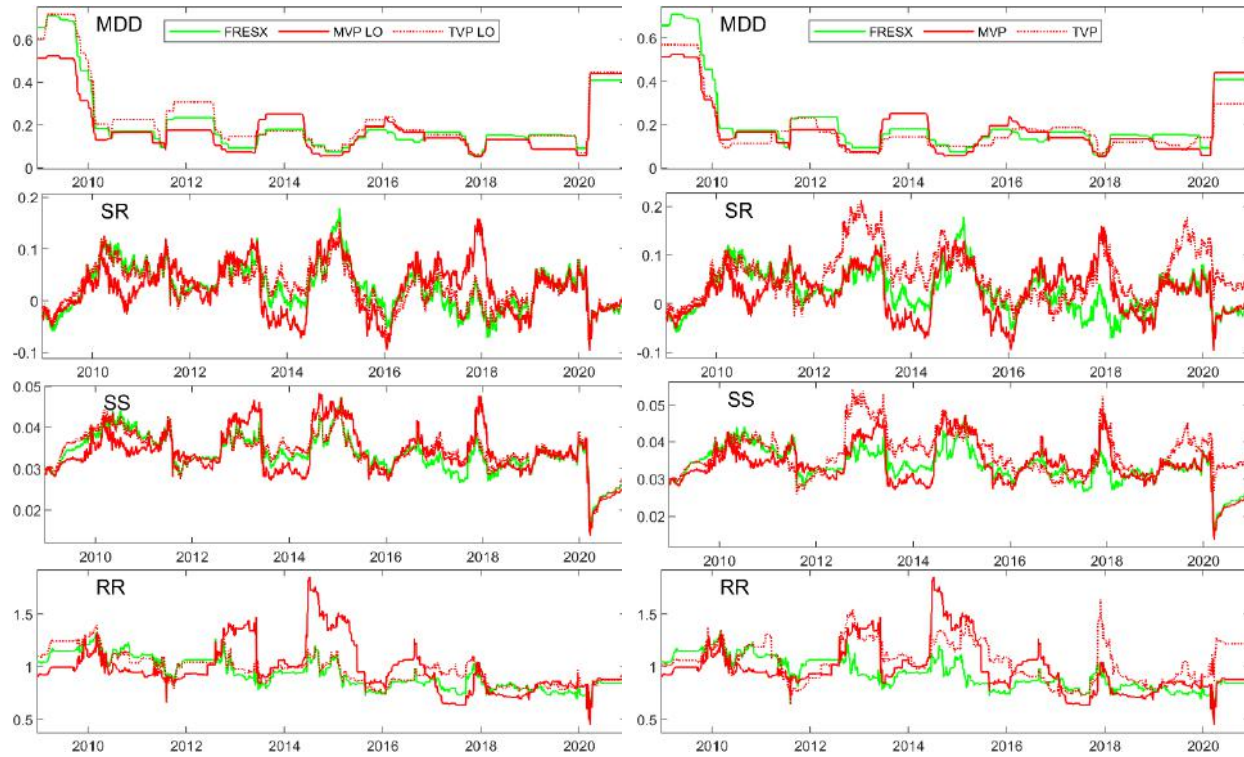


Figure 6.3 Risk-measure time series for the BL MVP and BL TVP (left) and MVP and TVP (right) long-only domestic portfolio optimizations subject to a 0.4% turnover constraint.

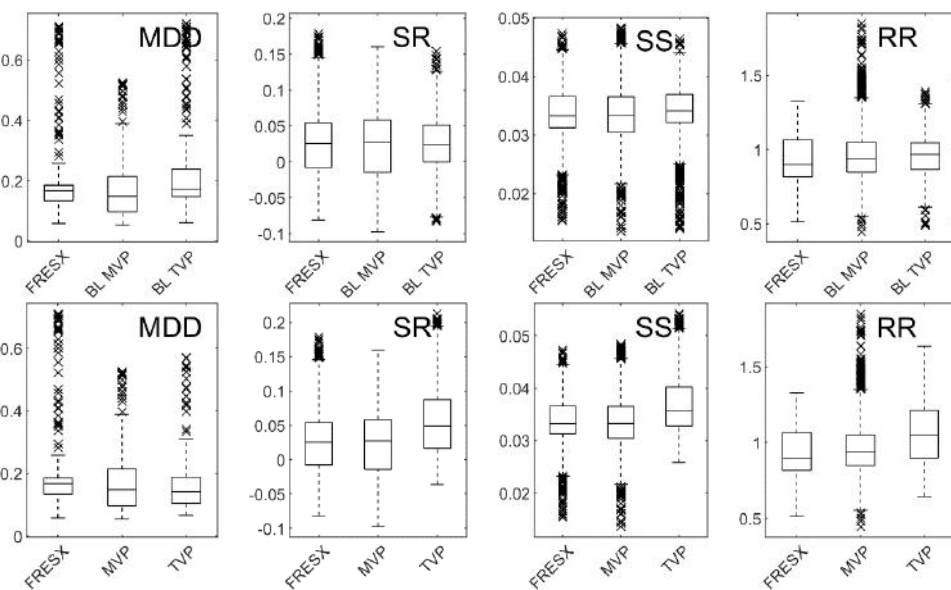


Figure 6.4 Box-whisker summaries of the risk measures for FRESX and the BL MVP and BL TVP (top) and MVP and TVP (bottom) long-only domestic portfolios optimized subject to a 0.4% turnover constraint.




Figure 6.5 Performance measure values computed for the time period 12/16/2008 through 12/18/2020 for FRESX and the BL MVP and BL TVP (left) and MVP and TVP (right) long-only domestic portfolio optimizations subject to a 0.4% turnover constraint.

Fig. 6.5 compares the risk measures computed for the total time period for FRESX, BL MVP, BL TVP, MVP, and TVP. The overall risk measures for BL MVP and MVP are the same; both have the best MDD value but the worst values for overall SR and RR. TVP has the best SR, SS, and RR values. The effect of the Black–Litterman scheme is to align the overall value of every risk measure more closely to that of FRESX. As a result, although BL TVP still retains better SR, SS, and RR values than BL MVP and FRESX, BL TVP develops the worst overall MDD value.

6.2 Global Portfolios

We investigate the combined effect of diversification by adding the international assets and imposing the Black–Litterman method. As noted in Chapter 5, the time period for portfolio optimizations is reduced to 04/30/2018 through 12/18/2020, and risk-measure statistics cover the period 07/30/2018 through 12/18/2020. As in section 6.1, we use FRESX as the market benchmark and consider long-only optimization under the turnover constraint $C_{TO} = 0.004$. Fig. 6.6 compares the performance of MVP and TVP optimizations for the domestic and global portfolios with and without the application of the Black–Litterman scheme. Starting with the MVP and TVP optimizations of the domestic long-only portfolios (Fig. 6.6(a)), it is evident that adding the international REITs worsens the price performance of TVP and slightly improves that of MVP (Fig. 6.6(b)). Applying the Black–Litterman method without diversification (Fig. 6.6(c)) dramatically worsens the performance of TVP, aligning it strongly with FRESX, but does not markedly change the performance of MVP. Combining diversification and the Black–Litterman method (Fig. 6.6(d)) results in strong alignment of the optimization schemes both to each other and to FRESX prior to the pandemic. During the pandemic, the performance of BL MVP and BL TVP once again begins to approach that of FRESX near the end of 2020.

Fig. 6.7 explores the effects of diversification and the Black–Litterman method on MDD. Diversification alone generally reduces MDD for MVP. The addition of the Black–Litterman method does not affect MDD statistics for MVP but generally increases MDD for TVP. Fig. 6.7 also explores the effects of diversification and the Black–Litterman method on SR. Diversification

6. Black-Litterman Optimization Results

tends to slightly increase the SR values for MVP, whereas the addition of the Black–Litterman method has a weaker effect. Diversification has less influence on the SR values for TVP, whereas the addition of the Black–Litterman scheme reduces TVP SR values. For brevity, the effects on SS and RR are not shown. The effects on SS are similar to those on SR. Diversification tends to increase RR for both MVP and TVP. Adding the Black–Litterman scheme tends to have little effect on MVP and generally reduces RR for TVP.

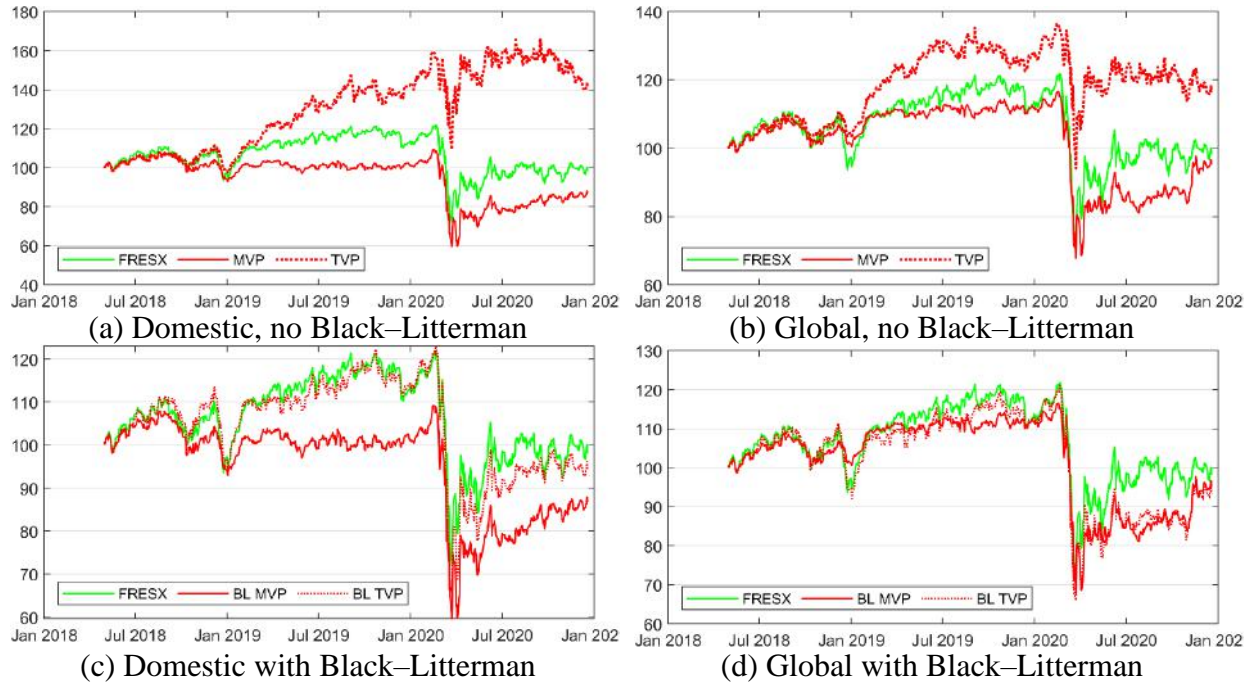
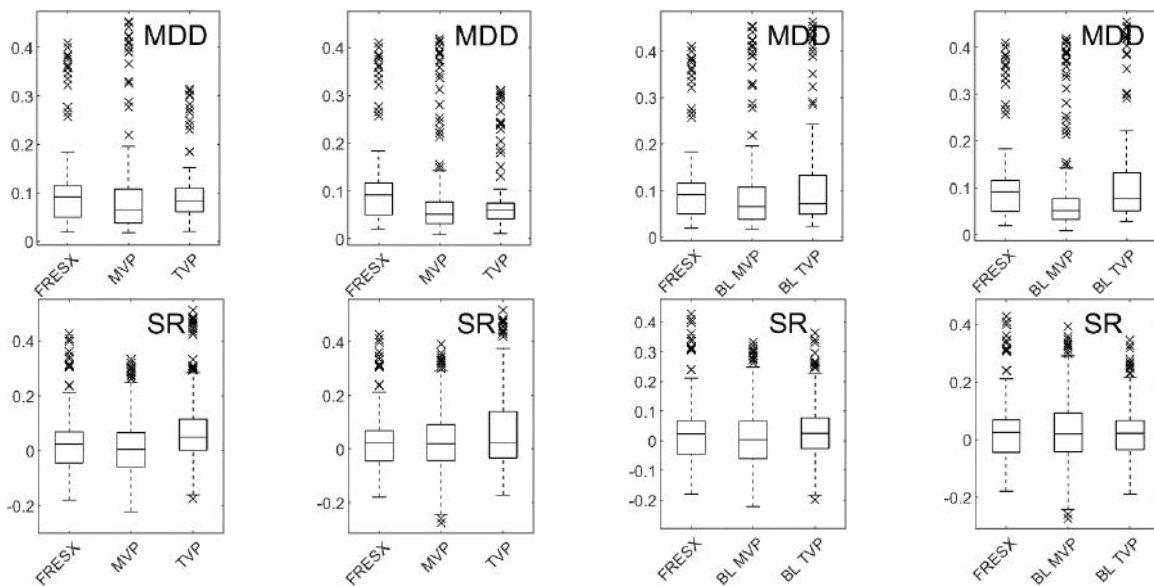


Figure 6.6 Cumulative price for the MVP, TVP, BL MVP, and BL TVP long-only domestic and global portfolio optimizations subject to a 0.4% turnover constraint. Price performance is compared to that of the market benchmark FRESX.



(a) Domestic, no
Black–Litterman(b) Global, no
Black–Litterman(c) Domestic with
Black–Litterman(d) Global with
Black–Litterman

Figure 6.7 Box-whisker summary of MDD (top) and SR (bottom) statistics for FRESX and the MVP, TVP, BL MVP, and BL TVP long-only domestic and global portfolio optimizations subject to a 0.4% turnover constraint.

Fig. 6.8 summarizes the total-time-period risk-measure values. The Black–Litterman method of alignment with the market portfolio FRESX dramatically worsens the total-risk-measure values for TVP, whereas the impact on MVP values is slight.

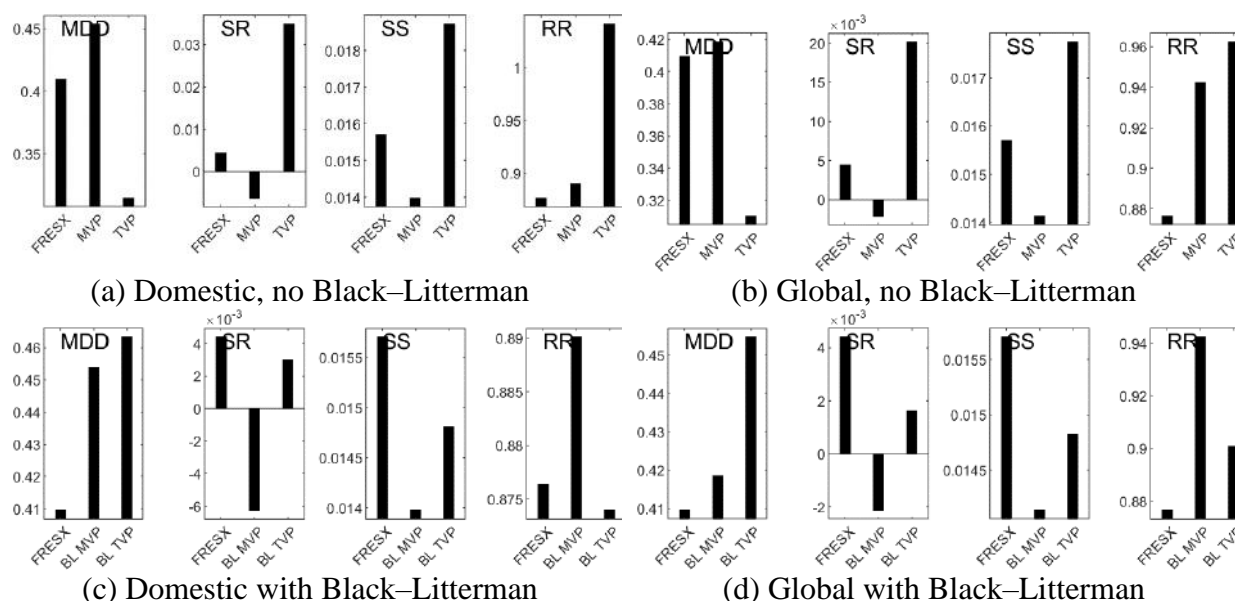


Figure 6.8 Performance measure values computed for the time period 07/30/2018 through 12/18/2020 for FRESX and the BL MVP and BL TVP (left) and MVP and TVP (right) long-only domestic and global portfolio optimizations subject to a 0.4% turnover constraint.

Chapter 7

Dynamic Portfolio Optimization: Beyond MPT

The historical optimization method discussed in section 3.4 and illustrated in Chapters 4–6 samples return data sequentially by moving windows from a fixed historical period containing a finite sample of market activity and world events. As is usually noted in any fund prospectus, *past performance may not be indicative of future returns*. Rather than simply using the samples of historical asset returns, dynamic optimization expands on the informational content measurable from the historical data. Specifically, the historical returns are assumed to come from a dynamic multivariate distribution (dynamic in the sense that its moments, such as the covariance, may change over time), of which the historical returns are a representative *but limited* sample. Dynamic optimization attempts to discover the character of this distribution and then to generate *very large* predictive samples of correlated asset returns that sample much more of the tail behavior (i.e., extreme events) of the distribution. The result is a portfolio optimization that is more “attuned” to possible dramatic changes in market performance.

The form of dynamic optimization we implement here contains the following elements. In each moving window, a very general ARMA(1,1)–GARCH(1,1) model is used to fit the return data. A Student’s t -distribution, which allows for fatter tails than a normal distribution, is utilized as the model for the innovations in the ARMA–GARCH fit. More inventively, the innovations from the fit are empirically transformed into copula “space.” Under a copula transformation, all regions of the distribution, including tail regions, are given equal weighting. The copula-transformed innovations are fit to a multivariate t -copula distribution, which captures the covariance behavior among the asset innovations. A very large sample of asset-innovation values is then generated from the t -copula distribution. After applying inverse copula transformations to the sample set, these innovations are utilized in the ARMA–GARCH fit to generate a very large set of portfolio asset-return values, which are fed into the optimization routine to generate the next-day weights for the portfolio.

We discuss the individual steps of this optimization in greater detail in sections 7.1.1–7.1.3. In section 7.1.4, we describe the method for combining dynamic optimization with the Black–Litterman method. In section 7.2, we illustrate dynamic long-only and long–short optimization of the domestic, international, and global portfolios. In section 7.3, we illustrate the combination of the Black–Litterman method with dynamic optimization of the long-only domestic portfolio.

7.1 Dynamic Optimization

Consider the mean-variance optimization method of section 3.1. As diagrammed in Fig. 3.3, it requires $n + n(n + 1)/2 = n(n + 3)/2$ input values, consisting of the entries of the vector $\bar{\mathbf{r}}$ of mean returns and the entries of the diagonal covariance matrix $\mathbf{\Sigma}$. Consider a moving window of T days. It contains nT asset-return values, $r_{i,t}$. Thus, whereas $\bar{r}_i = T^{-1} \sum_{t=1}^T r_{i,t}$ provides an estimate for \bar{r}_i with an error $\sim T^{-1/2}$, estimates of $\Sigma_{ij} = \sum_{t=1}^T (r_{i,t} - \bar{r}_i)(r_{j,t} - \bar{r}_j)$ exhibit a higher degree of error, because the same set of T values for $(r_{i,t} - \bar{r}_i)$ is used to compute the values for Σ_{ij} , $j = 1, \dots, n$, $j \neq i$. Because this observation holds for each asset i , a restricted set of independent observations goes into the computation of $\mathbf{\Sigma}$. Stated equivalently, for any practically sized moving window, there is the potential for insufficient sampling of the (range of) covariances

experienced between the assets. Dynamic optimization seeks to provide a statistically accurate large sampling. We do so in a manner that combines a historical window $[t - T + 1, t]$ of return data with a dynamic prediction of returns for day $t + 1$ to feed into an optimization for the weights on day $t + 1$.

7.1.1 ARMA(1,1)–GARCH(1,1) with Student's *t*-Distribution

If a time series, r_t , of returns is stationary, a useful general model for describing the time series is the synthesis of the autoregressive moving-average (ARMA) model and the generalized autoregressive conditional heteroscedasticity (GARCH) model. The ARMA (Engel, 1982) component explicitly models the behavior of the return, whereas the GARCH (Bollersley, 1986) component explicitly models its variance. Both models contain (theoretically infinite) parameters; the variations in the models are denoted by the (finite) number of parameters employed. The ARMA(p, q) model (Tsay, 2010) is

$$r_t = \varphi_0 + \sum_{i=1}^p \varphi_i r_{t-i} + a_t + \sum_{j=1}^q \theta_j a_{t-j}, \quad (7.1)$$

where each *shock*, a_t , is a zero-mean random variable. The first two terms in (7.1) describe the autoregressive dependence of r_t on previous returns; the second two terms add the influence of a weighted (moving) average of shocks, a_t . The GARCH(m, s) model relates a_t to, and provides a model for, the variance σ_t^2 of the series

$$a_t = \sigma_t \epsilon_t, \quad \sigma_t^2 = \alpha_0 + \sum_{i=1}^m \alpha_i a_{t-i}^2 + \sum_{j=1}^s \beta_j \sigma_{t-j}^2. \quad (7.2)$$

Here, the so-called *innovations*, ϵ_t , are zero-mean, unit-variance, independent, identically distributed random variables. The GARCH model is clearly autoregressive in both σ_t^2 and a_t^2 . Identifying the daily variance as the volatility of the time series, (7.2) captures the property of conditional heteroscedasticity, that is, the property that the volatility is nonconstant relative to that of prior days.

With six parameters, the ARMA(1,1)–GARCH(1,1) model,

$$\begin{aligned} r_t &= \varphi_0 + \varphi_1 r_{t-1} + a_t + \theta_1 a_{t-1}, \\ a_t &= \sigma_t \epsilon_t, \\ \sigma_t^2 &= \alpha_0 + \alpha_1 a_{t-1}^2 + \beta_1 \sigma_{t-1}^2, \end{aligned} \quad (7.3)$$

provides enough generality to model many return time series.⁶³ However, providing a fit to a particular time series requires the specification of the distribution governing the innovation random variables. In the dynamic optimization method, we assume that the innovations, ϵ_t , are governed by the Student's *t*-distribution,

⁶³ To be sure, there are variations of the GARCH model (Bollersley, 1986) designed to include additional behaviors. Some examples are the I(ntegrated)GARCH (Engel and Bollersley, 1986); F(ractionally)I(ntegrated)GARCH (Baillie et al. 1996); E(xponential)GARCH (Nelson 1991); GARCH-M(ean) (Engel et al., 1986); T(hreshold)GARCH (Zakoian, 1994); and G(lost)J(agannathan)R(unkle)GARCH (Glosten et al., 1993) models.

$$t_v(x) = \frac{\Gamma((1+\nu)/2)}{\sqrt{\nu\pi} \Gamma(\nu/2)} \left(1 + \frac{x^2}{\nu}\right)^{-(1+\nu)/2}, \quad (7.4)$$

where $\Gamma()$ is the gamma function. This distribution is symmetric, but it is fat-tailed relative to the normal distribution. The parameter ν is known as the number of degrees of freedom. The distribution has finite variance for $\nu > 2$, zero skewness for $\nu > 3$, and finite kurtosis for $\nu > 4$. Otherwise, these moments are either infinite or undefined. As $\nu \rightarrow \infty$, the Student's t -distribution approaches the normal distribution. (In practice, the normal distribution is a sufficiently accurate approximation when $\nu \geq 30$.)

7.1.2 Multivariate t -Distribution and t -Copulas

The Student's t -distribution (7.4) is applicable to univariate random variables. The extension to random vectors, whose elements are random values that may be correlated, is the multivariate Student's t -distribution. Although there are many candidates for the multivariate generalization of the Student's t -distribution, we utilize the common form

$$t_v(\mathbf{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma}) = \frac{\Gamma\left(\frac{p+\nu}{2}\right)}{(v\pi)^{p/2} \Gamma\left(\frac{\nu}{2}\right) |\boldsymbol{\Sigma}|^{1/2}} \left(1 + \frac{(\mathbf{x} - \boldsymbol{\mu})^T \boldsymbol{\Sigma}^{-1} (\mathbf{x} - \boldsymbol{\mu})}{\nu}\right)^{-(p+\nu)/2}. \quad (7.5)$$

Here, \mathbf{x} is a p -dimensional random vector having mean $\boldsymbol{\mu}$ (for $\nu > 1$). (Because the distribution is symmetric, $\boldsymbol{\mu}$ is also the median and mode values of the distribution.) The $p \times p$ matrix $\boldsymbol{\Sigma}$ is proportional to the covariance matrix of the elements of \mathbf{x} ; specifically, the covariance is $\nu/(\nu - 2) \boldsymbol{\Sigma}$ for $\nu > 2$. The notation $|\boldsymbol{\Sigma}|$ denotes the determinant of $\boldsymbol{\Sigma}$. The parameters of the model are $\boldsymbol{\Sigma}$, $\boldsymbol{\mu}$, and ν . An attractive feature of this version of the multivariate t -distribution is that ν continues to be a scalar measure of the degrees of freedom.⁶⁴

The CDF for this multivariate distribution,

$$T_v(\mathbf{y}) = \Pr(\mathbf{X} \leq \mathbf{y}) \text{ where } \mathbf{X} \sim t_v(\mathbf{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma}), \quad (7.6)$$

does not have an analytic form but can be approximated using numerical integration.

As noted in section 10.1.3.1, the copula is a multivariate cumulative distribution function. From (7.6), the t -copula is the multivariate cumulative distribution given by

$$C_v^t(u_1, \dots, u_N) = T_v(t_v^{-1}(u_1), \dots, t_v^{-1}(u_N)), \quad (7.7)$$

where t_v^{-1} is the inverse of the Student's t -distribution (7.4).

7.1.3 Generation of Dynamic Returns

⁶⁴ In practice, $\boldsymbol{\Sigma}$ and $\boldsymbol{\mu}$ are estimated from sample data. Thus, fitting the distribution $t_v(\mathbf{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma})$ to a sample of multivariate data reduces to finding a best-fit value for ν .

As illustrated by the schematic in Fig. 7.1, dynamic optimization introduces an additional module that converts a window of historical returns into a dynamic set of returns that are then passed to the portfolio-optimizing routine.

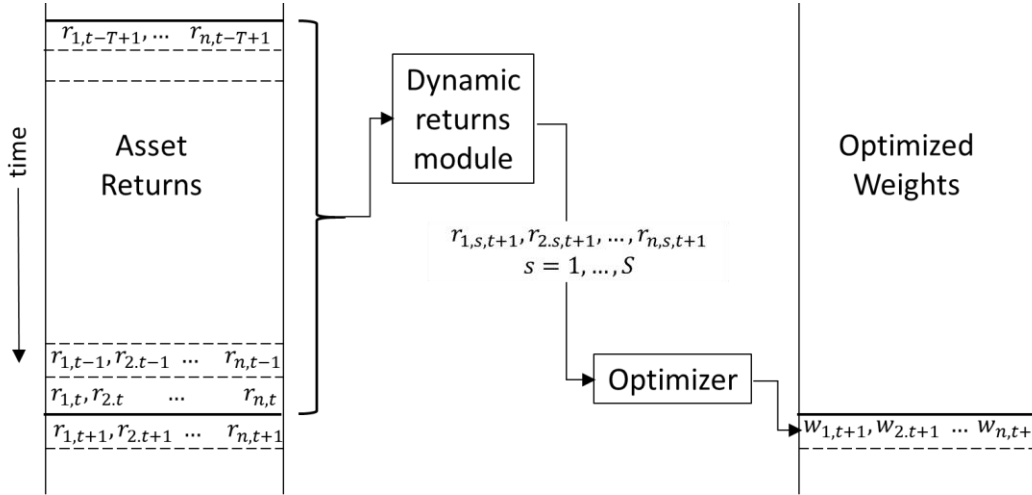


Figure 7.1. Schematic of dynamic portfolio optimization.

As noted in the introduction to this section, the dynamic module seeks to provide a statistically accurate larger sampling of returns for subsequent optimization. This is achieved by the following steps.

- S1. Let $\{r_{i,t-T+1}, \dots, r_{i,t}\}$ denote the historical return series for asset i , $i = 1, \dots, n$ in the window $[t - T + 1, t]$. Fit an ARMA(1,1)–GARCH(1,1)–Student’s t -distribution model to the time series of each asset, generating the parameters $\varphi_{0,i}, \varphi_{1,i}, \theta_{1,i}, \alpha_{0,i}, \alpha_{1,i}, \beta_{1,i}, \nu_i$, $i = 1, \dots, n$.
- S2. Compute the shock series $\{a_{i,t-T+1}, \dots, a_{i,t}\}$ as the difference between the historical time series $\{r_{i,t-T+1}, \dots, r_{i,t}\}$ and the values predicted by the ARMA(1,1) model in (7.3). From the shock series and the variances σ_k^2 , $k = t - T + 1, \dots, t$, predicted from the fitted GARCH(1,1) model in (7.3), compute the innovation series $\{\varepsilon_{i,t-T+1}, \dots, \varepsilon_{i,t}\}$, $i = 1, \dots, n$.
- S3. Recall that the innovations have been fit to a Student’s t -distribution in step S1. Perform the copula transformations, $u_{i,t-T+1}, \dots, u_{i,t} = \tau_{\nu_i}(\varepsilon_{i,t-T+1}, \dots, \varepsilon_{i,t})$, where τ_{ν} is the CDF for the univariate Student’s t -distribution (7.5).
- S4. Fit the transformed innovations $\{u_{i,t-T+1}, \dots, u_{i,t}\}$, $i = 1, \dots, n$ to a t -copula (7.7) using an objective function that approximates the log-likelihood for the degrees-of-freedom parameter, ν , of the multivariate t -distribution.
- S5. Generate $S \times n$ correlated samples, $\{u_{i,1}, \dots, u_{i,S}\}$, $i = 1, \dots, n$ from the t -copula, where $S \gg T$. (Recall that T is the size of the historical window.)
- S6. Perform the inverse transformations $\varepsilon_{i,1}, \dots, \varepsilon_{i,S} = \tau_{\nu_i}^{-1}(u_{i,1}, \dots, u_{i,S})$, $i = 1, \dots, n$.

7. Dynamic Portfolio Optimization

S7. Use these transformed, correlated innovations in the ARMA(1,1)–GARCH(1,1) model of step S1 to generate a dynamic ensemble $\{r_{i,s,t+1}, s = 1, \dots, S\}$ of S predicted returns for each asset i for day $t + 1$:

$$\begin{aligned}\sigma_{i,t+1}^2 &= \alpha_0 + \alpha_1 a_{i,t}^2 + \beta_1 \sigma_{i,t}^2, \\ a_{i,s,t+1} &= \sigma_{i,t+1} \epsilon_{i,s}, \\ r_{i,s,t+1} &= \varphi_0 + \varphi_1 r_{i,t} + a_{i,s,t+1} + \theta_1 a_{i,t}, \\ s &= 1, \dots, S; \quad i = 1, \dots, n.\end{aligned}\tag{7.8}$$

The values $\sigma_{i,t}$ and $\sigma_{i,t}$ used in (7.8) are generated in step S2.

The ensemble of returns $\{r_{i,s,t+1}, s = 1, \dots, S; i = 1, \dots, n\}$ represents the output from the dynamic module, which is then fed to the portfolio optimizer, as illustrated in Fig. 7.1. The CVaR-based optimizations require the entire ensemble of return values, $\{r_{i,s,t+1}, s = 1, \dots, S; i = 1, \dots, n\}$, whereas the mean-variance optimizations require only the mean return of each asset and the covariance matrix computed from this ensemble. The goals of the dynamic formulation outlined in the introduction to this section are therefore accomplished as follows. The ARMA(1,1)–GARCH(1,1)–Student’s t -distribution model in step S1 facilitates the extraction (step S2) of the set of zero-mean, unit-variance, random innovations that characterize the behavior of the (correlated) asset returns over the historical window $[t - T + 1, t]$. As noted in section 7.2.2.1, mapping the innovations to copula space (step S3) preserves the correlation behavior while weighting tail events with larger probability. Because tail information is precisely the area where samples are lacking in the historical window data set, transforming to copula-space enables better sampling of the tails of the innovation distribution. This is accomplished in steps S4–S6. The effect of steps S3–S6 is to take a $T \times n$ sample set of innovations and replace it with a much larger $S \times n$ sample that retains the same statistical distribution. The final, notable idea in step S7 is to use the ARMA(1,1)–GARCH(1,1) model to generate a large ensemble of *predicted* return values for day $t + 1$ that is fed into the portfolio optimizer to generate optimal weight for use during day $t + 1$.

7.1.4 Combining the Dynamic Approach with Black–Litterman Optimization

Fig. 7.2 illustrates the modification of Fig. 7.1 used to combine the dynamic approach with the Black–Litterman optimization. The Black–Litterman optimization outlined in equations (3.36)–(3.50) in section 3.6 requires $\boldsymbol{\pi}$ and $\boldsymbol{\Sigma}$ (in addition to the managerial-team views) as input. In the historical optimizations presented in section 3.6, $\boldsymbol{\Sigma}$ is estimated from the $n \times T$ returns in the historical window $[t - T + 1, t]$. Similarly, the market weights required for computing $\boldsymbol{\pi}$ from (3.49) are obtained by linearly regressing $n \times T$ returns against returns (covering the same time period) from a market benchmark. Under the dynamic approach, the dynamic return ensemble $\{r_{i,s,t+1}, s = 1, \dots, S; i = 1, \dots, n\}$ does not correlate with any historical returns from a benchmark. This problem is solved by including the returns of the market benchmark as an $n + 1$ st asset in step S1 and performing the entire series S1–S7 for these $n + 1$ assets. The result is an $(n + 1) \times S$ data set of returns. The $n \times S$ ensemble of returns for the portfolio assets can then be linearly regressed against the corresponding S ensemble of returns for the benchmark to generate

the weights, \mathbf{w}_{mkt} , of the market portfolio. The ensemble benchmark returns and \mathbf{w}_{mkt} are used in the computation of the risk-aversion coefficient γ in (3.50).

The second important change is that the values of the correlation matrix, Σ , appearing in (3.36)–(3.50) are now computed from the $n \times S$ ensemble of asset returns.

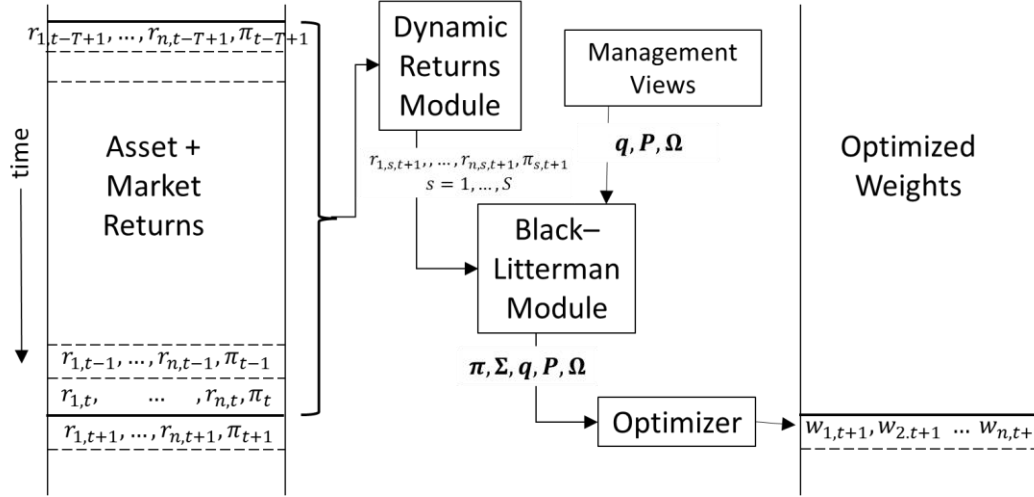


Figure 7.2. Schematic of the implementation of the Black–Litterman approach with dynamic portfolio optimization.

7.2 Portfolio Optimization Using Dynamic Returns

We consider the impact that the dynamic returns presented in section 7.1 have on the optimized portfolios examined in Chapters 4–6. We first consider dynamic optimization without the addition of the Black–Litterman scheme (Fig. 7.1).

7.2.1 Dynamic Long-Only Portfolios

Fig. 7.3 shows the cumulative price and log-return of the six MPT-based long-only domestic portfolios computed with dynamic optimization subject to no turnover constraint ($C_{\text{TO}} = \infty$) and to a 4% daily constraint ($C_{\text{TO}} = 0.04$). The price and return plots for the corresponding historical (nondynamic) optimizations are given in Figs. 4.1 and 4.9. Compared with the historical optimizations of Fig. 4.1, the prepandemic performance of all portfolios is enhanced, especially the performance of the tangent portfolios, which develop strong returns in the period 2014–2016, roughly double that of their historical counterparts during this period. This 2014–2016 behavior is reminiscent of, though much more pronounced than, that of the long–short historical portfolios of Figs. 4.4, 4.6, and 4.15. However, during the pandemic, the dynamic tangent portfolios suffer greater drawdowns than their historical counterparts. Under a 4% daily turnover constraint, the dynamic tangent portfolios significantly outperform their historical counterparts both before and during the pandemic, although the 2014–2016 performance enhancement is replaced by one of steady growth. Under the 4% turnover constraint, the dynamic minimum-risk portfolios do not

7. Dynamic Portfolio Optimization

keep up with their historical counterparts, though their performance still equals or exceeds that of VNQ.




Figure 7.3 Cumulative price (left) and log-return (right) for the benchmarks and the dynamic long-only domestic portfolio optimizations subject to no turnover constraint (top) and a 4% turnover constraint (bottom).

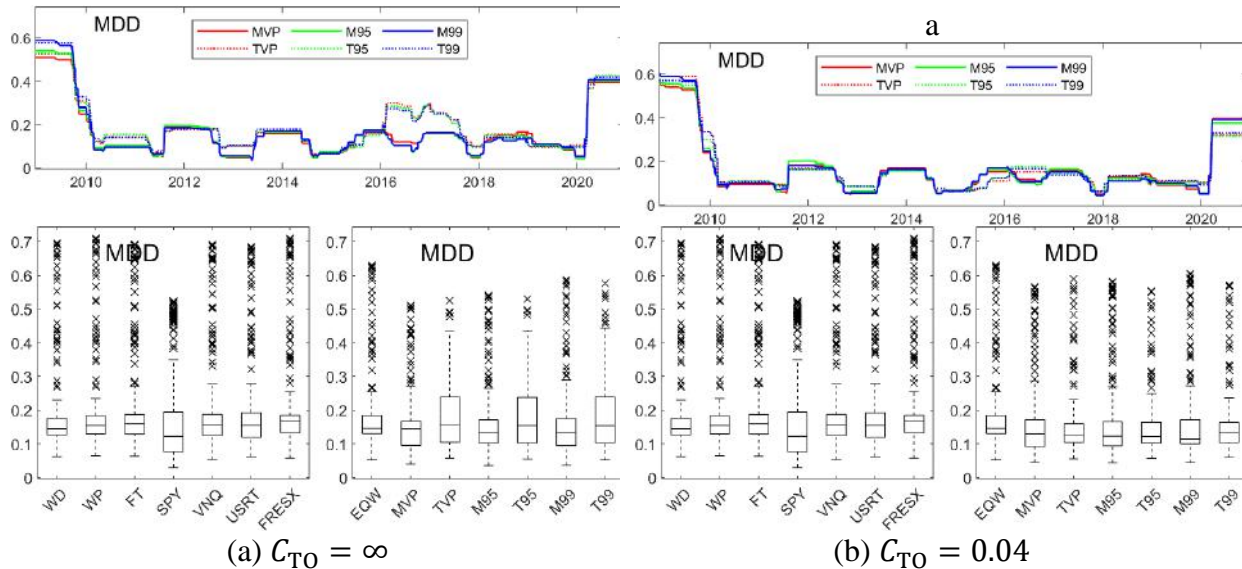


Figure 7.4 MDD time-series and box-whisker statistics (computed from one-year moving windows) for the dynamic long-only portfolio optimizations subject to (a) no turnover constraint and (b) a 4% turnover constraint.

Fig. 7.4 plots the MDD time-series and box-whisker summaries for these two dynamic portfolios. These can be compared to the results, shown in Figs. 4.18 and 4.19, for the historical long-only portfolio subject to 4% daily turnover (the results for the historical long-only portfolio subject to no turnover constraint are very similar to those shown in Fig. 4.18). The time series in Fig. 7.4 reveal that during the pandemic, the dynamic tangent portfolios under no turnover constraint exhibit a larger drawdown and that their MDD is reduced under the 4% constraint. In addition, the time series with no turnover constraint reflect the large MDDs for the three tangent portfolios during the 2016–2017 period accompanying the decline from the 2014–2016 period of high returns. The box-whisker plots show large MDD IQRs for the tangent portfolios under no turnover constraint and a strong reduction of that range under the 4% daily turnover constraint.




Figure 7.5 Box-whisker summary statistics of TO (4.5) and the L_2 -norm (4.7) for the dynamic long-only domestic portfolio optimizations subject to (a) no turnover constraint and (b) a 4% turnover constraint.

Fig. 7.5 shows the TO and L_2 -norm box-whisker summary statistics for these dynamic optimizations. A comparison to Fig. 4.12, which shows the same for the corresponding historical optimizations, indicates that under dynamic optimization, there is a very strong shift to higher daily turnover values. Unconstrained, the TO IQRs for the dynamic tangent portfolios span from a Q1 value of 55% to a Q3 value of 100%, whereas for the historical tangent portfolios, TO Q3 values are under 6%. Under a 4% daily constraint, all six dynamic portfolios have $Q1 = Q3 = 4\%$; by contrast, among the corresponding six historical portfolios, only T95 and T99 have TO Q3 values approaching 4%.

As noted in the discussion of our implementation of turnover constraints following (4.6), there are days on which no optimized weight solution can be found under the imposed value of C_{TO} . This is the case for $C_{TO} = 0.04$ in Fig. 7.5, where for such days, the value of C_{TO} is increased through the sequence $\{0.04, 0.05, 0.06, 0.08, 0.12\}$ until an optimized weight solution can be found. If no solution is obtained under this sequence, then asset-weight assignments for day t remain unchanged from those for day $t - 1$ (generating TO and L_2 -norm values of zero for such days). The L_2 -norm distinguishes among outlier values having the same TO value. The L_2 -norm box-whisker plot therefore gives a better view of the number and separation of the outlier TO values that exceed 0.04. Table 7.1 documents the frequency of increase in C_{TO} value required to obtain weight solutions for the portfolios corresponding to Fig. 7.5(b). For the MVP portfolio, 100% of the daily optimizations succeed under the 4% turnover constraint. For M95, only 95.3%

7. Dynamic Portfolio Optimization

succeed under the 4% constraint; an additional 0.6% succeed with the turnover constraint raised to 5%; etc. For 1.5% of the dates, no optimized solution with a turnover constraint of 12% can be found, and the portfolio weights for that date remain unchanged from the previous trading day. For the remaining portfolios, optimization under the 4% turnover constraint succeeds on over 95% of the dates. On 1.5% or less of the dates for the M95 and M99 portfolios, no optimizing solution with turnover less than 12% can be found. For the T95 and T99 portfolios, this “failure” rate increases to 2.2% and 2.5%, respectively.

Table 7.1 Percent of daily optimizations succeeding at each value of turnover constraint for the dynamic long-only domestic portfolio optimizations.

C_{TO}	0.4	0.5	0.6	0.8	0.12	$w(t) = w(t - 1)$
MVP	100.0	0.0	0.0	0.0	0.0	0.0
TVP	99.7	0.3	0.0	0.0	0.0	0.0
M95	95.3	0.6	0.6	0.9	1.0	1.5
T95	95.1	0.4	0.4	0.5	0.9	2.5
M99	95.6	0.6	0.5	1.0	1.0	1.3
T99	95.5	0.4	0.4	0.5	1.0	2.2

Fig. 7.6 shows the time series for the SR, SS, and RR risk measures for these dynamic portfolios under a 4% turnover constraint and under no constraint. Prior to 2013, risk measures for the tangent portfolios are better under no turnover constraint; the situation reverses after 2013, with noticeable improvement during the pandemic. In comparison, the minimum-risk portfolios show relatively small changes in risk measure under the change in turnover constraint.

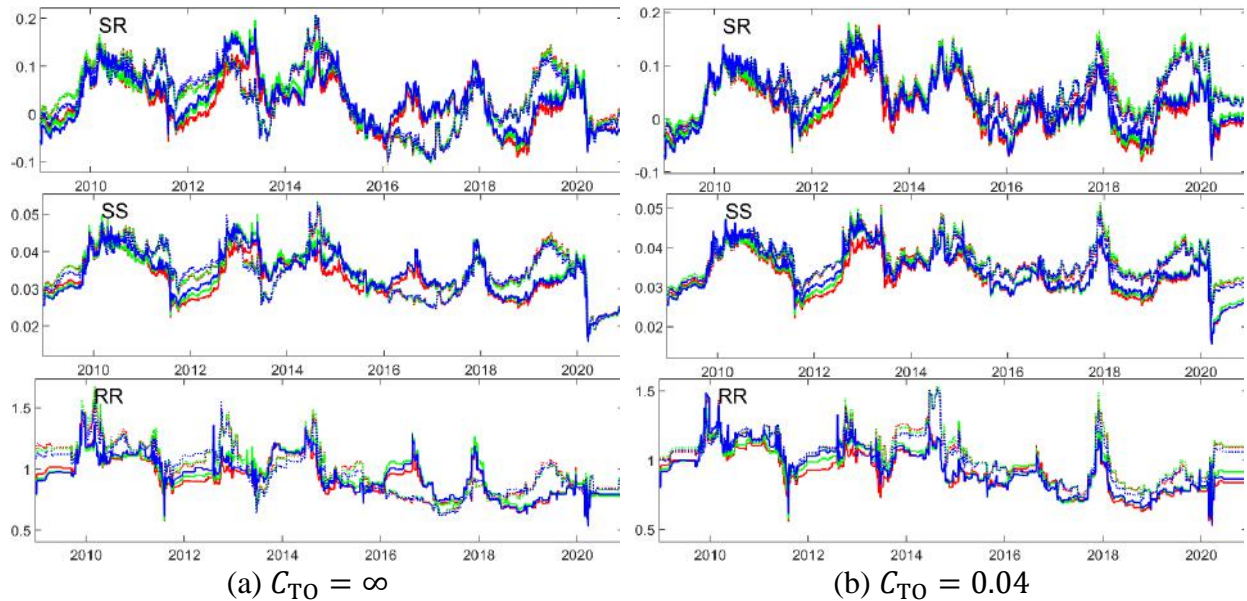


Figure 7.6 SR, SS, and RR time series (computed from one-year moving windows) for the dynamic long-only portfolio optimizations subject to (a) no turnover constraint and (b) a 4% turnover constraint.

Fig. 7.7 summarizes the total value for each performance measure over the time period 12/17/2008 through 12/18/2020 for the benchmarks and $C_{TO} = 0.04$ constrained dynamic portfolios of Figs. 7.3–7.6. SPY has the best MDD, followed by T95. The tangent portfolios have the best SR and SS values. The equal-weighted portfolio has the best RR values, followed by the benchmarks (except for SPY) and T95. The remaining two tangent portfolios have competitive RR values.

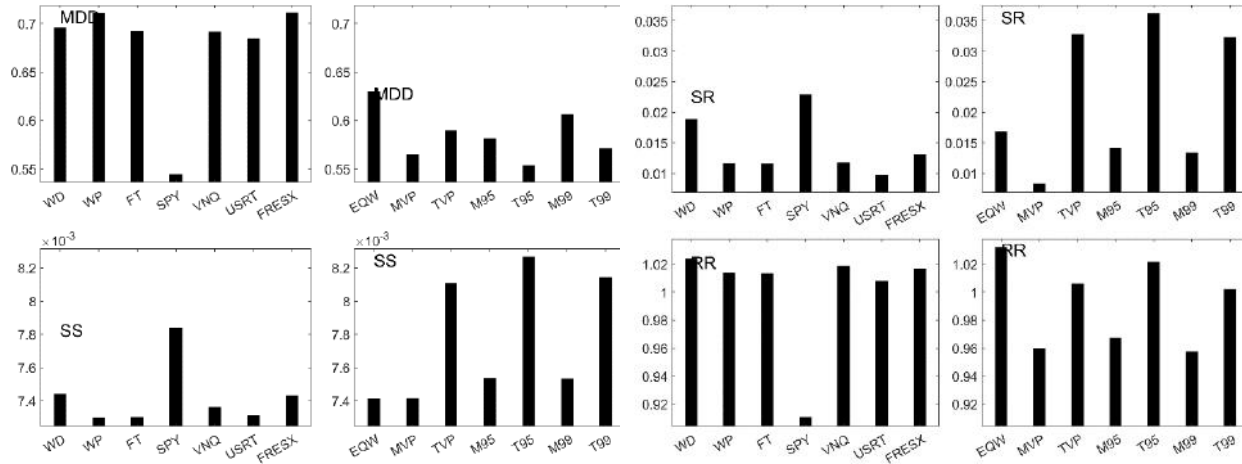


Figure 7.7 Performance measure values computed for the total period 12/17/2008 through 12/18/2020 for the benchmarks and the dynamic long-only domestic portfolio optimizations subject to a 4% turnover constraint.

Evaluating the performance of optimization methods by cumulative price and return and performance measures, such as the four used here (MDD, SR, SS, and RR), both in time series and in total values over a time period, involves dissecting a rich, multidimensional view. In such a multidimensional view, how does one determine which optimization is best? There is no single answer. Statistically, there are a number of rank correlation tests.⁶⁵ These are designed to test whether two rankings (scoring systems) of a set of items are correlated (i.e., likely to give the same ranking of scores), but they do not address the question of which item does the best over the set of different rankings. Ostensibly less sophisticated alternatives for addressing this question use methods that base a final rank upon some weighted combination of the individual rank scores for each item. These methods include

- taking the average score from each system and then ranking the averages;
- taking the median score from each system and then ranking the medians;
- taking the number of first-place scores for each item and ranking accordingly; and
- ranking (inversely) based on the number of last-place scores.

We use the first option to compare the performance of the historical and dynamic optimizations of the long-only domestic portfolio subject to no turnover constraint and to a 4% daily turnover constraint. We rank each of the six optimizations (MVP, M95, M99, TVP, T94, and T99) as well as the EQW portfolio and each of the seven benchmarks according to their scores in five ranking categories, as follows:

⁶⁵ Popular rank correlation tests include Spearman's ρ , Kendall's τ , Goodman and Kruskal's γ , and Somers' D .

7. Dynamic Portfolio Optimization

1. $P(T)/P(0)$, where $P(T)$ is the closing price on 12/18/2020 and $P(0)$ is the closing price on 12/19/2007;
2. Total values of MDD, SR, SS, and RR over the time period 12/17/2008 through 12/18/2020. Because each of these five measures is differently scaled, each measure is assigned a score from the unit interval $[0,1]$. For measure M_i for optimization/benchmark $i, i = 1, \dots, 14$, its score, m_i , is computed as

$$m_i = \frac{\left(M_i - \min_{j=1,\dots,14} (M_j)\right)}{\left(\max_{j=1,\dots,14} (M_j) - \min_{j=1,\dots,14} (M_j)\right)}, \quad (7.9)$$

with the exception of the scores for MDD, which are computed as $1 - m_i$, so that the optimization/benchmark with the lowest MDD value gets the highest numerical score.

Table 7.2 Average score over the five ranking categories for select long-only optimizations.

H-LO-NO		H-LO-0.04		D-LO-NO		D-LO-0.04	
TVP	0.86	T95	0.87	TVP	0.98	T95	0.97
T95	0.85	TVP	0.86	T95	0.97	T99	0.81
T99	0.72	T99	0.72	T99	0.73	TVP	0.80
SPY	0.52	SPY	0.50	SPY	0.58	SPY	0.47
EQW	0.43	M95	0.43	EQW	0.40	EQW	0.41
M95	0.41	EQW	0.42	WD	0.38	M95	0.36
WD	0.39	WD	0.39	M95	0.38	WD	0.35
MVP	0.37	M99	0.38	MVP	0.32	M99	0.31
M99	0.33	MVP	0.36	M99	0.30	MVP	0.28
FRESX	0.26	FRESX	0.25	FRESX	0.24	FRESX	0.25
VNQ	0.24	VNQ	0.24	VNQ	0.21	VNQ	0.25
FT	0.21	FT	0.21	FT	0.18	FT	0.22
WP	0.19	WP	0.19	WP	0.16	WP	0.20
USRT	0.19	USRT	0.19	USRT	0.16	USRT	0.20

Table 7.2 summarizes the final rankings of averaged scores over these five ranking categories for the benchmarks, the EQW portfolio, the six long-only domestic portfolios under historical optimization with no turnover constraint (H-LO-NO) and with $C_{TO} = 0.04$ (H-LO-0.04), and the six long-only domestic portfolios under dynamic optimization with no turnover constraint (D-LO-NO) and with $C_{TO} = 0.04$ (D-LO-0.04). Note that the order of the benchmarks and EQW cannot change with the optimization scheme;⁶⁶ they rank relative to each other as follows: SPY, EQW, WD, FRESX, VNQ, FT, PW, and USRT (the last two tie). The scores assigned to these eight portfolios change slightly from one optimization scheme to the next as the minimum and maximum values used in (7.9) change slightly with the optimization scheme.

⁶⁶ This is another reason why rank correlation tests are not useful for this data set. The ranking of this subsample of 8 of the 14 items is perfectly correlated from one optimization scheme to the next.

The three tangent portfolios (TVP, T95, and T99) consistently occupy the top three ranks. SPY and EQW consistently occupy ranks 4 and 5. The three minimum-risk portfolios (MVP, M95, and M99, together with WD) consistently occupy ranks 6–9. FRESX, VNQ, FT, WP, and USRT consistently occupy the bottom ranks, 10–14. T95 occupies rank 1 twice and rank 2 twice, followed by TVP (rank 1 twice, rank 2 once). The rank-1 and rank-2 averaged scores for T95 are consistently higher (by 12%–15%) in dynamic optimization compared to each historical counterpart (i.e., D-LO-NO compared with H-LO-NO and D-LO-0.04 compared with H-LO-0.04). Similarly, the averaged TVP score is 13% higher for D-LO-NO than for H-LO-NO. The dynamic optimization scores for T99 are also equal to or better than their historical counterparts. In contrast to the scores of the tangent optimizations, those of the minimum-risk optimizations decrease when they move from historical to dynamic optimization.

The strong evidence that dynamic long-only tangent optimization offers superior performance is tempered by the fact that this improvement comes with higher transaction costs (turnover). Fortunately, our results indicate that the imposition of turnover constraints as one control over such costs does not degrade performance.

7.2.2 Dynamic Jacobs et al. Long–Short Portfolios

Fig. 7.8 displays the results of the cumulative price and log-return achieved by the dynamic Jacobs et al. long–short strategy (4.2) allowing for shorting up to 10% of the total weight by any stock ($s = 0.1$) and subject to a 4% daily turnover constraint ($C_{TO} = 0.04$). Paralleling the results for dynamic long-only optimization, there is marked improvement in the performance of the tangent portfolios compared to historical Jacobs et al. long–short optimization (Fig. 4.15(a)). The box-whisker summaries of TO and L_2 -norm values in Fig. 7.9 show the same concentration of daily turnover values hitting the 4% daily limit virtually every day. A relaxation of the daily limit using the sequence $C_{TO} = \{0.04, 0.05, 0.06, 0.08, 0.12\}$ reveals that the percentage values, shown in Table 7.3, are very similar to those shown in Table 7.1 for the dynamic long-only simulations. The assignment $\mathbf{w}(t) = \mathbf{w}(t - 1)$ had to be performed on fewer than 2% of the days for any of these six, dynamic, long–short optimizations.





Figure 7.8 Cumulative price (left) and log-return (right) for the benchmarks and the dynamic Jacobs et al. long–short domestic portfolio optimizations subject to $s = 0.1$ and $C_{TO} = 0.04$.

7. Dynamic Portfolio Optimization



(b) $C_{TO} = 0.04$

Figure 7.9 Box-whisker summary statistics of TO (4.5) and the L_2 -norm (4.7) for the dynamic Jacobs et al. long-short domestic portfolio optimizations subject to $s = 0.1$ and $C_{TO} = 0.04$.

Table 7.3 Percent of daily optimizations succeeding at each range of turnover constraint.

	[0,0.4]	(0.4,0.12]	Failed
MVP	100.0	0.0	0.0
TVP	99.8	0.2	0.0
M95	95.1	3.5	1.4
T95	96.1	2.0	1.9
M99	95.4	3.4	1.2
T99	96.1	2.3	1.6

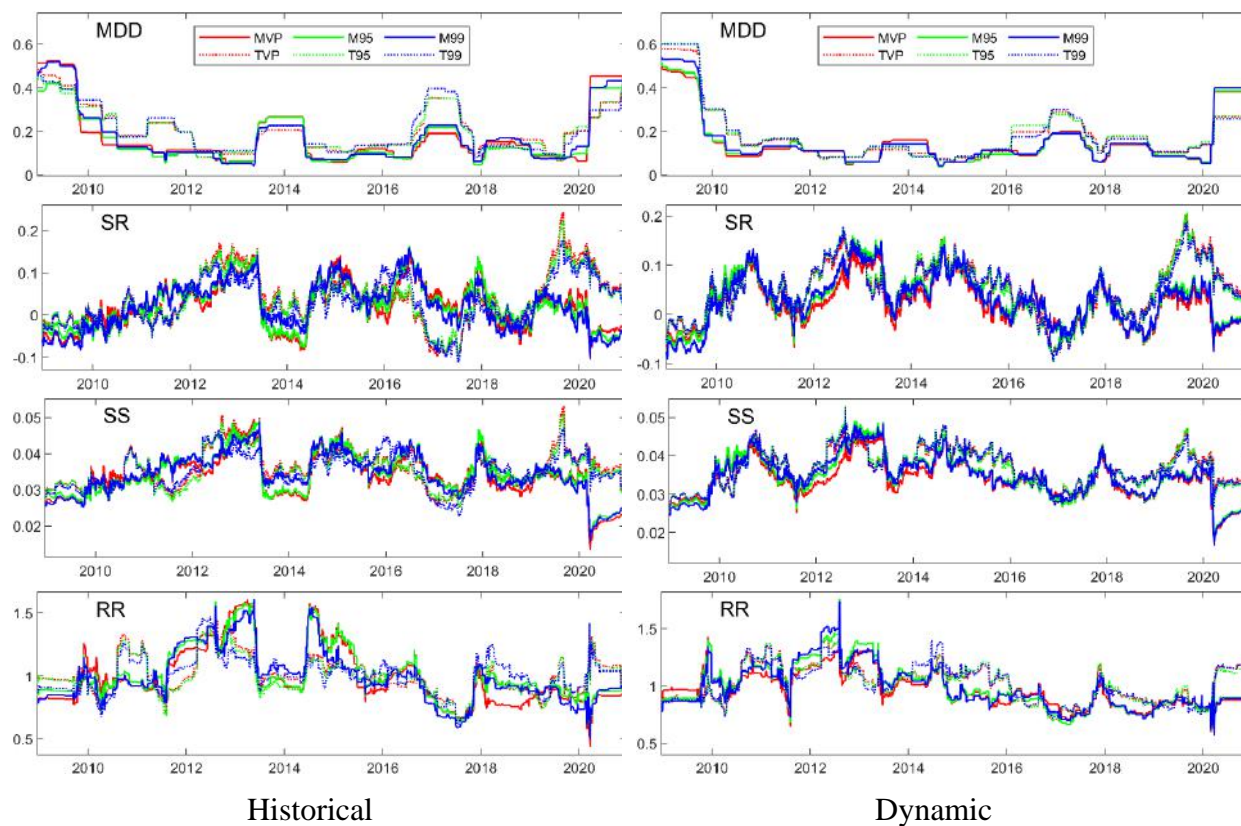


Figure 7.10 Comparison of the performance-measure time series for the historical (left) and dynamic (right) Jacobs et al. long–short domestic portfolio optimizations subject to $s = 0.1$ and $C_{TO} = 0.04$.

Fig. 7.10 compares the time series for the performance measures for the historical and dynamic Jacobs et al. long–short optimizations of the domestic portfolio. After the Great Recession period, and even into the pandemic period, there is clear improvement in MDD values for the tangent portfolios under dynamic optimization. Improvements in SR and SS are also evident, but RR improvement is questionable. We utilize the average-of-ranking-scores approach used in section 7.2.1 to evaluate the overall rankings of the historical and dynamic Jacobs et al. long–short optimizations of the domestic portfolio subject to $s = 0.1$ and $C_{TO} = 0.04$. Table 7.4 presents the results. As in the long-only case, the three tangent portfolios occupy the top three ranks. However, in the dynamic optimization, these tangent portfolios usually exhibit a decrease in average score. SPY and EQW also occupy ranks 4 and 5. The remaining rankings differ from those in the long-only case. WD and M95 occupy ranks 6 and 7. Other rankings are less consistent between historical and dynamic optimization. Interestingly, MPV and M99 occupy the lowest ranks in the historical case, but their rankings move up (particularly for MVP) in the dynamic case.

Table 7.4 Average scores over the five ranking categories for select Jacob et al. long–short optimizations.

	H-LS-0.04		D-LS-0.04
TVP	0.783	TVP	0.779
T95	0.769	T99	0.711
T99	0.693	T95	0.666
SPY	0.660	SPY	0.531
EQW	0.561	EQW	0.461
WD	0.555	WD	0.427
M95	0.404	M95	0.352
FRESX	0.387	MVP	0.311
VNQ	0.364	FRESX	0.298
FT	0.331	VNQ	0.286
WP	0.315	FT	0.257
USRT	0.302	M99	0.255
M99	0.238	WP	0.239
MVP	0.221	USRT	0.236

7.2.3 Dynamic Lo–Patel Long–Short Portfolios

We consider dynamic optimization of the domestic portfolio using the Lo–Patel long–short method subject to $lev = 0.1$. This dynamic optimization shows the greatest cumulative price/return improvement compared to the corresponding historical optimization. A comparison of Fig. 7.11 with Fig. 4.4(b) shows a cumulative price improvement of a factor of four for the tangent

7. Dynamic Portfolio Optimization




Figure 7.11 Cumulative price (left) and log-return (right) for the dynamic Lo–Patel long–short domestic portfolio optimizations subject to $lev = 0.1$.

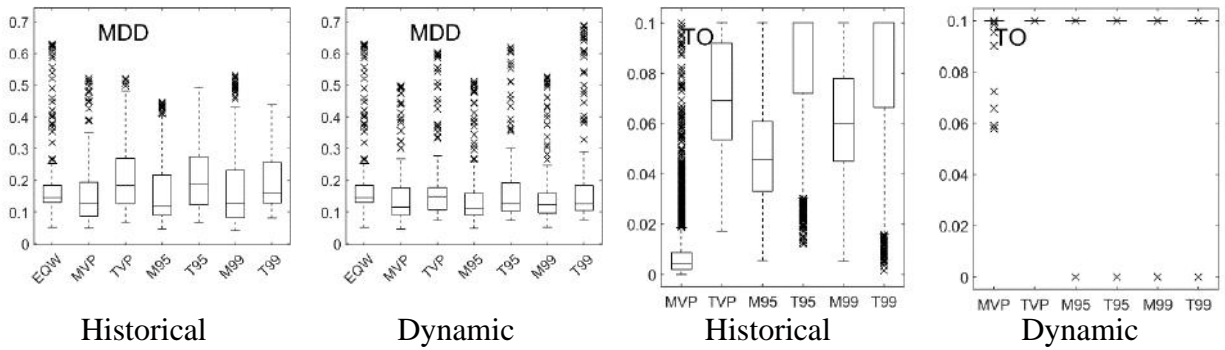


Figure 7.12 Box-whisker summaries of the MDD and TO distributions for the historical and dynamic Lo–Patel long–short domestic portfolio optimizations subject to $lev = 0.1$.




Figure 7.13 MDD time series of the historical (left) and dynamic (right) Lo–Patel long–short domestic portfolio optimizations subject to $lev = 0.1$.

portfolios. The MDD box-whisker summaries in Fig. 7.12 show a substantial lowering of Q3 values for the tangent portfolios under dynamic optimization compared to historical optimization. The time series in Fig. 7.13 show the general improvement in MDD values for the dynamic tangent portfolios after 2010.

One advantage of the Lo–Patel method is that it naturally imposes a turnover constraint through the value of lev . Fig. 7.12 also compares the historical and dynamic box-whisker summaries of the TO distribution. Virtually every day, the dynamic optimization requires a 10% turnover of asset weights. Days with zero turnover correspond to days on which no optimized weight solution that would satisfy the $lev = 0.1$ constraint can be found (in which case the assignment $\mathbf{w}(t) = \mathbf{w}(t - 1)$ is made). Table 7.5 summarizes the frequency of occurrence of these assignments. Dynamic T95 and T99 optimizations have the lowest success rate, at just under 95%.

Table 7.5 Percent of daily Lo–Patel optimizations succeeding with the $lev = 0.1$ constraint.

	Historical	Dynamic
MVP	100.0	100.0
TVP	100.0	100.0
M95	100.0	97.8
T95	100.0	94.7
M99	100.0	97.4
T99	100.0	94.9

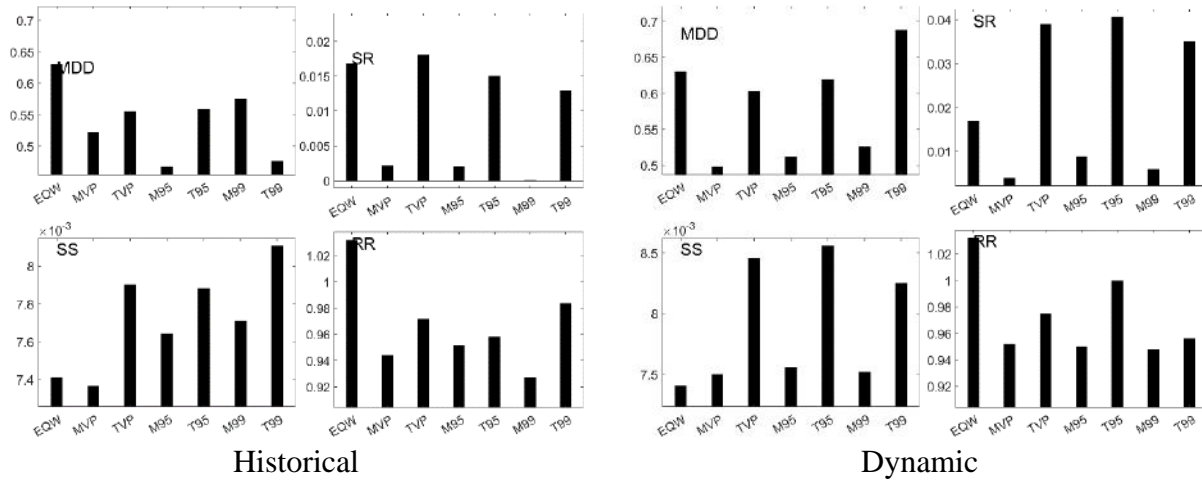


Figure 7.14 Total-time-period performance-measure values for EQW and the historical (left) and dynamic (right) Lo–Patel long–short domestic portfolio optimizations subject to $lev = 0.1$.

Table 7.6 Average scores over the five ranking categories for select Jacob et al. long–short optimizations.

	H-LP-0.1		D-LP-0.1
SPY	0.67	T95	0.83
T95	0.66	TVP	0.60
T99	0.63	SPY	0.55
EQW	0.55	EQW	0.48
WD	0.55	WD	0.46
TVP	0.54	MVP	0.44
MVP	0.38	T99	0.42
M95	0.38	M95	0.33
FRESX	0.36	M99	0.30
VNQ	0.34	FRESX	0.29
FT	0.30	VNQ	0.28
WP	0.29	FT	0.25
USRT	0.27	WP	0.23
M99	0.25	USRT	0.23

Fig. 7.14 compares the total-time-period performance-measure values for EQW and the historical and dynamic Lo–Patel optimizations of the domestic portfolio. Moving from historical to dynamic optimization, the overall MDD values for the tangent portfolios worsen, but their overall SR and SS values show strong improvement. Table 7.6 compares the average-ranking scores for the benchmarks, EQW, and the historical (H-LP-0.1) and dynamic (D-LP-0.1) Lo–Patel optimizations. Interestingly, the benchmark SPY is now competitive with two of the tangent portfolios under Lo–Patel optimization. Whereas SPY ranks first under historical optimization, it ranks (significantly) lower than TVP and T95 under dynamic optimization. Whereas M99 ranks the lowest among historically optimized portfolios, its rank moves up five places under dynamic optimization. Conversely, T99 suffers a strong drop in ranking, from 3 to 7, between H-LP-0.1 and D-LP-0.1.

7.3 Dynamic Optimization with the Black–Litterman Model

We consider the effect of adding the Black–Litterman approach to the dynamic approach. As noted in section 4.5, we run our prototype portfolios without the specification of analyst views. In this case, the primary result of adding the Black–Litterman approach is to include the effect of an influencing market benchmark. As in section 4.5, we use the returns of the real-estate-oriented mutual fund FRESX for this purpose. The Black–Litterman model is therefore implemented as discussed in section 4.5 and added to the dynamic approach sketched in Fig. 7.2 and discussed in section 7.1.4.

We consider the BL MVP and BL TVP long-only optimizations of the domestic portfolio with a 4% turnover constraint. Fig. 7.15 shows the cumulative price and return performance. Compared with the historical optimization (Fig. 6.2(b)), there is a noticeable improvement in the TVP return. Fig. 7.16 compares the box-whisker summaries of the TO and L_2 -norm distributions of these two dynamic portfolios to those of their historical counterparts. Again, it is evident that the large number of samples generated by the dynamic simulation forces the majority daily turnover to reach the 4% limit, in contrast to the historical simulation, where the majority daily turnover is $< 0.5\%$ and $< 2\%$ for MVP and TVP, respectively. Fig. 7.17 shows the total-time-period performance measures. For BL TVP, MDD and SS improved as they moved from historical to dynamic optimization, whereas SR experienced essentially no change and RR dropped. For BL MVP, all the performance measures worsened under dynamic optimization.




Figure 7.15 Cumulative price (left) and log-return (right) for FRESX and the dynamic BL MVP

and BL TVP long-only domestic portfolio optimizations subject to a 4% turnover constraint.

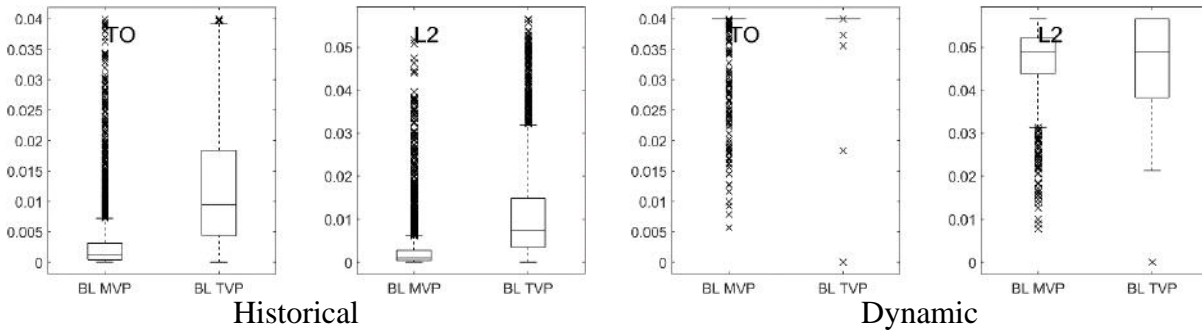


Figure 7.16 Box-whisker summaries of TO and the L_2 -norm for the historical (left) and dynamic (right) BL MVP and BL TVP long-only domestic portfolio optimizations subject to a 4% turnover constraint.




Figure 7.17 Total-time-period performance-measure values for FRESX and the historical (left) and dynamic (right) BL MVP and BL TVP long-only domestic portfolios optimizations subject to a 4% turnover constraint.

Table 7.7 compares the average-ranking scores for FRESX and these historical (H-BL-0.04) and dynamic (D-BL-0.04) Black–Litterman optimizations. Consistent with prior results, BL TVP outranks FRESX. Interestingly, BL MVP is top ranked under historical optimization but drops to rank 3 under dynamic optimization (with a very low average score).

Table 7.7 Average scores over the five ranking categories for select Black–Litterman optimizations.

	H-BL-0.040		D-BL-0.04
MVP	0.45	TVP	0.51
TVP	0.37	FRESX	0.31
FRESX	0.34	MVP	0.14

References

- Baillie, R.T., Bollerslev, T. & Mikkelsen, H.O. (1996). Fractionally integrated generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 74(1), 3–30.
- Bollersley, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31(1), 307–327.

7. Dynamic Portfolio Optimization

- Engle, R. F. (1982). Autoregressive conditional heteroskedasticity with estimates of the variance of U.K. inflation. *Econometrica*, 50, 987–1008.
- Engle, R.F. & Bollersley, T. (1986). Modeling the persistence of conditional variances. *Econometric Reviews*, 5, 1–50.
- Engle, R.F., Lilien, D. & R. Robins, R. (1986). Estimating time varying risk premia in the term structure: The ARCH-M model. *Econometrica*, 55, 391–407.
- Glosten, L., Jagannathan, R. & Runkle, D. (1993). Relationship between the expected value and volatility of the nominal excess returns on stocks. *Journal of Finance*, 48, 1779–1802.
- Nelson, D. B. (1991). Conditional heteroskedasticity in asset returns: A new approach. *Econometrica*, 59(2), 347–370.
- Tsay, R. S. (2010). *Analysis of financial time series*. John Wiley & Sons, Hoboken, NJ.
- Zakoian, J.-M. (1994). Threshold heteroskedastic models. *Journal of Economic Dynamics and Control*, 18, 931–955.

Chapter 8

Backtesting

Under regulatory guidelines, banks with substantial trading activity are required to set aside capital (the market-risk capital requirement) to insure against extreme portfolio loss. The size of the capital requirement is determined by the VaR of the portfolio. The VaR at confidence level $1 - \alpha$, $\text{VaR}_\alpha(x)$ is defined in (3.21). If x is profit/loss measured in dollars, then a daily $\text{VaR}_{0.05}$ value of \$1 million means there is a 5% chance that a portfolio will lose \$1 million or more during a one-day period (assuming the portfolio undergoes no changes over the day). If x is return values, then a daily $\text{VaR}_{0.05}$ value of 0.017 means there is a 5% chance that the daily return of the portfolio will be negative, with a magnitude exceeding 1.7%. As of 2005 (Campbell, 2005), banks have been required to report their 1% VaR over a 10-day horizon, $\overline{\text{VaR}}_{0.01}(t)$. The market-risk capital requirement was then set as the larger of the current horizon estimate and a multiple, S_t , of the average estimate over the previous 60 trading days (backtesting):

$$\text{MRC}(t) = \max \left(\overline{\text{VaR}}_{0.01}(t), S(t) \frac{1}{60} \sum_{i=0}^{59} \overline{\text{VaR}}_{0.01}(t-1) \right) + c.$$

The multiplying factor, $S(t)$, is itself determined by backtesting by classifying the number of 1% VaR violations, x , in the previous 250 trading days.

Table 8.1 Basel II guidelines for $S(t)$.

Zone	x	$S(t) - 3$	$F_B(X \leq x 250, 0.01)$
Green	0	0.00	0.0811
	1	0.00	0.2858
	2	0.00	0.5432
	3	0.00	0.7581
	4	0.00	0.8922
Yellow	5	0.40	0.9588
	6	0.50	0.9863
	7	0.65	0.9960
	8	0.75	0.9989
	9	0.85	0.9997
Red	≥ 10	1.00	0.9999

The last column of Table 8.1 gives the binomial cumulative probability distribution for the occurrence of $\leq x$ 1% VaR violations over 250 days. These cumulative probabilities determine the boundaries of the Basel II green, yellow, and red (traffic light) zones.

Risk-based capital requirements depend on the performance of the portfolio and the accuracy of the VaR model, which are coupled to each other. That is, portfolio daily returns are viewed as a random variable, and at time t , the distribution governing these daily returns will be characterized by a VaR at a desired confidence level (e.g., 99%). As Christoffersen (1998) points out, if the management of the portfolio is responding adequately to market-risk conditions and if the VaR estimate is sufficiently accurate, the sequence of VaR violations should satisfy two distinct properties:

8. Backtesting

1. Unconditional coverage: The probability of realizing a violation (i.e., of exceeding $\text{VaR}_\alpha(t)$) should be α . Thus, the unconditional-coverage property restricts how often violations should occur.
2. Independence: The probability of a violation occurring on day t should be independent of when any previous violation occurred; that is, the previous history of violations should not affect their future occurrences. Because market volatility is known to occur in clusters, the requirement of independence is a condition not only on the accuracy of the VaR model but also on the rapid risk-adjusted active management of the portfolio.

A number of tests have been devised to examine one or both of these properties; these are described in section 8.1. In this chapter, we concentrate on VaR backtesting and demonstrate that dynamic simulation (Chapter 7) can improve VaR backtest results over those of historical simulation (Chapters 4–6). For the Basel II convention, testing based on $\text{VaR}_{0.05}$ is required; for Basel III, testing based on $\text{VaR}_{0.01}$ is required. In addition to considering the 95% and 99% quantile levels, we also consider testing based on the 99.5% quantile level ($\text{VaR}_{0.005}$).

As discussed above, market-risk capital requirements require an accurate estimation of VaR as well as tests of the portfolio's performance measured by VaR. The core binomial question to be answered is, Does today's portfolio return $-r_p(t)$ exceed today's (projected) value of $\text{VaR}_\alpha(r_p, t)$ at quantile level $1 - \alpha$? Such an occurrence is viewed as a “failure” (violation). Because the portfolio return, $r_p(t)$, is known by the end of the day, the key problem in constructing a backtest is the estimation of the value of $\text{VaR}_\alpha(r_p, t)$ relevant for day t . There are a number of methods for performing this estimation, all of which, in one form or another,⁶⁷ require knowledge of the portfolio's performance over the previous τ days.

- For the historical optimizations discussed in Chapters 4–6, we compute $\text{VaR}_\alpha(r_p, t)$ empirically, based on synthetic portfolio returns for the previous days $t - \tau, \dots, t - 1$. Let $r_{i,k}$, $k = t - \tau, \dots, t - 1$; $i = 1, \dots, n$ denote the individual-asset returns over this period, and let $w_i(t)$, $i = 1, \dots, n$ denote the optimized weights used for day t . Then, the set of returns, $r_{p,k} = \sum_{i=1}^n w_i(t) r_{i,k}$, $k = t - \tau, \dots, t - 1$, is representative of returns expected for the portfolio on day t . We estimate values of $\text{VaR}_\alpha(r_p, t)$ based on the distribution of the synthetic values $r_{p,k}$, $k = t - \tau, \dots, t - 1$.
- For the dynamic optimizations discussed in Chapter 7, recall that we compute a very large sample of dynamic returns $\{r_{i,s,t}, s = 1, \dots, S\}$ for day t for each of the $i = 1, \dots, n$ assets in the portfolio. Applying the actual weights $w_i(t)$, $i = 1, \dots, n$ used for the optimized portfolio on day t to the sample of dynamic asset returns gives a sample of “predicted” portfolio returns $r_{p,s,t} = \sum_{i=1}^n w_i(t) r_{i,s,t}$, $s = 1, \dots, S$. We estimate the values of $\text{VaR}_\alpha(r_p, t)$ based on the distribution of the values $r_{p,s,t}$, $s = 1, \dots, S$.

8.1 VaR Tests

⁶⁷ For example, prior portfolio returns could be assumed to follow a particular distribution, such as Gaussian. The previous $t - T, \dots, t - 1$ daily returns would then be fit to said distribution, and $\text{VaR}_\alpha(r_p, t)$ would then be computed using the distribution parameters computed from the fit.

Assume individual-asset-return data is available for days $1, \dots, T$. Because our optimization methods use a moving window of τ days to compute optimized portfolio-weight values, the backtest data set consists of $N = T - \tau$ values (days $\tau + 1, \dots, T$) of observed portfolio returns and estimated values-at-risk or, equivalently, a time series of N values, each value being either “success” or “failure.” Standard backtesting involves the following tests on this data set.

8.1.1 Binomial Test

For each of the N days, either the (negative of the) portfolio return exceeds the $\text{VaR}_\alpha(r_p, t)$ value (a failure) or it does not (a success). The number of failures should follow a binomial distribution with α being the probability of failure. Thus, we expect to see $N\alpha$ failures, with standard deviation $\sqrt{N\alpha(1 - \alpha)}$. Let x denote the observed number of failures. Using the z -score

$$z(x) = \frac{x - N\alpha}{\sqrt{N\alpha(1 - \alpha)}} \quad (8.1)$$

as the test statistic, and assuming N and $N\alpha$ and $N(1 - \alpha)$ are sufficiently large,⁶⁸ the z -score is approximately $\mathcal{N}(0,1)$, where \mathcal{N} denotes the Gaussian (normal) distribution. Thus, the (tail) probability that this z -score is exceeded is $1 - F_G(z(x))$, where $F_G(\cdot)$ is the standard Gaussian cumulative probability distribution. Because the binomial test (BIN) is a two-sided test, the p -value, $p_{\text{BIN}}(z(x))$, is twice the tail probability. The binomial test accepts the null hypothesis (that the probability of failure is α) if

$$p_{\text{BIN}}(z(x)) = 2 [1 - F_G(z(x))] > p_{\text{test}}, \quad (8.2)$$

where $1 - p_{\text{test}}$ is the test confidence level, usually set to 0.95; that is, $p_{\text{test}} = 0.05$. (If $p_{\text{BIN}}(z(x)) < p_{\text{test}}$, then $z(x)$ is further out in the distribution tail than the value determined by $F_G^{-1}(p_{\text{test}})$.)

8.1.2 Traffic Light Test

The Basel framework (1996) for backtesting defines the traffic light test, which is again based on a binomial distribution. As in section 8.1.1, N is the number of observations, α is the probability of failure, and x is the observed number of failures. The probability of observing up to x failures in N “tries” is

$$P(X \leq x|N, \alpha) = F_B(X \leq x|N, \alpha), \quad (8.3)$$

where $F_B(x|N, \alpha)$ is the binomial cumulative probability distribution with parameters N and α . The Basel framework defines three zones:

⁶⁸ As long as $\alpha \neq 0$ or 1 , then $N\alpha = \mathcal{O}(N)$ and $N(1 - \alpha) = \mathcal{O}(N)$ for sufficiently large N . In practice, it is sufficient to ensure $\min(N\alpha, N(1 - \alpha)) \gtrsim 5$, which requires $N \geq 5 \max(1/\alpha, 1/(1 - \alpha))$.

$$\begin{aligned}
&\text{green,} & F_B(X \leq x|N, \alpha) \leq 0.95, \\
&\text{yellow,} & 0.95 < F_B(X \leq x|N, \alpha) \leq 0.9999, \\
&\text{red,} & 0.9999 < F_B(X \leq x|N, \alpha).
\end{aligned} \tag{8.4}$$

The probability of a false positive (type 1 error or rejection of a true null hypothesis) is $P(X \geq x|N, \alpha) = F_B(X \geq x|N, \alpha)$.⁶⁹ The tests (8.4) are known collectively as the traffic light (TL) test.

The Basel framework also includes a multiplication factor to guide the size of the potential increase in a firm's capital requirements (relative to a baseline) resulting from the traffic light test. This is reported as a scaling factor, s , which is 0 for the green zone and 1 for the red zone and increases over the yellow zone based on the relative difference between the assumed VaR quantile level, $1 - \alpha$, and the observed quantile level, $1 - x/N$. Assuming a standard Gaussian distribution, the z -score (z_{assumed}) corresponding to the cumulative value $1 - \alpha$ and the z -score (z_{observed}) corresponding to the cumulative value $1 - x/N$ can be computed. The scaling factor, s_{yellow} , for the yellow zone is then computed as

$$s_{\text{yellow}} = 3 \left(\frac{z_{\text{assumed}} - z_{\text{observed}}}{z_{\text{observed}}} \right), \text{ s. t. } 0 \leq s_{\text{yellow}} \leq 1. \tag{8.5}$$

8.1.3 Kupiec's Tests

Kupiec (1995) introduced two tests, proportion of failures (PoF) and time until first failure (TUFF). The PoF test assesses whether x/N is consistent with α . If the number of failures in N tries follows a binomial distribution with probability of failure α , then the probability of seeing x failures in N tries is

$$P(k|N, \alpha) = \binom{N}{x} \alpha^x (1 - \alpha)^{N-x}.$$

The PoF test assesses whether x/N is consistent with α by testing whether the likelihood ratio

$$R_{\text{PoF}} = \frac{\binom{N}{x} \alpha^x (1 - \alpha)^{N-x}}{\binom{N}{x} \left(\frac{x}{N}\right)^x \left(1 - \left(\frac{x}{N}\right)\right)^{N-x}} = \frac{\alpha^x (1 - \alpha)^{N-x}}{\left(\frac{x}{N}\right)^x \left(1 - \left(\frac{x}{N}\right)\right)^{N-x}}$$

is significantly different from 1 or, equivalently, whether the logarithm of this ratio (the log-likelihood ratio),

$$LR_{\text{PoF}} = -2 \ln \left(\frac{\alpha^x (1 - \alpha)^{N-x}}{\left(\frac{x}{N}\right)^x \left(1 - \frac{x}{N}\right)^{N-x}} \right) = -2 \left[x \ln \left(\frac{N\alpha}{x} \right) + (N - x) \ln \left(\frac{N(1 - \alpha)}{N - x} \right) \right], \tag{8.6}$$

is sufficiently different from 0. The test statistic LR_{PoF} has the limiting values

⁶⁹ Note that $P(X \leq x|N, \alpha) + P(X \geq x|N, \alpha) = 1 + P(X = x|N, \alpha)$.

$$\begin{aligned}\lim_{x \rightarrow 0} LR_{\text{PoF}} &= -2N \ln(1 - \alpha), \\ \lim_{x \rightarrow N} LR_{\text{PoF}} &= -2N \ln(\alpha).\end{aligned}$$

For large N , this statistic follows a chi-square distribution with one degree of freedom. The p -value, p_{PoF} , of the PoF test is the tail probability for values that exceed LR_{PoF} , that is,

$$p_{\text{PoF}} = 1 - F_{\chi_1^2}(LR_{\text{PoF}}), \quad (8.7)$$

where $F_{\chi_1^2}(\cdot)$ is the cumulative distribution for a chi-square distribution with one degree of freedom. The PoF test accepts the null hypothesis (that x/N is consistent with α) if

$$p_{\text{PoF}} > p_{\text{test}}, \quad (8.8)$$

where $1 - p_{\text{test}}$ is a test level set by the user, typically chosen to be 0.95 ($p_{\text{test}} = 0.05$).

In an identical fashion to PoF, the TUFF test assesses whether the number of trading days, n_1 , until first failure is consistent with α . The test again employs the log-likelihood-ratio test statistic

$$LR_{\text{TUFF}} = -2 \ln \left(\frac{\alpha(1 - \alpha)^{n_1 - 1}}{\left(\frac{1}{n_1}\right) \left(1 - \frac{1}{n_1}\right)^{n_1 - 1}} \right) = -2 \left[\ln(n_1 \alpha) + (n_1 - 1) \ln \left(\frac{n_1(1 - \alpha)}{(n_1 - 1)} \right) \right]. \quad (8.9)$$

Equation (8.9) has the limiting value $\lim_{n_1 \rightarrow 1} LR_{\text{TUFF}} = -2 \ln(\alpha)$. For large N , this statistic also follows a chi-square distribution with one degree of freedom. Thus, its p -value, p_{TUFF} , is given by

$$p_{\text{TUFF}} = 1 - F_{\chi_1^2}(LR_{\text{TUFF}}), \quad (8.10)$$

and the TUFF test accepts the null hypothesis (that n_1 is consistent with α) if

$$p_{\text{TUFF}} > p_{\text{test}}. \quad (8.11)$$

Note that LR_{TUFF} is undefined if there are no observed failures ($n_1 = 0$). If this happens, two scenarios are considered:

1. If $N > 1/\alpha$ and if the TUFF test fails when $n_1 = N + 1$ is considered (i.e., for the earliest possible value that n_1 could have, given that no failure occurred in N days), then the test rejects the null hypothesis (and values for LR_{TUFF} and p_{TUFF} are reported for $n_1 = N + 1$).
2. Otherwise, it is impossible for the TUFF test to accept or reject the null hypothesis.

8.1.4 Christoffersen's Tests

Christoffersen (1998) introduced the conditional coverage independence (CCI) test, which measures whether the probability of observing a failure on day t depends on observing a failure on the previous day. Thus, the test considers all possible pairs of days, $(t - 1, t)$, and accumulates the following values:

- n_{00} : number of pairs in which no failure occurred on both days;
- n_{10} : number of pairs in which a failure occurred on day $t - 1$ followed by no failure on day t ;
- n_{01} : number of pairs in which no failure occurred on day $t - 1$ followed by a failure on day t ;
- n_{11} : number of pairs in which a failure occurred on both days.

8. Backtesting

From these values, the following probabilities are estimated:

- $\pi_0 = n_{01} / (n_{00} + n_{01})$: conditional probability of a failure on the second day given no failure on the first;
- $\pi_1 = n_{11} / (n_{10} + n_{11})$: conditional probability of a failure on the second day given a failure on the first;
- $\pi = (n_{01} + n_{11}) / (n_{00} + n_{01} + n_{10} + n_{11})$: unconditional probability of a failure on the second day.

The CCI test employs a log-likelihood ratio as a test statistic:

$$\begin{aligned} LR_{CCI} &= -2 \ln \left(\frac{\pi^{n_{01}+n_{11}} (1-\pi)^{n_{00}+n_{10}}}{\pi_0^{n_{01}} (1-\pi_0)^{n_{00}} \pi_1^{n_{11}} (1-\pi_1)^{n_{10}}} \right) \\ &= -2 \ln(\pi^{n_{01}+n_{11}} (1-\pi)^{n_{00}+n_{10}}) + 2 \ln(\pi_0^{n_{01}} (1-\pi_0)^{n_{00}}) \\ &\quad + 2 \ln(\pi_1^{n_{11}} (1-\pi_1)^{n_{10}}). \end{aligned} \quad (8.12)$$

Again, for large N , this statistic follows a chi-square distribution with one degree of freedom and has a p -value, p_{CCI} , given by

$$p_{CCI} = 1 - F_{\chi^2_1}(LR_{CCI}). \quad (8.13)$$

The CCI test accepts the null hypothesis (that the probability of failure on day t is independent of the probability of failure on day $t - 1$) if

$$p_{CCI} > p_{\text{test}}. \quad (8.14)$$

The test level, $1 - p_{\text{test}}$, is typically chosen to be 0.95.

As shown in (8.12), the test statistic LR_{CCI} is the sum of three log-likelihood functions, each of the form

$$L = \ln(p^{n_a}(1-p)^{n_b}), \quad (8.15)$$

where p , n_a , and n_b are functions of the n_{ij} ; $i = 0, 1$; $j = 0, 1$. Consider, for example,

$$L = \ln(\pi_0^{n_{01}} (1-\pi_0)^{n_{00}}) = \ln \left[\left(\frac{n_{01}}{n_{00} + n_{01}} \right)^{n_{01}} \left(\frac{n_{00}}{n_{00} + n_{01}} \right)^{n_{00}} \right].$$

This term has the property $\lim_{n_{00} \rightarrow 0} L = \lim_{n_{01} \rightarrow 0} L = 0$. This property holds for all three terms of the form (8.15) in (8.12). Thus, if any of n_{00}, n_{01}, n_{10} , or n_{11} are zero, the corresponding log-likelihood functions in (8.12) are set to 0, and (8.12) remains well-defined.

Kupiec's PoF test can be combined with Christoffersen's CCI test to get the conditional coverage (CC) mixed⁷⁰ test having the test statistic

$$LR_{CC} = LR_{\text{PoF}} + LR_{CCI}, \quad (8.16)$$

which for large N follows a chi-square distribution with two degrees of freedom. This statistic has a p -value, p_{CC} , given by

$$p_{CC} = 1 - F_{\chi^2_2}(LR_{CC}), \quad (8.17)$$

⁷⁰ The CC test is "mixed" because it combines a frequency test (PoF) and an independence test (CCI).

where $F_{\chi^2_2}(\cdot)$ is the cumulative probability distribution for a chi-square distribution with two degrees of freedom. The CC test is accepted if

$$p_{CC} > p_{\text{test}}. \quad (8.18)$$

8.1.5 Haas's Tests

Haas (2001) extended the TUFF test to incorporate the time (number of trading days) between each successive failure, $i - 1$ and i (using TUFF for the case $i = 1$). This is referred to as the time between failures independence (TBFI) test. Let x be the number of failures and $n_i, i = 1, \dots, x$ denote the number of days between failure $i - 1$ and i . (Here, n_1 is the time to first failure as defined in section 8.1.3.) Using time between failures as a metric, the TBFI test determines whether the failures are independent of each other. A log-likelihood based on the TUFF formula (8.9) is defined for each n_i :

$$LR_{\text{TBFI},i} = -2 \ln \left(\frac{\alpha(1-\alpha)^{n_i-1}}{\left(\frac{1}{n_i}\right)\left(1-\frac{1}{n_i}\right)^{n_i-1}} \right) = -2 \left[\ln(n_i\alpha) + (n_i-1) \ln \left(\frac{n_i(1-\alpha)}{(n_i-1)} \right) \right]. \quad (8.19)$$

Assuming failures are independent, an overall log-likelihood ratio is then the sum of the individual log-likelihood ratios,

$$LR_{\text{TBFI}} = \sum_{i=1}^x LR_{\text{TBFI},i}, \quad (8.20)$$

which for large N follows a chi-square distribution with x degrees of freedom. Consequently, the p -value, p_{TBFI} , of the test statistic is

$$p_{\text{TBFI}} = 1 - F_{\chi^2_x}(LR_{\text{TBFI}}), \quad (8.21)$$

where $F_{\chi^2_x}(\cdot)$ is the cumulative probability distribution for a chi-square distribution with x degrees of freedom. The TBFI test is accepted if

$$p_{\text{TBFI}} > p_{\text{test}}. \quad (8.22)$$

If no failures are observed ($x = 0$), the scenarios discussed in the TUFF test apply.

Kupiec's PoF test can be combined with Haas's TBFI test to get the time between failures (TBF) mixed test having the test statistic

$$LR_{\text{TBF}} = LR_{\text{PoF}} + LR_{\text{TBFI}}, \quad (8.23)$$

which for large N follows a chi-square distribution with $x + 1$ degrees of freedom. This statistic has a p -value, p_{TBF} , given by

$$p_{\text{TBF}} = 1 - F_{\chi^2_{x+1}}(LR_{\text{TBF}}). \quad (8.24)$$

The TBF test is accepted if

$$p_{\text{TBF}} > p_{\text{test}}. \quad (8.25)$$

If there are no failures, the TBF test result is “to accept” only if the separate PoF and TBF test results are “to accept”.

Campbell (2005) provides a review of various VaR tests. Of these eight tests, the binomial, traffic light, and Kupiec’s PoF and TUFF are unconditional-coverage tests. Christoffersen’s CCI and Haas’s TBF are independence tests. The CC and TBF tests are attempts to measure both unconditional coverage and independence under one test. For the backtest results discussed in section 8.2, we use the value $p_{\text{test}} = 0.05$ as the basis for accepting or rejecting each test (except for TL).

8.2 Backtest Results

8.2.1 Historical Optimization

We first discuss typical backtesting results for the historically optimized portfolios of Chapters 4 and 5. Fig. 8.1 shows the time series of daily returns, $r_p(t)$, versus value-at-risk projections, $\text{VaR}_{1-\alpha}(r_p, t)$, for $\alpha = 0.05, 0.01$, and 0.005 for the TVP long-only domestic portfolio subject to a 4% turnover constraint. (See Figs. 4.9–4.12 for the other performance and risk measures for this portfolio.) The $\text{VaR}_{1-\alpha}(r_p, t)$ is computed using an eight-year (2,016 trading days) rolling window.




Figure 8.1 Daily returns, $r_p(t)$, versus value-at-risk projections, $\text{VaR}_{1-\alpha}(r_p, t)$, for $\alpha = 0.05, 0.01$, and 0.005 for the TVP long-only domestic portfolio optimizations subject to a 4% turnover constraint.

Table 8.2 summarizes the failure data for this portfolio. The variables defining each column are defined in sections 8.1.1–8.1.3. The number of observed failures, x , exceeds the number of expected failures, $N\alpha$, for all three values of α , by ratios ranging from 1.16 to 2.08.

Table 8.2 Failure statistics for the TVP long-only domestic portfolio.

$1 - \alpha$	$1 - x/N$	N	x	$N\alpha$	$x/(N\alpha)$
0.95	0.942	3273	190	163.6	1.16
0.99	0.98	3273	61	32.7	1.86
0.995	0.987	3273	34	16.4	2.08
	n_1	n_{00}	n_{10}	n_{01}	n_{11}
0.95	5	2925	157	158	32
0.99	10	3160	51	51	10
0.995	10	3209	29	29	5

Table 8.3 summarizes the p -values computed for all eight tests described in sections 8.1.1–8.1.5. For all the tests except TL, which prescribes zone boundaries, we use a test level of 0.95 (i.e., $p_{\text{test}} = 0.05$) to determine acceptance. For visual clarity, we extend the red-yellow-green color scheme of the TL test to the other tests, using red for reject and green for accept. The first four tests are for unconditional coverage, the next two are the independence tests, and the last two are the mixed tests. As might be anticipated from Table 8.2, almost all the tests fail, with the exception of a “yellow zone” result for the TL test for $\alpha = 0.05$ and “acceptance” for the TUFF tests for $\alpha = 0.05$ and 0.01. With the exception of p_{BIN} and p_{POF} for $\alpha = 0.05$ and p_{TUFF} for $\alpha = 0.005$, all the p -values are $\ll p_{\text{test}}$, indicating “strong” rejection of the backtest results.

Table 8.3 Summary of p -values for the backtests run on the TVP long-only domestic portfolio returns.

$1 - \alpha$	TL	BIN	PoF	TUFF	CCI	TBFI	CC	TBF
	$P(X \leq x N, \alpha)$	p_{BIN}	p_{POF}	p_{TUFF}	p_{CCI}	p_{TBFI}	p_{CC}	p_{TBF}
0.95	0.982693	0.035	0.039	0.237	$1.8 \cdot 10^{-8}$	$< 10^{-40}$	$1.5 \cdot 10^{-8}$	$< 10^{-40}$
0.99	0.999997	$6.8 \cdot 10^{-7}$	$9.2 \cdot 10^{-6}$	0.089	$9.2 \cdot 10^{-8}$	$< 10^{-30}$	$3.4 \cdot 10^{-11}$	$< 10^{-33}$
0.995	0.999960	$1.2 \cdot 10^{-5}$	$1.4 \cdot 10^{-4}$	0.041	$1.6 \cdot 10^{-5}$	$< 10^{-25}$	$6.5 \cdot 10^{-8}$	$< 10^{-27}$

A subsidiary question we investigate is the role of REIT diversification in VaR performance. For this purpose, we perform a comparison of the historical TVP long-only domestic, international, and global portfolios. We therefore reconsider the backtest results for the historical TVP long-only domestic portfolio over the time period 5/1/2018 through 12/18/2020, the same time period as for the international and global

8. Backtesting




Figure 8.2 Daily returns, $r_p(t)$, versus value-at-risk projections, $\text{VaR}_{1-\alpha}(r_p, t)$, for $\alpha = 0.05, 0.01$, and 0.005 for the TVP long-only domestic portfolio optimization subject to a 4% turnover constraint over the time period 5/1/2018 through 12/18/2020.

portfolios. For this time period, Fig. 8.2 plots the daily returns and $\text{VaR}_{1-\alpha}(r_p, t)$ values for $\alpha = 0.05, 0.01$, and 0.005 . Table 8.4 shows the failure data, and Table 8.5 summarizes the test results. In contrast to the longer time period covered by Table 8.3, in the shorter time period covered by Table 8.5, there are better results at $1 - \alpha = 0.995$ than at $1 - \alpha = 0.95$. Although this international portfolio still fails all the independence and mixed tests, the measured p -values are better than for the domestic portfolio.

Table 8.4 Failure statistics for the TVP long-only domestic portfolio optimization over the time period 5/1/2018 through 12/18/2020.

$1 - \alpha$	$1 - x/N$	N	x	$N\alpha$	$x/(N\alpha)$
0.95	0.904	666	64	33.3	1.92
0.99	0.973	666	18	6.7	2.70
0.995	0.986	666	9	3.3	2.70
	n_1	n_{00}	n_{10}	n_{01}	n_{11}
0.95	31	551	50	51	13
0.99	115	634	13	13	5
0.995	460	650	6	6	3

Table 8.5 Summary of p -values for the backtests run on the TVP optimized long-only domestic portfolio returns over the time period 5/1/2018 through 12/18/2020.

	TL	BIN	PoF	TUFF	CCI	TBFI	CC	TBF
$1 - \alpha$	$P(X \leq x N, \alpha)$	p_{BIN}	p_{POF}	p_{TUFF}	p_{CCI}	p_{TBFI}	p_{CC}	p_{TBF}
0.95	1.0000	$4.8 \cdot 10^{-8}$	$1.1 \cdot 10^{-6}$	0.629	0.0052	$2.0 \cdot 10^{-7}$	$1.4 \cdot 10^{-7}$	$< 10^{-10}$

0.99	0.9999	$1.0 \cdot 10^{-5}$	$2.6 \cdot 10^{-4}$	0.886	$4.1 \cdot 10^{-5}$	$< 10^{-10}$	$2.8 \cdot 10^{-7}$	$< 10^{-12}$
0.995	0.9977	0.0018	0.0102	0.333	$7.9 \cdot 10^{-5}$	$< 10^{-10}$	$1.5 \cdot 10^{-5}$	$< 10^{-10}$

Fig. 8.3 and Tables 8.6 and 8.7 present the results for the TVP optimized long-only international portfolio.




Figure 8.3 Daily returns, $r_p(t)$, versus value-at-risk projections, $\text{VaR}_{1-\alpha}(r_p, t)$, for $\alpha = 0.05, 0.01$, and 0.005 for the TVP long-only international portfolio optimized subject to a 4% turnover constraint.

Table 8.6 Failure statistics for the TVP optimized long-only international portfolio.

$1 - \alpha$	$1 - x/N$	N	x	$N\alpha$	$x/(N\alpha)$
0.95	0.924	662	50	33.1	1.51
0.99	0.974	662	17	6.62	2.56
0.995	0.985	662	10	3.31	3.02
	n_1	n_{00}	n_{10}	n_{01}	n_{11}
0.95	31	569	40	41	11
0.99	109	624	16	16	5
0.995	457	642	7	7	5

Table 8.7 Summary of p -values for the backtests run on the TVP optimized long-only international portfolio returns.

$1 - \alpha$	TL	BIN	PoF	TUF _F	CCI	TBFI	CC	TBF
	$P(X \leq x N, \alpha)$	p_{BIN}	p_{POF}	p_{TUFF}	p_{CCI}	p_{TBFI}	p_{CC}	p_{TBF}
0.95	0.9982	0.0026	0.0049	0.629	0.231	0.0045	0.0094	0.0010

8. Backtesting

0.99	0.9998	$5.0 \cdot 10^{-5}$	$7.1 \cdot 10^{-4}$	0.930	0.071	$9.2 \cdot 10^{-9}$	$6.3 \cdot 10^{-4}$	$< 10^{-10}$
0.995	0.9994	$2.2 \cdot 10^{-4}$	0.0030	0.337	$\frac{0.006}{6}$	$< 10^{-10}$	$3.1 \cdot 10^{-4}$	$< 10^{-11}$

optimized subject to a turnover constraint of 4%. This portfolio (of seven REITs) is in the TL yellow zone and passes the TUFF test at all three $1 - \alpha$ levels. It passes the CCI independence test at two $1 - \alpha$ levels. The p -values for the failed tests are “better” at the $1 - \alpha = 0.95$ level than at the other two levels.

Fig. 8.4 and Tables 8.8 and 8.9 present the results for the TVP long-only global portfolio optimized subject to a turnover constraint of 4%. This global portfolio, of 33 REITS, comprises both the domestic and international REIT assets. Although the test results are not quite as good as those for the international portfolio, the diversification does result in better performance (Table 8.9) with respect to the unconditional-coverage tests compared to that for the domestic portfolio on its own (Table 8.5).




Figure 8.4 Daily returns, $r_p(t)$, versus value-at-risk projections, $\text{VaR}_{1-\alpha}(r_p, t)$, for $\alpha = 0.05, 0.01$, and 0.005 for the TVP long-only global portfolio optimized subject to a 4% turnover constraint.

Table 8.8 Failure statistics for the TVP optimized long-only global portfolio.

$1 - \alpha$	$1 - x/N$	N	x	$N\alpha$	$x/(N\alpha)$
0.95	0.921	662	52	33.1	1.57
0.99	0.980	662	13	6.62	1.96
0.995	0.986	662	9	3.31	2.72
	n_1	n_{00}	n_{10}	n_{01}	n_{11}
0.95	11	568	41	41	11
0.99	113	640	8	8	5

0.995 456 646 6 6 3

Table 8.9 Summary of p -values for the backtests run on the TVP optimized long-only global portfolio returns.

$1 - \alpha$	TL	BIN	PoF	TUFF	CCI	TBFI	CC	TBF
	$P(X \leq x N, \alpha)$	p_{BIN}	p_{POF}	p_{TUFF}	p_{CCI}	p_{TBFI}	p_{CC}	p_{TBF}
0.95	0.9994	$7.5 \cdot 10^{-4}$	0.0018	0.574	0.0015	$4.1 \cdot 10^{-8}$	$4.8 \cdot 10^{-5}$	$3.2 \cdot 10^{-9}$
0.99	0.9920	0.0127	0.0277	0.900	$8.0 \cdot 10^{-7}$	$< 10^{-10}$	$4.6 \cdot 10^{-7}$	$< 10^{-11}$
0.995	0.9978	0.0017	0.0098	0.339	$8.0 \cdot 10^{-5}$	$< 10^{-11}$	$1.5 \cdot 10^{-5}$	$< 10^{-12}$

8.2.2 Dynamic Optimizations

We now consider backtesting results for the dynamically optimized portfolios of Chapter 7. Fig. 8.5 shows the time series of daily returns, $r_p(t)$, versus value-at-risk projections, $\text{VaR}_{1-\alpha}(r_p, t)$, for $\alpha = 0.05, 0.01$, and 0.005 for the dynamic TVP long-only domestic portfolio optimized subject to a 4% turnover constraint. (See Figs. 7.3 through 7.7 for other performance and risk measures for this portfolio.) The $\text{VaR}_{1-\alpha}(r_p, t)$ was computed using a sample of 10,000 projections of $r_p(t)$ based on correlated statistics estimated from the asset returns over an eight-year (2,016 trading days) rolling window.




Figure 8.5 Daily returns, $r_p(t)$, versus value-at-risk projections, $\text{VaR}_{1-\alpha}(r_p, t)$, for $\alpha = 0.05, 0.01$, and 0.005 for the dynamic TVP long-only domestic portfolio optimized subject to a 4% turnover constraint.

8. Backtesting

Table 8.10 summarizes the failure data for VaR backtests at the quantile levels $1 - \alpha = 0.95, 0.99$, and 0.995 for this portfolio. The variables defining each column are defined in sections 8.1.1, 8.1.3 and 8.1.3. The number of observed failures, x , exceeds the number of expected failures, $N\alpha$, for all three values of α , by ratios ranging from 1.15 to 1.41.

Table 8.10 Failure statistics for the dynamic TVP long-only domestic portfolio optimized subject to a 4% turnover constraint.

$1 - \alpha$	$1 - x/N$	N	x	$N\alpha$	$x/(N\alpha)$
0.95	0.942	3273	189	163.6	1.15
0.99	0.987	3273	42	32.7	1.28
0.995	0.993	3273	23	16.4	1.41

	n_1	n_{00}	n_{10}	n_{01}	n_{11}
0.95	5	2912	171	171	18
0.99	10	3191	39	39	3
0.995	116	3226	23	23	0

Table 8.11 Summary of p -values for the backtests run on the dynamic TVP long-only domestic portfolio optimized subject to a 4% turnover constraint.

$1 - \alpha$	TL	BIN	PoF	TUFF	CCI	TBFI	CC	TBF
	$P(X \leq x N, \alpha)$	p_{BIN}	p_{POF}	p_{TUFF}	p_{CCI}	p_{TBFI}	p_{CC}	p_{TBF}
0.95	0.979	0.042	0.047	0.237	0.035	6.810^{-7}	0.015	4.010^{-7}
0.99	0.952	0.103	0.119	0.089	0.017	0.004	0.017	0.003
0.995	0.955	0.100	0.121	0.616	0.568	0.123	0.256	0.096

Table 8.11 summarizes the test results. The results are dramatically improved over those for the corresponding historically optimized portfolio (shown in Table 8.3). More specifically, all the TL tests are in the yellow, and all the remaining tests are passed at the 0.995 level. Moreover, the three unconditional-coverage tests, BIN, PoF, and TUFF, are passed at the 0.99 level.

Based on the results presented, we make the following observations.

1. Because volatility occurs in clusters, the unconditional-coverage tests are easier to pass than the independence tests. Active portfolio management requires additional weight adjustments in response to short-term shocks. To emphasize this point, we rerun the tests presented in Table 8.11 but restrict the time period to 1/1/2010 through 12/31/2019, avoiding the market turbulence of the 2008 Great Recession and the 2020 pandemic. The results, shown in Table 8.12, are much improved. There is still enough volatility clustering to make it difficult to pass the independence tests.
2. No single improvement in optimization (e.g., moving from historical to dynamic) can fully improve the backtest performance of a portfolio. Dynamic optimization must be combined with an active management approach that incorporates the risk-management techniques of Chapter 10. For example, the results of incremental and component risk (section 5.4) must be used continuously to monitor the portfolio, identifying its greatest risk-contributing and risk-diversifying assets and applying additional weighting techniques to capitalize on their risk.

Table 8.12 Summary of p -values for the backtests run on the dynamic TVP long-only domestic portfolio optimized subject to a 4% turnover constraint for the time period 1/1/2010 through 12/31/2019.

$1 - \alpha$	TL	BIN	PoF	TUFF	CCI	TBFI	CC	TBF
	$P(X \leq x N, \alpha)$	p_{BIN}	p_{POF}	p_{TUFF}	p_{CCI}	p_{TBFI}	p_{CC}	p_{TBF}
0.95	0.285	0.534	0.530	0.352	0.0015	$1.4 \cdot 10^{-7}$	0.0054	$1.6 \cdot 10^{-7}$
0.99	0.810	0.442	0.453	0.132	0.0039	0.0084	0.011	0.010
0.995	0.865	0.334	0.354	0.108	0.651	0.376	0.587	0.389

References

- Basel Committee on Banking Supervision, Supervisory Framework for the Use of “Backtesting” in Conjunction with the Internal Models Approach to Market Risk Capital Requirements. (1996, January). <https://www.bis.org/publ/bcbs22.htm>.
- Campbell, S. D. (2005). A review of backtesting and backtesting procedures. Finance and Economic Discussion Series 2005-21. Division of Research & Statistics and Monetary Affairs, Federal Reserve Board, Washington, DC.
- Christoffersen, P. (1998). Evaluating interval forecasts. *International Economic Review*, 39, 841–862.
- Haas, M. (2001). New methods in backtesting. Financial Engineering, Center of Advanced European Studies and Research, Bonn.
- Kupiec, P. (1995). Techniques for verifying the accuracy of risk management models. *Journal of Derivatives*, 3, 73–84.

Chapter 9

Diversification with Real Estate Stocks

In Chapter 5, we considered the diversification of our domestic REIT portfolio based on the addition of international REITs traded over the counter as ADRs. In this sense, the foreign REITs are in the same investment class as the domestic REITs. We now consider diversification based on the addition of securities from a different investment class – namely, real estate stocks – to our domestic REIT portfolio. Following the requirement that daily prices for each stock be available since 12/13/1999, we choose the five stocks described in section 2.2. Moreover, adding only five stocks to our existing portfolio of 26 REITs has the virtue of exerting an incremental, rather than dominant, influence on the portfolio.

For demonstration purposes, we compare the M95 and T95 optimizations of the REIT-only (R) and REITs-plus-stocks (R+S) portfolios using both historical (H) and dynamic (D) methods. All eight portfolios are run long-only and subjected to a 0.4% daily turnover constraint. Fig. 9.1 shows the cumulative price performance of the four CVaR_{95} -minimizing portfolios and the four tangent portfolios minimizing the CVaR_{95} ratio. Under both historical and dynamic optimization, adding the stocks significantly improves the price performance of the portfolio. The reverse holds for the tangent optimizations, under which adding the stocks dramatically worsens the price performance.





Figure 9.1 Cumulative price for the domestic REITs-only (R) and REITs-plus-stocks (R+S) portfolios optimized under the historical (H) and dynamic (D) M95 (left) and T95 (right) methods. The optimizations are long-only and subject to a 0.4% turnover constraint.


We postulate that the improvement seen in the mean-risk portfolios is due to the increased number of assets (diversification). Addition of assets potentially reduces the risk contribution of any single asset and thus the overall risk exposure. This postulate is supported by the results shown in Figs. 9.2–9.4. Fig. 9.2 shows the relative weight given to each asset as a function of time for each portfolio and optimization. For instance, a comparison of the historical M95 optimization of the R and R+S portfolios indicates how the optimization gives weight to the stocks TRC, NTP, MLP, and BVH (but reduces the weights of the REITs IRM, AMT, NLY, and CO). Fig. 9.3 quantifies this further by plotting the fraction of the portfolio weight contributed by the real estate stocks in the historical and dynamic M95 optimizations of the R+S portfolio. For the historical M95 optimization, between 20% and 40% of the portfolio weight is assigned to stocks. For the dynamic M95 optimization, this fraction varies with time from 0% to 50%.

Although a discussion based on weights is suggestive (and, as we will show, incorrect in the case of T95 optimization), more conclusive evidence of the impact of stock diversification on


minimum




Historical, REITs only



Historical, REITs + stocks



Dynamic, REITs only



Dynamic, REITs + stocks

Figure 9.2 Percent asset-weight distributions for the indicated long-only portfolios optimized under the M95 (left) and T95 (right) methods. The optimizations are subject to a 0.4% turnover constraint.




Figure 9.3 Fraction of portfolio weight held by stocks in the REITs-plus-stocks portfolio optimized under the historical (H) and dynamic (D) M95 and T95 methods.

9. Diversification with Real Estate Stocks

-risk optimization comes from Fig. 9.4, which compares values-at-risk for the domestic R and R+S portfolios subject to historical M95 optimization. This figure shows that, uniformly over time, adding the stocks reduces both the 95% and 99% quantile value-at-risk of the portfolio. Fig. 9.5 quantifies the changes between the R and R+S values-at-risk in the form of box-whisker plots for both the historical and dynamic optimizations.




Figure 9.4 VaR_{95} and VaR_{99} values-at-risk for the domestic REITs-only (R) and REITs-plus-stocks (R+S) portfolios optimized under the historical (H) and dynamic (D) M95 methods. The optimizations are long-only and subject to a 0.4% turnover constraint.




Figure 9.5 VaR_{95} and VaR_{99} values-at-risk for the domestic REITs-only (R) and REITs-plus-stocks (R+S) portfolios optimized under the historical (H) and dynamic (D) M95 methods. The optimizations are long-only and subject to a 0.4% turnover constraint. Not all outlier values are visible in the D,R and D,R+S plots.

As Fig. 9.1 illustrates, adding the stocks diminishes the performance of the T95 portfolio. In this case, it is possible that adding arbitrary stocks to the portfolio worsens performance because the REITs comprising the portfolio are already optimized to the market. Comparing the VaR_{95} and VaR_{99} values-at-risk for the historical T95 optimizations (Fig. 9.6) with those for the M95 optimizations (Fig. 9.4) indicates that including the stocks leads to a worsening of the portfolio's value-at-risk for significant time periods.




Figure 9.6 VaR_{95} and VaR_{99} values-at-risk for the domestic REITs-only (R) and REITs-plus-stocks (R+S) portfolios optimized under the historical (H) and dynamic (D) T95 methods. The

optimizations are long-only and subject to a 0.4% turnover constraint.

We now consider how adding the stocks affects the risk measures of the portfolio. Fig. 9.7 compares MDD values (computed using a one-year moving window) for the portfolios/optimizations. For the M95 optimization, over the period 2009–2010, MDD is worse for the R+S portfolio than for the R portfolio. For the period 2013–2017, the MDD of the R+S portfolio is generally better (historical optimization) or no worse (dynamic optimization) than that of the R portfolio. Interestingly, during the pandemic, there are no significant differences in MDD among the various portfolios/optimizations. For the T95 optimization, except for a period in 2017, the MDD of the R+S portfolio is always worse than that of the corresponding R portfolio, with the significant differences occurring during the pandemic period.




Figure 9.7 MDD (one-year moving average) for the domestic REITs-only (R) and REITs-plus-stocks (R+S) portfolios optimized under the historical (H) and dynamic (D) M95 (top) and T95 (bottom) methods.

Fig. 9.8 compares the time evolution of SR (computed using a one-year moving window) as well as the total-time-period SR for these portfolios/optimizations. For M95, during the Great Recession and pandemic periods, the R+S portfolio has smaller SRs. However, between these two periods, the R+S portfolio outperforms in terms of SR. The net effect over the total time period is improved SR performance in the R+S portfolio compared to the R portfolio for the same (historical or dynamic) M95 optimization. For the T95 optimization, the R+S portfolio outperforms the R portfolio in terms of SR only during the 2017 period; for the remainder of the time, the R portfolio has better SRs. Thus, over the entire time period, the R+S portfolio undergoes an SR decrease compared to the R portfolio under the same (historical or dynamic) T95 optimization.

Our discussion of SS (see Fig. 9.9) is similar to that of the SR, with the exception that the former exhibits a slight decrease under the dynamic M95 optimization in moving from the R to the R+S portfolio. Fig. 9.10 shows that for both the M95 and T95 optimizations, the total RR

9. Diversification with Real Estate Stocks

increases as a result of diversifying the REIT portfolio with real estate stocks, with the historical optimization producing greater increases than the dynamic optimization.




Figure 9.8 SR versus time (one-year moving window) (left) and total-time-period SR (right) for the domestic REITs-only (R) and REITs-plus-stocks (R+S) portfolios optimized under the historical (H) and dynamic (D) M95 (top) and T95 (bottom) methods.




Figure 9.9 SS versus time (one-year moving window) (left) and total-time-period SS (right) for the domestic REITs-only (R) and REITs-plus-stocks (R+S) portfolios optimized under the historical (H) and dynamic (D) M95 (top) and T95 (bottom) methods.

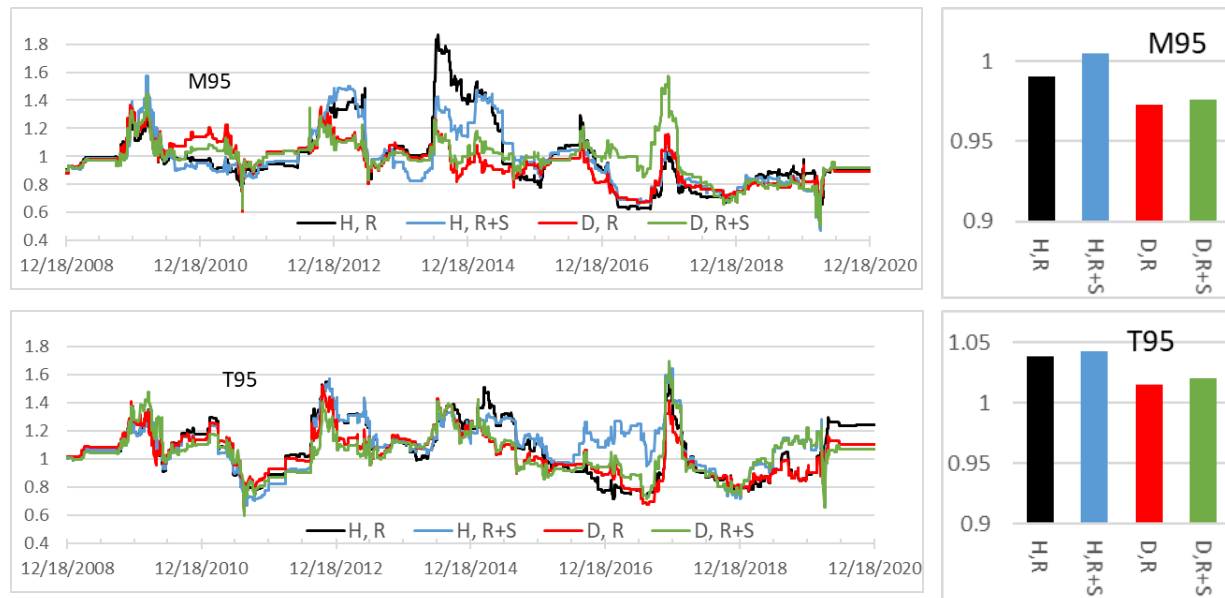


Figure 9.10 RR versus time (one-year moving window) (left) and total-time-period RR (right) for the domestic REITs-only (R) and REITs-plus-stocks (R+S) portfolios optimized under the historical (H) and dynamic (D) M95 (top) and T95 (bottom) methods.

Chapter 10

Risk Information and Management

Focusing on risk-management tools, this chapter covers three topics that concern risk information and management: early warning systems (section 10.1); risk budgets, incremental risk, and component risk (section 10.3); and factor analysis (section 10.4). To complement the discussion of component risk, in section 10.2 we discuss the time evolution of asset weights.

Timely and accurate information is essential in managing investment risk. An early warning system continually tests market risk and has the potential to forecast distressed market periods. There are open questions regarding the capability of early warning systems, including concerns about signal-to-noise ratio and how early detection can be achieved. We consider two warning systems, the first based on the tail-loss ratio (TLR) and the second on Mahalanobis distance. We apply these warning systems to historical data for the one general market index (SPY: SPDR S&P 500 ETF Trust) and three REIT indices (WILLREITIND, WILLREITPR, and FNRE) used in Chapter 3 and compare the results to known global-market upheavals.

Changes made to a portfolio (either in normal investment mode or in response to anticipated market disruption) alter investment risk. The change in risk accompanying the change in some factor in a portfolio is referred to as incremental risk. Measures of incremental risk facilitate risk–return decision-making. Closely aligned with the concept of incremental risk is the decomposition of a portfolio into its constituent (or component) risks. Such a decomposition is useful for identifying both high and low sources of risk, setting position limits, determining capital requirements, etc. We illustrate these two risk-management techniques by applying them to the various long-only and long–short portfolios of domestic REIT ETF assets developed in Chapter 3. In our analysis, we use VaR as the proxy measure of risk. Thus, our discussion centers on incremental VaR and component VaR.

The principal question addressed by a factor model is whether, given a time series of returns for a portfolio, there are (a few) common factors responsible for (most of) the observed behavior. Factor analysis provides the answer to this question in terms of the relative influence (loading or “beta”) of each common factor on the return of each asset as well as in terms of an estimate of the portion of the return behavior (the latent error) that remains unexplained by the common factors. We apply a factor analysis to an equally weighted portfolio tracking 20 indices consisting of (i) the MVP long-only domestic portfolio (interpreted as an index) of section 4.1.1, (ii) the ETF SPY, (iii) the three REIT indices listed in section 2.3.1, and (iv) the 15 market-representative assets listed in section 2.4.

10.1 Early Warning Systems

Research has shown that *structural breaks* in financial time series can predict market disruptions (Andreou and Ghysels, 2009). The Chow test (Chow, 1960) is used to determine whether such structural breaks are present. We therefore consider an early warning system that utilizes the Chow test to look for structural breaks in the (time-averaged) behavior of a time series.

10.1.1 Chow Test for a Structural Break

Consider an observed time series y_t , $t = 1, \dots, T$. A common approach is to try to fit a predictive regression model with a small number of parameters to the time series. However, there may be a significant structural break at some intermediate time $1 < T_1 < T$ such that the time series is better fit by two regression models,⁷¹ the first over the time interval $1 \leq t \leq T_1$ and the second over the interval $T_1 + 1 \leq t \leq T$. The Chow test (Chow, 1960) is an application of the F -test that has the following null hypothesis: The coefficients of the two regression models are the same; that is, no structural break exists. *For the early-warning-system computations in this chapter, we utilize a linear regression model.*

To examine a time series for structural breaks, we employ a moving window of $N < T$ trading days. For time t , the window isolates the subset of the time series $(t - N + 1, t)$. Let $N_2 < N$. We look for a break in this subset between the successive times $t - N_2$ and $t - N_2 + 1$. Then, $N_1 = N - N_2$ is the length of the time-period $[t - N + 1, t - N_2]$. Let $RSS(t)$ be the sum of squared residuals of the regression model fit to the time series of length N in the entire window, and let $RSS_1(t)$ and $RSS_2(t)$ be the sum of squared residuals from the regression models over the two separate time intervals of lengths N_1 and N_2 . Assume the regression model has K parameters. The Chow test statistic is

$$F_{\text{Chow}}(t) = \frac{(RSS(t) - RSS_1(t) - RSS_2(t))/K}{(RSS_1(t) + RSS_2(t))/(N_1 + N_2 - 2K)} \sim F_{K, N_1 + N_2 - 2K}. \quad (10.1)$$

As noted in (10.1), the Chow statistic follows an F -distribution with K and $N_1 + N_2 - 2K$ degrees of freedom. We use the hypothesis test to generate a probability value $p(t)$ ⁷² to indicate whether a breakpoint occurs between $t - N_2$ and $t - N_2 + 1$. We use a threshold value of 5% to declare the probable existence of a breakpoint. Thus, we can define the “hypothesis acceptance/rejection” time series as follows:

$$h(t) = \begin{cases} 1 & \text{if } p(t) \leq 0.05, \\ 0 & \text{if } p(t) > 0.05. \end{cases} \quad (10.2)$$

A value of $h(t) = 1$ indicates a rejection of the null hypothesis, that is, $h(t) = 1$ indicates the existence of a breakpoint at a 95% confidence level. A value of $h = 0$ indicates acceptance of the null hypothesis, no breakpoint, with 95% confidence.

10.1.2 Early Warning Based on Tail-Loss Ratio

A TLR has been proposed as a risk measure for determining financial market behavior before and during periods of high volatility (see Shirvani et al., 2019). The ratio (defined below) was originally based on the normalized difference between values of VaR (3.21) and CVaR (3.22). We also consider some alternatives. Intuitively, a CVaR value that begins to differ considerably from the VaR

⁷¹ We restrict to the case in which both regression models have the same form but have different values for the parameters.

⁷² Under the null hypothesis of no structural break, the value $p(t)$ is the largest probability to have a value of $F_{K, N_1 + N_2 - 2K}$ that is as extreme (or more extreme) than the value of $F_{\text{Chow}}(t)$ observed. If $p(t)$ is very small, then an unusual event relative to the null hypothesis is considered to have occurred.

value indicates the growing presence of low-probability, high-impact losses. Because the VaR and CVaR of an asset (portfolio, market index) will change over time, for given tail probabilities α and β where $0 < \alpha < \beta < 1$, we define the tail-loss-ratio time series $\text{TLR}(t; \alpha, \beta)$ as


$$\text{TLR}(t; \alpha, \beta) = \frac{\text{CVaR}_\alpha(t) - \text{VaR}_\beta(t)}{\text{VaR}_\beta(t)} = \frac{\text{CVaR}_\alpha(t)}{\text{VaR}_\beta(t)} - 1, \quad 1 \leq t \leq N. \quad (10.3a)$$

We also consider this variant of (10.3a),

$$\text{MTLR}(t; \alpha, \beta) = \frac{\text{MVar}_\alpha(t)}{\text{VaR}_\beta(t)} - 1 = \frac{\text{VaR}_{\alpha/2}(t)}{\text{VaR}_\beta(t)} - 1, \quad 1 \leq t \leq N, \quad (10.3b)$$

where $\text{MVar}_\alpha(t)$ is the median value of VaR values corresponding to all the probabilities between 0 and α .⁷³ MTLR is an improvement on TLR in the sense that the median value MVar is much less influenced by outlier values than is the average value CVaR.

To test the models (10.3a) and (10.3b), we apply them to the historical returns for the four market indices listed in section 2.3.1: SPY, WILLREITIND, WILLREITPR, and FNRE. Return data for all four indices is available beginning 12/13/1999. We track return data from that day through 7/26/2019. For each index, empirical daily VaR and CVaR values are estimated from the daily log-return data using a rolling-window analysis with a window of four years⁷⁴ (1,008 trading days). As a result, VaR and CVaR values are available from 1/9/2004 through 7/26/2019. The daily returns, $\text{TLR}(t; 0.01, 0.10)$ and $\text{MTLR}(t; 0.01, 0.10)$ values computed for the WILLREITIND index are shown in Fig. 10.1. Because $\text{CVaR}_\alpha(t) \geq \text{MVar}_\alpha(t) \geq \text{VaR}_\alpha(t) \geq \text{VaR}_\beta(t)$ for $\alpha < \beta$, we have $\text{TLR}(t; \alpha, \beta) \geq \text{MTLR}(t; \alpha, \beta)$ for any time value t , as Fig. 10.1 indicates. Fig. 10.1 also shows the major global-market disruptions, indicated by the date commonly associated with the disruption. The correlation between volatility in the returns and some of these disruptions is plainly evident.



a: Chinese stock bubble; b: U.S. bear market; c: Great Recession; d: Dubai debt standstill; e: European sovereign debt crisis; f: U.S. flash crash; g: U.S. Aug. 2011 decline; h: China crash; i: U.S. market selloff; j: cryptocurrency selloff.


Figure 10.1 Daily returns and computed TLR and MTLR values for the WILLREITIND index.

However, the tail-loss ratio (10.3a) is an unstable measure when $\text{VaR}_\beta(t) \sim 0$ (because $\text{TLR}(t; \alpha, \beta) \rightarrow +\infty$ as $\text{VaR}_\beta(t) \rightarrow 0^+$ while $\text{TLR}(t; \alpha, \beta) \rightarrow -\infty$ as $\text{VaR}_\beta(t) \rightarrow 0^-$) resulting in spurious breakpoints when $\text{VaR}_\beta(t)$ moves even slightly between small positive and negative

⁷³ Therefore, $\text{MVar}_\alpha(t)$ is just $\text{VaR}_{\alpha/2}(t)$.

⁷⁴ As per Basel II.

values. Fig. 10.2 illustrates this instability for $\text{TLR}(t; \alpha = 0.01, \beta = 0.50)$ computed for the WILLREITIND return data. Because it has the same denominator, MTLR is subject to the same instability.



a: Chinese stock bubble; b: U.S. bear market; c: Great Recession; d: Dubai debt standstill; e: European sovereign debt crisis; f: U.S. flash crash; g: U.S. Aug. 2011 decline; h: China crash; i: U.S. market selloff; j: cryptocurrency selloff.


Figure 10.2 The unstable character of $\text{TLR}(t; \alpha = 0.01, \beta = 0.50)$ as computed for the WILLREITIND return data.

To address this unstable behavior, we replace (10.3a) and (10.3b) by the “tail-loss spread” between $\text{CVaR}_\alpha(t)$ or $\text{MVar}_\alpha(t)$ and $\text{VaR}_\beta(t)$, which, again for $0 < \alpha < \beta < 1$, are defined as

$$\text{TLS}(t; \alpha, \beta) = \text{CVaR}_\alpha(t) - \text{VaR}_\beta(t), \quad 1 \leq t \leq N, \quad (10.3c)$$

$$\text{MTLS}(t; \alpha, \beta) = \text{MVar}_\alpha(t) - \text{VaR}_\beta(t) = \text{VaR}_{\alpha/2}(t) - \text{VaR}_\beta(t), \quad 1 \leq t \leq N. \quad (10.3d)$$

Note that $\text{TLS}(t; \alpha, \beta) \geq \text{MTLS}(t; \alpha, \beta)$ for any time value t . Fig. 10.3 presents plots of $\text{TLS}(t; \alpha = 0.01, \beta = 0.50)$ and $\text{MTLS}(t; \alpha = 0.01, \beta = 0.50)$ computed for the WILLREITIND return data.



a: Chinese stock bubble; b: U.S. bear market; c: Great Recession; d: Dubai debt standstill; e: European sovereign debt crisis; f: U.S. flash crash; g: U.S. Aug. 2011 decline; h: China crash; i: U.S. market selloff; j: cryptocurrency selloff.


Figure 10.3 Daily returns and comparisons of TLS to MTLS for the WILLREITIND index.

In the computations performed in this chapter, when $\beta = 0.5$, we use the measures TLS and MTLS. For the values of $\beta \leq 0.25$ that we consider, we use the TLR and MTLR measures. However, switching between measures (10.3a), (10.3b) and (10.3c), (10.3d) depending on the value of β is cumbersome and undesirable. For this reason, we propose the following measure to avoid the instability issue:

$$\text{TLR}_{\text{WBL}}(t; \alpha) = \frac{\text{CVaR}_\alpha(t) - \text{VaR}_{0.5}(t)}{\text{MVar}_{0.5}(t) - \text{VaR}_{0.5}(t)} = \frac{\text{CVaR}_\alpha(t) - \text{VaR}_{0.5}(t)}{\text{VaR}_{0.25}(t) - \text{VaR}_{0.5}(t)}. \quad (10.3e)$$

$$0 < \alpha < 0.25, \quad 1 \leq t \leq N.$$


The denominator of TLR_{WBL} is $\text{MTLS}(t; \alpha = 0.5, \beta = 0.5)$, which is also half of the IQR of the distribution of VaR values. Fig 10.4 shows the daily returns and $\text{TLR}_{\text{WBL}}(t; \alpha = 0.01)$ values for the WILLREITIND index.



a: Chinese stock bubble; b: U.S. bear market; c: Great Recession; d: Dubai debt standstill; e: European sovereign debt crisis; f: U.S. flash crash; g: U.S. Aug. 2011 decline; h: China crash; i: U.S. market selloff; j: cryptocurrency selloff.

Figure 10.4 Daily returns and $\text{TLR}_{\text{WBL}}(t; \alpha = 0.01)$ for the WILLREITIND index.

Equations (10.3a)–(10.3d) each provide a two-parameter (α, β) model of an early warning system.⁷⁵ Fig. 10.5 compares $\text{TLR}(t; 0.01, 0.10)$ for the four indices. Plots for MTLR, TLS, and MTLs for this and other values of α, β are qualitatively similar, except that as the difference $\beta - \alpha$ increases (decreases), the range of TLR values increases (decreases). To summarize the results displayed in Figs. 10.1–10.5 for the four indices, although there is some evidence of rapid change in the values of the time series (10.3a)–(10.3e) around the dates of certain market disruptions, there are market disruption dates for which no rapid change in TLR value is apparent, and there are rapid changes in TLR that are not correlated with a market disruption. We therefore employ the Chow test to look for structural breaks in the time series (10.3a)–(10.3d).



a: Chinese stock bubble; b: U.S. bear market; c: Great Recession; d: Dubai debt standstill; e: European sovereign debt crisis; f: U.S. flash crash; g: U.S. Aug. 2011 decline; h: China crash; i: U.S. market selloff; j: cryptocurrency selloff.


Figure 10.5 Comparison of $\text{TLR}(t; 0.01, 0.10)$ across the four indices.

Fig. 10.6 plots the time series $p(t)$ and $h(t)$ for Chow tests run on the WILLREITIND index for the values $\alpha = 0.01, \beta = 0.50$ using the MTLs spread values (10.3d). We choose a window

⁷⁵ (5e) provides a one-parameter model.

size of $N = 1,008$ trading days and consider three values for N_2 : 126 trading days (six months), 63 trading days (three months), and 21 trading days (one month). Note that as N_2 decreases (N_1 increases), the start and end dates of the time series move accordingly. As a result, the 11/10/2007 date, normally associated with the beginning of the 2007–2009 U.S. bear market, moves out of the range of breakpoint dates tested when $N_1 = 987$.

Concentrating on the hypothesis acceptance/rejection function $h(t)$, when $N_2 = 126$ trading days, we see a strong signal for breakpoints in $MTLS(t; \alpha = 0.01, \beta = 0.50)$ occurring as early as



b: U.S. bear market; c: Great Recession; d: Dubai debt standstill; e: European sovereign debt crisis; f: U.S. flash crash; g: U.S. Aug. 2011 decline; h: China crash; i: U.S. market selloff; j: cryptocurrency selloff.

Figure 10.6 $p(t)$ and $h(t)$ for the Chow test on (10.3d) with varying values of N_2 for the WILLREITIND index.

164 trading days before 9/16/2008 – the failure date of large financial institutions in the United States and the start of the Great Recession. The signal lasts for 10 trading days following the Dubai debt deferment. For the European sovereign debt crisis, the U.S. flash crash, or the U.S. August 2011 decline, no breakpoint signal appears. Breakpoints commensurate with the China crash and U.S. market selloff are detected; some of these occur as early as 45 trading days before the China crash. No breakpoint signal associated with the 2018 cryptocurrency crash appears.

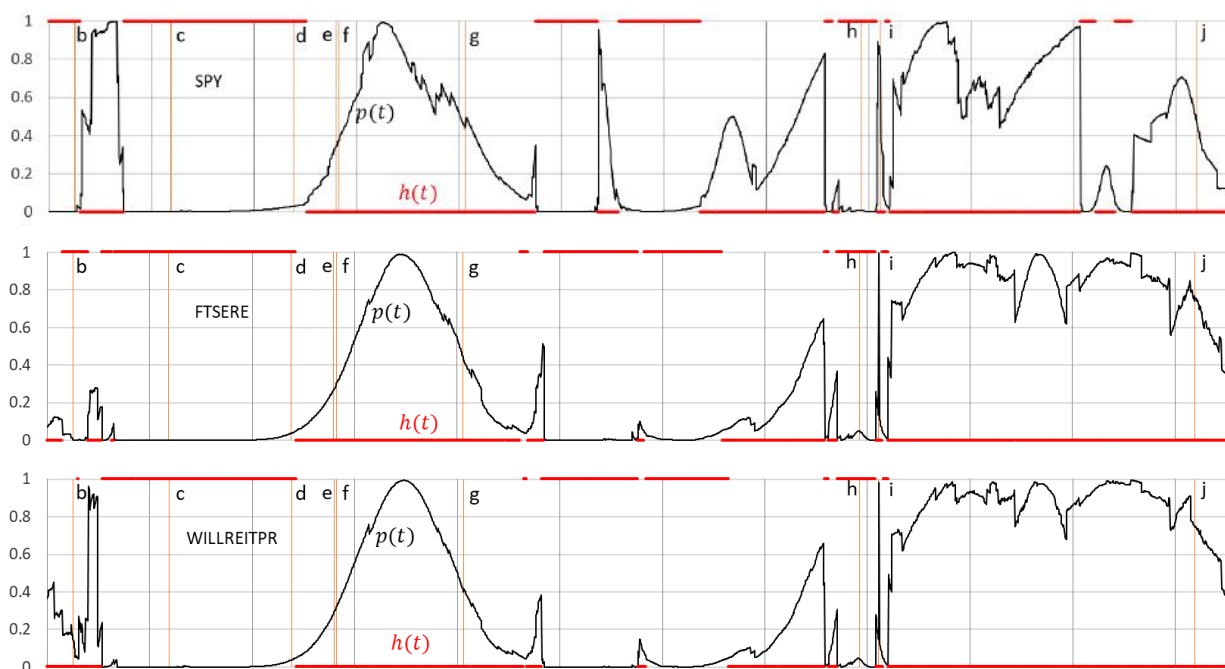
When N_2 is shortened to 63 trading days, the association between breakpoints and market disruptions continues as noted for $N_2 = 126$; however, the “strength” of the signal (measured in terms of how early and over what duration of time) is somewhat weaker. For example, a continuous set of breakpoints associated with the Great Recession appears as early as 112 trading days before 9/16/2008, but it disappears before the date of the Dubai debt deferment.

When N_2 is shortened to 21 trading days, the breakpoint signals become even weaker. The

10. Risk Information and Management

behavior for the other three market indices is qualitatively similar to that displayed for the WILLREITIND index.

Curiously, we detect a strong breakpoint signal over the period 5/8/2012 through 4/15/2013, a period not associated with any global-market disruption. As Fig. 10.7 indicates, this signal is just as strong for the other three indices. Several factors could have contributed to this observed signal.^{76,77}




b: U.S. bear market; c: Great Recession; d: Dubai debt standstill; e: European sovereign debt crisis; f: U.S. flash crash; g: U.S. Aug. 2011 decline; h: China crash; i: U.S. market selloff; j: cryptocurrency selloff.

⁷⁶ “At the property level, the period in question was characterized by the recovery stage of the cycle, compressing rates of return and rising property prices. I find no reason for these signals at the property level. I therefore conclude that the signals were at the REIT level, not at the property level. My surmise, certainly no more than that, is that the signals were driven by political concerns/risks, that is, the prospect of Obama’s reelection and concomitant tax increases, additional regulations making it more difficult do business, etc. Signals after the election likely reinforced these concerns. I have not examined REIT share prices before, during, and after this period.” (Stephen Crosson, Integra Realty Resources, 7/9/2020)

⁷⁷ Kimberly Amadeo’s The Balance article, updated May 13, 2019 (<https://www.thebalance.com/u-s-economy-2012-3305742>), provides an analysis of the economic uncertainties of 2012. She notes that the greatest uncertainty was the 2012 presidential election, which featured “two candidates with radically different approaches to stimulating economic growth.” Consequently, economic growth slowed “as businesses waited to see what direction the country would take.” She attributes the second-largest factor to the uncertainty surrounding the January 1, 2013, economic cliff of impending tax increases and spending cuts. This uncertainty “kept \$1 trillion of corporate expenditures on the sidelines.” The continuing eurozone debt crisis also “wreaked havoc with the U.S. stock market.” Amado notes that, countering these negative forces, (i) the Federal Reserve continued to implement monetary stimulus measures, (ii) the housing market improved as “federal courts settled with banks over the robo-signing accusations,” and (iii) “consumers waded through their debt and resumed shopping.” As a result, “the economy ended the year with a 2.2 percent growth rate” – a healthy rate, but not high enough to absorb the component of the workforce still unemployed from the 2008 financial crisis.

Figure 10.7 $p(t)$ and $h(t)$ for the Chow test on (10.3d) with $N_2 = 126$ for the other three indices.

Rather than vary the size of the subwindows, N_1 and N_2 , we examine the question of varying α and β . We employ the MTLR measure (10.3b) and use the values $N_1 = 882$, $N_2 =$



b: U.S. bear market; c: Great Recession; d: Dubai debt standstill; e: European sovereign debt crisis; f: U.S. flash crash; g: U.S. Aug. 2011 decline; h: China crash; i: U.S. market selloff; j: cryptocurrency selloff.

Figure 10.8 $p(t)$ and $h(t)$ for the Chow test on (10.3b) with varying β for the WILLREITIND index.

126, which produce the strongest “signal” in Fig. 10.6. We hold $\alpha = 0.01$ and vary β over the values 0.25, 0.10, and 0.05. Fig. 10.8 shows the results. Comparing these results with the $\alpha = 0.01$, $\beta = 0.50$ results presented in Fig. 10.6, we see that the duration of the signals generally appears to decrease as $\beta \rightarrow \alpha$. There is one exception: the signal associated with 9/16/2008 and the start of the Great Recession. For $\beta = 0.25$, the breakpoint sequence begins 128 trading days before 9/16/2008 and lasts 290 trading days. For $\beta = 0.10$, the breakpoint sequence begins 173 trading days before 9/16/2008 and lasts 331 trading days. For $\beta = 0.05$, the breakpoint sequence begins 180 trading days before 9/16/2008 and lasts 332 trading days. For comparison, for $\beta = 0.50$ (Fig. 10.6), the breakpoint sequence begins 164 trading days before 9/16/2008 and lasts 479 trading days. However, the results in Fig. 10.6 are for the MTLs measure rather than the MTLR measure employed in Fig. 10.8, which may account for the difference. What is clear is that the financial crisis leading to the Great Recession produces a strong precursor signal that predates the 9/16/2008 crisis point and lasts well into the Great Recession.

10.1.3 Early Warning Based on Mahalanobis Distance

In this section, we consider an early warning system based on Mahalanobis distance. Again, we seek to capture market signals from tail dependencies (extreme events). To capture the tail dependence, we employ the concept of the copula. After estimating copulas empirically, we calculate the Mahalanobis distance between successive copulas. Finally, we analyze structural breaks in the Mahalanobis-distance time series using the Chow test. We begin by expanding on these concepts.

10.1.3.1 Copulas

Let X_i be a random variable. $F_i(x) = \Pr[X_i \leq x]$ is the cumulative distribution function (CDF) for X_i . Whereas values of X_i are generally distributed over the range $(-\infty, \infty)$ or $(0, \infty)$, values of the random variable $U_i = F_i(X_i)$ are uniformly distributed over $[0, 1]$. In particular, information in the long tails of the distribution for X_i is compressed into information for U_i in the compact interval $[0, 1]$. As a result, tail events are weighted with larger probability in the transformed variable.

Consider N random variables, X_1, \dots, X_N , which may be correlated with each other. The copula of the transformed random variables is

$$\begin{aligned} C(u_1, \dots, u_N) &= \Pr(U_1 \leq u_1, \dots, U_N \leq u_N), \\ \text{i.e. } C(u_1, \dots, u_N) &= \Pr(F_1(X_1) \leq u_1, \dots, F_N(X_N) \leq u_N). \end{aligned} \quad (10.4)$$

Thus, the copula C is a multivariate cumulative distribution function. It retains all the information on the dependence structure between the components X_i . Of particular significance is Sklar's theorem, which states that any multivariate joint distribution can be written in terms of univariate marginal distribution functions and a copula describing the dependence structure between the variables. In practice, an empirical copula is computed based on a (large) sample of observed values for each X_i . Note that if each CDF, $F_i(x)$, can be inverted, then the inverse map

$$X_1, \dots, X_N = F_1^{-1}(U_1), \dots, F_N^{-1}(U_N) \quad (10.5)$$

transforms observations of the U_i into observations of the X_i . Equation (10.4) can be rewritten in terms of the inverses, F_i^{-1} :

$$C(u_1, \dots, u_N) = \Pr(X_1 \leq F_1^{-1}(u_1), \dots, X_N \leq F_N^{-1}(u_N)). \quad (10.6)$$

Our random variables X_i will be the daily log-return r_i^t for the $i = 1, \dots, N$ assets in a portfolio. The CDFs, $F_i^t(r)$, $i = 1, \dots, N$, are estimated empirically using the historical log-return values r_i^k , $i = 1, \dots, N$; $k = t - (\tau_1 - 1), \dots, t$, for a window size $\tau_1 < T$:

$$F_i^t(r) = \frac{1}{\tau_1} \sum_{k=t-\tau_1+1}^t \mathbf{1}(r_i^k \leq r). \quad (10.7)$$

Here, $\mathbf{1}(a \leq b)$ is the indicator function. The empirically transformed return values are $u_i^k = F_i^t(r_i^k)$, $i = 1, \dots, N$; $k = t - \tau_1 + 1, \dots, t$, and the empirically evaluated copula is


$$C^t(u_1, \dots, u_N) = \frac{1}{\tau_1} \sum_{k=t-\tau_1+1}^t \mathbf{1}(u_1^k \leq u_1, \dots, u_N^k \leq u_N). \quad (10.8)$$

We are interested in a specific value of the empirical copula

$$C^t(u_1^t, \dots, u_N^t) = \frac{1}{\tau_1} \sum_{k=t-\tau_1+1}^t \mathbf{1}(u_1^k \leq u_1^t, \dots, u_N^k \leq u_N^t), \quad (10.9)$$

which, loosely speaking, is the average probability that the log-returns of the assets over the last τ_1 days are less than today's returns.

We apply the computation of the empirical copula (10.9) to the returns of a portfolio consisting of the 26 domestic REITs and the four indices, SPY, WD, WP, and FNRE. (Thus, $N = 30$.) Return values for all the assets are available over the common time period 01/03/2000 through 07/26/2019. We use a window size of $\tau_1 = 1,008$ trading days. Fig. 10.1 shows the time series of the daily returns beginning from 01/09/2004 for one of the assets in this portfolio, the index WD. The empirical copula probabilities $C^t(u_1^t, \dots, u_N^t)$ computed for the entire portfolio are plotted in Fig. 10.9. Correlations between the perceived volatility in copula values and most of these disruption dates are plainly evident.



a: Chinese stock bubble; b: U.S. bear market; c: Great Recession; d: Dubai debt standstill; e: European sovereign debt crisis; f: U.S. flash crash; g: U.S. Aug. 2011 decline; h: China crash; i: U.S. market selloff; j: cryptocurrency selloff.

Figure 10.9 Empirical copula values for the portfolio of 26 domestic REITs and four market indices.

10.1.3.2 Mahalanobis Distance

The Mahalanobis distance, d (Mahalanobis, 1936), is a generalization of a z -score in that it is designed to measure a (normalized) distance between an observation of a random variable and the mean value of a distribution (presumably the distribution governing that random variable). The Mahalanobis generalization extends the z -score concept to the case of an observation $x = (x_1, \dots, x_N)$ on N random variables governed by a joint-probability distribution with a mean value $\mu = (\mu_1, \dots, \mu_N)$ and covariance matrix S as


$$d^2 = (x - \mu)S^{-1}(x - \mu)', \quad (10.10)$$

where x' denotes vector transpose. Equation (10.10) computes the square of the Mahalanobis distance.

We apply this measure to the observations of the daily copulas computed in equation (10.4) as follows. Consider a window of τ_2 trading days and the vector $\mathbf{C}^t = (C^{t-\tau_2+1}, \dots, C^t)$ of computed copula values in that window. Let u_t and σ_t^2 denote the mean and variance computed from this set of copula values. Moving the window, we can compute a daily sequence of vectors with corresponding mean values and variances (e.g., $\mathbf{C}^{t+1} = (C^{t-\tau_2+2}, \dots, C^{t+1})$, u_{t+1} , σ_{t+1}^2). We use the Mahalanobis measure to determine a distance between \mathbf{C}^{t+1} and the mean value u_t estimated for the distribution of \mathbf{C}^t :

$$d_t^2(u_t, \mathbf{C}^{t+1}) = \frac{1}{\tau_2} \sum_{k=t-(\tau_2-1)}^t \left[\frac{(C^{k+1} - u_t)^2}{\sigma_t^2} \right]. \quad (10.11)$$

Employing a window length of $\tau_2 = 252$ trading days (one trading year) for the copula data plotted in Fig. 10.9 gives the time series of Mahalanobis distances $d(u_t, \mathbf{C}^{t+1}) = \sqrt{d^2(u_t, \mathbf{C}^{t+1})}$ shown in Fig. 10.10. The lower plot shows the Mahalanobis distance over a restricted range of y-axis values to provide greater detail regarding the behavior around the value $d = 1$.⁷⁸



a: Chinese stock bubble; b: U.S. bear market; c: Great Recession; d: Dubai debt standstill; e: European sovereign debt crisis; f: U.S. flash crash; g: U.S. Aug. 2011 decline; h: China crash; i: U.S. market selloff; j: cryptocurrency selloff.


Figure 10.10 Mahalanobis distance for the empirical copula probabilities of Fig. 10.9.

We use the Chow test to look for structural breaks in the time series of Mahalanobis distance computed in equation (10.11). Again, we employ a moving window of $N = 1,008$ trading days

⁷⁸ The fact that the Mahalanobis time series is centered around 1.0 appears to be insignificant.

and vary N_1 and $N_2 = N - N_1$ to test for early warning signal “strength” in searching for possible structural breakpoints in the time series $d(u_t, \mathbf{C}^{t+1})$. In our test, we vary from $N_2 = 882$ to $N_2 = 126$ in increments of 126 trading days (six trading months). (It is clearly preferable to maximize the warning time of a market downturn in order to adjust very large positions.) Because $d(u_t, \mathbf{C}^{t+1})$ is a function of the single variable t , a linear regression model with $K = 1$ parameter (slope) is computed. As in the TLR-based early warning system, we use the hypothesis test to generate a probability value $p(t)$ of a structural break at each point tested. Again, we use a threshold value of 5% to declare the probable existence of a breakpoint and generate the hypothesis accept/reject function $h(t)$.

Fig. 10.11 shows the probability values $p(t)$ and hypothesis function $h(t)$ for two choices of N_1 and N_2 .⁷⁹ The date axis indicates how the choice of N_1 and N_2 affects the total length of time over which potential breakpoints can be detected. For example, for $N_1 = 504$, the earliest possible date for detecting a breakpoint is 01/09/2007, 504 trading days after 01/07/2005 (the first date for which the Mahalanobis distance is computable), and the last possible detection date is 07/25/2017. Consequently, when $N_1 = 504$, we have enough data to look for potential breakpoints prior to the collapse of the Chinese stock bubble on 02/27/2007. When $N_1 = 630$, we lose any view of that stock market event.



a: Chinese stock bubble; b: U.S. bear market; c: Great Recession; d: Dubai debt standstill; e: European sovereign debt crisis; f: U.S. flash crash; g: U.S. Aug. 2011 decline; h: China crash; i: U.S. market selloff.

Figure 10.11 Computed values for $p(t)$ and $h(t)$ for $d(u_t, \mathbf{C}^{t+1})$.

Using Mahalanobis distance, we detect structural breaks in our portfolio commensurate with every major global-market upheaval during this time period except the Chinese stock bubble collapse and the U.S. bear market that began on 10/11/2007. It is important to run the early warning system with varied values of N_1 and N_2 to look for potential signals from “multiple views.”

10.2 Asset Weighting

⁷⁹ These provided the best predictive performance.

To lay the groundwork for our discussion of risk budgets and component risk in section 10.3, we first compare asset weighting under different optimization schemes. Fig. 10.12 plots asset weight as a function of time for the historical optimizations of the MVP, TVP, M99, and T99 long-only domestic portfolios subject to a 4% turnover constraint. The MVP and TVP portfolios minimize portfolio standard deviation, whereas M99 and T99 minimize CVaR at the 99% quantile level (1% tail risk). The most significant difference in terms of which assets are weighted most heavily occurs between the global risk-minimizing portfolios (MVP, M99) and the tangent portfolios (TVP, T99), which maximize the Sharpe (TVP) or Sortino (T99) ratio. Whereas MVP and M99 weight the assets IRM, NLY, AMT, WY, CCI, and PSA most heavily, TVP and T99 concentrate weights on SBAC, AMT, and VTR and to a lesser degree on PSA, ESS, and CCI. The clearest differences between MVP and M99 are M99's more abrupt and pronounced changes in asset weights. The same statement holds for the differences between TVP and T99.

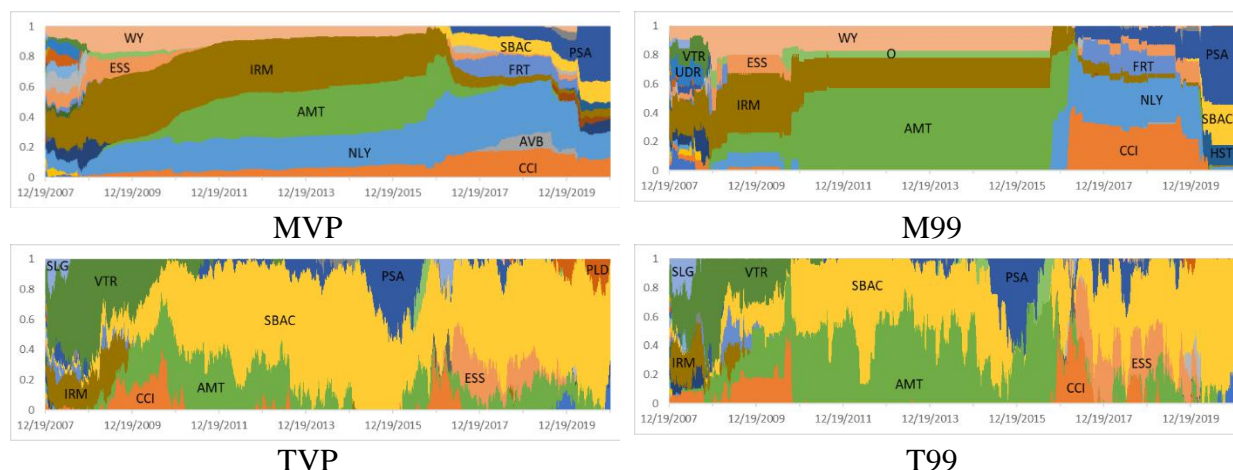


Figure 10.12 Asset weights as a function of time for the historical MVP, TVP, M99, and T99 long-only domestic portfolio optimizations subject to a 4% turnover constraint.

Fig. 10.13 plots asset weight as a function of time for the dynamic long-only optimizations of the MVP, TVP, M99, and T99 domestic portfolios subject to a 4% turnover constraint. Because a dynamic optimization samples a much larger statistical set of asset returns than does a historical optimization, the following changes in asset allocation occur:

- Significant changes in asset weight occur more rapidly in the dynamic optimization.
- Weight changes occur more abruptly an optimization that minimizes CVaR (M99) than for an optimization that minimizes standard error (MVP).
- Some assets (e.g., IRM and AMT for MVP/M99) in the historical optimization are weighted much less significantly in the dynamic optimization.
- Additional assets (e.g., AVB in TVP/T99) are weighted more significantly in the dynamic optimization.

Fig. 10.14 illustrates the impact of asset-weight assignment for the TVP optimization subject to a 4% turnover constraint when applied to the domestic, international, and global portfolios using either the historical or dynamic optimization schemes. The main consideration here is the effect on the weight distribution of adding the international assets to the domestic portfolio. The results

cover the common time period 4/30/2018 through 12/18/2020. Over this period, the weights for the historically optimized domestic portfolio are concentrated on the REIT ETF SBAC, with secondary admixtures of ESS, AMT, and PLD. The weights for the historically optimized international portfolio are roughly equally split between LKREF and CDPYF. When the historical optimization is applied to the global portfolio, which combines the domestic and international assets, the two international REIT ETFs LKREF and CDPYF together comprise 25%–60% of the portfolio weight. The weighting of SBAC is greatly reduced until the beginning

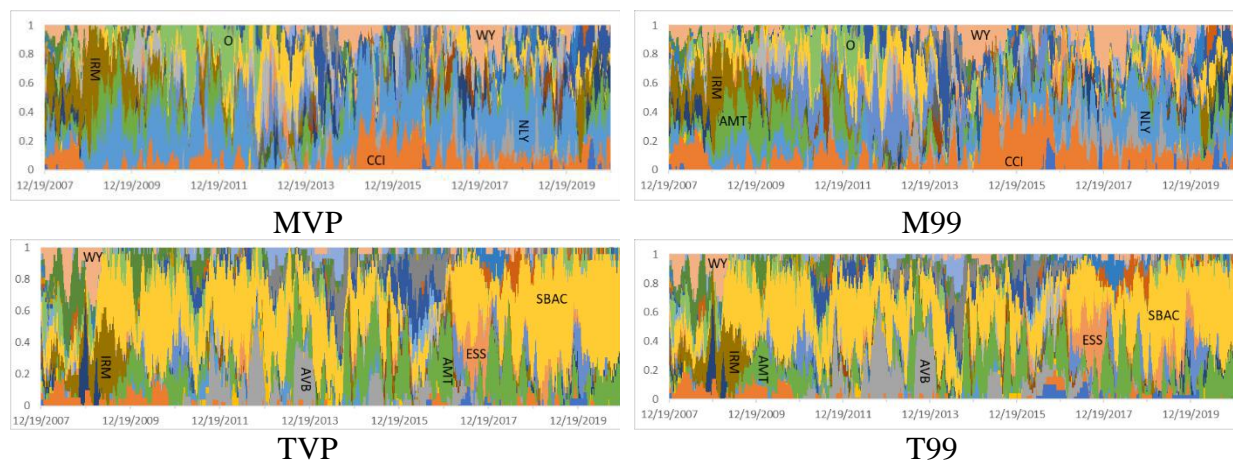


Figure 10.13 Asset weights as a function of time for the dynamic MVP, TVP, M99, and T99 long-only domestic portfolio optimizations subject to a 4% turnover constraint.

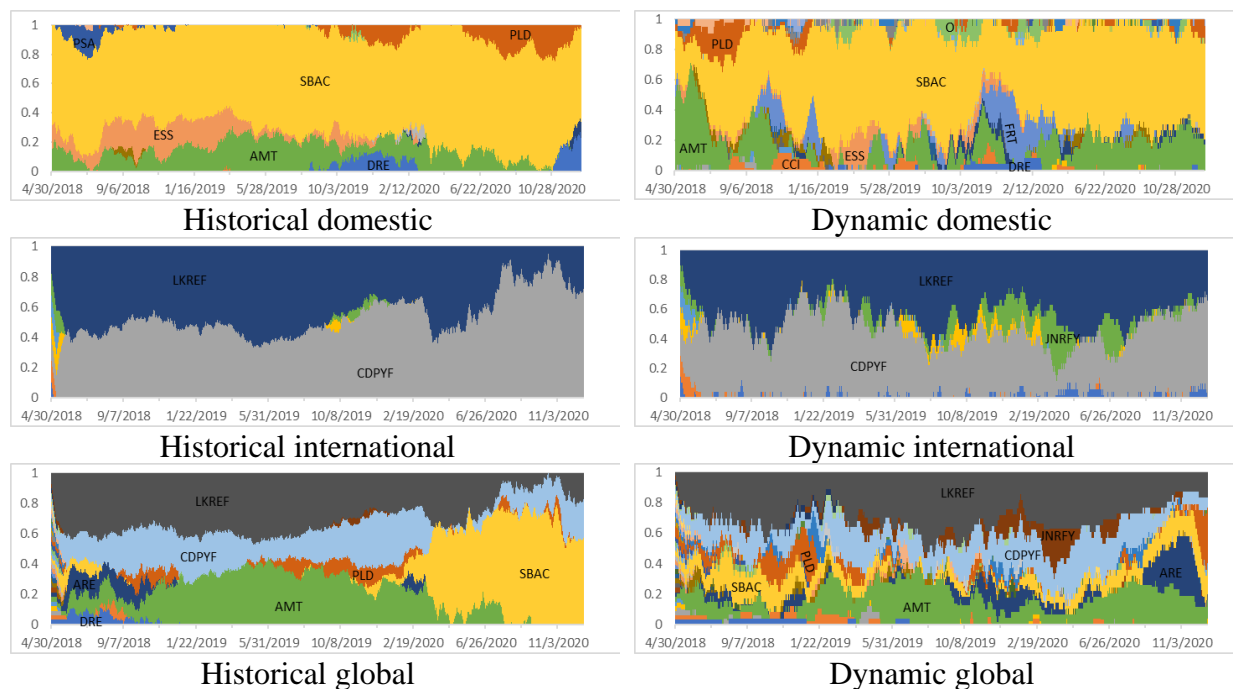


Figure 10.14 Asset weights as a function of time for the historical (left) and dynamic (right) TVP long-only domestic, international, and global portfolio optimizations subject to a 4% turnover constraint.

of the 2020 pandemic. In the case of dynamic optimization, the weighting of SBAC in the global portfolio is greatly reduced (in comparison to the domestic portfolio) over the entire time period, and the international assets LKREF and CDPYF, and to some extent JNRFY, play a consistently stronger role.

We now consider long–short optimizations. Let

$$w_{-}(t) = \sum_{i=1}^n \min(w_i(t), 0) \quad \text{and} \quad w_{+}(t) = \sum_{i=1}^n \max(w_i(t), 0), \quad (10.12)$$

where $w_{-}(t)$ is the total weight of shorted REITs in the portfolio. Because the portfolio is fully invested in REITs, $w_{+}(t) = 1 - w_{-}(t)$. For long–short portfolios, the fractions

$$\frac{w_i(t)}{|w_{-}(t)| + w_{+}(t)}, \quad i = 1, \dots, n, \quad (10.13)$$

provide the appropriate data for asset-weight plots that are analogs of the plots displayed for long-only portfolios in Figs. 10.12–10.14. Using (10.13) reveals a remarkable similarity between the Jacobs et al. and Lo–Patel optimizations. Fig. 10.15 illustrates this similarity for the historical T95 portfolios optimized via Jacobs et al. with $s = 0.04$ and via Lo–Patel with $lev = 0.04$. The weight distributions are virtually identical. This also characterizes comparisons of their respective MVP, TVP, M95, M99, and T99 portfolios as long as $s = lev$ (at least over the range of values $(\infty, 0.04]$).

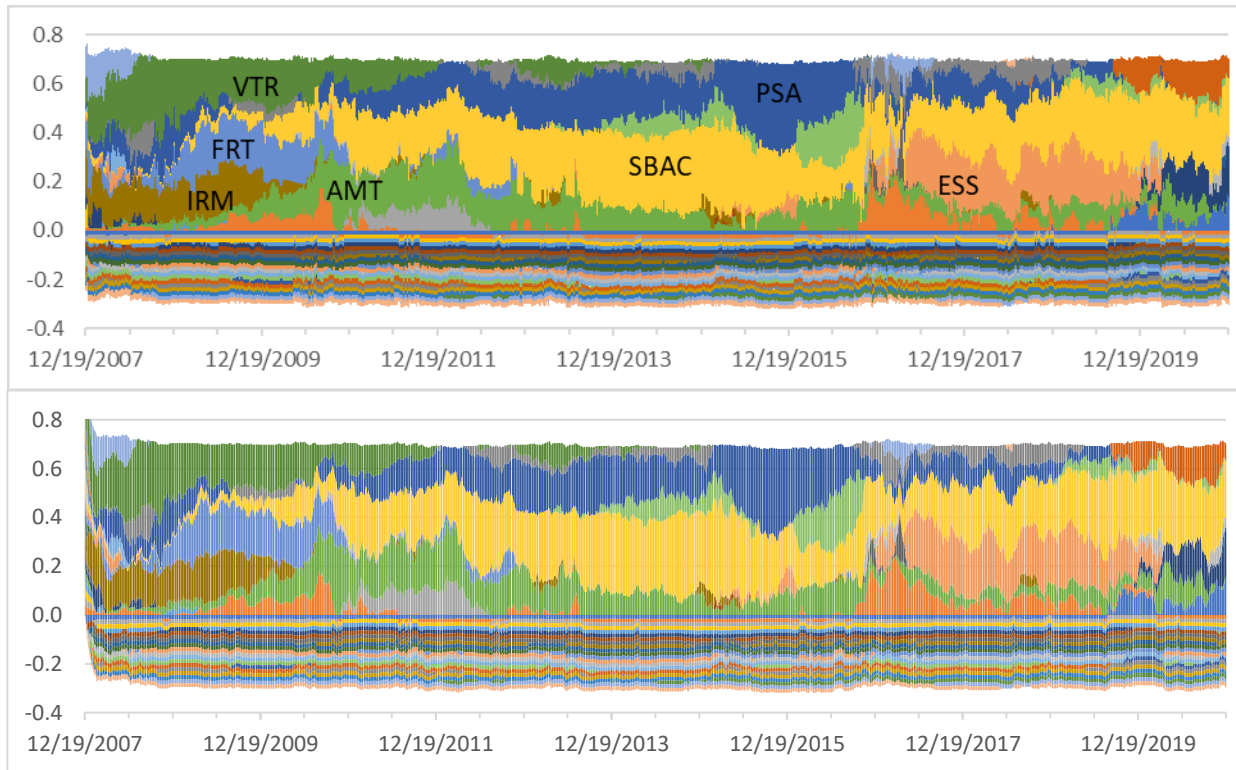


Fig. 10.15 Asset weights as a function of time for the historical T95 long–short domestic portfolio optimizations via Jacobs et al. with $s = 0.04$ (top) and via Lo–Patel with $lev = 0.04$ (bottom). Asset colors are identical to those of Figs. 10.12 and 10.13.

Equally striking is the relatively constant value of the fraction

$$f_{\text{short}}(t) \equiv \frac{w_{-}(t)}{|w_{-}(t)| + w_{+}(t)}, \quad (10.14)$$

which represents a relative total weight of shorted assets in the portfolio. Fig. 10.16 compares the plots of $f_{\text{short}}(t)$ for the six historical long–short domestic portfolios (MVP through T99) optimized via Jacobs et al. with $s = 0.04$ and via Lo–Patel with $lev = 0.04$. Also shown are the quartile values, Q_1, Q_2 , and Q_3 , of the f_{short} time series, which emphasize the high degree of similarity between the Jacobs et al. and Lo–Patel f_{short} time series and quantify the variation over time of each series. In particular, the tangent optimizations, TVP, T95, and T99, develop an extremely similar $Q_2 = 30\%$ total weight of shorted assets. In contrast, the total weight of shorted assets for the minimum-risk portfolios decreases slowly from $Q_2 = 26\%$ for MVP to $Q_2 = 28\%$ for M99. The $f_{\text{short}}(t)$ time series show deviation from these constant values at the onset of the Great Recession. In addition, $f_{\text{short}}(t)$ for MVP changes before the pandemic; after the pandemic “crash,” its f_{short} value returns to the range of 26%. M95 and M99 exhibit similar, but less pronounced, changes. Interestingly, the tangent portfolio $f_{\text{short}}(t)$ values are relatively unaffected by the pandemic.




Fig. 10.16 Fractional weighting (10.14) of the shorted assets for the historical long–short domestic portfolios optimizations via Jacobs et al. with $s = 0.04$ (top) and via Lo–Patel with $lev = 0.04$ (bottom). Quartile values, Q_1, Q_2 , and Q_3 , for each time series are presented on the right.

We now examine the effect of dynamic optimization on $f_{\text{short}}(t)$. Fig. 10.17 plots $f_{\text{short}}(t)$ for the six dynamic long–short domestic portfolios, MVP through T99, optimized via Lo–Patel with $lev = 0.04$. In contrast to the historic optimizations of Fig. 10.16, the statistics of the $f_{\text{short}}(t)$ time series for the three minimum-risk dynamic portfolios are virtually identical, as are the statistics for the three dynamic tangent portfolios. The day-to-day variation around the median

value is larger for dynamic optimization than for historical. Under dynamic optimization, there is no hint of the pandemic in the tangent portfolio time series. *Based on our results, it would appear that in the Jacobs et al. and Lo–Patel long–short optimizations, the minimum-risk optimizations seek to establish a constant median value for the fraction, f_{short} , of the total weight of the shorted assets in the portfolio. The same statement holds for the tangent portfolios. The value of this median value changes with the value of s (Jacobs et al.) or lev (Lo–Patel) and with historical versus dynamic simulation.* Fig. 10.18 demonstrates the change in quartile values under changing values of lev for the dynamic Lo–Patel optimizations of the domestic portfolio.




Fig. 10.17 Fractional weighting (10.14) of the shorted assets for the dynamic Lo–Patel long–short domestic portfolio optimizations with $\text{lev} = 0.04$. Quartile values, Q_1 , Q_2 and Q_3 , for the time series are presented on the right.




Fig. 10.18 Quartile values, Q_1 , Q_2 and Q_3 , for $f_{\text{short}}(t)$ under changing values of lev for the dynamic Lo–Patel long–short domestic portfolio optimizations.

10.3 Risk Budgets: Incremental and Component Risk

A *risk budget* decomposes the total portfolio risk into the risk contribution of each component asset. It is the basis of *risk budgeting*, a portfolio-management technique with an optimization goal of ensuring that target risk-budget constraints are met for different assets, or classes of assets, in the portfolio. Changes in a portfolio alter investment risk. The change in risk accompanying the change in some factor in a portfolio is referred to as *incremental risk*. Measures of incremental risk aid in risk–return decision-making. Closely aligned with the concept of incremental risk is the concept of marginal risk as well as the decomposition of a portfolio into its constituent (or *component*) risks.⁸⁰ Such a decomposition is useful for identifying both high and low sources of risk, setting position limits, determining capital requirements, etc. Portfolio standard deviation (Std), VaR, and CVaR are the most popular measures of risk used in such budgeting strategies. Std quantifies the contributions of each asset to the center risk of the portfolio, whereas VaR and CVaR quantify the contributions to tail risk. Chow and Kritzman (2001), Litterman (1996),

⁸⁰ The nomenclature for this incremental/marginal/component risk is inconsistent in the literature.

Maillard et al. (2010), and Peterson and Boudt (2008) have investigated the use of Std and VaR in risk budgeting. Boudt et al. (2013) reviews risk budgeting based on CVaR (also known as expected tail loss).

Below, we illustrate these three risk-management measures by applying them to select portfolios of REIT ETF assets developed in Chapters 4 and 5. In our analysis, we use Std to quantify center risk and VaR as the proxy measure of tail risk. Thus, our discussion centers on incremental Std (IStd), marginal Std (MStd), component Std (ciStd),⁸¹ incremental VaR (IVaR), marginal VaR (MVaR), and component VaR (ciVaR). (See, e.g., Chapter 11 of Dowd (2005) for a further discussion of incremental and component risk.)

10.3.1 Incremental, Marginal, and Component VaR

In this section, we develop the mathematical concepts behind incremental, marginal, and component risk measures. For brevity, we use VaR as the illustrative risk measure. Developments for Std or CVaR can be inferred from the discussion here, and they are discussed in Bruder and Roncalli (2012), among other studies. Consider increments in VaR accompanying a change in the weighting of assets in a portfolio p . Let $\mathbf{w} = \{w_i\}$, $i = 1, \dots, N$ denote the position of a fully invested portfolio of N assets. Let $\boldsymbol{\delta} = \{\Delta w_1, \Delta w_2, \dots, \Delta w_N\}$ denote a set of asset-weight changes. Consider the changed position $\hat{\mathbf{w}} = \mathbf{w} + \boldsymbol{\delta} = \{w_1 + \Delta w_1, w_2 + \Delta w_2, \dots, w_N + \Delta w_N\} \equiv \{\hat{w}_i\}$, where the weights \hat{w}_i also satisfy the required portfolio constraints (e.g., $\sum_{i=1}^N \hat{w}_i = 1$; $0 \leq \hat{w}_i \leq 1$ for long-only investing; etc.). Then, the *incremental value-at-risk* (IVaR) at confidence level α , of the portfolio is the change

$$\text{IVaR}_\alpha(\mathbf{w}) = \text{VaR}_\alpha(\hat{\mathbf{w}}) - \text{VaR}_\alpha(\mathbf{w}). \quad (10.15)$$

Note that this general definition (10.15) makes no assumption regarding the particulars or size of the change $\boldsymbol{\delta}$.

One common practice is to consider IVaR_α as a result of adding or subtracting a single asset from the portfolio.⁸² A second practice is to consider (10.15) as the result of a “small” change in a single asset weight. If the position change $\boldsymbol{\delta}$ is “small” (e.g., $|\Delta w_i| \ll 1$, $i = 1, \dots, N$), then Taylor’s series, truncated to first order, gives a reasonably good approximation of (10.15):

$$\text{IVaR}_\alpha(\mathbf{w}) = \text{VaR}_\alpha(\hat{\mathbf{w}}) - \text{VaR}_\alpha(\mathbf{w}) \approx \sum_{i=1}^N \frac{\partial \text{VaR}_\alpha(\mathbf{w})}{\partial w_i} \Delta \hat{w}_i, \quad (10.16)$$

where $\Delta \hat{w}_i = \hat{w}_i - w_i$, $i = 1, \dots, N$. The vector of derivatives,

$$\nabla \text{VaR}_\alpha(\mathbf{w}) = \left\{ \frac{\partial \text{VaR}_\alpha(\mathbf{w})}{\partial w_1}, \frac{\partial \text{VaR}_\alpha(\mathbf{w})}{\partial w_2}, \dots, \frac{\partial \text{VaR}_\alpha(\mathbf{w})}{\partial w_N} \right\}, \quad (10.17)$$

⁸¹ Although it is tempting to use the abbreviations “CoStd” and “CoVaR” for component Std and VaR, “CoVaR” is already commonly used to refer to the VaR of a financial system conditioned on institutions being under distress.

⁸² VaR change as a result of removing or adding a single asset has been used in the literature either as the definition of incremental VaR or as the definition of marginal VaR.

is known as the gradient of $\text{VaR}_\alpha(\mathbf{w})$ evaluated at the value \mathbf{w} . The i th element of the gradient vector,

$$\text{MVar}_\alpha^{(i)}(\mathbf{w}) = \frac{\partial \text{VaR}_\alpha(\mathbf{w})}{\partial w_i}, \quad (10.18)$$

is known as the *marginal value-at-risk* (MVar) with respect to asset i . As (10.16) shows, IVaR and MVar values are inherently linked under small weight changes in the portfolio.

As a function operating on a random variable X (e.g., the return on a portfolio), $\text{VaR}_\alpha(X)$ has the property of linear positive homogeneity; specifically, if $\lambda > 0$ is a positive constant, then $\text{VaR}_\alpha(\lambda X) = \lambda \text{VaR}_\alpha(X)$. Loosely speaking, this property states, for example, that if you double the size of your investment, then you double your risk. Because the return $r_p(t)$ of a portfolio is a linear sum of the weighted returns of each asset, $r_p(t) = \sum_{i=1}^N w_i(t)r_i(t)$, this linear homogeneous property holds relative to the weights; that is, $\text{VaR}_\alpha(\lambda r_{p,t}) = \text{VaR}_\alpha(\sum_{i=1}^N \lambda w_i(t)r_i(t)) = \lambda \text{VaR}_\alpha(r_p(t))$. Consequently, a theorem by Euler⁸³ for linear homogeneous functions ensures that

$$\text{VaR}_\alpha(r_p(t)) = \sum_{i=1}^N \frac{\partial \text{VaR}_\alpha(r_p(t))}{\partial w_i(t)} w_i(t) = \sum_{i=1}^N \text{MVar}_\alpha^{(i)}(\mathbf{w}(t)) w_i(t). \quad (10.19)$$

Equation (10.19) expresses the VaR of a portfolio as a sum of component risks. Each element of this sum is referred to as a component VaR (ciVaR):

$$\text{ciVaR}_\alpha^{(i)}(r_p(t)) = \text{MVar}_\alpha^{(i)}(\mathbf{w}(t)) w_i(t). \quad (10.20)$$

Thus, the concepts of ciVaR and MVar values are inherently linked. Using (10.20), equation (10.19) can be rewritten as

$$1 = \sum_{i=1}^N \frac{\text{ciVaR}_\alpha^{(i)}(r_p(t))}{\text{VaR}_\alpha(r_p(t))}, \quad (10.21)$$

which represents the ciVaR for each asset i as a fraction of the total portfolio VaR.

Gourieroux et al. (2000) consider the CVaR version of (10.19) (replacing $\text{VaR}_\alpha(r_p(t))$ with $\text{CVaR}_\alpha(r_p(t))$). Each element of the sum then becomes a component CVaR (ciCVaR). They show that

$$\text{ciCVaR}_\alpha^{(i)}(r_p(t)) = -E \left[w_i(t)r_i(t) | r_p(t) = -\text{VaR}_\alpha(r_p(t)) \right]. \quad (10.22)$$

Equation (10.22) states that $\text{ciCVaR}_\alpha^{(i)}(r_p(t))$ is (the negative of) the asset's expected contribution to $r_p(t)$ when the portfolio return is equal to (the negative of) the portfolio's VaR.⁸⁴

10.3.2 Computing VaR, IVaR, MVar, and ciVaR

⁸³ One of many theorems developed by the prolific 18th-century mathematician Leonhard Euler.

⁸⁴ All the negative signs in (5.22) stem from the fundamental definition of VaR_α as having a positive value to represent a loss.

For a portfolio, $\text{VaR}_\alpha(\mathbf{w}(t))$ is generally estimated from a distribution of *synthetic* portfolio returns, $r_{p,1}(t), \dots, r_{p,m}(t)$. If computed from historical data, $\text{VaR}_\alpha(\mathbf{w}(t))$ is estimated from the time series of *synthetic* portfolio returns, $r_p(t-m), \dots, r_p(t-1)$:

$$r_{p,k}(t) = r_p(t-k) = \sum_{j=1}^N w_j(t) r_j(t-k) \text{ for } k = 1, \dots, m, \quad (10.23)$$

where $r_j(t-k)$ is the historical return of asset j at time $t-k$ and $\{w_j(t)\}$ are the weights applied to the portfolio assets at time t . If computed using Monte-Carlo simulations,

$$r_{p,k}(t) = \sum_{j=1}^N w_j(t) r_{j,k}(t) \text{ for } k = 1, \dots, m, \quad (10.24)$$

where $r_{j,k}$ is the generated return of asset j during the k th Monte-Carlo simulation. There are generally two classes of methods for estimating $\text{VaR}_\alpha(\mathbf{w}(t))$ (see (3.25)) from the distribution of *synthetic* returns. The most straightforward, if the window size m is sufficiently large, is to compute the VaR value as the negative,

$$\text{VaR}_\alpha(\mathbf{w}(t)) = -r_{p,\alpha}(t), \quad (10.25)$$

of the $1-\alpha$ quantile value of the return distribution. Otherwise, parametric methods can be used by employing a distribution model (e.g., Gaussian), particularly one (e.g., log-normal, generalized Pareto) that allows for a more realistic fit to the fat-tailed nature of the loss tail of the return distribution. If the form of $\text{VaR}_\alpha(\mathbf{w})$ is known analytically for the distribution, it can then be estimated using the parameters that provide the best fit to the distribution of returns $r_{p,1}(t), \dots, r_{p,m}(t)$. For example, if the asset returns in (10.23) or (10.24) are normally distributed with mean vector $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$, then (Haugh et al., 2017; Boudt et al., 2008; McNeil et al., 2005)

$$\text{VaR}_\alpha(\mathbf{w}) = -\boldsymbol{\mu}^T \mathbf{w} - z_N(\alpha) \sqrt{\mathbf{w}^T \boldsymbol{\Sigma} \mathbf{w}}, \quad (10.26)$$

where $z_N(\alpha) = \sqrt{2} \text{erf}^{-1}(2\alpha - 1)$ is the $1-\alpha$ quantile value of the normal distribution, $N(0,1)$. Some parametric estimators include higher-order moments to improve tail estimates of values for VaR. For example, Zangari (1996) and Favre and Galeano (2002) provide a VaR estimate using the Cornish–Fisher expansion to include skewness, S , and kurtosis, K , through third order in the probabilist's Hermite polynomials. The result of their estimate leads to the form (10.26), but with $z_N(\alpha) \rightarrow z_{\text{CF}}(\alpha)$,

$$\text{VaR}_\alpha(\mathbf{w}) = -\boldsymbol{\mu}^T \mathbf{w} - z_{\text{CF}}(\alpha) \sqrt{\mathbf{w}^T \boldsymbol{\Sigma} \mathbf{w}}, \quad (10.27a)$$

where

$$z_{\text{CF}}(\alpha) = He_1(z_N(\alpha)) + \frac{He_2(z_N(\alpha))S}{6} + \frac{He_3(z_N(\alpha))K}{24} - \frac{[2He_3(z_N(\alpha)) + He_1(z_N(\alpha))]S^2}{36} \quad (10.27b)$$

$$= z_N(\alpha) + \frac{(z_N(\alpha)^2 - 1)S}{6} + \frac{(z_N(\alpha)^3 - 3z_N(\alpha))K}{24} - \frac{[2z_N(\alpha)^3 - 5z_N(\alpha)]S^2}{36}.$$

The parametric estimates (10.26) and (10.27) explicitly show that the covariance structure (which captures both volatility and correlation effects) of the asset returns impacts VaR value (and consequently IVaR, MVar, and ciVaR values⁸⁵).

Note that estimation of VaR can be susceptible to outlier values, particularly for estimates, such as (10.27), that use higher moments (Cont, Deguest, and Scandolo, 2007). Methods for cleaning outliers, such as winsorization (see, e.g., Boudt, Peterson, and Croux, 2008), can be employed to improve robustness.

Computation of IVaR from the general form (10.15) can be performed with any of these VaR-estimation methods. Equation (10.15) involves evaluating differences of the form $\text{VaR}_\alpha(\hat{\mathbf{w}}) - \text{VaR}_\alpha(\mathbf{w})$, in which both the changed portfolio weights \hat{w}_i and the original weights w_i must satisfy portfolio constraints. One practical approach is to distribute the weight change proportionally to the risk represented by each individual asset, giving less risky assets a greater portion of the weight change, that is, proportionally to asset ciVaR values. This weight distribution technique is the basis of the risk budgeting discussed in section 10.2.3.

For small weight changes, computation of IVaR from (10.16) requires estimation of the MVar values. This can be done through numerical derivatives, such as the first-order derivative estimate

$$\text{MVar}_\alpha^{(i)}(\mathbf{w}) \approx \frac{\text{VaR}_\alpha(\hat{\mathbf{w}}^{(i)}) - \text{VaR}_\alpha(\mathbf{w})}{\Delta w_i}, \quad (10.28)$$

where $\hat{\mathbf{w}}^{(i)} = \{w_1, \dots, w_{i-1}, w_i + \Delta w_i, w_{i+1}, \dots, w_N\}$. The $\text{MVar}_\alpha^{(i)}$ values in (10.19) may be computed under, for example, a 1% change in each respective weight, $\Delta w_i = 0.01, i = 1, \dots, N$ (irrespective of the values of the small changes, $\Delta \hat{w}_i$, to be used in evaluating $\text{IVaR}_\alpha(\mathbf{w})$ in (10.16)). However, numerical estimation via (10.28) is prone to roundoff error as $\Delta w_i \rightarrow 0$. Generally, the more robust approach is to use a parametric method to estimate values of $\nabla \text{VaR}_\alpha(\mathbf{w})$. From either (10.26) or (10.27), we have

$$\nabla \text{VaR}_\alpha(\mathbf{w}) = \boldsymbol{\mu} + z(\alpha) \frac{\boldsymbol{\Sigma} \mathbf{w}}{\sqrt{\mathbf{w}^T \boldsymbol{\Sigma} \mathbf{w}}}, \quad (10.29)$$

where $z(\alpha) = z_N(\alpha)$ or $z(\alpha) = z_{\text{CF}}(\alpha)$, respectively. Equation (10.29) can be used to provide parametric estimates for $\text{MVar}_\alpha^{(i)}(\mathbf{w})$ (see (10.18)), $\text{IVaR}_\alpha(\mathbf{w})$ (see (10.16)), and $\text{ciVaR}_\alpha^{(i)}(\mathbf{w})$ (see (10.20)).

For the case in which IVaR is to be computed as a result of the removal (similarly for the addition) of a single asset, Mina and Xiao (2001) develop the parametric estimate

⁸⁵ An asset with a large position size may have a small ciVaR due to volatility or correlation effects; similarly, an asset with large volatility may have a small ciVaR due to position size or correlation effects. Analogous statements might hold for assets with small position or volatility.

$$\begin{aligned}
& \text{VaR}_\alpha(r_p) - \text{VaR}_\alpha(r_p - r_j) \\
&= [\text{VaR}_\alpha^2(r_p - r_j) + \text{VaR}_\alpha^2(r_j) + 2\rho \text{VaR}_\alpha(r_j)\text{VaR}_\alpha(r_p - r_j)]^{\frac{1}{2}} \\
&\quad - \text{VaR}_\alpha(r_p - r_j) \\
&= \text{VaR}_\alpha(r_j) \frac{1}{\xi} [\sqrt{\xi^2 + 2\rho\xi + 1} - 1].
\end{aligned} \tag{10.30}$$


Here, $r_p - r_j$ represents the portfolio return without asset j ; $\text{VaR}_\alpha(r_j)$ is the VaR value for the asset return r_j ; $\text{VaR}_\alpha(r_p - r_j)$ is the VaR value for the portfolio return computed without asset j ; ρ is the correlation between asset j and the rest of the portfolio $r_p - r_j$; and $\xi = \text{VaR}_\alpha(r_j)/\text{VaR}_\alpha(r_p - r_j)$.

If Std is used as the risk measure, then computations for Std, MStd, and ciStd are (Bruder and Roncelli, 2012)

$$\text{Std}(\mathbf{w}) = \sqrt{\mathbf{w}^T \boldsymbol{\Sigma} \mathbf{w}}, \quad \text{MStd}(\mathbf{w}) = \frac{\boldsymbol{\Sigma} \mathbf{w}}{\sqrt{\mathbf{w}^T \boldsymbol{\Sigma} \mathbf{w}}}, \quad \text{ciStd}^{(i)}(\mathbf{w}) = w_i(t) \frac{(\boldsymbol{\Sigma} \mathbf{w})_i}{\sqrt{\mathbf{w}^T \boldsymbol{\Sigma} \mathbf{w}}}. \tag{10.31}$$

10.3.3 Portfolio Results

Component risk. We start with a discussion of the results for the computation of ciStd and VaR_α . Unless noted otherwise, all VaR_α computations are performed at the 95% quantile level ($\alpha = 0.95$). Because component risk (see (10.20) and (10.31)) is a product of an asset's weight with its marginal risk, computing component risk for an equal-weighted portfolio gives an indication of the *inherent* risk provided by each asset, whereas computing component risk for a weight-optimized portfolio gives an indication of the *actual* risk of each asset. It is useful to evaluate inherent risk when considering the addition of a new asset to a portfolio. We contrast these two approaches by considering the long-only domestic portfolios: EQW, MVP subject to $C_{\text{TO}} = 0.04$, and TVP subject to $C_{\text{TO}} = 0.04$. We consider both center risk (ciStd) and tail risk (ciVaR).



10. Risk Information and Management




Figure 10.19 Risk-budget *ciStd* (top) and *ciVaR* (bottom) values as a function of time for the 26 assets of the long-only domestic portfolio EQW. The top-down ordering of assets is the same in both plots.

When portfolio managers use risk budgeting, they are interested primarily in the component risk of each asset measured over a recent time interval (i.e., the *risk budgets*). Fig. 10.19 provides a graphical view of the change in risk budgets, measured in terms of *ciStd* and *ciVaR*, as a function of time for the long-only domestic portfolio EQW. The plots are qualitatively similar. Using boxplots showing minimum, Q_1 , Q_3 , and maximum values, Fig. 10.20 presents the statistical differences between the time variations of *ciStd* and *ciVaR* over this 15-year period. The average contributions of assets vary from 2.2% to 5.0% for *ciStd* and 2.4% to 4.9% for *ciVaR*. The IQRs, which vary from 0.05% to 1.1% for *ciStd* and 0.12% to 1.1% for *ciVaR*, and the maximum–minimum differences, which vary from 0.56% to 7.6% for *ciStd* and 0.53% to 6.2% for *ciVaR*, provide perspective on asset *ciVaR* variation over time. The assets are ordered from left to right by decreasing Q_2 value. Between *ciStd* and *ciVaR*, there are minor local variations in the ordering




Figure 10.20 Boxplot statistics of the risk budgets of *ciStd* and *ciVaR* for the long-only domestic portfolio EQW. (The red line indicates the value $1/26$; \times indicates minimum and maximum values; — indicates Q_1 and Q_3 values.)

of the assets. With 26 assets in the portfolio, the overall average value of $1/26 = 0.385$ (red line) can be used to divide the assets into two classes, that is, greater and lesser risk contributors. These risk measures identify the same sets of assets in each class. Three assets in particular, SBAC, AMT, and CCI, show the greatest change in *ciVaR* over time, changing from large-risk contributors during 2018 to low-risk contributors for the remaining 14 years.

In contrast to Fig. 10.20, Fig. 10.21 presents the component risk of each asset computed over the entire 15-year period. With the exception of SBAC under *ciVaR*, the ordering is fairly consistent with that presented in Fig. 10.20. Also plotted are the Q_1 and Q_3 values from Fig. 10.21, which show how the 15-year values align with the IQR ranges. Alignment is generally good, with the exception of the three assets, SBAC, AMT, and CCI, whose component risk changes considerably over the time period.




Figure 10.21 Risk-budget *ciStd* (top) and *ciVaR* (bottom) values for the long-only domestic portfolio EQW computed for the entire 15-year period. (The red line indicates the value $1/26$; the dashed values are the Q_1 and Q_3 values from Fig. 10.20.)




Figure 10.22 Risk-budget $ciStd$ values as a function of time for the 26 assets of the MVP long-only domestic portfolio optimized subject to a 4% turnover constraint. Asset colors are identical to those of Fig. 10.19.

We now consider the actual component risk of each asset in the optimized portfolios. To provide a contrast with Figs. 10.19–10.21, we consider component risk for the MVP and TVP long-only domestic portfolios optimized subject to a 4% daily turnover constraint. Because the risk measure in both of these optimizations is portfolio variance, we concentrate on asset $ciStd$ values and consider $ciVaR$ values secondarily. Fig. 10.22 plots a graphical view of the change in $ciStd$ values as a function of time for this portfolio. Because asset weights differ drastically, both from each other and over time, $ciStd$ values are dramatically different from those displayed in Fig. 10.19 for the EQW portfolio. Now, *CCI*, *IRM*, *AMT*, *NLY*, and *PSA* dominate the risk over long periods of time, whereas they are minimal inherent risk contributors in the EQW portfolio, which is exactly why they receive heavy weightings in the minimum-variance optimization of MVP. The results provide insight into how the mean-variance optimization adjusts to changing market conditions. Whereas the $ciStd$ changes for the equal-weighted portfolio show a distinction between 2008 and the remaining years, the $ciStd$ changes for the MVP fall into roughly six time periods: 2008, 2009–2011 (Great Recession), 2011–2016, 2017–2019(Q2), 2019(Q3)–2020(Q1), and 2020(Q2–Q4) (COVID-19 pandemic). Fig. 10.23 summarizes the minimum, maximum, Q_1 , and Q_3 values for $ciStd$ for each asset over this period. (For comparison, the $ciVaR$ is also shown. The plots are very similar.) Assets are presented from left to right in order of decreasing maximum value. The large IQRs and/or maximum–minimum values for the assets *NLY*, *IRM*, *PSA*, and *AMT* reflect the time-varying values shown in Fig. 10.22.




Figure 10.23 Boxplot statistics of risk-budget $ciStd$ and $ciVaR$ values for the MVP long-only domestic portfolio optimized subject to a 4% turnover constraint. (The red line indicates the value $1/26$; \times indicates minimum and maximum values; $-$ indicates Q_1 and Q_3 values.)

It is instructive to examine the correlation between asset weights and asset component risk values under a historical long-only investing scheme. Clearly, if an asset has very little to no weight in a portfolio, it will have little to no component risk; similarly, an asset with a large weight will have a larger component risk. In addition, the portfolio-optimization techniques we use here seek to




Figure 10.24 The difference, asset weight – asset ciStd, for two consecutive dates for the MVP long-only domestic portfolio optimized subject to a 4% turnover constraint.

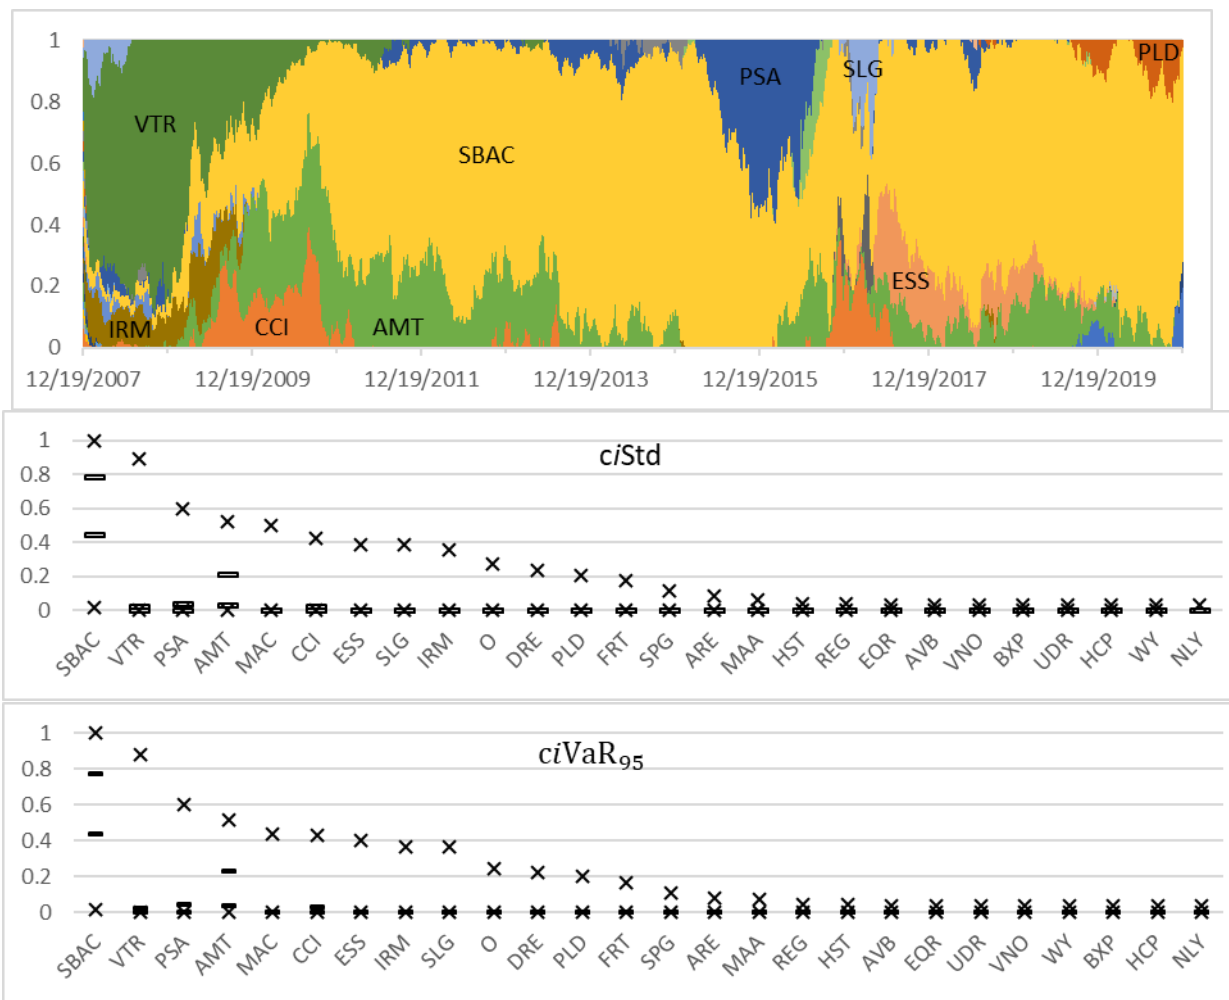


Figure 10.25 ciStd values as a function of time for the 26 assets of the TVP long-only domestic

portfolio optimized subject to a 4% turnover constraint (top). Asset colors are identical to those of Figs. 10.19 and 10.22. Boxplot statistics of asset $ciStd$ (middle) and $ciVaR_{95}$ (bottom) values. (\times indicates minimum and maximum values; $-$ indicates Q_1 and Q_3 values.)

reduce risk by reducing the weights of riskier assets. For the (nonoptimizing) equal-weighted domestic portfolio, each asset weight is $1/26$; thus, the statistical differences between each asset's $ciStd$ and weight can be inferred directly from the difference between the red line and the boxplot values presented in Fig. 10.20. We contrast these results with those of the MVP portfolio. Fig. 10.12 displays the asset weights versus time for the MVP portfolio. The differences between asset weight and the $ciStd$ values plotted in Fig. 10.22 are very subtle. Fig. 10.24 emphasizes this by plotting the difference $w_i(t) - ciStd^{(i)}(t)$; $i = 1, \dots, 26$ for two successive days, 3/11/2020 and 3/12/2020, at the onset of the pandemic-related market disruption. Differences are $\leq \pm 0.6\%$.

We contrast the component risk of the tangent portfolio optimization, TVP, with that of both the minimum-risk portfolio optimization, MVP, and the equal-weighted portfolio, EQW. Fig. 10.25 displays asset $ciStd$ versus time and boxplot statistics for the asset $ciStd$ and $ciVaR$ values for the TVP long-only domestic portfolio subject to a 4% turnover constraint. The SR-maximizing tangent portfolio optimizes assets very differently than the minimum mean-variance portfolio. The EQW portfolio $ciStd$ and $ciVaR$ values shown in Fig. 10.20 indicate that NLY, IRM, CCI, AMT, SBAC, WY, O, PSA, MAA, and ESS are the 10 lowest-risk-contributing assets. The MVP and TVP optimizations weight these differently over this 15-year period. Although all 10 of these assets play a nontrivial role in the MVP portfolio (Fig. 10.22), the greatest prominence is given to NLY, IRM, and AMT. In contrast, only six of these play a nontrivial role in the TVP portfolio (Fig. 10.25), with the greatest prominence given to SBAC.




Figure 10.26 Boxplot statistics of the Sharpe ratios for the individual assets in the domestic portfolio over the time period 12/19/2007 through 12/18/2020. (\times indicates minimum and maximum values; $-$ indicates Q_1 and Q_3 values.)

To illustrate the effect of TVP optimization, Fig. 10.26 displays the boxplot statistics of the SRs of the individual assets over this time period. (These would be the asset SRs in an equal-weighted portfolio.) For consistency, the SRs are computed using the same 2,016-day moving window used in the MVP and TVP optimizations. Comparison of Fig. 10.26 with Fig. 10.25 demonstrates that the TVP optimization gives greater weight to those assets having larger SRs over this period.

Incremental VaR. Dowd (2005) classifies IVaR performance into three categories, which we summarize in terms of the discussion here.

- (1) High IVaR: An asset with a large positive IVaR value indicates that the increased weight (i.e., “increased position”) of this asset in the portfolio adds substantially to portfolio risk. If the asset position in the portfolio increases further, the IVaR associated with this asset will likely grow at an increasing rate.
- (2) Moderate IVaR: An asset with a moderate IVaR value indicates that the increased position of this asset adds moderately to portfolio risk. However, as the asset position continues to grow, its IVaR will likely grow at an increasing rate.
- (3) Negative IVaR: An asset with a negative IVaR value indicates that the increased position of this asset reduces overall portfolio risk; that is, the new position is a natural hedge against the existing portfolio. However, if the relative position of this asset in the portfolio grows too large, its hedging ability will likely be diminished by decreasing portfolio diversification, and ultimately the IVaR associated with this increasing position must become positive. There is thus an optimal value of weight at which such an asset provides maximum hedging effect (a “best hedge”).

We illustrate these categories of performance with two computations of IVaR: (i) the incremental change in portfolio VaR as a result of the daily weight changes in the portfolio and (ii) the incremental change in portfolio VaR that would accompany the removal of each asset.




Figure 10.27 $\text{VaR}_{95}(r_p(t))$ and $\text{IVaR}_{95}(r_p(t))$ for the EQW (top) long-only domestic portfolio and the MVP (middle) and TVP (bottom) long-only domestic portfolios optimized subject to a 4% turnover constraint.

Fig. 10.27 shows the daily values of portfolio VaR^{86} and the daily incremental change in VaR, $[\text{VaR}_{95}(r_p(t)) - \text{VaR}_{95}(r_p(t-1))]/\text{VaR}_{95}(r_p(t-1))$ over the 15-year time period for the domestic portfolio using (i) the EQW strategy, (ii) the MVP optimization subject to a 4% turnover constraint, and (iii) the TVP optimization subject to a 4% turnover constraint. VaR values are steadiest for the MVP portfolio and most variable for TVP. The EQW and MVP portfolios show the strongest correlation with market events: with the rising VaR during the 2008 financial crisis, with the decreasing VaR over Q3 and Q4 of 2016, and with the increasing VaR during the 2020 pandemic. The first two trends are more protracted for the TVP portfolio than for the others. IVaR is consistently higher (more negative) for the TVP portfolio and consistently lowest (least negative) for MVP.

Fig. 10.28 shows the percent change, $[\text{VaR}_{95}(r_p) - \text{VaR}_{95}(r_p - r_i)]/\text{VaR}_{95}(r_p)$, in $\text{IVaR}^{(i)}$ computed using (10.30) that results from removing each asset i from the EQW portfolio. The width of each color band represents the percent change in $\text{VaR}_{95}(r_p)$ produced by removing that asset. Note that a positive change in Fig. 10.28 indicates that *removal* of the asset *decreases* the portfolio




Figure 10.28 Percent change in $\text{VaR}_{95}(r_p(t))$ produced by removing each asset in the long-only domestic EQW portfolio. The width of each band indicates the percent change. Asset colors are identical to those of Figs. 10.19, 10.22, and 10.25.

VaR or that, conversely, *adding* the asset *increases* the portfolio VaR. This sign convention allows for direct association with Dowd's high, moderate, and negative categories. In terms of Dowd's scheme, an overall classification of the assets from the EQW portfolio might be the following:

- high: SLG, VNO, PLD, HST, MAC, REG, DRE
- moderate: AVB, BXP, ARE, EQR, UDR, SPG, VTR, HCP, SPG
- negative: SBAC, IRM, AMT, NLY, CCI, WY, O, PSA, MAA, ESS

These results confirm that SBAC, IRM, AMT, NLY, and CCI are the principal risk-diversifying

⁸⁶ Following convention, VaR is plotted as a negative value for losses in Fig. 10.27 rather than as the positive value defined in (3.21).

assets (in each case, adding the asset decreases the portfolio VaR). No asset in Dowd's high category produces as strong a magnitude of VaR change as do the main risk diversifiers. Note that Dowd's asset classification varies over time. For example, SBAC is a strong risk diversifier over the 12-year period 2010–2020 and a strong risk contributor (high category) during 2008, and other asset IVaR values (e.g., CCI, PSA, ARE) change their sign over time. Because all the assets in EQW have equal weight, Fig. 10.28 also sheds light on possible hedging strategies, for example, by balancing the addition of a risk-contributing asset, such as PLD, with the removal of an equal risk-contributing asset, such as SLG.




Figure 10.29 Time series of IVaR values produced by the daily removal of each asset in the MVP (left) and TVP (right) long-only domestic portfolio optimizations subject to a 4% turnover constraint. The width of each band indicates the value in percent. Assets colors are identical to those of Fig. 10.28.

Fig. 10.29 plots the IVaR values that would result from removing individual assets each day from the MVP and TVP long-only domestic portfolio optimizations subject to a 4% turnover constraint. The IVaR plot for the MVP portfolio reemphasizes that this portfolio concentrates on those assets that produce the greatest risk diversification. One exception is the presence of the risk contributor SBAC starting in 2016. For the TVP portfolio, the exception is more pronounced: SBAC, which represents the largest component of VaR in the portfolio, is a positive risk contributor for most of the period. The same is true for VTR in 2008. Fig. 10.29 shows that the TVP optimization balances IVaR more effectively than MVP because critical assets under TVP are weighted such that their removal would generate opposite-signed IVaR values over the 15-year period. This rarely occurs for assets under MVP optimization.

10.4 Factor Analysis

It can be difficult to determine why a portfolio is performing better or worse than the market. Factor analysis enables an investor to pinpoint factors (the underlying exposures) that are contributing most strongly to the observed return performance. The principal objective addressed by a financial factor model can be described as follows:

“Given a time series of returns for a set of assets, quantify the relationship between the observed asset returns and the time series of the returns (or ‘scores’) of a set of latent factors that jointly (‘in common’) influence the asset returns.”

Inherent in the usefulness of factor analysis is the assumption that the number of these common latent factors is much smaller than the number of assets in the portfolio. In practice, financial factor analysis usually assumes a multilinear relationship between the portfolio return and the factor returns, resulting in a relative weight of influence (or “loading” or “beta”) of each common

factor on the return of each asset, as well as an estimate of the portion of the return behavior (the latent error) that remains unexplained by the common factors. The most familiar model is the single-factor capital asset pricing model (CAPM) developed by Sharpe (1970):

$$r_{it} = r_{ft} + \beta_i r_{mt} + \epsilon_{it}, \quad i = 1, \dots, N, \quad t = 1, \dots, T.$$

Here, r_{it} is the return of asset i ; r_{ft} is the risk-free rate; the single factor, r_{mt} , is the excess (relative to the risk-free rate) return of the market; β_i is the loading of asset i on the excess market return; and ϵ_{it} is the latent error for asset i .⁸⁷ Note that the loadings are modeled as time independent. CAPM performs a factor analysis assuming a single factor (e.g., a market index) that represents the returns of a market. If r_{it} , on the left-hand side, is the return of a managed portfolio, then a large value of β_i indicates that the portfolio return is tracking the market return well. If β_i is small, then other factors are influencing performance, and a multifactor model is needed.

Let $\mathbf{r}_t = (r_{1t}, \dots, r_{Nt})^T$ be a vector of returns at time t of N assets with (time-averaged) mean vector $\boldsymbol{\mu} = (\mu_1, \dots, \mu_N)^T$. Then, the general multilinear factor model is

$$\mathbf{r}_t = \boldsymbol{\mu} + \boldsymbol{\beta} \mathbf{f}_t + \boldsymbol{\epsilon}_t, \quad t = 1, \dots, T, \quad (10.32)$$

where $\mathbf{f} = (f_{1t}, \dots, f_{mt})^T$, $m < N$, is the vector of returns at time t for each of m common factors; $\boldsymbol{\beta} = \{\beta_{ij}\}$ is the matrix of loadings, with β_{ij} being the load of asset i on the j th factor; and $\boldsymbol{\epsilon}_t = (\epsilon_{1t}, \dots, \epsilon_{Nt})^T$ is the vector of latent errors.

Conceptually, one would like to perform the following analysis: “Here are the daily returns for a managed portfolio; run some analysis to identify a set of factors driving those returns, as well as the loadings and latent errors in model (10.32).” Unfortunately, there is not enough information in the daily returns of a single portfolio to do so. There are, therefore, two approaches to using factor models. In the first, as in the CAPM model or the familiar Fama–French (1993, 2015) three- and five-factor models, *the number of factors and factor returns are known, and the loadings are computed via regression on (10.32)*. This approach addresses the following question: “The set of factors $\{S\}$ is believed to be significant in setting market value. How much of the performance of my portfolio can be explained by the return performance of these specific factors?” In the second approach, as in the BARRA factor model (see Grinhold and Kahn, 2000), only the number of factors, m , is assumed to be known. To use this approach, one has to provide return data on a larger set of “potentially contributing” attributes. Thus, the central question addressed by this second approach is the following: “Is the performance of this attribute data being significantly driven by m latent factors?” The factor analysis will then estimate the loadings, latent errors, and returns for m factors that drive relationships in this return data set. In the second approach, the identity of each latent factor can be surmised only by examining the attribute loadings on the factor. Because the number of factors is unknown, the factor analysis can be run several times assuming different values for m .

Implementing the second approach involves the following additional assumptions. If the latent errors are assumed to be uncorrelated with each other (i.e., $\sum_t \epsilon_{it} \epsilon_{jt} = 0$, for $i \neq j$), then their covariance matrix is diagonal:

⁸⁷ In the CAPM as originally derived by Sharpe, it is assumed that $\epsilon_{it} = 0$, $i = 1, \dots, N$.

$$\text{Cov}(\epsilon_t) = \mathbf{D} = \text{diag}\{\sigma_1^2, \dots, \sigma_N^2\}. \quad (10.33)$$

If the common factors are assumed to be uncorrelated with the latent errors (i.e., $\sum_t f_{at}\epsilon_{jt} = 0$, $a = 1, \dots, m$; $j = 1, \dots, N$), the covariance of the return series \mathbf{r}_t is then

$$\text{Cov}(\mathbf{r}_t) = \boldsymbol{\beta}\boldsymbol{\Sigma}_f\boldsymbol{\beta}^T + \mathbf{D}, \quad (10.34)$$

where $\boldsymbol{\Sigma}_f = \mathbf{f}_t\mathbf{f}_t^T$ is the covariance matrix of the common factors. If the factors are assumed to be uncorrelated and their returns are appropriately normalized (i.e., $\sum_t f_{at}f_{bt} = \delta_{ab}$), then $\boldsymbol{\Sigma}_f = \mathbf{I}$ (the $m \times m$ identity matrix) and

$$\text{Cov}(\mathbf{r}_t) = \boldsymbol{\beta}\boldsymbol{\beta}^T + \mathbf{D}. \quad (10.35)$$

Given the number m of factors, values for β_{ij} and σ_i^2 , $i = 1, \dots, N$, $j = 1, \dots, m$ can be obtained from a maximum-likelihood analysis of (10.35). The factors \mathbf{f} can then be estimated using cross-sectional regression on (10.32) for each time t .⁸⁸

Thus, (10.32) captures the movement of the portfolio return through a predictive linear dependence, $\boldsymbol{\mu} + \boldsymbol{\beta}\mathbf{f}_t$, on the factor returns combined with a factor-independent random component ϵ_t . Equally important, (10.35) describes what fraction of the variability of the portfolio return, $\text{Cov}(\mathbf{r}_t)$, results from the variability of the factor loadings, $\boldsymbol{\beta}\boldsymbol{\beta}^T$, and what fraction results from the latent random component, \mathbf{D} .

We illustrate the second approach by applying it to the returns of the TVP long-only domestic portfolio subject to a 4% turnover constraint. We consider the following as a set of *exogenous* attributes that may be influenced by the same common factors as our REIT portfolio: the stock ETF SPY, the REIT market index WD,⁸⁹ and the 13 market-representative assets listed in section 2.4. These 13 assets are representative of the equity, fixed-income, cash, currency, and commodity classes. The price data for each asset covering the period 12/18/2007 through 12/17/2020 was obtained from Bloomberg Professional Services. We consider an additional set of *endogenous* attributes that consists of the one-day lagged returns of the 26 REIT assets in the domestic portfolio and the one-day lagged return of the example portfolio itself. In part, the purpose of adding the endogenous assets is to determine whether the common factors include predictability of returns within the portfolio.

⁸⁸ See, for example, Linton (2019, section 8.8.1).

⁸⁹ Each pairwise correlation between the REIT indices WD, WP, and FNRE exceeds the value 0.999. Because they are so highly linearly correlated, only WD was chosen for the factor analysis.




Figure 10.30 Plot of the loading vector (component 1, component 2) = (β_{i1}, β_{i2}) in the space defined by the factors \mathbf{f}_1 and \mathbf{f}_2 obtained from a two-factor analysis of 42 assets that are exogenous and endogenous to the TVP long-only domestic portfolio optimized subject to a 4% turnover constraint.

As a first illustration, we consider the entire set of exogenous and endogenous attributes. For graphical purposes, we look for only two factors. Fig. 10.30 plots the loading vectors (β_{i1}, β_{i2}) in the space defined by the factors \mathbf{f}_1 and \mathbf{f}_2 obtained from a two-factor analysis of $i = 1, \dots, 42$ assets that are exogenous and endogenous to the example portfolio. Except for the treasury-rate-tracking attributes, SPBDUB3T and TIP, the data has a strong positive orientation with one of the two factors. Fig. 10.31 uses a bar graph to present the same data; the top plot shows the component-loading values β_{i1} and the bottom shows the values β_{i2} . There are two advantages to the bar plot presentation: it is extendable to more than two factors, and it is easy to identify the individual assets. Fig. 10.31 identifies the two factors as \mathbf{f}_1 : endogenous data and \mathbf{f}_2 : exogenous data. It is also clear that \mathbf{f}_2 comprises the class of exogenous data, excepting those assets heavily weighted to treasury rates (AGG SPBUD3T, TIP) and the gold commodity (XAU).




Figure 10.31 Bar plots of the component-loading data, β_{i1} (top) and β_{i2} (bottom), in Fig. 10.30 (top).

The specific variances for each attribute (bottom).

Fig. 10.31 also plots the specific variances (the diagonal elements σ_t^2 of the latent random component \mathbf{D}) for each attribute. A small specific variance indicates that the factors are explaining a great deal of the observed variance for that attribute (e.g., SPY, RUA, VBINX). In contrast, the factor model is explaining little of the observed variance for gold and the three treasury-rate-related attributes. The *commonality* h^2 is defined as the fraction of the total variance in the data set that is explained by the factors. Summing the specific variances and normalizing by the number of attributes gives $1 - h^2$, the fraction of the total variance that is unaccounted for by the factors. For this two-factor model, $h^2 = 62.9\%$.

Because the factor analysis returns a return time series for each identified factor, a regression can be run for the portfolio returns against the factor returns. Running the regression

$$r_p = \bar{r}_p + \alpha_1 f_1 + \alpha_2 f_2$$

produces

$$r_p = 0.00063 - 0.0027f_1 + 0.014f_2 \quad (10.36)$$

with an R^2 value of 55%. The portfolio return has a positive dependence on the exogenous factor and a negative (and 80% smaller) dependence on the endogenous factor. These two factors account for 55% of the variance exhibited by the portfolio-return series. Fig. 10.32 presents a 3D plot of the portfolio-return data compared to the regression plane (10.36).




Figure 10.32 3D plot of the return data for the TVP long-only domestic portfolio optimized subject to a 4% turnover constraint compared to the regression plane (10.36).




Figure 10.33 Bar plots of the component-loading data, β_{i1} , β_{i2} , and β_{i3} (top to bottom), resulting from a three-factor model run on the exogenous and endogenous attributes (top). The specific variances for each attribute (bottom).

Exploring a three-factor analysis produces the component loadings and specific variances shown in Fig. 10.33. Again, we identify the first two factors, respectively, as “endogenous” and “exogenous – with the exception of gold and treasury rates.” The third latent factor emerging from the endogenous data has strong loadings on the one-day lagged REITs CCI, AMT, and SBAC, as well as on the one-day lagged portfolio returns. Examining the business descriptions in section 2.1.1 leads to the conclusion that *factor 3 has an association with wireless infrastructure REITs*. Regressing the portfolio return against these three factors produces

$$r_p = 0.00063 - 0.0025f_1 + 0.014f_2 - 0.0012f_3 \quad (10.37)$$

with an R^2 value of 55%. With a coefficient half the size of the first factor, the portfolio return has a negative correlation with this third factor.




Figure 10.34 Bar plots of the component-loading data, β_{i1} to β_{i6} (top to bottom), resulting from a six-factor model run on the exogenous and endogenous attributes (top). The specific variances for each attribute (bottom).

Fig. 10.34 displays the results for a six-factor analysis. The regression of the portfolio returns against these factor returns results in

$$\begin{aligned} r_p = & 0.00063 - 0.0026f_1 + 0.014f_2 - 0.00085f_3 - 0.00066f_4 - 0.000845f_5 \\ & - 0.0021f_6 \end{aligned} \quad (10.38)$$

with an R^2 value of 55%. The first three factors correspond to those identified in the three-factor analysis. The fourth factor is exogenous and has dominant loadings by the treasury-rate-related attributes AGG and TIP (though, interestingly, not SPBDUB3T). Thus treasury-rate dependence enters as a lower-influence factor. The fifth factor is an interesting combination of (i) retail REITS (MAC, FRT, REG, and SPG) with negative loadings on the factor, (ii) two REITS (VNO and SLG), also with negative loadings, that have the common feature of investing in office real estate in New York City, and (iii) residential and residentially correlated REITS (AVB, EQR, ESS, MAA, PSA, and UDR) with positive loadings. The sixth factor has positive dominant loadings by EEM and EFA (international markets) and by USO and XAU (international commodities) as well as negative loading attributes by a set of attributes whose commonality is more obscure.

This process can be iterated to search for additional latent factors. Objective tests (the Kaiser criterion (Kaiser, 1960); Horn's parallel analysis (Horn, 1965); and Velicer's MAP test (Velicer,

1976)) have been suggested for determining the number of factors. Because they are based on different criteria, however, they may not agree on the number of factors. A simple threshold method consisting of finding enough factors to achieve a predetermined value of the commonality h^2 could be employed, but it is generally difficult to establish a predetermined threshold. A compromise is to employ a *scree plot* (Cattell, 1966), which plots the commonality, h^2 , as a function of the number of factors. Fig. 10.35 presents a scree plot summarizing factor analyses using two through six factors on this data set. The figure implies that little additional explanatory variance will be accomplished by increasing the number of factors much beyond six. An interpretation regarding when the value of h^2 is “leveling off” is subjective and can be susceptible to user bias.




Figure 10.35 Scree plot of the commonality, h^2 , as a function of the number of factors assumed in the analysis of the 42 exogenous and endogenous attributes considered for the TVP long-only domestic portfolio optimized subject to a 4% turnover constraint.

Note that the R^2 values of the regression fits (10.36)–(10.39) do not change with the number of factors, suggesting that the size of the attribute set, rather than the number of factors, must be expanded in order to obtain an improved regression.

We use a small set of 42 exogenous and endogenous factors for this analysis. A reference such as the *Handbook for the Barra Risk Model* (MSCI Barra Inc., 2007) provides some insight into the huge number of potential attributes that can be considered in any factor analysis. Using stepwise regression, possibly followed by computation of pairwise correlations, can be an effective way of winnowing an initial list. We illustrate these two techniques by adding the REIT index FNRE to the set of 42 exogenous and endogenous attributes examined above. We begin by running a stepwise regression of the daily returns of the example portfolio against the returns of these 43 factors. We use the adjusted R^2 criterion, which adds an attribute to the regression only if it increases the regression value of the adjusted R^2 . Doing so reduces the data set to 28 attributes. We then compute the correlation matrix of these 28 attributes. The REIT indices WD and FNRE are highly correlated, with a coefficient value of 0.9990. A linear regression between these two attributes reveals

$$r_{\text{FNRE}} = 0.0000 - 0.9990r_{\text{WD}}$$

with an R^2 value of 99.8% and an error variance of 0.0019. This is a very remarkable linear dependence between two series involving 3,741 daily return values. Including both attributes in the factor analysis will lead to spurious specific variances (Heywood case). We therefore remove FNRE from the data

set and rerun the stepwise regression on the original 42 attributes.⁹⁰ The stepwise regression results in 26 attributes.




Figure 10.36 Bar plots of the component-loading data, β_{i1} to β_{i5} (top to bottom), resulting from a five-factor model run on the reduced set of exogenous and endogenous attributes.

For brevity, we consider the results of running a five-factor analysis on this reduced (by 38%) attribute set. Fig. 10.36 displays the factor loadings. Labeling the five factors from this figure f'_1 through f'_5 to distinguish them from the factors in Fig. 10.34, we see the clear correspondences $f'_1 \leftrightarrow f_1$; $f'_2 \leftrightarrow f_2$; $f'_3 \leftrightarrow f_3$; $f'_4 \leftrightarrow f_5$ (The fact that individual-attribute loadings have the opposite signs between f'_4 and f_5 is not relevant, because the $+/-$ direction of a factor in factor space is arbitrary.) Finally, f'_5 corresponds to f_4 and f_6 . The commonality of this factor analysis is 71%, the same as for the five-factor analysis using all 42 attributes. Regressing the portfolio returns against the five factors, f'_1 to f'_5 , yields

$$\mathbf{r}_p = 0.00063 - 0.0025\mathbf{f}'_1 + 0.014\mathbf{f}'_2 - 0.0014\mathbf{f}'_3 - 0.0012\mathbf{f}'_4 - 0.0019\mathbf{f}'_5 \quad (10.39)$$

with an R^2 value of 55% and an error variance of 0.00016.

Other acceptance/rejection criteria can be used in a stepwise regression to winnow the attribute set. The possibilities include the following:

SSE: running an F -test on the null hypothesis of no significant change in the sum-of-squares error under the addition (or removal) of an attribute

⁹⁰ Clearly one could simplify the stepwise regression–correlation matrix–stepwise regression sequence by starting with the computation of the correlation matrix on the unwinnowed list of attributes, eliminating those having strong linear dependence, and then running a single stepwise regression to reduce the data set. Anticipating a difference in the size of the correlation matrices in the unwinnowed and winnowed case, we suggest the three-step sequence.

AIC: requiring a reduction in the Akaike information criterion

BIC: requiring a reduction in the Bayesian information criterion

Each criterion will lead to a different level of reduction of the starting data set. Table 10.1 compares the size of the winnowed data set using the adjusted R^2 (ADJR) criterion, the SSE criterion (at 95% significance level), the AIC, and the BIC. The BIC provides the most aggressive reduction of our data set, followed by SSE.

Table 10.1 Size of reduced attribute set obtained from stepwise regression under the four acceptance criteria.

ADJR	SSE (95%)	AIC	BIC
26	14	23	5

Regardless of the identity of the latent factors, additional valuable information may be gleaned from the loading dependencies on the factors. For example, any set of attributes that is positively loaded on a factor provides for opportunities to hedge against any set of attributes that is negatively loaded on the same factor.

References

- Andreou, E. & Ghysels, E. (2009). Structural breaks in financial time series. In T. G. Anderson, R. A. Davis, J.-P Kreiß & T. Mikosch (Eds.), *Handbook of financial time series*. Springer-Verlag, Heidelberg.
- Bruder, B. & Roncalli, T. (2012). Managing risk exposures using the risk budgeting approach. *SSRN Electronic Journal*. <http://dx.doi.org/10.2139/ssrn.2009778>
- Boudt, K., Carl, P. & Peterson, B. G. (2013). Asset allocation with conditional value-at-risk budgets. *Journal of Risk*, 15, 39–68.
- Boudt, K., Peterson, B. G. & Christophe, C. (2008). Estimation and decomposition of downside risk for portfolios with non-normal returns. *Journal of Risk*, 11, 79–103.
- Cattell, R. (1966). The scree test for the number of factors. *Multivariate Behavioral Research*, 1, 245–76.
- Chow, G. (1960). Tests of equality between sets of coefficients in two linear regressions. *Econometrica*, 28, 591–605.
- Chow, G. & Kritzman, M. (2001). Risk budgets. *Journal of Portfolio Management*, 27(2), 56–60.
- Cont, R., Deguest, R. & Giacomo S. (2007). Robustness and sensitivity analysis of risk measurement procedures. Financial Engineering Report No. 2007-06. Columbia University Center for Financial Engineering.
- Dowd, K. (2005). *Measuring market risk* (2nd ed.). John Wiley & Sons, West Sussex.
- Fama, E. F. & French, K. R. (1993). Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics*, 33, 3–56.
- Fama, E. F. & French, K. R. (2015). A five-factor asset pricing model. *Journal of Financial Economics*, 116, 1–22.

- Favre, L. & Galeano, J.-A. (2002). Mean-modified value-at-risk optimization with hedge funds. *Journal of Alternative Investment*, 5, 2–21.
- Gourieroux, C., Laurent, J.-P. & Scaillet, O. (2000). Sensitivity analysis of value-at-risk. *Journal of Empirical Finance*, 7, 225–245.
- Grinhold, R. C. & Kahn, R. N. (2000). *Active portfolio management* (2nd ed.). McGraw Hill, New York.
- Haugh, M., Iyengar, G. & Song, I. (2017). A generalized risk budgeting approach to portfolio construction. *Journal of Computational Finance*, 21(2), 29–60.
- Horn, J. L. (1965). A rationale and test for the number of factors in factor analysis. *Psychometrika*, 30, 179–185.
- Kaiser, H. F. (1960). The application of electronic computers to factor analysis. *Educational and Psychological Measurement*, 20, 141–151.
- Linton, O. (2019). *Financial econometrics: Models and methods*. Cambridge University Press, Cambridge.
- Litterman, R. (1996). Hot SpotsTM and hedges. *Journal of Portfolio Management*, 23(5), 52–75.
- Mahalanobis, P. C. (1936). On the generalized distance in statistics. *Proceedings of the National Institute of Science of India*, 2, 49–55.
- Maillard, S., Roncalli, T. & Tëiletche, J. (2010). The properties of equally weighted risk contribution portfolios. *Journal of Portfolio Management*, 36(4), 60–70.
- McNeil, A. J., Frey, R. & Embrechts, P. (2005). *Quantitative risk management: Concepts, techniques, and tools*. Princeton University Press, Princeton, NJ.
- Mina, J. & Xiao, J. Y. (2001). Return to RiskMetrics: The evolution of a standard. RiskMetrics Group, Inc., New York.
- MSCI Barra Inc. (2007). *Barra risk model handbook*. RV 10-2007.
- Peterson, B. G. & Boudt, K. (2008). Component VaR for a non-normal world. *Risk*, 21(11), 78–81.
- Sharpe, W. (1970). *Portfolio theory and capital markets*. McGraw-Hill, New York.
- Shirvani, A., Stoyanov, S., Rachev, S. T. & Fabozzi, F. J. (2020). A new set of financial instruments. *Frontiers in Applied Mathematics and Statistics*, 26. <https://doi.org/10.3389/fams.2020.606812>
- Velicer, W. F. (1976). Determining the number of components from the matrix of partial correlations. *Psychometrika*, 41, 321–327.
- Zangari, P. (1996). A VaR methodology for portfolios that include options. *RiskMetrics Monitor*, JP Morgan-Reuters, First Quarter, 4–12.

Chapter 11

Optimization with Performance-Attribution Constraints

This chapter demonstrates performance-attribution measures as a basis for constraining portfolio optimization. We employ optimizations that minimize CVaR and investigate two performance attributes, asset allocation (AA) and the selection effect (SE), as constraints on asset weights. The test portfolio consists of stocks from the Dow Jones Industrial Average index. Values for the performance attributes are established relative to a benchmark consisting of equal-weighted portfolios of the same stocks. The performance of the optimized portfolios is evaluated using comparisons of cumulative price and the risk measures MDD (4.8), SR (4.9), SS (4.10), and RR (4.11). The results suggest that achieving SE performance thresholds requires larger turnover values than those required for achieving comparable AA thresholds. The results also suggest that the imposition of constraints on AA and SE plays a positive role in price and risk-measure performance.

How well a portfolio performs is always the primary concern for investors, and it is usually the metric that best reflects investor confidence in the portfolio's management. In common terms, a good portfolio delivers satisfactory return with low risk. Attribution analysis provides measures for how well a portfolio is being managed. Paraphrasing Bacon (2008), performance attribution is a technique used to quantify the excess return (relative to a benchmark) of a portfolio and explain that performance in terms of investment strategy and market conditions. From a management perspective, attribution analysis has been used to monitor performance, identify early indications of underperformance, and gain investor confidence by demonstrating a thorough understanding of the performance drivers. To the best of our knowledge, performance-attribution measures are currently used exclusively as a diagnostic tool, in the sense that if today's attribute values underperform, changes are implemented in the portfolio with the goal of improving tomorrow's attribute values. In this chapter, we investigate the approach of imposing performance-attribution constraints to guarantee that tomorrow's portfolio will achieve the required attribute values.

Following the foundational studies of performance attribution by Brinson and Fachler (1985) and Brinson et al. (1986), we decompose excess return into two quantities that reflect investment strategy: AA, which measures the contribution of each asset class in a portfolio to the portfolio's total performance, and SE, which measures the impact of asset choice within each class in the portfolio. As their definitions in the next section indicate, AA and SE measure the differences between the mean performance of asset classes in a managed portfolio and that of asset classes in a market benchmark and are therefore "blind" to volatility effects, that is, to tail risk. Motivated by the further work of Biglova and Rachev (2007) and Rachev et al. (2009), we investigate the impact on portfolio optimization of using AA and SE as additional constraints on asset weights as a method of combining performance and tail-risk control. We apply this methodology to a test portfolio of stocks comprising a major market index, namely the Dow Jones Industrial Average. Optimization is performed by minimizing CVaR at a specified quantile level, α . For the required market benchmark, we consider two options, an equal-weighted portfolio and a price-weighted portfolio, comprised of the same assets. The performance of the resulting optimal portfolios is measured in terms of cumulative portfolio price and standard risk measures.

11.1 Performance-Attribute Constraints

Consider a managed portfolio p comprised of N assets consisting of M asset classes with n_i assets in class i ; $i = 1, \dots, M$ such that $\sum_{i=1}^M n_i = N$. Let b denote a benchmark portfolio composed of Q assets comprising the same M asset classes, with q_i assets in class i ; $i = 1, \dots, M$ such that $\sum_{i=1}^M q_i = Q$. Let the index pair, ij ; $i = 1, \dots, M$; $j = 1, \dots, n_i$, identify portfolio asset j in class i , with the analogous identification for benchmark assets. Denote the daily closing price of an asset as $S_{ij}(t)$ and its corresponding log-return as $r_{ij}(t) = \ln(S_{ij}(t)/S_{ij}(t-1))$. For brevity, we suppress the time variable for most of the discussion in this section. Let $w_{ij}^{(p)}$ denote the weight of asset ij in portfolio p and $w_{ij}^{(b)}$ denote asset weight in the benchmark. We assume all weights are nonnegative; that is, all the portfolios considered take long-only positions. Let $w_i^{(p)} = \sum_{j=1}^{n_i} w_{ij}^{(p)}$ and $w_i^{(b)} = \sum_{j=1}^{q_i} w_{ij}^{(b)}$ represent the total weights of the assets in class i in the portfolio and benchmark, respectively. Note that for any portfolio fully invested in its component assets (which we assume is the case), $\sum_{i=1}^M w_i^{(p)} = \sum_{i=1}^M w_i^{(b)} = 1$.

The quantities AA and SE for asset class i are defined as follows (Biglova and Rachev, 2007):⁹¹

$$AA_i = (w_i^{(p)} - w_i^{(b)})(R_i^{(b)} - R^{(b)}), \quad (11.1)$$

$$SE_i = w_i^{(b)}(R_i^{(p)} - R_i^{(b)}), \quad (11.2)$$

where

$$R_i^{(p)} = \sum_{j=1}^{n_i} \frac{w_{ij}^{(p)}}{w_i^{(p)}} \mathbb{E}[r_{ij}], \quad R_i^{(b)} = \sum_{j=1}^{q_i} \frac{w_{ij}^{(b)}}{w_i^{(b)}} \mathbb{E}[r_{ij}], \quad R^{(b)} = \sum_{i=1}^M \sum_{j=1}^{q_i} w_{ij}^{(b)} \mathbb{E}[r_{ij}], \quad (11.3)$$

and $\mathbb{E}[\cdot]$ denotes expected value. In (11.3), the ratio $w_{ij}^{(p)}/w_i^{(p)}$ represents the fractional weight held by asset j in class i in portfolio p . (That is, $\sum_{j=1}^{n_i} (w_{ij}^{(p)}/w_i^{(p)}) = 1$.) Thus, $R_i^{(p)}$ (and similarly, $R_i^{(b)}$) represents an expected log-return of asset class i viewed as a fully invested portfolio in itself. In contrast, $R^{(b)}$ represents the usual expected log-return of the entire benchmark portfolio.⁹² From (11.3), we have $R^{(b)} = \sum_{i=1}^M w_i^{(b)} R_i^{(b)}$. Similarly, we have the usual expected log-return of portfolio p ,

⁹¹ In the original formulation developed by Brinson et al. (1986) (see also Bacon, 2008, Chapter 5), AA_i is defined as $AA_i = (w_i^{(p)} - w_i^{(b)})R_i^{(b)}$. The definition in Biglova and Rachev (2007), which we follow here, uses the excess return $R_i^{(b)} - R^{(b)}$ for benchmark class i relative to the entire benchmark return in the definition (8.1) of AA_i . We note that although this modifies the values for AA_i relative to that of the original Brinson et al. formulation, the total value, $AA = \sum_{i=1}^M (w_i^{(p)} - w_i^{(b)})(R_i^{(b)} - R^{(b)}) = \sum_{i=1}^M (w_i^{(p)} - w_i^{(b)})R_i^{(b)} - \sum_{i=1}^M (w_i^{(p)} - w_i^{(b)})R^{(b)} = \sum_{i=1}^M (w_i^{(p)} - w_i^{(b)})R_i^{(b)} - 0$, is in agreement with the total value of AA in the Brinson et al. approach.

⁹² If $R_i^{(p)}$, $R_i^{(b)}$, $R^{(b)}$, and r_{ij} were simple (i.e., discrete) returns, the formulas in (11.3) would be exact. However, because they are log-returns, such formulas are approximate. For example, to leading order in a Taylor series expansion $R^{(b)} - \sum_{i=1}^M \sum_{j=1}^{q_i} w_{ij}^{(b)} \mathbb{E}[r_{ij}] \approx \frac{1}{2} \left[\sum_{i=1}^M \sum_{j=1}^{q_i} w_{ij}^{(b)} \mathbb{E}[r_{ij}]^2 - \left(\sum_{i=1}^M \sum_{j=1}^{q_i} w_{ij}^{(b)} \mathbb{E}[r_{ij}] \right)^2 \right]$.

11. Optimization with Performance Attribution Constraints

$$R^{(p)} = \sum_{i=1}^M \sum_{j=1}^{n_i} w_{ij}^{(p)} \mathbb{E}[r_{ij}] = \sum_{i=1}^M w_i^{(p)} R_i^{(p)}. \quad (11.4)$$

The excess return, $S = R^{(p)} - R^{(b)}$, can be viewed as the value added by portfolio management. From (11.1) through (11.4),

$$S = \sum_{i=1}^M (AA_i + SE_i + I_i) = AA + SE + I, \quad (11.5)$$

where $I_i = (w_i^{(p)} - w_i^{(b)})(R_i^{(p)} - R_i^{(b)})$ is an “interaction” term. AA, SE, and I are, respectively, the total AA, total SE, and total interaction terms for portfolio p . AA_i represents the contribution to the total value added to the excess return, S , from asset class i , whereas SE_i represents the contribution to S determined by the choice of assets within class i . To understand these interpretations, consider first the sign of the value of AA_i in (11.1).

- If $R_i^{(b)} - R^{(b)} > 0$, the expected return of asset class i in the benchmark is outperforming the total expected return of the benchmark. Therefore, if $w_i^{(p)} - w_i^{(b)} > 0$, the weight of asset class i in portfolio p is larger than in the benchmark, capitalizing further on the better return of class i . Otherwise, if $w_i^{(p)} - w_i^{(b)} < 0$, the class i weighting in portfolio p is detrimental to the potential performance of that class (as determined by the benchmark).
- If $R_i^{(b)} - R^{(b)} < 0$, the expected return of asset class i in the benchmark is underperforming the total expected return of the benchmark. Therefore, if $w_i^{(p)} - w_i^{(b)} < 0$, the weight of asset class i in portfolio p is smaller than in the benchmark, further suppressing the poorer return of that class. Otherwise, if $w_i^{(p)} - w_i^{(b)} > 0$, the class i weighting in portfolio p is overweighting the poor performance of that class.

Thus, a positive sign for the value of AA_i indicates a “correct” decision in the management of portfolio p relative to the benchmark, whereas a negative sign indicates a “poor” decision. The magnitude of AA_i quantifies how correct or poor the decision is.

Similarly, because we assume⁹³ $w_i^{(b)} > 0$, $i = 1, \dots, M$, a positive sign for the value of SE_i in (11.2) indicates that the expected return of the choice of assets in class i in portfolio p is outperforming that class in the benchmark, whereas a negative sign indicates that the expected return of the choice of assets in class i in portfolio p is underperforming.

The interaction term, I_i , captures the part of the excess return unexplained by AA and SE. Written as

$$I_i = \frac{AA_i SE_i}{w_i^{(b)} (R_i^{(b)} - R^{(b)})}, \quad (11.6)$$

it can be viewed as the product of the AA and SE contributions of class i to portfolio p compared to the weighted excess return of class i in the benchmark b . Alternatively, written as

⁹³ Also a requirement for class i to be in the portfolio.

$$I_i = \left(\frac{w_i^{(p)}}{w_i^{(b)}} - 1 \right) SE_i, \quad (11.7)$$

it can be interpreted as the product of SE and the over- or underweighted part of asset class i . The relationship (11.7) between I_i and SE_i reveals a simple form for the sum of SE and the interaction terms:

$$\overline{SE}_i := SE_i + I_i = w_i^{(p)} \left(R_i^{(p)} - R_i^{(b)} \right). \quad (11.8)$$

Equation (11.8) provides a way to incorporate a constraint on the sum of SE and the interaction effect for class i ; however, we do not consider such a combined constraint here.

Portfolio optimizations that maximize return while minimizing risk (subject to additional constraints) require the specification of a proxy measure for risk. Common examples of risk proxies are the variance of the portfolio (Markowitz, 1952), VaR (JP Morgan, 1996), expected tail loss, CVaR (Rockafellar and Uryasev, 2000),⁹⁴ and mean absolute deviation (Konno and Yamazaki, 1991). Measures that focus on tail risk, such as VaR and CVaR, have become very popular as a result of the need to understand exposure to loss under “extreme” market events. (Gava et al. (2021) has recently demonstrated that consideration of tail risk can successfully reduce sharp losses in multiasset portfolios.) However, VaR has undesirable mathematical characteristics; except when the underlying random process is Gaussian, VaR is not a coherent risk measure, because it lacks the properties of subadditivity and convexity (Artzner et al., 1999). As a risk measure, CVaR is coherent (Pflug, 2000); its use as a standard has grown to the point that the Basel III regulatory framework for banks requires it. We therefore use CVaR as the risk measure for our portfolio optimizations.

Rockafellar and Uryasev’s (2000) approach to portfolio optimization, which is based on minimizing CVaR at quantile level $1 - \alpha$, is discussed in section 3.2.3 and leads to the minimization problem (3.31). When evaluated for a portfolio consisting of a finite sample of asset returns, $\mathbf{r}(t), t = 1, \dots, T$, with $f(\mathbf{w}, \mathbf{r}) = \mathbf{w}'\mathbf{r}$,⁹⁵ the discrete form of (3.31) results in the minimization problem

$$\min_{\mathbf{w}} \text{CVaR}_{\alpha}(\mathbf{w}) = \min_{\mathbf{w}, \gamma} \left\{ \gamma + \frac{1}{\alpha T} \sum_{t=1}^T (-\mathbf{w}'\mathbf{r}(t) - \gamma)^+ \right\}. \quad (11.9)$$

As noted in (3.32), Rockafellar and Uryasev’s approach proceeds by converting (11.9) into a linear objective function by introducing the variable $y_t \geq -\mathbf{r}_t^T \mathbf{w} - \gamma \geq 0$. This conversion is particularly appropriate if all the constraints are also linear, in which case the constrained minimization problem can be solved by linear programming. Because we deal with nonlinear constraints, we leave the objective function in the form (11.9) and solve using nonlinear optimization.

⁹⁴ If the underlying profit–loss distribution is continuous, then the definitions of ETL (also known as tail conditional expectation (TCE) or tail value-at-risk (TVaR)) and CVaR (also known as expected shortfall (ES) or average value-at-risk (AVaR)) coincide. In the general case, however, CVaR is a coherent risk measure whereas ETL is not (Stoyanov, 2005).

⁹⁵ In contrast to Chapter 3, here we use the notation \mathbf{w}' to refer to vector transpose to avoid confusion with the time value T .

11. Optimization with Performance Attribution Constraints

We describe our approach for solving (11.9) with a general constraint here and discuss the specific constraints below. Consider optimization of (11.9) under the constraint $c_1(\mathbf{w}) \leq 0$. If, on day t , the feasible set for the constrained optimization is null, the constraint is removed and replaced for that day by adding a quadratic penalty term to (11.9),

$$\min_{\mathbf{w}} \text{CVaR}_{\alpha}(\mathbf{w}) = \min_{\mathbf{w}, \gamma} \left\{ \gamma + \frac{1}{\alpha T} \sum_{t=1}^T (-\mathbf{w}'\mathbf{r}(t) - \gamma)^+ + \beta ((c_1(\mathbf{w}))^+)^2 \right\}. \quad (11.10)$$

The coefficient β can be set by the user. If k constraints need to be removed, they are replaced in (11.10) by the sum $\beta \sum_{i=1}^k ((c_i(\mathbf{w}))^+)^2$.

We consider four portfolio-optimization problems, P_{α}^k , $k = 0, \dots, 3$, based on constrained minimization of (11.9) or, in the case of a null feasible set, (11.10):

P_{α}^0 : (a) $w_{ij}^{(p)} \geq 0$, $\sum_{ij} w_{ij}^{(p)} = 1$; and (b) $\text{TO} \leq C_{\text{TO}}$;

P_{α}^1 : (a) $w_{ij}^{(p)} \geq 0$, $\sum_{ij} w_{ij}^{(p)} = 1$; (b) $\text{TO} \leq C_{\text{TO}}$; and (c) $a_1 \leq \text{AA} \leq b_1$.

P_{α}^2 : (a) $w_{ij}^{(p)} \geq 0$, $\sum_{ij} w_{ij}^{(p)} = 1$; (b) $\text{TO} \leq C_{\text{TO}}$; and (d) $a_2 \leq \text{SE} \leq b_2$.

P_{α}^3 : (a) $w_{ij}^{(p)} \geq 0$, $\sum_{ij} w_{ij}^{(p)} = 1$; (b) $\text{TO} \leq C_{\text{TO}}$; (c) $a_1 \leq \text{AA} \leq b_1$; and (d) $a_2 \leq \text{SE} \leq b_2$.

Here, $\text{TO} \leq C_{\text{TO}}$ is a turnover constraint,

$$\text{TO} = \frac{1}{2} \sum_{i=1}^M \sum_{j=1}^{n_i} |w_{ij}^{(p)}(t) - w_{ij}^{(p)}(t-1)| \leq C_{\text{TO}}, \quad (11.11)$$

used as a proxy to control transaction costs. The “base case” portfolio P_{α}^0 considers no performance-attribution constraints and is therefore independent of the benchmark portfolio. The optimization problems P_{α}^1 through P_{α}^3 successively add further performance-attribution constraints to the long-only, fully invested, CVaR_{α} -minimized base portfolio.

The user can specify the constants a_i, b_i to meet particular goals. For example, the constraint $\text{AA} \geq 0$ requires that, on average, the asset classes in the optimized portfolio p equal or outperform those in the benchmark. A constraint $\text{SE}_i \geq 0$ requires that the weights of the portfolio assets in class i be adjusted to perform as well as or better than class i in the benchmark. Because individual asset weights can be zero, adjusting weights is equivalent to choosing the set of assets within in the class. The constraint $\text{SE} \geq 0$ requires that improved performance be achieved as an average over classes. Because $R_i^{(p)}(t)$ involves the ratio $w_{ij}^{(p)}(t)/w_i^{(p)}(t)$, constraints involving SE_i terms are nonlinear. In contrast, constraints involving terms AA_i and $\overline{\text{SE}}_i$ are linear.⁹⁶ Our implementation is performed in MatLab using the constrained, nonlinear multivariate function *fmincon()* and the solver *sqp()*. The performance of these four optimized portfolios relative to one another is evaluated on the basis of cumulative price and the four risk measures discussed in section 4.3.

11.2 Application to Domestic REIT Portfolio

Using both historical and dynamic long-only optimization, we apply CVaR minimization under performance-attribution constraints to the domestic portfolio. The REIT asset classes assigned are presented in Table 11.1. As equations (11.1)–(11.6) indicate, the results of an attribution analysis

⁹⁶ This assumes that benchmark weight values can be obtained in a timely manner and are not part of the optimization.

depend on the choice of benchmark. For ease of assignment to asset classes, we use the equal-weighted domestic portfolio EQW as the benchmark for determining values for AA_i and SE_i .⁹⁷ The 10-year U.S. Treasury yield-curve rate is used for the risk-free rate.

Table 11.1 Asset classes in the domestic portfolio.

Class	REIT								
Office/Industrial	VNO	ARE	PLD	BXP	DRE	SLG	PSA		
Residential	AVB	EQR	MAA	UDR	ESS				
Retail	SPG	MAC	FRT	O	REG				
Specialty	HCP	VTR	NLY	AMT	CCI	IRM	SBAC	WY	HST

We optimize at two separate quantile levels, $\alpha \in \{0.95, 0.99\}$. Daily turnover constraints are set to one of three values: $C_{TO} = \infty$ (no turnover constraint), 4%, and 0.4%. For the attribution constraints in optimizations P_α^1 , P_α^2 , and P_α^3 , we set the lower bounds $a_1 = a_2 = 0$ and set no upper bounds ($b_1 = b_2 = \infty$). Thus, for example, optimization P_α^1 minimizes $CVaR_\alpha$ for the long-only portfolio while requiring that, on average, its asset classes outperform the benchmark.

If the constrained optimization problem results in a null feasible set for day t , constraints are replaced by penalty terms in the following order:

P_α^0 : The turnover constraint is replaced by a penalty term.

P_α^1 : The turnover constraint is replaced by a penalty term. If the feasible set is still null, the AA constraint is then additionally replaced by a penalty term.

P_α^2 : The turnover constraint is replaced by a penalty term; if necessary, the SE constraint is also replaced.

P_α^3 : The order of additional conversion into penalty terms is as follows: turnover constraint, SE, and finally AA.

If the feasible space is still null after all the indicated hard constraints are converted into penalty terms for day t , the optimized weights obtained for day $t - 1$ are used for day t .

Table 11.2 Frequency of conversion of hard constraints into penalty terms (expressed as a percent of total time steps).

Portfolio	Historical				Dynamic			
	TO	TO +AA	TO +SE	TO +AA +SE	TO	TO +AA	TO +SE	TO +AA +SE
		TO < ∞				TO < ∞		
$P_{0.95}^0$		0.0				0.0		
$P_{0.95}^1$			8.7 [†]				4.2	
$P_{0.95}^2$			7.2	0.0			3.3	0.06
$P_{0.95}^3$		TO \leq 4%				TO \leq 4%		
$P_{0.95}^0$	0.0				36.6			
$P_{0.95}^1$	1.2	0.0			67.6	0.0		
$P_{0.95}^2$	7.4		3.8		43.4		1.0	

⁹⁷ Thus, $Q = N$ and $q_i = n_i$, $i = 1, \dots, M$.


11. Optimization with Performance Attribution Constraints

$P_{0.95}^3$	7.9	3.5	0.03	73.1	0.8	0.0
	TO $\leq 0.4\%$			TO $\leq 0.4\%$		
$P_{0.95}^0$	11.2			97.1		
$P_{0.95}^1$	30.9	0.0		98.6	0.0	
$P_{0.95}^2$	37.2		1.6	95.4	1.0	
$P_{0.95}^3$	53.2		2.1	98.1	0.3	0.0

[†] 2/3274 = 0.06% of the timesteps resulted in a null feasible set for time t

For the optimizations that minimize $\text{CVaR}_{0.95}$, Table 11.2 summarizes the frequency of conversion of a hard constraint into a penalty term. For example, for historical optimization of $P_{0.95}^3$ under the turnover constraint $\text{TO} \leq 4\%$, 88.6% of the timesteps result in a feasible solution to the fully constrained problem; 7.9% of the timesteps require converting the turnover constraint into a penalty term; 3.5% of the timesteps require converting both the turnover and SE constraints into penalty terms; 0.03% of the timesteps (i.e., one timestep) require the conversion of turnover, SE, and AA constraints. A (non-empty) feasible set was always obtained when solving with a (combined) penalty term.⁹⁸ We emphasize the following points:

- When the turnover-constraint limit is decreased (a heavier constraint imposed), the frequency of conversion from a hard turnover constraint into a penalty term increases significantly.
- The frequency of conversion of a hard turnover constraint into a penalty term also increases significantly in moving from historical to dynamic optimization. For the dynamic optimization, over 95% of the timesteps require the conversion of the $\text{TO} \leq 0.4\%$ constraint into a penalty term.
- Considering all of the optimizations performed in Table 11.2, the $\text{AA} \geq 0$ constraint had to be converted into a penalty term only once: for a single timestep in the $P_{0.90}^1$ optimization under the constraint $\text{TO} \leq 0.4\%$.
- The frequency of conversion of the $\text{SE} \geq 0$ constraint into a penalty term decreases as the turnover-constraint limit decreases, and it also decreases in moving from historical to dynamic optimization.
- The results presented in Table 11.2 provide support for the choice of the TO–SE–AA order of conversion of hard constraints into penalty terms.



⁹⁸ In all the optimizations, whether for $\alpha = 0.95$ or $\alpha = 0.99$, only the historical, $P_{0.95}^2$, optimization with no turnover constraint recorded a timestep (and in fact, only two timesteps) in which an empty feasible set was obtained when attempting a solution using penalty terms. For those two timesteps, optimum weight values from the previous timestep were used.




Figure 11.1 Box-whisker summary statistics of the distribution of TO values for the historical and dynamic optimizations with $TO < \infty$ and $TO \leq 0.004$ (0.4%).

Fig. 11.1 displays the box-whisker summaries of the distribution of TO values in the resulting historical and dynamic optimized portfolios for $TO < \infty$ and $TO \leq 0.4\%$. By definition, $0 \leq TO \leq 1$. For many of the optimizations, and especially P_{α}^2 and P_{α}^3 , outlier TO values approach 1. To illustrate the distribution statistics more clearly, the y-axis range is limited in Figs. 11.1(a), (b), and (d), whereas the full range is displayed in (c).

- For optimizations with no turnover constraint, the imposition of performance-attribute constraints increases turnover values, with the constraint $SE \geq 0$ producing the most significant turnover-value increases under historical optimization. For dynamic optimization, the imposition of $AA \geq 0$ produces the most significant increase in turnover value.
- As noted in Table 11.2, when a turnover constraint is imposed, the TO constraint is replaced by a penalty term very frequently, resulting in the Q_{25} value frequently exceeding the value of C_{TO} .
- For a fixed value of i , the IQR and Q_{50} for P_{α}^i generally increase in moving from $\alpha = 0.95$ to $\alpha = 0.99$.
- For fixed values of i and α , the IQR and Q_{50} for P_{α}^i increase in moving from historical to dynamic optimization.
- Q_{50} increases from P_{α}^0 to P_{α}^3 in the historical optimizations; in the dynamic optimizations, the increase, from P_{α}^0 to P_{α}^1 and then from P_{α}^2 to P_{α}^3 , suggests that the increase has a stronger correlation with the imposition of the AA constraint.

Fig. 11.2 displays the box-whisker summaries of the distribution of AA values observed in the resulting historical and dynamic portfolios for $TO < \infty$ and $TO \leq 0.4\%$. In all the runs, the imposition of the $AA \geq 0$ constraint in P_{α}^1 and P_{α}^3 results in a positive distribution of AA values.⁹⁹ For the historical optimizations, the AA distribution shifts from $\sim 22.5\%$ negative values for P_{α}^0 and P_{α}^2 to no negative values for P_{α}^1 and P_{α}^3 . For the dynamic optimizations, the shift is from 35% ($TO < \infty$) and 40% ($TO \leq 0.4\%$) for P_{α}^0 and P_{α}^2 to no negative values for P_{α}^1 and P_{α}^3 .

⁹⁹ To within a relative constraint tolerance of $|10^{-6}|$ imposed by the *sqp* solver.

11. Optimization with Performance Attribution Constraints

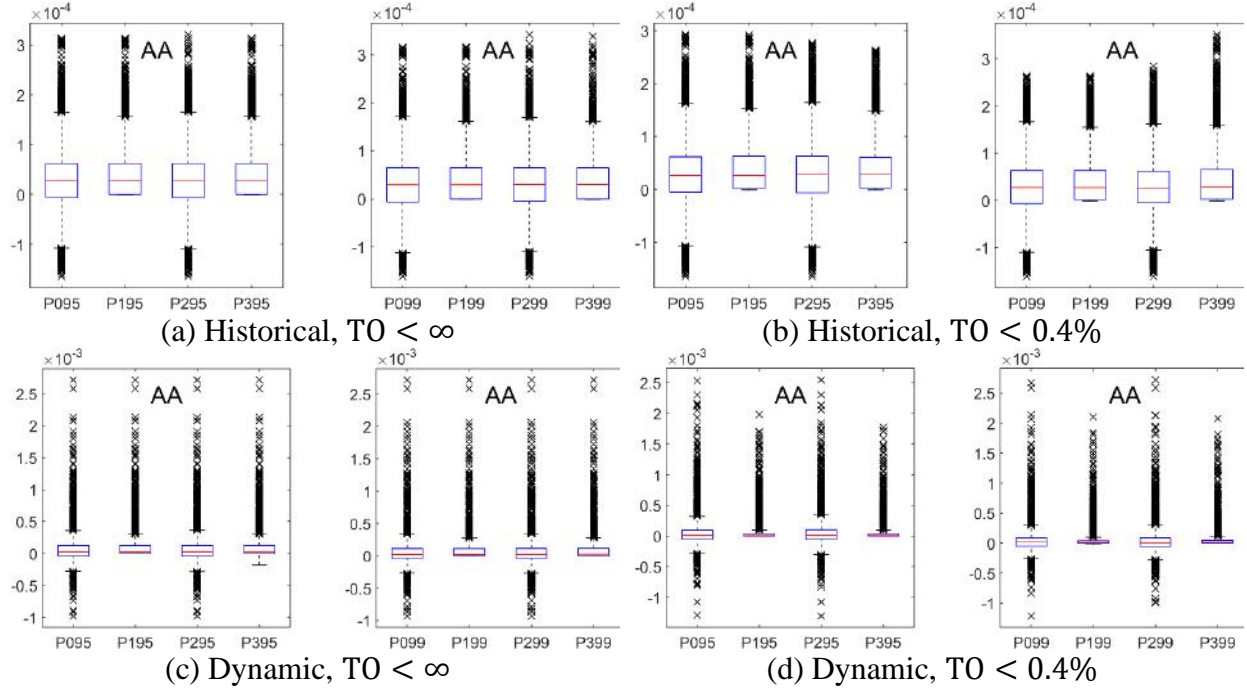


Figure 11.2 Box-whisker summary statistics of the distribution of AA values for the historical and dynamic optimizations with $TO < \infty$ and $TO \leq 0.4\%$.

- For a fixed value of i , the distribution statistics of AA for P_{α}^i do not differ significantly in moving from $\alpha = 0.95$ to $\alpha = 0.99$ under historical or dynamic optimization. The most notable difference is in the (positive) outlier behavior for $TO \leq 0.4\%$ under historical optimization (Fig. 11.2(b)).
- For fixed values of i and α , the distribution statistics of AA for P_{α}^i exhibit relatively minor differences (in the behavior of positive outliers) in moving from $TO < \infty$ to $TO \leq 0.4\%$ under historical optimization (Figs. 11.2(a) and (b)).
- For fixed values of i and α , the distribution statistics of AA for P_{α}^i exhibit a decrease in the IQR for P_{α}^1 and P_{α}^3 in moving from $TO < \infty$ to $TO \leq 0.4\%$ under dynamic optimization (Figs. 11.2(c) and (d)).

Based on these observations, we note that requiring $AA \geq 0$ is a “strong” requirement that can be achieved through the straightforward imposition of a hard constraint.

Fig. 11.3 displays the box-whisker summaries of the distribution of SE values observed in the historical portfolios for $TO < \infty$. Imposition of the $SE \geq 0$ constraint in P_{α}^2 and P_{α}^3 causes a significant increase in the positive value of SE for some timesteps. This behavior is seen in both dynamic and historical optimization, with the size of the increase becoming less prominent as the value of C_{TO} decreases. Under a restricted range of y-axis values, Fig. 11.4 displays the box-whisker summaries of the distribution of SE values observed in the resulting historical and dynamic portfolios for $TO < \infty$ and $TO \leq 0.4\%$. In all the runs, the imposition of the $SE \geq 0$ constraint in P_{α}^2 and P_{α}^3 results in a much more positive distribution of SE values.⁹⁹ For the historical optimizations, the SE distribution shifts from 82%–86% negative values for P_{α}^0 and P_{α}^1 to 1%–5% negative values for P_{α}^2 and P_{α}^3 (with the smaller percentage values occurring for $TO \leq 0.4\%$). For the dynamic optimizations, the shift is from 34%–43% for P_{α}^0 and P_{α}^1 to less than 1.3% for P_{α}^2 and P_{α}^3 (again with the smaller percentage values occurring for $TO \leq 0.4\%$). Unlike for

the AA constraint, the negative SE values occurring for P_{α}^2 and P_{α}^3 correspond to the percent of time the SE constraint has to be converted into a penalty term in order to achieve a non-null feasible set. Attempting to impose $SE \geq 0$ as a hard constraint results in more frequent conversion into a penalty term than does the imposition of $AA \geq 0$ as a hard constraint.




Figure 11.3 Box-whisker summary statistics of the distribution of SE values for the historical optimizations with $TO < \infty$.

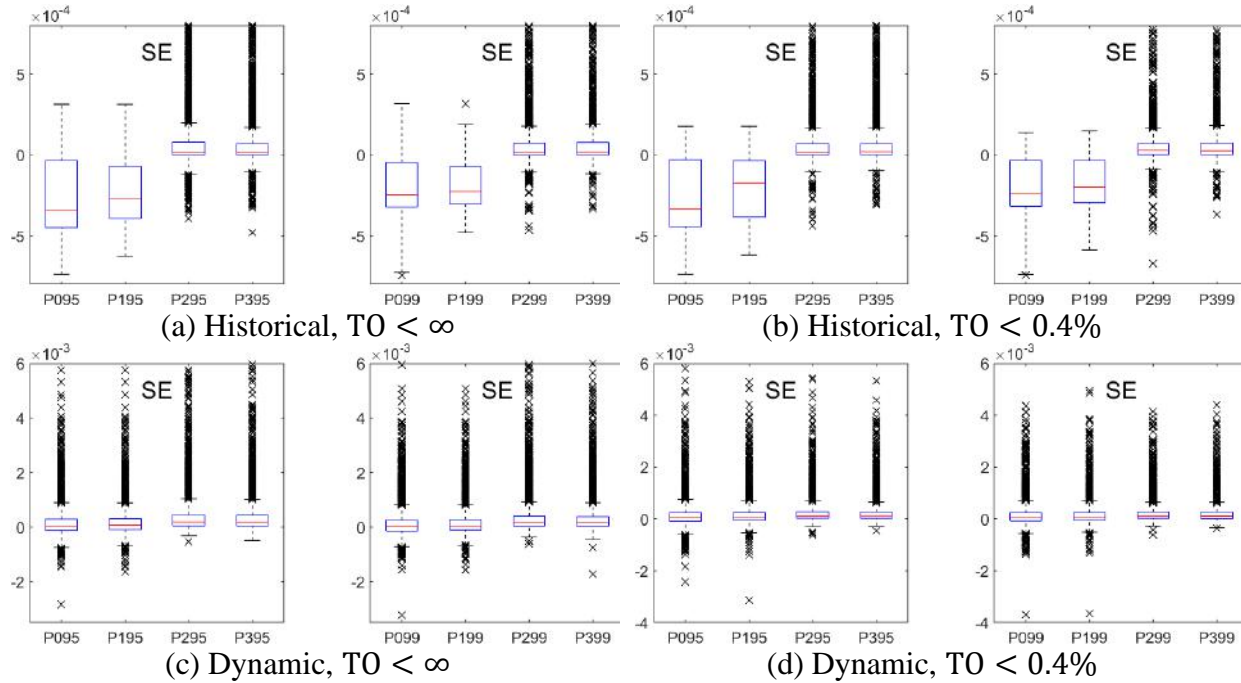


Figure 11.4 Box-whisker summary statistics (under a restricted range of y-axis values) of the distribution of SE values for the historical and dynamic optimizations with $TO < \infty$ and $TO \leq 0.4\%$.

For both historical and dynamic optimization, Fig. 11.5 shows the price performance of the base portfolios, $P_{0.95}^0$ and $P_{0.99}^0$, under changing turnover constraint. With no turnover constraint, the dynamic optimization outperforms the historical optimization. Dynamic optimization is more sensitive to increases in the turnover-constraint level; with a tighter daily turnover constraint, the price performance becomes more constrained. In contrast, under historical optimization, the price performance improves slightly under increased turnover constraint.

11. Optimization with Performance Attribution Constraints

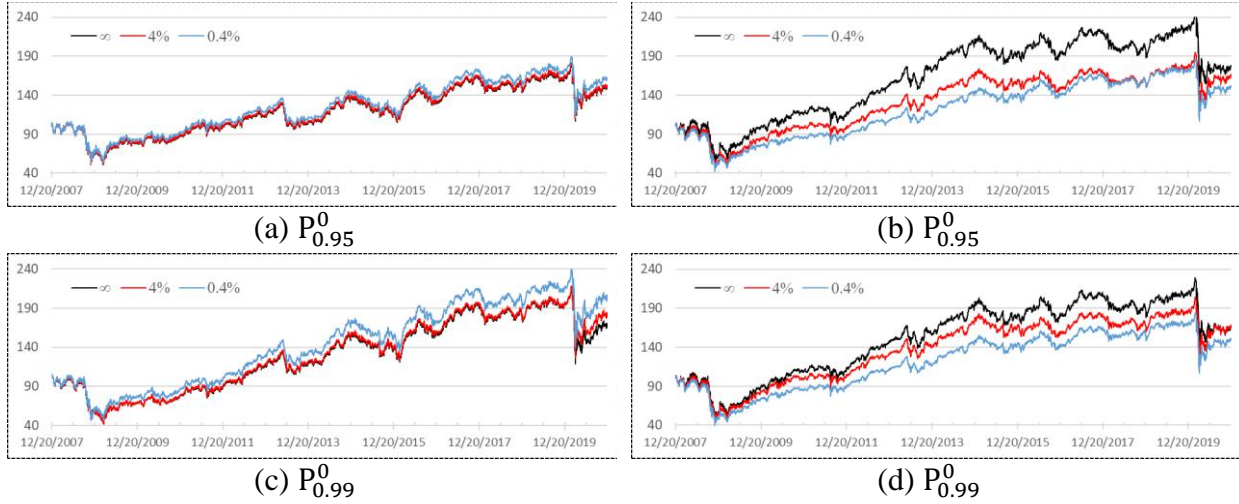


Figure 11.5 Price performance of the base portfolios $P_{0.95}^0$ and $P_{0.99}^0$ as a function of changing turnover constraint for historical (left) and dynamic (right) optimization.

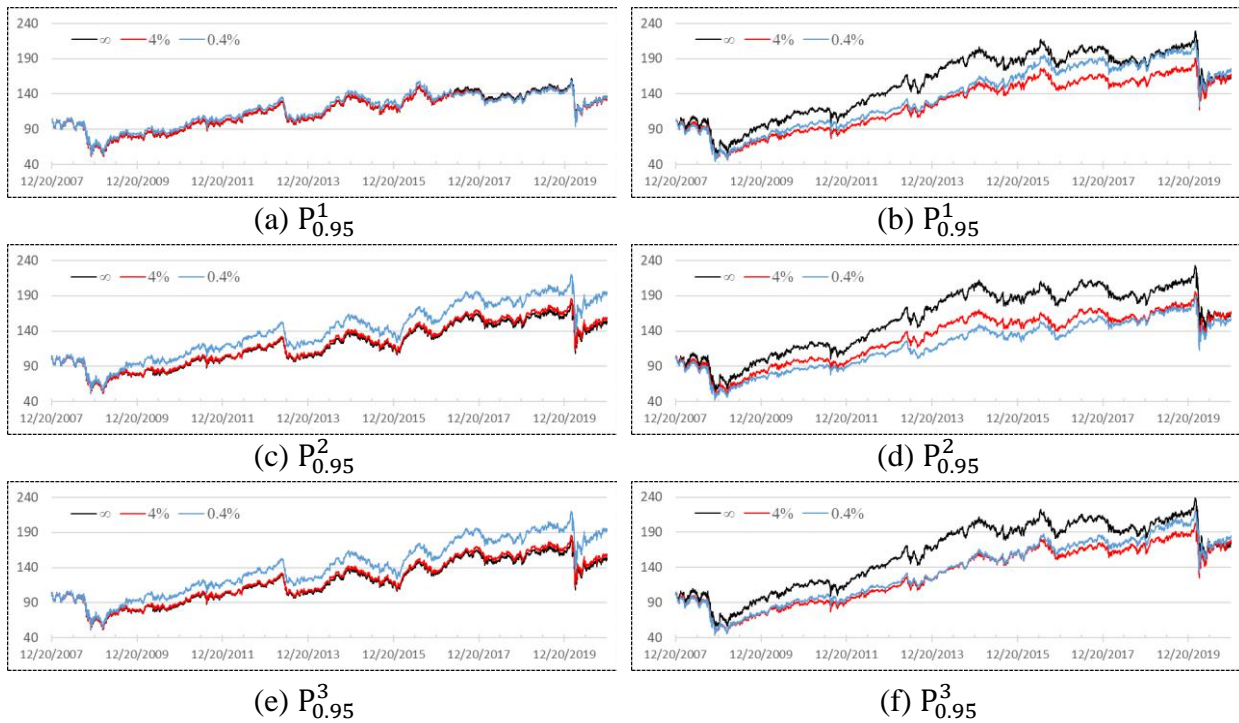


Figure 11.6 Price performance of the performance-attribute-constrained portfolios $P_{0.95}^1$, $P_{0.95}^2$, and $P_{0.99}^0$ as a function of changing turnover constraint for historical (left) and dynamic (right) optimization.

Fig. 11.6 shows the cumulative price performance of each of the performance-attribute-constrained $\text{CVaR}_{0.95}$ -minimized portfolios under changing turnover-constraint level computed using both the historical and dynamic optimizations. As with the base portfolio, in the absence of a turnover constraint, the price performance of the dynamic optimizations is superior to that of the historical optimizations. Again, the historical optimizations show price improvement under an

increase in the turnover constraint. For the dynamic optimizations, price performance worsens under the 4% daily turnover constraint, but the AA-constrained optimizations $P_{0.95}^1$ and $P_{0.95}^3$ show some improvement when the daily constraint is further reduced to 0.4%.




Figure 11.7 Price performance of the historical (left) and dynamic (right) $CVar_{0.95}$ - and $CVar_{0.99}$ -minimized portfolios under a 0.4% turnover constraint.

Fig. 11.7 summarizes the cumulative price performance of the historical and dynamic $CVar_{0.95}$ - and $CVar_{0.99}$ -minimized portfolios under a daily turnover constraint of 0.4%. Compared with the base portfolio, for the historical optimization, the AA-constrained portfolio, P_{α}^1 , produces a degraded price performance, whereas the SE-constrained P_{α}^2 improves performance. (The price performance of the doubly constrained P_{α}^3 is intermediate between that of P_{α}^1 and P_{α}^2 .) For the dynamic $CVar_{0.95}$ minimization, the AA-constrained $P_{0.95}^1$ portfolio produces improved price performance, whereas the SE-constrained $P_{0.95}^2$ produces no significant price improvement compared to the base portfolio $P_{0.95}^0$. In contrast, for the dynamic $CVar_{0.99}$ minimization, all three performance-attribute-constrained optimizations show improved price performance compared to $P_{0.99}^0$.

Fig. 11.8 summarizes the risk measures for the historical and dynamic $CVar_{0.95}$ -minimized portfolios under a daily turnover constraint of 0.4%. The measures are computed using a one-year moving window. Fig. 11.9 presents box-whisker summaries of the statistics of the plots in Fig. 11.8. The worst MDD occurs most frequently for the EQW portfolio. The difference is most discernible for dynamic optimization, in which $P_{0.95}^1$ and $P_{0.95}^3$ have low-value, narrow-spread IQRs. In terms of SR, SS, and RR, examination of the median values in Fig. 11.9 shows that $P_{0.95}^0$ and $P_{0.95}^2$ tend to rank comparably and $P_{0.95}^1$ and $P_{0.95}^3$ also tend to rank comparably. Under historical optimization, the $P_{0.95}^0$ – $P_{0.95}^2$ pair tends to perform the best, whereas under dynamic optimization, the $P_{0.95}^1$ – $P_{0.95}^3$ pair tends to perform the best. (The exception is the performance of the median value of RR for $P_{0.95}^3$, which has the lowest value.) For dynamic optimization, EQW has the lowest median SR, SS, and RR values, whereas for historical optimization, it has the lowest median RR value but has intermediate rank in median SR and SS values. Interestingly, all four portfolios $P_{0.95}^i, i = 1, \dots$, have a large number of high-value RR outliers under historical optimization.

11. Optimization with Performance Attribution Constraints

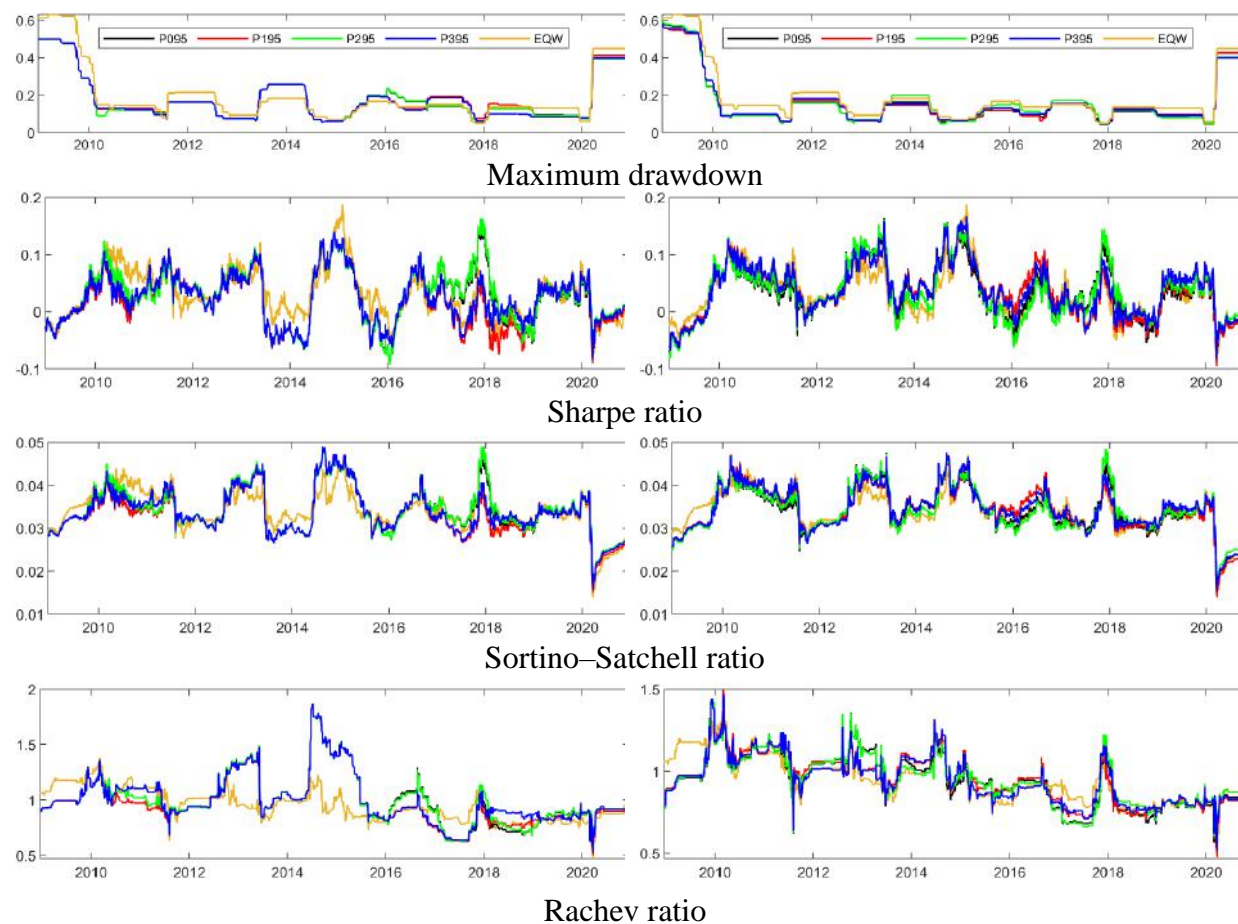
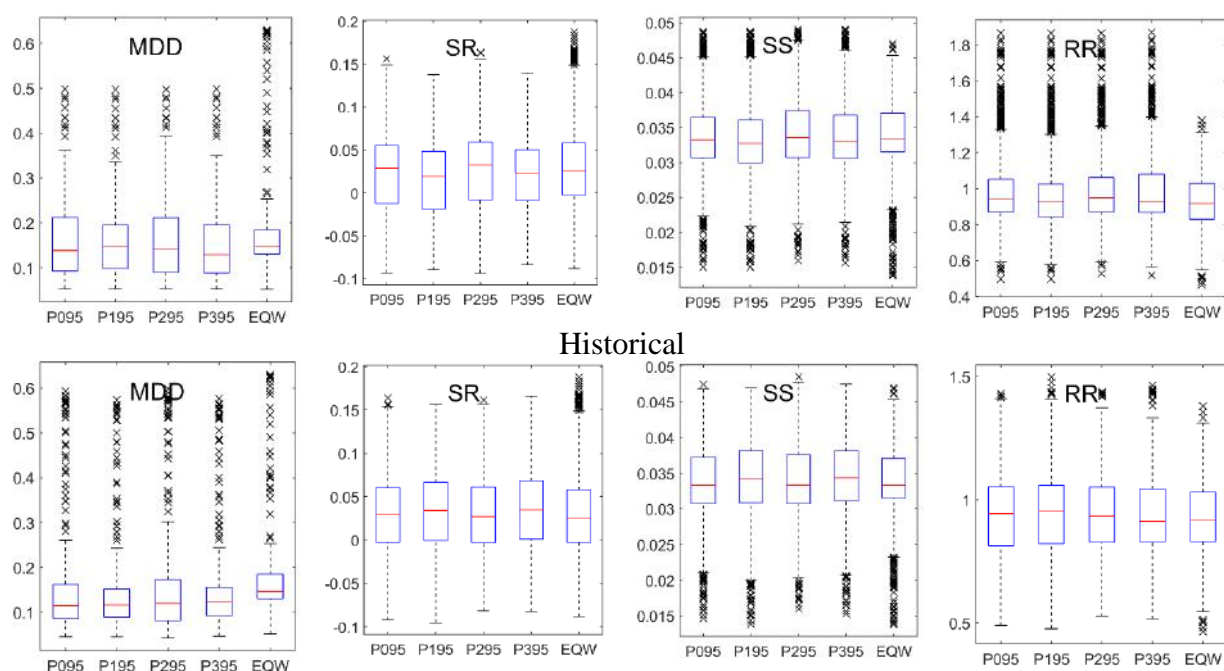


Figure 11.8 Performance of annual MDD, SR, SS, and RR for historical (left) and dynamic (right) $\text{CVar}_{0.95}$ -minimized portfolios under a 0.4% turnover constraint.



Dynamic

Figure 11.9 Box-whisker summary of annual MDD, SR, SS, and RR for historical and dynamic $\text{CVaR}_{0.95}$ -minimized portfolios under a 0.4% turnover constraint.

Fig. 11.10 summarizes the total-time-period (12-year) risk measures for the $\alpha = 0.95$ portfolios constrained by a daily turnover of 0.4%. All the MDDs reflect behavior related to the onset of the COVID-19 pandemic. The EQW portfolio has the largest drawdown, and the historical portfolios have the smallest. The drawdown values for the various historical portfolios differ by at most 0.1%, whereas those for the dynamic portfolios differ by at most 2.5%, with $P_{0.95}^1$ and $P_{0.95}^3$ performing the best in the dynamic case. Interestingly, the EQW portfolio has the best 12-year SR and RR, though it has the worst SS. Under dynamic optimization, the performance-attribute-constrained portfolios outperform the base case $P_{0.95}^0$ with respect to SR, SS, and RR. Under historical optimization, $P_{0.95}^2$ consistently outperforms $P_{0.95}^0$ over this 12-year period.




Figure 11.9 Twelve-year MDD, SR, SS, and RR for historical and dynamic $\text{CVaR}_{0.95}$ -minimized portfolios under a 0.4% turnover constraint.

References

- Artzner, P., Delbaen, F., Eber, J.-M. & Heath, D. (1999). Coherent measures of risk. *Risk*, 10, 203–228.
- Bacon, C. R. (2008). *Practical portfolio performance measurement and attribution* (2nd ed.). John Wiley & Sons, West Sussex.
- Biglova, A. & Rachev, S. T. (2007). Portfolio performance attribution. *Investment Management and Financial Innovations*, 4, 7–22.
- Brinson, G. P. & Fachler, N. (1985). Measuring non-United-States equity portfolio performance. *Journal of Portfolio Management*, 11, 73–76.
- Brinson, G. P., Hood, L. R. & Beebower, G. L. (1986). Determinants of portfolio performance. *Financial Analysts Journal*, 42, 39–44.
- Cherdito, P. & Kromer, E. (2013). Reward-risk ratios. *Journal of Investment Strategies*, 3, 1–16.
- Gava, J., Guevara, F. & Turc, J. (2021). Turning tail risks into tailwinds. *Journal of Portfolio Management*, 47(4), 41–70.
- JP Morgan (1996). *Risk metrics technical manual* (4th ed.). JP Morgan, New York.
- Konno, H. & Yamazaki, H. (1991). Mean-absolute deviation portfolio optimization model and its application to Tokyo stock market. *Management Science*, 37, 519–531.

11. Optimization with Performance Attribution Constraints

- Markowitz, H. (1952). Portfolio selection*. *Journal of Finance*, 7, 77–91.
- Pflug, G. C. (2000). Some remarks on the value-at-risk and the conditional value-at-risk. In S. P. Uryasev (Ed.), *Probabilistic constrained optimization: Methodology and applications*, pp. 272–281. Springer, Boston.
- Rachev, S. T., Martin, R. D., Racheva, B. & Stoyanov, S. (2009). Stable ETL optimal portfolios and extreme risk management. In G. Bol, S. T. Rachev, R. Würth (Eds.), *Risk assessment*, pp. 235–262. Physica-Verlag HD, Heidelberg.
- Rockafellar, R. T. & Uryasev, S. (2000). Optimization of conditional value-at-risk. *Journal of Risk*, 2, 21–41.

Chapter 12

Option Pricing

In this chapter, we simultaneously address three innately coupled topics: asset pricing, option pricing, and the prediction of future asset volatility. The intimate coupling of these three can be made apparent by a brief recapitulation of the Black–Scholes–Merton (Black and Scholes, 1973; Merton, 1973) model. (Here we follow Duffie, 2001.) The Black–Scholes–Merton option pricing model has the following components: a risky asset \mathcal{S} ; a riskless asset \mathcal{B} (a bond); an option \mathcal{C} whose underlying asset is \mathcal{S} ; and a trading strategy (a_t, b_t) for \mathcal{S} and \mathcal{B} that is self-financing and, relative to the price of \mathcal{C} , arbitrage-free. Assume the price process for \mathcal{S} is given by geometric Brownian (equivalently log-normal) motion,

$$S_t = S_0 e^{X_t}, \quad t \geq 0, \quad S_0 > 0, \quad (12.1)$$

having normally distributed cumulative log-return (“return” for brevity),

$$X_t = \ln(S_t) - \ln(S_0) = \mu t + \sigma W_t, \quad \mu \in \mathbb{R}, \quad \sigma > 0, \quad (12.2)$$

where $\mathbb{B} = (W_t, t \geq 0)$ is a standard Brownian motion. The constant μ is the drift term and the standard deviation σ is the volatility of the price. The stock price is assumed to be an Ito process obeying

$$dS_t = (\mu + \sigma^2/2)dt + \sigma dW_t. \quad (12.3)$$

Assume \mathcal{B} has the price process

$$B_t = B_0 e^{r_f t}, \quad t \geq 0, \quad B_0 > 0, \quad (12.4)$$

where r_f is the risk-free rate. The bond price is trivially an Ito process obeying the ODE

$$dB_t = r_f B_t dt. \quad (12.5)$$

Consider a European call option with strike price K and maturity time T having the price process $C(S_t, t; K, T)$. The trading strategy is self-financing if

$$a_t S_t + b_t B_t = a_0 S_0 + b_0 B_0 + \int_0^t a_w dS_w + \int_0^t b_w dB_w \quad (12.6)$$

for all $t \geq 0$, and guarantees an arbitrage-free market $(\mathcal{S}, \mathcal{B}, \mathcal{C})$ provided that

$$a_t S_t + b_t B_t = C(S_t, t, T, K) \quad (12.7)$$

for all $t \in [0, T]$, where (12.7) satisfies the boundary condition at maturity, $a_T S_T + b_T B_T = \max(S_T - K, 0)$. This boundary condition guarantees that any losses or gains undertaken by a long position in \mathcal{S}, \mathcal{B} are, at maturity, precisely hedged by a short position in the call option. Under these conditions¹⁰⁰, Black and Scholes (1973) and Merton (1973) derived the value of the call option for all $t \in [0, T]$,

$$C(S_t, t, T, K) = C(S_t, T - t, K) = S_t \Phi(U) - e^{-r_f(T-t)} K \Phi(U - \sigma \sqrt{T-t}), \quad (12.8)$$

¹⁰⁰ There are additional assumptions implied in the Black–Scholes–Merton formulation, one being that investors are risk-neutral (indifferent to risk) when making investment decisions.

12. Option Pricing

$$U = \frac{\ln(S_t/K) + (r_f + \sigma^2/2)(T - t)}{\sigma\sqrt{T - t}},$$

where $\Phi(\cdot)$ is the cumulative standard normal distribution function. Given market values for: the risk-free rate r_f ; the asset price S_t ; and the price of a call option with strike price K and maturity time T ; (12.8) can be solved numerically for the volatility σ of the risky asset at any future time $t \in [0, T]$. Such a computed value is referred to as an *implied volatility* and is an example of volatility metrics that attempt to utilize current market sentiment (pricing) to measure future risk. Equations (12.1) through (12.8) demonstrate the tight coupling between asset prices, option prices, and implied volatility.

In addition to continuum models such as the Black–Scholes–Merton formula, discrete models based on binomial or trinomial trees (Cox et al., 1979), discrete stochastic volatility models (Nögel, 2004), and Monte Carlo simulation (Boyle, 1977; Broadie and Glasserman, 1996; Carriere, 1996) have all been used to price derivative contracts. All contain the same set of fundamental components: a risky asset \mathcal{S} ; a riskless asset \mathcal{B} ; a derivative \mathcal{C} whose underlying asset is \mathcal{S} ; and a trading strategy (a_t, b_t) for \mathcal{S} and \mathcal{B} that is self-financing and, relative to the price of \mathcal{C} , arbitrage-free. These methods require more sophisticated methods to achieve an equivalent (risk-neutral) martingale measure either through Esscher transform (Gerber and Shiu, 1994) or mean correction (Schoutens, 2003; Yao et al., 2016) to compute the option pricing under a risk-neutral formulation.

Computations of the implied volatility from (12.8) (or from any of the alternate derivative pricing methods) using market data reveal immediately that the implied volatility for any asset varies with K (equivalently with “moneyness” K/S_t) in a manner that evokes either the shape of a “smile” or a “smirk”. Consequently, implied volatility has been criticized as follows: “A smiley implied volatility is the wrong number to put [*sic*] in the wrong formula to obtain the right price.” (Rebonato, 1999; page 78). The presence of the characteristic smile/smirk seen in implied volatility computations, which characterizes different volatility values for options that are differentially in- and out-of-the-money, contradicts the fundamental option-pricing model result (12.8) that the volatility should be constant. Taking one approach to correct this, Hull and White (1987) and Heston (1993) randomized the volatility parameter (asset return standard deviation) in the Black–Scholes–Merton model.

As they are irrevocably coupled, improved option pricing models (and hence the potential for improved volatility prediction) must begin with improved asset pricing models. A new approach to asset pricing was adopted by Carr et al. (2003) based upon ideas by Geman et al. (2001) and Clark (1973). Clark conjectured that price processes are controlled by a random clock, which dictates arrival time of financial information and hence provides the underlying measure of economic activity. Clark used transaction volume as a proxy for this measure. We briefly expand upon this random clock concept of intrinsic or business time before proceeding with our discussion of these improved models.

Intrinsic time. In finance, the phrase “intrinsic time” was first used by Mandelbrot and Taylor (1967) to refer to the fact that market information does not arrive continuously in time but rather arrives in a series of events which, when measured in the continuum of physical time, occur differentially spaced. The prototypical example is market orders (“ticks”) for any asset, which

display sizeable variation in terms of timing and number of orders throughout the trading day.¹⁰¹ Asset value (price) information therefore arrives as discrete events spaced at different time intervals. In addition, the price change information may be more-or-less informative: small price changes provide less information than larger changes; consecutive price changes having opposite signs provide less information than consecutive price changes having the same sign. Thus, intrinsic time is characterized by the occurrence, magnitude and sign of each event. Several works have been devoted to developing the concepts of intrinsic time and applying them to financial time series; see Guillaume et al. (1997), Tsang (2010), and Aloud et al. (2011).

Under the intrinsic time viewpoint, no information arrives and hence “no time has passed” between events. One way to conceptualize intrinsic time is to think of the seconds-hand of an analog clock that doesn’t move (tick) regularly, but rather advances randomly and makes a ticking noise of varying volume. No time passes until the hand advances, and the volume of the tick quantifies the informational content.

Subordinated pricing models. The assumption that the log-price of the underlying risky asset has a normal distribution provides the basis of the Black–Scholes–Merton model. There are, so-called, stylized facts (asymmetry, heavy tails) characterizing asset returns that contradict that assumption. Unfortunately, distributions such as Student’s t , which exhibits heavy tails, and stable-non-Gaussian, which exhibit skewness and heavy tails, while potentially useful for asset pricing, are not amenable for use by the classic option pricing methods using either the Esscher transform (Gerber and Shiu, 1994) or mean correction (Schoutens, 2003; Yao et al., 2016) because they cause required integrals to diverge (Cassidy et al., 2010). To deal with the non-Gaussian behavior of asset returns, models based on subordinated Brownian motion have been introduced (Mandelbrot and Taylor, 1967; Clark, 1973). In a subordinated Brownian motion model, the return process (12.2) is replaced by

$$= \mu t + \rho \mathcal{T}(t) + \sigma W_{\mathcal{T}(t)}. \quad (12.9)$$

Here, $\mathbb{T} = (\mathcal{T}(t), t \geq 0, \mathcal{T}(0) = 0)$ is a Lévy subordinator; $W_{\mathcal{T}(t)}$ is a singly subordinated Brownian motion; and X_t is therefore a singly subordinated return process. The stochastic process \mathbb{T} defined on a stochastic basis $(\Omega, \mathcal{F}, \mathbb{F} = (\mathcal{F}_t, t \geq 0), P)$ is said to be a Lévy process if the following hold (Bochner, 1955; Sato, 1999; Schoutens, 2003).

- $\mathcal{T}(0) = 0$ P – almost surely. (I.e., $P(\mathcal{T}(0) = 0) = 1$.)
- \mathbb{T} is a process with independent increments; for any partition $0 = t_0 < t_1 < \dots < t_n < \infty$, the increments $\mathcal{T}(t_i) - \mathcal{T}(t_{i-1})$, $i = 1, 2, \dots, n$ are independent.
- \mathbb{T} is a process with stationary increments; for any $0 \leq s < t$, the increment $\mathcal{T}(t) - \mathcal{T}(s)$ has the same distribution as $\mathcal{T}(t - s)$, that is, $\mathcal{T}(t) - \mathcal{T}(s) \sim \mathcal{T}(t - s)$. (We use \sim to denote “equal in distribution” or “equal in probability law”.)
- \mathbb{T} is continuous in probability process; for every $\varepsilon > 0$ and $t \geq 0$, there exists $h_{\varepsilon, t} > 0$ such that $P(|\mathcal{T}(t + h_{\varepsilon, t}) - \mathcal{T}(t)| > \varepsilon) < \varepsilon$.

Because $\mathcal{T}(0) = 0$, the trajectories of a Lévy process can take only nonnegative values. Any Lévy process \mathbb{T} with nondecreasing trajectories is called a Lévy subordinator. The subordinated return

¹⁰¹ “The different evolution of price series on different days is due to the fact that information is available to traders at a varying rate. On days when no new information is available, trading is slow, and the price process evolves slowly. On days when new information violates old expectations, trading is brisk, and the price process evolves much faster.” (Clark, 1973)

12. Option Pricing

process X_t in (12.9) is a Lévy process (Sato, 1999; Chapter 6), and thus a martingale. By the fundamental asset pricing theorem (Delbaen and Schachermayer, 1994), the market $(\mathcal{S}, \mathcal{B}, \mathcal{C})$ is free of arbitrage.

Each subordinator introduces additional parameters to the return model, which can be used to model additional behaviors. As the Lévy subordinator takes the role of a time parameter, it acts as a second-hand, defining and providing information (direction and magnitude) on each event. Considering various distribution models for the subordinator, Hurst et al. (1997) used the method of subordination to model the leptokurtic characteristics of market index returns. Geman et al. (2001) investigated pure jump processes for the subordinator. Carr et al. (2003) extended this line of investigation using normal inverse Gaussian and variance gamma examples of pure jump Lévy processes. Carr and Wu (2004) carried the subordinated asset pricing concept into option pricing using time-changed Lévy processes.¹⁰² Klingler et al. (2013) introduced two new six-parameter processes based on time-changed tempered stable distributions and developed the corresponding option pricing models.

Studies have demonstrated that single subordinated models cannot resolve some aspects of observed market behavior. For example, Lundtofte and Wilhelmsson (2013) and Shirvani et al. (2021a) show that single subordination fails to explain the equity premium puzzle. To resolve this and to provide greater flexibility to model higher moments (skewness and kurtosis) observed in asset return behavior, Shirvani et al. (2020) introduced the concept of double subordination. The first subordinator is interpreted as being responsible for transforming a Gaussian distribution of returns in unit time to a distribution of skewed and heavy tailed returns. The second subordinator transforms the now skewed, heavy-tailed-distributed return events from unit-time spacing to random spacing. (In their 2020 article, these subordinators were interpreted in terms of time-varying investor views, as per behavioral finance dictum that investors view positive and negative returns differently, which – in turn – affects investor actions and, consequently, asset and option pricing.) This work demonstrated that double subordinated log-return processes produce heavy-tailed distributions with heavier tails than single subordinated models and are capable of capturing the skewness and kurtosis observed in distributions of asset log-returns. Shirvani et al. (2021a) further demonstrated that the high value for the risk-aversion coefficient, which gives rise to the equity premium puzzle, can be obtained from a return process driven by a double subordinated model. In another application, Shirvani et al. (2021b) used a double subordinated Lévy process to model the high volatility of bitcoin.

Volatility. Volatility refers to the degree of unpredictability in a market. The volatility can be quantified in a variety of ways. *Historical volatility* measures the dispersion of a portfolio's returns using the (moving-window averaged) standard deviation of past daily log-returns. The concept of measuring an implied volatility was introduced above in the context of the Black–Scholes–Merton model. Unlike historical volatility, implied volatility is forward looking. The best-known (and first practical) implementation of implied volatility is the Cboe VIX, developed to reflect volatility expectations for the S&P 500 stock index (SPX). The intention behind the VIX model (Cboe

¹⁰² Fallahgoul and Nam (2020) have pointed out a subtle error in the Carr and Wu paper. By introducing a correlation between the driving Lévy process and its time-subordinator, the resulting Carr–Wu option pricing model is no longer arbitrage-free. For this reason, each subordinator should be independent of each other and of the driving (Brownian or Lévy) process.

Exchange, Inc., 2019) is to reflect all available market information determining SPX future volatility, provided that option markets are efficient¹⁰³. Specifically, the value of the VIX is intended to measure the volatility of SPX expected in 30 days and is given by a weighted portfolio of near-term and next-term, out-of-the-money, European put and call options with more than 23 days and less than 37 days to expiration. The options include those with standard 3rd Friday expiration dates and weekly options that expire on the remaining Fridays of each month. As an example of a different approach, Petrov et al. (2019) develop an *instantaneous volatility* measure based upon intrinsic time analysis of price time series and used the measure to explore seasonality patterns in the volatility of selected exchange rates.

Section 12.1 summarizes the double subordinated model for the case in which both subordinators are independent, inverse Gaussian processes. The model is developed in terms of the characteristic function of the log-return process. Section 12.2 describes option pricing under this model using the mean-correction martingale measure of Yao et al. (2016) and the fast Fourier transform (FFT) method of Carr and Madan (1998). Section 12.3 provides a numerical application where the double subordinated model is fit to the price process of one of the REIT portfolios developed in Chapter 4, specifically the domestic, long-only, T95 portfolio subjected to 4% daily turnover constraint. For brevity in this chapter, this portfolio will be referred to as the T95 portfolio. An in-depth discussion is presented for estimating the model parameters as well as for setting numerical parameters needed in the Carr-Madan FFT method. Prices are developed for European call and put options with the T95 portfolio as the underlying risky asset. The double subordinated method leads naturally to a new measure of volatility, which we refer to as the *NDIG volatility*. This volatility relies on asset prices (and not on option prices). We compute the NDIG volatility for the T95 portfolio over an in-sample data set. The results are compared to those obtained using historical volatility as well as those obtained by using the Cboe VIX implied volatility methodology. As the VIX methodology requires option prices, call and put option prices computed from the in-sample data using the double-subordinated method are used.

12.1 Double Subordinated Pricing Models

Under double subordination, the return process (12.2) is replaced by

$$X_t = \mu t + \gamma U(t) + \rho \mathcal{T}(U(t)) + \sigma W_{\mathcal{T}(U(t))}, \quad (12.10)$$

where the triplet $(W_t, \mathcal{T}(t), U(t), t \geq 0)$ consists of independent processes generating a stochastic basis $(\Omega, \mathcal{F}, \mathbb{F} = (\mathcal{F}_t, t \geq 0), P)$ representing the natural world. Here $\{W_t, t \geq 0\}$ is again a standard Brownian motion while $\mathcal{T}(t)$ with $\mathcal{T}(0) = 0$ and $U(t)$ with $U(0) = 0$ are independent Lévy subordinators. The process $\mathcal{T}(U(t))$ is a double subordinator process, hence (12.10) is a double subordinated return process.

Consider the case when the subordinators $\mathcal{T}(t)$ and $U(t)$ are inverse Gaussian (IG) Lévy processes¹⁰⁴; that is, $\mathcal{T}(1) \sim \text{IG}(\lambda_{\mathcal{T}}, \mu_{\mathcal{T}})$ having the probability density function (PDF)

¹⁰³ The accuracy, and indeed the ability, of the VIX to predict future volatility is a matter of debate. See Goldstein and Taleb (2007) and Adhikari and Hilliard (2014).

¹⁰⁴ We have already indicated that stable, non-Gaussian processes are inappropriate choices for subordinators due to their heavy tails. Use of (finite variation) gamma processes as subordinators is appropriate but may require additional Brownian motions, which would increase the number of parameters in the model.

12. Option Pricing

$$f_{\mathcal{T}(1)}(x) = \sqrt{\frac{\lambda_{\mathcal{T}}}{2\pi x^3}} \exp\left(-\frac{\lambda_{\mathcal{T}}(x - \mu_{\mathcal{T}})^2}{2\mu_{\mathcal{T}}^2 x}\right), \quad \lambda_{\mathcal{T}} > 0, \quad \mu_{\mathcal{T}} > 0, \quad x \geq 0, \quad (12.11)$$

and $U(1) \sim \text{IG}(\lambda_U, \mu_U)$. Shirvani et al. (2020) referred to $\mathcal{T}(U(t))$ as the double IG subordinator and the resulting X_t as the normal double IG (NDIG) log-price process. The probability distribution for X_t can be written as a double integral (Shirvani et al., 2020; Supplementary Appendix A.10), which is unfortunately analytically intractable. However, the characteristic function (CF) (the Fourier transform of the probability density function) $\varphi_{X_t}(v)$ is amenable to analytic evaluation. Using the identity

$$\varphi_{X_t}(v) = E^{(P)}[e^{ivX_t}] = (E^{(P)}[e^{ivX_1}])^t = (\varphi_{X_1}(v))^t, \quad (12.12)$$

the CF for X_1 is given by (Shirvani et al., 2020; supplementary material Appendix A.11)

$$\begin{aligned} \varphi_{X_1}(v) &= E^{(P)}[e^{ivX_1}] \stackrel{\text{def}}{=} \exp\{\psi_{X_1}(v)\} \\ &= \exp\left\{iv\mu + \frac{\lambda_U}{\mu_U} \left[1 - \left(1 - \frac{2\mu_U^2}{\lambda_U} \left\{\frac{\lambda_{\mathcal{T}}}{\mu_{\mathcal{T}}} \left[1 - \sqrt{1 - \frac{\mu_{\mathcal{T}}^2}{\lambda_{\mathcal{T}}} (2iv\rho - \sigma^2 v^2)}\right] + iv\gamma\right\}\right)^{1/2}\right]\right\}. \end{aligned} \quad (12.13)$$

Here, $E^{(P)}[\cdot]$ denotes expectation under the probability P and $\psi_{X_1}(v)$ is referred to as the characteristic exponent.

The NDIG model (12.10), (12.13) has eight parameters, $\mu, \sigma, \gamma, \rho, \mu_U, \lambda_U, \mu_{\mathcal{T}}, \lambda_{\mathcal{T}}$, which provides a challenge for fitting to data. Appendix A provides the proof that only six of these parameters are identifiable within the model. To set the remaining two, consider the expectation

$$E^{(P)}[X_1] = \mu + \gamma E^{(P)}[U(1)] + \rho E^{(P)}[\mathcal{T}(U(1))].$$

As the processes U and \mathcal{T} are independent IG, to uniquely identify γ and ρ we can require

$$E^{(P)}[U(1)] = \mu_U = 1 \text{ and } E^{(P)}[\mathcal{T}(U(1))] = \mu_U \mu_{\mathcal{T}} = 1 \rightarrow \mu_{\mathcal{T}} = 1. \quad (12.14)$$

Using (12.14), the set of identifiable model parameters becomes $\mu, \sigma, \gamma, \rho, \lambda_U, \lambda_{\mathcal{T}}$.

One source of data for fitting these six parameters are the central moments of the NDIG. The moment generating function (MGF) $M_{X_1}(w)$ for X_1 , which generates the moments of its probability distribution, is obtained by evaluating $E^{(P)}[e^{wX_1}]$, $w \in \mathbb{R}$. This can be obtained from (12.13) using $w = iv$. The MGF is written in terms of the cumulant generating function $K_{X_1}(w)$,

$$M_{X_1}(w) = E^{(P)}[e^{wX_1}] = \exp\{\psi_{X_1}(-iw)\} \stackrel{\text{def}}{=} \exp\{K_{X_1}(w)\}. \quad (12.15)$$

Paralleling (12.12), note that $M_{X_t}(w) = \exp\{K_{X_1}(w)t\}$. From (12.13), $K_{X_1}(w)$ can be written as

$$\begin{aligned} K_{X_1}(w) &= \mu w + \lambda_U \left(1 - \sqrt{g(w)}\right), \\ g(w) &= 1 - 2\frac{\lambda_{\mathcal{T}}}{\lambda_U} \left(1 - \sqrt{h(w)}\right) - 2\frac{\gamma}{\lambda_U} w, \\ h(w) &= 1 - 2\frac{\rho}{\lambda_{\mathcal{T}}} w - \frac{\sigma^2}{\lambda_{\mathcal{T}}} w^2. \end{aligned} \quad (12.16)$$

The first four central derivatives of $K_{X_1}(w)$ are

$$\begin{aligned}
 d K_{X_1}(w)/dw|_{w=0} &= \mu + s, \\
 d^2 K_{X_1}(w)/dw^2|_{w=0} &= \sigma_1 + \frac{s^2}{\lambda_U}, \\
 d^3 K_{X_1}(w)/dw^3|_{w=0} &= 3 \left(\sigma_1 c + \frac{s^3}{\lambda_U^2} \right), \\
 d^4 K_{X_1}(w)/dw^4|_{w=0} &= 3 \left[\left(\frac{1}{\lambda_T} + \frac{1}{\lambda_U} \right) \sigma_1^2 + 2\sigma_1 \left(c^2 + \left(\frac{\rho}{\lambda_T} \right)^2 + 2 \left(\frac{s}{\lambda_U} \right)^2 \right) + 5 \frac{s^4}{\lambda_U^3} \right], \\
 s &= \gamma + \rho, \quad \sigma_1 = \frac{\rho^2}{\lambda_T} + \sigma^2, \quad c = \frac{\rho}{\lambda_T} + \frac{s}{\lambda_U}.
 \end{aligned} \tag{12.17}$$

The first four centered moments of X_1 are

$$\begin{aligned}
 E^{(P)}[X_1] &= \mu + s, \\
 \text{Var}[X_1] &= \sigma_1 + \frac{s^2}{\lambda_U}, \\
 \text{Skew}[X_1] &= 3 \frac{\left(\sigma_1 c + \frac{s^3}{\lambda_U^2} \right)}{\left(\sigma_1 + \frac{s^2}{\lambda_U} \right)^{3/2}}, \\
 \text{Kurt}[X_1] &= 3 \frac{\left[\left(\frac{1}{\lambda_T} + \frac{1}{\lambda_U} \right) \sigma_1^2 + 2\sigma_1 \left(c^2 + \left(\frac{\rho}{\lambda_T} \right)^2 + 2 \left(\frac{s}{\lambda_U} \right)^2 \right) + 5 \frac{s^4}{\lambda_U^3} \right]}{\left(\sigma_1 + \frac{s^2}{\lambda_U} \right)^2}.
 \end{aligned} \tag{12.18}$$

Equation (12.18) provides four conditions for fitting the six parameters of the model. The numerical example in section 12.3 discusses the further conditions needed for the parameter fits.

12.2 Option Pricing under the Double Subordinated IG Model

We assume the price process (12.1) for the T95 portfolio is driven by a double subordinated IG process; i.e., its return process (12.10) is NDIG. To price a European call option with $\mathcal{S} = \text{T95}$ as the underlying asset, we search for an equivalent martingale measure (EMM) Q of P such that the discounted price process, $e^{-r_f t} S_t$, is a martingale under Q (Duffie, 2001, Chapter 6). Following Yao et al. (2016; Theorem 3.2), the mean-correction martingale measure method can be used to find an EMM for X_t . Parametrizing this measure as $Q^{(m)}$, the mean-correction martingale method requires

$$Q^{(m)}(X_t - mt \leq x) = P(X_t \leq x), \quad m \in R, \tag{12.19}$$

which has the unique solution $Q = Q^{(m_0)}$ with $m_0 = r_f - K_{X_1}(1)$. If the price and return processes for \mathcal{S} are given by (12.1) and (12.10) respectively, using mean-correction martingale measure the risk-neutral price $S_t^{(Q)}$ for \mathcal{S} is given by

$$S_t^{(Q)} = \frac{B_t S_t}{M_{X_t}(1)} = S_0 \exp\{(r_f - K_{X_1}(1))t + X_t\}, \quad t \in [0, T]. \quad (12.20)$$

The CF of the log-price, $\ln(S_t^{(Q)})$, is

$$\varphi_{\ln(S_t^{(Q)})}(v) = S_0^{iv} \exp\left\{\left(iv[r_f - K_{X_1}(1)] + \psi_{X_1}(v)\right)t\right\}. \quad (12.21)$$

Then, for the underlying asset S having price S_t on day t , a European call option \mathcal{C} expiring on day $t + T$ (maturing T days after day t) has the fair price

$$C(S_t, T, K) = e^{-r_f T} E^{(Q)} \left[\max(S_T^{(Q)} - K, 0) \right]. \quad (12.22)$$

Carr and Madan (1998) show how to use the fast Fourier transform (FFT) to value options when the CF of the log-price of the underlying asset is known analytically. They considered the modified option price $C_a(S_t, T, k) = e^{ak} C(S_t, T, K)$ where $k = \ln(K)$, which, for a range of values $a > 0$, guarantees that $C_a(S_t, T, k)$ is square integrable over $k \in (-\infty, \infty)$. The starting point for applying the FFT is their derived relationship

$$C(S_t, T, k) = \frac{e^{-r_f T - ak}}{\pi} \int_0^\infty e^{-ivk} \frac{\varphi_{\ln(S_T^{(Q)})}[-i(1 + a + iv)]}{(a + iv)(1 + a + iv)} dv. \quad (12.23)$$

Numerical solution of this integral involves two fundamental concerns, an optimum value for a and control over the error produced by truncating the integral (12.23) over a finite domain $v \in [0, v_{\max}]$. These concerns, which are dependent on the price process, are addressed in section 12.3.1.

Constructing the implied volatility surface for call prices requires estimating (12.23) over a discrete mesh of values of (k, T) . Put options can then be valued assuming put–call parity holds. Implementing an FFT requires a numerical discretization of (12.23) of the form

$$\hat{Z}_p = \sum_{j=0}^{M-1} \exp\left[-i \frac{2\pi}{M} jp\right] Z_j, \quad p = 0, \dots, M-1, \quad (12.24)$$

which the FFT method can solve in $O(M \ln_2(M))$ operations. Writing (12.23) as

$$C(S_t, T, k) = \frac{e^{-r_f T - ak}}{\pi} \Psi(k), \quad \text{where } \Psi(k) = \int_0^{v_{\max}} e^{-ivk} h(S_t, r_f, v) dv, \quad (12.25)$$

discretization of $\Psi(k)$ using the trapezoid rule gives

$$\Psi(k) = \sum_{j=0}^{N-1} \exp(-iv_j k) h(S_t, r_f, v_j) w_j \Delta v, \quad (12.26)$$

where $v_j = j\Delta v$, $j = 0, \dots, N-1$ with $\Delta v = v_{\max}/(N-1)$. The w_j are the trapezoid rule weights, $w_0 = w_{N-1} = 1/2$, $w_j = 1$, $j = 1, \dots, N-2$. Consider the discretization of k over a range $[-\bar{k}_l, \bar{k}_h]$ with N equally spaced points, $k_p = -\bar{k}_l + p\Delta k$, $p = 0, \dots, N-1$ with $k_{N-1} = \bar{k}_h$.¹⁰⁵ For simplicity, consider $\bar{k}_l = \bar{k}_h = \bar{k}$. On this grid, (12.26) becomes

¹⁰⁵ This discretizes strike prices $K = e^k$ over the range $[e^{-\bar{k}_l}, e^{\bar{k}_h}]$. As the k -grid is evenly spaced, strike prices will

$$\Psi(k_p) = \sum_{j=0}^{N-1} \exp(-i j \Delta v p \Delta k) \exp(i \bar{k} v_j) h(S_t, r_f, v_j) w_j \Delta v, \quad (12.27)$$

$$p = 0, \dots, N-1,$$

which is identical to (12.24) with the identifications $M = N$, $2\pi/M = \Delta v \Delta k$, $\hat{Z}_p = \Psi(k_p)$, and $Z_j = \exp(i \bar{k} v_j) h(S_t, r_f, v_j) w_j \Delta v$. The second identification specifies the size of the k -grid spacing $\Delta k = 2\pi/(N \Delta v)$ and reflects the familiar FFT tradeoff between the spans covered in the “space” and “frequency” domains,

$$v_{\max} \bar{k} = \frac{\pi(N-1)^2}{N} \sim \pi N. \quad (12.28)$$

The computations in section 12.3 use $N = 2^{10} = 1024$.

12.3 Empirical Example

We illustrate the double subordinated method of sections 12.1 and 12.2 using the T95 REIT portfolio as the underlying asset. To further reduce the parameter set $\mu, \sigma, \gamma, \rho, \lambda_U, \lambda_T$, assume that the subordinator $U(t)$ is used to model the intrinsic time of the return process, while the subordinator $\mathcal{T}(t)$ is used to model the return skewness and heavy tailed behavior. In this model it is reasonable to require no $\gamma U(t)$ term in (12.10), i.e. $\gamma = 0$. Under this requirement, the first four moments of X_1 reduce to

$$\begin{aligned} E^{(P)}[X_1] &= \mu + \rho, \\ \text{Var}[X_1] &= \sigma^2 + d\rho^2, \\ \text{Skew}[X_1] &= 3 \frac{d\rho(\sigma^2 + \rho^2)}{(\sigma^2 + d\rho^2)^{3/2}}, \\ \text{Kurt}[X_1] &= \frac{3\sigma_1^2 d + 6\rho^2 \sigma_1 \left[d^2 + \frac{1}{\lambda_T^2} + \frac{2}{\lambda_U^2} \right] + 15\rho^4 \frac{1}{\lambda_U^3}}{(\sigma^2 + d\rho^2)^2}, \\ \sigma_1 &= \frac{\rho^2}{\lambda_T} + \sigma^2, \quad d = \frac{1}{\lambda_T} + \frac{1}{\lambda_U}. \end{aligned} \quad (12.29)$$

The five remaining parameters $\mu, \sigma, \rho, \lambda_U, \lambda_T$ can be estimated by the minimization¹⁰⁶

$$\begin{aligned} \min_{\mu, \sigma, \rho, \lambda_U, \lambda_T} \{ & (\Delta M_1)^2 + (\Delta M_2)^2 + (\Delta M_3)^2 + (\Delta M_4)^2 + (\Delta CF)^2 \}, \\ \Delta M_1 &= 1 - E^{(P)}[X_1]/E^{(P)}[r_{T95}], \quad \Delta M_2 = 1 - \text{Var}[X_1]/\text{Var}[r_{T95}], \\ \Delta M_3 &= 1 - \text{Skew}[X_1]/\text{Skew}[r_{T95}], \quad \Delta M_4 = 1 - \text{Kurt}[X_1]/\text{Kurt}[r_{T95}], \end{aligned} \quad (12.30)$$

be exponentially spaced.

¹⁰⁶ An alternate minimization that uses the first five normalized moment differences $\Delta M_1, \dots, \Delta M_5$ must begin to deal with the fact that higher order moments of r_{T95} get noisier and noisier, providing no valid signal beyond a certain order.

12. Option Pricing

$$(\Delta CF)^2 = \int_{-\infty}^{\infty} |1 - \varphi_{X_1}(v)/\text{ecf}(v, r_{T95})|^2 dv.$$

In (12.30), r_{T95} denotes the observed return time series $r_{T95,j}$, $j = 1, \dots, n = 3273$ for the T95 portfolio and $\text{ecf}(v, r_{T95}) = n^{-1} \sum_{j=1}^n \exp(i v r_{T95,j})$ is the empirical characteristic function. The inclusion of the term $(\Delta CF)^2$ relies on the one-to-one correspondence between the cumulative distribution function and the CF (since the PDF is the inverse Fourier transform of the CF). The integral for $(\Delta CF)^2$ can be estimated as described by Yu (2004).




Figure 12.1 T95 portfolio time series of (top) price (assuming an initial investment of \$100) and (bottom) log-return.

Table 12.1 Preset and estimated parameter values for the fit of the NDIG distribution to daily log-return values of the T95 portfolio

	μ	σ	γ	ρ	μ_U	λ_U	μ_T	λ_T
global	$9.0 \cdot 10^{-4}$	$1.9 \cdot 10^{-2}$	0	$-4.4 \cdot 10^{-4}$	1	0.192	1	9.87
mean	$1.2 \cdot 10^{-3}$	$1.4 \cdot 10^{-2}$	0	$-5.9 \cdot 10^{-4}$	1	0.45	1	9.77
std	$0.4 \cdot 10^{-3}$	$0.5 \cdot 10^{-2}$	NA	$4.4 \cdot 10^{-4}$	NA	0.13	NA	0.10

As noted in Chapter 4, 3,273 daily log-returns for the T95 portfolio were computed covering the period from 12/19/2007 through 12/18/2020. The price and return time series for this portfolio are shown in Fig. 12.1. Using the vector of log-returns $r = \{r_1, \dots, r_{3273}\}$ in (12.30) produced the parameter values given in the row labelled “global” in Table 12.1. To test the variation of the values over time, the fits (12.30) were performed using a moving window of 4 years (1,008 days). The resultant time series, covering the period 12/19/2011 through 12/18/2020, are shown in Fig. 12.2, while mean and standard deviation values computed for the time series are summarized in Table 12.1. Comparison of the global and rolling window mean values, as well as the standard deviation values, indicates that λ_T has the smallest variation with time. The global values measured

for μ and σ are within one standard deviation from their respective rolling window means while the global value for λ_U is within two standard deviations from the rolling window mean. The parameter ρ has the smallest value and a large standard deviation. Fig. 12.3 compares the first four empirical moments computed from the time series in Fig. 12.1 and the moments computed from the fitted parameters using (12.18).




Figure 12.2 Four-year moving window fits to the parameter values $\mu, \sigma, \rho, \lambda_U, \lambda_T$.




Figure 12.3 Comparison of the first four empirical (“emp”) moments computed from the return time series in Fig. 12.1 and the moments (“th”) computed from the fitted parameters.

12. Option Pricing




Figure 12.4 Fitted NDIG, normal and Student's t -distributions compared with the empirical kernel density of the log-returns with (left) linear and (right) log-scale y-axis.

By evaluating the CF at these parameter values, the resultant PDF can be computed via FFT. This PDF is compared with the empirical density of the T95 log-returns in Fig. 12.4. For comparison, the pdfs of the normal and the Student's t -distribution are also fit to the empirical density (using the standard maximum likelihood estimator for pdfs). Except for a narrow range around the modal return value, the NDIG distribution provides the best fit, with the t -distribution a close second. The normal distribution is incompatible with the heavy tails of the return distribution. The t -distribution provides a reasonable tail fit over this tail range but its power-law tail becomes too heavy beyond this range.

In section 12.3.2, the NDIG model is used to price a European call option \mathcal{C} with the underlying risky asset \mathcal{S} being the T95 portfolio. The dynamics of \mathcal{S} under Q is given by (12.20) and the CF of the log-price by (12.21). Option prices for \mathcal{C} are determined by evaluating (12.23) using the FFT for different values of strike price K and maturity time T . Put option prices are computed using the put-call parity. Section 12.3.1 discusses how to set the numerical parameters a and v_{\max} required for the option price computations.

12.3.1 Choice of a and v_{\max}

Carr and Madan (1998) introduce the parameter a to ensure that the call pricing function (12.22) is square integrable as $K \rightarrow 0$ (i.e. as $k = \ln(K) \rightarrow -\infty$). They note that a sufficient condition for square integrability is provided by the requirement that

$$\varphi_{\ln(S_T^{(Q)})}[-i(1+a)] < \infty. \quad (12.31)$$

From (12.21),

$$\begin{aligned} & \varphi_{\ln(S_T^{(Q)})}(-i(1+a)) \\ &= S_t^{1+a} \exp\{(1+a)[r_f - K_{X_1}(1)]T\} \left\{ \exp(\psi_{X_1}(-i(1+a))) \right\}^T. \end{aligned} \quad (12.32)$$

From (12.13) and (12.15) note that $\psi_{X_1}(-iw) = K_{X_1}(w)$ for $w \in \mathbb{R}$. Hence (12.31) and (12.32) can be combined as the requirement

$$\exp\{(1+a)[r_f - K_{X_1}(1)]T\} \left\{ \exp(K_{X_1}(1+a)) \right\}^T < \infty. \quad (12.33)$$

As the cumulant generating function should remain real-valued, requirement (12.33) reduces to positivity requirements for the arguments of the square root evaluations in (12.16) for $K_{X_1}(1+a)$, $a \in [0, a_{\max})$. From (12.15), assuming $\gamma = 0$, this reduces to the requirements

$$h(w) = 1 - \frac{\sigma^2}{\lambda_{\mathcal{T}}} w^2 - 2 \frac{\rho}{\lambda_{\mathcal{T}}} w \geq 0 \text{ and } g(w) = 1 - 2 \frac{\lambda_{\mathcal{T}}}{\lambda_U} (1 - \sqrt{h(w)}) \geq 0. \quad (12.34)$$

Solving the second of these for $h(w)$ gives

$$h(w) \geq \left(1 - \frac{\lambda_U}{2\lambda_{\mathcal{T}}}\right)^2 \stackrel{\text{def}}{=} d^2. \quad (12.35)$$

Combining (12.35) with the first equation in (12.34) gives

$$1 - \frac{\sigma^2}{\lambda_{\mathcal{T}}} w^2 - 2 \frac{\rho}{\lambda_{\mathcal{T}}} w \geq d^2. \quad (12.36)$$

From the roots of this quadratic, (12.36) is satisfied for $w = 1 + a$ when

$$-\frac{\rho}{\sigma^2} - \frac{\sqrt{\rho^2 + \lambda_{\mathcal{T}}\sigma^2(1-d^2)}}{\sigma^2} \leq 1 + a \leq -\frac{\rho}{\sigma^2} + \frac{\sqrt{\rho^2 + \lambda_{\mathcal{T}}\sigma^2(1-d^2)}}{\sigma^2}. \quad (12.37)$$

Since $a > 0$, we must have

$$a \leq a_{\max} = -1 - \frac{\rho}{\sigma^2} + \frac{\sqrt{\rho^2 + \lambda_{\mathcal{T}}\sigma^2(1-d^2)}}{\sigma^2}. \quad (12.38)$$

Rewriting (12.38) using (12.35) gives

$$a_{\max} = \frac{1}{\sigma^2} \sqrt{\rho^2 + \lambda_U \left(1 - \frac{\lambda_U}{4\lambda_{\mathcal{T}}}\right) \sigma^2} - \frac{\rho}{\sigma^2} - 1. \quad (12.39)$$

Using the global parameter values given in Table 12.1, this produces a bound of $a_{\max} \lesssim 23.2$. Fig. 12.5 shows values of a_{\max} computed over the time-period 12/19/2011 through 12/18/2020 using the parameter values plotted in Fig. 12.2. The smallest values of $\lesssim 20$ for a_{\max} occur during the end of the great recession. They are also small during the pandemic.




Figure 12.5 Values for the upper bound a_{\max} computed over the period 12/19/2011 through 12/18/2020.

A stricter determination for the value of $a \in [0, a_{\max})$ is based on the requirement that option prices remain within allowed bounds. For a European call option, the upper and lower bounds on the price are given by (See, e.g., Hull, 2018)

12. Option Pricing

$$\max(S_t - Ke^{-rfT}, 0) \leq C(S_t, T, K) \leq S_t. \quad (12.40)$$

From put–call parity,

$$C(S_t, T, K) + Ke^{-rfT} = P(S_t, T, K) + S_t,$$

the bounds on the price $P(S_t, T, K)$ of a European put option are

$$\max(0, Ke^{-rfT} - S_t) \leq P(S_t, T, K) \leq Ke^{-rfT}. \quad (12.41)$$

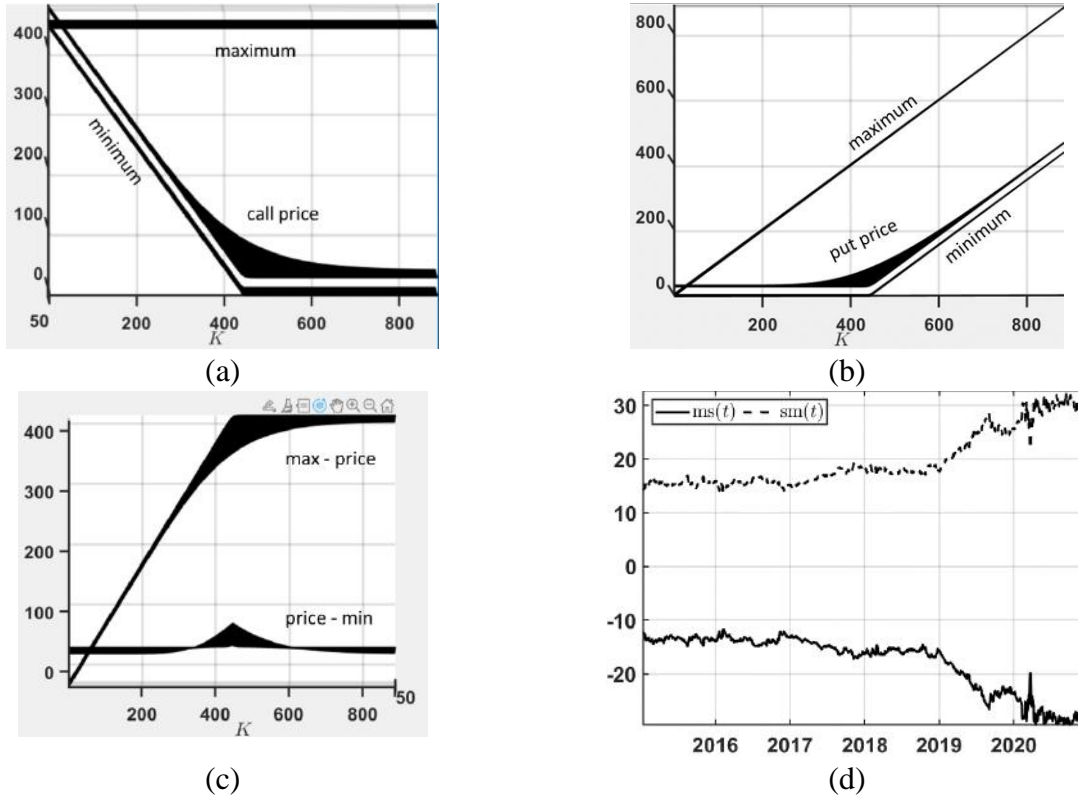


Figure 12.6 The maximum, minimum, and price surfaces computed for (a) call and (b) put options from (12.40) and (12.41). (c) The maximum – price and price – minimum difference surfaces computed from (12.42). (d) The minimum difference time series $ms(t)$ and $sm(t)$ computed from (12.43) over the period 12/19/2011 through 12/18/2020.

Fig. 12.6 shows call and put price surfaces, as well as the maximum and minimum threshold surfaces, computed for a value of $a = 0.2$ using the value of the T95 portfolio on 12/18/2020 and the NDIG parameters from Table 12.1. While these are three-dimensional plots of the surfaces as a function of K and T , they have been projected into the price– K plane to show most clearly how the price surfaces approach the minimum and maximum surfaces. In this projection, it is clear that the price surfaces exceed the maximum threshold for low values of K . As the put option prices in this numerical example are computed using put–call parity, it is unnecessary to examine put and call prices separately; the difference surfaces

$$S_t - C(S_t, T, K) \text{ and } C(S_t, T, K) - \max(S_t - Ke^{-rfT}, 0), \quad (12.42)$$

are respectively identical to the difference surfaces

$$Ke^{-rfT} - P(S_t, T, K) \text{ and } P(S_t, T, K) - \max(0, Ke^{-rfT} - S_t).$$

Fig. 12.6(c) summarizes the results of Figs. 12.6(a) and (b) by showing the difference surfaces from (12.42). The “maximum – price” surface, $S_t - C(S_t, T, K)$, has negative values for the lowest strike price values. Fig 12.6(d) plots the values

$$\begin{aligned} ms(t) &= \min[S_t - C(S_t, T, K)], \\ sm(t) &= \min[C(S_t, T, K) - \max(S_t - Ke^{-rfT}, 0)], \end{aligned} \quad (12.43)$$

obtained from the same computations that produced the results shown in Fig. 12.5. Using the value $a = 0.2$ results in price surfaces that exceed maximum price threshold every day over this period.

Further computations show that choosing a too small results in option prices that exceed the maximum threshold, while choosing a too large produces prices that go below the minimum threshold (and may be negative). It can also be deduced from Fig. 12.6(d) that the range of threshold-obeying a values changes with the price of the underlying. For the computations in this section, which use, as the underlying asset, the T95 portfolio having prices given by Fig. 12.1, computations of option prices for the period 12/19/2011 through 12/18/2020 reveal that it is sufficient to uniformly bound a in the range

$$a \in [0.40, 0.90] \quad (12.44)$$

to ensure that threshold prices are never exceeded. For the computations in the remainder of this section, we utilize the smallest possible value, $a = 0.40$.

To address the question of the error produced by truncating the integration in (12.23) to the range $v \in [0, v_{\max}]$, again two considerations come into play. From (12.28), we see that by using an FFT to perform the integration in (12.23), the value of v_{\max} determines the range of strike prices $K = e^k$ examined. In practice, the range of strike prices must cover at least the range $(0, 2S_t)$, where S_t is the price of the underlying on the date t for which option prices are to be computed. Under the requirement $\bar{k} \geq \ln(2S_t)$, from (12.28) this necessitates

$$v_{\max}(t) \leq \frac{(N-1)^2}{N} \frac{\pi}{\ln(2S_t)} \stackrel{\text{def}}{=} \bar{v}(t). \quad (12.45)$$

As the FFT requires an equally spaced k -grid, the spacing between strike prices K grows exponentially, reducing coverage for large values of K . Choice of $v_{\max}(t)$ involves a compromise between the upper bound (12.45) and a sufficiently large value needed to control the truncation error produced by approximating (12.23) by (12.25). From (12.23), this truncation error is

$$\begin{aligned} \varepsilon(S_t, T, k) &= \frac{e^{-rfT-ak}}{\pi} \int_{v_{\max}(t)}^{\infty} e^{-ivk} \frac{\varphi_{\ln(S_t^{(Q)})}[-i(1+a+iv)]}{(a+iv)(1+a+iv)} dv \\ &\equiv \frac{e^{-rfT-ak}}{\pi} \int_{v_{\max}(t)}^{\infty} g(v) dv, \end{aligned}$$

which can be bounded by

$$\varepsilon(S_t, T, k) \leq \frac{e^{-rfT-ak}}{\pi} \int_{v_{\max}(t)}^{\infty} |g(v)| dv, \quad (12.46)$$

12. Option Pricing

where

$$|g(v)| = \frac{|\varphi_{\ln(S_T^{(Q)})}[-i(\bar{a} + iv)]|}{\sqrt{(a^2 + v^2)(\bar{a}^2 + v^2)}}, \quad \bar{a} \equiv 1 + a. \quad (12.47)$$

From (12.21)

$$\varphi_{\ln(S_T^{(Q)})}(-iw) = S_t^w \exp(wT \bar{r} + TK_{X_1}(w)), \quad (12.48)$$

with $w \equiv \bar{a} + iv$ and $\bar{r} \equiv r_f - K_{X_1}(1) \in R$. Starting from (12.48) and (12.16), a few pages of computation (assuming $\gamma = 0$) lead to the result

$$\begin{aligned} |\varphi_{\ln(S_T^{(Q)})}(-iw)| &= (S_t)^{\bar{a}} \exp\{T[\bar{a}(\bar{r} + \mu) + \lambda_U - h(v)]\}, \\ h(v) &= \lambda_U(c_1^2(v) + c_2^2(v))^{1/4} \cos\left(\frac{\theta_c(v)}{2}\right). \end{aligned} \quad (12.49)$$

The terms appearing in $h(v)$ are defined in Appendix B of this chapter. Noting that $\exp[-Th(v)] < 1$, we have the bound

$$|\varphi_{\ln(S_T^{(Q)})}(-iw)| < (S_t)^{\bar{a}} \exp\{T[\bar{a}(\bar{r} + \mu) + \lambda_U]\} \equiv A,$$

and, for sufficiently large values of $v_{\max}(t)$, $|g(v)| < A/v^2$. Thus, the truncation error is bounded by¹⁰⁷

$$\varepsilon(S_t, T, k) \leq \frac{e^{-r_f T - ak}}{\pi} \frac{A}{v_{\max}(t)}. \quad (12.50)$$

Imposing $h(v) < 1$ is a rather crude bound. With $v_{\max}(t)$ sufficiently large and using the large v approximation derived in equation (12.B) of Appendix B, we have $h(v) \approx \lambda_U \sqrt{c_1}$ giving

$$\exp\{-T\lambda_U \sqrt{c_1}\} < \exp\{-\sqrt{2}T\lambda_J^{1/4}\lambda_U^{1/2}\sigma^{1/2}[v_{\max}(t)]^{1/2}\}, \quad (12.51)$$

for $v > v_{\max}(t)$. Under this improved bound (12.46) becomes

$$\begin{aligned} \varepsilon(S_t, T, k) &\leq \frac{e^{-r_f T - ak}}{\pi} \frac{(S_t)^{\bar{a}} \exp\{T[\bar{a}(\bar{r} + \mu) + \lambda_U - \sqrt{2}\lambda_J^{1/4}\lambda_U^{1/2}\sigma^{1/2}[v_{\max}(t)]^{1/2}]\}}{v_{\max}(t)}. \end{aligned} \quad (12.52)$$

For sufficiently large values of $v_{\max}(t)$, the term $\exp\{-T\sqrt{2}\lambda_J^{1/4}\lambda_U^{1/2}\sigma^{1/2}[v_{\max}(t)]^{1/2}\}$ dominates the numerator of (12.52) and the truncation error decreases exponentially with maturity time T .

From (12.45), the truncation error is maximally reduced by setting $v_{\max}(t) = \bar{v}(t)$. As $v_{\max}(t)$ varies with day t (that is, with S_t), from (12.28) we note that \bar{k} (which sets the range of strike prices) varies with t . For simplicity, the computations for this empirical example are based upon the use of the constant value $\bar{v} = \min_t[v_{\max}(t)]$. Based on the range of underlying asset prices over this period (Fig. 12.1), $\bar{v} = 465.65$. The value $\bar{v} = 465$ was used in the computations. Fig.

¹⁰⁷ This result using the bound $h(v) < 1$ is presented in Carr and Madan (1998).

12.7 shows the daily bound on the truncation error $\varepsilon(S_t, 1, k)$ for the smallest maturity time, $T = 1$ day. The bound for $T > 1$ can be reasonably approximated by $\varepsilon(S_t, 1, k)^T$. During asset price down-turns, the truncation error increases. The daily truncation bound can be improved by using a time-varying $v_{\max}(t)$ rather than the constant value \bar{v} used in these calculations.




Figure 12.7 Maximum truncation error $\varepsilon(S_t, 1, k)$ in the call option price computation.

12.3.2 Option Price and Implied Volatility Surfaces

Fig. 12.8 shows example call and put price surfaces, $C(S_t, T, K)$ and $P(S_t, T, K)$, computed on 12/18/2020 ($S_t = \$443.95$) using the parameter values from Table 12.1. Note that for close-to-the-money values ($K \approx S_t$), the call and put option prices increase with T reflecting increased future uncertainty. Fig. 12.8 also plots the implied volatility surface computed from (12.8) as a function of T and moneyness, $M = K/S_0$. Since put option prices were computed from call option prices using put–call parity, the implied volatility surface is the same for both call and put options. As is typically observed, at constant values of T the implied volatility (future uncertainty) increases as strike prices move away from the value S_t (the volatility “smile”). At constant values of M , implied volatility decreases as time to maturity increases.




Figure 12.8 (a) Call and (b) put prices for European options with the T95 portfolio as the underlying risky asset. (c) The implied volatility surface is the same for both call and put options.

12.4 Volatility Measures

As illustrated in Fig. 12.1, REIT portfolio returns (and indeed the entire market) undergo periods of significant volatility. As a result, quantifying volatility and providing estimates of future

12. Option Pricing

volatility are critical in portfolio management and financial markets. As option pricing provides (practical) estimates of the current value of future prices, determining volatility based upon option prices provides one method for estimating future volatility. Indeed, this is the impetus behind implied volatility computations using the Black–Scholes–Merton model. This section develops a volatility measure based upon NDIG option pricing and compares it to two existing volatility measures: historical volatility and one based upon the Chicago Board Options Exchange (Cboe) VIX formulation. The three measures will be compared based upon their performance on historical data. The historical volatility measure will be used as a proxy for the true volatility, to which results from the other two measures will be compared.

Historical volatility. A time series of historical volatility is defined as the standard deviation of log-returns computed in a moving window of fixed time-length. Fig. 12.9 shows the historical volatility (in percent) using the log-return data of Fig. 12.1 for the T95 portfolio and a moving window of 4 years. Historical daily volatility over this period varied from a low of 1.0% (Q3–Q4, 2015) to a high of 2.9% (Q4, 2012).




Figure 12.9 Comparison of (left) unnormalized (in percent) and (right) normalized volatility measures for the log-returns for the T95 portfolio over the period 12/19/2011 through 12/18/2020. Note the log-scale y-axis in the left figure. The STD and NDIG plots are indistinguishable.

Implied volatility measured under Q . By their very nature, option prices attempt to capture future uncertainty. Deducing an implied volatility from option prices provides a proxy measure for future volatility. The VIX was the first practical index developed to reflect volatility expectations for the S&P 500 stock index (SPX). The value of the VIX index is

$$\text{VIX} = 100 \sqrt{W_1 \sigma_1^2 + W_2 \sigma_2^2}. \quad (12.53)$$

Subscript “1” denotes near-term options while “2” denotes next-term options. As SPX options expire on Fridays, depending on the trading day, near-term options have 23 to 30 days until expiration, while next-term options have 31 to 37 days until expiration. The weights W_j , $j = 1, 2$ express these expiration times in a normalized form, accurate to the minute,

$$W_1 = \frac{M_{T_1}}{M_{30}} \left(\frac{M_{T_2} - M_{30}}{M_{T_2} - M_{T_1}} \right), \quad W_2 = \frac{M_{T_2}}{M_{30}} \left(\frac{M_{30} - M_{T_1}}{M_{T_2} - M_{T_1}} \right). \quad (12.54)$$

Here M_{T_j} is the number of minutes to settlement of term j options while M_{30} denotes the number of minutes in 30 days. Hence $0 \leq W_1 \leq 1$, $0 \leq W_2 \leq 1$, and $W_1 + W_2 = 1$. The values of the near- and next-term volatilities are given by (Demeterfi et al., 1999)

$$\sigma_j^2 = \frac{2e^{r_f T_j}}{T_j} \sum_i \frac{\Delta K_i}{K_i^2} Q(K_i) - \frac{1}{T_j} \left(\frac{F_j}{K_{0j}} - 1 \right)^2, \quad (12.55)$$

where: $T_j = M_{T_j}/M_{365}$ is the time to expiration measured in fractions of a year; r_f is the annualized risk-free rate; F_j is the forward index level derived from index option prices; K_{0j} is the first strike price below F_j ; K_i is the strike price of the i^{th} out-of-the-money option (a call if $K_i > K_0$, a put if $K_i < K_0$, both call and put if $K_i = K_0$); $\Delta K_i = (K_{i+1} - K_{i-1})/2$ is the interval between strike prices; and $Q(K_i)$ is the midpoint of the bid–ask spread for each option with strike price K_i . Details on the computation of each of these parameters can be found in the Cboe white paper (Cboe Exchange, Inc., 2019).

This methodology can be applied to calculate a “VIX-like” volatility, which we refer to as REVIX, for the T95 portfolio using the NDIG model to compute prices for European put and call options having between 23 and 37 days to expiration. To compare with the historical volatility, REVIX values were calculated using historical data as follows. For each moving window of length 4 years:

1. Fit the NDIG model to the T95 log-return data and estimate the model parameters using the minimization (12.30). This results in the parameter values displayed in Fig. 12.2.
2. Determine the dynamics of the T95 price S_t^Q in the risk-neutral world and, subsequently, the CF of the log price $\ln S_t^Q$ using (12.21).
3. Using the portfolio spot price S_t for the last date of the moving window, call option prices with appropriate expiration dates are computed using the FFT formulation (12.25) to (12.28). Put option prices are computed using put–call parity.
4. REVIX values are then computed using (12.53) through (21.58). As there are no traded options on our T95 portfolio, the following minor modifications were made to the VIX formulation.
 - a. The risk-free interest rate used was the annualized bond-equivalent (coupon-equivalent) yield for 4-week U.S. treasury bills¹⁰⁸ published for day t .
 - b. Closing times for option evaluation are 4:00 PM on day t .
 - c. Options expire on near- and next-term Fridays at 4:00 PM.
 - d. As there is no bid–ask spread in the NDIG option price computations, $Q(K_i)$ is computed directly as the NDIG option price.
 - e. As NDIG computed option prices do not go to zero, the range of strike prices considered in the REVIX computation was $0.75 S_t < K < 1.5 S_t$.

As noted in point 4d, as our underlying asset portfolio is not traded, no market sentiment is setting bid–ask prices for these call and put options, and the REVIX volatility directly reflects the NDIG option price computations. There is therefore no risk premium (market price of risk) that is traditionally captured by the VIX.

Fig. 12.9 (left) shows the results of the REVIX volatility computation. At first sight, the larger REVIX volatilities appear inconsistent with the historical volatilities; REVIX values vary from a

¹⁰⁸ <https://www.treasury.gov/resource-center/data-chart-center/interest-rates/Pages/TextView.aspx?data=billrates>

12. Option Pricing

minimum of 30.6% (Q3–Q4, 2015) to a maximum of 52.7% (Q1, 2012), and are 19 to 30 times larger than the historical volatility. This inconsistency is ultimately related to the fact that the VIX formulation is defined for use with SPX. The scale of the VIX value is set by the scale of the variances σ_1^2 and σ_2^2 in (12.53), which are defined in (12.55). We note without further comment that, in practice, the first of the two terms in (12.55) dominates. This first term approximates the integral

$$\sigma^2 = \frac{2e^{r_f T}}{T} \int_{K_L}^{K_H} \frac{Q(K)}{K^2} dK \quad (12.56)$$

for the appropriate near- or next-term options. The qualitative form of $Q(K)$ can be approximated by the piecewise linear function

$$Q(K) = \begin{cases} \Delta Q \frac{K - K_L}{K_0 - K_L} + Q_L, & K_L < K \leq K_0, \\ \Delta Q \frac{K_H - K}{K_H - K_0} + Q_L, & K_0 \leq K < K_H, \end{cases} \quad (12.57)$$

where $\Delta Q = Q_H - Q_L$ characterizes the range of out-of-the-money call and put prices. Using (12.57), the evaluation of integral in (12.56) can be written

$$\int_{K_L}^{K_H} \frac{Q(K)}{K^2} dK = \frac{Q_H}{K_L} \left\{ (1-b) \left[\frac{\ln(x)}{x-1} + \frac{c \ln(cx)}{1-cx} \right] + b(1-c) \right\} \equiv \frac{Q_H}{K_L} g(c, x; b), \quad (12.58)$$

$$b \equiv \frac{Q_L}{Q_H}, \quad 0 < b < 1; \quad c \equiv \frac{K_L}{K_H}, \quad 0 < c < 1; \quad x \equiv \frac{K_0}{K_L}, \quad 1 < x < \frac{1}{c}.$$

The surface $g(c, x; b)$ parameterized by b is a convex surface with the limits $0 < g(c, x; b) < 1$ for all $b \in (0, 1)$. Three examples of this surface are shown in Fig. 12.10 spanning the range of b .

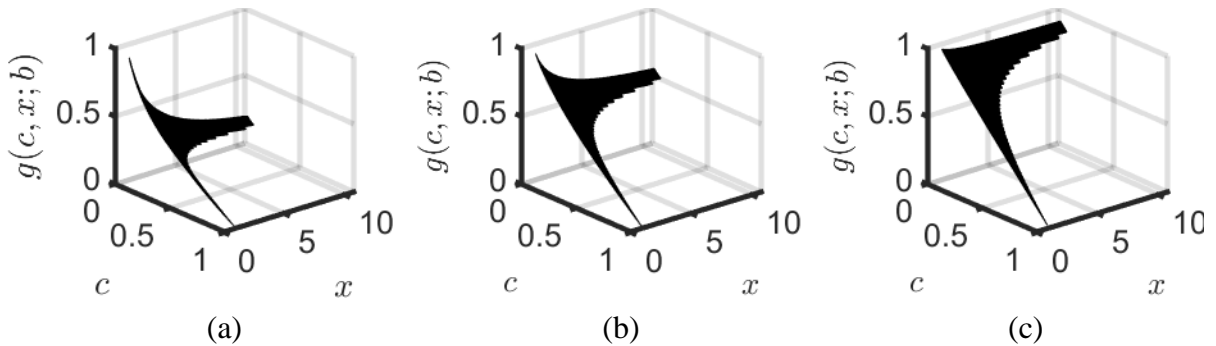


Figure 12.10 The surface $g(c, \bar{K}; b)$ for (a) $b = 0.01$, (b) $b = 0.46$ and (c) $b = 0.97$.

Thus, the scale of the integral (12.58) is set by the ratio Q_H/K_L , and hence the range of the variance (12.56) is

$$0 < \sigma^2 < \frac{2e^{r_f T}}{T} \frac{Q_H}{K_L} = 2e^{r_f T} \frac{M_{365}}{M_T} \frac{Q_H}{K_L} \approx 24 \frac{Q_H}{K_L}. \quad (12.59)$$

There is no guarantee that the upper bound in (12.59) is less than unity! As noted in point 4e above, evaluation of the integral (12.58) was restricted to the range $K_L = 0.75 S_t$, $K_H = 1.5 S_t$. The

resultant surface $Q_1(K, t)$ computed from the near-term options is shown in Fig. 12.11. (The surface for the next-term options is very similar, with slightly larger values for $Q_2(K, t)$.) The accuracy of the approximation (12.57) can be judged by viewing this surface along constant values of t .




Figure 12.11 The surface $Q_1(K, t)$ computed from the near-term options over the period 12/19/2011 through 12/18/2020. The surface is truncated to lie within the range $K \in [0.75 S_t, 1.5 S_t]$.

To adjust for the different scales, the historical and REVIX volatility time series can be separately normalized (by subtracting the series mean and dividing by the series standard deviation). These normalized plots, also shown in Fig. 12.9, exhibit close agreement, with the REVIX volatility exhibiting a small daily fluctuation that is absent in the historical volatility.¹⁰⁹ The results of Fig. 12.9 underscore the fact that the only source of variance in this numerical example is that of the daily returns of the underlying risky asset, the T95 portfolio. This variance is directly quantified by the historical volatility. As all option prices in this example are computed using the Carr–Madan formulation and the NDIG model, no additional volatility is added to these option prices through, for example, option trader sentiment (bid–ask pricing). Thus, the REVIX volatility formulation encapsulates only the original asset return volatility.

NDIG volatility under P . The double subordinated method leads naturally to a new measure of volatility under the measure P . From (12.18) with $\gamma = 0$, the variance of the unit increment $\text{Var}[X_1]$ is

$$\text{Var}(X_1) = \sigma^2 + \rho^2 \left(\frac{1}{\lambda_T} + \frac{1}{\lambda_U} \right). \quad (12.60)$$

The volatility $\sqrt{\text{Var}[X_1]}$, which we refer to as the NDIG volatility, reflects the Brownian motion (Gaussian) and Lévy subordinator components of the model. The NDIG volatility is also shown in

¹⁰⁹ Fig. 12.9 provides evidence that the VIX methodology, used carefully, does reproduce historical (in-sample) volatility. In practice and as intended, with actively traded options the VIX computation will also capture additional volatility induced by option trader sentiment.

Fig. 12.9. The unnormalized (and hence the normalized) NDIG volatility tracks the historical volatility veery accurately.

Appendix A

Although (12.13) implies that the model (12.10) has eight parameters, $\mu, \gamma, \rho, \sigma, \mu_U, \lambda_U, \mu_T, \lambda_T$, only six are identifiable. Using the relationship (12.15) between the CF and the MGF, from (12.13) $K_{X_1}(w)$ can be written as

$$\begin{aligned} K_{X_1}(w) &= \mu w + c_K \left(1 - \sqrt{g(w)}\right), \\ g(w) &= 1 - 2c_g \left(1 - \sqrt{h(w)}\right) - 2d_\gamma w, \\ h(w) &= 1 - 2d_\rho w - d_\sigma w^2, \end{aligned} \quad (12.A1)$$

where $c_K = \lambda_U/\mu_U$, $c_g = \mu_U^2 \lambda_T/(\lambda_U \mu_T)$, $d_\gamma = \gamma \mu_U^2/\lambda_U$, $d_\rho = \rho \mu_T^2/\lambda_T$, and $d_\sigma = \sigma^2 \mu_T^2/\lambda_T$. From (12.A1) we have the following derivatives with respect to w :

$$\begin{aligned} h'(w) &= -2d_\rho - 2d_\sigma w, \\ h''(w) &= -2d_\sigma, \\ h^{(p)}(w) &= 0, \quad p = 3, 4, \dots; \end{aligned} \quad (12.A2)$$

$$\begin{aligned} g'(w) &= c_g [h^{-1/2} h'] - 2d_\gamma, \\ g''(w) &= c_g \left[-\frac{1}{2} h^{-3/2} (h')^2 + h^{-1/2} h'' \right], \\ g'''(w) &= c_g \left[\frac{3}{4} h^{-5/2} (h')^3 - \frac{3}{2} h^{-3/2} h' h'' + h^{-1/2} h''' \right], \\ g^{(4)}(w) &= c_g \left[-\frac{15}{8} h^{-7/2} (h')^4 + \frac{9}{2} h^{-5/2} (h')^2 h'' - \frac{1}{2} h^{-3/2} [3(h'')^2 + 4h' h'''] \right. \\ &\quad \left. + h^{-1/2} h^{(4)} \right]; \end{aligned} \quad (12.A3)$$

and

$$\begin{aligned} K_{X_1}'(w) &= \mu - \frac{c_K}{2} [g^{-1/2} g'], \\ K_{X_1}''(w) &= -\frac{c_K}{2} \left[-\frac{1}{2} g^{-3/2} (g')^2 + g^{-1/2} g'' \right], \\ K_{X_1}'''(w) &= -\frac{c_K}{2} \left[\frac{3}{4} g^{-5/2} (g')^3 - \frac{3}{2} g^{-3/2} g' g'' + g^{-1/2} g''' \right], \\ K_{X_1}^{(4)}(w) &= -\frac{c_K}{2} \left[-\frac{15}{8} g^{-7/2} (g')^4 + \frac{9}{2} g^{-5/2} (g')^2 g'' \right. \\ &\quad \left. - \frac{1}{2} g^{-3/2} [3(g'')^2 + 4g' g'''] + g^{-1/2} g^{(4)} \right]. \end{aligned} \quad (12.A4)$$

Noting that $h(0) = g(0) = 1$, the dependence of the derivatives in (12.A2) through (12.A4), when evaluated at $w = 0$, can be summarized as follows:

$$\begin{aligned} h'(0) &= h'(d_\rho), \\ h''(0) &= h''(d_\sigma), \\ h^{(p)}(w) &= 0, \quad p = 3, 4, \dots; \end{aligned} \quad (12.A5)$$

$$\begin{aligned} g'(0) &= g'(c_g, h'(0), d_\gamma) = g'(c_g, d_\rho, d_\gamma), \\ g^{(p)}(0) &= g^{(p)}(c_g, h'(0), h''(0)) = g^{(p)}(c_g, d_\rho, d_\sigma), p = 2, 3, \dots; \end{aligned} \quad (12.A6)$$

and

$$\begin{aligned} K_{X_1}'(0) &= K_{X_1}'(\mu, c_K, g'(0)) = K_{X_1}'(\mu, c_K, c_g, d_\rho, d_\gamma), \\ K_{X_1}^{(p)}(0) &= K_{X_1}^{(p)}(c_K, g'(0), g''(0), \dots, g^{(p)}(0)) = K_{X_1}^{(p)}(c_K, c_g, d_\rho, d_\gamma, d_\sigma), \\ p &= 2, 3, \dots. \end{aligned} \quad (12.A7)$$

Thus, all derivatives of $K_{X_1}(w)|_{w=0}$, or equivalently all moments of the process X_1 , depend only on the six identifiable parameters $\mu, c_K, c_g, d_\rho, d_\gamma, d_\sigma$.

Appendix B

Using the definitions,

$$\bar{a} = 1 + a, \quad d_\lambda = \frac{\lambda_{\mathcal{T}}}{\lambda_U}, \quad d_\rho = \frac{\varrho}{\lambda_{\mathcal{T}}}, \quad d_\sigma = \frac{\sigma^2}{\lambda_{\mathcal{T}}},$$

the terms in (12.49) are

$$\begin{aligned} c_1 &= 1 - 2d_\lambda + 2d_\lambda(b_1^2 + b_2^2)^{1/4} \cos\left(\frac{\theta_b}{2}\right), \quad c_2 = 2d_\lambda(b_1^2 + b_2^2)^{1/4} \sin\left(\frac{\theta_b}{2}\right), \\ \theta_c &= \text{atan}\left(\frac{c_2}{c_1}\right) \in \left(-\frac{\pi}{2}, \frac{\pi}{2}\right], \end{aligned}$$

where

$$\begin{aligned} b_1 &= 1 - 2d_\rho \bar{a} - d_\sigma(\bar{a}^2 - v^2), \quad b_2 = -2(d_\rho + \bar{a}d_\sigma)v, \\ \theta_b &= \text{atan}\left(\frac{b_2}{b_1}\right) \in \left(-\frac{\pi}{2}, \frac{\pi}{2}\right]. \end{aligned}$$

For sufficiently large values of v , we have the approximations:

$$\begin{aligned} b_1 &\approx d_\sigma v^2, \quad b_2 = -2(d_\rho + \bar{a}d_\sigma)v, \quad (b_1^2 + b_2^2)^{1/4} \approx d_\sigma^{1/2} v, \\ \cos\left(\frac{\theta_b}{2}\right) &\approx 1, \quad \sin\left(\frac{\theta_b}{2}\right) \approx -\frac{d_\rho + \bar{a}d_\sigma}{d_\sigma v}, \\ c_1 &\approx 2d_\lambda \sqrt{d_\sigma} v, \quad c_2 \approx -2d_\lambda \frac{d_\rho + \bar{a}d_\sigma}{\sqrt{d_\sigma}}, \quad (c_1^2 + c_2^2)^{1/4} \approx \sqrt{c_1}, \\ \cos\left(\frac{\theta_c}{2}\right) &\approx 1. \end{aligned} \quad (12.B)$$

References

12. Option Pricing

- Adhikari, B. K. & Hilliard, J. E. (2014). The VIX, VXO and realised volatility: a test of lagged and contemporaneous relationships. *International Journal of Financial Markets and Derivatives*, 3(3), 222–240
- Aloud, M., Tsang, E., Olsen, R. B., & Dupuis, A. (2012). A directional-change event approach for studying financial time series. *Economics: The Open-Access, Open-Assessment E-Journal*, 6(2012-36), 1–17.
- Black, F. & Scholes, M. (1973). The pricing of options and corporate liabilities. *The Journal of Political Economy*, 81(3), 637–654.
- Bochner, S. (1995). *Harmonic analysis and the theory of probability*. University of California Press, Berkeley, CA.
- Boyle, P. P. (1977). Options: A Monte Carlo approach. *Journal of Financial Economics*. 4 (3), 323–338.
- Broadie, M. & Glasserman, P. (1996). Estimating security price derivatives using simulation. *Management Science*, 42 (2), 269–285.
- Carr, P. & Madan, D. (1998). Option valuation using the fast Fourier transform. *Journal of Computational Finance*, 2, 61–73.
- Carr, P., Geman, H., Madan, D. B., & Yor, M. (2003). Stochastic volatility for Lévy processes. *Mathematical Finance*, 13(3), 345–382.
- Carr, P. & Wu, L. (2004). Time-changed Lévy processes and option pricing. *Journal of Financial Economics*, 71(1), 113–141.
- Carriere, J. (1996). Valuation of the early-exercise price for options using simulations and nonparametric regression. *Insurance: Mathematics and Economics*. 19, 19–30.
- Cassidy, D. T., Hamp, M. J., & Ouyed, R. (2010). Pricing European options with a log Student's *t*-distribution: A Gosset formula. *Physica A: Statistical Mechanics and its Applications*, 389(24), 5736–5748.
- Clark, P. (1973). A subordinated stochastic process model with fixed variance for speculative prices. *Econometrica*, 41, 135–156.
- Cox, J. C., Ross, S. A., & Rubinstein, M. (1979). Option pricing: A simplified approach. *Journal of Financial Economics*, 7(3), 229–263.
- Delbaen, F., & Schachermayer, W. (1994). A general version of the fundamental theorem of asset pricing. *Mathematische Annalen*, 300(1), 463–520.
- Demeterfi, K., Derman, E., Kamal, M., & Zou, J. (1999). More than you ever wanted to know about volatility swaps. *Quantitative Strategies Research Notes*. Goldman, Sachs & Co., New York.
- Duffie, D. (2001). *Dynamic asset pricing theory* (3rd ed.). Princeton University Press, Princeton.
- Fallagoul, H. A. & Nam, K. (2020). Time-changed Lévy processes and option pricing: a critical comment. *SSRN Electronic Journal*, <https://ssrn.com/abstract=3226748>.
- Geman, H., Madan, D., & Yor, M. (2001). Time changes for Lévy processes. *Mathematics of Finance*, 11, 79–96.
- Gerber, H. U. & Shiu, E. S. W. (1994). Option pricing by the Esscher transform. *Transactions of Society of Actuaries*, 46, 99–191.

- Goldstein, D. G. & Taleb, N. N. (2007). We Don't Quite Know What We are Talking About When We Talk About Volatility. *Journal of Portfolio Management*, 33(4), 84–86.
- Guillaume, D. M., Dacorogua, M. M., Davé, R. R., Müller, U. A., Olsen, R. B., & Pictet, O. V. (1997). From the bird's eye to the microscope: A survey of new stylized facts of the intra-daily foreign exchange markets. *Finance and Stochastics*, 1, 95–129.
- Heston, S. L. (1993). A closed-form solution for options with stochastic volatility with applications to bond and currency options. *The Review of Financial Studies*, 6(2), 327–343.
- Hull, J. C. (2018). *Options, futures and other derivatives* (10th ed.). Pearson, London.
- Hull, J. C. & White, A. (1987). The pricing of options on assets with stochastic volatilities. *The Journal of Finance*, 42(2), 281–300.
- Lundtofte, F. & Wilhelmsson, A. (2013). Risk premia: Exact solutions vs. log-linear approximations. *Journal of Banking & Finance*, 37, 4256–4264.
- Mandelbrot, B. B. (1963). The variation of certain speculative prices. *The Journal of Business*, 36(4), 394–419.
- Mandelbrot, B. B. & Taylor, H. M. (1967). On the distribution of stock price differences. *Journal of Operations Research*, 15, 1057–1062.
- Merton, R. C. (1973). Theory of rational option pricing. *The Bell Journal of Economics and Management Science*, 4(1), 141–183.
- Nögel U. (2004). Option Pricing Using Stochastic Volatility Models. In A. Buikis, R. Čiegis, A. D. Fitt (Eds.) *Progress in industrial mathematics at ECMI 2002*, The European Consortium for Mathematics in Industry, vol 5, pp. 221–225. Springer, Berlin, Heidelberg.
- Paulson, A. S., Holcomb, E. W., & Leitch, R. A. (1975). The estimation of the parameters of the stable laws. *Biometrika*, 62, 163–170.
- Petrov, V., Golub, A., & Olsen, R. (2019). Instantaneous volatility seasonality of high-frequency markets in directional-change intrinsic time. *Journal of Risk and Financial Management*, 12(2), 1–31.
- Rebonato, R. (1999). *Volatility and correlation in the pricing of equity, FX and interest-rate options*. Wiley & Sons, Chichester.
- Sato, K.-I. (1999). *Lévy processes and infinitely divisible distributions*. Cambridge Studies in Advanced Mathematics, Cambridge.
- Schoutens, W. (2003). *Lévy processes in finance: Pricing financial derivatives*. John Wiley & Sons, Hoboken, NJ.
- Shirvani, A., S. T. Rachev, S. T., & Fabozzi, F.J. (2020). Multiple subordinated modeling of asset returns: Implications for option pricing. *Econometric Reviews*, 40(3), 290–319.
- Shirvani, A., Stoyanov, S., Fabozzi, F. J., & Rachev, S. T. (2021a). Equity premium puzzle or faulty economic modelling? *Review of Quantitative Finance and Accounting*, 56, 1329–1342.
- Shirvani, A, Mitnik, S., Lindquist, W. B. & Rachev, S. T. (2021b). Bitcoin volatility and intrinsic time using double subordinated Lévy processes. arXiv:2109.15051.
- Tsang, E. (2010). Directional changes, definitions. Working Paper WP050-10, Center for Computational Finance and Economic Agents, University of Essex.

12. Option Pricing

- Yao, L., Yang, G., & Yang, X. (2016). The mean correcting martingale measures for exponential additive processes. *Applied Mathematics – A Journal of Chinese Universities*, 31, 81–88.
- Yu, J. (2004). Empirical characteristic function estimation and its applications. *Econometric Reviews*, 23, 93–123.

Chapter 13

Inclusion of ESG Ratings in Optimization

Abstract

A Bill Is Coming Due for Greener Offices (The Wall Street Journal, 5/6/2021¹¹⁰) “Energy-inefficient buildings will need to be upgraded if countries are to meet net-zero commitments. Signs that landlords can claw back the cost are mixed.”

Investors’ planned holding periods (terms of ownership) are likely to be a factor in this issue. For U.S. institutional investors, the decision to pay efficiency premiums is less clear than for European investors who typically invest for longer terms. More rigorous analytics are needed to reduce or eliminate collinearity of new building premiums and that for energy efficiency. (S.T.C. 5/23/2021)

There is substantial developer/investor uncertainty relative to the cost benefit analysis of energy efficient features for all forms of real estate. Due to the competitive nature of the real estate market after the reopening of the economy, it is highly speculative and unknown if there will be sufficient rental premiums to justify the substantial initial investment related to all energy efficient features. The tax credit advantage to developers appear to be the main compelling driver of this type of investment at this moment in time. (J.H.J. 5/23/2021)

ESG Data Is All Over the Place. How to Fix It, From Two Who Work with the Numbers (Barron’s, 9/21/2021¹¹¹) “Sustainable investing keeps getting more attention—and investors and corporations more frustrated. There’s more than \$35 trillion invested in a manner that takes into account environmental, social, and governance concerns, known as ESG, according to the Global Sustainable Investment Alliance. ... How these money managers use ESG data, such as greenhouse gas emissions or workplace safety, is a scattershot process, however. There is no single framework for reporting data; there is no one agency that collects or audits it; there is no requirement as to how companies are to collect and report ESG data if at all.”

The ESG Reporting Endgame: Lessons from Human Capital Reporting (Forbes, 9/21/2021¹¹²) “As CFOs prepare new disclosures related to environmental, social and governance (ESG) issues—either voluntarily or due to evolving regulatory requirements—their experiences complying with the U.S. Securities and Exchange Commission’s (SEC) new human capital disclosure requirements can help.”

The Morningstar Sustainable Investing Framework: Clear terminology leads to greater understanding of an evolving area (Morningstar, 9/21/2021¹¹³) “To many investors, however, the world of sustainable investing can be a confusing mix of terms and approaches. That’s mainly

¹¹⁰ https://www.wsj.com/articles/a-bill-is-coming-due-for-greener-offices-11620385382?st=p21c2o1nd3i3o56&reflink=article_gmail_share

¹¹¹ <https://www.barrons.com/articles/esg-data-how-to-fix-51632173139>

¹¹² <https://www.forbes.com/sites/jimdeloach/2021/09/21/the-esg-reporting-endgame-lessons-from-human-capital-reporting/>

¹¹³ <https://www.morningstar.com/articles/1058990/the-morningstar-sustainable-investing-framework>

13. ESG Ratings in Optimization

because of a lack of consensus over terminology and, perhaps more significantly, because sustainable investing does not represent a single, distinct investment approach.”

Mortgage market is unprepared for climate risk, says industry report (CNBC, 9/23/2021¹¹⁴) “With numerous stakeholders in housing finance, climate change will send significant stress down a long financial line, according to a Mortgage Bankers Association report. The report said it could increase mortgage default, increase the volatility of house prices and produce significant climate migration.”

References

Artzner, P., Delbaen, F., Eber, J.-M., & Heath, D. (1999). Coherent measures of risk. *Risk*, 10, 203–228.

¹¹⁴ <https://www.cnbc.com/2021/09/23/mortgage-market-is-unprepared-for-climate-risk-says-industry-report.html>

Chapter 14

Inclusion of ESG Ratings in Option Pricing

14. ESG Ratings in Option Pricing

Odd number page holder for Chapter 14