A sparsity constrained dynamic trading framework to identify stochastic arbitrage opportunities with market index options

Benjamin Weeks

THE UNIVERSITY OF SYDNEY

School of Economics

University of Sydney

Thesis submitted in partial fulfilment of the award course requirements of the Bachelor of Economics (Honours)

Supervised by Dr Brendan K. Beare

STATEMENT OF ORIGINALITY

I hereby declare that this submission is my own work and to the best of my knowledge it contains no material previously published or written by another person. Nor does it contain any material which has been accepted for the award of any other degree or diploma at the University of Sydney or at any other educational institution, except where due acknowledgement is made in this thesis. Any contributions made to the research by others with whom I have had the benefit of working at the University of Sydney is explicitly acknowledged.

I also declare that the intellectual content of this study is the product of my own work and research, except to the extent that assistance from others in the project's conception and design is acknowledged.

[redacted]

Benjamin Weeks

10 November 2025

Acknowledgements

First and foremost, I would like to thank Dr Brendan K. Beare for your supervision and invaluable guidance. Over the course of this year your constructive feedback and suggestions kept this project moving forward – I could not have done it without your support. I am also grateful for the feedback and advice from the members of the econometrics group, in particular Simon Kwok and Ye Lu following my thesis and proposal presentations. To Mum and Dad, thank you for your quiet encouragement behind every step of this degree. And finally, to my partner who was always there providing encouragement when I needed it the most, it made all the difference.

Abstract

Options markets provide investors a comprehensive array of state-contingent payoffs. Stochastic arbitrage represents the opportunity to construct negatively priced portfolios of long and short positions in options, that when combined with a passive investment in the market constructs a payoff that stochastically dominates the market. This paper formulates a mixed-integer linear program that applies sequential stochastic dominance in a dynamic, daily-rebalanced trading framework to identify stochastic arbitrage opportunities with options written on the S&P 500 index. This paper finds that a dynamic trading framework succeeds at identifying stochastic arbitrage opportunities more frequently than static trading frameworks explored in recent research, but both strategies invariably return payoffs inferior to the market. This paper shows that the application of sparsity constraints to both static and dynamic trading strategies yields improved performance.

Contents

1 Introduction

The pursuit of portfolios and trading strategies that deliver superior to market risk-adjusted returns are central to financial econometrics. Since Markowitz (1952), portfolio optimisation and financial markets have evolved significantly. Introduced to US exchanges in 1973, options provide investors the ability to construct state-contingent payoffs. A core principle in options markets is that if correctly priced, it should not be possible to make sure profits by constructing portfolios of long and short positions (Black and Scholes, 1973).

Research by Ait-Sahalia and Lo (2000), Jackwerth (2000) and Rosenberg and Engle (2002) identified evidence for non-monotonic pricing kernels in options written on the Standard and Poors 500 index (SPX). A non-monotonic pricing kernel implies the existence of a stochastic arbitrage opportunity, allowing an investor to increase expected utility by engaging in zero-cost trades (Constantinides, Jackwerth and Perrakis, 2009). This discovery commenced a series of empirical investigations to firstly validate whether this systematic mispricing does in fact, exist, and if it does, then formulate an appropriate strategy to capitalise on it.

Adjacent research suggests that these findings may be driven by incorrect specification of the conditioning set of state variables when constructing risk-neutral probability measures (Chabi-Yo et al., 2008). Recently, Post and Longarela (2021) and Beare et al. (2025) formulate linear programs to identify stochastic arbitrage opportunities with market index options. The realised performance of portfolios selected in these investigations is underwhelming, suggesting that evidence of mispricing is attributable to incorrect specification in state probabilities. However, what is yet to be explored is a dynamic trading strategy and the evaluation of optimal portfolio diversification (sparsity) to identify stochastic arbitrage opportunities. Importantly, Beare et al. (2025) define a stochastic arbitrage opportunity as a mix of long and short positions in options (a "layover portfolio") with two properties. Firstly, the layover portfolio generates positive option premium (has a negative price). And secondly, an investor who holds the market portfolio can improve their return by second-order stochastic dominance (SOSD) by also holding the layover portfolio. In this paper, the "enhanced portfolio", the market plus the layover portfolio, is required to stochastically dominate the combined market and option payoff distribution from all prior days. This is the central concept in the dynamic trading framework introduced and discussed in this paper.

Markowitz's expected return-variance (EV) rule implies the importance of portfolio diversification, a departure from prior theory that concentrated on assets with the highest discounted expected return (see Williams, 1938). There is extensive research into the optimal level of diversification in portfolio constructed with equities. Evans and Archer (1968) showed that portfolio standard deviation is a rapidly decreasing asymptotic function in portfolio size. Beyond 10 assets, the marginal cost of adding an additional asset outweighs the marginal benefit, highlighting the cost of overdiversification. This paper defines overdiversification as taking both sensible positions in mispriced options (a stochastic arbitrage opportunity), but also less rational positions driven by state probability misspecification. Ideally, sparsity is appropriately constrained such that portfolio weights are concentrated only in positions that present clear opportunity for stochastic arbitrage ensuring that realised performance is not undermined.

This paper has two main objectives. Its primary aim is to formulate and evaluate the suitability of a dynamic trading strategy to identify stochastic arbitrage opportunities with market index options. This paper finds that a dynamic trading strategy does identify a larger quantity of stochastic arbitrage opportunities, however, the realised performance of selected portfolios yields returns moderately worse than passively holding the market portfolio. This provides further evidence to suggest the apparent identification of stochastic arbitrage is a consequence of misspecification in state probabilities as opposed to systematic mispricing. The paper's secondary aim is to identify optimal portfolio sparsity through the application of cardinality constraints in an option portfolio context. By evaluating the realised performance of selected enhanced portfolios against the market return across a range of sparsity values, it is determined that limiting the distinct number of option positions between 4 and 10 yields portfolios with favourable characteristics in both a static and dynamic trading context.

In pursuit of these objectives, this paper makes the following contributions. It builds upon the literature exploring the existence of statistically and economically significant stochastic arbitrage opportunities with options written on the SPX by applying a novel dynamic trading approach. The methodological contributions include the following. Firstly, the paper formulates and applies a mixed-integer linear program (MILP) that applies sequential stochastic dominance and sparsity constraints. Two key design choices distinguish the dynamic trading framework applied in this paper to the broader literature on active portfolio management. Namely, restricting the set of available option contracts to be those with the same expiry in each month. Fixing the expiry in this way facilitates direct comparison with static strategies applied in recent empirical investigations. The second design choice is the application of sequential stochastic dominance. Which in principle, allows a theoretical investor to rebalance their portfolio in response to information arrival during the month period, such that the cumulative portfolio payoff distribution only stochastically improves, where possible. A component of the dynamic trading framework is the introduction of dynamic moneyness constraints that tighten as time-to-maturity decreases to insulate against bets on tail events, that in practice, become increasingly less probable on shorter time spans.

The rest of this paper is structured as follows. Section 2 examines the two strands of related literature: stochastic dominance and stochastic arbitrage. Section 3 discusses data. Section 4 discusses the methodology, state probability distributions and dynamic portfolio choice model. Section 5 examines the results and Section 6 concludes the paper.

2 Literature review

2.1 Stochastic dominance theory and applications

First introduced to economics by Hadar and Russell (1969), stochastic dominance provides a criterion by which a preference order can be assigned to uncertain prospects. Stochastic dominance compares the cumulative distribution functions of two (or more) random variables to establish superiority. Before its introduction, expected utility theory (see Von Neumann and Morgenstern, 1944) was a popular framework used for preference ordering, an approach contingent on the specification of a utility function. Advantageously, stochastic dominance provides a generalised framework that does not rely on the presupposition of a utility function. The key definitions from Hadar and Russell (1969), Hanoch and Levy (1969), Rothschild and Stiglitz (1970) and Whitmore (1970) can be summarised as follows:

Let X and Y be random variables with cumulative distribution functions F_X and F_Y respectively. With "$\preceq$" denoting a partial order, X is k-th order stochastically dominated by Y, $X \preceq^{(k)} Y$ if and only if:

$$F_X^{(k)}(t) \geq F_Y^{(k)}(t) \quad \text{for all } t \in \mathbb{R}$$

Holding with strict inequality for at least one t. Where for $k = 1$, $F_X^{(1)}(t) := F_X(t)$ and for $k \geq 2$, the k-th repeated interval is $F_X^{(k)}(t) := \int_{-\infty}^{t} F_X^{(k-1)}(t)$. Note that first-order stochastic dominance (FOSD) implies second-order stochastic dominance (SOSD) which implies third-order stochastic dominance (TOSD) and so on. Since its introduction to economics, stochastic dominance has had broad applications, in areas such as estimating the probability of bankruptcy (Broske and Levy, 1989) and even agricultural economics (Harris and Mapp, 1986) (see Levy (1992) for survey).

Ogryczak and Ruszczynski (2002) apply FOSD and SOSD to mean-risk models. They highlight that mean-risk models may select an efficient solution that is stochastically dominated by another feasible solution. Dentcheva and Ruszczynski (2006) develop a system of linear inequalities to enforce stochastic dominance in a finite asset context. Their empirical investigation found that if the benchmark portfolio contains more assets in order to stochastically dominate it a utility function was required that penalised diversification. Later, Luedtke (2008) provides more compact and efficient formulations for FOSD and SOSD for a generalised optimisation setting. The approach taken in this thesis is a modification of the formulation provided in Leudtke (2008), which is explored in Section 4. Roman, Mitra and Zverovich (2013) use SOSD to generate portfolios that outperform the FTSE 100, S&P 500 and Nikkei 225 indexes. An interesting observation was that SOSD models do not require the imposition of cardinality constraints and that little or no rebalancing is necessary to attain superior performance.

Post and Kopa (2017) extend SD to higher orders by refining superconvex dominance conditions introduced in Bawa et al. (1985) and taking a quadratically constrained programming approach to formulate TOSD. In doing so they capturing investor prudence; the preference for positive skewness. This yielded improved returns relative to mean-variance models, driven by greater concentration in past winner industries, generating positive skewness but sacrificing diversification. Later Fang and Post (2022) investigate fourth and fifth-order stochastic dominance. These investigations suggest benefits in the application of higher-order stochastic dominance in constructing portfolios that minimise kurtosis and amplify skewness.

Recently, Mei, Chen, Lui & Ji (2022) explore dynamic SOSD in a multi-stage asset-allocation model where the investor's entire wealth path must SOSD the return sequence of a benchmark asset. Their model utilises Monte-Carlo to simulate a large panel of return paths then use K-means clustering to make the problem tractable. This approach yielded promising results, pointing to efficiency gains and performance that beat single-stage SOSD baselines. Consigli, Dentcheva, Maggioni & Micheli (2025) apply a sequential SOSD framework to asset-liability management and solve recursively finding that SOSD constraints significantly affect the optimal investment policy. This thesis differs in that it explores index options and the approach to dynamic second-order stochastic dominance is to allow the reference variable to be time variant. However, the results of these papers underpin the applicability of SOSD in time-varying contexts.

2.2 Stochastic arbitrage and the pricing kernel puzzle

The asset pricing kernel captures investor risk and time preferences for payoffs in different states. Ait-Sahalia and Lo (2000), Jackwerth (2000) and Rosenberg and Engle (2002) identified evidence for non-monotonic pricing kernels in options written on the S&P500 index (SPX). This discovery contradicts financial theory that posits that pricing kernels should be positive, monotonically decreasing in wealth. Later work by Chabi-Yo et al. (2008) challenge these findings by arguing that misspecification of the conditioning set of state variables is responsible for the evidence of non-monotonic pricing kernels.

Constantinides, Jackwerth and Perrakis (2009) motivated by the findings of Jackwerth (2000), explore option mispricing in options written on the SPX. They report widespread violations of stochastic dominance (i.e., violations of monotonic pricing kernels) even in the presence of transaction costs and market incompleteness, meaning an investor can increase expected utility by engaging in zero-cost trades, exploiting systematic mispricing. Beare (2011) showed that in the existence of a non-monotonic pricing kernel there exists an optimal measure-preserving derivative that first-order stochastically dominates the index's payoff. Indeed, the existence of a nonmonotonic pricing kernel is equivalent to the existence of a *first-order stochastic arbitrage opportunity.* A stochastic arbitrage opportunity, in the context of the SPX, as defined in Beare et al. (2025) refers to a layover portfolio that provides (1) a positive option premium and (2) the payoff distribution of the layover portfolio combined with a unit investment in the market index k-th order stochastically dominates the payoff distribution of the market index alone. A layover portfolio refers to some combination of long and short, call and put options. In this paper we slightly modify the definition of stochastic arbitrage in Beare et al. (2025) such that the enhanced portfolio (layover portfolio + unit position in market) must second-order stochastically dominate the market **and** option payoff distribution from all prior days.

Recently, Post and Longarela (2021) build upon the linear inequalities developed by Constantinides, Jackwerth and Perrakis (2009) to construct a LP that identifies all stochastic arbitrage opportunities. They report systematic stochastic arbitrage opportunities but these are limited by market depth constraints (portfolio size is limited by maximum capital of the investor, ranging from retail investor to institutional). Beare et al. (2025) use the efficient linear formulation in Luedtke (2008) to improve Post and Longarela's (2021) LP for computing second-order stochastic arbitrage (SOSA) opportunities. Additionally, they present a MILP to compute first-order stochastic (FOSA) opportunities by modifying the formulation presented in Luedtke (2008). Interestingly, Beare et al. (2025) rarely identify stochastic arbitrage opportunities in the market for SPX options. This is particularly the case when only considering options with moderate moneyness (strike between 90 to 105 percent of current SPX value). When allowing for more extreme strikes more stochastic arbitrage opportunities are identified but the realised performance of these portfolios is worse than the market.

Beare et al. (2025) explain that this is likely a consequence of difficulty in accurately specifying state probabilities for options far out-of-the-money (OTM), leading to the algorithm to believe there is mispricing when in fact there is not.

3 Data

This research utilises SPX call and put option data from January 2004 to November 2021 (213 months[1]), sourced from the Chicago Board of Options Exchange (CBOE) DataShop. For each trading day, the options data includes the time-to-expiry, best bid and ask prices, and associated volumes at 2:45 Central Time (30 minutes before market close). These options are of the European style (Black and Scholes, 1973) preserving clean terminal-payoff valuation at maturity, a feature that is complicated by the early-exercise characteristic of American style options (Brennan-Schwartz, 1977). This fixed-exercise date helps to accommodate comparison between the dynamically rebalanced portfolio and a static investment strategy. In particular, we restrict the available options to be those with an expiry as the third Friday or Saturday in each month. Doing so ensures that there is a consistent maximum time-to-expiry of 18-21 trading days, dependent on weekends and public holidays. We also ensure that on any given day the state probability distribution is constructed for a single time-to-maturity. This prevents an observed problem where a trading day can have options available for two distinct maturity dates, meaning the used state probability distribution will not have appropriate horizon parametrisation. Figure 1 visualises this expiry-anchored daily rebalanced trading framework.

Figure 1: Dynamic trading timeline

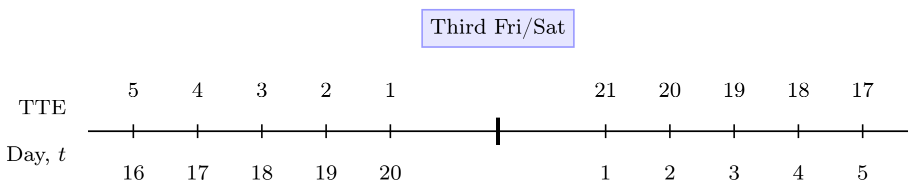

The option strike range available for analysis is restricted *ex ante* to a forward-moneyness band of **at most** 90-105% of the SPX value on trading day t. As the time-to-maturity shortens we tighten the moneyness bounds by computing the ratio

[1]The total number of months is 215 but 2 months were excluded due to data inconsistency

of the scale parameter (see Section 4.2.2) on day t to day t_1 (the first day in that month) and multiply the upper (5%) and lower (-10%) bounds by this scale factor. Figure 2 plots the average upper and lower bounds for each time-to-maturity.

Figure 2: Average dynamic moneyness bounds by horizon

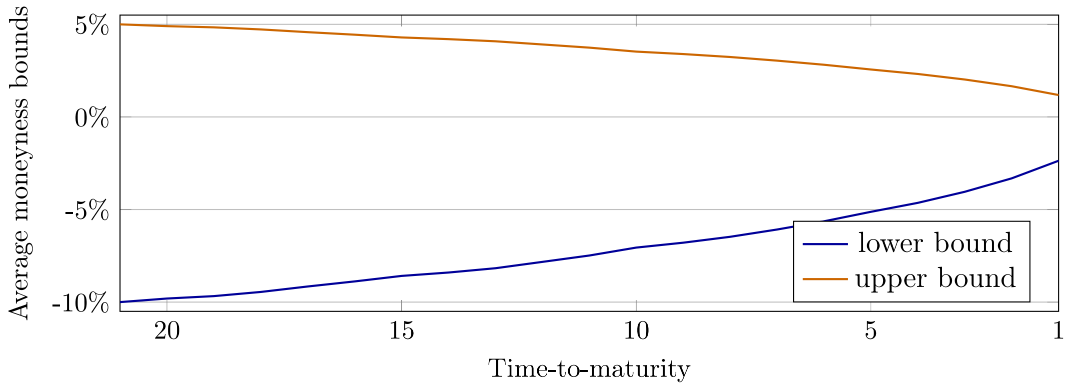

We confine our attention to moderate strikes for a few reasons. Firstly, as is shown in Figure 3, 89% of the months in our sample see monthly changes within this band, showing that during this period the SPX rarely saw month changes beyond -10% and 5%. On shorter time horizons the frequency of extreme tail events is more sparse, justifying the increasing restrictiveness as horizon length shrinks. Secondly, even when an ask price is quoted for an OTM option, the bid volume is often zero (Post and Longarela, 2021). Indeed, far-tail quotes are the first to create butterfly and vertical spread violations (giving rise to pure arbitrage opportunities) as observed in Beare et al. (2025). This occurs due to the difficulty in reliably pricing extreme states, characterised by the 'Peso problem' (see Barro and Liao (2020)). This limitation also ensures that our state probability distributions are reliable across the options book available to our theoretical investor.

Figure 3: Histogram of 1-month SPX returns (%)

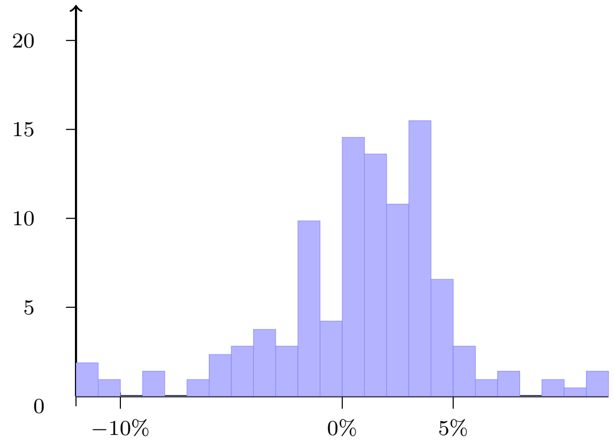

Notes: Tail bins group all months ≤ −12% and ≥ 12%; mid bins are 1% wide.

Carr and Madan (2005) define sufficient conditions to guarantee no pure arbitrage. These conditions are nonnegative cost of butterfly spreads and vertical spreads. These pure arbitrage opportunities represent standalone non-systematic mispricing and therefore do not provide evidence for nonomonotonic pricing kernels, nor the capability to consistently construct portfolios which in a probabilistic sense, outperform the market payoff. Additionally, these opportunities are both infrequent and small in size (as seen in Beare et al. (2025)), meaning that in a live-trading context, it may be impossible to identify and execute these trades before the opportunity disappears. We compute results with pure-arbitrage excluded.

Any premia generated by the layover portfolio is invested at the risk-free rate, set as the London Interbank Offered Rate (LIBOR), until month-end. Using the LIBOR as the risk-free rate and SPX dividend yield data we impute a forward price for a unit of the SPX at each fixed-expiry. LIBOR data is sourced from FRED and SPX dividend yield data is obtained from the multpl online database (www.multil.com). Daily realised variances over the horizons are computed from 5-minute SPX data sourced from LSEG Data and Analytics. The implied volatility is given by the VIX, which is sourced from Yahoo Finance.

4 Methodology

This section details the approach from market data to state probability distributions and presents the optimisation model and mixed integer linear program.

4.1 Market data

Let T_m be the fixed expiry date for month m. On any trading day t within month m, define the time-to-expiry in trading days as $\tau_t := \{\#\text{ of trading days from } t \text{ to } T_m\} \in \{21, 20, ..., 1\}$. Hence, on each day all objects are recomputed with the same calendar expiry T_m while τ_t shrinks during the month. To compare the dynamic portfolio strategy to a static trading strategy we compute for only $\max\{\tau_m\}\ \forall m$, the first day in each month. Let S_t be the SPX close on date t. The risk-free rate, r_t^f and dividend yield, d_t are used to compute an annualised forward price as defined in Hull (2021):

$$F_{t,T_m} = S_t \cdot e^{(r_t^f - d_t)\cdot\tau_t/252}$$

The τ-day log annualised market return, $r_{t,\tau} := \log(S_{t,\tau}/S_t)\cdot 252/\tau$ is used to compute the excess-market return $r_{t,\tau}^x := r_{t,\tau} - r_t^f \cdot \tau/21$. Here we scale the risk-free rate by the ratio of time-to-expiry and 21 trading days. This approach is taken to preserve consistency with the static investment strategy as the risk-free rate is a monthly measure. To compute an adjusted realised volatility measure we follow the approach taken in Bollerslev & Todorov (2023). First, using the 5-minute SPX data, let $P_{t,j}$ be the SPX level at the j-th 5-minute interval on day t, $j = 0, 1, ..., M$. Hence, the intraday log return sequence is $\rho_{t,j} := \log P_{t,j} - \log P_{t,j-1}$ and the intraday realised variance, RV_t^{intra} is the sum of squared intraday returns:

$$RV_t^{intra} = \sum_{j=1}^{M} \rho_{t,j}^2$$

Index volatility is not confined to trading hours, macroeconomic announcements and changes in global markets contribute materially to close-to-open moves. Martens (2002) and Hansen & Lunde (2005) show that overnight variance is much noisier than intraday realised variance. As such, simple approaches like adding the overnight variance to the intraday variance tends to worsen precision of the variance measure.

However, the option-implied variance (VIX in our case) is the risk-neutral expectation of total return variation over the month period, not limiting itself to trading hour changes. As such, our realised variance measure should capture close-to-close variance across the period too. To minimise added noise from the overnight returns we follow the scaled approach taken in Bollerselv and Todorov (2023). Let C_{t-1} be the prior close and O_t be the current open. We define close-to-open move as $\Delta_t^{ON} = \log O_t - \log C_{t-1}$ and the overnight realised variance as $RV^{ON} = (\Delta_t^{ON})^2$. Next we compute the rolling overnight adjustment factor, $\widehat{\psi}_t$:

$$\widehat{\psi}_t = \frac{1}{252} \sum_{s=t-252}^{t} \frac{RV_s^{intra} + RV_s^{ON}}{RV_s^{intra}}$$

This scaling factor allows us to define our daily adjusted realised variance, $\widehat{RV}_t^{adj} = \widehat{\psi}_t RV_t^{intra}$ which is used to define our adjusted annualised realised variance and volatility ($\widehat{\sigma}$) predictors for each horizon τ, based on Prokopczuk and Simen (2014):

$$\widehat{RV}_{t,\tau}^{PS,adj} = \frac{252}{\tau} \sum_{s=t-\tau+1}^{t} RV_s^{adj}, \qquad \widehat{\sigma}_{t,\tau}^{PS} = \sqrt{\widehat{RV}_{t,\tau}^{PS,adj}}$$

4.2 State probability distributions

The return distribution of the SPX is assumed to be a skewed generalised t-distribution (SGT) as introduced by Theodossiou (1998) and utilised in Beare et al. (2025). The rationale for using the SGT is its ability to fit the skewness and excess kurtosis observed in the SPX's empirical distribution (Theodossiou, 1998). The normalised returns $X_{t,\tau} = \frac{r_{t,\tau} - \mu_{t,\tau}}{\sigma_{t,\tau}}$, are fitted with the SGT $X_{t,\tau} \sim \text{SGT}(\hat{k}_\tau, \hat{v}_\tau, \hat{\lambda}_\tau; \mu_{t,\tau}, \sigma_{t,\tau})$. Where parameters, shape ($\hat{k}_\tau$), degrees of freedom ($\hat{v}_\tau$), and asymmetry ($\hat{\lambda}_\tau$), are computed by maximum likelihood. The location parameter ($\mu_{t,\tau}$) and scale parameter ($\sigma_{t,\tau}$) vary with both time and horizon. We discretise the continuous terminal SPX value distribution into a finite state space at 5 index point intervals.

4.2.1 Location parameter and constant market risk premium

We set the location parameter ($\mu_{t,\tau}$) to the current risk-free rate plus an additive market risk premium (MKRP). To accommodate horizon-specific parametrisation, we compute the MKRP for each horizon as the annualised mean τ-period excess return:

$$\text{MKRP}_\tau = \frac{1}{\tau}\sum_{t\in\tau} r^x_{t,\tau}$$

As such the location parameter is $\mu_{t,\tau} = (1 + r^f_t + \text{MKRP}_\tau)^{\frac{\tau}{252}} - 1$. Note that the MKRP for all horizons falls within a $\pm 0.6\%$ range of 7.2%. Beare et al. (2025) found that setting this premium at 5% or 8% was inconsequential for portfolio selection, an observation supported by Post and Longarela (2021).

4.2.2 Scale parameter and multiplicative variance risk premium

The scale parameter ($\sigma_{t,\tau}$) of the SGT distribution is set as the horizon specific forward-looking variance. The VIX embeds a variance risk premium, leading to it overstating variance when measured against realised variance. As such, to construct an accurate forward-looking variance measure we follow Prokopczuk and Simen (2014) who argue a multiplicative adjustment is optimal. We compute the multiplicative variance risk premium (MVRP) to be both time and horizon dependent:

$$\text{MVRP}_{t,\tau} = \sqrt{\frac{1}{\tau}\sum_{t\in\tau}\left(\frac{\text{VIX}_t \cdot \tau/21}{\widehat{\sigma}^{PS}_{t,\tau}}\right)^2}$$

We use a simple linear scaling for the VIX so that we minimise exaggerating implied variance on shorter time horizons. Note that since 2011, CBOE reports 9-day volatility index (VIX9D). Since our data period begins in 2004, the VIX9D was not applicable, but is likely to provide a better estimate of implied volatility over shorter horizons. Indeed, the simple scaling of the VIX may contribute to poor goodness-of-fit of the SGT distributions at shorter time horizons as discussed in the following subsections. CBOE has published the full methodology to construct a VIX1D measure which could be adapted to other horizons (CBOE, 2025). With the necessary data future investigations may benefit from following this approach.

Figure 4 shows the MVRP at each horizon length, with an expected decreasing relationship, that ensures that the VIX is more heavily adjusted for shorter time spans. The MVRP is used to compute an annualised, horizon-scaled volatility ($\sigma_{t,\tau}$)

used in the SGT distribution:

$$\sigma_{t,\tau} = \frac{\text{VIX}_t \cdot \tau/21}{\text{MVRP}_{t,\tau}} \cdot \sqrt{\frac{\tau}{252}}$$

Figure 4: Multiplicative variance risk premium by horizon

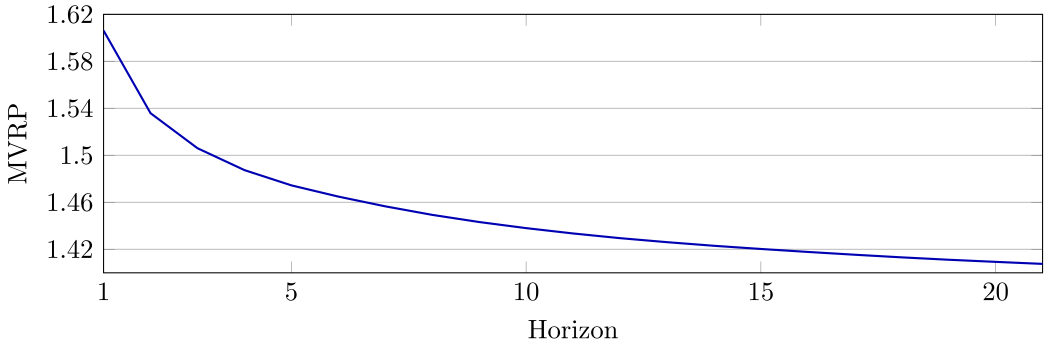

4.2.3 Shape, degrees of freedom and asymmetry

We use maximum likelihood to estimate parameters, shape ($\hat{k}_\tau$), degrees of freedom ($\hat{v}_\tau$), and asymmetry ($\hat{\lambda}_\tau$) using daily observations of normalised horizon-specific SPX returns. These parameters are time-invariant but horizon-specific. The degrees of freedom ($\hat{v}_\tau$) is set to be at minimum 5, ensuring a finite fourth moment and following the approach taken in Beare et al. (2025). Figure 5 shows the relationship between the estimates of shape, degrees of freedom and asymmetry respectively. For horizon 1, we see a significant outlier in the degrees of freedom of 140. For horizon 1, we see that the shape, degrees of freedom and asymmetry are closer to the parameter values of the standard normal distribution, $k = 2$, $v = \infty$ and $\lambda = 0$.

Figure 5: Estimated SGT parameters by horizons

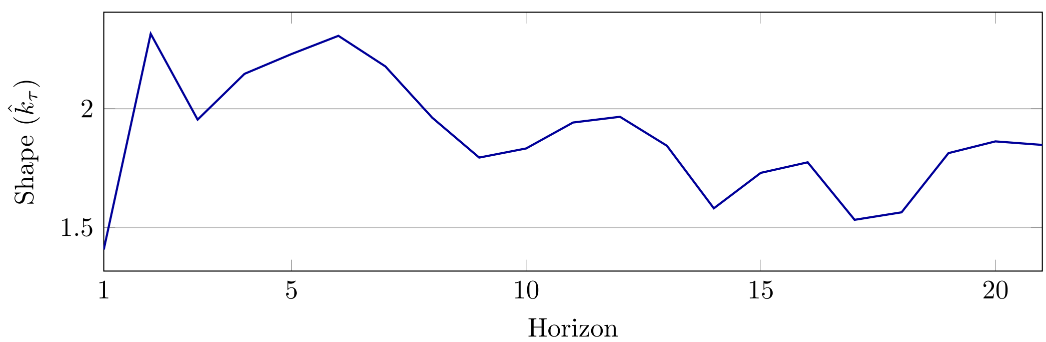

Figure 5 continues on next page.

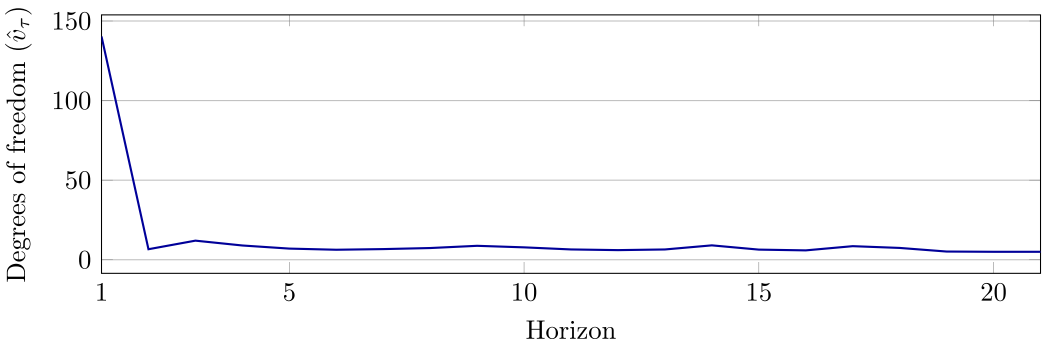

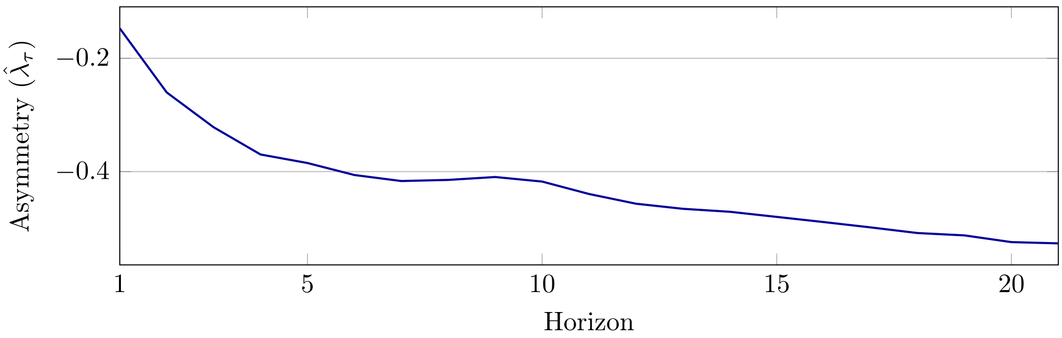

4.2.4 The goodness-of-fit of horizon-dependent distributions

Deibold et al. (1998) show that if a one-step ahead predictive cumulative distribution function (CDF), F_t is correctly specified then the probability integral transform (PIT) $z_t = F_t(x_t)$ is uniformly distributed: $\{z_t\}_{t=1}^{m} \sim iid\ U(0,1)$. To apply this approach to the horizon-specific state probability distributions, on each day t and for all horizons τ we construct horizon-specific, forward-looking (to date t, from $t - \tau$) CDF $F_{t,\tau}$. This is done using the parameters defined both on date t (for the location parameter $(\mu_{t,\tau})$ and scale parameter $(\sigma_{t,\tau})$) and horizon, τ, for the estimated parameters. As defined previously, let S_t be the SPX value on date t. Hence, we compute the PIT $z_{t,\tau} = F_{t,\tau}(S_t)$. By binning the PITs by decile we can analyse the uniformity of our predictions. Figure 6 presents a histogram of daily observations of location and scale normalised horizon-based SPX returns overlaid by the density functions of the standard normal distribution and the specified SGT distribution for that horizon. Additionally, it presents the PITs for both the standard normal and specified SGT.

As mentioned, if state probability distributions are correctly specified, each unit interval in the histogram of PITs should contain approximately 10% of the observed PITs. To analyse the suitability of our approach to state probability section, we split the 21 horizons into three subsets, short horizons (1-7), intermediate horizons (8-14) and long horizons (15-21). **Short horizons.** The left panels presenting the histogram of normalised SPX returns exhibit negative skewness and excess kurtosis. The center of the distributions are close to Gaussian but the tails are heavier and biased towards losses. The fitted SGT (red line), with the exception of the 1 day horizon, demonstrate improved fit in capturing asymmetry when compared to the normal distribution (blue dashed line). Interestingly, the SGT distribution for the 1 day horizon, has a very sharp peak, driven by the significantly higher degrees of freedom estimated by maximum likelihood and presented in Figure 5. When examining the middle and right panel, which present the histograms of PITs for SGT and standard Normal distributions respectively, two observations can be made. Firstly, at the extrema, the SGT significantly underestimates the probability of significantly negative and significantly positive returns. Similarly, the standard Normal distribution exhibits both less uniformity across the inner-deciles and an underestimation of significant negative and position returns (though to a lesser extent). This observation is consistent with theory, indeed, Campbell, Lo and MacKinlay (1997) document that shorter horizon returns exhibit significant departures from normality. But evidently the specification of the SGT distribution may be inappropriate at this horizon. For horizon 2 there is less underprediction of significant positive and negative returns, and improved but yet imperfect uniformity suggesting inappropriate specification of the SGT distribution. **Intermediate horizons.** Across these horizons the distribution of SPX returns become more peaked, with greater concentration around zero and the slightly positive center. The distribution retains the skewness characteristics that justify the application of the SGT distribution as explained in Theodossiou (1998). **Long horizons.** While to a lesser extent than the very short horizons, the underprediction of very large positive and negative returns reemerges. However, the uniformity in inner deciles is significantly improved. As explained in Beare et al. (2025) the underpredition of tail events is of little consequence, given the moderate strike range restriction, which ensures that this will have minimal effect on the results. Hence, the uniformity in inner deciles indicates appropriate fit.

4.3 Model

In this section we present the central portfolio choice model and describe the algorithm used in this investigation. Consider an investor who takes a unit position in the market index. One month later, she liquidates her unit position, realising a nonnegative random payoff $x_1 < \cdots < x_n$ with probabilities $\mu_1, \ldots, \mu_n$. The investor is perceptive to information arrival each day following her initial investment. As market conditions change, the investor can augment her positions by taking additional long and short positions. The investor observes m call and put options written on the market index, each delivering some nonnegative random payoff in a month $\theta_i(x)$ $\forall i \in m$ options, dependent on the market index value at expiry, x. Hence, the payoff of a long position in call option i with strike s_i is $\theta_i(x) = (x - s_i)^+$ and similarly for a long put position $\theta_i(x) = (s_i - x)^+$. The investor can take a long position (buy) in the ith option at price p_i per unit, or a short position (sell) at price $q_i \geq 0$ per unit, where $p_i \geq q_i$.

We have $\mathbf{\Theta}_{m\times n}$ with entries θ_{ij}. $\boldsymbol{\mu}$ and $\boldsymbol{x}$ are $n \times 1$ vectors with entries μ_j and x_j. $\mathbf{p}$ and $\mathbf{q}$ are $m \times 1$ vectors with entries p_i and q_i. $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ are $m \times 1$ vectors representing, respectively, the long and short positions taken by the investor during the current trading day. Let $\bar{\boldsymbol{\alpha}}$ and $\bar{\boldsymbol{\beta}}$ be $m \times 1$ vectors that store net long and short positions chosen on all prior trading days (within a trading month). We define $\mathbf{v}$ as an $m \times 1$ vector with entries being maximum long positions available in each option. Similarly, $\mathbf{w}$ is an $m \times 1$ vector with maximum short positions in each option. The investor seeks to maximise premium; the negative of the price of the layover portfolio. If the payoff distribution of the unit position augmented by the layover portfolio stochastically dominates the payoff distribution of the unit position, the investor expects to outperform (*or perform at least as well as*) the market while collecting upfront premium. This represents a stochastic arbitrage opportunity, in which the maximal value of the objective function (Equation [1]) is positive, yielding option premium, **and** the enhanced portfolio stochastically dominates the market. Beyond the first trading day, the changes to the enhanced portfolio must stochastically dominate both the unit position in the market and positions taken previously (Equation [2]). In principle, this should ensure that the optimiser (Gurobi) only takes positions that probabilistically improves the total portfolio relative to updating market conditions. In summary, the investor chooses a layover portfolio $(\boldsymbol{\alpha}, \boldsymbol{\beta})$ on every trading day, that solves the optimisation problem:

$$\text{maximise} \quad -\mathbf{p}^\top \boldsymbol{\alpha} + \mathbf{q}^\top \boldsymbol{\beta} \tag{1}$$

$$\text{subject to} \quad x + \sum_{i=1}^{m} \bar{\alpha}_i \bar{\theta}_i(x) + \sum_{i=1}^{m} \alpha_i \theta_i(x) - \sum_{i=1}^{m} \bar{\beta}_i \bar{\theta}_i(x) - \sum_{i=1}^{m} \beta_i \theta_i(x) \tag{2}$$

$$\gtrsim x + \sum_{i=1}^{m} \bar{\alpha}_i \bar{\theta}_i(x) - \sum_{i=1}^{m} \bar{\beta}_i \bar{\theta}_i(x),$$

$$(\boldsymbol{\alpha}, \boldsymbol{\beta}) \in \mathcal{P}. \tag{3}$$

Where $\mathcal{P}$ is a polytope that defines the set of feasible long and short positions.

$$\mathcal{P} = \{(\boldsymbol{\alpha}, \boldsymbol{\beta}) \in \mathbb{R}_+^m \times \mathbb{R}_+^m : \mathbf{A}\boldsymbol{\alpha} + \mathbf{B}\boldsymbol{\beta} \leq \mathbf{c}\} \tag{4}$$

Where,

$$\mathbf{A} = \begin{bmatrix} \mathbf{I}_m \\ \mathbf{0}_{m\times m} \\ \mathbf{a}_{2m+1} \\ -\mathbf{a}_{2m+1} \\ \mathbf{a}_{2m+3} \\ -\mathbf{a}_{2m+3} \\ \mathbf{a}_{2m+5} \\ -\mathbf{a}_{2m+5} \\ \mathbf{a}_{2m+7} \\ -\mathbf{a}_{2m+7} \end{bmatrix}, \quad \mathbf{B} = \begin{bmatrix} \mathbf{0}_{m\times m} \\ \mathbf{I}_m \\ -\mathbf{a}_{2m+1} \\ \mathbf{a}_{2m+1} \\ -\mathbf{a}_{2m+3} \\ \mathbf{a}_{2m+3} \\ -\mathbf{a}_{2m+5} \\ \mathbf{a}_{2m+5} \\ -\mathbf{a}_{2m+7} \\ \mathbf{a}_{2m+7} \end{bmatrix}, \quad \mathbf{c} = \begin{bmatrix} \mathbf{v} \\ \mathbf{w} \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}.$$

In basic terms, $\mathbf{A}$ and $\mathbf{B}$ are $\ell \times m$ matrices and $\mathbf{c}$ is an $\ell \times 1$ vector that, firstly, enforce that $\mathbf{v}$ and $\mathbf{w}$, the maximum available long and short positions in each option are not exceeded. And secondly, corresponding to the final 8 entries, ensures that the layover portfolio delivers zero payoff outside the range of strikes as derived in Beare et al. (2025). As previously mentioned, the zero payoff constraint ensures that bets on extreme tail events are prevented.

4.3.1 Mixed integer linear program

To solve the portfolio choice problem presented in Equations 1, 2, and 3, we need to translate the stochastic dominance constraint into a linear program. This allows readily available linear programming solvers, Gurobi in our case, to solve the problem directly. Additionally, because we wish to impose integer restrictions on long, $\boldsymbol{\alpha}$ and short, $\boldsymbol{\beta}$ positions, to limit portfolio sparsity, the linear formulation can be solved as a mixed integer linear program (MILP). To achieve this we modify the linear formulation of SOSD presented in equations (3.2), (3.4), (3.6) and (3.7) in Leudtke (2008).

Firstly, Theorem 3.2 in Leudtke (2008) establishes that for random variables W, Y with probabilities $p_i\ \forall i \in \mathcal{N} := \{1, ..., N\},\ q_k\ \forall k \in \mathcal{D} := \{1, ..., D\}$ respectively, we have $W \succeq_{(2)} Y$ if and only if there exists $\pi \in \mathbb{R}_+^{ND}$ which satisfies (3.2), (3.4), (3.6) and (3.7). Equation 3.2 is $\sum_{j=1}^{D} \pi_{ij} = 1\ \forall i \in \mathcal{N}$. We define a nonnegative matrix $\boldsymbol{\Psi} \in \mathbb{R}^{n \times n}$ as the equivalent to π such that we have $\sum_{j=1}^{n} \Psi_{ij} = 1,\ i = 1, 2, ..., n$ which we express in matrix form $\boldsymbol{\Psi}\mathbf{1}_n = \mathbf{1}_n$. Note that n, as previously defined, is the number of payoff states. In equation 3.6 Leudtke introduces variable $v \in \mathbb{R}^D$ and constraint $v_j - \sum_{i=1}^{N} p_i \pi_{ij} = 0,\ j \in \mathcal{D}$. We introduce nonnegative vector $\xi \in \mathbb{R}^n$ as an equivalent to v, and with μ as our probability vector over states we express the constraint as $\xi_j - \sum_{i=1}^{n} \mu_i \Psi_{ij} = 0$ and in matrix form $\boldsymbol{\xi} - \boldsymbol{\Psi}^\top \boldsymbol{\mu} = \mathbf{0}_n$. Equation 3.7 is:

$$\sum_{j=1}^{k-1} (y_k - y_j) v_j \leq \sum_{j=1}^{k-1} (y_k - y_j) q_j, \quad k \in \mathcal{D}$$

Following Beare et al. (2025) the equivalent formulation in matrix notation is $S\xi \leq S\mu$ where S is an $n \times n$ strictly lower triangular matrix with each element being one. The equivalence here is less direct. In place of $y_1 < y_2 < ... < y_D$ we have the payoff states $x_1 < x_2 < ... < x_n$. Due to the discretisation of the continuous SPX return distribution into increments of 5 index points over the specified strike range, the scaling factor represented by $\sum_{j=1}^{k-1} (y_k - y_j)$ is simply 1 in our application. As such we have $\sum_{j=1}^{k-1} v_j \leq \sum_{j=1}^{k-1} q_j$ where in place of v we have $\boldsymbol{\xi}$ and in place of q we have $\boldsymbol{\mu}$. And therefore, in matrix form $\mathbf{S}\boldsymbol{\xi} \leq \mathbf{S}\boldsymbol{\mu}$. To complete the system, Luedtke defines constraint 3.4 as $w_i \geq \sum_{j=1}^{D} y_j \pi_{ij}$. Beare et al. (2025) express this as

$\mathbf{x} + \mathbf{\Theta}^T(\boldsymbol{\alpha} - \boldsymbol{\beta}) \geq \mathbf{\Psi}\mathbf{x}$ where w is replaced with the layover portfolio, $\mathbf{x} + \mathbf{\Theta}^T(\boldsymbol{\alpha} - \boldsymbol{\beta})$, which is required to be greater than the market payoff $\mathbf{x}$ weighted by the auxiliary choice variable $\mathbf{\Psi}$. To accommodate the dynamic trading framework, we require the reference variable to incorporate both the market and positions taken previously, but evaluated according to today's state probability matrix. Hence, with $\bar{\boldsymbol{\alpha}}$ and $\bar{\boldsymbol{\beta}}$ representing the long and short positions taken previously, the constraint becomes $\mathbf{\Psi}(\mathbf{x} + \mathbf{\Theta}^T(\bar{\boldsymbol{\alpha}} - \bar{\boldsymbol{\beta}})) \leq \mathbf{x} + \mathbf{\Theta}^T(\boldsymbol{\alpha} - \boldsymbol{\beta}) + \mathbf{\Theta}^T(\bar{\boldsymbol{\alpha}} - \bar{\boldsymbol{\beta}})$. Thus, we define a sequential stochastic dominance problem unique to both Leudtke (2008) and Beare et al. (2025). Formally the SOSD MILP is:

Algorithm. Let $\mathbf{1}_n$ be an $n \times 1$ vector of ones, $\mathbf{S}$ be an $n \times n$ lower triangular matrix where all entries below the diagonal are one. Choose two $m \times 1$ nonnegative vectors $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$, two binary vectors $\boldsymbol{\alpha}_b$ and $\boldsymbol{\beta}_b$, a sparsity scalar k, an $n \times 1$ nonnegative random vector $\boldsymbol{\xi}$ and a $n \times n$ nonnegative matrix $\mathbf{\Psi}$ to maximise Equation 1 subject to the linearisation of Equation 2 and three constraints to limit portfolio sparsity:

$$\mathbf{\Psi}\mathbf{1}_n = \mathbf{1}_n, \tag{5}$$

$$\boldsymbol{\xi} - \mathbf{\Psi}^\top \boldsymbol{\mu} = \mathbf{0}_n, \tag{6}$$

$$\mathbf{S}\boldsymbol{\xi} \leq \mathbf{S}\boldsymbol{\mu}, \tag{7}$$

$$\mathbf{\Psi}(\mathbf{x} + \bar{\mathbf{\Theta}}^\top(\bar{\boldsymbol{\alpha}} - \bar{\boldsymbol{\beta}})) \leq \mathbf{x} + \bar{\mathbf{\Theta}}^\top(\bar{\boldsymbol{\alpha}} - \bar{\boldsymbol{\beta}}) + \mathbf{\Theta}^\top(\boldsymbol{\alpha} - \boldsymbol{\beta}), \tag{8}$$

$$\mathbf{A}\boldsymbol{\alpha} + \mathbf{B}\boldsymbol{\beta} \leq \mathbf{c}, \tag{9}$$

$$\boldsymbol{\alpha} \leq \mathbf{v} \odot \boldsymbol{\alpha}_b, \tag{10}$$

$$\boldsymbol{\beta} \leq \mathbf{w} \odot \boldsymbol{\beta}_b, \tag{11}$$

$$\mathbf{1}_n^\top(\boldsymbol{\alpha}_b + \boldsymbol{\beta}_b) \leq k. \tag{12}$$

This MILP algorithm contains a total of $4n+2m+\ell+1$ scalar equality, inequality and binary constraints and $2n + 2m$ choice variables. To limit portfolio sparsity, we have introduced [10], [11] and [12]. Equation [12] imposes that the combination of unique long and short options do not exceed the set sparsity parameter. Equations [10] and [11] allow the positions in the unique options to extend to the maximum volume available where optimal. These binary constraints add $2m$ choice variables and an additional 1 constraint. This increases computational complexity, leading to larger runtime, especially for lower values of k where the portfolio is most constrained.

4.3.2 Alternate approaches to limiting sparsity

When researching optimal approaches to limit portfolio sparsity, we initially applied type 1 special-ordered-set (SOS1) constraints rather than the cardinality constraints as shown above. SOS1 constraints enforce that only one variable in a set can take a nonzero position. Both two and four SOS1 constraints were applied, one (two) on long and one (two) on short positions. Whilst the realised performance of portfolios using the approach were underwhelming (this is a consequence of SOS1 constraints themselves), this investigation yielded some notable observations. Firstly, when computing portfolios with highly constrained sparsity (one or two long and short positions), SOS1 constraints were markedly more efficient than using cardinality constraints. This is likely attributable to the SOS1 constraints appealing to Gurobi's branch and bound algorithm (SOS functionality built into Gurobi directly), and the relative restriction of the SOS1 constraints is less than for the cardinality constraints, as Gurobi does not have the freedom to choose the composition of positions as with the cardinality constraints[2]. Gurobi had trouble computing six SOS1 constraints simultaneously, hence our decision to implement cardinality constraints instead. However, we formulated an iterative approach that would increase complexity linearly. All within a given month, Gurobi would optimise within a loop where the first iteration, one long and one short position is selected, then for a specified number of subsequent loops, Gurobi would continue to take one long and one short position, but now the stochastic dominance would be in reference to the market and the positions taken in prior iterations. The motive being that Gurobi would find optimal sparsity at the point at which there were no positions that stochastically dominate the portfolio already selected. This concept of iterative portfolio optimisation according to stochastic dominance inspired the changes to Equation [8] which in a static optimisation context is $\boldsymbol{\Psi}\mathbf{x} - \boldsymbol{\Theta}^\top(\boldsymbol{\alpha} - \boldsymbol{\beta}) \leq \mathbf{x}$.

[2]The cardinality constraints can be reformulated to constrain long and short positions separately to achieve something comparable to SOS1 constraints

5 Results

We compute day-level outcomes before aggregating to a single month. Note that the collected premia is invested at the risk-free rate for only the remaining time until option expiry before in-month summation (to prevent exaggerating the return of the enhanced portfolios). Hence, we report the results of dynamic trading strategy for the 213 months for sparsity values $K = [4 : 100, \infty]$. Where $K = \infty$ is the portfolio with no sparsity constraint. The results focus particularly on sparsity values $K = 4$ and $K = 100$ as they represent heavily constrained, and effectively unconstrained portfolios.

5.1 Portfolio dynamics

Before discussing the overall performance of the selected enhanced portfolios, we discuss the portfolio dynamics which shed light on the process that the optimiser takes. Namely, the way in which the cumulative portfolio is rebalanced throughout the month with respect to the sequential SOSD constraint:

$$\boldsymbol{\Psi}(\mathbf{x} + \bar{\boldsymbol{\Theta}}^{\top}(\bar{\boldsymbol{\alpha}} - \bar{\boldsymbol{\beta}})) \leq \mathbf{x} + \bar{\boldsymbol{\Theta}}^{\top}(\bar{\boldsymbol{\alpha}} - \bar{\boldsymbol{\beta}}) + \boldsymbol{\Theta}^{\top}(\boldsymbol{\alpha} - \boldsymbol{\beta})$$

Figures 11, 12 present the portfolio payoff functions for an arbitrary month (August 2004) for sparsity constraints $K = 4$ and $K = 100$ respectively. Figure 11 shows no positions taken for the first 16 days of the month. On day 17, the optimiser identifies a layover portfolio that firstly, generates positive option premium to be invested at the risk-free rate and satisfies the SOSD constraint. The piecewise linear geometry of the payoff function relates directly to Equation [9] which enforces zero payoff outside of the specified strike range (forward moneyness bounds). The effect of this constraint can be seen in practice on day 17, where the 'N' shaped payoff function is constructed such that outside of the range of strikes (which is at most -10% to +5%) payoff is zero. A minimum of four distinct strikes are required to replicate the 'N' shape that ensures zero-payoff outside of the strike range. Therefore, the minimum number of distinct positions is the maximally sparsity constrained portfolio, $K = 4$. On day 18 and again on day 19 the optimiser takes additional positions but with negligible weight to make large changes to the payoff function. Day 21 is the last day in the trading month, the red dotted line shows the terminal value of the SPX. Evidently the layover portfolio payoff is zero at expiry, an outcome commensurate with the objective of the optimisation: to generate positive option premium, exploiting a stochastic arbitrage opportunity.

Figure 12 shows the portfolio payoff functions for $K = 100$. We see that the optimiser enters the market earlier, on day 11, and then proceeds to reconstitute the cumulative portfolio on subsequent days. With sparsity less constrained, the optimiser has more flexibility with respect to satisfying the zero payoff constraint. Therefore, as the month progresses, the investor/optimiser observes changes to the SPX value and state probability distributions and sequentially stochastically improves the cumulative portfolio and generates more option premia. After day 17 the optimiser takes no new positions and we observe the terminal SPX value after day 21. Whilst the layover portfolio expires OTM, significantly more option premia is generated at $K = 100$ relative to $K = 4$ offsetting the payoff loss. Table 6 reports the results of August 2004 for $K = 4$ and $K = 100$. In the month the SPX returned a small 0.21%. For $K = 4$ the option payoff is 0, the few positions taken at $K = 4$ generate premium of 0.3 SPX units. Comparatively, for $K = 100$ the option payoff is -3.37 SPX units but the premium generated is 3.81. The net return of the enhanced portfolio (SPX return + option payoff + net premium) is marginally higher for $K = 100$ at 0.25% compared to 0.23%. This elucidates two fundamental components of this research. Firstly, the risk of unconstrained portfolios in the pursuit of stochastic arbitrage. When unconstrained the optimiser can take positions which are less rational. In this month, the premia generated outweighs the negative payoff, but generally this is not always the case. Secondly, at $K = 4$ where the optimiser takes more sensible positions (in principle) the resulting stochastic arbitrage opportunity is economically insignificant in magnitude. Indeed, for this month, the excess return achieved of 0.02%, is insignificant, and in practice, would likely be absorbed by trading costs.

Nevertheless, in a static trading framework the optimiser would have taken no positions this month (as no positions were taken on day 1), missing the stochastic arbitrage opportunities that exist closer to month expiry. This outlines the advantages of implementing a dynamic trading framework in this context.

5.2 Realised performance of enhanced portfolios

Tables 1, 2 report the characteristics of the options portfolios selected by the MILP for $K = 4$, the maximum sparsity constraint, for the dynamic trading strategy and a static trading strategy respectively. Columns $S = 1, 10, 100, 1000$ correspond to size of the investor where $S = 1$ is a retail investor of size $190,000 and $S = 1000$ is a large institutional investor of size $19,000,000. As S increases, each option trade becomes more capacity constrained as market depth is relatively tighter the larger the investor is. In practice, the available bid and ask sizes are scaled by S, the scale constraint. The first line in Table 1 and 2, "Pctg. premia 0.1% mkt. investment" report the percentage of the 213 months in which the premia generated by the layover portfolio exceeds 0.1% of the value of the SPX unit investment (SPX spot at month start). This metric is a measure of the frequency of stochastic arbitrage opportunities.

For the dynamic trading strategy (Table 1) this threshold is exceeded in 35.7% (76 months) of the 213 months for the unit scale constraint. At the tightest scale constraint we still see a promising 18.8% (40 months) of the 213 months with material premia generated. Comparatively, the static trading strategy (Table 2) yields considerably fewer months with significant premia generation of at most 1.9% for the unit scale and 0% for $S = 100$ and $S = 1000$. This is in line with expectations, given the dynamic trading strategy allows the investor approximately 21 times more trading opportunities (though available strikes shrinks according to the dynamic moneyness bounds). Although importantly, this provides strong evidence for the appropriateness of a dynamic trading framework to identify notable stochastic arbitrage opportunities.

However, the realised mean excess return of the enhanced portfolios reported in Table 1 is consistently below the SPX return over the 213 months of 9.39%. As expected, the return is higher for $S = 1000$ at 8.67%, as there are less options taken that can expire out-of-the-money. The static portfolios (Table 2) perform better by similar merit, even exceeding the market at scales $S = 10$ and $S = 100$, however, because there is so few months with premia generated and very minimal options positions taken as evidenced in rows titled "Calls bought", "Calls written", "Puts bought" and "Puts written" the majority is this return is driven by the unit position in the market with select layover portfolios expiring in-the-money. Furthermore, the p-values indicate that the slightly superior returns are not statistically significant. The p-values correspond to tests of the null hypothesis that the enhanced portfolio SOSD dominates the SPX return and the null hypothesis that the enhanced portfolio and the SPX have equal mean.

Tables 3, 4 report the characteristics of the options portfolios selected for sparsity $K = 100$ for the dynamic trading strategy and static trading strategy respectively. With the portfolio effectively unconstrained the optimiser can take more positions each day. As expected we see higher percentage of months with material premia generated, at 39% (83 months) for $S = 1$ and 18.8% (40 months) for $S = 1000$. Similar to $K = 4$ the realised return of the enhanced portfolio is worse than the market, but is marginally better than the results for $K = 4$. In Tables 1, 2, 3, and 4 we report the Sortino ratio, which is similar to the Sharpe ratio but it only penalises downside volatility (see Sortino and Price, 1994). Consistently, the reported Sortino ratios across sparsity values and investor size are worse than market, suggesting higher exposure to downside risk.

Table 5 reports the results of a dynamic linear program (no binary constraints to limit portfolio sparsity) for comparison. There is negligible difference between the unconstrained portfolios and portfolios with sparsity $K = 100$ with respect to months with material premia generated and quantity of calls and puts bought and written. It should be highlighted that the calls and puts bought and written are the net value of options as a percentage of the SPX. The reason that the percentages are similar between $K = 4$, $K = 100$ and $K = \infty$ is because the optimiser, when sparsity constrained takes larger weights in fewer distinct positions, as opposed to small weightings in a large number of distinct options contracts. When summed the values are comparable, but there are some notable advantages to holding fewer distinct contracts ($K = 4$) in the practical application of this strategy. Firstly, as an investor progresses through the month, taking positions each day, the total portfolio of options becomes progressively less tractable, especially if taking small weights across dozens of distinct contracts. Secondly, a retail investor is less likely to have access to the low brokerage/exchange fee arrangements that institutional investors do. Hence, while trading costs are not explicitly built into the portfolio choice model, limiting sparsity is advantageous from the perspective of lower trading fees from fewer distinct contracts traded.

Tables 2, 4 (static trading strategy) provide the most direct comparison to the results reported in Table 2 in Post and Longarela (2021) and Table 3 in Beare et al. (2025). Our results differ substantially from Post and Longarela (2021) who report portfolio mean return of 28% at the unit investor scale. As discussed in Beare et al. (2025) this is likely driven by the existence of pure arbitrage opportunities. Notably, for investor scales $S = 10, 100, 1000$ Post and Longarela (2021) report negative realised return signalling that the pure arbitrage opportunities are present only at small scales. Additionally, our results, while more modest at $S = 1$ at 8.93% (pure arbitrage is excluded) are more consistent at scales $S = 10, 100, 1000$. This suggests that the skewed generalised t-distribution provides an improved specification of the SPX return distribution relative to the Cox-Ross-Rubinstein binomial tree applied in their investigation. Moreover, Post and Longarela (2021) are more permissive of far OTM quotes, only limiting the range to those within the support of the estimated probability distribution which is $\pm 10\sigma$, stating that because bid and ask sizes are explicitly taken into account the selected portfolios are robust to excessive weighting in tail events. However, the improved mean return of enhanced portfolios in Table 4 and in Table 2 in Beare et al. (2025) contend that directly restricting forward moneyness and applying the zero payoff constraint offers enhanced protection against betting on tail events. This is further supported by the materially higher Sortino ratios reported in Table 4.

The top panel of Table 2 in Beare et al. (2025) corresponds to Table 4 (with the addition of sparsity constraint $K = 100$). Our results are, as expected, very similar, with minor differences driven by the exclusion of 2 months of the data period. Neither the results in Beare et al. (2025) nor those reported here yield statistically significant superior to market realised return, whether that be for a static or dynamic trading strategy and irrespective of portfolio sparsity. However, Figure 15 does show that a static trading strategy performs best when sparsity constrained across all investor scales. Similar is the case with the dynamic trading strategy, although with significantly more volatility at low sparsity values $K = [4, 10]$ see Figure 7. Post and Longarela (2021) highlight a design choice that contributes to the poor realised mean return. Common to Post and Longarela (2021), Beare et al. (2025) and this thesis, the available positions are constrained to the best bid/ask prices and volumes. This has two consequences, firstly, the volume of the best quoted bid/ask may be low. Greater volumes may be available at worse prices that could satisfy the objective of maximising the negative cost of the layover portfolio. Secondly, trading at the quoted bid/ask spread can adversely impact algorithmic traders by amplifying effective costs of trading options (Muravyev and Pearson, 2020). These spread frictions disproportionately burden large scale investors, who in practice could provide liquidity by posting limit orders inside the spread. Post and Longarela (2021) explore a counterfactual by assuming trades execute at the bid/ask midpoint. This yielded improved returns but not significantly enough for the portfolios to outperform the market.

The source of portfolio underperformance in the dynamic trading strategy introduced in this thesis is at least in part attributable to imperfect specification of state probability distributions. As discussed in Section 4.2.4, Figure 6 shows unsatisfactory goodness-of-fit of horizon-specific state probability distributions, especially at shorter horizons. The poor uniformity of the observed PITs is particularly problematic for horizons below 10. This is somewhat mitigated by the dynamic moneyness bounds, that progressively restrict the optimiser from taking bets on misspecified tail probabilities, but unfortunately the misspecification is not isolated to the tails. Future research may benefit from exploring alternate approaches to computing the market risk premium and variance risk premium at shorter horizons and assessing alternatives to the SGT distribution that may better fit the observed empirical SPX return distribution at shorter horizons.

5.3 Optimal portfolio sparsity

Figures 7, 8, 9, 10, and 15 present the results per metric (mean return, premia, Sortino, standard deviation) for all sparsity values and investor sizes. On each plot the red dotted line corresponds to the unconstrained portfolio. Figure 7 (Figure 15 for static comparison) shows the mean excess return of the enhanced portfolios across the range of sparsity values. Note that across all investor scales, portfolio performance is better when constrained and optimal performance from a mean return perspective happens at sparsity values lower than 10. Figure 8 presents the relationship between the percentage of months with material premia versus sparsity. We see that across all scale constraints the percentage of months is lower at $K = 4$ compared to $K = \infty$ but by around $K = 10$ there is convergence with $K = \infty$. With respect to downside risk, Figure 10 presents the enhanced portfolio Sortino ratio by sparsity. Consistently, constricting sparsity results in lower downside risk, evidenced by the higher Sortino. Across all scale constraints there seems to be an optimal level of sparsity within the $K = [4, 10]$ range. However, there is notable volatility with spikes and troughs happening at successive sparsity values. As previously mentioned, to satisfy the zero payoff constraint at least 4 distinct positions are required. Additionally, the optimiser can not take 5 distinct positions and still satisfy this constraint. Having the additional space (1 extra position) means the optimiser considers more before realising infeasibility due to requiring an even number of positions to satisfy the constraint. This may manifest itself as the optimiser spending too much time considering possibilities (in the 60 second run time) rather than selecting the most optimal portfolio according to SOSD. This volatility in the small sparsity values is consistent with the standard deviation plots in Figure 10, but by around $K = 10$ the standard deviation converges with the unconstrained portfolio. Note that this volatility is specific to the dynamic strategy and is not observed in the static strategy (Figure 15).

Ergo, a sparsity value in the range of $K = [4, 10]$ is optimal. Within this range, portfolio mean return is highest, there is less downside risk and an insignificant reduction in the percentage of months where material premia is generated. Combined with the practical advantages of sparser portfolios, cardinality constraints are effective in this context. This is consistent with results from Evans and Archer (1968) and Arvanitis et al. (2024) who find that there is no benefits from expanding portfolio sparsity beyond 10-20 assets and 45 assets respectively.

5.4 Sub-sample analysis - COVID-19

The COVID-19 pandemic resulted in an extreme 30% crash in the SPX in a single month. This provides an opportunity to assess the merits of the dynamic trading strategy during periods of exceptional price moves. Figures 13, 14 present payoff functions for March 2020 for $K = 4$ and $K = 100$ respectively. Figure 13 shows a large position taken on day 1. Observing the rapidly decreasing price additional positions are taken on day 2, augmenting the layover portfolio to have greater payoff in the left tail. Due to the rapid price moves, the zero payoff constraint effectively moves leftwards with the SPX value. After day 2 no further positions are taken. At expiry the option payoff is zero, protected by the zero payoff outside of strike range constraints. Figure 14, shows significant trading activity on a few days in the month. We see the zero payoff line begin around 80% of the SPX month start level. The rapidly decreasing price allowed the optimiser to take short meaningful short positions that generated significant amounts of option premia. This is summarised in Table 7. While for both sparsity levels, the layover portfolio payoff is 0 there is significant premia generated of 48.4 and 129.5 in SPX units for $K = 4$ and $K = 100$ respectively. Of course, because we hold a passive position in the market, we are exposed to the entirety of the market loss, however, the option premium softened the downturn. Instead of -30.1%, the loss was moderately reduced to -28.7% and -26.3%.

Generally, positive layover payoffs correspond to months with moderately negative market returns, a result consistent with Beare et al. (2025). For extreme price moves like COVID-19, there are benefits from dynamic portfolio rebalancing, as it can capitalise on generating premia but because of the zero payoff constraint, the net layover portfolio is unlikely to expire ITM. Importantly, the payoff function shows that the investor is exposed to significant losses if the market rebounded during the month. This is especially problematic if the portfolio is not rebalanced in response to price changes. If the rebound begins to occur toward the tail end of the month, the dynamic moneyness bounds are likely to prove problematic in not allowing the optimiser to take extreme strikes to balance the existing positions. Though, this problem is limited to extremely volatile months.

6 Concluding remarks

This empirical research pursues two primary objectives. Firstly, to formulate and evaluate the suitability of a dynamic trading strategy to identify stochastic arbitrage opportunities with market index options. Secondly, to identify optimal portfolio sparsity to maximise realised portfolio performance by avoiding overdiversification[3]. This paper achieves these objectives by contributing a dynamic trading framework, that employs a MILP that applies sequential stochastic dominance and incorporates cardinality constraints.

The results contribute to the literature exploring the pricing kernel puzzle, corroborating the claim that perceived mispricing is a consequence of misspecification in state probabilities rather than systematic mispricing. This is evidenced by both the static and dynamic trading strategies performing worse relative to a passive investment in the market. However, this paper does identify an optimal sparsity range to avoid overdiversification, improving portfolio characteristics. Improved specification of state probability distributions, particularly in a sparsity constrained dynamic trading context may yield portfolios that outperform the market.

[3]Taking both sensible positions in mispriced options (a stochastic arbitrage opportunity), but also less rational positions driven by state probability misspecification

References

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

Arvanitis, S., Scaillet, O., & Topaloglou, N. (2024). Sparse spanning portfolios and under-diversification with second-order stochastic dominance. arXiv.Org.

Barro, R. J., & Liao, G. Y. (2021). Rare disaster probability and options pricing. Journal of Financial Economics, 139(3), 750–769. https://doi.org/10.1016/j.jfineco.2020.10.001

Bawa, V. S., Bodurtha JR, J. N., Rao, M. R., & Suri, H. L. (1985). On Determination of Stochastic Dominance Optimal Sets. The Journal of Finance (New York), 40(2), 417–431. https://doi.org/10.1111/j.1540-6261.1985.tb04965.x

Beare, B. K. (2011). Measure preserving derivatives and the pricing kernel puzzle. *Journal of Mathematical Economics*, *47*(6), 689–697. https://doi.org/10.1016/j.jmateco.2011.09.005

Beare, B. K., Seo, J., & Zheng, Z. (2025). Stochastic arbitrage with market index options. *Journal of Banking & Finance*, *173*, 107395-. https://doi.org/10.1016/j.jbankfin.2025.107395

Black, F., & Scholes, M. (1973). The pricing of options and corporate liabilities. *The Journal of political economy*, *81*(3), 637–654.

Bollerslev, T., & Todorov, V. (2023). The jump leverage risk premium. *Journal of Financial Economics*, *150*(3), 103723-. https://doi.org/10.1016/j.jfineco.2023.103723

Brennan, M. J., & Schwartz, E. S. (1977). The Valuation of American Put Options. The Journal of Finance (New York), 32(2), 449. https://doi.org/10.2307/2326779

Broske, M. S., Levy, H., Fomby, T. B., & Seo, T. K. (1989). The Stochastic Dominance Estimation of Default Probability. In *Studies in the Economics of Uncertainty* (pp. 91–112). Springer New York. https://doi.org/10.1007/978-1-4613-8922-4_6

Campbell, J. Y., Lo, A. W., & MacKinlay, A. C. (2012). The Econometrics of Financial Markets. Princeton University Press,. https://doi.org/10.1515/9781400830213

Carr, P., & Madan, D. B. (2005). A note on sufficient conditions for no arbitrage. *Finance Research Letters*, *2*(3), 125–130. https://doi.org/10.1016/j.frl.2005.04.005

Chabi-Yo, F., Garcia, R., & Renault, E. (2008). State Dependence Can Explain the Risk Aversion Puzzle. *The Review of Financial Studies*, *21*(2), 973–1011. https://doi.org/10.1093/rfs/hhm070

Consigli, G., Dentcheva, D., Maggioni, F., & Micheli, G. (2025). Asset liability management under sequential stochastic dominance constraints. arXiv:2505.16486. https://doi.org/10.48550/arXiv.2505.16486

Constantinides, G. M., Jackwerth, J. C., & Perrakis, S. (2009). Mispricing of S&P 500 Index Options. *The Review of Financial Studies*, *22*(3), 1247–1277. https://doi.org/10.1093/rfs/hhn009

Dentcheva, D., & Ruszczyński, A. (2006). Portfolio optimization with stochastic dominance constraints. Journal of Banking & Finance, 30(2), 433–451. https://doi.org/10.1016/j.jbankfin.2005.04.024

Diebold, F. X., Gunther, T. A., & Tay, A. S. (1998). Evaluating Density Forecasts with Applications to Financial Risk Management. International Economic Review (Philadelphia), 39(4), 863–883. https://doi.org/10.2307/2527342

Evans, J. L., & Archer, S. H. (1968). Diversification and the Reduction of Dispersion: An Empirical Analysis. *The Journal of Finance (New York)*, *23*(5), 761-. https://doi.org/10.2307/2325905

Fang, Y., & Post, T. (2017). Higher-degree stochastic dominance optimality and efficiency. *European Journal of Operational Research*, *261*(3), 984–993. https://doi.org/10.1016/j.ejor.2017.03.035

Fang, Y., & Post, T. (2022). Optimal portfolio choice for higher-order risk averters. *Journal of Banking & Finance*, *137*, 106429-. https://doi.org/10.1016/j.jbankfin.2022.106429

Hadar, J., & Russell, W. R. (1969). Rules for Ordering Uncertain Prospects. *The American Economic Review*, *59*(1), 25–34.

Hanoch, G., & Levy, J. H. (1969). The Efficiency Analysis of Choices Involving Risk. *The Review of Economic Studies*, *36*(3), 335–346. https://doi.org/10.2307/2296431

Hansen, P. R., & Lunde, A. (2005). A forecast comparison of volatility models: does anything beat a GARCH(1,1)? Journal of Applied Econometrics (Chichester, England), 20(7), 873–889. https://doi.org/10.1002/jae.800

Harris, T. R., & Mapp, H. P. (1986). A Stochastic Dominance Comparison of Water-Conserving Irrigation Strategies. *American Journal of Agricultural Economics*, *68*(2), 298–305. https://doi.org/10.2307/1241431

Hull, J. C. (2021). Options, futures, and other derivatives (11th ed.). Pearson.

Jackwerth, J. C. (2000). Recovering Risk Aversion from Option Prices and Realized Returns. *The Review of Financial Studies*, *13*(2), 433–451. https://doi.org/10.1093/rfs/13.2.433

Levy, H. (1992). Stochastic Dominance and Expected Utility: Survey and Analysis. *Management Science*, *38*(4), 555–593. https://doi.org/10.1287/mnsc.38.4.555

Luedtke, J. (2008). New Formulations for Optimization under Stochastic Dominance Constraints. *SIAM Journal on Optimization*, *19*(3), 1433–1450. https://doi.org/10.1137/070707956

Markowitz, H. (1952). Portfolio Selection. *The Journal of Finance (New York)*, *7*(1), 77-. https://doi.org/10.2307/2975974

Martens, M. (2002). Measuring and forecasting S&P 500 index-futures volatility using high-frequency data. The Journal of Futures Markets, 22(6), 497–518. https://doi.org/10.1002/fut.10016

Mei, Y., Chen, Z., Liu, J., & Ji, B. (2022). Multi-stage portfolio selection problem with dynamic stochastic dominance constraints. Journal of Global Optimization, 83(3), 585–613. https://doi.org/10.1007/s10898-021-01113-z

Ogryczak, Wl., & Ruszczynski, A. (2002). Dual Stochastic Dominance and Related Mean-Risk Models. *SIAM Journal on Optimization*, *13*(1), 60–78. https://doi.org/10.1137/S1052623400375075

Post, T., & Kopa, M. (2017). Portfolio Choice Based on Third-Degree Stochastic Dominance. Management Science, 63(10), 3381–3392. https://doi.org/10.1287/mnsc.2016.2506

Post, T., & Longarela, I. R. (2021). Risk Arbitrage Opportunities for Stock Index Options. *Operations Research*, *69*(1), 100–113. https://doi.org/10.1287/opre.2020.2012

Prokopczuk, M., & Wese Simen, C. (2014). The importance of the volatility risk premium for volatility forecasting. Journal of Banking & Finance, 40(1), 303–320. https://doi.org/10.1016/j.jbankfin.2013.12.002

Roman, D., Mitra, G., & Zverovich, V. (2013). Enhanced indexation based on second-order stochastic dominance. *European Journal of Operational Research*, *228*(1), 273–281. https://doi.org/10.1016/j.ejor.2013.01.035

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

Rothschild, M., & Stiglitz, J. E. (1970). Increasing risk: I. A definition. *Journal of Economic Theory*, *2*(3), 225–243. https://doi.org/10.1016/0022-0531(70)90038-4

Sortino, F.A. and Price, L.N. (1994) Performance Measurement in a Downside Risk Framework. Journal of Investing, 3, 59-64. https://doi.org/10.3905/joi.3.3.59

Theodossiou, P. (1998). Financial Data and the Skewed Generalized T Distribution. *Management Science*, *44*(12-Part-1), 1650–1661. https://doi.org/10.1287/mnsc.44.12.1650

Von Neumann, J., & Morgenstern, O. (1944). *Theory of games and economic behavior.* Princeton University Press.

Whitmore, G. A. (1970). Third-Degree Stochastic Dominance. *The American Economic Review*, *60*(3), 457–459.

Williams, J. B. (1938). The theory of investment value. *Harvard University Press.*

Table 1: Portfolio characteristics with sparsity $K = 4$

				Enhanced portfolio			
			SPX	$S = 1$	$S = 10$	$S = 100$	$S = 1000$
MILP	Pctg. premia > 0.1% mkt. investment			35.7	27.7	18.8	9.4
	Value of options as a percentage of market investment (average)	Calls bought		2.8	2.0	0.9	0.2
		Calls written		3.2	2.3	1.1	0.2
		Puts bought		3.0	1.6	0.9	0.2
		Puts written		3.3	1.9	1.0	0.3
	Realized moments of excess returns	Mean	9.39	7.92	7.53	7.94	8.67
		Std. dev.	17.64	18.15	17.54	17.55	17.55
		Skew	-1.43	-0.71	-1.16	-1.29	-1.41
		Sortino	0.69	0.60	0.57	0.59	0.64
	Stochastic dominance p-values			0.068	0.015	0.014	0.039
	Equal mean excess return p-value			0.402	0.050	0.043	0.062

Notes. The realised mean and standard deviation are reported in annualised percentage points. The p-values correspond to the null hypothesis that the enhanced portfolio excess returns second order stochastically dominate the SPX excess returns, and to the null hypothesis that the enhanced portfolio excess returns and SPX excess returns have equal mean.

Table 2: **Static** portfolio characteristics with sparsity $K = 4$

				Enhanced portfolio			
			SPX	$S = 1$	$S = 10$	$S = 100$	$S = 1000$
MILP	Pctg. premia > 0.1% mkt. investment			1.9	1.4	0.0	0.0
	Value of options as a percentage of market investment (average)	Calls bought		0.2	0.2	0.1	0.0
		Calls written		0.2	0.2	0.1	0.0
		Puts bought		1.1	0.2	0.1	0.0
		Puts written		1.1	0.2	0.1	0.0
	Realized moments of excess returns	Mean	9.39	9.27	9.46	9.57	9.39
		Std. dev.	17.64	17.60	17.56	17.56	17.63
		Skew	-1.43	-1.42	-1.43	-1.44	-1.43
		Sortino	0.69	0.68	0.70	0.71	0.69
	Stochastic dominance p-values			0.351	0.911	0.892	0.892
	Equal mean excess return p-value			0.730	0.856	0.493	0.960

Table 3: Portfolio characteristics with sparsity $K = 100$

				Enhanced portfolio			
			SPX	$S = 1$	$S = 10$	$S = 100$	$S = 1000$
MILP	Pctg. premia > 0.1% mkt. investment			39.0	31.0	24.4	18.8
	Value of options as a percentage of market investment (average)	Calls bought		2.8	2.4	1.9	1.1
		Calls written		4.0	3.3	2.7	1.5
		Puts bought		3.7	2.4	1.9	1.2
		Puts written		3.5	2.2	1.7	1.0
	Realized moments of excess returns	Mean	9.39	8.41	7.55	7.63	8.23
		Std. dev.	17.64	18.06	17.30	17.34	17.52
		Skew	-1.43	-0.49	-0.94	-0.97	-1.25
		Sortino	0.69	0.65	0.59	0.59	0.62
	Stochastic dominance p-values			0.129	0.030	0.015	0.023
	Equal mean excess return p-value			0.593	0.087	0.058	0.080

Table 4: **Static** portfolio characteristics with sparsity $K = 100$

				Enhanced portfolio			
			SPX	$S = 1$	$S = 10$	$S = 100$	$S = 1000$
MILP	Pctg. premia > 0.1% mkt. investment			5.2	5.2	4.7	2.8
	Value of options as a percentage of market investment (average)	Calls bought		0.3	0.3	0.3	0.2
		Calls written		0.5	0.4	0.4	0.3
		Puts bought		1.4	0.4	0.3	0.2
		Puts written		1.2	0.3	0.2	0.2
	Realized moments of excess returns	Mean	9.39	8.93	8.64	8.52	8.74
		Std. dev.	17.64	17.60	17.59	17.60	17.59
		Skew	-1.43	-1.40	-1.39	-1.39	-1.40
		Sortino	0.69	0.66	0.64	0.63	0.65
	Stochastic dominance p-values			0.251	0.119	0.090	0.113
	Equal mean excess return p-value			0.477	0.203	0.132	0.152

Table 5: Portfolio characteristics with sparsity $K = \infty$ (no sparsity constraints)

				Enhanced portfolio			
			SPX	$S = 1$	$S = 10$	$S = 100$	$S = 1000$
LP	Pctg. premia > 0.1% mkt. investment			39.0	31.0	24.4	18.8
	Value of options as a percentage of market investment (average)	Calls bought		2.8	2.4	1.9	1.1
		Calls written		4.0	3.3	2.7	1.6
		Puts bought		3.7	2.4	1.9	1.2
		Puts written		3.5	2.2	1.7	1.1
	Realized moments of excess returns	Mean	9.39	8.30	7.43	7.60	8.21
		Std. dev.	17.64	18.05	17.30	17.35	17.52
		Skew	-1.43	-0.49	-0.94	-0.97	-1.24
		Sortino	0.69	0.65	0.58	0.59	0.62
	Stochastic dominance p-values			0.114	0.020	0.014	0.022
	Equal mean excess return p-value			0.549	0.064	0.054	0.076

Figure 6: Goodness-of-fit of horizon-specific SPX return distributions

Horizon 1

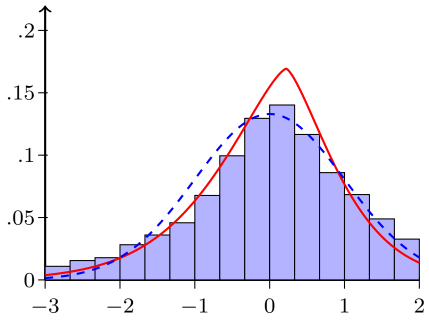

(a) Normalized SPX return histogram

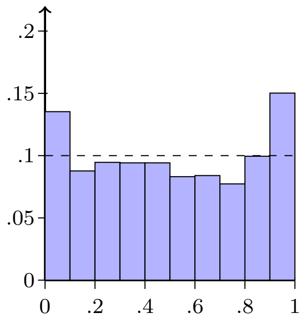

(b) PIT histogram: SGT

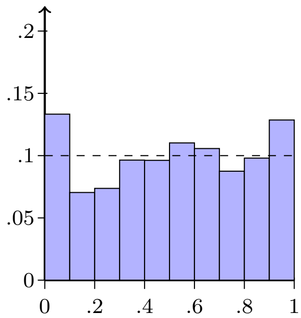

(c) PIT histogram: $\mathcal{N}(0,1)$

Horizon 2

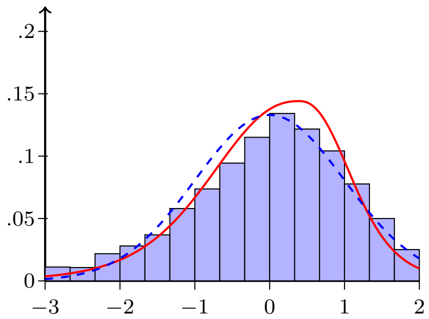

(a) Normalized SPX return histogram

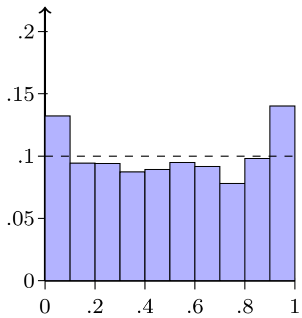

(b) PIT histogram: SGT

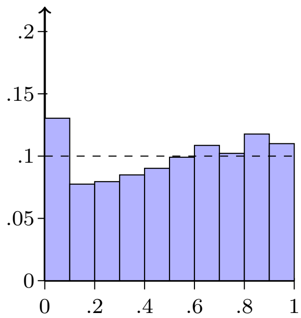

(c) PIT histogram: $\mathcal{N}(0,1)$

Horizon 3

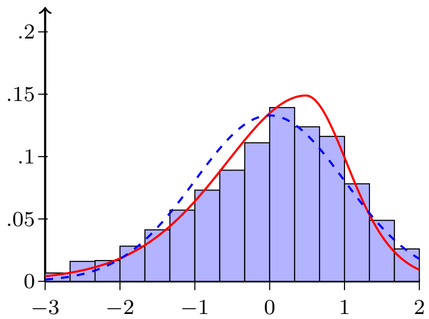

(a) Normalized SPX return histogram

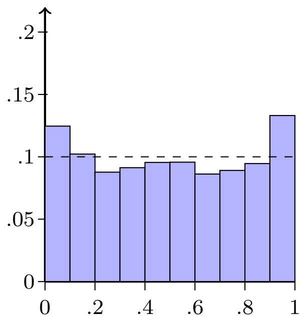

(b) PIT histogram: SGT

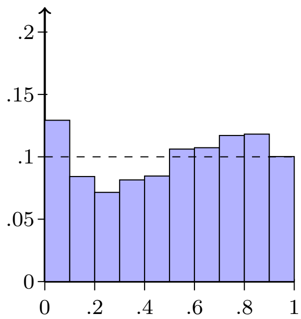

(c) PIT histogram: $\mathcal{N}(0,1)$

Figure 6 continues on next page

Horizon 4

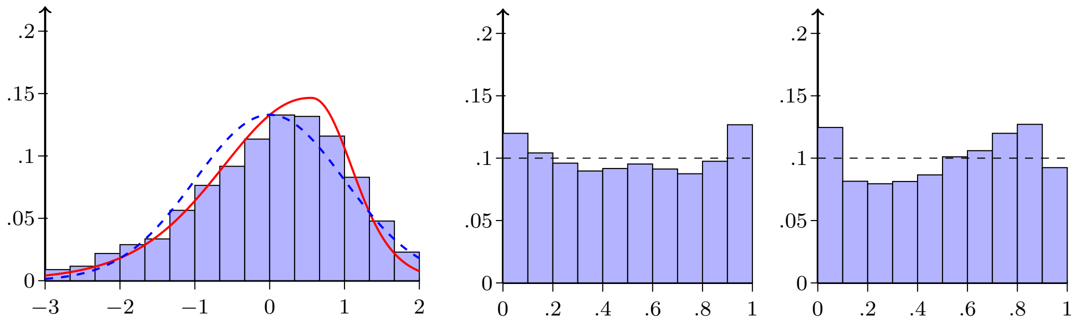

(a) Normalized SPX return histogram (b) PIT histogram: SGT (c) PIT histogram: $\mathcal{N}(0,1)$

Horizon 5

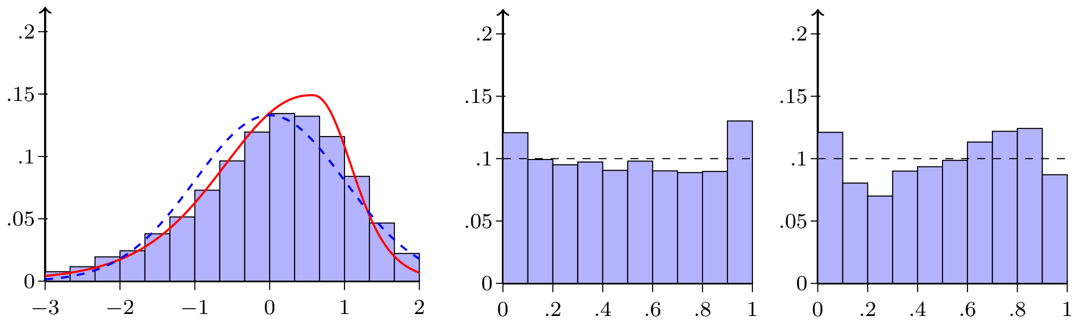

(a) Normalized SPX return histogram (b) PIT histogram: SGT (c) PIT histogram: $\mathcal{N}(0,1)$

Horizon 6

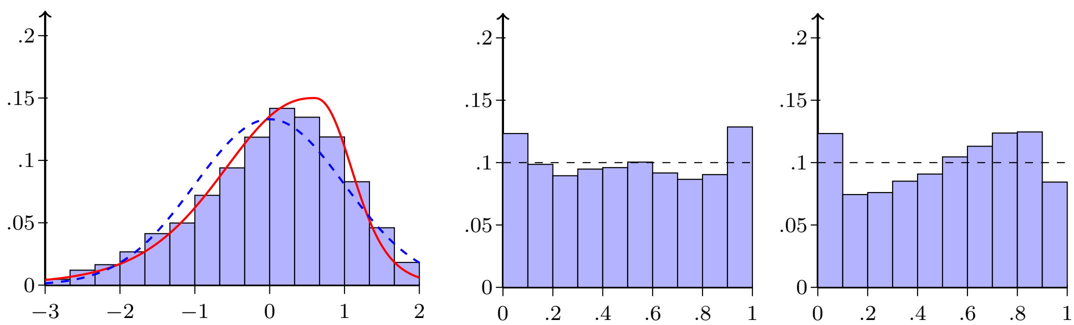

(a) Normalized SPX return histogram (b) PIT histogram: SGT (c) PIT histogram: $\mathcal{N}(0,1)$

Figure 6 continues on next page

Horizon 7

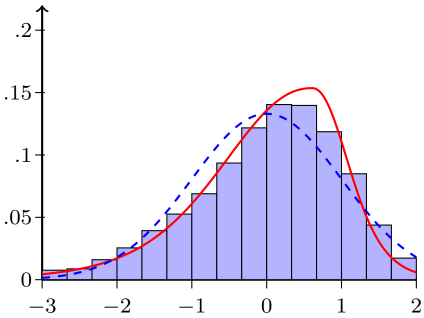

(a) Normalized SPX return histogram

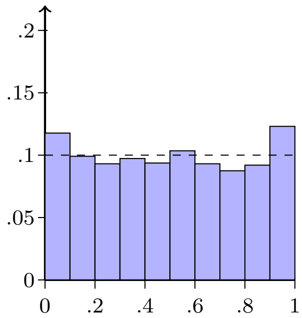

(b) PIT histogram: SGT

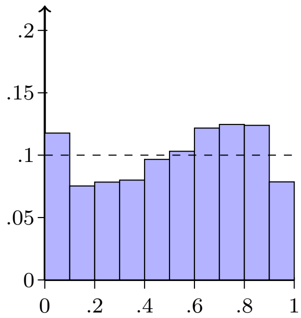

(c) PIT histogram: $\mathcal{N}(0,1)$

Horizon 8

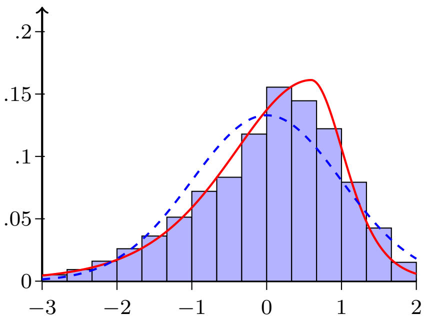

(a) Normalized SPX return histogram

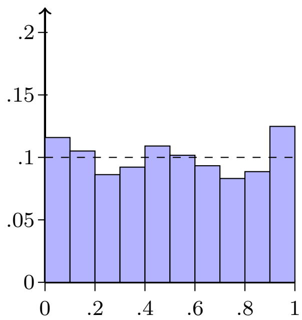

(b) PIT histogram: SGT

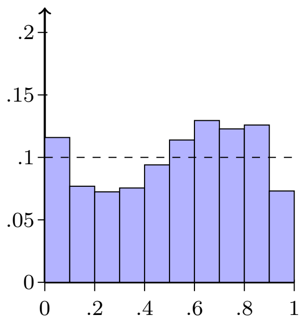

(c) PIT histogram: $\mathcal{N}(0,1)$

Horizon 9

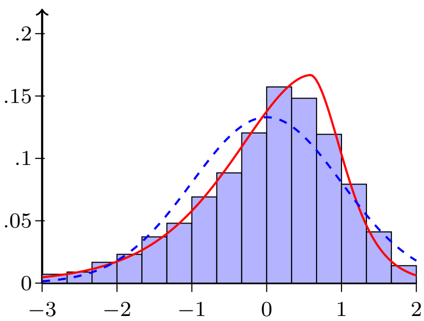

(a) Normalized SPX return histogram

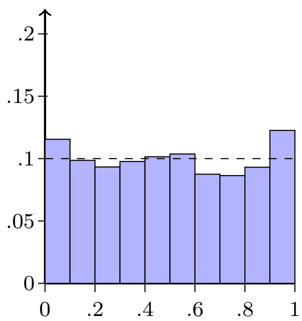

(b) PIT histogram: SGT

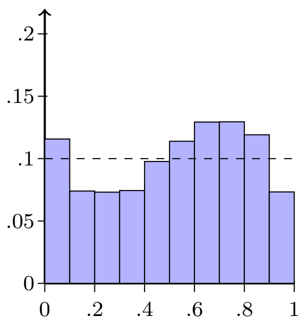

(c) PIT histogram: $\mathcal{N}(0,1)$

Figure 6 continues on next page

Horizon 10

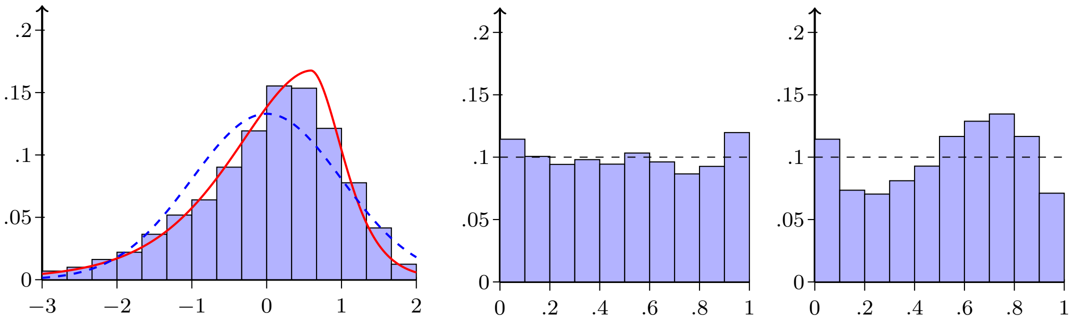

(a) Normalized SPX return histogram (b) PIT histogram: SGT (c) PIT histogram: $\mathcal{N}(0,1)$

Horizon 11

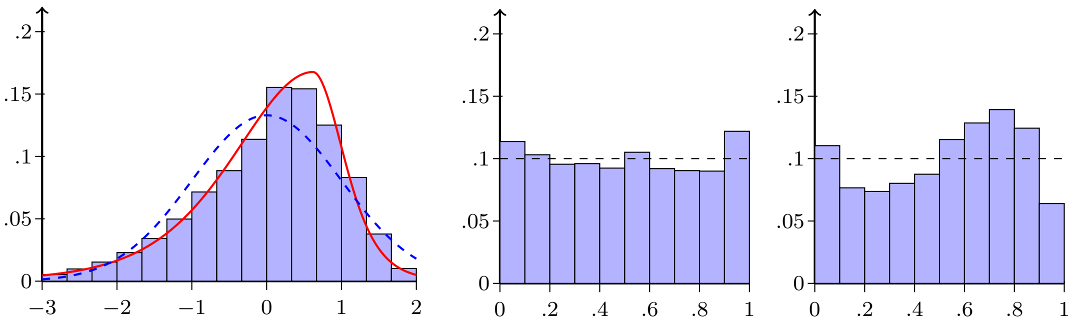

(a) Normalized SPX return histogram (b) PIT histogram: SGT (c) PIT histogram: $\mathcal{N}(0,1)$

Horizon 12

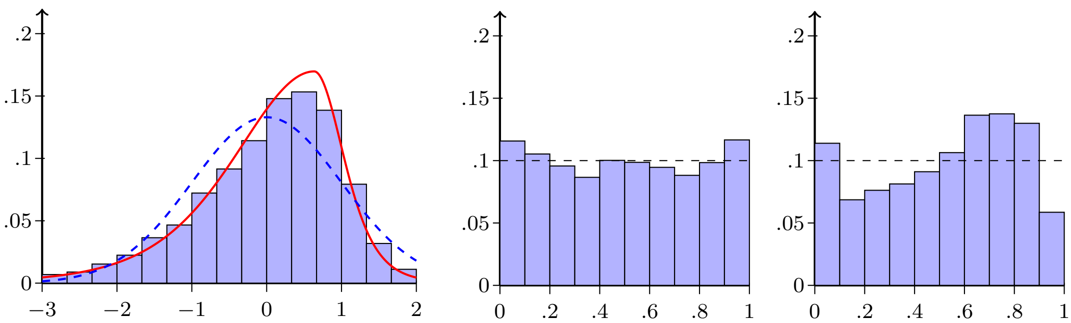

(a) Normalized SPX return histogram (b) PIT histogram: SGT (c) PIT histogram: $\mathcal{N}(0,1)$

Figure 6 continues on next page

Horizon 13

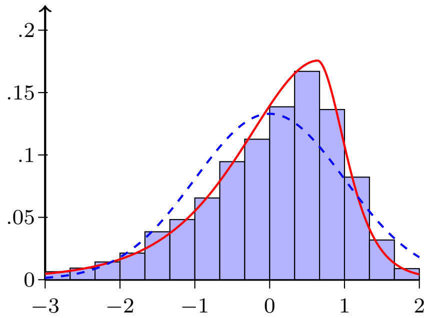

(a) Normalized SPX return histogram

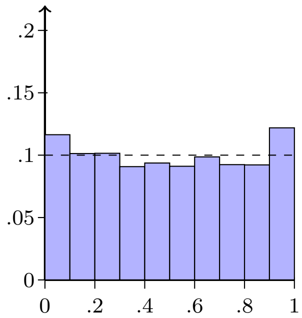

(b) PIT histogram: SGT

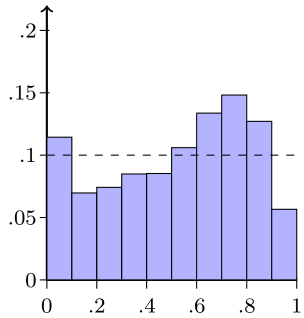

(c) PIT histogram: $\mathcal{N}(0, 1)$

Horizon 14

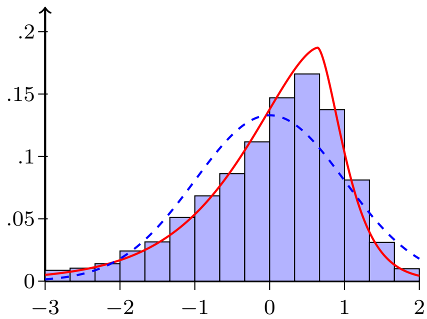

(a) Normalized SPX return histogram

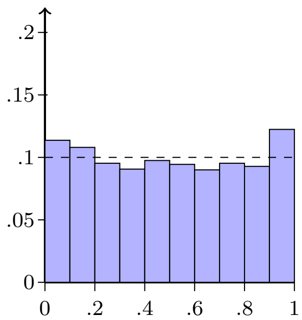

(b) PIT histogram: SGT

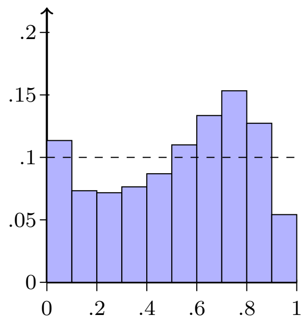

(c) PIT histogram: $\mathcal{N}(0, 1)$

Horizon 15

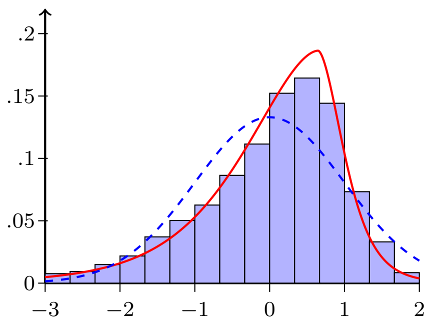

(a) Normalized SPX return histogram

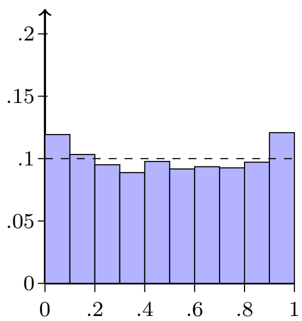

(b) PIT histogram: SGT

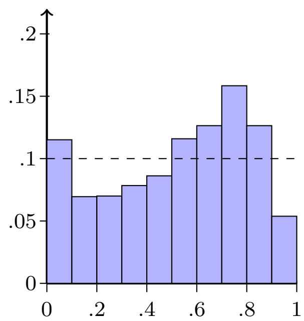

(c) PIT histogram: $\mathcal{N}(0, 1)$

Figure 6 continues on next page

Horizon 16

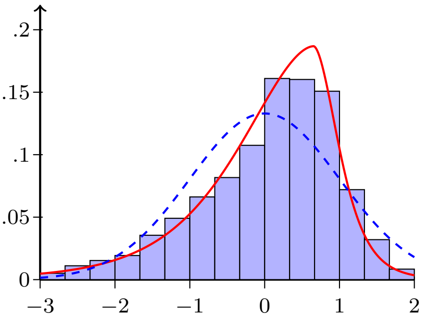

(a) Normalized SPX return histogram

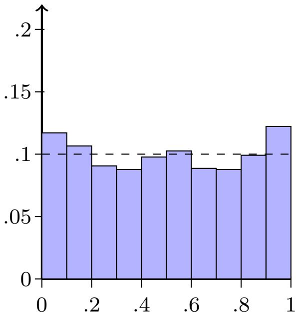

(b) PIT histogram: SGT

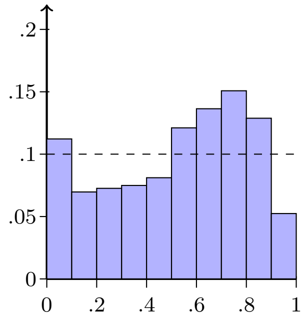

(c) PIT histogram: $\mathcal{N}(0,1)$

Horizon 17

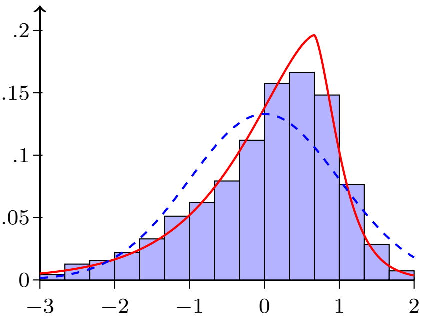

(a) Normalized SPX return histogram

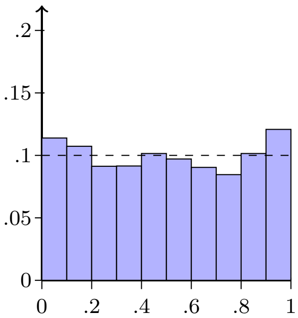

(b) PIT histogram: SGT

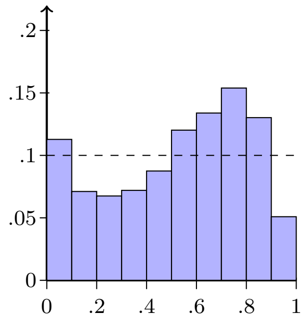

(c) PIT histogram: $\mathcal{N}(0,1)$

Horizon 18

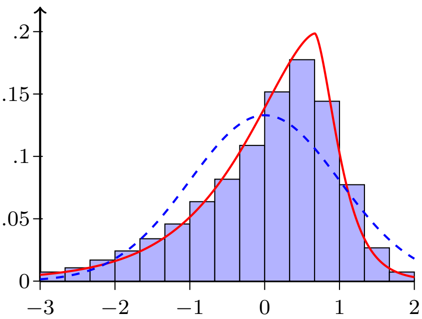

(a) Normalized SPX return histogram

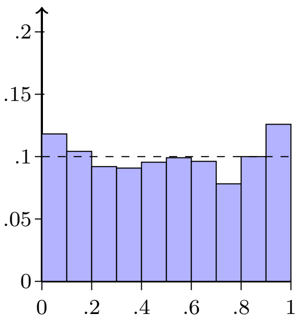

(b) PIT histogram: SGT

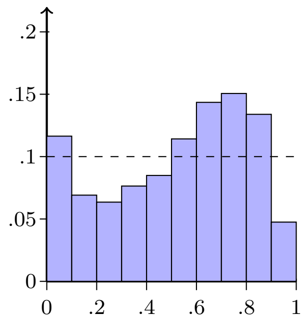

(c) PIT histogram: $\mathcal{N}(0,1)$

Figure 6 continues on next page

Horizon 19

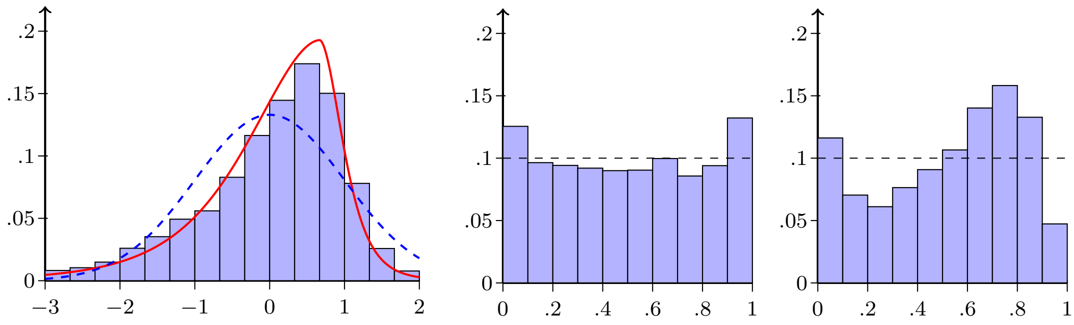

(a) Normalized SPX return histogram (b) PIT histogram: SGT (c) PIT histogram: $\mathcal{N}(0,1)$

Horizon 20

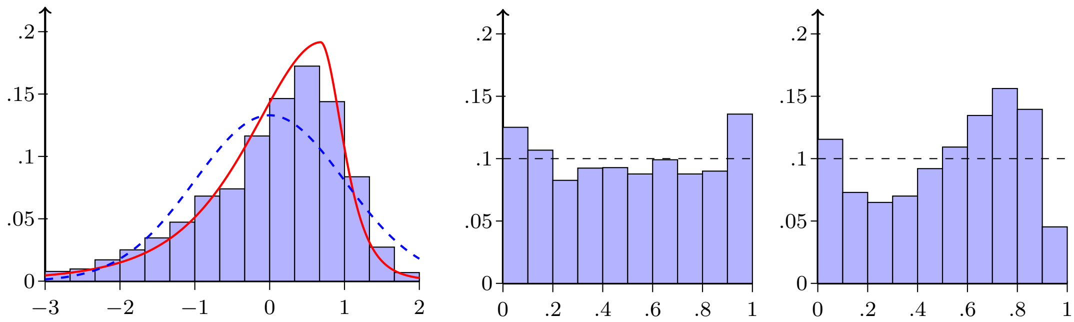

(a) Normalized SPX return histogram (b) PIT histogram: SGT (c) PIT histogram: $\mathcal{N}(0,1)$

Horizon 21

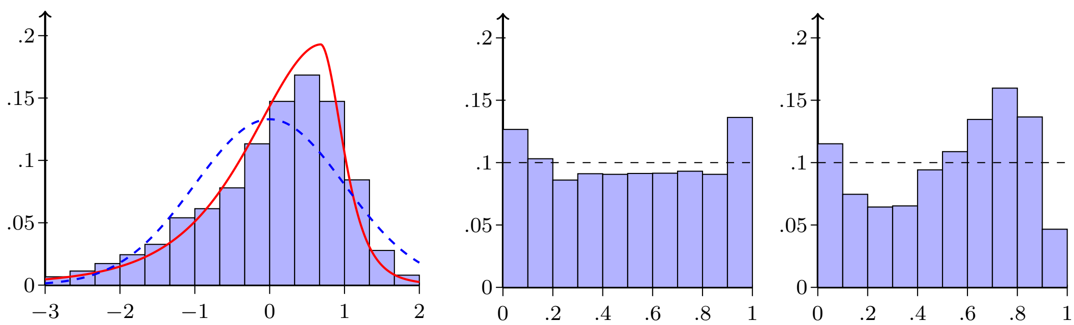

(a) Normalized SPX return histogram (b) PIT histogram: SGT (c) PIT histogram: $\mathcal{N}(0,1)$

Notes. Panels (a) present a histogram of normalised SPX returns for each horizon overlaid with the specified SGT distribution (red line) and the standard Normal for comparison (blue dashed line). Panels (b) display histograms of the observed PITs for the SGT distribution and Panels (c) display the histograms of the observed PITs for the standard Normal distribution for goodness-of-fit comparison.

Figure 7: Mean excess return by sparsity K and investor scale S

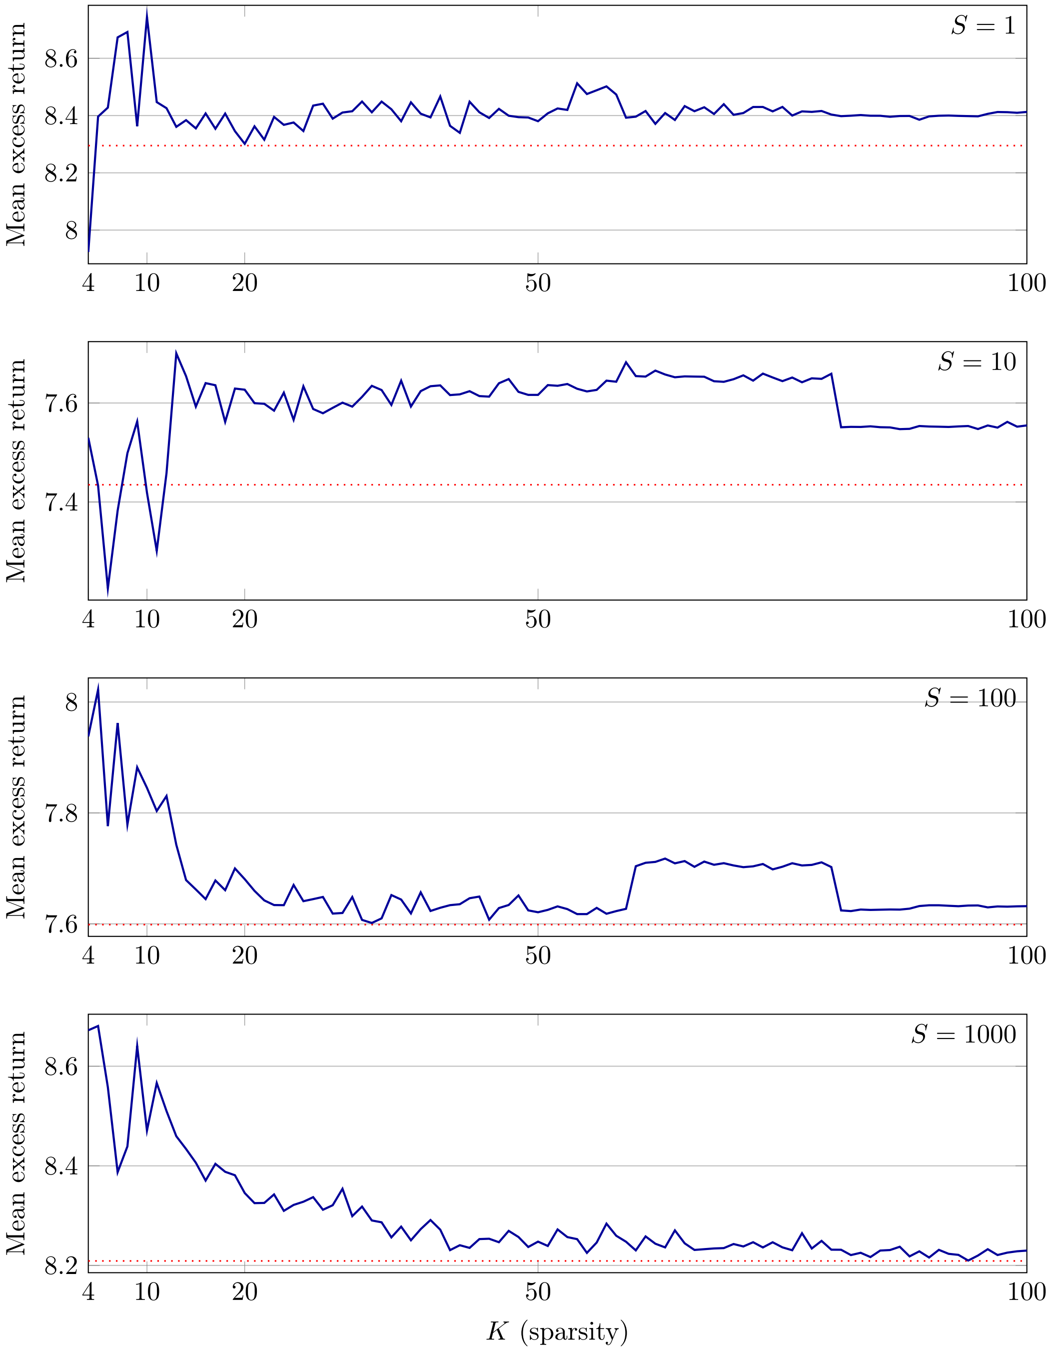

Notes: The red dotted line corresponds to the portfolio with no sparsity constraints.

Figure 8: Enhanced portfolio % of months with premia > 0.1% by K and investor scale S

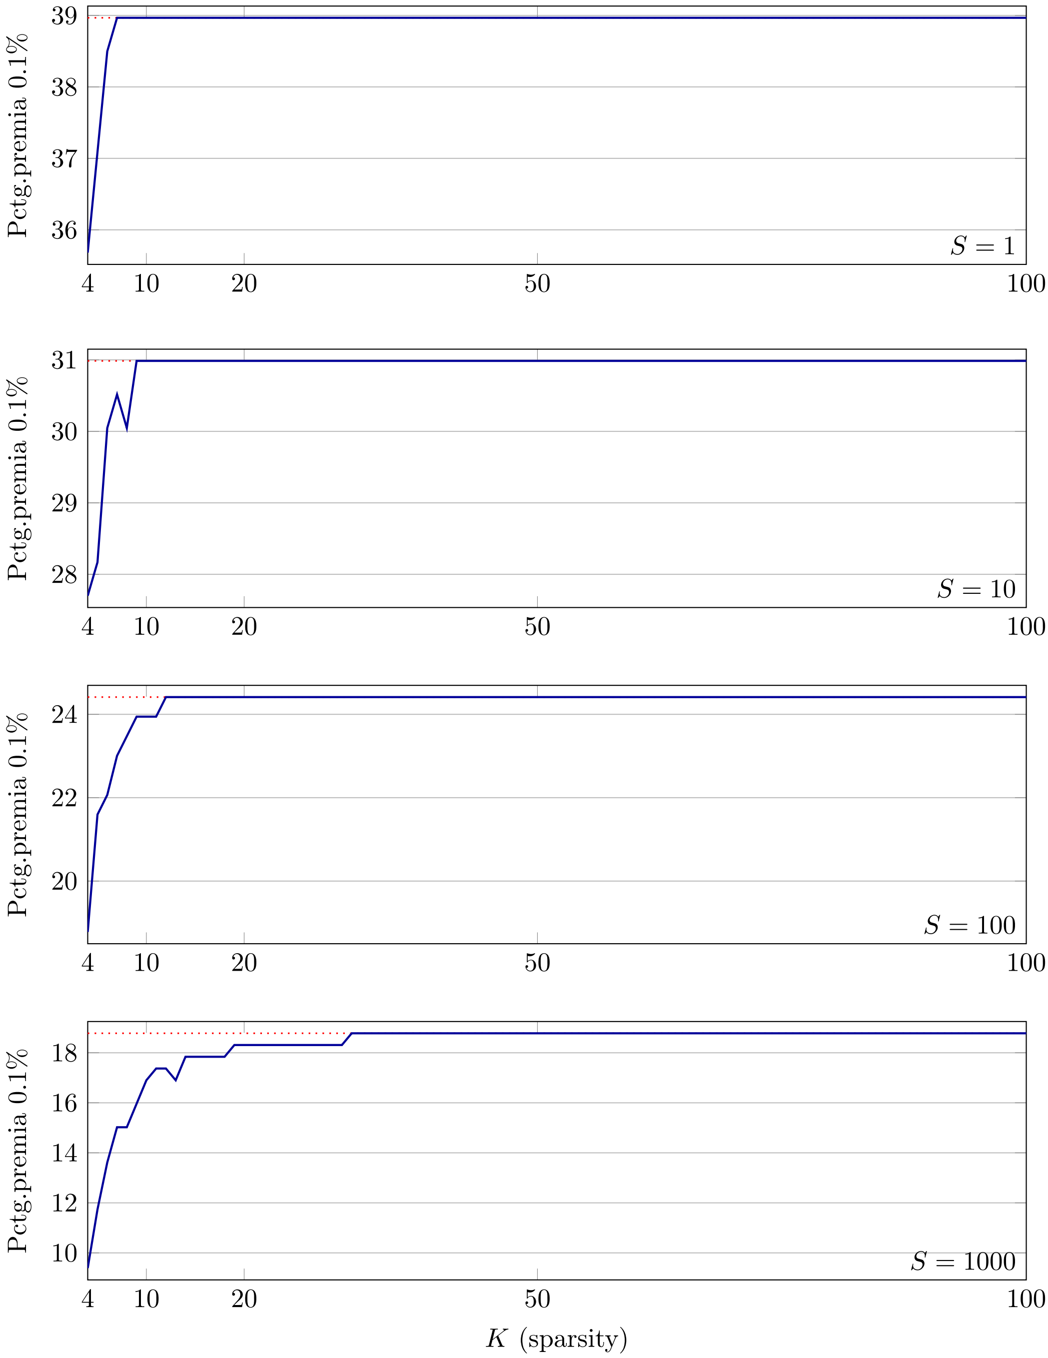

Notes: The red dotted line corresponds to the portfolio with no sparsity constraints.

Figure 9: Enhanced portfolio Sortino by K and investor scale S

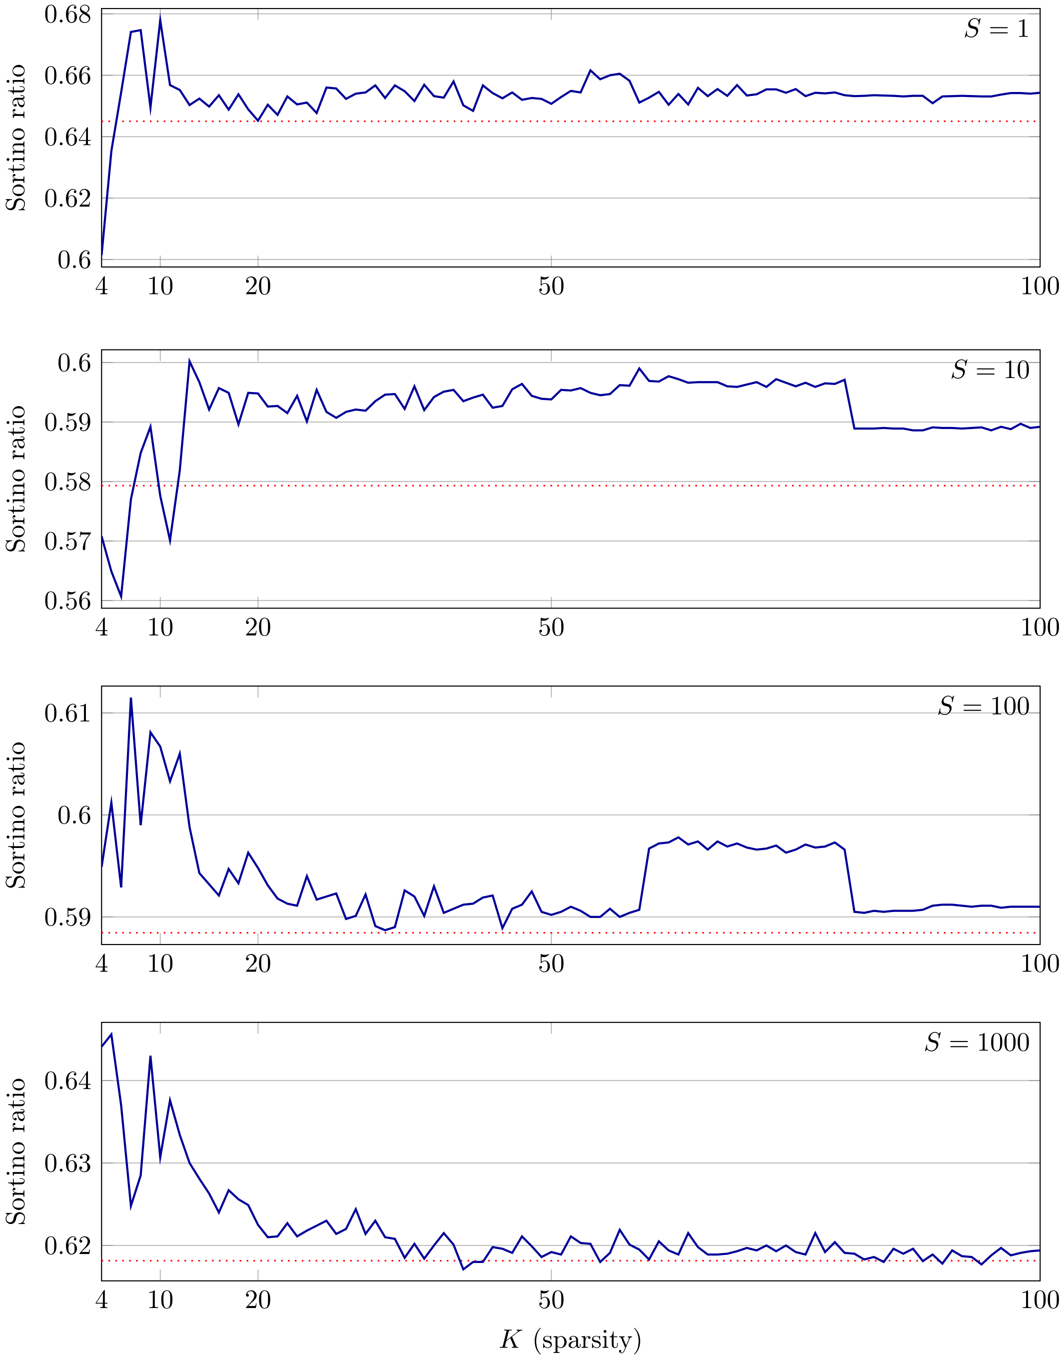

Notes: The red dotted line corresponds to the portfolio with no sparsity constraints.

Figure 10: Enhanced portfolio standard deviation by K and investor scale S

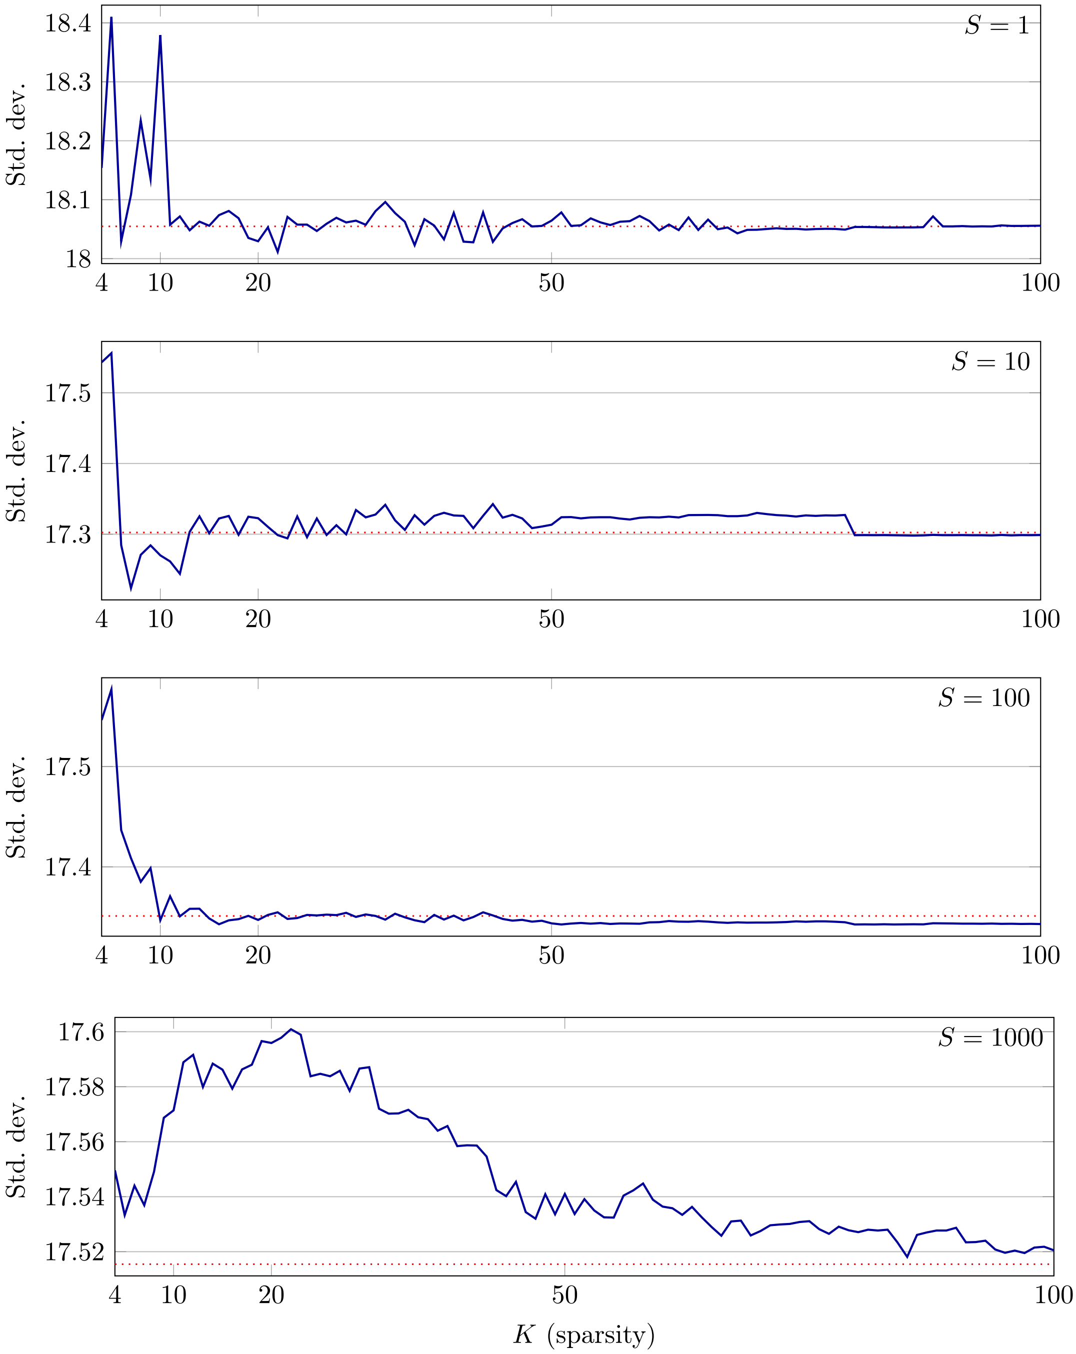

Notes: The red dotted line corresponds to the portfolio with no sparsity constraints.

Figure 11: Daily (cumulative) payoff functions ($K = 4$) for Aug 2004

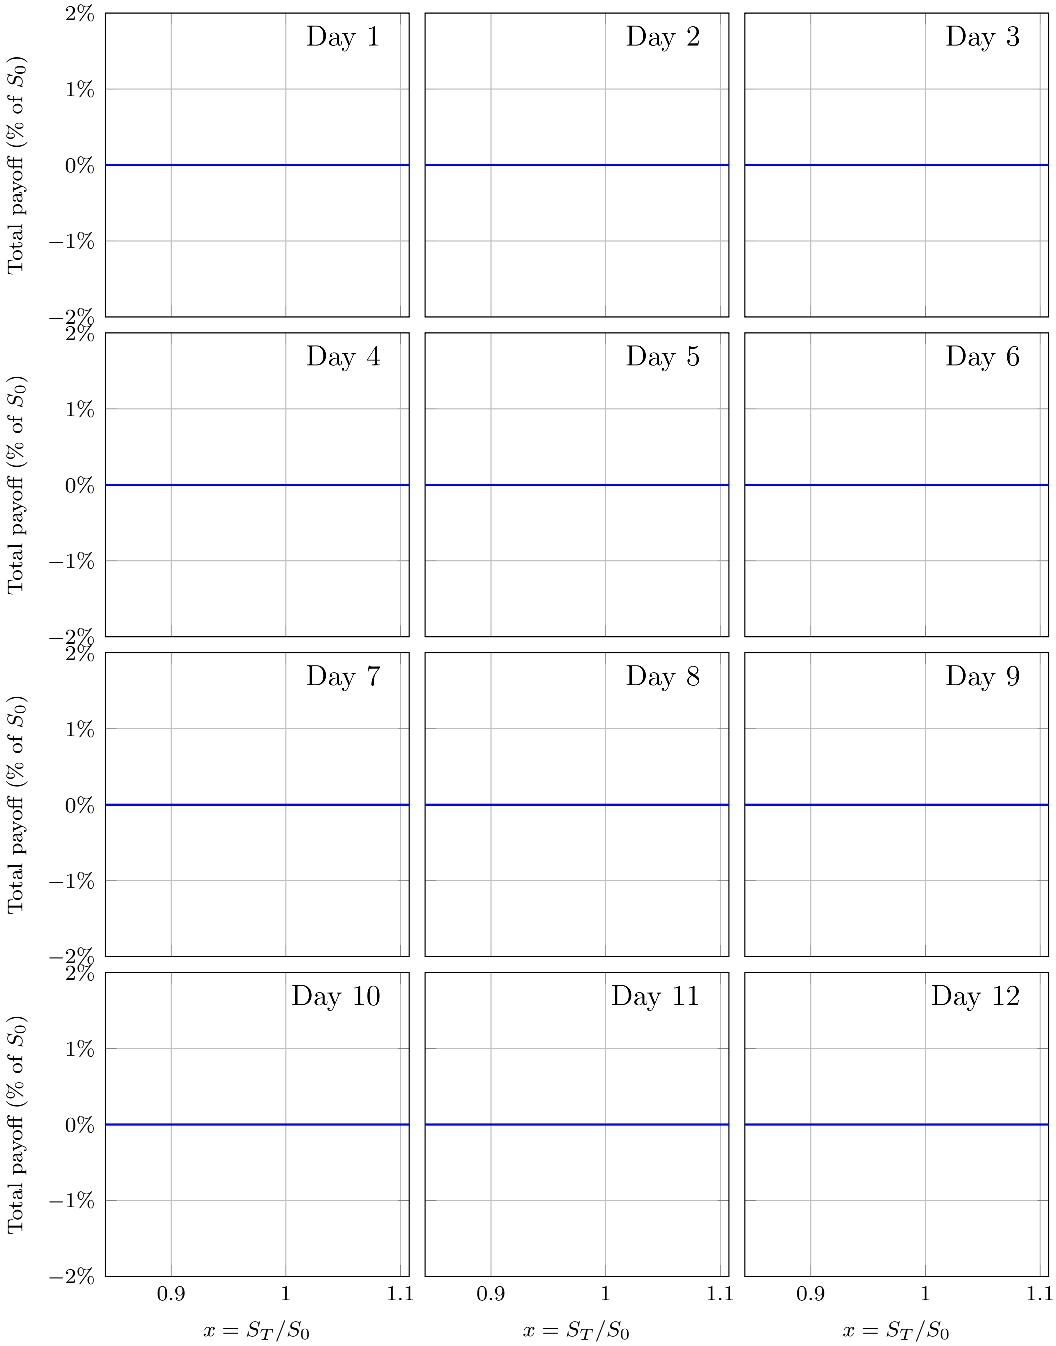

Figure 11 continues on next page

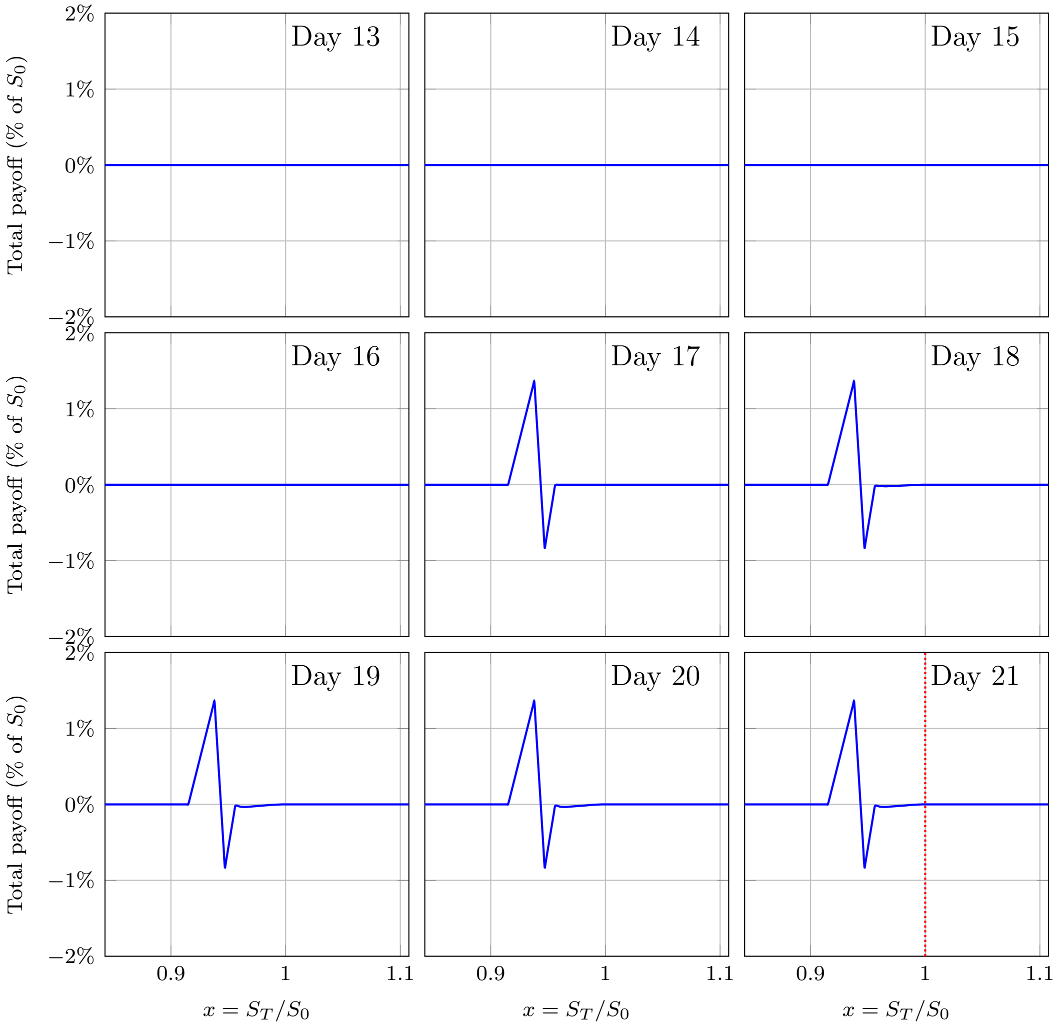

Notes. These payoff functions correspond the maximally constrained portfolio $K = 4$ for the trading month August 2004. S_T is the terminal SPX value and S_0 is the SPX value at the beginning of the month. The red dotted line corresponds to the terminal SPX value (as a ratio to month start value).

Figure 12: Daily (cumulative) payoff functions ($K = 100$) for Aug 2004

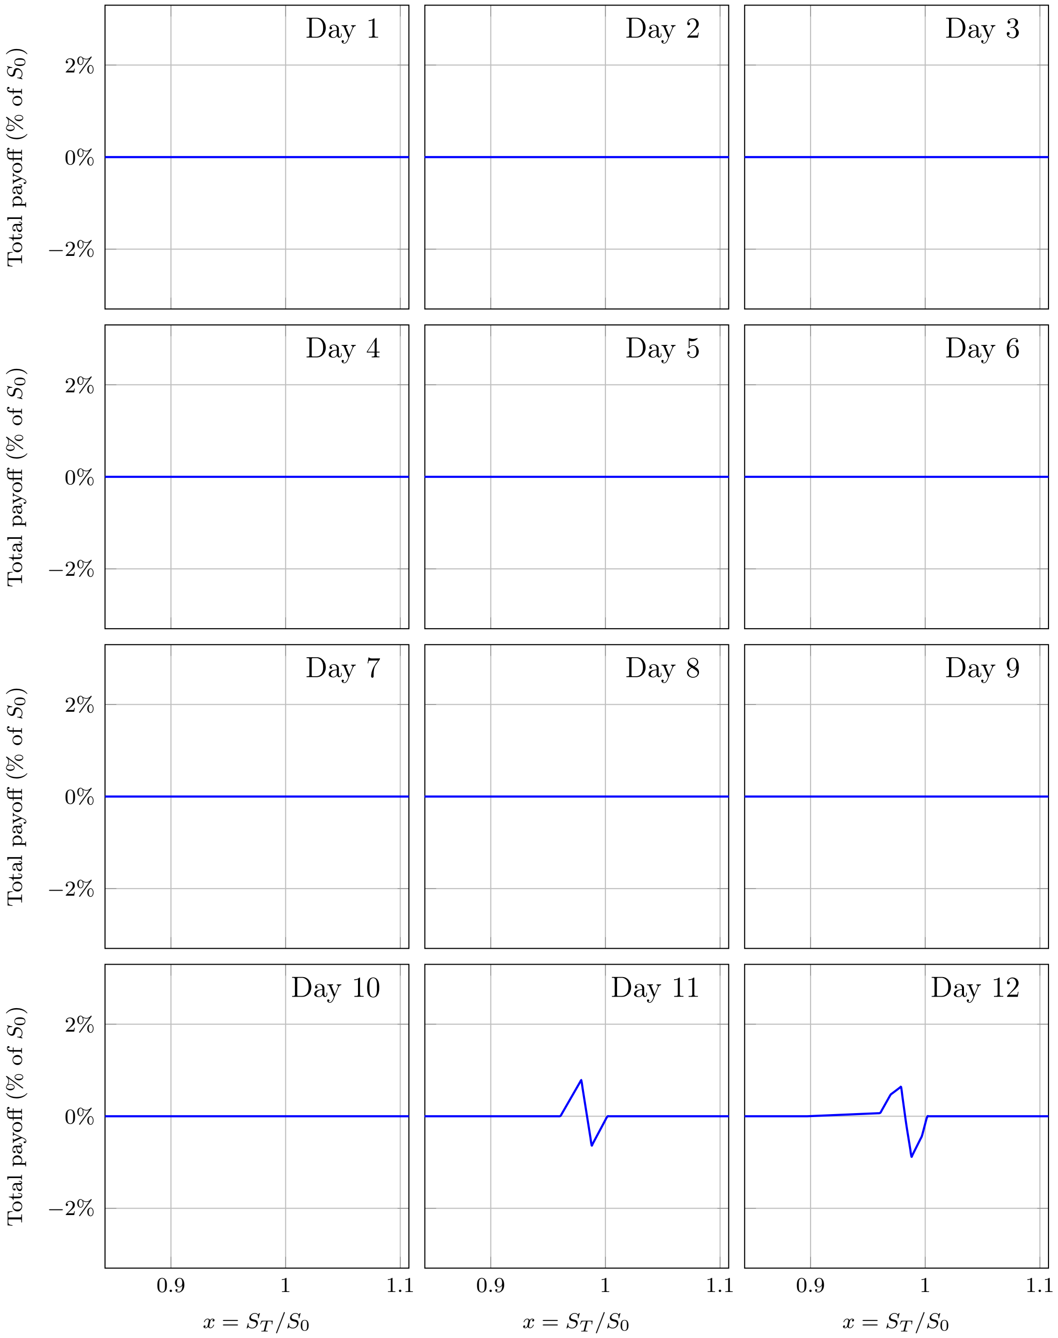

Figure 12 continues on next page

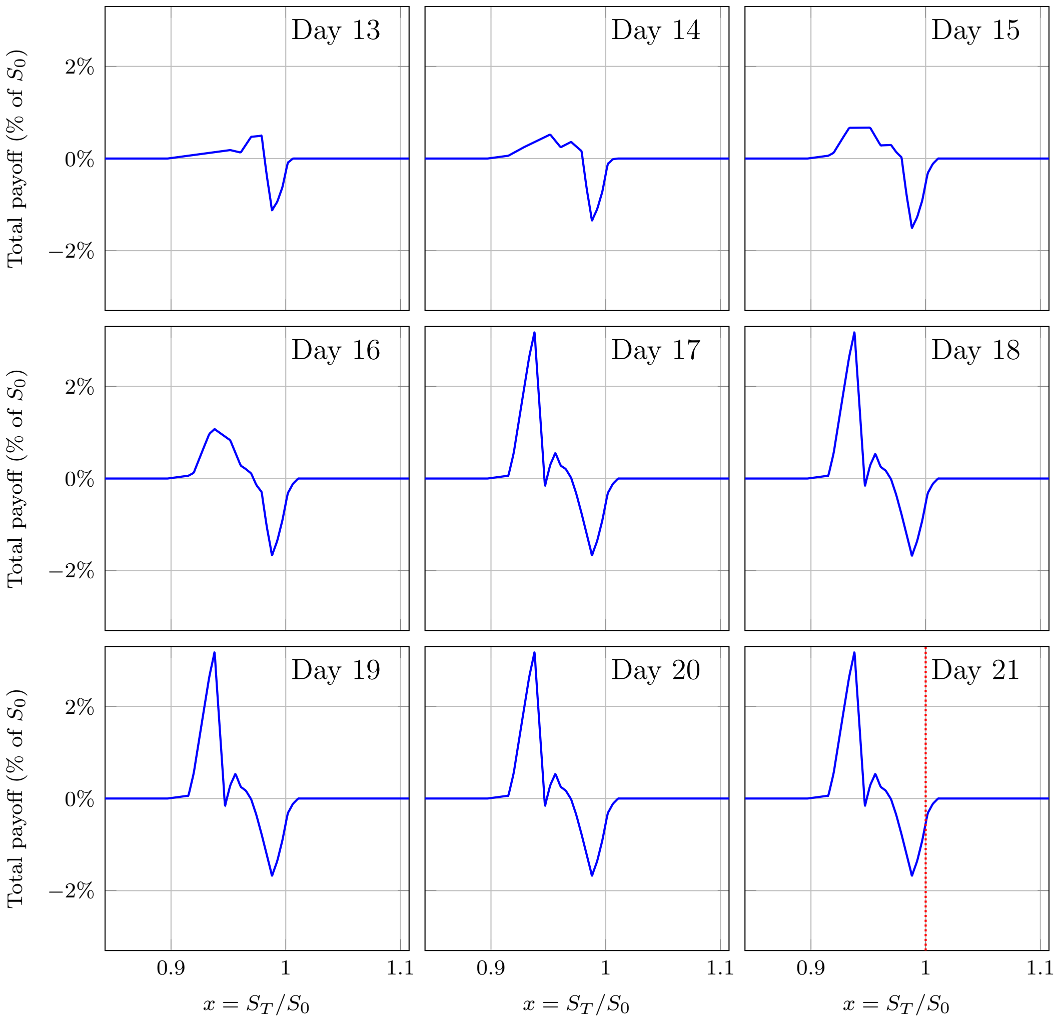

Table 6: Sparsity comparison, August 2004

Sparsity, K	SPX return (%)	Option payoff (S_0)	Premium (S_0)	Portfolio return (%)
4	0.21	0.0	0.30	0.23
100	0.21	-3.37	3.81	0.25

Notes: Portfolio return is SPX return plus option payoff and net premium. Option payoff and premium are in S_0 units.

Figure 13: Daily (cumulative) payoff functions ($K = 4$) for March 2020 (COVID-19)

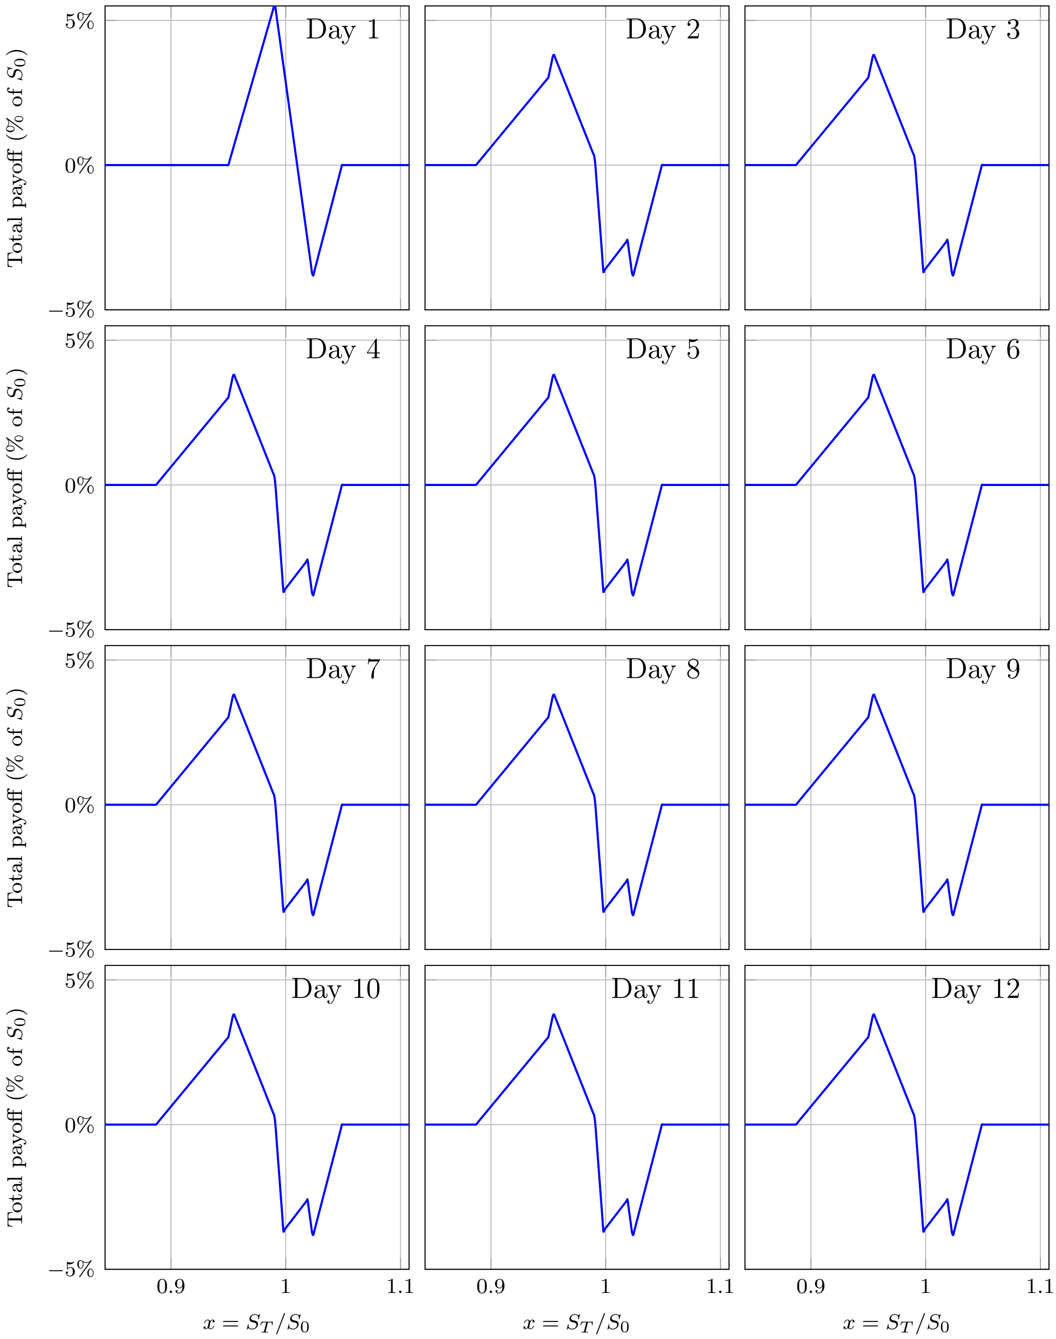

Figure 13 continues on next page

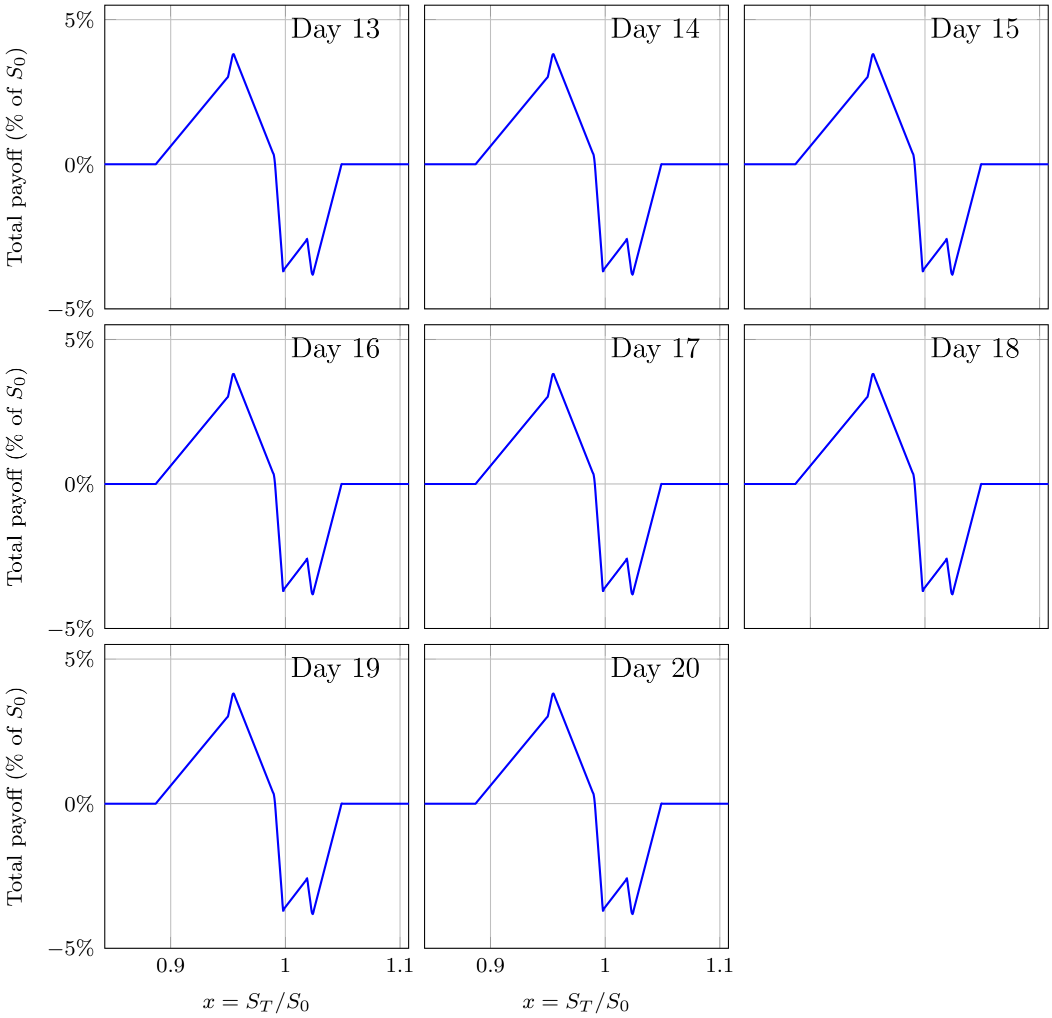

Notes. These payoff functions correspond the maximally constrained portfolio $K = 4$ for the trading month March 2020, corresponding to the COVID-19 pandemic. Note that there is no red line to specify the terminal SPX value because during this month the SPX began at 3332 and ended at 2327 a roughly 30% drop, which is far outside of the available strike range and hence, the zero profit constraints.

Figure 14: Daily (cumulative) payoff functions ($K = 100$) for March 2020 (COVID-19 crash)

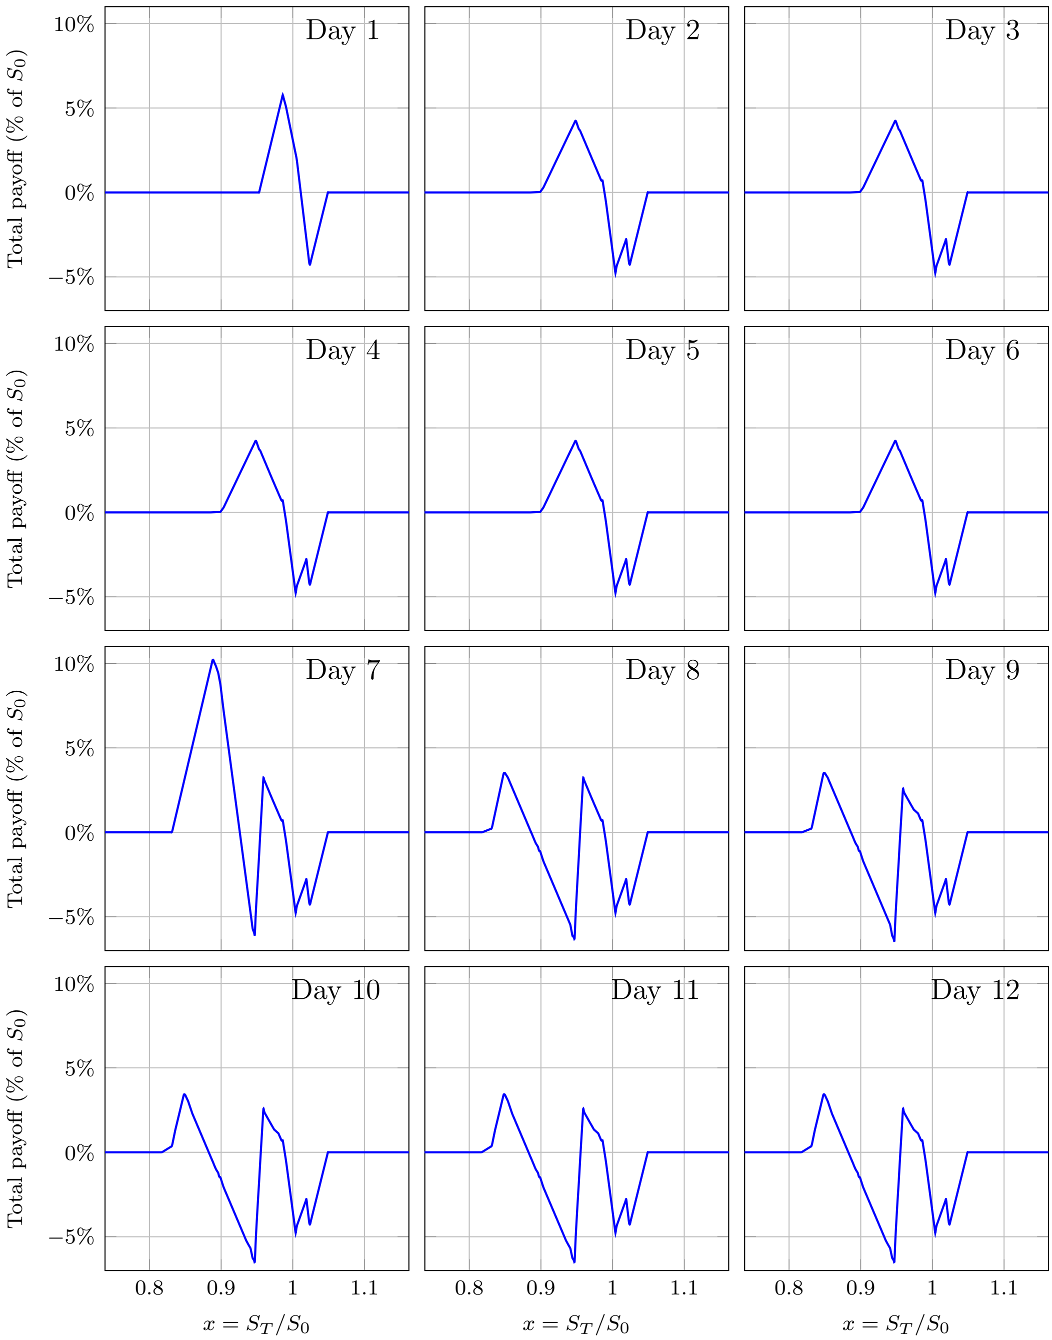

Figure 14 continues on next page

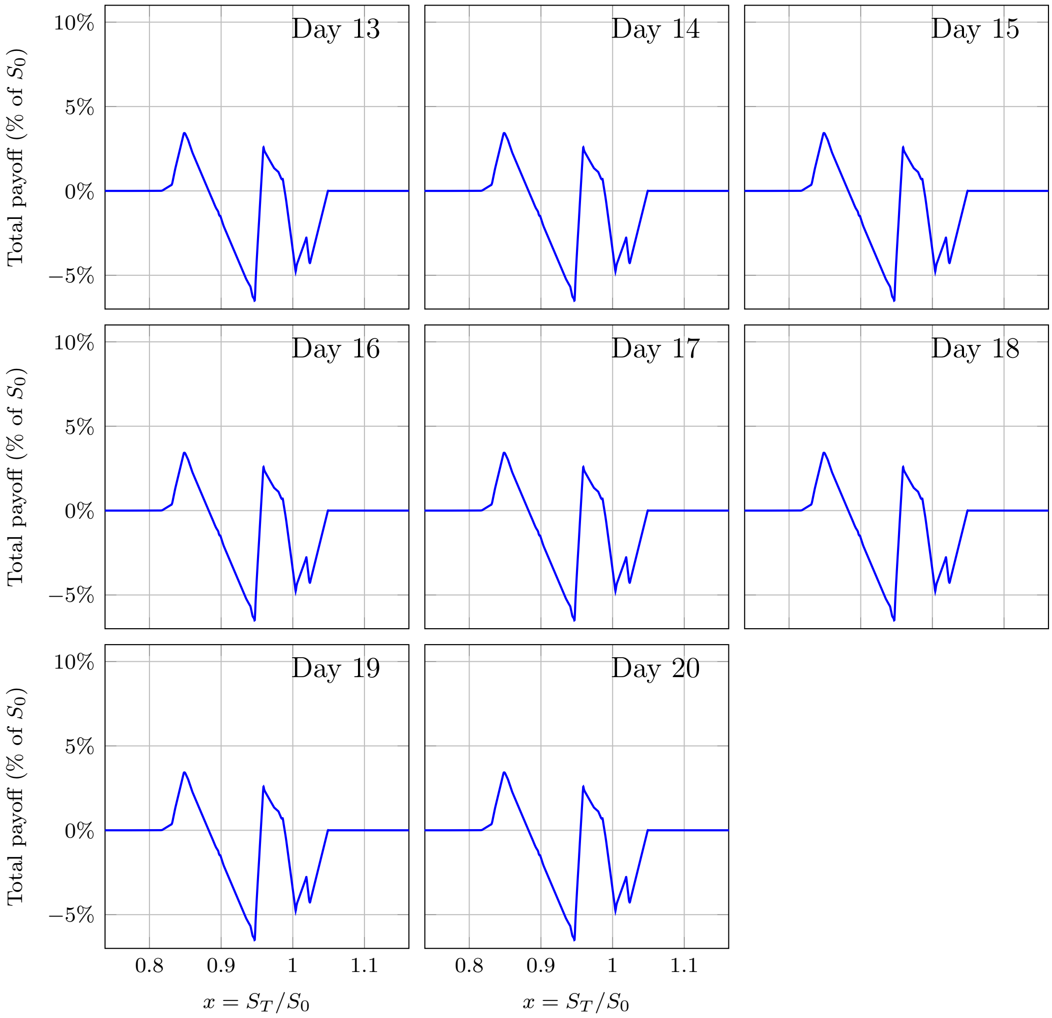

Table 7: Sparsity comparison, March 2020

Sparsity, K	SPX return (%)	Option payoff (S_0)	Premium (S_0)	Portfolio return (%)
4	-30.1	0.0	48.4	-28.7
100	-30.1	0.0	129.5	-26.3

Notes: Portfolio return is SPX return plus option payoff and net premium. Option payoff and premium are in S_0 units.

Figure 15: Enhanced **(static)** portfolio excess return by K and investor scale S