Stochastic Dominance and GARCH Option Pricing: a New Approach*

by

Russell Davidson

Department of Economics and CIREQ
McGill University
Montreal, Quebec, Canada
H3A 2T7

Aix-Marseille Université
CNRS, EHESS, AMSE
13205 Marseille cedex 01, France

email: **russell.davidson@mcgill.ca**

Hamed Ghanbari

University of Lethbridge
email: **hamed.ghanbari.uleth@gmail.com**

and

Stylianos Perrakis

John Molson School of Business
Concordia University
email: **stylianos.perrakis@concordia.ca**

Keywords:

JEL classification: G10, G11, G13, G23

* The authors wish to thank Michal Czerwonko for his contributions to the numerical estimation of several results derived in this paper.

March 2025

Abstract

We present benchmark values for one-month SPX index option prices when the underlying asset dynamics follow a GARCH model with empirical innovations. We fit these dynamics to ex-dividend option return data and find very few jumps that exceed a chosen volatility threshold. We apply the model-free stochastic dominance bounds on option values consistent with a monotone pricing kernel in a frictionless world. We find that the observed option market data lies above the Stochastic Dominance upper bound for several liquid option moneyness categories. We present dynamic policies that exploit these overpriced options in a frictionless world and show with out-of-sample tests that these policies produce significant risk-adjusted profits. We also find that a monotone kernel does not pass through the observed option bid-ask spread in a majority of the cross sections. We conclude that the options are mispriced for both frictionless markets and in the presence of frictions and demonstrate in both cases the mispricing by out-of-sample empirical tests.

1. Introduction

Several theoretical studies in the financial literature have derived the risk-neutral dynamics for the valuation of index options in the frictionless world, based on the General Autoregressive Conditional Heteroscedasticity or GARCH models of the ex-dividend index returns. In these models the random factors in the current index daily return enter under various formulations into the volatility of future index returns. The pioneering such study is by Duan (1995), followed by Kallsen and Taqqu (1998), and Heston and Nandi (HN, 2000). These studies used daily returns, normal innovations, and various alternative volatility updating equations. There is no intraday trading and risk neutralization was achieved under various assumptions, including particular types of trader utility functions. An alternative approach was pioneered by Rosenberg and Engle (RE, 2002), in which the innovations were empirically defined, and risk neutralization took place by relying on the observed option market data and extracting empirically defined pricing kernels.

On the other hand, the theoretical derivation of the risk-neutral dynamics is, to our knowledge, not feasible when the GARCH models are mixed with independent Poisson jumps (JGARCH), or when the innovations are not normal. Christoffersen *et al* (2010) analyzed in detail a series of theoretical models based on no-arbitrage equilibrium (NAE) that covered most conceivable discrete-time index return dynamics. Their models examined separately risk neutralization through admissible pricing kernels and the equilibrium setups through which the kernels were derived. For empirical applications, however, the incompleteness that arises in the option market when there are jumps or when the innovations are non-normal has been handled by assuming the existence of an arbitrary pricing kernel, which is generally extracted from observed option price data. This is what was done in Duan, Ritchken and Sun ((2006), Hsieh and Ritchken (2005), Christoffersen, Heston and Jacobs (2006), Barone-Adesi, Engle and Mancini (BEM, 2008), Christoffersen, Jacobs and Orthanalai (2012, 2013), Christoffersen, Heston and Jacobs (2013), Orthanalai (2014), Linn, Shive and Shumway (2018), Babaoglou *et al* (2018), and Barone-Adesi *et al* (2020). These empirical studies have generated controversial results concerning the pricing kernel and the ability to represent the values of certain types of options. The controversies started with Jackwerth (2000) and were summarized extensively in Cuesdeanu and Jackwerth (2018), and Perrakis (2022).

In this paper we tackle a more fundamental question that concerns the possibility of extracting the NAE from the observable option market data, as assumed in all the above studies. This question arises from the fact that the frictionless NAE is not observable, insofar as the option market data produce a pair of bid and ask prices, and the NAE lies within these prices only under strong assumptions, discussed further on in this section. We test these assumptions by applying the stochastic dominance (SD) approach to index option pricing under GARCH dynamics with empirically defined innovations. Our major contribution lies in the fact that the observed option market prices are not used in deriving upper and lower bounds (UB and LB) on admissible frictionless option values consistent with the index dynamics. We can thus check whether the NAE assumption about the existence of frictionless equilibrium option values within the bid-ask spread is justified for the universe of traded options in our sample.

In the SD approach there is partial equilibrium between the underlying and the option markets that does not allow an individual investor holding the index and a riskless asset to realize superior risk-adjusted returns by including zero-net-cost individual options or option portfolios in her holdings. The only assumption necessary for such a partial equilibrium in the option market is the existence of a class of traders holding index futures or an index tracking fund, and the riskless asset. Such traders obviously exist in the economy since the S&P 500 index has in most studies the interpretation of the market portfolio and there are several exchange-traded funds (ETF) indexed on it. Further, all traded options use designated market makers, who are employees rewarded on the realized profits and so maximize their wealth at a target date after the option maturities. An immediate corollary of such an assumption is that the kernel must be monotone decreasing in the index return. As Barone Adesi *et al* (2020, p. 431) point out, violations of monotonicity "are inconsistent with the neoclassical definition of investor risk-aversion and market equilibrium".

Because of its structure, SD allows without exception the risk neutralization of all GARCH models that have appeared in the literature, without and with independent rare events and with normal or empirically defined innovations. There are explicit and model-free expressions transforming the index dynamics into risk-neutral distributions without using option market data, which we adapt to the GARCH dynamics. These transformations define a unique option value under all GARCH models with normal innovations without any additional assumptions about trader utility. When the innovations are empirically defined and/or when there are rare events the SD theory derives bounds that contain the admissible option values, those that are consistent with a monotone pricing kernel. While the simple GARCH risk-neutral distributions were already derived under normal innovations in the above-cited references, their derivations under SD are novel, and they lead directly to the bounds under empirically defined innovations and/or jump GARCH or JGARCH. Since the SD bounds use only the index dynamics for their derivation, they can be used to assess the appropriateness of fitting kernels to observed option market data and thus explain the resulting controversies.

We present such a generalized application of SD and illustrate it with specific examples of GARCH models with empirically defined innovations. We use a long time series of observed SPX monthly data and apply as our base cases the RE and BEM GARCH models with empirically derived innovations to extract the SD bounds. Volatility updating takes place under the Glosten, Jagannathan and Runkle (GJR, 1993) model. We compare these bounds to the observed bid-ask spreads in the option market and check whether there is partial overlap which includes the bid-ask midpoint, widely used as the correct equilibrium price in the above empirical studies. The tests are carried out at the level of individual options.

As it turns out, we observe that there is systematic no overlap between the bid-ask spread of the option market, and the SD bounds for several highly liquid options defined by their degree of moneyness in most of the observed cross sections. This non-overlap occurs from above, implying that the market overprices the options under frictionless trading conditions. The overpricing is extensive for the highly liquid zones near the at-the-money (ATM) category for both calls and puts, as well as one out-of-the money (OTM) category for calls. It is especially pronounced in all the OTM zones for puts, including the liquid deep OTM (DOTM) category.

For cross sections in which such an overlap does not exist for significant portions of the support of the index distribution, there are two separate but related implications, referring to trading without and with frictions. Individual traded options are overvalued (undervalued) in the frictionless market if their bid-ask spread lies entirely above (below) the SD bound interval. In such a case we derive and apply frictionless trading strategies at the level of the individual options to exploit these inconsistencies. Since we do not observe any underpricing in our options data, the strategies are for overpriced options, separately for calls and puts. These strategies are dynamic, in the sense that the zero-net-cost portfolios are established at the option trade date and closed at option maturity, but they are also optimally rebalanced along the path, depending on the realized index return. Such rebalancing is costless since we are in a frictionless world.

The profitability of such strategies can then be tested with out-of-sample tests on their realized payoffs, which should stochastically dominate the portfolios of the index and riskless asset. These tests compare the realized returns of two generic investors termed the index trader (IT), who holds the index plus riskless asset portfolio, and an identical trader named the option trader (OT) who adds to the IT portfolio a zero-net-cost option position. The null of the tests is of non-dominance, following Davidson and Duclos (DD, 2013). In the particular case of the frictionless dynamic strategies the DD tests reject the null for both OTM calls and puts, weakly for calls and strongly, with major excess returns, for puts. By contrast, the rejection weakens or disappears when the ATM zone overpricing is included in the short option portfolios.

An important advantage of the SD approach is the fact that it can also be applied to markets in which there are costs in trading the index, as well as bid-ask spreads in trading the options. With very few exceptions noted below, all empirical index option studies have focused on the modelling of the frictionless NAE, which is supposed to be located within the observed bid-ask spread. Nonetheless, the conditions for this for an entire option cross section are that there should be perfect competition with free entry and no market power, as shown in an early but rarely cited theoretical study by Jouini and Kallal (1995) in a market with frictions for any type of traded asset. While the observed no overlap of the option bid-ask spread and the SD bounds implies overpricing in the frictionless market, it may also imply that there is no monotone kernel that passes through the bid-ask spread in the corresponding cross section.

In such cases we may invoke recently derived theoretical results by Constantinides, Jackwerth and Perrakis (2009), Post and Longarela (2021), and Beare (2011, 2023), according to which there are SD opportunities in the form of suitably derived zero-net-cost option portfolios in the market with frictions. These strategies are generally buy-and-hold and are generated by a search algorithm based on linear programming

(LP), introduced by Constantinides, Czerwonko and Perrakis (CCP, 2020). This latter study found several mispriced option cross sections, particularly in short-term options. The time series of their realized payoffs can be translated into added returns to the index and subjected to the out-of-sample Davidson-Duclos (DD, 2013) tests, as it was done in Constantinides *et al* (2011) CCP, and Post and Longarela (2021). Since the types of mispricing observed in the frictionless world suggest strongly that the markets for calls and puts are segmented and a monotone kernel does not pass through the bid-ask spread, we describe the CCP algorithm and apply it to the observed option market data.

The algorithm first identifies the mispriced cross sections in which no monotone kernel passes through the bid-ask spreads of the options. We observe that this occurs for the overwhelming majority of the cross sections in the 25-year period of our option market data. The application of the algorithm then identifies the set of portfolios that maximize certain well-known criteria of the excess OT returns such as the Sharpe, Sortino, and information ratios, and applies to them the DD test. The results reject the non-dominance null for most of the criteria, thus confirming the mispricing of the options in the world with frictions.

In the next section we present the main theoretical content of the paper, starting with the specification of the GARCH model and the risk-neutral transformations of the dynamics corresponding to the SD upper bound. We also define the dynamic strategies to exploit the violations of the corresponding bounds, separately for calls and puts. Section 3 describes our index and option data and presents the GARCH parameter estimates, as well as the corresponding SD upper bounds, which it compares with the observed option data and notes the massive overpricing of the options in the frictionless world. Section 4 applies the dynamics frictionless strategies and confirms the mispricing of both calls and puts in several highly liquid moneyness zones. It also applies the CCP algorithm to the mispriced cross sections. Section 5 concludes.

2. The General Model

Let S_t denote the value of the index at time t, σ_{t+1}^2 the variance of the realized returns at time $t+1$, assumed known at t, r_t the logarithm of the riskless return, and ε_{t+1} the random term. In all cases we assume that within each period $[t, t+1]$, generally taken to be equal to one day, the return follows either empirically defined innovations, or a simple lognormal diffusion augmented with an independent jump process.

Empirical innovations

We consider empirically defined innovations, implying as in RE that ε_{t+1} has an empirically defined distribution of mean 0 and variance 1, $G(0,1)$. The daily returns are then given by the following general model, in which both the ex-dividend[1] risk premium on the index and the volatility updating equation have assumed various forms

[1] Hereafter we omit the dividends, which are taken as given in all empirical applications.

in existing studies:

$$\begin{aligned} \log S_{t+1}/S_t &= r_t + b(\sigma_{t+1}^2) + \sigma_{t+1}\varepsilon_{t+1}, \;\; \sigma_t^2 = F(\sigma_{t-1}^2), \\ \sigma_t^2 &= F(\sigma_{t-1}^2, \sigma_{t-2}^2, \ldots, \sigma_1^2) \end{aligned} \qquad t = 0, 1, \ldots, T-1. \quad (2.1)$$

In our empirical applications this model will be specialized to the Glosten, Jagannathan and Runkle (1993), or GJR GARCH case, the one used in RE and BEM, eqs (13)-(14) and (3) respectively, in which the variance updating differs between positive and negative index return innovations. All undefined symbols represent constant parameters of the model. As in all the empirical GARCH option studies, we limit ourselves to single lag cases as in the empirical version of the corresponding paper, with the extension to multiple lags following similar paths.

$$b(\sigma_{t+1}^2) = b, \qquad \sigma_{t+1}\varepsilon_{t+1} = u_{t+1} \Rightarrow u_{t+1}/\sigma_{t+1} = \varepsilon_{t+1}; \tag{2.2}$$
$$F(\sigma_{t-1}^2, \sigma_{t-2}^2, \ldots, \sigma_1^2) = \alpha + (\gamma + \theta I_{t-1})u_{t-1}^2 + \delta\sigma_{t-1}^2, \;\; I_{t-1} = \begin{cases} 0 & \text{if } u_{t-1} \geq 0 \\ 1 & \text{if } u_{t-1} < 0 \end{cases}.$$

Our main results can be shown to hold for all models that conform to the general formulation (2.1) such as, for instance, the Duan (1995), Heston and Nandi (2000) and Nelson (1991) models, as well as for all variants of the Duan model presented in Christoffersen and Jacobs (2004).

For risk neutralization, assume that the parameters $(\alpha, b, \gamma, \theta, \delta)$ in (2.2) have been estimated as described in Appendix I, following the method introduced by Bollerslev and Wooldridge (1992) as elaborated in Chapter 13 of Davidson and MacKinnon (2004). Given our time series of ex-dividend index returns $t = 0, \ldots, T$ and the initial variance σ_1^2, we filter out the daily observed innovations as follows:

$$\log(S_t/S_{t-1}) - (r_t + b) = u_t, \;\; u_t/\sigma_t = \hat{\varepsilon}_t, \;\; t = -, 1, \ldots, T \tag{2.3}$$

These data points generate an empirical distribution of mutually independent innovations that can be ordered into values $\hat{\varepsilon}_t = z_t \in [z_{\min}, z_{\max}]$ with mean 0 and unit variance denoted by the cumulative distribution (CDF) $P(z)$, which is the key input in deriving the risk-neutral Q-distribution.

For the risk neutralization of these dynamics, suppose that we have T daily returns in one observation point of our sample of SPX option values. Let $S_\tau, \tau \in [t, T-1]$ denote the ex-dividend index value in any given day, K the strike price, $C(S_t, K, T)$ and $P(S_t, K, T)$ the *frictionless* option values consistent with the index dynamics, and $Y(z)$ the monotone pricing kernel. These frictionless values are distinct from the observed option market data at time t, consisting of the bid and ask prices (C_{bt}, C_{at}) and (P_{bt}, P_{at}) for calls and puts respectively. The relationship between these observed data and the theoretically "correct" frictionless values, which is a major component of this paper, is to our knowledge novel, since most empirical option market studies have used the option market data to value the frictionless options.

Define $Z_\tau = \exp(r_\tau + b + \sigma_\tau z)$. The equilibrium equations for every day till maturity are as follows, with the expectations taken with respect to $P(z)$:

$$\mathrm{E}_\tau[Y(z)] = 1, \;\; \mathrm{E}_\tau[Y(z)Z_\tau] = \mathrm{E}_\tau[Y(z)\exp(r_\tau + b + \sigma_\tau z)] = \exp r_\tau. \tag{2.4}$$

The SD theory for frictionless markets (Perrakis and Ryan (1984), Ritchken (1985), Levy (1985), Perrakis (1986), Ritchken and Kuo (1988)) has shown that the maximum option value that can be supported by the equilibrium (2.4) for any monotone kernel is given by the expectation at day τ of the option payoff with the following risk-neutral return distribution, where we have defined:

$$\hat{Z}_\tau = \mathrm{E}_\tau[Z_\tau], \;\; Z_{\tau\min} = \exp(r_\tau + b + \sigma_\tau z_{\min}). \tag{2.5}$$

$$Q(z) = \begin{cases} P(z) & \text{with probability } \dfrac{\exp(r_\tau) - Z_{\tau\min}}{\hat{Z}_\tau - Z_{\tau\min}} \\ 1_{z_{\min}} & \text{with probability } \dfrac{\hat{Z}_\tau - \exp(r_\tau)}{\hat{Z}_\tau - Z_{\tau\min}} \end{cases} \tag{2.6}$$

An SD lower bound can also be derived, but it can be safely conjectured that it will never be violated.[2] As shown in Ghanbari *et al* (2021), in the case of jumps this lower bound lies below the Merton (1976) option value that assumes a non-priced jump risk.

Equation (2.6) can generate numerically the SD upper bound given the empirical distribution of the observed innovations, the GARCH and other return parameters, and the riskless rate. It suffices to apply Monte Carlo (MC) simulation of the paths along the 21 days till option maturity and apply the transformation (2.6) to generate the risk-neutral path probability for every path. The comparison of this derived upper bound with the observed bid and ask prices of the corresponding option will enable us to derive inferences about the intermediate market in which these observed prices have been formed.

Exploiting mispriced options in the frictionless world

Suppose we observe at time t a call option in a cross section with maturity T, whose bid price lies above the SD upper bound, or $C_{bt}(S_t, K, T) > \bar{C}(S_t, K, T)$. We derive the strategy that exploits the overpricing for all cases of discretized index dynamics, including the empirical GARCH innovations. The strategy consists in shorting one call per unit index at the bid price and allocating $\beta_t C_{bt}$ and $(1-\beta_t)C_{bt}$ in the riskless bond and the index, respectively. At time t, the allocation β_t is chosen so that at the lowest value of the return $Z_{t\min} \equiv Z_{1t}$ that corresponds to the left tail of the return distribution, the portfolio payoff from t to $t+1$ as in (2.5) will be zero. These zero-net-cost call option positions must be rebalanced at each intermediate time $\tau \in [t, T-1]$ and closed at maturity at the option payoff $(S_T - K)^+$. Making the definitions

[2] Both SD upper and lower bound distributions are shown in Perrakis (2019, p. 24), together with two alternative proofs in pp. 24-26 and 30-35.

$Z_\tau = \exp(r_\tau + b + \sigma_\tau z_{\tau+1})$, $R_\tau = \exp(r_\tau)$, we have the following optimal allocation at time t and the corresponding portfolio value $\Pi_t^C(Z_{t+1})$:

$$\begin{aligned} \beta_t = \frac{\bar{C}(S_t Z_{1t}, K, T) - Z_{1t} C_{bt}(S_t, K, T)}{R_t C_{bt}(S_t, K, T) - Z_{1t} C_{bt}(S_t, K, T)} \equiv \beta_t^*, \\ \Pi_t^C(Z_{t+1}) = \beta_t^* C_{bt} R_t + (1 - \beta_t^*) C_{bt} Z_{t+1}. \end{aligned} \tag{2.7}$$

Observe that by construction we have $E_t^Q[\Pi_t^C(Z_{t+1})] = \beta_t^* C_{bt} R_t + (1-\beta_t^*) C_{bt} \mathrm{E}_t^Q(Z_{t+1})$ $= C_{bt} R_t$. The rebalancing at any time $\tau \in (t+1, T-1]$ is done on the same principle, namely that the portfolio value $\Pi_{\tau+1}^C(Z_\tau)$ must be 0 at the lowest value $Z_{\tau+1\min} = Z_{1\tau+1}$ and the expectation at τ must be equal to the product of the successive risk-neutral allocations, namely that $\mathrm{E}_t^Q[\Pi_{\tau-1}^C(Z_\tau)] = C_{bt} \prod_{i=t}^{i=\tau} R_i$. For this we set

$$\begin{aligned} \beta_\tau = & \frac{\bar{C}(S_\tau, Z_{1\tau+1}, K, T) - Z_{1\tau+1} \Pi_{\tau-1}^C(Z_\tau)}{[R_\tau - Z_{1\tau+1}] \Pi_{\tau-1}^C(Z_\tau)} \equiv \beta_\tau^*, \\ \mathrm{E}_\tau^Q \left[\Pi_{\tau-1}^C(Z_\tau) \beta_\tau^* R_\tau + (1 - \beta_\tau^*) Z_{\tau+1} \right] \right] = & R_\tau \Pi_{\tau-1}^C(Z_\tau). \end{aligned} \tag{2.8}$$

It can be easily shown by induction that if $C_{bt} > \bar{C}(S_t, K, T)$ then the rebalancing strategy (2.6)-(2.8)yields for any $t > T$ a frictionless SD opportunity in the sense that its expected payoff is positive if there are no trading frictions and the short position is closed at option maturity. Note that at $T-1$ the option payoff $(S_T - K)^+$ replaces the SD upper bound in (2.8), yielding

$$\frac{(S_{T-1} Z_{1T} - K)^+ - Z_{1T} \Pi_{T-2}^C(Z_{T-1})}{[R_{T-1} - Z_{1T}] \Pi_{T-2}^C Z_{T-1}} \equiv \beta_{T-1}^*.$$

For the implementation of the ex-post tests the null hypothesis is the following relation, which is then used in out-of-sample dominance tests:

$$\Pi_{T-2}^C(Z_{T-1})[\beta_{T-1}^* R_{T-1} + (1-\beta_{T-1}^*) Z_\tau] = C_{bt} \prod_{i=t}^{i=T-1} [\beta_i^* R_i + (1-\beta_i^*) Z_{i+1}] \geq (S_T - K)^+. \tag{2.9}$$

For an overpriced put option, at time t, we write a put option at its bid price P_{bt}, short $\beta_t S_t - P_{bt}$ units of the index and invest $\beta_t S_t$ in the riskless asset. The portfolio payoff if the short position is closed in the frictionless market at $t+1$ is $\beta_t S_t R_t - [\beta_t S_t - P_{bt}] Z_{t+1} - P(S_t Z_{t+1}, K, T)$, whose lowest value is when the put is at its upper bound at $t+1$, or $P(S_T Z_{t+1}, K, T) = \bar{P}(S_{t+1}, K, T)$. This payoff is clearly increasing in the put bid price P_{bt} for every β_t. At the lowest value of the index return support $Z_{t+1} = Z_{1t+1}$, the payoff should be nonnegative, implying that the optimal allocation at time t is:

$$\beta_t \geq \frac{\bar{P}(S_t Z_{1t+1}, K, T) - P_{bt} Z_{1t+1}}{S_t [T_t - Z_{1t+1}]} \equiv \beta_t^*. \tag{2.10}$$

Setting β_t at its optimal level at time t, the expected portfolio payoff if the position is closed at the upper bound at $t+1$ is $\beta_t^* S_t R_t - \mathrm{E}^Q[(\beta_t^* S_t - P_{bt}) Z_{t+1} - \bar{P}(S_t Z_{t+1}, K, T)]$. This expectation should be zero at $P_{bt} = \bar{P}(S_t Z_{t+1}, K, t)$, which is verified using (2.6).
If the position is not closed at $t+1$, we have realized the short position $\Pi_t^P(Z_{t+1}) = \beta_t^* S_t(R_t - Z_{t+1}) + P_{bt} Z_{t+1}$, of which the payoff at $t+2$ is

$$\beta_t^* S_t(R_t R_{t+1} - S_{t+1} Z_{t+2}) + P_{bt} Z_{t+1} Z_{t+2}.$$

Suppose that we decide at $t+1$ to close the position at $t+2$ at the prevailing put price $P(S_{t+1} Z_{t+2}, K, T)$. Then the payoff is at its lowest value when the put is at its upper bound at $t+2$, or $P(S_{t+1} Z_{t+2}, K, T) = \bar{P}(S_{t+1}, K, T)$. If the index position is not readjusted, then at $t+2$ we have

$$\beta_t^* S_t R_t R_{t+1} - \beta_t^* S_{t+1} Z_{t+2} + P_{bt} Z_{t+1} Z_{t+2} - \bar{P}(S_{t+1} Z_{t+2}, K, T).$$

This payoff is clearly increasing in the put bid price P_{bt} for every β_t. For SD, at the lowest value of the support of the index return $Z_{t+2} = Z_{1t+2}$ the payoff should be nonnegative, implying that the optimal reallocation at time $t+1$ would be to short an additional amount $\Delta\beta_{t+1}^* = \beta_{t+1}^* - \beta_t^*$. By setting

$$\Pi_{t+1}^P(Z_{t+2} = \beta_t^*(S_t R_t R_{t+1} - S_{t+1} Z_{t+2}) + P_{bt} Z_{t+1} Z_{t+2} + \Delta\beta_{t+1}^* S_{t+1}(R_{t+1} - Z_{t+2})$$

we must have a nonnegative result when $Z_{t+2} = Z_{1t+2}$, namely $\Pi_{t+1}^P(Z_{1t+2} = \bar{P}(S_{t+1} Z_{1t+2}, K,)$. Solving, we get

$$\Delta\beta_{t+1}^* = \frac{\bar{P}(S_{t+1} Z_{1t+2}, K, T) - \beta_t^*(R_t R_{t+1} - Z_{t+1} Z_{t+2}) - Z_{1t+2} Z_{t+1} P_{bt}}{S_{t+1}[R_{t+1} - Z_{1t+2}]} \tag{2.11}$$

If $\Delta\beta_{t+1}^*$ is the optimal level at time $t+1$ according to (2.11), the portfolio payoff at $t+2$ is

$$\Pi_{t+1}^P(Z_{t+2}) \equiv \beta_t^* S_t[R_t R_{t+1} - Z_{t+1} Z_{t+2}] + \Delta\beta_{t+1}^* S_{t+1}(R_{t+1} - Z_{t+2}) + P_{bt} Z_{t+1} Z_{t+2}$$

if it is closed when $P(S_{t+1} Z_{t+2}, K, T) = \bar{P}(S_{t+1}, K, T)$, and its Q-expectation at $t+1$ should be 0 when

$$Z_{t+1} P_{bt} \equiv P_{bt+1} = \frac{\mathrm{E}^Q[\bar{P}(S_{t+1} Z_{t+1}, K, T)]}{R_{t+1}} = \bar{P}(S_{t+1}, K, T),$$

which is verified using (2.6).
In general, for any $\tau \in [t+1, T-1]$, we define

$$P_{b\tau+1} = P_{bt} \prod_{i=t+1}^{i=\tau+1} Z_\tau, \quad \Delta\beta_{\tau+1}^* = \beta_{\tau+1}^* - \beta_\tau^*,$$

$$\text{solving } \Pi_\tau^P(Z_{1\tau+1}) = \bar{P}(S_\tau Z_{1\tau+1}, K, T) \tag{2.12}$$

$$\Pi_\tau^P(Z_{\tau+1} = \beta_\tau^* S_\tau[\prod_{i=t}^{\tau} R_i - \prod_{i=t+1}^{\tau+1} Z_i] + \sum_{i=t}^{\tau} \Delta\beta_{i+1}^* S_{i+1}[\prod_{j=i+1}^{\tau} R_j - \prod_{j=i+2}^{\tau+1} Z_j] + P_{b\tau+1}$$

At $T-1$ the upper bound is equal to $\mathrm{E}^Q[(K-S_{T-1}Z_T)^+]$ and $(K-S_{T-1}Z_{1T})^+$ replaces the upper bound at the lowest return value in evaluating $\Delta\beta^*_{T-1}$ and $\Pi^P_{T-1}(Z_T)$.

For the implementation, we apply (2.10)-(2.12) at every step of the path. The cumulative proceeds must at maturity satisfy the following on average:

$$\beta^*_t S_t\Big[\prod_{i=t}^{T-1} R_t - \prod_{i=t+1}^{T} Z_i\Big] + \sum_{i=t}^{T-2}\Delta\beta^*_{i+1}S_{i+1}\Big[\prod_{i=t}^{T-1} R_i - \prod_{i=t+1}^{T} Z_i\Big] + P_{bt}\prod_{i=t+1}^{T} Z_i - (K-S_T)^+ \geq 0. \tag{2.13}$$

It is easy to see that the expectation of these cumulative proceeds is positive for $P_{bt} > \bar{P}$, since their Q-expectation is zero if the short put price is equal at t to the SD upper bound.

Several variants of these strategies can also be applied to verify the profitability in frictionless markets of trading in options whose bid prices lie above the frictionless upper bound. For instance, it is possible to incorporate statistical errors in the estimation of the SD upper bound, by filtering out option bid prices that exceed the SD bound multiplied by a factor greater than one. Similarly, we may also close the put position at the upper bound $\bar{P}(S_\tau, K, T)$, $\tau \in (t, T-1]$, and carry the proceeds to option expiration. Such strategies are obviously not feasible in the market with frictions, since the ask price lies above the bid price. In our empirical applications, we liquidate the portfolio and close the option positions at maturity by following the rebalancing (2.10)-(2.12) and assuming trading is frictionless, or at the prevailing ask price at any step along the path in the presence of frictions.[3]

Observe that if there exist many cross sections in which the observed bid prices lie above the SD upper bound then there is no monotone kernel passing through the bid-ask spread in those cross sections. In such cases the equilibrium models (2.4) or (2.7) need to be drastically modified in order to recognize frictions in trading the underlying. The pricing kernel becomes a bivariate function depending on the traders' holdings of both the index and the riskless asset, since it is no longer costless to transfer money from the index to the bond account or *vice versa*. To our knowledge, the equilibrium determination of the bid and ask prices of the options under transaction costs is available only theoretically in terms of the reservation prices of the traders in Perrakis (2019, pp. 233-241), and only for CRRA investors. There is, however, an important theoretical result as noted in the introduction, according to which there are option portfolios yielding profitable SD opportunities for all risk averse investors if a monotone kernel does not pass through the bid-ask spreads of all the options in a cross section. We present the LP that derives these portfolios in Appendix II and discuss their results in the following sections.

[3] Note, however, that frictions were not taken into account in establishing the option portfolios, implying that the resulting profits under frictions, if any, are upper bounds in exploiting the observed frictionless violations of the SD bounds.

3. Data and Estimations for the GJR GARCH with Empirical Innovations

Estimating the GARCH parameters

Since the parameters vary to some extent depending on the data from which they were extracted, we use as our base case separate estimations for the 1960-2000 and 2000-2023 series, corresponding to 10096 and 5781 observations of ex dividend log returns of the index respectively. Apart from taking into account the first order serial correlation in expressions (3.1)-(3.2) below, our estimation method corresponds to the one used by RE. The extracted parameters corresponding to the two series are denoted by the subscripts 1 and 2. Since there is some evidence of first-order serial correlation in the series of returns, the variable denoted by u_t is in fact given by the residuals from the regression of the log-returns on a constant and the lagged log-return, as follows:

$$\log(S_{t+1}/S_t) = c + \rho \log(S_t/S_{t-1}) + u_{t+1}. \tag{3.1}$$

The estimated parameters are

$$c_1 = 0.00028283, \quad \rho_1 = 0.1051757, \quad c_2 = 0.000185, \quad \rho_2 = -0.1042065. \tag{3.2}$$

Next, we estimate a GARCH (1,1) model by Gaussian maximum likelihood (ML). Since the ML estimation has been shown not to be robust to the presence of jumps, we wish to filter out observations for which the residuals are too large in absolute value. Then ML estimates can be obtained by re-estimating a GARCH (1,1) model on the sample from which the jumps have been filtered out. We find that if the criterion for deciding that an observation contains a jump finds too many jumps, the second ML estimation either fails to converge or converges to something that allows for negative variances. It appears that, to avoid this, a residual must exceed a threshold of at least 4.5 in absolute value before the corresponding observation is filtered out.

The first ML application yields a grand total of 10 days in the first series, in which there were residuals in excess of 4.5 times the standard deviation. Six of the residuals were negative and four positive. The largest negative return, equal to -10.967 standard deviations, occurred on October 13, 1989. As a proportion of the total, these days with jumps were less than 0.001 of the 10096 observations and are omitted in the parameter estimations. Qualitatively similar results also obtained for the second series, in which the ML application identified 5 days with residuals exceeding 4.4 times the standard deviation, again less than 0.001 of the total. All of them were negative and the largest one occurred on January 4, 2000, equal to -9.462 standard deviations.

For the GJR-GARCH parameters, the estimates from the samples for which the jumps had been filtered out are as follows.

$$\begin{aligned} &b_1 = -0.000251, \alpha_1 = 0.367547 \times 10^{-6}, \gamma_1 = 0.0547, \theta_1 = 0.03109; \delta_1 = 0.9272 \\ &b_2 = 0.93 \times 10^{-4}, \alpha_2 = 0.21461 \times 10^{-5}, \gamma_2 = 0.13191, \theta_2 = 0.009367, \delta_2 = 0.8351. \end{aligned} \tag{3.3}$$

As we verify numerically, the filtering out of the jumps had an insignificant impact on the average volatility of the returns, which remained approximately the same for the full and filtered samples. From (2.3) and (3.1) the risk premium, the excess return of the ex-dividend index over the riskless rate, is under the assumption of stationarity equal to

$$\mathrm{E}\big[\log(S_{t+1}/S_t)\big] - r_t = c/(1-\rho) + b. \tag{3.4}$$

Substituting the parameter values from (3.2) and (3.3) and multiplying by 254 in order to annualize it we get the values of 2.57% for the first sample and 6.62% for the second one. We defer the discussion of these estimates for the following section.

The SPX option data

Our sample consists of end-of-day standard SPX option market bid and ask prices for the third Friday of the month from Optionmetrics, covering the period from 1996/01/04 to 2022/02/18, for a total of 312 cross sections. Table 1 shows the characteristics of the sample, separately for call and put options as functions of the degree of moneyness.

The results show a strong moneyness effect, as expected from earlier studies. Trading for the SPX options is concentrated in the at-the-money (ATM) and out-of-the-money (OTM) zones. The ATM are defined arbitrarily for both types of options as those for which $K/S_t \in [0.98, 1.02]$.

The OTM zones are $K/S_t > 1.02$ for calls and $K/S_t < 0.98$ for puts. We use the volume of trade as an indicator of liquidity. As the table indicates, the volume for puts at 34,803 traded contracts is much larger than the 18,215 traded call contracts. Almost 49% of the call volume is in the ATM zone, as distinct from about 22.5% for puts. The calls also have a non-negligible in-the-money (ITM) volume of trade of 5% in the 0.90-0.98 zone, while for the puts the corresponding proportion is 3% in the 1-02-1-07 zones. The tables also show a strong dependence of the relative bid-ask spread as a proportion of its midpoint on the degree of moneyness, with the spread widening as the options get deeper OTM. Since most empirical studies uses the bid-ask midpoint as a proxy for the "true" frictionless option value, it is not surprising that there is significant uncertainty in the derived results with respect to the tails of the index return distribution.[4]

Table 1 also shows the width of the SD bounds as a function of the degree of moneyness for both calls and puts. This width increases as the corresponding options become more OTM, as shown in Perrakis (2019, p. 29), paralleling the increase in the bid-ask spread.

The Monte Carlo (MC) estimation of the SD upper bounds

The procedure described here is applicable to all options, puts as well as calls, and the same MC procedure can generate all the SD upper bounds in each cross section. The estimations of the bounds by MC are applied separately for the periods 1960-2000 and 2000-2022, using the parameters in (3.2) and the corresponding empirical

[4] See, for instance, Ross (2015) and Anderson, Fusari and Todorov (2017).

distributions. For every day till option maturity T these innovations will be drawn independently from the same distribution $\hat{\varepsilon}_t = z_t \in [z_{\min}, z_{\max}]$.

Assume that we are at the start of the month t, with a maturity T, and that we have filtered out the starting variance σ_t^2. From the corresponding empirical distribution we draw randomly a set of innovations and arrange them in terms of order of appearance as representing the innovations of the days $z_1, \ldots, z_T$. Their P-probabilities $p(z_i)$, $i = 1, \ldots, T$, are found from the empirical distribution, generating a probability of the path in the P-world equal to $\prod_1^T p(z_i)$.

Next, we transform this sequence of independent innovations into a path of ex-dividend index returns $Z_1, \ldots, Z_T$. At each step of the path, we compute $Z_i = \exp(r_i + b + \sigma_i z_i)$, $i = 1. \ldots, T$. Unlike the sequence of innovations, the returns are not mutually independent since the volatility σ_i depends on the innovation z_{i-1}, but the path probability in the P-world continues to be $\prod_1^T p(z_i)$. Since the length of the path from t to T is the same under P and Q, the Q-probability of the path is found from the equilibrium (2.4) and the transformations (2.5)-(2.6). The path of ex-dividend index returns $Z_1, \ldots, Z_T$ remains the same, but the corresponding probabilities are now equal to

$$q(z_i) = \begin{cases} p(z_i) & \text{with probability } \dfrac{\exp(r_i) - Z_{i\min}}{\hat{Z}_i - Z_{i\min}} \\ 1_{z_{\min}} & \text{with probability } \dfrac{\hat{Z}_i - \exp(r_i)}{\hat{Z}_i - Z_{i\min}} \end{cases}, i = 1, \ldots, T. \tag{3.5}$$

The corresponding probability of the path in the Q-world is $\prod_i^T q(z_i)$. Multiplying this probability by the option payoff for a large number of randomly selected paths, adding the results and discounting by the product $\prod_1^T r_t$ of daily riskless return values, yields the option upper bound value.

4. Results: Option Market Data and the SD Upper Bounds

Consistency of the option data with the SD upper bound

Table 2 presents the results of the comparison of the observed option market data in Table 1 with the SD upper bounds estimated by MC and based on the P-parameter estimates (3.1)-(3.4). The results are presented separately for calls and puts. The table focuses on the existence of an (unobservable) frictionless equilibrium price consistent with the P-dynamics, represented by the SD upper bound at the level of each individual option.

We focus separately on calls and puts and for each liquid moneyness zone. Recall that the NAE-based studies assume that the frictionless equilibrium prices of all options in a given cross section lie within the observed bid-ask spread, most often at the spread midpoint. The power of the SD paradigm is that it allows us to test this fundamental feature of NAE, and also to identify the moneyness zones where this feature does not hold.

For the call options, the table shows that the entire ATM zone, as well as the two liquid ITM categories with 2% of the traded volume each, are fundamentally inconsistent with NAE. For a large percentage of the cross sections the entire bid-ask spread lies entirely above the SD bound, implying that these options are overvalued in the frictionless world. Even for those cross sections where the SD upper bound lies within the spread, the midpoint of the spread is within the SD bounds only for a small percentage, less than 20% of the total. Overpricing is also observed, albeit less pronounced, in the highly liquid OTM zones 1.02-1.04 and 1.04-1.07, that account for 20% and 17% of the traded volume respectively. Even at the deep OTM zone of 1.04-1.07 where in a majority of the cross sections the SD upper bound lies within the spread, the bid-ask midpoint lies below the SD upper bound for only 18% of the cross sections, in spite of the fact that both the SD bounds and the quoted spread widen. Since a monotone kernel obviously cannot pass through the bid-ask spread for a majority of the cross sections, we conclude that a NAE frictionless equilibrium cannot be extracted from the option data.

The overpricing results are, if anything, even more pronounced for the put options. The main difference with the calls lies in the fact that the option data is much more concentrated in the OTM and deep OTM zones, that account for about 72% of the total traded volume, almost twice what is observed for the corresponding OTM zones for the calls. This OTM concentration of put option trading largely accounts for the much larger volume of trade in put rather than in call options, since the ATM and ITM volumes are comparable in size. The inconsistency of the observed data with the SD upper bound is observable across the entire OTM and ATM moneyness zones, for which a majority of the cross sections show no overlap with the bid-ask spread of the options. Only for the 3% of ITM trades is there such an overlap for a majority of the cross sections, but even there the spread midpoint lies outside the SD bounds.

These results imply that the observed option market data are inconsistent with a frictionless equilibrium value of the options that relies on the estimated dynamics of the index. As noted in the introduction, such an inconsistency can be tested in the frictionless world, by comparing the realized time series of returns of the IT and OT portfolios. The appropriate strategies were described for individual options in equations (2.7)-(2.13). We describe below their implementation.

The ex-post SD tests in the frictionless world

For expository purposes and because of the fact that the option market is obviously partially segmented between puts and calls given the large discrepancy in volume, the tests are carried out separately for call and put options, following (2.7)-(2.9) and (2.10)-(2.13) respectively. Starting with the calls, let $C_{bt}^{ji} > \bar{C}(S_t, K_j^i, T)$, $j = 1, \ldots, N_i$ denote one of the N_i overpriced call options in a given cross section of the subset $i = 1, \ldots, I$ of the cross sections that have overpriced call options. We select the OT portfolio by aggregating the total OT excess return over ranges of moneyness as follows, with the weights specific to the cross section:

$$\Big\{\sum_{j=1}^{N_i} w_j^i \big[C_{bt}^{ji} - \bar{C}(S_t, K_j^i, T)\big]\Big\}/S_t^i,\ w_j^i \geq 0,\ \sum_{j=1}^{N_i} w_j^i = 1. \tag{4.1}$$

Once the optimal weights w_j^{i*} have been determined at time t they are kept constant till maturity. At each time $\tau \in [t, T-1]$ we define the optimal allocation $\beta_{j\tau}^{i*}$, $j = 1, \ldots, N_i$, as in (2.8). At $T-1$ we verify the terminal condition (2.9)] for the entire portfolio of short call options.

$$\sum_{j=1}^{N_i} w_j^{i*} C_{bt}^j \prod_{k=t}^{T-1} [\beta_k^* R_k + (1 - \beta_k^*) Z_{k+1}] - \sum_{j=1}^{N} w_j^{i*} (S_T^i - K_j^i)^+ \geq 0 \tag{4.2}$$

Dividing the right-hand side of (4.2) by the ex-dividend index value S_t^i, we transform the OT option payoff into an excess return for the subset of cross sections $i = 1, \ldots, I$ that have overpriced calls. This allows the comparison of the time series of IT and OT returns that verify the SD, as described in the following subsection.

An identical procedure also holds for the put options. The optimal weights w_j^{i*} are determined at time t from the following relation, replacing (4.1):

$$\Big\{\sum_{j=1}^{N_i} w_j^i \big[P_{bt}^{ji} - \bar{P}(S_t, K_j^i, T)\big]\Big\}/S_t^i,\ w_j^i \geq 0,\ \sum_{j=1}^{N_i} w_j^i = 1. \tag{4.3}$$

The key optimal parameters $\beta_{j\tau}^{i*}$, $j = 1, \ldots, N_i$, are determined from (2.10)-(2.12), while the cumulative proceeds at maturity are determined from the following application of (2.13):

$$\begin{aligned}\sum_{j=1}^{N_i} w_j^i \Big[\beta_{jt}^{i*} S_t^i [\prod_{i=t}^{T-1} R_i - \prod_{i=t+1}^{T} Z_i] + \sum_{i=t}^{T-2} \Delta\beta_{j,i+1}^{i*} S_{t+1}^i [\prod_{i=t}^{T-2} R_i - \prod_{i=t+2}^{T} Z_i] \\ + P_{bt}^j \prod_{i=t+1}^{T} Z_i - (K_j^i - S_T^i)^+\Big] \geq 0. \end{aligned} \tag{4.4}$$

As with the call, dividing the right-hand side of (4.4) by S_t^i we get the excess return of OT over IT for the given cross section. The time series of the returns of the cross sections with overpriced puts is the input for the next phase of the empirical work.

Results: the ex-post Davidson-Duclos tests in the frictionless world

Tables 3 and 4 below show the results of the frictionless DD tests for calls and puts respectively. These tests add the time series of the excess returns of the overpriced calls and puts to the returns of the index and compare them to the returns of the index. For the call options, in spite of the widespread overpricing in many cross sections in the liquid ATM and OTM zones the average realized excess returns are small and quite volatile, as shown in the first panel of Table 3. Only the OTM zones survive these stringent tests with their non-dominance null, which are significant only after 5% and 10% trimming in the right tail of the realized returns, in addition to the 10% trimming in the left tail that is always applied.

By contrast, the results in Table 4 are clearcut and reject the non-dominance null in all OTM zones unequivocally and for all aggregations. The inclusion of the two ATM zones to the OTM weakens and eventually nullifies the aggregate results. We conclude that the OTM put option price data is fundamentally inconsistent with the existence of a frictionless competitive equilibrium within the observed spreads. The difference between the call and put results is probably due to the much larger percentage of calls than puts in the ATM zone of [0.98,1.02] accounting for 49% as compared to 25%.

The CCP algorithm for the market with frictions

In Appendix II we describe the search algorithm applied to each cross section in which there is no monotone kernel passing through the bid-ask spread of the options. We verify this condition by a linear program (LP), initially presented in Constantinides, Jackwerth and Perrakis (CCP, 2009) and repeated in Perrakis (2019, pp. 179-180). We define the kernel $Y(S_T)$ as cum dividend till maturity and assume a single trading period $[t, T]$. Letting now δ denote the dividend rate, we check whether $Y(S_T)$ passes through the bid-ask spread of both calls and puts. In CCP we considered an interval $0.6S_t \leq S_T < 1.15S_t$, with the left tail obeying the restriction $\min\{\min(Z_T), K/S_t, 0.6\}$; it can be easily seen that it includes virtually the entire universe of liquid options. We then choose a strike price near ATM, or $K_j \approx S_t$ and examine the feasibility of the following LP:

$$
\begin{aligned}
&R_t^{-1}\begin{cases}\max_Y \mathrm{E}_t\left[Y(S_T)\left(S_T/(1+\delta)-K_j\right)^+\right]\\ \min_Y \mathrm{E}_t\left[Y(S_T)\left(S_T/(1+\delta)-K_j\right)^+\right]\end{cases} \text{subject to non-increasing } Y(S_T) \text{ and}\\
&R_t(1+k)^{-1} \leq \mathrm{E}_t[Y(S_T)] \leq (1-k)^{-1},\ S_tR_t(1-k) \leq \mathrm{E}_t[Y(S_)S_T] \leq S_tR_t(1+k)\\
&C_{ib} \leq \mathrm{E}_t\Big[Y(S_T)\big(S_T/(1+\delta)-K_i\big)^+\Big] \leq C_{ia}, \qquad (4.5)\\
&P_{ib} \leq \mathrm{E}_t\Big[Y(S_T)\big(K_i - S_T/(1+\delta)\big)^+\Big] \leq P_{ia}\\
&\text{for } i = 1,\dots,j-1, j+1,\dots N,\ i \neq j
\end{aligned}
$$

Those cross sections for which (4.5) is infeasible will be subjected to the search algorithm described in Appendix II.

Once this maximum feasible value for $\hat{S}$ is found, we partition the segment $[S_t, \hat{S}]$ and maximize the excess return to OT for each value of this partition to find the complete set $\Omega(\hat{S})$. Finally, the optimal portfolio is defined as the one for which a given criterion selection reaches its supremum. We consider six such criteria, the Sharpe ratio (SR), the Gain/Loss ratio (GLR), the Sortino ratio (SR), the maximum feasible , the Information Ratio (IR) and the maximum expected excess return.

Results: the ex-post Davidson-Duclos tests in the world with frictions

For this application we use a somewhat different fitting of the GARCH dynamics in a more restricted data that covers the entire periods 1960-2000, as well as from January 6, 2000 to March 25, 2022, with 10098 and 5595 observations respectively. In this slightly shorter total period, there are 302 option cross sections. There are 11 jumps, all negative, in the first subperiod, and 6 negative jumps in the second subperiod. The estimations are slightly more accurate but not significantly different from the previous larger data set. Instead of (3.3) we now have the following parameter estimates:

$$\begin{aligned} b_1 &= -0.000164, \rho_1 = 0.14287, \alpha_1 = 3.6728 \times 10^{-7}, \\ &\gamma_1 = 0.04635, \theta_1 = 0.02475; \delta_1 = 0.93664 \\ b_2 &= 0.000353, \rho_2 = -0.05886, \alpha_2 = 0.1961 \times 10^{-5}, \\ &\gamma_2 = 0.04473, \theta_2 = 0.1271, \delta_2 = 0.87218. \end{aligned} \tag{4.6}$$

With these two sets of parameters we apply the CCP search algorithm described in Appendix II. Out of 302 dates there were 300 dates with violations of the kernel criteria, the infeasibility of the system (4.5). The portfolios are scaled in terms of one option per unit index and are derived in terms of the maximization of each one of the above six criteria, all expressed in terms of excess returns for the OT portfolios, whose significance is assessed on the basis of the DD tests, shown in Table 5.

The columns of the table show the average excess return of the optimal portfolio as per the corresponding criterion, the bootstrapped P value of the excess return, the volatility of the excess return, the IR of the excess return, and the DD test P values under the three conditions, no trimming at the right tail, 5% trimming and 10% trimming. On average, the OT trades contained 0.74 puts and 0.26 calls, a proportion that is almost the opposite of the one observed in the CCP Table 3 proportions. The OT portfolios contained short index positions in only 9 cross sections.

We note that all criteria have excess OT returns close to each other, but their volatilities vary considerably. Trimming does not seem to have much impact on the significance of the results, unlike the CCP Table 2 results; this is probably due to the different compositions of the OT option portfolios, since the right tail has little effect on put option payoffs and the portfolios contain many more puts than calls in our data, unlike CCP. In two cases the rejection of the non-dominance null is not significant at the 5% and 10% levels, while it is strongly significant in three and significant at better than 9% for the IR criterion. We conclude that the option prices are at the very least inconsistent with the competitive standard in the presence of frictions. We discuss plausible economic interpretations of our results in the following section.

Last, we note that the CCP algorithm, when it yields significant results in terms of the DD tests for any criterion, will automatically confirm the option mispricing at the frictionless level as well. Indeed, the derived OT portfolios contain a mixture of long and short positions in call and put options, purchased (written) at the corresponding ask (bid) prices. In the frictionless world these option positions would be established at the bid-ask midpoints of the options. Hence, the OT trader will collect more money from the short positions and pay less money for the long ones. We carried this exercise in the Table 5 results and observed that the gain shown in the second column of the average excess return of OT for the various criteria varies from a low of 0.17 to a high of 0.27, improving slightly the significance of the DD test statistics.

5. Conclusions

In this paper we have presented an alternative modelling of the risk-neutral GARCH dynamics for the GARCH models on the S&P 500 index returns with empirical innovations pioneered by RE. Our fundamental contributions are twofold. First, we examine separately option valuation in a frictionless world, which has been the focus of virtually the entire literature on empirical option research, from the observed option market data that is far from frictionless. Second, we show that with the same dynamics the observed market data is mispriced in the market with frictions as well, insofar as zero net cost option portfolios generate SD of OT over IT as confirmed by powerful out of sample tests.

For the frictionless world, using the SD theory for frictionless markets introduced by studies published in the 1980s but not well known and never before applied empirically, we derive admissible bounds for frictionless option prices for the RE and BEM GARCH dynamics. These bounds are then compared to observed option market bid and ask prices. We observe that a large proportion of these prices in the liquid moneyness zones for both calls and puts lies entirely above the corresponding SD upper bounds for most cross sections. We then develop dynamic strategies for exploiting these overpriced options that involve shorting them at their bid prices and allocate the proceeds to the index and the riskless asset for calls, and to the index for puts hedged by short positions in the index, both suitably rebalanced along every node of the discretized path to maturity. Applying these strategies separately for calls and puts at every overpriced cross section, we show with the out-of-sample DD tests that OT dominates IT over the entire sample of observed cross sections. We conclude that a frictionless option market equilibrium consistent with the estimated index dynamics cannot be extracted from the observed market data.

For the market with frictions, since there is obviously no monotone pricing kernel passing through the option bid-ask spread of most cross sections, we apply a slightly modified version of the CCP (2020) search algorithm based on LP that identifies profitable zero net cost OT portfolios. After standardizing these portfolios to one option per unit index, we select the portfolio that optimizes six different performance standards. In all cases the OT excess return is positive, and the out of sample tests reject the non-dominance null in four of the 6 cases, three of them strongly significant. Setting the bid-ask spread to 0 for these portfolios increases the expected payoffs, although the improvement is modest due to the tightness of the spreads.

How does one interpret these results? First of all, we have presented strong proofs of the incompatibility of the observed option market data with the index dynamics in the frictionless world, with the data overpricing the options. This explains why in RE the fitted non-monotone kernel, which is inconsistent with economic theory, gives a better fit to the data than the monotone one. Second, we have shown that there are profitable zero net cost portfolios in the real market, with the observed bid-ask spreads. The obvious question that arises is why such portfolios exist and have not been realized by new entrants.

This question requires further research that transcends the scope of this paper. One possible explanation is that the portfolios cannot be executed, since there may not be sufficient depth in the standing quotes of the options in the zero-net-cost portfolios. This will certainly put these portfolios out of reach for retail investors. In such a case the option market equilibrium is not consistent with perfect competition, and a model that relies on market power is certainly a possibility. This is also buttressed by preliminary data on the market maker shares in these SPX monthly options that are quite high, implying that a large percentage of the trades takes place at the actual quoted prices. A formal modelling of the intermediate market equilibrium is certainly a worthwhile project.

Appendix I

Estimation of (2.2)

We rewrite (2.2) for our GARCH(1,1) model as follows:

$$\begin{aligned} \sigma_t^2 &= \alpha + (\gamma + \theta I_{t-1})u_{t-1}^2 + \delta\sigma_{t-1}^2 \quad \text{or} \\ \sigma_t^2 &= (1-\delta\mathrm{L})^{-1}(\alpha + (\gamma+\theta I_{t-1})u_{t-1}^2) \end{aligned} \tag{I.1}$$

where L is the lag operator. (I.1) can be implemented if the infinite series implied by it can be truncated by specifying the value of σ_1^2 then (I.1) becomes

$$\sigma_t^2 = \delta^{t-1}\sigma_1^2 + \frac{\alpha(1-\delta^{t-1})}{1-\delta} + \gamma\sum_{j=1}^{t-1}\delta^{j-1}u_{t-j}^2 + \theta\sum_{j=1}^{t-1}\delta^{j-1}u_{t-j}^2 I_{t-1} \tag{I.2}$$

The model can be estimated by Gaussian Maximum Likelihood with the assumption that the ε_t are normally distributed. Bollerslev and Wooldridge (1992) give conditions that allow this technique to provide consistent parameter estimates even when the true innovation density is not normal. Further, with a long time series of the sort we use here, making the simplification that $\delta^{t-1} = 0$ has very little effect on the ML estimates. Then (I.2) simplifies to

$$\sigma_t^2 = \frac{\alpha}{1-\delta} + \gamma\sum_{j=1}^{t-1}\delta^{j-1}u_{t-j}^2 + \theta\sum_{j=1}^{t-1}\delta^{j-1}u_{t-j}^2 I_{t-j} \tag{I.3}$$

The contribution from observation t to the loglikelihood is

$$\ell_t = -\frac{1}{2}\Big(\log(2\pi\sigma_t^2) + \frac{u_t^2}{\sigma_t^2}\Big),$$

so that

$$\frac{\partial \ell_t}{\partial \sigma_t^2} = -\frac{1}{2}\Big(\frac{1}{\sigma_t^2} - \frac{u_t^2}{\sigma_t^4}\Big) \quad \text{and} \quad \frac{\partial \ell_t}{\partial u_t} = -\frac{u_t}{\sigma_t^2}.$$

Now

$$\frac{\partial \sigma_t^2}{\partial \beta} = -\sum_{j=1}^{t-1} \frac{\partial \sigma_t^2}{\partial u_{t-j}} = -2\gamma \sum_{j=1}^{t-1} \delta^{j-1} u_{t-j},$$

and so

$$\frac{\partial \ell_t}{\partial \beta} = \frac{u_t}{\sigma_t^2} + \gamma\Big(\frac{1}{\sigma_t^2} - \frac{u_t^2}{\sigma_t^4}\Big)\sum_{j=1}^{t-1} \delta^{j-1} u_{t-j}. \tag{I.4}$$

The following relationships can be shown:

$$\begin{aligned}
\frac{\partial \ell_t}{\partial \alpha} &= \frac{\partial \ell_t}{\partial \sigma_t^2}\frac{\partial \sigma_t^2}{\partial \alpha} = \frac{1}{2}\Big(\frac{u_t^2}{\sigma_t^4} - \frac{1}{\sigma_t^2}\Big)\frac{1}{1-\delta}, \\
\frac{\partial \ell_t}{\partial \gamma} &= \frac{\partial \ell_t}{\partial \sigma_t^2}\frac{\partial \sigma_t^2}{\partial \gamma} = \frac{1}{2}\Big(\frac{u_t^2}{\sigma_t^4} - \frac{1}{\sigma_t^2}\Big)\sum_{j=1}^{t-1} \delta^{j-1} u_{t-j}^2, \\
\frac{\partial \ell_t}{\partial \theta} &= \frac{\partial \ell_t}{\partial \sigma_t^2}\frac{\partial \sigma_t^2}{\partial \theta} = \frac{1}{2}\Big(\frac{u_t^2}{\sigma_t^4} - \frac{1}{\sigma_t^2}\Big)\sum_{j=1}^{t-1} \delta^{j-1} u_{t-j}^2 \mathrm{I}(u_{t-j} < 0), \\
\frac{\partial \ell_t}{\partial \delta} &= \frac{\partial \ell_t}{\partial \sigma_t^2}\frac{\partial \sigma_t^2}{\partial \delta} = \frac{1}{2}\Big(\frac{u_t^2}{\sigma_t^4} - \frac{1}{\sigma_t^2}\Big)\Big(\frac{\alpha}{(1-\delta)^2} + \gamma\sum_{j=1}^{t-1}(j-1)\delta^{j-2} u_{t-j}^2.
\end{aligned} \tag{I.5}$$

Setting the partial derivatives (I.4) and (I.5) equal to zero gives the likelihood equations, which are solved by the ML estimates of β, α, γ, δ, and θ.

Appendix II

The Constantinides, Czerwonko and Perrakis (2020) algorithm

The method is presented for a scale of trading of one option per unit of the index. The actual scale of trading for a trader that wishes to apply the method will depend upon the depths of the quotes of individual options that will determine the IT wealth for comparison with OT. We consider two time periods, current time t and maturity T. Let $A(S_T)$ denote the payoff at maturity of a suitably selected zero-net-cost portfolio of call and put options. We want this portfolio to have a positive expected payoff, $E_t[A(S_T)] > 0$. Since these payoffs are in cash, for IT-OT comparison purposes we transform them to index units by dividing them by $1 + k \quad (1 - k)$ when they are positive (negative), with k denoting the transaction cost rate. We limit the left tail of the return at $0.6S_t$ and choose the option portfolio so that it has exactly one zero at some value $\hat{S}$, so that $A(S_T) > 0$ for $0.6S_t \leq S_T \leq \hat{S}$ and $A(S_T) \leq 0$ for $S_T > \hat{S}$. Let $\Omega(\hat{S})$ denote the set of such zero-net-cost portfolios with positive expected payoffs.

We find the set $\Omega(\hat{S})$ by solving the following linear program (LP). Let w_i, $i = 1, \ldots, 2N$, denote the number of options O_i , both calls and puts, entering into the OT portfolio from the N available options in a given cross-section ordered in ascending strike price. We treat long and short option positions as separate options, allowing us to restrict the total option position linearly. We also include a short position in the index as an extra option, yielding $2N + 1$ options in the OT portfolio. Then if Π denotes the value of the option portfolio, we must have

$$\Pi = \sum_{i=1}^{2N+1} w_i O_i, \quad 0 \leq w_i \leq 1. \tag{II.1}$$

Let also $g_i(S_T)$ denote the payoff of the i^{th} option, the total payoff at expiration is equal to $-\Pi R_t + \sum_{i=1}^{2N+1} w_i g_i(S_T)$. We then have:

$$A(S_T) = \begin{cases} -\Pi R_t + \dfrac{\sum_{i=1}^{2N+1} w_i g_i(S_T)}{1+k}, S_T \leq \hat{S} \\ -\Pi R_t + \dfrac{\sum_{i=1}^{2N+1} w_i g_i(S_T)}{1-k}, S_T > \hat{S} \end{cases} \tag{II.2}$$

$A(S_T)$ is a piecewise linear function with constant slope $\partial A/\partial S_T$ in any interval $[K_j, K_{j+1}]$ of two successive strike prices K_j, $j = 1, \ldots, N$, of the available strike prices in the option cross-section. We add the fundamental SD constraints:

$$A(S_T) > 0 \text{ for } 0.6S_t \leq S_T \leq \hat{S}, \quad A(S_T) \leq 0 \text{ for } S_T > \hat{S}, \quad \mathrm{E}[A(S_T)] > 0. \tag{II.3}$$

These constraints need only be verified at the strike prices to the left of $\hat{S}$, while at the right, we simplify the search by adding the constraint that the payoff should be nonincreasing. Finally, we find the OT portfolio by solving the following LP:

$$\max_{w_i}\{\mathrm{E}_t[A(S_T)]\} \quad \text{given } \hat{S}, \text{ subject to (II.1)-(II.3).}$$

If this program is feasible, then the set of optimal weights and corresponding options $\{w_i^* \neq 0, O_i^*\}$ belongs to the ex-ante stochastically dominant set $\Omega(\hat{S})$ of OT portfolios. In the search we vary $\hat{S}$ until the LP becomes infeasible for some maximum value of $S_T = \hat{S}$, arbitrarily restricted to $1.15S_t$.

Tables

Table I

DTM=28 Days, Vol>0, k/s>.6, bid>.5

Calls

K/S	0.6-0.84	0.84-0.90	0.90-0.93	0.93-0.96	0.96-0.98	0.98-1.0	1.0-1.02	1.02-1.04	1.04-1.07	1.07-1.40	Total
# Contracts	440	519	636	1,223	1,320	1,761	1,854	1,804	1,637	1,126	12,320
% Contracts	3.6%	4.2%	5.2%	9.9%	10.7%	14.3%	15.0%	14.6%	13.3%	9.1%	100%
Avg IV	49.2%	30.8%	25.3%	21.7%	18.6%	16.8%	15.2%	14.1%	15.9%	26.3%	19.4%
Avg Moneyness	0.76	0.88	0.92	0.95	0.97	0.99	1.01	1.03	1.05	1.12	1.00
Avg Mid-Quotes	474.76	250.66	176.17	120.88	77.73	49.18	25.57	10.96	6.12	4.73	70.67
Relative BA Spread	0.01	0.02	0.02	0.03	0.04	0.05	0.07	0.14	0.18	0.24	0.09
Relative Bound Spread	0.00	0.01	0.02	0.03	0.06	0.11	0.25	0.45	0.38	0.16	0.20
Avg Strike	1,468	1,745	1,853	1,946	2,013	2,147	2,247	2,291	2,435	2,392	2,153
Sum Volume /1000	183	85	107	350	436	3,322	5,582	3,642	3,014	1,493	18,215
% Volume	1%	0%	1%	2%	2%	18%	31%	20%	17%	8%	92%
Sum Open Interest /100C	3,276	3,842	6,395	14,136	18,684	27,708	28,710	24,798	20,038	13,558	161,147
% Open Interest	2%	2%	4%	9%	12%	17%	18%	15%	12%	8%	92%

Puts

K/S	0.6-0.84	0.84-0.90	0.90-0.93	0.93-0.96	0.96-0.98	0.98-1.0	1.0-1.02	1.02-1.04	1.04-1.07	1.07-1.40	Total
# Contracts	2,735	3,502	2,496	2,664	1,848	1,900	1,644	1,092	853	1,043	19,777
% Contracts	13.8%	17.7%	12.6%	13.5%	9.3%	9.6%	8.3%	5.5%	4.3%	5.3%	100%
Avg IV	46.5%	30.3%	24.8%	21.5%	19.1%	17.2%	15.5%	15.9%	18.3%	35.1%	26.0%
Avg Moneyness	0.78	0.87	0.92	0.95	0.97	0.99	1.01	1.03	1.05	1.17	0.94
Avg Mid-Quotes	2.75	4.62	7.55	12.48	19.76	30.11	46.21	69.05	103.67	297.58	36.39
Relative BA Spread	0.20	0.23	0.17	0.12	0.08	0.06	0.05	0.05	0.04	0.02	0.13
Relative Bound Spread	1.64	1.47	1.20	0.90	0.57	0.27	0.07	0.02	0.01	0.00	0.85
Avg Strike	2,106	2,043	2,096	2,136	2,154	2,186	2,163	2,013	1,922	2,010	2,096
Sum Volume /1000	2,728	5,470	6,121	6,026	4,460	6,001	2,865	591	228	314	34,803
% Volume	8%	16%	18%	17%	13%	17%	8%	2%	1%	1%	99%
Sum Open Interest /100C	28,662	45,395	38,709	46,561	33,503	30,287	16,902	8,988	7,625	11,993	268,623
% Open Interest	11%	17%	14%	17%	12%	11%	6%	3%	3%	4%	96%

Table 2

Calls

K/S	0.6-0.84	0.84-0.90	0.90-0.93	0.93-0.96	0.96-0.98	0.98-1.0	1.0-1.02	1.02-1.04	1.04-1.07	1.07-1.40	Total
Mid-quotes within	18	155	189	245	91	170	256	301	300	70	1,795
Bounds	4%	30%	30%	20%	7%	10%	14%	17%	18%	6%	15%
Overpriced Contracts	136	213	308	770	1,066	1,454	1,213	742	288	43	6,233
	31%	41%	48%	63%	81%	83%	65%	41%	18%	4%	51%
Cross-sections > 1	73	96	141	219	265	263	229	176	105	26	289
Option Overpriced	38%	43%	52%	74%	90%	89%	77%	59%	39%	20%	97%
Cross-sections >50%	66	90	133	190	237	253	213	152	88	21	193
Options Overpriced	34%	40%	49%	64%	81%	85%	72%	51%	32%	16%	65%
Cross-sections 100%	63	88	129	169	217	234	190	123	74	19	22
Options Overpriced	32%	39%	48%	57%	74%	79%	64%	41%	27%	14%	7%

Puts

K/S	0.6-0.84	0.84-0.90	0.90-0.93	0.93-0.96	0.96-0.98	0.98-1.0	1.0-1.02	1.02-1.04	1.04-1.07	1.07-1.40	Total
Mid-quotes within	154	1,669	1,108	698	166	244	215	200	103	31	4,588
Bounds	6%	48%	44%	26%	9%	13%	13%	18%	12%	3%	23%
Overpriced Contracts	2,518	1,695	1,296	1,859	1,623	1,481	952	248	61	75	11,808
	92%	48%	52%	70%	88%	78%	58%	23%	7%	7%	60%
Cross-sections >1	158	163	194	248	284	269	221	107	31	23	297
Overpriced	84%	58%	66%	84%	96%	91%	74%	36%	12%	10%	100%
Cross-sections >50%	148	149	172	222	270	241	188	81	27	17	198
Overpriced	79%	53%	58%	75%	91%	81%	63%	27%	10%	8%	67%
Cross-sections 100%	123	142	166	192	238	227	145	60	25	15	11
Overpriced	66%	51%	56%	65%	80%	77%	49%	20%	9%	7%	4%

Table 3

Calls

K/S	Avg Ex Ret	Std	Min	25%	50%	75%	Max	Skew	Kurt	%CS Ret>0	#CS	#Contracts
1.02-1.07	0.54%	0.1145	-66.90%	0.96%	2.27%	5.49%	18.65%	-2.94	14.17	81%	176	1,030
1.02-1.40	0.44%	0.1114	-66.90%	0.96%	2.27%	4.82%	18.65%	-3.06	15.12	81%	176	1,073
1.00-1.07	-0.08%	0.1570	-67.90%	-4.69%	3.79%	8.52%	30.65%	-1.58	6.19	66%	229	2,243
0.98-1.07	-1.14%	0.1837	-73.36%	-9.79%	3.06%	10.02%	34.08%	-1.09	4.34	62%	263	3,697
0.93-1.07	-2.44%	0.2085	-82.56%	-13.89%	2.44%	10.08%	50.85%	-0.83	4.04	57%	287	5,533

K/S	Avg Ex Ret		OT $\succ_2$ IT	OT $\succ_2$ IT (5% Trim)	OT $\succ_2$ IT (10% Trim)
1.02-1.07	0.54%	0.2520	0.2651	0.0110	0
1.02-1.40	0.44%	0.2860	0.3007	0.0250	0
1.00-1.07	-0.08%	0.5253	1	1	1
0.98-1.07	-1.14%	0.8429	1	1	1
0.93-1.07	-2.44%	0.9807	1	1	1

Table 4

Puts

K/S	Avg Ex Ret	Std	Min	25%	50%	75%	Max	Skew	Kurt	%CS Ret>0	#CS	#Contracts
0.93-1.00	35.28%	0.3933	-0.85%	15.45%	22.33%	41.29%	401.06%	4.84	36.84	100%	296	4,963
0.93-0.98	54.38%	0.3628	14.74%	37.24%	44.00%	60.63%	401.06%	5.12	40.51	100%	290	3,482
0.84-0.98	52.32%	0.3782	-11.15%	34.85%	41.76%	57.93%	383.47%	4.10	28.53	99%	291	6,473
0.84-1.00	33.11%	0.4073	-31.56%	13.64%	21.13%	41.41%	383.47%	4.00	27.07	97%	297	7,954
0.90-1.02	20.60%	0.4483	-34.09%	-4.27%	6.60%	33.79%	401.85%	3.93	26.84	61%	297	7,211

K/S	Avg Ex Ret		OT $\succ_2$ IT	OT $\succ_2$ IT (5% Trim)	OT $\succ_2$ IT (10% Trim)
0.93-1.00	35.28%	0	0	0	0
0.93-0.98	54.38%	0	0	0	0
0.84-0.98	52.32%	0	0	0	0
0.84-1.00	33.11%	0	0.0751	0.0751	0.0751
0.90-1.02	20.60%	0	1	1	1

Table 5

Criterion	mean	P value	std	IR	no trim	5% trim	10% trim
SHR	4.33	0	4.90	0.88	0.2661	0.2661	0.2661
GL	4.14	0.001	6.85	0.61	0	0	0
SOR	4.02	0.002	6.85	0.59	0	0	0
max zh	4.33	0	4.90	0.88	0.2651	0.2651	0.2651
IR	4.05	0.001	6.51	0.62	0.0853	0.0847	0.0798
max E[r]	4.46	0.001	6.98	0.64	0.0001	0.0001	0.0001

References

Babaoglou, K., P. Christoffersen, S. Heston and K. Jacobs, (2018), "Option Valuation with Volatility Components, Fat tails and Non-Monotonic Pricing Kernels", *Review of Asset Pricing Studies* 8, 183-231.

Barone-Adesi G., R. F. Engle, and L. Mancini, (2008), "A GARCH Option Pricing Model with Filtered Historical Simulation," *Review of Financial Studies* 21, 1223-1258.

Barone-Adesi, G., N. Fusari, A. Mira, and C. Sala, (2020), "Option Market Trading Activity and the Estimation of the Pricing Kernel: a Bayesian Approach", *Journal of Econometrics* 216, 430- 449.

Beare, B. K., (2011), "Measure Preserving Derivatives and the Pricing Kernel Puzzle," *Journal of Mathematical Economics* 47, 689-697.

Beare, B. K., (2023), "Optimal Measure Preserving Derivatives Revisited", *Mathematical Finance* 33, 370-388.

Bollerslev, T., and J. W. Wooldridge, (1992), "Quasi-maximum likelihood estimation and inference in dynamic models with time-varying covariances", *Econometric Reviews* 11, 143-172.

Christoffersen, P., R. Elkamhi, B. Feunou and K. Jacobs, (2010), "Option Valuation with Conditional Heteroscedasticity and Nonnormality", *Review of Financial Studies* 23, 2139-2183.

Christoffersen, P., S. L. Heston, and K. Jacobs, (2006), "Option Valuation with Conditional Skewness", *Journal of Econometrics* 131, 253-284.

Christoffersen, P., S. L. Heston, and K. Jacobs, (2013), "Capturing Option Anomalies with a Variance-Dependent Pricing Kernel," *Review of Financial Studies* 26, 1963-2006.

Christoffersen, P., and K. Jacobs, (2004), "Which GARCH Model for Option Valuation?", *Management Science* 50, 1204-1221.

Christoffersen, P., K. Jacobs and C. Orthanalai, (2012), "Dynamic Jump Intensities and Risk Premiums: Evidence from S&P 500 Returns and Options", *Journal of Financial Economics* 106, 447-472.

Christoffersen, P., K. Jacobs and C. Orthanalai, (2013), "GARCH Option Valuation: Theory and Evidence", *Journal of Derivatives* 21, 8-41.

Constantinides, G. M., M. Czerwonko, J. C. Jackwerth, and S. Perrakis, (2011), "Are Options on Index Futures Profitable for Risk Averse Investors? Empirical Evidence," *Journal of Finance* 66, 1407-1437.

Constantinides, G. M., M. Czerwonko, and S. Perrakis, (2020), "Mispriced index option portfolios," *Financial Management*, 49, 2, 297-330.

Constantinides, G. M., J. C. Jackwerth, and S. Perrakis, (2009), "Mispricing of S&P 500 Index Options," *Review of Financial Studies*, 22, 1247-1277.

Cuesdeanu, H., and J. C. Jackwerth, (2018), "The Pricing Kernel Puzzle: Survey and Outlook", *Annals of Finance* 14, 289-329.

Davidson, R and J-Y Duclos, (2013), "Testing for Restricted Stochastic Dominance", *Econometric Reviews* 32 84-125.

Davidson, R. and J. G. MacKinnon (2004), *Economic Theory and Methods*, Oxford University Press, New York.

Duan, J.-C. (1995), "The GARCH Option Pricing Model", *Mathematical Finance*, 5, 13-32.

Duan, J.-C., P. Ritchken and Z. Sun, (2006), "Approximating Jump-GARCH Models, Jump Diffusion Processes and Option Pricing", *Mathematical Finance 16*, 21-52.

Duan, J.-C., and J-G. Simonato, (2001), "American Option Pricing under GARCH by a Markov Chain Approximation", *Journal of Economic Dynamics and Control* 25, 1689-1718.

Ghanbari, H., M. Oancea and S. Perrakis, (2021), "Shedding Light in a Dark Matter: Jump-Diffusion and Option-Implied Investor Preferences", *European Financial Management*— 27, 244-286.

Glosten, L. R., R. Jaganathan, and D. Runkle, (1993), "On the Relation between the Expected Value and the Volatility of the Nominal Excess Return on Stocks", *Journal of Finance* 48, 1779-1801.

Heston, S. L., and S. Nandi, (2000), "A Closed-form GARCH Option Valuation Model," *Review of Financial Studies*, 13, 585-625.

Hsieh, K., and P. Ritchken, (2005), "An Empirical Comparison of GARCH Option Pricing Models", *Review of Derivatives Research* 8, 129-150.

Jackwerth, J. C., (2000), "Recovering Risk Aversion from Option Prices and Realized Returns", *Review of Financial Studies* 13, 433-451.

Jouini, E. and H. Kallal, (1995), "Martingales and Arbitrage in Securities Markets with Transaction Costs", *Journal of Economic Theory* 66, 178-197.

Levy, H., (1985), "Upper and Lower Bounds of Put and Call Option Value: Stochastic Dominance Approach", *Journal of Finance* 40, 1197-1217.

Linn, M., S. Shive, and T. Shumway, (2018), "Pricing Kernel Monotonicity and Conditional Information", *Review of Financial Studies* 31, 493-531.

Merton, R. C., (1976), "Option Pricing when the Underlying Stock Returns are Discontinuous", *Journal of Financial Economics* 3, 125-144.

Orthanalai, C., (2014), "Levy Jump Risk: Evidence from Options and Returns", *Journal of Financial Economics* 112, 69-90.

Perrakis, S., (1986), "Option Bounds in Discrete Time: Extensions and the Pricing of the American Put", *Journal of Business* 59, 119-141.

Perrakis, S., (2019), *Stochastic Dominance Option Pricing: An Alternative Approach to Option Market Research*, Palgrave Macmillan.

Perrakis, S., (2022), "From Innovation to Obfuscation: Continuous Time Finance Fifty Years Later", *Financial Markets and Portfolio Management* 36, 369-401.

Perrakis, S. and P. J. Ryan, (1984). "Option Pricing Bounds in Discrete Time", *Journal of Finance* 39, 519-525.

Post, T and I. Longarela, (2021), "Stochastic Arbitrage Opportunities for Stock Index Options", *Operations Research 69*, 100-113.

Ritchken, P. H., (1985), "On Option Pricing Bounds", *Journal of Finance* 40, 1219-1233.

Ritchken, P. H. and S. Kuo, (1988), "Option Bounds with Finite Revision Opportunities", *Journal of Finance* 43, 301-308.

Rosenberg, J V., and R. F. Engle, (2002), "Empirical Pricing Kernels", *Journal of Financial Economics* 64, 341-372.