Options Portfolio Selection with Position Limits.

November 27, 2023

Abstract

This paper introduces a new method to maximize the Sharpe ratio of options' portfolios, using a constrained optimization approach that incorporates position limits, transaction costs, and volatility persistence. By regularizing portfolio weights, this approach effectively mimics the leverage limits imposed by margin requirements and mitigates price impacts from large trades in individual contracts. The out-of-sample performance of optimal portfolios of monthly options from 1996 to 2020 yields Sharpe ratios of index-neutral strategies between one and two for the S&P 500 and Nasdaq 100, but less than half for the Dow Jones. Constraining portfolios to be solvent on all past index' returns reduces Sharpe ratios by a third in the S&P 500 and Nasdaq 100 and by two thirds in the Dow Jones. All strategies suffer significant losses from the coronavirus shock of March 2020, underscoring the vulnerability of options' strategies to rare events.

Keywords: options, portfolio choice, position limits, quadratic optimization

MSC (2020): 91G10, 91G80.

JEL: G11, G12.

1 Introduction

The options' market offers significant investment opportunities,[1] but the optimization of options' portfolios is fraught with technical challenges.[2] Options' trading – unlike equities' — entails transaction costs of several percentage points even in the most liquid contracts, substantial margin requirements, and extreme correlation among contracts with similar strikes. Incorporating such market frictions is thus critical in both the construction and the evaluation of the performance and feasibility of options' strategies.

This paper presents a parsimonious method for selecting portfolios of options that maximize the Sharpe ratio, subject to position limits, and accounting for separate prices for buying and selling. Our method is akin to a regularization of portfolio weights, a technique used to mitigate estimation error in portfolio optimization in traditional asset classes, but hitherto unexplored for options' portfolios.[3] *A priori*, we constrain the norm of portfolio weights in order to mimic the leverage limits implied by both the margin requirements imposed by brokers and exchanges, and to avoid price impact from large trades in individual contracts. *A posteriori*, we find that such constraints also substantially reduce the impact of high collinearity in options' payoff in portfolio selection. In the absence of transaction costs, optimal portfolios admit an explicit solution, which shows that the constraint has a similar effect to the use of a shrinkage estimator for the covariance matrix, combined with an increase in risk aversion linked to the shrinkage weight. Explicit solutions become infeasible with transaction costs, but the above intuition remains valid and the problem retains its computational tractability, as every option contract is replaced in the optimization problem with two contracts, one for long and one for short positions, each of them commanding a separate price and constrained to be held in a positive quantity.

The contribution of this paper is twofold: First, our methodology is the first one to specify explicitly position limits in options' portfolio optimization: Santa-Clara and Saretto (2009)

[1]See Jackwerth (2000), Coval and Shumway (2001), Bakshi and Kapadia (2003), Jones (2006), Driessen and Maenhout (2007), Santa-Clara and Saretto (2009), Bondarenko (2014) and Schneider and Trojani (2015).

[2]Empirical studies that document significant risk-adjusted returns to options' strategies include Coval and Shumway (2001), Bakshi and Kapadia (2003), Santa-Clara and Saretto (2009), Eraker (2013), Driessen and Maenhout (2013), Schneider and Trojani (2015) and Faias and Santa-Clara (2017).

[3]For portfolio optimization with constraints and penalties on portfolio weights, see Ledoit and Wolf (2003), Jagannathan and Ma (2003), and DeMiguel et al. (2009). Korsaye et al. (2021) apply regularization to the estimation of pricing kernels.

find that margin requirements are a significant hurdle in executing ostensibly profitable options' strategies because they limit investors' positions. Faias and Santa-Clara (2017) mimic the effect of position limits by using a higher risk aversion for optimization than for performance evaluation. By including position limits directly in the optimization, we find that the tension between mean-variance efficiency and low portfolio weights not only reduces the scale of the overall portfolio, but it also equalizes weights across different strikes. In the small-position limit, which admits an explicit solution even with transaction costs, we see that the asymptotically optimal portfolio entails equal weights in all contracts for which the expected return is positive, and zero otherwise. As the constraint becomes more lenient, optimal options' portfolios typically imply higher weights on more extreme (but also more profitable) strike prices.

Second, we offer a novel – and fully replicable – empirical study of the performance of optimal options portfolios. Unlike previous studies, we include all call and put options with available strikes in the optimization, while limiting the overall options' position. Numerous studies document high Sharpe ratios in option-writing strategies, even controlling for exposure to known risk factors. A related strand of literature explores the optimization of portfolios including a limited number of options' contracts, such as one out-of-the-money (OTM) option and the index (Liu and Pan, 2003), at-the-money (ATM) straddles and OTM calls and puts (Eraker, 2013), two ATM and two OTM options (Faias and Santa-Clara, 2017). (Table 1.1 summarizes the empirical literature on options' performance.) Our approach makes optimization over contracts with all available strikes feasible by resolving collinearity issues, incorporating transaction costs, and imposing position limits. These features overcome the tendency of mean-variance optimization to generate large long-short positions in highly correlated securities with minor differences among in-sample returns.

In addition, while most empirical work focuses on options on the S&P 500 index and its futures contract (with the exception of Driessen and Maenhout (2013), considering also the FTSE 100 and Nikkei 225, and Malamud (2014), considering the Nasdaq 100 and Russel 2000), this paper also examines options on the Nasdaq 100 index and Dow Jones Industrial Average index, highlighting some important differences.

Options' strategies on the Nasdaq 100 yield even higher Sharpe ratios (up to 2.40) than those on the S&P 500 index (up to 1.87) after controlling for exposure to the respective indexes. By contrast, strategies in Dow Jones Industrial Average index options have much

Paper	Index(es)	Period	Options
Coval and Shumway (2001)	S&P 500 S&P 100	1990-01 – 1995-10 1986-01 – 1995-12	Options with strike prices 10 to 15 points below the index level to those with strike prices 5 to 10 points above the index.
Plyakha and Vilkov (2008)	S&P 100	1996-01 – 2004-12	Three moneyness and two maturity buckets for calls and puts.
Constantinides et al. (2013)	S&P 500	1986-04 – 2012-01	Calls or puts with one of nine target moneyness ratios.
Eraker (2013)	S&P 500	1996-01 – 2003-02	ATM straddles and OTM calls and puts.
Driessen and Maenhout (2013)	S&P 500 Nikkei 225 FTSE 100	1992-04 – 2001-06	OTM put and ATM straddle, each with 'crash-neutral' variant, shorting a deep OTM put.
Bondarenko (2014)	S&P 500	1987-01 – 2000-12	ATM puts, deep OTM puts.
Malamud (2014)	S&P 500 Nasdaq 100 Russel 2000	2004-01 - 2013-08	OTM strikes with moneyness in $[0.8, 1.2]$.
Faias and Santa-Clara (2017)	S&P 500	1996-01 – 2013-08	One ATM call, one ATM put, one 5% OTM call, one 5% OTM put.
Chan et al. (2021)	S&P 500	2017-05 – 2021-05	10% OTM calls compared with a 5% OTM call, a 5% OTM put and a 10% either-side strangle.
This paper	S&P 500 Nasdaq 100 Dow Jones	1996-01 – 2020-12	All non-zero bid calls and puts.

Table 1.1: Summary of the empirical literature on options' investing.

lower Sharpe ratios (up to 0.33), which do not significantly outperform their index' return over the same period.

The risk-adjusted returns of these strategies remain significantly positive even after enforcing the constraint that they remain solvent over any index return previously realized. However, such a constraint reduces the Sharpe ratios on S&P 500 and Nasdaq 100 options by about a third and on Dow Jones options by about two thirds, suggesting that a substantial component of options' returns is explained by the aversion of options' writers to the risk of large losses.

The rest of this paper is organized as follows: Section 2 describes the model and the optimization algorithm, including theoretical results that guarantee its convergence. (All proofs of are in the appendix.) Section 3 explains data analysis, volatility prediction and its use for historical simulation. Section 4 contains the empirical results, including the description of the estimation procedure.

2 Model and Method

Recall the familiar Markowitz setting, whereby an investor seeks the portfolio of d risky assets and one safe asset that maximizes the mean-variance tradeoff. Denoting by μ the vector of expected excess returns on risky assets and by Σ their covariance matrix, the risky portfolio weights $w = (w_1, \ldots, w_d)$ that maximize

$$w^\top \mu - \frac{\gamma}{2} w^\top \Sigma w,$$

where γ denotes the risk-aversion parameter, are given by the formula

$$\hat{w} = \frac{1}{\gamma} \Sigma^{-1} \mu.$$

Bringing this approach to bear on the construction of optimal options' portfolios raises multiple related issues: First, the mean-variance optimal portfolio is very sensitive to errors in the estimator of the covariance matrix (Best and Grauer, 1991), especially in combination with errors in the expected returns. For example, two options with nearby strikes are almost perfectly correlated, rendering a scenario-based estimator of Σ either singular or near-singular. As a result, any minor difference in their estimated expected returns is perceived by mean-variance optimization as a near-arbitrage opportunity, to be exploited with large long positions in one option, combined with large short positions in the other option.

Second, options incur high transaction costs (typically, from 5% for at-the-money options to 10% for out-of-the-money options), ignored by the familiar mean-variance objective, which make long-short positions very costly. Third, options portfolios entail substantial margin requirements, also ignored by the above objective, which *de facto* limit the size of long and short positions in options.[4] Fourth, though the market on index options is rather liquid, such

[4]Strategy-based options' margins are studied by Hitzemann et al. (2021) and Bali et al. (2023), who consider the specific rules described in `https://www.cboe.com/us/options/strategy_based_margin`. For the options' portfolios discussed in this paper the relevant margins would be calculated with portfolio-based methodologies, such as SPAN, TIMS, or STANS.

The first two methods compute margins as worst-case losses based on a prespecified set of scenarios. The STANS methodology, available to registered broker-dealers who are members of the Options Clearing Corporation, calculates a portfolio's margin as the 99% expected shortfall on a two-day horizon, using a Monte-Carlo simulation accounting for current market conditions. For customers' accounts, SEC rules require minimum margins that essentially follow the TIMS method, see `https://www.ecfr.gov/current/title-17/chapter-II/part-240/subpart-A/subject-group-ECFR541343e5c1fa459/section-240.15c3-1a`.

liquidity is spread over hundreds of individual contracts, each of them with limited depth and hence sensitive to price impact from large orders.

In summary, the challenge of options portfolio optimization is to depart enough from the familiar mean-variance objective to incorporate options' specific features, while at the same time retaining computational tractability.

2.1 Position Limits

To incorporate position limits in mean-variance optimization, we specify the objective function as

$$L(w, \lambda) = \mu^\top w - \frac{\gamma}{2} w^\top \Sigma w - \frac{\lambda}{2} w^\top w, \tag{2.1}$$

where $\lambda > 0$. The significance of this variant of the mean-variance objective is threefold. First, it can be interpreted as a robust version of the mean-variance tradeoff: suppose that the investor has limited confidence in the estimated expected return μ, and therefore evaluates the performance of a portfolio w as the minimum performance across the range of all possible views

$$\min_{\hat{\mu} \in \mathbb{R}^d} \left(w^\top \hat{\mu} - \frac{\gamma}{2} w^\top \Sigma w + \frac{1}{2\lambda} (\hat{\mu} - \mu)^\top \Omega^{-1} (\hat{\mu} - \mu) \right)$$

where the term $\frac{1}{2\lambda}(\hat{\mu} - \mu)' \Omega^{-1} (\hat{\mu} - \mu)$ downplays the negative performance under those views $\hat{\mu}$ that differ the most from the reference estimate μ. This expression reaches its minimum for $\hat{\mu} = \mu - \lambda \Omega w$, thereby yielding the formula

$$\mu^\top w - \frac{\gamma}{2} w^\top \Sigma w - \frac{\lambda}{2} w^\top \Omega w$$

which coincides with (2.1) in the baseline setting of $\Omega = I$ (I is the $d \times d$ identity matrix), whereby errors in all expected returns are considered equally important and uncorrelated with each other.

Second, the objective (2.1) can be understood as the unconstrained version of the con-

We are indebted to John Dodson for explaining the complexities of options' margins conventions. For details, see `https://www.cmegroup.com/clearing/risk-management/span-overview.html` (SPAN), `https://www.theocc.com/Risk-Management/Margin-Methodology` (TIMS), and `https://www.theocc.com/risk-management/customer-portfolio-margin` (STANS).

strained problem

$$\max_{w\in\mathbb{R}^d}\left\{\mu^\top w-\frac{\gamma}{2}w^\top\Sigma w : w^\top w=\sum_{i=1}^d w_i^2\le\frac{L^2}{d}\right\} \tag{2.2}$$

and coincides with the solution of the latter when the constraint is binding, for a suitable choice of the penalty λ. The upper bound on the sum of squared weights is conveniently specified as $\frac{L^2}{d}$, as it implies that L is an upper bound on the aggregate weights of options' portfolios, because

$$\sum_{i=1}^d |w_i|\le\sqrt{d}\sqrt{\sum_{i=1}^d w_i^2}=\sqrt{d\cdot w^\top w}\le L.$$

Thus, one can choose the parameter λ in relation to the bound L that it implies.

Third, as the objective (2.1) favors small portfolio weights in several contracts rather than large weights in few contracts, it provides a tractable approximation of the combined effect of margin requirements (without entering the minutiae of brokers' and exchanges' conventions) and the limited depth of options' markets. Indeed, the bound L above represents the maximum fraction of capital comprising the options' premia paid by the investor or to the investor. Conversely, the sum of squared weights discourages portfolios with large differences in the weights of nearby contracts (such as two options with the same expiration and similar strike prices), thereby mitigating the price impact that may result from large orders in the same contract.

The first-order condition of (2.1) identifies the optimal portfolio as

$$w=(\gamma\Sigma+\lambda I)^{-1}\mu. \tag{2.3}$$

The position limit then identifies the Lagrange multiplier $\lambda\in(0,\infty)$ as the solution to

$$\mu^\top(\gamma\Sigma+\lambda I)^{-2}\mu=\frac{L^2}{d}.$$

Thus, upon reflection, formulating position limits in terms of the objective (2.1) has another significant advantage: it is tantamount to the regularization of the covariance matrix estimator, combined with a change in risk aversion. This device remedies the typical near-singularity of the usual covariance estimator and links the regularization parameter to the position limit. Indeed, the optimal portfolio in (2.3) admits the equivalent expression

$$w = \frac{1}{\gamma+\lambda}\left(\frac{\gamma}{\gamma+\lambda}\Sigma + \frac{\lambda}{\gamma+\lambda}I\right)^{-1}\mu,$$

which means that when the constraint is binding, solving the constrained optimization problem (2.2) is equivalent to solving the unconstrained problem, while simultaneously (i) increasing the risk aversion from γ to $\gamma+\lambda$, and (ii) replacing the covariance matrix estimator Σ with the shrinkage estimator $\frac{\gamma}{\gamma+\lambda}\Sigma + \frac{\lambda}{\gamma+\lambda}I$, which is always strictly positive definite (cf. Ledoit and Wolf (2003); Jagannathan and Ma (2003); DeMiguel et al. (2009)).

In practice, the constraint binds under the following condition, which is satisfied for all realistic combinations of position limits and risk aversion.[5]

Proposition 2.1. Suppose $\mu \neq 0$.[6] Let σ_∞ the largest eigenvalue of Σ. The constraint $w^\top w \leq L^2/d$ is binding if

$$L < \frac{\sqrt{d}\|\mu\|}{\gamma\sigma_\infty}. \tag{2.4}$$

2.2 Bid-Ask Spreads

The previous formulation incorporates position limits in the portfolio optimization problem, but does not address transaction costs, which are very significant for options. To reflect the difference between bid and ask prices, we follow Plyakha and Vilkov (2008), Eraker (2013), and Faias and Santa-Clara (2017) by specifying each option contract as two distinct securities, to be held in positive amounts: one with the long-option payoff, priced at the ask quote, and one with the short-option payoff, priced at minus the bid quote.

With this representation, a portfolio optimization problem with d securities and transaction costs translates to an optimization problem with $2d$ securities but short-sale constraints instead of no transaction costs. That is, the portfolio weights $\bar{w}$ in the enlarged security set are linked to the portfolio weights w in original set by the relations $\bar{w}_{2i-1} = w_i^+$, $\bar{w}_{2i} = w_i^-$,

[5] Lemma A.1 in the Appendix characterizes the conditions under which a suitable value of λ satisfies the constraint. If the vector of expected returns lies in the range of the covariance matrix, then the constraint is binding for any value of L. Otherwise, it is binding for $L^2/d \leq \mu^\top\Sigma^{-2}\mu$.

[6] If all average excess returns are zero, then the optimal strategy is $w = 0$, whence the limit position cannot be binding.

and $w_i = \bar{w}_{2i-1} - \bar{w}_{2i}$. Thus, the optimization problem with bid-ask prices is:

$$\max_{\bar{w}\in\mathbb{R}^{2d}} \left\{ \bar{w}^\top \mu - \frac{\gamma}{2}\bar{w}^\top \Sigma \bar{w} : \bar{w} \geq 0, \bar{w}^\top \bar{w} \leq L^2/d \right\}. \tag{2.5}$$

The correspondence between w and $\bar{w}$ is one-to-one, provided that one restricts the attention to only those strategies which do not take a long and short position simultaneously.[7] Due to the positivity constraint, the solution to the optimization problem in (2.5) is not available in closed form, but is nonetheless easy to obtain numerically: if the covariance matrix Σ is invertible, the quadratic program with linear constraints

$$\max_{\bar{w}\in\mathbb{R}^{2d}} \left\{ \bar{w}^\top \mu - \frac{\gamma}{2}\bar{w}^\top \Sigma \bar{w} : \bar{w} \geq 0 \right\}$$

yields the solution $\tilde{w}$. If $\tilde{w}^\top \tilde{w} \leq L^2/d$, then the constraint is not binding, and $\tilde{w}$ solves (2.5). Otherwise, the constraint is binding, and it suffices to solve the quadratic program with linear constraints

$$\max_{\bar{w}\in\mathbb{R}^{2d}} \left\{ \bar{w}^\top \mu - \frac{\gamma}{2}\bar{w}^\top \Sigma \bar{w} - \frac{\lambda}{2}\bar{w}^\top \bar{w} : \bar{w} \geq 0 \right\} \tag{2.6}$$

for both some large λ_+, for which the corresponding $\tilde{w}_{\lambda_+}$ satisfies $\tilde{w}_{\lambda_+}^\top \tilde{w}_{\lambda_+} < \frac{L^2}{d}$ and some small (or zero) λ_-, for which $\tilde{w}_{\lambda_-}^\top \tilde{w}_{\lambda_-} > \frac{L^2}{d}$. Then, updating λ_-, λ_+ through a binary search yields the value of λ for which $\tilde{w}_\lambda^\top \tilde{w}_\lambda = \frac{L^2}{d}$.

Explicit solutions are available with bid-ask spreads in the small-position limit, i.e., for L close to zero. In this special case, both the covariance matrix Σ and the risk aversion γ become irrelevant, as they are overridden by the strict position limit, and the optimal portfolio is equally weighted on all options that offer a positive expected return:

Proposition 2.2. Suppose there exists $1 \leq i \leq 2d$ for which $\mu_i > 0$.[8] For small L, the constraint $w^\top w \leq L^2/d$ is binding and, denoting by $\mu_+ = (\max(\mu_i, 0))_{1\leq i\leq 2d}$, the optimal

[7] Such a restriction is obviously inconsequential for a monotonic objective such as expected utility. For the mean-variance objective under consideration, simultaneous long and short positions are also suboptimal, as confirmed by Lemma A.2 below. Note that, for this reason, the constraint L^2/d does not to be replaced by $L^2/(2d)$.

[8] If all average excess returns are non-positive, then the optimal strategy is $w = 0$, whence the limit position cannot be binding.

portfolio satisfies

$$\tilde{w} = \frac{L}{\sqrt{d}} \frac{\mu_+}{\|\mu_+\|} + o(L). \tag{2.7}$$

2.3 Solvency Constraints

The mean-variance objective in (2.5) does not exclude potentially negative payoffs, therefore it may lead to in-sample insolvency, that is, on some of the outcomes included in the empirical distribution used for optimization. To understand how insolvency affects the performance of options' portfolios, it is useful to examine the performance of optimal options portfolios with in-sample solvency constraints, i.e.,

$$\max_{\bar{w} \in \mathbb{R}^{2d}} \left\{ \bar{w}^\top \mu - \frac{\gamma}{2} \bar{w}^\top \Sigma \bar{w} : \bar{w} \geq 0,\ \bar{w}^\top \bar{w} \leq \frac{L^2}{d},\ \bar{w}^\top R_i \geq -1,\ 1 \leq i \leq S \right\},$$

where $(R_i)_{1 \leq i \leq S}$ represents the sample of size S of all the returns used to estimate μ and Σ.[9] The additional constraint $\bar{w}^\top R_i \geq -1$ thus ensures that the optimizer yields positive wealth for all outcomes in the sample, reflecting the approach of an investor who requires a strategy to be solvent under any potential past return.[10]

3 Data

The analysis of options' portfolios performance relies on OptionMetrics' IvyDB database on European options on the S&P 500,[11] Nasdaq 100, and Dow Jones Industrial Average equity indexes from January 1996 to December 2020. The dataset includes daily closing bid and ask prices for each monthly option with the nearest expiration, daily closing prices of the indexes, settlement price of the indexes (used to calculate options' payoffs), and the daily

[9] While a priori each of the constraints $R_i \geq 0$ may be binding, a posteriori only $R_{\min} = \min_{1 \leq i \leq S} R_i$ and $R_{\max} = \min_{1 \leq i \leq S} R_i$ turn out to be binding.

[10] A priori, insolvency is possible out-of-sample, that is, on outcomes that are not included in the empirical distribution, unless the portfolio's payoff has no net short option position. A posteriori, the strategies considered remained always solvent.

[11] In the OptionMetrics' IvyDB database here are also quotes for SPX weekly options (SPXW), a new option different from SPX options but having the same Security ID. They are removed form the data frame to avoid repetition and because they settle at the closing price instead of the special opening quotation.

term structure of interest rates.[12]

3.1 Options' Filters

To ensure that optimal portfolios trade off return against risk, rather than pursuing spurious arbitrage opportunities arising from asynchronous or stale data, options are filtered by removing (i) observations with zero bids, (ii) bid-ask pairs where the bid is higher than the ask, (iii) bid-ask pairs that violate put-call parity, adjusted for transaction costs and dividends, and (iv) bid-ask pairs with zero or missing volume or open interest. Numerous studies document parity violations (Phillips and Smith Jr, 1980; Baesel et al., 1983; Santa-Clara and Saretto, 2009) and attribute them to the large transaction costs in options. The filter described below ensures that all remaining options' pairs satisfy the restrictions implied by put-call parity, after controlling for the effects of bid-ask prices and dividends (Van Binsbergen et al., 2012).

Recall the familiar parity relation $C - P = S - Ke^{-rT} - D_T$ among the price of a call option C and a put option P with the same strike price K and expiration T, when the safe rate is r, the price of the index S, and the present value of the dividend D_T (Stoll, 1969). (The estimation of dividends is discussed in the next subsection.) This relation holds when bid-ask spreads are absent: when they are present, the equality is replaced by the two simultaneous inequalities

$$C_{\text{bid}} - P_{\text{ask}} < S - Ke^{-rT} - D_T < C_{\text{ask}} - P_{\text{bid}},$$

which are equivalent to the restriction

$$C_{\text{bid}} - P_{\text{ask}} + Ke^{-rT} + D_T < S < C_{\text{ask}} - P_{\text{bid}} + Ke^{-rT} + D_T \tag{3.1}$$

on the index price. These inequalities result from absence of arbitrage, i.e., from the condition that there are no gains from riskless replications of long or short forward payoffs through options. Because (3.1) must hold simultaneously for all strike prices $(K^i)_{1\leq i\leq k}$, the joint

[12] The settlement price of monthly options is not the closing price of the index on the expiration date, but the special opening quotation (SOQ), which is calculated from the opening prices of the index' components.

Index	No Removal	Average Removal
S&P 500	96.33%	6.06%
Nasdaq 100	97.33%	9.76%
Dow Jones	99.64%	6.56%

Table 3.1: Percentage of months for which no options are removed due to put-call parity violations (second column) and average percentage of options removed (third column) when necessary, for each index (first column).

restriction is equivalent to

$$\max_{1\le i\le k}\left(C^i_{\text{bid}} - P^i_{\text{ask}} + K^i e^{-rT} + D_T\right) =: S_L < S < S_U := \min_{1\le i\le k}\left(C^i_{\text{ask}} - P^i_{\text{bid}} + K^i e^{-rT} + D_T\right). \tag{3.2}$$

This condition is a restriction on both the bid-ask prices of options' contracts, and on the price of the index. In the data, asynchronous observations of options' quotes and the closing price may cause such condition to fail: such failure is ascribed to the options' quotes if $S_L > S_U$, in which case no choice of S is compatible with the absence of arbitrage, or to the closing price if it lies outside (S_L, S_U), in which case it suffices to replace the closing price with the midpoint of the arbitrage-free interval (S_L, S_U).

These observations motivate the following procedure to eliminate parity violations:

(i) If $S_L \le S_U$ in (3.2) holds, then set S equal to the closing price if it lies in (S_L, S_U), or to the midpoint $(S_L + S_U)/2$ otherwise.

(ii) Else (if $S_L > S_U$) remove both options pairs with both the strikes K_l and K_u for which the lower and upper bounds S_L, S_U are achieved in (3.2) (thereby widening the arbitrage-free interval). Then, return to step (i).

Table 3.1 shows that no options need to be removed for over 95% of the months (second column). When removal is necessary, it averages from 5% to 10% of the strike prices available (third column).

3.2 Dividends

The underlying indexes considered do not include the dividends paid by their components to investors, which means that the put-call parity relation requires the dividend-adjustment discussed above, hence the estimation of dividends. The dividend yield on these indexes is typically small (an annual dividend yield of 1.9% from 1996 to 2007, 3.11% in 2008, and 1.97% from 2009 to 2020) and the dividends are usually announced one month before being paid, so their value is known with high accuracy at the time that the options' portfolio is set up. Thus, exclusively for the purpose of excluding options leading to parity violations, it is appropriate to estimate the value of the dividend from the difference in the returns of the index (which excludes dividends) and the exchange-traded fund (ETF) that tracks it (which includes dividends). Thus,[13]

$$R_T^{\text{index}} = \frac{S_T^{\text{index}} - S_0^{\text{index}}}{S_0^{\text{index}}}, \quad R_T^{\text{ETF}} = \frac{S_T^{\text{index}} - S_0^{\text{index}} + D_T}{S_0^{\text{index}}} = R_T^{\text{index}} + \frac{D_T}{S_0^{\text{index}}},$$

where $S_0^{\text{index}}, S_T^{\text{index}}$ denote the values of the ex-dividend index at the beginning (portfolio construction) and the end (options' expiration) of the trading period. Thus, the above relation yields the dividend estimate $D_T = (R_T^{\text{ETF}} - R_T^{\text{index}})S_0^{\text{index}}$, which is used to exclude put-call parity violations in options' data.

3.3 Volatility Forecasts

The prices of options reflect both their risk premia and the market's expectations about future realized volatility until the options' expiration. As volatility is highly persistent from one month to the next, it is thus critical for an options' investor to forecast realized volatility at the beginning of a trading period, so as to estimate the expected returns of options with different strikes. This paper adopts the specification

$$\log(\sigma_{t,T}) = \log(\alpha) + \beta_{\text{imp}} \log(\sigma_t^{\text{imp}}) + \beta_{\text{hist}} \log(\sigma_t^{\text{hist}}) + \varepsilon_{t,T}, \tag{3.3}$$

where $\sigma_{t,T}$ is the realized volatility over the next trading period and $\varepsilon_{t,T}$ is the forecasting error, while σ_t^{imp} and σ_t^{hist} are respectively the index' implied and historical volatility at the beginning of the period.

The specification in (3.3) is linear in the logarithms of volatilities rather than volatilities

[13] In principle, the ETF return also reflects the management fee deducted by the fund's issuer. As such fees are less than 0.1% per year while options expire in one month, the effect of fees is negligible and not included in dividends' estimation.

themselves both to avoid negative forecasts, and because the unconditional distribution of volatility is closer to a lognormal than to a normal distribution. Indeed, (3.3) is equivalent to the mixture of powers

$$\sigma_{t,T} = \alpha(\sigma_t^{\text{imp}})^{\beta_{\text{imp}}}(\sigma_t^{\text{hist}})^{\beta_{\text{hist}}} e^{\varepsilon_{t,T}}.$$

By construction, the logarithmic regression in (3.3) yields the unbiased estimator of $\log(\sigma_{t,T})$

$$\log(\hat{\sigma}_{t,T}) = \log(\hat{\alpha}) + \hat{\beta}_{\text{imp}} \log(\sigma_t^{\text{imp}}) + \hat{\beta}_{\text{hist}} \log(\sigma_t^{\text{hist}})$$

and the corresponding unbiased estimator of realized variance follows.

Proposition 3.1. If the forecasting errors $\varepsilon_{t,T}$ are conditionally (with respect to σ_t^{imp} and σ_t^{hist}) normal and homoskedastic, then

$$\tilde{\sigma}_{t,T} = \hat{\sigma}_{t,T} e^{(1-R^2)\,\text{Var}(\log \hat{\sigma}_{t,T})}$$

is an unbiased estimator of realized variance, in that $E[\tilde{\sigma}_{t,T}^2 - \sigma_{t,T}^2 | \sigma_t^{\text{imp}}, \sigma_t^{\text{hist}}] = 0$.

In the volatility forecast, σ^{imp} is the implied volatility index, on the day before trading, for the equity index considered (VIX for the S&P 500, VXN for the Nasdaq 100, and VXD for the Jow Jones), once such index has been available for at least two months. Before, VIX is used instead. Likewise, σ^{hist} is the historical volatility, also on the day before trading, obtained from the Oxford Man daily realized volatility database, once such index has been available for at least two months. Before, the volatility estimator from daily returns on the past monthly period is used instead. Table 3.3 describes in detail the use of data for portfolio optimization, including volatility forecasting.

Table 3.2 reports the model parameter estimates for the entire dataset, i.e., the parameters that would be used to forecast the volatility at the end of 2020 for the two different historical volatility estimators. When daily realized volatility is used (right column), the volatility estimator is approximately a weighted geometric mean, in that the weights for the implied and historical volatilities are both positive, and their sum is close to one. Instead, when daily returns are used to estimate historical volatility, the implied probability weight rises above one to compensate for the loss in significance of the historical probability weight. The scaling constant also decreases, as to offset the effect of the increased implied volatility weight, which would otherwise overestimate future realized volatility.

$$\log(\sigma_{t,T}) = \log(\alpha) + \beta_{\text{imp}} \log(\sigma_t^{\text{imp}}) + \beta_{\text{hist}} \log(\sigma_t^{\text{hist}}) + \varepsilon_{t,T}$$

	Daily returns	Oxford-Man
(Intercept)	-0.564^{***}	-0.320^{*}
	(0.138)	(0.171)
$\log(\sigma_t^{\text{imp}})$	1.132^{***}	0.901^{***}
	(0.086)	(0.096)
$\log(\sigma_t^{\text{hist}})$	-0.043	0.149^{**}
	(0.063)	(0.064)
No. Observations	371	251
R^2	0.624	0.656

Table 3.2: Volatility forecast (3.3) on the last trading date 2020-12-21, with maximal historical data available. Significance codes: *(0.1), **(0.05), ***(0.01).

3.4 Returns' Distribution

While volatility controls the time-varying scale of the distribution of returns, its location and shape are captured by the normalized empirical distribution, as follows.

At the beginning of each trading period, first consider the set $\mathcal{R}$ of all returns, realized before the trading date, over 21 consecutive business days. That is, $R_i = S_i/S_{i-21} - 1$, $21 < i \leq t$, where t is the last business day before trading. (Data availability varies by index and is reported in Table 3.3.) Each return $R_i \in \mathcal{R}$ is then standardized, by calculating

$$\tilde{R}_i = \frac{R_i - \bar{\mu}\Delta_i}{\bar{\sigma}\sqrt{\Delta_i}},$$

where $\bar{\mu}$ and $\bar{\sigma}$ are the annualized mean return and volatility of the index in the whole period and Δ_i is the number of calendar days spanning the return R_i, divided by 360. The set of normalized returns is $\tilde{\mathcal{R}} = (\tilde{R})_{i=1}^n$. (Because 21 business days lead to a variable number of calendar days, this adjustment rescales different calendar periods accordingly.)

At the beginning of each period, the investor estimates the distribution of the index' return over the next period by considering equally likely the samples

$$\hat{R}_i = \tilde{R}_i \tilde{\sigma}_{[t,T]} \sqrt{\Delta} + \bar{\mu}\Delta$$

where $\tilde{\sigma}_{[t,T]}$ is the volatility estimator in the previous subsection, while Δ denotes the number

Table 3.3: Data Sources

<table>
<tr><th></th><th>SPX</th><th colspan="2">NDX</th><th colspan="2">DJX</th></tr>
<tr><td rowspan="2">Sample Return</td><td rowspan="2">S&P 500 (^SPX)
1927:12 – 2020:12
Yahoo Finance</td><td>Nasdaq 100 (^NDX)
1985:10 – 2020:12
Yahoo Finance</td><td>Nasdaq Composite (^IXIC)
1971:02 – 1985:09
Yahoo Finance</td><td colspan="2" rowspan="2">Dow Jones Industrial (^DJI)
1927:12 – 2020:12
Bloomberg</td></tr>
<tr><td colspan="2">Nasdaq 100 available from 1985.
Nasdaq Composite included from 1971 to 1985.[a]</td></tr>
<tr><td rowspan="2">Dividend Adjustment</td><td rowspan="2">SPDR S&P 500
ETF Trust (SPY)
1996:01 – 2020:12
Yahoo Finance</td><td colspan="2">Invesco QQQ ETF Trust (QQQ)
1999:03 – 2020:12
Yahoo Finance</td><td colspan="2">SPDR Dow Jones Industrial
Average ETF Trust (DIA)
1998:01 – 2020:12
Yahoo Finance</td></tr>
<tr><td colspan="2">No dividend adjustment for parity from 1996:01 to 1999:02.</td><td colspan="2">No dividend adjustment for parity from 1997:10 to 1997:12.</td></tr>
<tr><td rowspan="2">Implied Volatility</td><td rowspan="2">CBOE Volatility
Index (^VIX)
1990:01 – 2020:12
Yahoo Finance</td><td>CBOE Volatility
Index (^VIX)
1990:01 – 2001:04
Yahoo Finance</td><td>CBOE Nasdaq 100
Volatility Index (^VXN)
2001:01 – 2020:12
Yahoo Finance</td><td>CBOE Volatility
Index (^VIX)
1990:01 – 1998:01
Yahoo Finance</td><td>CBOE DJIA Volatility
Index (^VXD)
1997:10 – 2020:12
Bloomberg</td></tr>
<tr><td colspan="2">VIX used for two months before VXN available.
Then, VXN backfilled from linear regression of VIX,
estimated in overlapping period.</td><td colspan="2">VIX used for two months before VXD available.
Then, VXD backfilled from linear regression of VIX,
estimated in overlapping period.</td></tr>
<tr><td>Historical Volatility</td><td colspan="5">Realized Variance (5-minutes intervals)
2000:01 – 2020:12
Oxford Man Institute of Quantitative Finance (rv5)
Standard deviation of index' daily returns over previous month from 1996:01 to 1999:12.</td></tr>
<tr><td>Option Price</td><td colspan="3">1996:01 – 2020:12
OptionMetrics</td><td colspan="2">1997:10 – 2020:12
OptionMetrics</td></tr>
</table>

[a] In the overlapping period 1985–1995, the regression of Nasdaq 100 monthly returns against the Nasdaq Composite yields an alpha of 20 basis points, a beta of 1.14, and $R^2 = 0.9271$, corresponding to a correlation of 96.3%.

of calendar days spanning the next return, divided by 360. Thus, $\hat{\mathcal{R}} = (\hat{R}_i)_{i=1}^n$ is the empirical distribution used to evaluate the expected performance of options' portfolios over the next period.

4 Results

This section discusses the performance of mean-variance optimal options' portfolios on the monthly options on the S&P 500, Nasdaq 100 and Dow Jones Industrial Average equity indexes for the 25-year period 1996:01-2020:12 (1997:10-2020:12 for the Dow Jones, see Table 3.3). Options on these indexes are European, hence can be exercised only at their expiration.

All the empirical results present of the paper are out-of-sample, i.e., the options' portfolio is constructed at each period using only data available before such period. In particular, at each period, only past index' returns are used to forecast volatility (including the parameters of the volatility model), to estimate the expected returns and covariance matrix of the options' portfolio (hence its weights), and even to estimate dividends. In-sample estimates are never used for portfolio construction throughout the paper.

Each month, on the Monday following monthly options' expiration[14] we construct a portfolio of options expiring the following month through the procedure described in the previous section: First, forecast realized volatility over the following month, obtaining a rescaled empirical distribution of returns. Second, from this empirical distribution and current options' prices, calculate options' expected returns and covariances. Third, combine expected returns and covariances with the risk-aversion and position-limit parameters in the constrained optimization problem, obtaining the optimal portfolio weights. Fourth, record the portfolio performance at the end of the month from the index' ex-post return. (To compute excess returns, the term-structure of interest rates in OptionMetrics is interpolated linearly.) To examine the effect of solvency requirements, the same optimization is also performed with the additional constraint that the portfolio remain in-sample solvent (i.e., over all scenarios in the empirical distribution).

Three position limits L on the total portfolio weights of long and short options premia are considered, corresponding to 3%, 5%, and 10% of the portfolio value each month, combined with risk aversions of 1, 3, and 5, yielding nine benchmark options' portfolios. These

[14] The third Friday of the month, or Thursday is Friday is a holiday, such as Good Friday.

parameters are chosen to obtain portfolios with annual volatility in the typical range of 10% to 20%. Performance analysis includes the annualized Sharpe ratios, alphas, and betas from regressions of portfolios' excess returns against the index's excess return, as well as the standard deviation of the hedged (i.e., zero beta) portfolio and its Sharpe ratio (i.e., the Appraisal ratio).[15]

4.1 S&P 500

Figure 4.1 displays the cumulative return on a hypothetical portfolio starting with a dollar at the beginning of 1996, reinvested in the mean-variance optimal strategies with risk aversion $\gamma = 3$ and various position limits, with (dashed) and without (solid) the solvency constraint, in comparison with the index' performance (black). For example, a large loss is observed on all options' portfolios in correspondence of the coronavirus crisis of March 2020: as the options' portfolios typically sell significant amounts of put options, as shown below, they are vulnerable to large losses on extreme events. Overall, all portfolios significantly outperform the index, though their risk is concentrated in large infrequent drawdowns rather than dispersed among the small frequent losses typical of the index. Looser position limits lead to higher returns but imply commensurate losses in unfavorable states. The solvency constraint is a significant drag on portfolio performance, as each dashed line remains always below the solid line with the same color, even immediately after rare large losses.

Table 4.1 assesses quantitatively the strategies' performance: note that the most significant determinant of strategies' volatility is the position limit L, indicating that the constraint is indeed binding for the parameters' combinations considered. Without the solvency constraints (top panel), the Sharpe ratios are all in the range 0.93 to 1.90, clearly outperform the index', but some component of such performance is explained by the positive exposure (beta) to the index. Controlling (i.e., hedging) for such exposure, the residual Sharpe ratio (the Appraisal ratio) is from 0.83 to 1.87, with the top of the range achieved for the loosest position limit. The bottom panel of Table 4.1 shows that the inclusion of solvency

[15] Denoting by R_o the excess return of the options' portfolio and by R the excess return of the index, the variance of the beta-neutral portfolio is then

$$\sigma_0^2 = \mathrm{Var}(R_o - \beta R) = \mathrm{Var}(R_o) - \beta^2 \mathrm{Var}(R).$$

Accordingly, the Appraisal ratio is $\frac{\alpha}{\sigma_0^2}$, where α is the alpha of the options' portfolio.

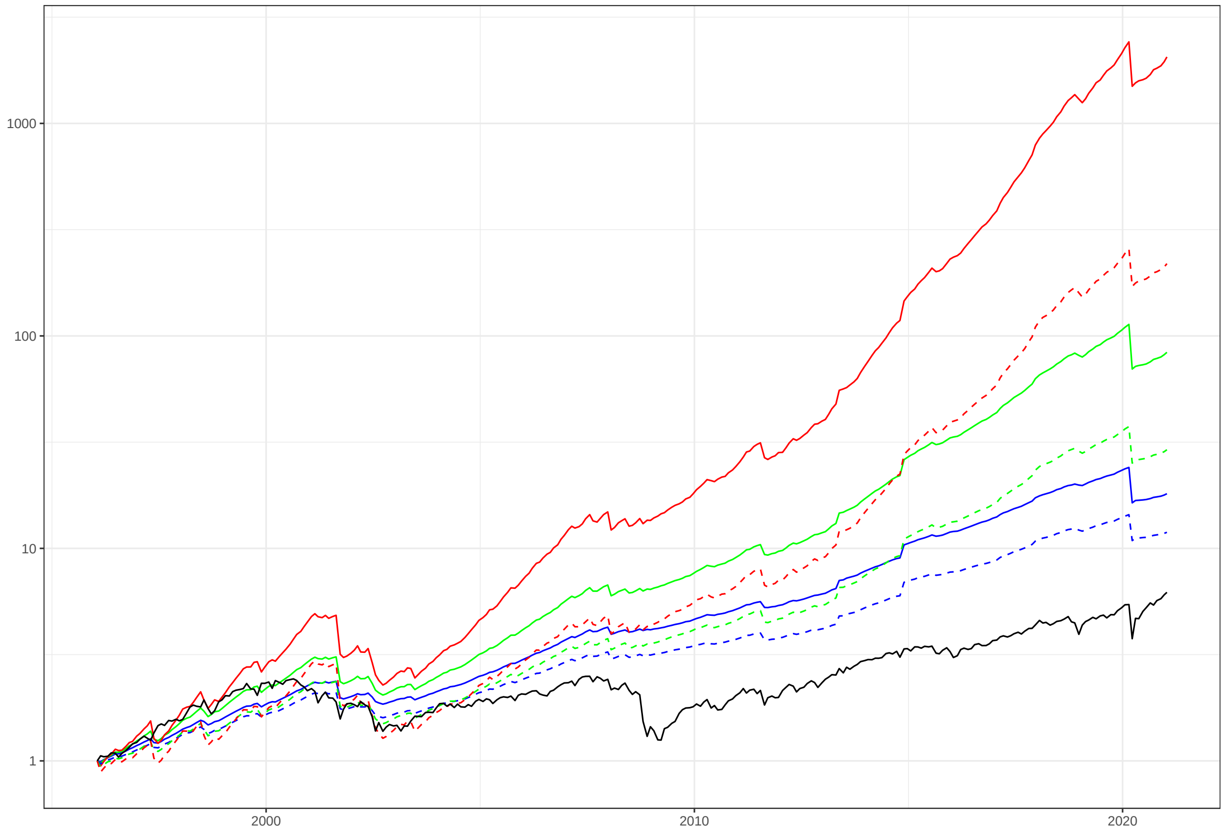

Figure 4.1: Performance of mean-variance optimal option strategies with (dashed) or without (solid) the solvency constraint, and position limits $L = 3\%$ (blue), $L = 5\%$ (green) and $L = 10\%$ (red) for risk-aversion $\gamma = 3$, compared to the performance of the index (black), the S&P 500.

constraints results in a reduction in the Sharpe Ratio of around 0.3 on average, as Appraisal ratios range from 0.65 to 1.39. (Index exposure also declines slightly.)

Overall, the mean-variance optimal options strategies display abnormal Sharpe ratios, in excess of one, over the entire 25-year period. Although the strategies considered have a positive exposure to the index, such exposure explains only a minor fraction of the positive excess returns. The abnormal Sharpe ratios persist also in the presence of solvency constraints, although they decrease significantly.

For L near zero, Proposition 2.2 demonstrates that the portfolio $w = \frac{L}{\sqrt{d}} \frac{\mu_+}{\|\mu_+\|}$ approximates well the mean-variance efficient portfolio, thereby raising the question of whether

γ	L	μ	σ	Sharpe	Alpha	Beta	Hedged σ	Appraisal
		Unconstrained						
1	3%	10.76%	10.19%	1.06	8.71%	0.28	8.84%	0.99
3	3%	9.88%	9.28%	1.07	8.11%	0.24	8.19%	0.99
5	3%	9.27%	7.37%	1.26	7.96%	0.18	6.61%	1.20
1	5%	18.16%	19.62%	0.93	14.32%	0.53	17.18%	0.83
3	5%	16.52%	12.68%	1.30	14.25%	0.31	11.37%	1.25
5	5%	15.44%	9.79%	1.58	13.89%	0.21	9.02%	1.54
1	10%	35.13%	31.53%	1.11	29.18%	0.82	27.89%	1.05
3	10%	30.51%	18.19%	1.68	27.80%	0.37	16.91%	1.64
5	10%	27.85%	14.64%	1.90	26.06%	0.25	13.96%	1.87
		Solvency Constrained						
1	3%	8.01%	10.96%	0.73	5.84%	0.30	9.55%	0.61
3	3%	8.09%	8.49%	0.95	6.48%	0.22	7.50%	0.86
5	3%	7.84%	7.23%	1.08	6.56%	0.18	6.50%	1.01
1	5%	12.42%	16.28%	0.76	9.26%	0.44	14.28%	0.65
3	5%	12.15%	12.29%	0.99	9.97%	0.30	11.05%	0.90
5	5%	12.07%	9.87%	1.22	10.52%	0.21	9.09%	1.16
1	10%	21.15%	26.79%	0.79	16.26%	0.68	23.90%	0.68
3	10%	21.28%	17.96%	1.18	18.65%	0.36	16.75%	1.11
5	10%	21.05%	14.56%	1.45	19.27%	0.25	13.88%	1.39
		Small Position Limit						
Small	3%	4.77%	6.34%	0.75	3.73%	0.14	5.79%	0.64
Small	5%	7.95%	10.56%	0.75	6.22%	0.24	9.65%	0.64
Small	10%	15.90%	21.12%	0.75	12.43%	0.48	19.30%	0.64

Table 4.1: Portfolios of options' on the S&P 500 index. The S&P 500 index has annualized average excess returns 7.22% and volatility of 17.85% (hence a Sharpe ratio of 0.40).

such an approximation, which does not depend on options' covariances and risk aversion, yields a performance that is comparable to that of portfolios obtained from the constrained maximization problem, which do depend on options' covariances and risk aversion. Note that the excess returns and volatility of the small-position limits scale linearly in L, thereby leading to the same Sharpe and Appraisal ratios.

Small-position approximations also yield positive excess returns, but they lag the performance of their finite counterparts, highlighting the importance of cross-hedging effects in generating high risk-adjusted returns.

Expected Utility

The crucial advantage of the mean-variance objective considered in this paper is the tractability characteristic of quadratic optimization problems. Its potential disadvantage is that, by focusing on the first two moments of a portfolio's payoff, quadratic optimization may not capture the effect of the whole distribution on the expected utility that the mean-variance objective approximates.

To address this issue, recall first the relation between mean-variance optimization and constant relative risk aversion utility functions

$$U(x) = \begin{cases} \frac{x^{1-\gamma}-1}{1-\gamma} & \gamma \geq 0, \gamma \neq 1 \\ \log x & \gamma = 1 \end{cases} .$$

The Equivalent Safe Rate $\mathrm{ESR}(R)$ of a return R is defined as the return on the certainty-equivalent, i.e., the hypothetical risk-free return that would make a utility-maximizer indifferent between R and such risk-free return, which is the solution to the equation

$$U(1 + \mathrm{ESR}(R)) = E[U(1 + R)]. \tag{4.1}$$

Likewise, the Mean-Variance Equivalent Safe Rate $\mathrm{MVR}(R)$ is obtained from the second-order Taylor expansion of (4.1), which is

$$\mathrm{MVR}(R) - \frac{\gamma}{2}\mathrm{MVR}(R)^2 = E[R] - \frac{\gamma}{2}E[R^2]$$

and has the solution

$$\mathrm{MVR}(R) = \frac{-1 + \sqrt{(\frac{\gamma}{2}E[R^2] - E[R])2\gamma + 1}}{-\gamma} \approx E[R] - \frac{\gamma}{2}E[R^2] = E[R] - \frac{\gamma}{2}\left(\mathrm{Var}(R) + E[R]^2\right).$$

Thus, the difference between ESR and MVR measures the additional equivalent safe rate obtained by a utility maximizer from a mean-variance optimal portfolio.

Table 4.2 reports the ESR and MVR – along with their absolute and relative differences – of the unconstrained portfolios in Table 4.1 on the whole period 1996-2020. Except for the riskiest combination ($\gamma = 1$, $L = 10$), the expected utility of each other portfolio is close to,

γ	L	ESR	MVR	ESR − MVR	ESR / MVR −1
1	3%	10.17%	10.19%	-0.02%	-0.21%
3	3%	8.10%	8.47%	-0.37%	-4.53%
5	3%	7.47%	7.73%	-0.27%	-3.56%
1	5%	14.93%	16.10%	-1.17%	-7.84%
3	5%	13.02%	13.77%	-0.75%	-5.80%
5	5%	12.29%	12.54%	-0.26%	-2.09%
1	10%	36.82%	29.64%	7.18%	19.49%
3	10%	23.69%	24.38%	-0.69%	-2.92%
5	10%	20.82%	20.87%	-0.06%	-0.26%

Table 4.2: Equivalent Safe Rate (ESR) and Mean-Variance Equivalent Safe Rate (ESR) for the S&P 500 options' portfolios in Table 4.1.

but marginally lower than, its mean-variance approximation, consistently with the familiar intuition that expected utility penalizes losses more than its quadratic expansion does.

The most negative difference of 1.17% occurs for $\gamma = 1$, $L = 5$, with the exact ESR of 14.93% against the approximate MVR of 16.10%. In all other cases, negative differences are less than 1%, and account for a few percentage points of the overall ESR. For $\gamma = 1$, $L = 10$, the exact ESR is significantly *higher* than the MVR: this particularly risky portfolio is prone to large fluctuations, and the high second moment in the observation period arises more from positive returns than from negative ones, leading the MVR to underestimate the ESR.

In summary, the comparison of the mean-variance performance MVR with the utility-based ESR tends to allay concerns that mean-variance analysis might substantially overestimate the certainty-equivalents of the options' portfolios considered. In all portfolio combinations considered, the mean-variance performance leads to either modest overestimation or, in one case, to underestimation of the ESR.

Performance on Subperiods

The remarkable performance of options' strategies over the entire period is rather uneven over time, as demonstrated by Tables 4.3 and 4.4: the extraordinary Sharpe ratios above two in the late nineties are followed by unimpressive 0.4-0.5 in the first half of the naughties. Sharpe ratios above one return in the following fifteen years, though performance at times varies significantly with the parameters, with low risk aversion associated with much lower Sharpe and Appraisal ratios. By contrast, such relation was inverted in the nineties, and absent in the naughties.

In summary, while mean-variance optimal options' strategies display significant abnormal returns over the whole observation period, the average returns vary considerably across five-year periods. Such variation is explained in part by the gap, within each period, between realized volatility, which reflects the actual variability of the index, and implied volatility, which reflects the price that options' buyers pay for such variability. Indeed, in 1996-2000 the average implied volatility (VIX) was 22.41%, over 7% higher than the realized volatility of 14.68%, leading to Sharpe and Appraisal ratios above 2.

By contrast, in 2001-2005 implied volatility averaged 20.65%, less than 2% above the 18.91% realized volatility, and the performance of options' portfolios was more modest. Similar remarks apply to the next two five-year periods, in which implied volatility averaged respectively 23.43% (2006-2010) and 17.42% (2011-2015), corresponding to gaps of 4.37% and 3.12% above realize volatility. This rule of thumb does not perform well in the last period (2016-2020), as realized volatility (21.40%) was in fact higher than average implied volatility (17.65%) and options' portfolios performance was mixed, ranging from -0.12 to 3.19. This phenomenon is partially explained by the several large daily changes in the index during the early stages of the coronavirus pandemic, which account for the unusually high realized volatility. The individual performance of options' portfolios during these critical months had a large impact on the period's overall performance and is a reminder of the impact of rare events on the performance of options' portfolios.

Effective Position

As the optimization on the squared sum of portfolio weights only ensures an upper bound on total portfolio weights, it is opportune to examine the effective position of options' portfolios over the observation period.

Denoting the option weights by $w = (w_i)_{i=1}^{2d}$, define the effective position as their sum (recall that all weights are positive, as short positions are encoded as positive positions in additional securities)

$$L_e = \sum_{i=1}^{2d} w_i.$$

By design, the effective position is less than or equal to L in each sample, in view of the Cauchy-Schwarz inequality and the constraint $w^\top w \leq L^2/d$. In fact, Table 4.5 shows that the average effective position $\hat{L}_e$ is between one half and two thirds of the upper bound L, across the parameter range $\gamma = 1, 3, 5$ and $L = 3\%, 5\%, 10\%$, and with modest deviations from these averages.

In practice, $L = 3\%, 5\%, 10\%$ lead to portfolios in which options premia paid and received account for approximately $L_e = 2\%, 3\%, 5\%$ of portfolio value, respectively.

Optimal Strategies

Figure 4.2 displays optimal options strategies on a representative date, December 2020. In each of the panels (a)-(f), the long put positions are in the upper left, long calls in the upper right, short puts in the lower left, and short calls in the lower right.

A clear pattern emerges across optimal strategies for different position limits, with or without the solvency constraint: the typical optimal strategy entails (i) short positions in out-of-the-money puts, with more aggressive (higher L) strategies loading more on deeper strikes, (ii) mixed long-short call positions, long in deep out-of-the-money calls, and short in calls that are nearer to the money.

The overall combination of these positions leads to a strategy that benefits from frequent small gains, while partly foregoing gains on months of unusually high returns, and suffering significant losses on months of negative returns.

4.2 Nasdaq 100

Figure 4.3 and Table 4.6 report the performance of mean-variance optimal strategies on options on the Nasdaq 100 index. As in the case of the S&P 500, the strategies produce abnormal positive Sharpe ratios, even controlling for index exposure.

Unlike the S&P 500, the reduction in performance from the solvency constraint is more modest, but this phenomenon may be due in part to the more limited historical sample, which starts in 1971 rather then 1927, and therefore is likely to include less extreme returns.

Rare events also appear to have a milder effect on strategies involving Nasdaq options' than on S&P 500 options. For example, the loss on the coronavirus month is much less noticeable, and is fully recovered by the end of the year. As in the case of the S&P, the

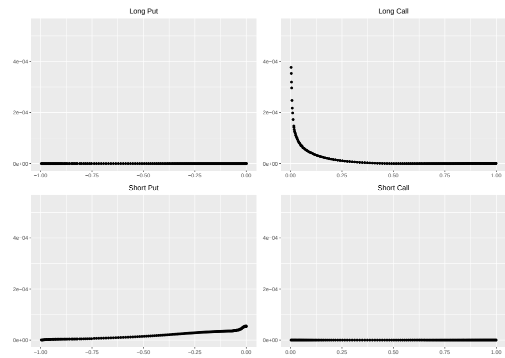

(a) Small-Position Limit, $L = 3\%$

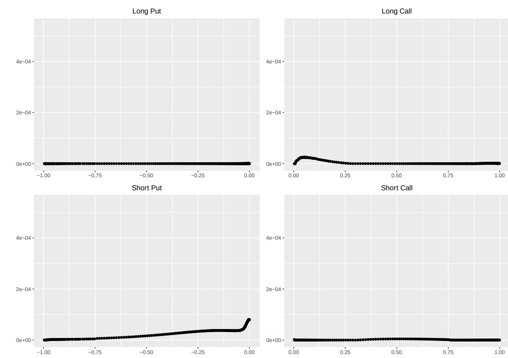

(b) $L = 3\%$

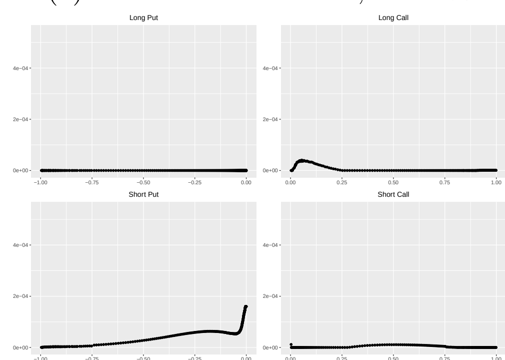

(c) $L = 5\%$

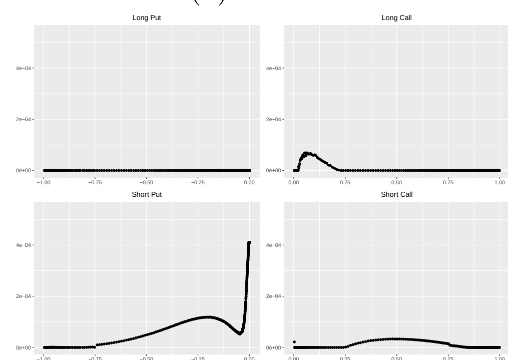

(d) $L = 10\%$

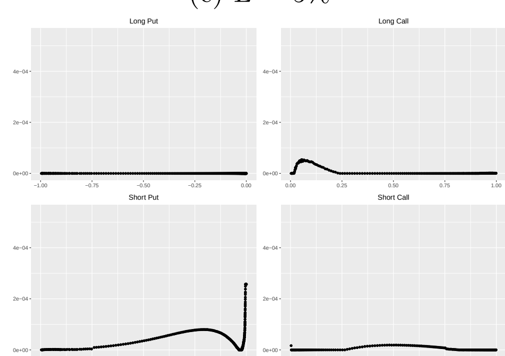

(e) Solvency Constrained, $L = 5\%$

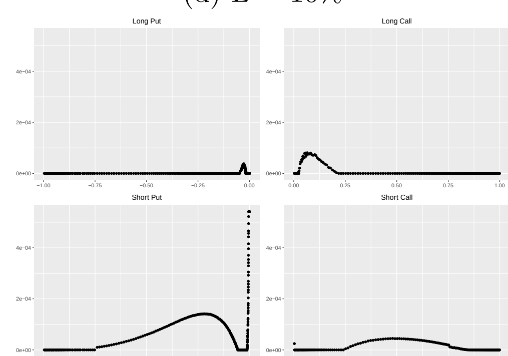

(f) Solvency Constrained, $L = 10\%$

Figure 4.2: Options' strategies in the S&P 500 index (SPX) on December 2020. Each panel displays portfolio weights (vertical) against the option delta of each contract (horizontal), for long puts (top left), long calls (top right), short puts (bottom left), and short calls (bottom right). (a): Small-position approximation, scaled for $L = 3\%$; (b), (c), (d): Unconstrained strategies for $L = 3\%, 5\%$, 10%, respectively; (e), (f): Solvency constrained strategies for $L = 5\%, 10\%$, respectively. Risk aversion $\gamma = 3$ in all strategies.

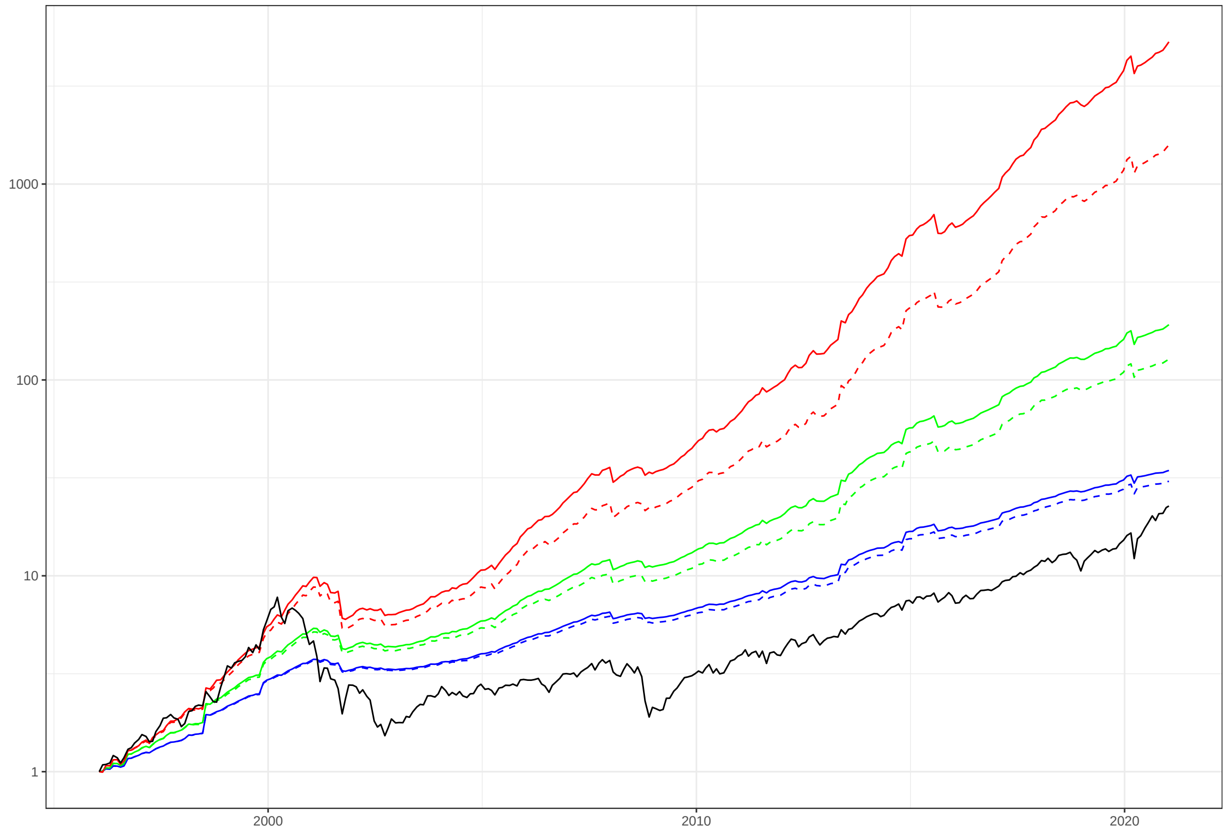

Figure 4.3: Performance of mean-variance optimal option strategies with (dashed) or without (solid) the solvency constraint, and position limits $L = 3\%$ (blue), $L = 5\%$ (green) and $L = 10\%$ (red) for risk-aversion $\gamma = 3$, compared to the performance of the index (black), the Nasdaq 100.

Nasdaq strategies also struggle in the first half of the naughties, regaining momentum in the years that follow. Index exposure is also significant for this index, but it varies considerably across parameter combinations.

Figure B.1 confirms that Nasdaq 100 options' strategies are qualitatively similar, in that they involve short positions in out-of-the-money calls and puts, long positions on near-the-money calls, and almost never long put positions. Table B.1 also confirms that actual portfolio weights are between half and two-thirds of the upper bound, as for the S&P 500 index.

4.3 Dow Jones Industrial Average

Options' strategies on the Dow Jones Industrial Average, which started trading in 1997, highlight some significant differences from the S&P 500 and Nasdaq 100.

Figure 4.4 is in stark contrast 4.1 and 4.3, showing a performance that lags behind by nearly two orders of magnitude over a period that is less than two years shorter, and a drastic loss on the coronavirus month, which erases five years of gains for most strategies, bringing their overall performance in line with that of the index itself. Note that such drastic difference in performance cannot be ascribed to data availability, as the historical sample for the Dow Jones dates back to 1927.

These considerations are summarized by Table 4.7, where the strategies' Sharpe ratios are about 0.4-0.5, and decline to 0.3 after controlling for index exposure, which is somewhat higher than for the other indexes. Solvency constraints further reduce the options' strategy to a meager 0.1-0.2, even below the index' performance.

While actual portfolio weights are similar to the other indexes' strategies (Table B.2), the options' strategies are qualitatively different, as it is clear from Figure B.2. Long positions in calls and puts are virtually absent, while short positions are mostly concentrated on near-the-money puts and, to a lesser extent, near-the-money calls.

Thus, the large abnormal returns present in the S&P 500 and Nasdaq 100 markets do not appear in options on the Dow Jones index. In isolation, this phenomenon appears to exempt such an index from the anomaly. But it does make the anomalies in the other indexes worse: because the Dow Jones is strongly correlated with the S&P 500 and the Nasdaq, an investor could use its options to hedge the short options' positions in the other indexes, potentially increasing the overall Sharpe and Appraisal ratios even further. Such a development is not discussed in this paper, which focuses on options' portfolio optimization on a single index.

5 Conclusion

This paper offers a novel method to construct optimal options' portfolios that incorporate the position limits implied by margin requirements, the substantial bid-ask spreads, and the limited depth typical of the options' market.

Position limits significantly distort optimal portfolio weights, spreading otherwise con-

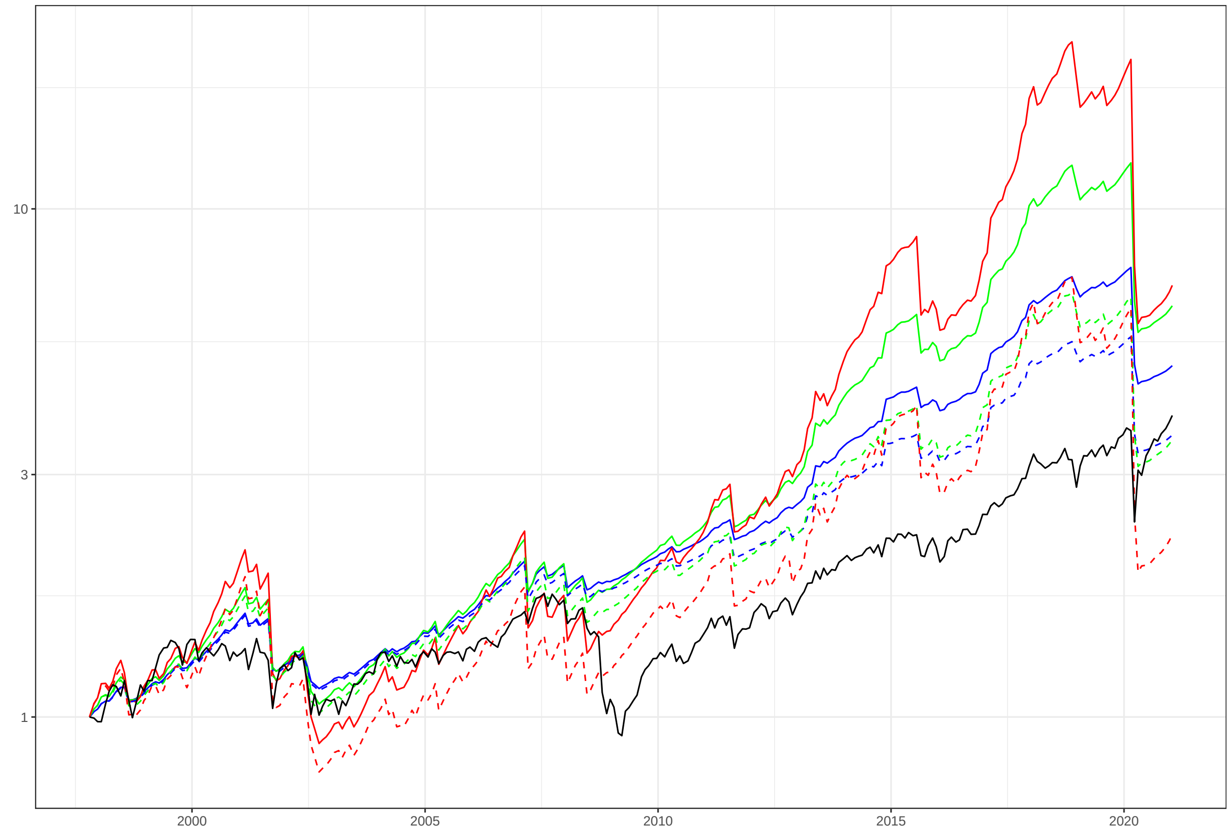

Figure 4.4: Performance of mean-variance optimal option strategies with (dashed) or without (solid) the solvency constraint, and position limits $L = 3\%$ (blue), $L = 5\%$ (green) and $L = 10\%$ (red) for risk-aversion $\gamma = 3$, compared to the performance of the index (black), the Dow Jones Industrial Average.

centrated options' positions across several strike prices, and mitigating the risks of spurious long-short positions in similar contracts. Overall, the effect of position limits is akin to that of a shrinkage estimation of the covariance matrix, combined with an increase in risk aversion.

An out-of-sample empirical study on twenty-five years of options' investments shows that typical options portfolios entail frequent small gains with rare large losses, similar to the risk profile of insurers.

A Appendix

Lemma A.1. Let $\mu \neq 0$ and $\text{rank}(\Sigma) \geq 1$. The function

$$\kappa(\lambda) := \mu^\top (\gamma\Sigma + \lambda I)^{-2}\mu, \quad \lambda > 0$$

is strictly positive, strictly decreasing, and it satisfies

$$\lim_{\lambda\to\infty} \kappa(\lambda) = 0, \qquad \lim_{\lambda\to 0} \kappa(\lambda) = \kappa_\star \in (0, \infty]. \tag{A.1}$$

The constraint $w^\top w \leq \frac{L^2}{d}$ is binding if and only if $\kappa_* \geq L^2/d$. Moreover:

(i) If $\det \Sigma > 0$, then $\kappa_\star = \mu^\top \Sigma^{-2}\mu > 0$.

(ii) If $\mu \notin \mathcal{R}(\Sigma)$, then $\kappa_\star = \infty$.

Proof. The function is strictly positive, as for any $\lambda > 0$, the inverse of $(\gamma\Sigma + \lambda I)$ exists, and $\mu \neq 0$. Differentiating, we obtain

$$\kappa'(t) = -2\mu^\top (\gamma\Sigma + \lambda I)^{-3}\mu < 0$$

for any $\lambda \in (0, \infty)$, which proves strict monotonicity and implies the limits in (A.1). Let now U be an orthogonal matrix such that $UDU^\top = \Sigma$, where D is a diagonal matrix with decreasing eigenvalues. If Σ is invertible, then clearly $\kappa_\star = \mu^\top\Sigma^{-2}\mu > 0$, and we have (i). To prove (ii), note that, as $\mu \notin \mathcal{R}(\Sigma)$, D must be degenerate, such that $D_{jj} = 0$ for $j = r + 1, \ldots, d$, where $1 \leq r = \text{rank}(D) < d$. $\mu \notin \mathcal{R}(\Sigma)$, implies that $U^\top\mu \notin \mathcal{R}(D)$[16], and thus there exists $r < j \leq d$ such that $(U^\top\mu)_j \neq 0$. It follows that

$$\kappa_\star = \lim_{\lambda\to 0} \mu^\top U(\gamma D + \lambda I)^{-2}U^\top\mu = \lim_{\lambda\to 0} \left(\frac{(U^\top\mu)_j}{\lambda}\right)^2 = \infty.$$

□

[16] In fact, $\mu \in \mathcal{R}(\Sigma)$ if and only if $U^\top\mu \in \mathcal{R}(D)$. This follows from the fact that $\mu = \Sigma\xi$ for some $\xi \in \mathbb{R}^d$ implies that $U^\top\mu = D\eta$, where $\eta = U^\top\xi$ (and vice versa).

Proof of Proposition 2.1. By rescaling the problem, the constraint is binding if the equation

$$\mu^\top \left(\gamma \Sigma L/\sqrt{d} + \lambda I\right)^{-2} \mu = 1$$

has a solution. By Lemma A.1 (strict monotonicity of κ and the first limit in (A.1)) it is enough to find $\lambda > 0$ such that

$$\mu^\top \left(\gamma \Sigma L/\sqrt{d} + \lambda I\right)^{-2} \mu > 1.$$

As

$$\mu^\top \left(\gamma \Sigma L/\sqrt{d} + \lambda I\right)^{-2} \mu \geq \frac{\|\mu\|^2}{(\gamma \sigma_\infty L/\sqrt{d} + \lambda)^2}$$

and by assumption (2.4), the right hand side becomes greater than 1 for sufficiently small λ. □

Lemma A.2. If $\bar{w}$ maximizes (2.5) or (2.6), then either $\bar{w}_{2k-1} = 0$ or $\bar{w}_{2k} = 0$ for all $1 \leq k \leq d$.

Proof of Lemma A.2. The return $\mathcal{R}_+$ from a long position equals

$$\mathcal{R}_+ = \frac{X}{A} - 1,$$

where X is the option payoff and $A > 0$ is the ask price. Denoting by $1 - \varepsilon$ the bid-to-ask ratio, the return on the short position is

$$\mathcal{R}_- = \frac{X}{A(1-\varepsilon)} - 1.$$

Without loss of generality, we consider the first option only. The optimal portfolio can be written as $\bar{w} := (w_+, w_-, w_\perp) \in \mathcal{R}_+^{2d}$. Its excess return is a random variable of the form

$$\mathcal{R} - r = w_+(\mathcal{R}_+ - r) + w_-(\mathcal{R}_- - r) + w_\perp^\top(\mathcal{R}_\perp - r).$$

Suppose, by contradiction, that the optimal portfolio $w_\delta(L)$ satisfies $w_+ > 0$ and $w_- > 0$. For $\delta < \min(w_+, w_-)$ consider the alternative strategy

$$w_\delta(L) = (w_+ - \delta, w_- - (1-\varepsilon)\delta, w_\perp). \tag{A.2}$$

Its excess return is of the form

$$\mathcal{R}^\delta - r = \mathcal{R} - r + \delta\varepsilon(r+1), \tag{A.3}$$

and thus the expected excess return for any choice $\delta < \min(w_+, w_-)$ is greater than the excess return of the optimal strategy. Furthermore, the variance of returns of the two strategies coincide by equation (A.3). The position limit also decreases because

$$w_\delta^\top w_\delta = (w_+ - \delta)^2 + (w_- - \delta(1-\varepsilon))^2 + w_\perp^\top w_\perp < w_+^2 + w_-^2 + w_\perp^\top w_\perp = w^\top w.$$

Thus, the alternative portfolio w_δ outperforms $\bar{w}$, which contradicts its optimality. Whence either $w_+ = 0$ or $w_- = 0$.

In view of equation (A.2) in the proof, the same conclusion also holds for the alternative objective (2.6). □

Proof of Proposition 2.2. Rescaling the problem in terms of the weights $v := \frac{\bar{w}}{L/\sqrt{d}}$, the problem is equivalent to maximizing

$$v^\top \mu - \frac{\gamma L}{2\sqrt{d}} v^\top \Sigma v \tag{A.4}$$

subject to

$$v \geq 0, \quad \|v\| \leq 1. \tag{A.5}$$

As the optimizer depends continuously on the parameters, the solutions $v = v(L)$ converge, as $L \to 0$, to the solution of

$$\max_v v^\top \mu, \quad \text{s.t.} \quad v \geq 0, \quad \|v\| \leq 1.$$

The limit solution is thus

$$\lim_{L\to 0} v(L) = \frac{\mu_+}{\|\mu_+\|},$$

whence (2.7) follows upon rescaling. This proves the second part of the Proposition.

It remains to show that for sufficiently small L, the constraints become binding.

By assumption, there exists $1 \leq i \leq 2d$ for which $\mu_i > 0$. (That is, a long or short option position has strictly positive excess return.) Let $w(L)$ be the optimal strategy satisfying the position limit constraint, and thus satisfying (A.4)–(A.5). By the first part of the proof, $w(0) = (w_1, \ldots, w_{2d})$ has positive weights $w_j > 0$, for every $1 \le j \le 2d$, where $\mu_j > 0$. We claim that this is a robust result, in the sense that for all sufficiently small L, the weight $w(L) = \mu_+ \times \kappa(L)$, where $\kappa(L)$ is a strictly positive $2d$-vector. To prove this claim, suppose, for a contradiction, that the optimal solution $w = w(L)$ satisfies $w_j > 0$ for some j, where $\mu_j = 0$. Let v be the alternative strategy defined by $v_j = w_j$, except for $v_i = 0$. Then, the difference between the objective evaluated at v and the maximum is

$$w_j \left(|\mu_j| + \gamma \left(\frac{L}{\sqrt{d}} \Sigma w \right)_j \right) \ge w_j \left(|\mu_j| - \gamma \frac{L}{\sqrt{d}} \|\Sigma\| \right) > 0,$$

where the last inequality holds provided that

$$L < \sqrt{d} \frac{\min\{|\mu_j| : \mu_j \neq 0\}}{\gamma \|\Sigma\|}$$

and $\|v\| < \|w\| \le 1$, which contradicts the optimality of $w(L)$.

We conclude that for sufficiently small L, the non-negativity constraints for the weights is redundant. Consider the reduced problem, where one strikes all indices j where $w_j = 0$, and uses the symbol $\star$ for such reduced quantities. This amounts to trade only those assets that have a strictly positive excess return. Since now $\mu_\star > 0$, the goal is equivalent to maximize

$$v^\top \mu - \frac{\gamma}{2} v^\top \frac{L}{\sqrt{d}} \Sigma_\star v - \frac{\lambda}{2} v^\top v$$

for $v = v(\lambda) \ge 0$ (using the Lagrange multiplier λ). Because, without non-negativity constraint, the problem has the explicit solution

$$v = \frac{1}{\gamma} \left(\frac{L}{\sqrt{d}} \Sigma + \frac{\lambda}{\gamma} I \right)^{-1} \mu_\star,$$

it follows that $v \ge 0$ for sufficiently small L (e.g., using the Neumann series). Therefore, it also solves the constrained problem $v \ge 0$ and the original problem (2.5), i.e., $v = w_\star$. Thus, for sufficiently small L, the map $\lambda \mapsto \|v^\lambda\|$ is strictly decreasing in λ, and Proposition 2.1 completes the proof. □

Proof of Proposition 3.1. Because the error $\varepsilon_{t,T}$ in (3.3) is uncorrelated with the estimator $\log \hat{\sigma}_{t,T}$, it follows that

$$\mathrm{Var}(\log \sigma_{t,T}) = \mathrm{Var}(\log \hat{\sigma}_{t,T}) + \mathrm{Var}(\varepsilon_{t,T})$$

Thus, the relation

$$\sigma_{t,T}^2 = \hat{\sigma}_{t,T}^2 e^{2\varepsilon_{t,T}}$$

implies that, as $\varepsilon_{t,T}$ is conditionally normal with zero mean,

$$E[\sigma_{t,T}^2 | \hat{\sigma}_{t,T}] = \hat{\sigma}_{t,T}^2 E[e^{2\varepsilon_{t,T}} | \hat{\sigma}_{t,T}] = e^{2\,\mathrm{Var}(\varepsilon_{t,T}|\hat{\sigma}_{t,T})}.$$

As $\varepsilon_{t,T}$ is conditionally homoskedastic, $\mathrm{Var}(\varepsilon_{t,T}|\hat{\sigma}_{t,T}) = \mathrm{Var}(\varepsilon_{t,T})$, and by the definition, the R^2 of the regression of $\log \sigma_{t,T}$ on $\log \hat{\sigma}_{t,T}$ equals

$$R^2 = \frac{\mathrm{Var}(\log \hat{\sigma}_{t,T})}{\mathrm{Var}(\log \sigma_{t,T})} = 1 - \frac{\mathrm{Var}(\varepsilon_{t,T})}{\mathrm{Var}(\log \sigma_{t,T})}$$

whence $\mathrm{Var}(\varepsilon_{t,T}) = (1 - R^2)\,\mathrm{Var}(\log \sigma_{t,T})$, and the claim follows. □

B Additional Tables

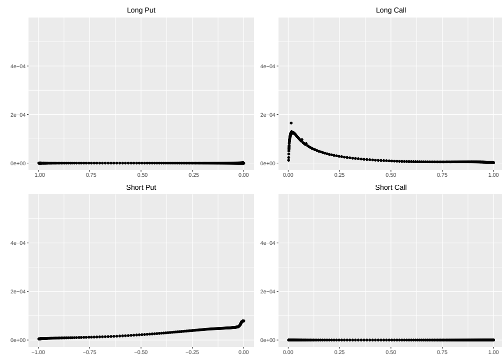

(a) Small-Position Limit, $L = 3\%$

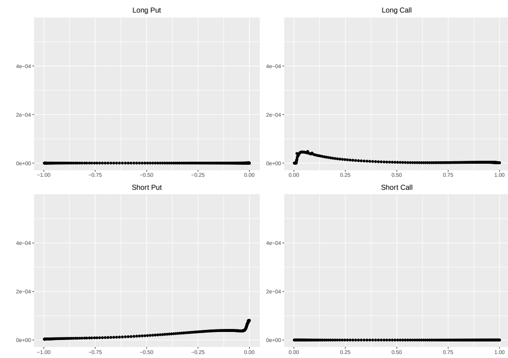

(b) $L = 3\%$

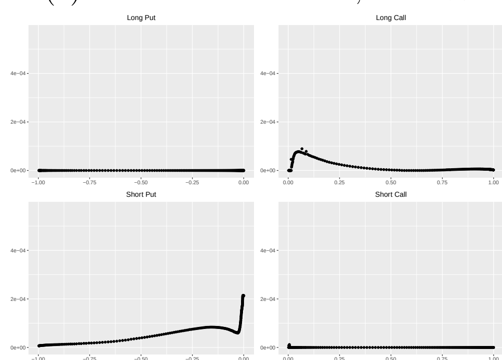

(c) $L = 5\%$

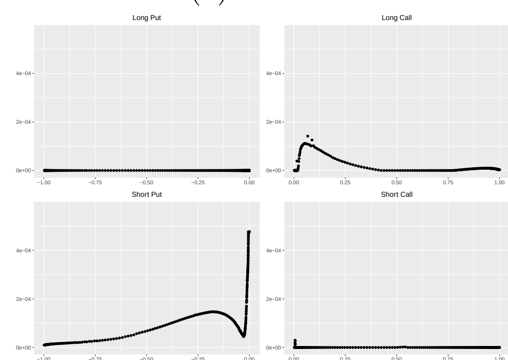

(d) $L = 10\%$

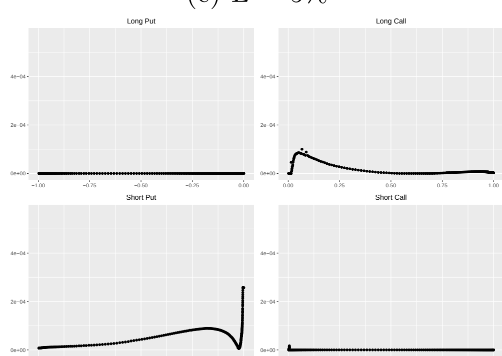

(e) Solvency Constrained, $L = 5\%$

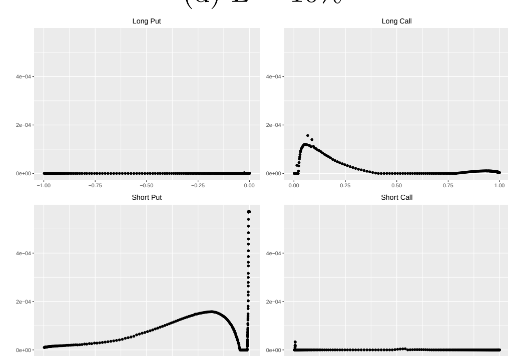

(f) Solvency Constrained, $L = 10\%$

Figure B.1: Options' strategies in the Nasdaq 100 index (NDX) on December 2020. Each panel displays portfolio weights (vertical) against the option delta of each contract (horizontal), for long puts (top left), long calls (top right), short puts (bottom left), and short calls (bottom right). (a): Small-position approximation, scaled for $L = 3\%$; (b), (c), (d): Unconstrained strategies for $L = 3\%, 5\%$, 10%, respectively; (e), (f): Solvency constrained strategies for $L = 5\%, 10\%$, respectively. Risk aversion $\gamma = 3$ in all strategies.

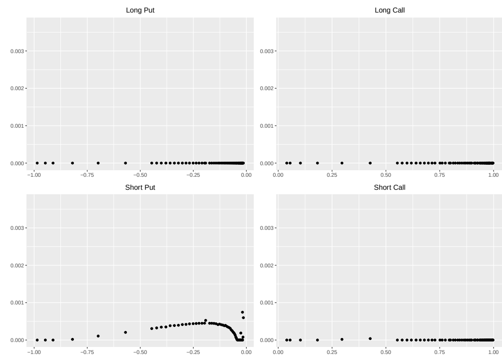

(a) Small-Position Limit, $L = 3\%$

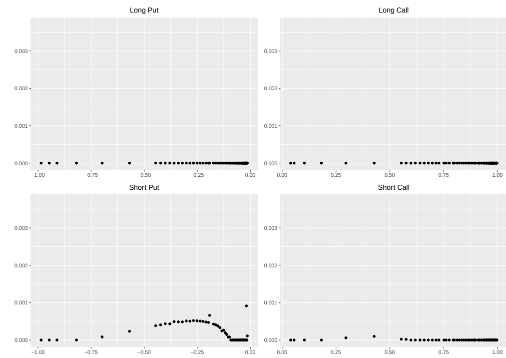

(b) $L = 3\%$

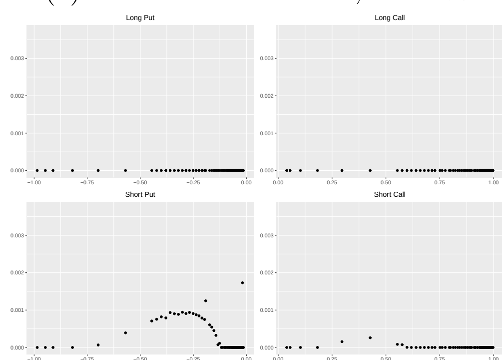

(c) $L = 5\%$

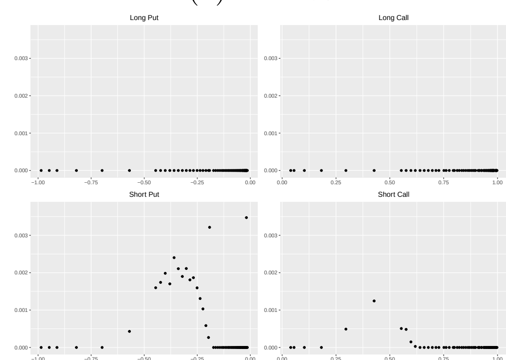

(d) $L = 10\%$

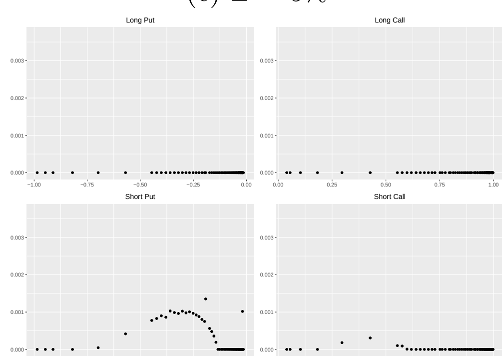

(e) Solvency Constrained, $L = 5\%$

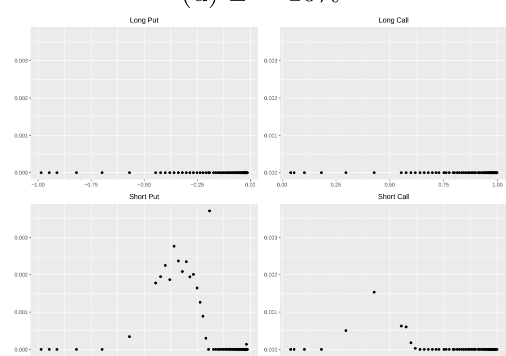

(f) Solvency Constrained, $L = 10\%$

Figure B.2: Options' strategies in the Dow Jones Industrial Average (DJX). Each panel displays portfolio weights (vertical) against the option delta of each contract (horizontal), for long puts (top left), long calls (top right), short puts (bottom left), and short calls (bottom right). (a): Small-position approximation, scaled for $L = 3\%$; (b), (c), (d): Unconstrained strategies for $L = 3\%, 5\%, 10\%$, respectively; (e), (f): Solvency constrained strategies for $L = 5\%, 10\%$, respectively. Risk aversion $\gamma = 3$ in all strategies.

γ	L	μ	σ	Sharpe	Alpha	Beta	Hedged σ	Appraisal
		1996 - 2000						
1	3%	13.65%	4.45%	3.07	14.26%	-0.06	4.37%	3.26
3	3%	11.77%	5.04%	2.34	12.85%	-0.10	4.82%	2.67
5	3%	10.37%	5.12%	2.03	11.68%	-0.12	4.80%	2.43
1	5%	21.90%	8.03%	2.73	23.32%	-0.13	7.80%	2.99
3	5%	17.43%	8.67%	2.01	19.65%	-0.21	8.13%	2.42
5	5%	14.81%	8.33%	1.78	17.26%	-0.23	7.64%	2.26
1	10%	38.86%	17.30%	2.25	42.77%	-0.36	16.46%	2.60
3	10%	27.84%	16.28%	1.71	32.84%	-0.46	14.79%	2.22
5	10%	22.79%	14.76%	1.54	27.91%	-0.47	13.02%	2.14
SPX	Index	10.80%	14.68%	0.74				
		2001 - 2005						
1	3%	4.33%	9.71%	0.45	3.97%	0.28	8.19%	0.49
3	3%	3.76%	9.16%	0.41	3.43%	0.25	7.85%	0.44
5	3%	3.33%	8.10%	0.41	3.05%	0.21	7.04%	0.43
1	5%	7.09%	16.37%	0.43	6.49%	0.46	13.89%	0.47
3	5%	5.66%	13.68%	0.41	5.19%	0.36	11.88%	0.44
5	5%	4.87%	11.65%	0.42	4.49%	0.29	10.24%	0.44
1	10%	12.60%	30.20%	0.42	11.54%	0.81	25.98%	0.44
3	10%	9.34%	21.85%	0.43	8.63%	0.54	19.30%	0.45
5	10%	8.40%	17.33%	0.48	7.86%	0.41	15.50%	0.51
SPX	Index	1.31%	18.91%	0.07				
		2006 - 2010						
1	3%	8.43%	5.94%	1.42	8.59%	0.16	5.14%	1.67
3	3%	8.61%	4.98%	1.73	8.74%	0.13	4.32%	2.02
5	3%	7.85%	4.61%	1.70	7.97%	0.12	4.05%	1.97
1	5%	15.16%	9.53%	1.59	15.42%	0.25	8.26%	1.87
3	5%	13.69%	7.79%	1.76	13.89%	0.20	6.83%	2.03
5	5%	12.26%	6.84%	1.79	12.43%	0.17	6.07%	2.05
1	10%	29.63%	17.22%	1.72	30.09%	0.44	15.02%	2.00
3	10%	23.50%	13.21%	1.78	23.83%	0.31	11.78%	2.02
5	10%	20.20%	11.20%	1.80	20.46%	0.25	10.13%	2.02
SPX	Index	-1.04%	19.06%	-0.05				

Table 4.3: Portfolios of options' on the S&P 500 index. Performance over the five-year periods 1996-2000, 2001-2005, and 2006-2010. The last row of each panel reports the index' performance over that period.

γ	L	μ	σ	Sharpe	Alpha	Beta	Hedged σ	Appraisal
		2011 - 2015						
1	3%	19.10%	10.72%	1.78	15.75%	0.37	9.32%	1.69
3	3%	16.81%	8.19%	2.05	14.08%	0.30	6.97%	2.02
5	3%	14.99%	6.57%	2.28	12.73%	0.25	5.52%	2.31
1	5%	32.48%	18.12%	1.79	26.77%	0.63	15.71%	1.70
3	5%	26.23%	11.01%	2.38	22.45%	0.42	9.26%	2.43
5	5%	23.51%	8.39%	2.80	20.73%	0.31	7.15%	2.90
1	10%	58.27%	27.49%	2.12	48.93%	1.03	23.19%	2.11
3	10%	46.10%	15.43%	2.99	41.18%	0.54	13.32%	3.09
5	10%	41.75%	11.90%	3.51	38.47%	0.36	10.72%	3.59
SPX	Index	9.04%	14.29%	0.63				
		2016 - 2020						
1	3%	8.29%	15.84%	0.52	0.10%	0.51	11.43%	0.01
3	3%	8.45%	15.06%	0.56	0.99%	0.47	11.29%	0.09
5	3%	9.81%	10.62%	0.92	4.73%	0.32	8.16%	0.58
1	5%	14.16%	34.18%	0.41	-3.14%	1.08	25.16%	-0.12
3	5%	19.62%	18.72%	1.05	10.68%	0.56	14.42%	0.74
5	5%	21.73%	12.02%	1.81	16.36%	0.34	9.64%	1.70
1	10%	36.29%	51.80%	0.70	10.91%	1.59	39.14%	0.28
3	10%	45.76%	20.94%	2.19	36.76%	0.56	17.12%	2.15
5	10%	46.10%	14.75%	3.13	40.48%	0.35	12.70%	3.19
SPX	Index	16.00%	21.40%	0.75				

Table 4.4: Portfolios of options' on the S&P 500 index. Performance over the five-year periods 2011-2015 and 2016-2020. The last row of each panel reports the index' performance over that period.

γ	L	$\hat{L}_e$, Unconstrained	$\hat{L}_e$, Solvency Constrained
1	3%	1.9% (0.3%)	1.8% (0.3%)
3	3%	1.8% (0.3%)	1.7% (0.3%)
5	3%	1.7% (0.3%)	1.6% (0.4%)
1	5%	3.2% (0.4%)	2.9% (0.5%)
3	5%	3.0% (0.5%)	2.8% (0.6%)
5	5%	2.8% (0.6%)	2.6% (0.6%)
1	10%	6.3% (1.0%)	5.5% (1.1%)
3	10%	5.5% (1.2%)	5.0% (1.2%)
5	10%	5.0% (1.2%)	4.7% (1.2%)

Table 4.5: Average position $\hat{L}_e$ (standard deviation in brackets), for unconstrained (third column) and solvency constrained (fourth column) options' portfolios on the S&P 500 index.

γ	L	μ	σ	Sharpe	Alpha	Beta	Hedged σ	Appraisal
		Unconstrained						
1	3%	12.27%	11.61%	1.06	8.72%	0.26	9.70%	0.90
3	3%	12.35%	8.65%	1.43	9.46%	0.22	6.91%	1.37
5	3%	11.40%	6.80%	1.68	9.06%	0.17	5.35%	1.69
1	5%	22.23%	18.83%	1.18	16.32%	0.44	15.55%	1.05
3	5%	19.60%	11.56%	1.70	15.63%	0.30	9.08%	1.72
5	5%	17.53%	9.02%	1.94	14.50%	0.23	7.19%	2.02
1	10%	42.43%	28.38%	1.49	32.81%	0.71	22.50%	1.46
3	10%	33.95%	16.59%	2.05	28.49%	0.41	13.37%	2.13
5	10%	29.71%	12.84%	2.31	25.80%	0.29	10.74%	2.40
		Solvency Constrained						
1	3%	11.68%	11.75%	0.99	8.08%	0.27	9.79%	0.82
3	3%	11.85%	8.79%	1.35	8.89%	0.22	6.99%	1.27
5	3%	11.19%	6.86%	1.63	8.83%	0.18	5.38%	1.64
1	5%	19.52%	19.17%	1.02	13.47%	0.45	15.77%	0.85
3	5%	17.97%	11.75%	1.53	13.91%	0.30	9.20%	1.51
5	5%	16.49%	9.04%	1.82	13.43%	0.23	7.17%	1.87
1	10%	33.91%	28.80%	1.18	24.09%	0.73	22.75%	1.06
3	10%	28.94%	16.46%	1.76	23.40%	0.41	13.11%	1.79
5	10%	26.42%	12.68%	2.08	22.44%	0.30	10.47%	2.14

Table 4.6: Portfolios of options' on the Nasdaq 100 index (NDX). The panels show the performance of mean-variance optimal strategies without (top) and with (bottom) the solvency constraint. The Nasdaq 100 index has annualized average excess returns 13.45% and volatility of 24.19% (hence a Sharpe ratio of 0.56).

γ	L	μ	σ	Sharpe	Alpha	Beta	Hedged σ	Appraisal
		Unconstrained						
1	3%	6.66%	16.09%	0.41	3.73%	0.49	13.38%	0.28
3	3%	5.64%	11.80%	0.48	3.43%	0.37	9.70%	0.35
5	3%	4.58%	9.96%	0.46	2.77%	0.30	8.31%	0.33
1	5%	10.98%	23.47%	0.47	6.61%	0.74	19.30%	0.34
3	5%	7.70%	16.72%	0.46	4.68%	0.51	13.95%	0.34
5	5%	5.79%	13.69%	0.42	3.47%	0.39	11.72%	0.30
1	10%	18.48%	38.26%	0.48	11.36%	1.20	31.50%	0.36
3	10%	10.42%	24.87%	0.42	6.34%	0.69	21.53%	0.29
5	10%	7.86%	17.72%	0.44	5.30%	0.43	15.91%	0.33
		Solvency Constrained						
1	3%	4.24%	16.41%	0.26	1.15%	0.52	13.43%	0.09
3	3%	4.30%	11.99%	0.36	2.02%	0.38	9.77%	0.21
5	3%	3.62%	10.08%	0.36	1.77%	0.31	8.35%	0.21
1	5%	6.34%	21.73%	0.29	2.13%	0.71	17.55%	0.12
3	5%	5.11%	16.95%	0.30	1.98%	0.53	13.99%	0.14
5	5%	4.06%	13.81%	0.29	1.67%	0.40	11.73%	0.14
1	10%	8.85%	31.64%	0.28	3.02%	0.98	26.17%	0.12
3	10%	5.42%	24.55%	0.22	1.24%	0.70	20.98%	0.06
5	10%	4.36%	17.97%	0.24	1.70%	0.45	16.03%	0.11

Table 4.7: Trading options on Dow Jones Industrial Average (DJX). The panels show the performance of mean-variance optimal strategies without (top) and with (bottom) the solvency constraint. The Dow Jones Industrial Average has annualized average excess returns 5.94% and a volatility of 18.13% (hence a Sharpe ratio of 0.33).

γ	L	$\hat{L}_e$, Unconstrained	$\hat{L}_e$, Solvency Constrained
1	3%	1.9% (0.3%)	1.9% (0.3%)
3	3%	1.9% (0.3%)	1.8% (0.3%)
5	3%	1.7% (0.4%)	1.7% (0.4%)
1	5%	3.3% (0.5%)	3.2% (0.5%)
3	5%	3.0% (0.6%)	2.9% (0.6%)
5	5%	2.7% (0.7%)	2.7% (0.7%)
1	10%	6.4% (1.1%)	6.0% (1.1%)
3	10%	5.2% (1.4%)	4.9% (1.3%)
5	10%	4.6% (1.4%)	4.4% (1.4%)

Table B.1: Average position $\hat{L}_e$ (standard deviation in brackets), for unconstrained (third column) and solvency constrained (fourth column) options' portfolios on the Nasdaq 100 index (NDX).

γ	L	$\hat{L}_e$, Unconstrained	$\hat{L}_e$, Solvency Constrained
1	3%	1.9% (0.3%)	1.8% (0.3%)
3	3%	1.8% (0.3%)	1.7% (0.3%)
5	3%	1.6% (0.3%)	1.6% (0.3%)
1	5%	3.1% (0.5%)	2.9% (0.5%)
3	5%	2.8% (0.6%)	2.7% (0.6%)
5	5%	2.5% (0.6%)	2.4% (0.6%)
1	10%	5.9% (1.1%)	5.4% (1.1%)
3	10%	4.8% (1.3%)	4.5% (1.3%)
5	10%	4.0% (1.3%)	3.8% (1.3%)

Table B.2: Average position $\hat{L}_e$ (standard deviation in brackets), for unconstrained (third column) and solvency constrained (fourth column) options' portfolios on the Dow Jones Industrial Average (DJX).

References

Jerome B Baesel, George Shows, and Edward Thorp. The cost of liquidity services in listed options: A note. *The Journal of Finance*, 38(3):989–995, 1983.

Gurdip Bakshi and Nikunj Kapadia. Delta-hedged gains and the negative market volatility risk premium. *The Review of Financial Studies*, 16(2):527–566, 2003.

Turan G Bali, Heiner Beckmeyer, Mathis Moerke, and Florian Weigert. Option return predictability with machine learning and big data. *The Review of Financial Studies*, 36 (9):3548–3602, 2023.

Michael J Best and Robert R Grauer. On the sensitivity of mean-variance-efficient portfolios to changes in asset means: some analytical and computational results. *The Review of Financial Studies*, 4(2):315–342, 1991.

Oleg Bondarenko. Why are put options so expensive? *The Quarterly Journal of Finance*, 4 (03):1450015, 2014.

Jonathan Raimana Chan, Thomas Huckle, Antoine Jacquier, and Aitor Muguruza. Portfolio optimisation with options. *arXiv preprint arXiv:2111.12658*, 2021.

George M Constantinides, Jens Carsten Jackwerth, and Alexi Savov. The puzzle of index option returns. *Review of Asset Pricing Studies*, 3(2):229–257, 2013.

Joshua D Coval and Tyler Shumway. Expected option returns. *The Journal of Finance*, 56 (3):983–1009, 2001.

Victor DeMiguel, Lorenzo Garlappi, Francisco J Nogales, and Raman Uppal. A generalized approach to portfolio optimization: Improving performance by constraining portfolio norms. *Management science*, 55(5):798–812, 2009.

Joost Driessen and Pascal Maenhout. An empirical portfolio perspective on option pricing anomalies. *Review of Finance*, 11(4):561–603, 2007.

Joost Driessen and Pascal Maenhout. The world price of jump and volatility risk. *Journal of Banking & Finance*, 37(2):518–536, 2013.

Bjørn Eraker. The performance of model based option trading strategies. *Review of Derivatives Research*, 16(1):1–23, 2013.

José Afonso Faias and Pedro Santa-Clara. Optimal option portfolio strategies: Deepening the puzzle of index option mispricing. *Journal of Financial and Quantitative Analysis*, 52 (1):277–303, 2017.

Steffen Hitzemann, Michael Hofmann, Marliese Uhrig-Homburg, and Christian Wagner. Margin requirements and equity option returns. *Available at SSRN 2789113*, 2021.

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

Ravi Jagannathan and Tongshu Ma. Risk reduction in large portfolios: Why imposing the wrong constraints helps. *The Journal of Finance*, 58(4):1651–1683, 2003.

Christopher S Jones. A nonlinear factor analysis of S&P 500 index option returns. *The Journal of Finance*, 61(5):2325–2363, 2006.

Sofonias A Korsaye, Alberto Quaini, and Fabio Trojani. Smart stochastic discount factors. *Available at SSRN 3878300*, 2021.

Olivier Ledoit and Michael Wolf. Improved estimation of the covariance matrix of stock returns with an application to portfolio selection. *Journal of Empirical Finance*, 10(5): 603–621, 2003.

Jun Liu and Jun Pan. Dynamic derivative strategies. *Journal of Financial Economics*, 69 (3):401–430, 2003.

Semyon Malamud. Portfolio Selection with Options and Transaction Costs. *Swiss Finance Working Paper Series*, 2014.

Susan M Phillips and Clifford W Smith Jr. Trading costs for listed options: The implications for market efficiency. *Journal of financial economics*, 8(2):179–201, 1980.

Yuliya Plyakha and Grigory Vilkov. Portfolio policies with stock options. *Available at SSRN 1101305*, 2008.

Pedro Santa-Clara and Alessio Saretto. Option strategies: Good deals and margin calls. *Journal of Financial Markets*, 12(3):391–417, 2009.

Paul Schneider and Fabio Trojani. Fear trading. Swiss Finance Institute Research Paper No. 15-03, Available at SSRN (1994454), 2015.

Hans R Stoll. The relationship between put and call option prices. *The Journal of Finance*, 24(5):801–824, 1969.

Jules Van Binsbergen, Michael Brandt, and Ralph Koijen. On the timing and pricing of dividends. *American Economic Review*, 102(4):1596–1618, 2012.