A Work Project presented as part of the requirements for the Award of a Master's Degree in Finance from the NOVA School of Business and Economics.

The Cost of Hedging Against Downside Market Risk

Stefano Puma and Alessandro Veneroni

Student numbers: 28522, 30422

A Project carried out for the Master's in Finance Program, under the supervision of:

Prof. Dr. Martijn Boons

January 14, 2019

Abstract

What part of the upside would an investor have to give up to obtain some form of continuous loss protection? In other words, what is the implicit cost of setting up a systematic option-based protective strategy on an equity position?

The acquisition of a 15% out-of-the-money put can be financed by selling a call at the same price. Our results suggest that such a strategy on an index can be costly and not necessarily convenient: hedging all drops bigger than 15% on the S&P 500 index starting from 2012 would have required to cap profits at 5.49%, on average.

Keywords: portfolio insurance, costless collar, self-financing, put protection

Note from the Authors

This research project has been carried out jointly by Stefano Puma and Alessandro Veneroni. Working together allowed us to conduct this research project on a topic that we encountered in several occasions during our Master track and that sparked our mutual interest: portfolio insurance techniques, which pertain to the broader subject of loss protection strategies.

Under the guidance of Professor Dr. Martijn Boons, we identified a specific research question that investigates a particular investment strategy, directly expanding some of the intuitions put forward by American economist Ivo Welch.

A lot of effort went into both data handling and the creation of an option price matching model able to handle the specific format of data downloaded from Datastream. As typically happens working together, we were able to benefit from the synergies of having complementary skills as well as different histories of working with option data, thus tackling tasks that would have resulted insurmountable for a single person.

We would like to thank Mr. Luca Giovanelli for assisting us during the initial phase of the project, providing us with useful insights from the asset management industry. His hands-on experience enabled us to gain familiarity with real-world business practices and to understand how portfolio insurance techniques are put in place today.

Table of Contents

1. Introduction

Ever since its inception in the early 1980s, portfolio insurance techniques have witnessed a mix of both popularity and misfortune that have made them largely studied and widely implemented. Blamed by many to be at the roots of the October 19, 1987 market crash, recently these trading strategies have been on the rise again amongst portfolio managers and academics – and although their effectiveness is still a matter under debate, today they keep on being employed, taking a plethora of different forms (see Tufano and Kyrillos, 1995, for a look at the origins of portfolio insurance techniques).

Compared to traditional hedging strategies, portfolio insurance strategies fundamentally differ due to their dynamic nature that offers protections against market downward moves while leaving investors exposed to upward moves – something that is sacrificed with traditional hedging strategies. Whereas traditional hedging typically involves holding protective positions in either future or forward contracts, the vast majority of the portfolio insurance techniques in use today make use of options to achieve different types of protection. For the purpose of this paper, we developed and analyzed the performance of a particular form of option-based portfolio insurance that satisfies given requirements; particularly, the strategy we are interested in must be self-financing and easily implementable.

1.1 Definition of Problem

In an attempt to trace how much of the historical 7% equity risk premium could have been risk compensation for black swan events, American economist Ivo Welch (2016) investigates the cost of deep-below-the-money index put options. Concentrating only on drops in the value of the underlying larger than 30%, Welch makes the claim that a long put strategy would offer a valid protection against disastrous events "for a very small absolute cost".

However, the fundamental question of whether a similar strategy can be easily implemented at a very small cost and for a broader protective effect still remains. A strategy that relies on some form of option-based portfolio insurance (OBPI) to limit exposure not only to extreme adverse market scenarios, but even to smaller sudden drops. Still, the attractiveness of such a strategy for small private investors would depend on factors such as ease of implementation, small transaction costs, and indeed inexpensive put option protection.

The cost of rolling over protective puts can be self-financed by writing out-of-the-money calls at the same price as the puts. However, by doing so, an investor would limit his exposure to potential value gains in the underlying in order to hedge against downward moves larger than a set threshold. This paper addresses the question of what percentage of the upside such a strategy would have historically given up in order to protect against drops larger than 15% in the S&P 500 index. In other words, we are interested in finding out the average moneyness of a call option that can be used to self-finance an OBPI protective strategy on the index.

2. Our Investment Strategy

Although applicable to any equity underlying for which options are frequently traded, this paper focuses on a strategy consisting on the following three simultaneous positions:

i. long on the S&P 500 index;

ii. long on an OTM protective put on the index;

iii. short on an OTM call on the index.

Out-of-the-money put options are dynamic hedging instruments that can protect against downward moves in the underlying bigger than any chosen threshold. The put acts as a form of protection against all drop bigger than x%, granting its buyer the right to sell the index in case its value falls below the strike price. While holding the underlying alone would expose the investor to potentially unlimited losses, the put assures that the maximum loss is always equal to the put's strike price – as long as the protection is rolled over at expiration.

To self-finance the cost of buying the put, a call with the same maturity is sold on the market at the same price of the put, making the strategy self-financing. By doing this, the investor is not explicitly paying the price of the put, but instead he is turning down all profits from upward moves in the underlying above the call's strike price. A crucial remark to keep in mind is that potentially unlimited profits are being capped in exchange for protection against value drops that are not potentially unlimited: at worst, equities can drop to a value of zero but cannot assume negative values.

Note that call options are out-of-the-money for strike prices set *above* the current value of the index; vice versa, put options are OTM for strike prices set *below* the value of the underlying at the time the contract is written.

2.1 The Cashless Collar

Protective strategies consisting of a long put combined with a short call position that use options with the same price and underlying are known as "zero-cost collar" or "costless collar" strategies. Interestingly, this name has been indicated as misleading by some and "cashless collar" has been suggested instead as a more appropriate definition.

The remark reflects the fact that the acquisition of a put (which typically serves protective purposes on the underlying) is being financed by simultaneously selling a call at the same price. Thus, while no money is explicitly required from the investor at the time of the acquisition, he or she is paying an implicit price: as writer of a call, the investor is obliged to sell the underlying in the event that its value surpasses the call's exercise price.

In other words, the investor is turning down all profits bigger than the moneyness level of the call in order to hedge against all drops greater than the moneyness level of the put.

Given their nature, cashless collar strategies involve OTM options only. The implementation of such strategies in real time, however, comes with a major challenge: finding put and call options that are being traded in the market at the same time and at the same price is a difficult task; because an exact price match is extremely rare, investors typically pick options with irrelevantly small price differences.

In order to pursue a cashless collar strategy over a timespan of few years, every time the put expires a call with the same maturity is sold. In our case, while observations are made at each week of a given month, the strategy assumes a 4-month option rollover: at each point in time, the matching put-call pair has expiration in four months.

Ideally, cashless collars strategies can be successfully implemented for volatile types of underlying that offer some upside potential but that are also prone to sudden drops in the short-term.

2.2 Option Margin Requirements

Option trading typically requires investors to deposit and maintain initial and maintenance margins that act as a collateral; such margins vary based on the characteristics of the options, the underlying and the exchange in which the options are traded. Typically, covered call and covered put positions only require holding the underlying and no other margins is required.

Assuming all options in our strategy were bought on the CBOE, the major U.S. option exchange, margins for our strategy can be computed as follows (CBOE, 2000):

i. no initial margin is required on either the long put or the short call: the investor is only required to pay for the long put and the long underlying in full;
ii. an initial margin of 50% on the long underlying position is required;
iii. no maintenance margin is required on either the put or the call; no maintenance margin is required on the underlying since the options are European style.

2.3 The Probability of a Market Downturn

For our strategy, a 15% level of protection has been chosen to reflect an observed annualized S&P 500 standard deviation of 15.82%.

Such value was obtained from monthly index returns ranging from January 1964 to November 2018; beginning-of-month prices were used to compute monthly returns (e.g., return for February = (1st weekly February price / 1st weekly January price) – 1). Over the roughly 54-year period, an annualized volatility of 15.82% was obtained (**Table 4** and **Figure 3, Appendix**).

As shown by **Figure 4, Appendix**, a visual analysis of the plotted probabilities of monthly index returns over the same time window reveals that they are, in fact, normally distributed around a mean value of 0.6486% with a standard deviation of 4.5645%.

Following the empirical rule for normal distributions, the 68.27% of the index returns is expected to lie within one standard deviation from the mean, in the interval [Mean – StDev; Mean + StDev]. Thus, in annual terms, it is reasonable to expect that approximately 68% of all returns from January 1964 to November 2018 will fall within ±15.82% from the annual average. Therefore, a protection level that hedges up to 15% losses has been chosen for our analysis.

3. Data and Methodology

The dataset comprises weekly closing prices for the S&P 500 and for OTM put and call options on the index at matching dates over a 7-year period ranging from 5 January 2012 to 27 September 2018. Prior to 2012, the option data available drastically decreases, thus undermining the validity of results derived on previous years. Each week of a given month, a put expiring at the end of the following third month is matched with a call that has the same price and maturity.

The price matching process starts by selecting the two put options whose strike price is the closest to being 15% out-of-the-money rounded to the nearest five and to the nearest ten[1]. Among the two puts, the one with strike price actually closest to being exactly 15% OTM is matched with a call with the same price and the same maturity. When there is no exact price match, the two call options with the nearest smaller and nearest greater price are selected – among these two, the call whose price is the closest to the put's price is favored.

Because an exact match between the price of the put and that of the call rarely happens, calls are sold at slightly higher or slightly lower prices than the amount needed to buy the protective puts. This gives rise to small amounts of extra-money being either generated or required each time new option contracts are written. Such amounts of money can be easily stored and withdrawn from a bank account; in our case, however, they are so small that they can be ignored without affecting any of the conclusions drawn from our model: for the whole period, negative and positive sums cancel each other out, resulting in an average $0.04 price difference per year and a positive total sum of $14.31 accumulated over the 7 years under analysis (**Table 5, Appendix**).

For every successful price match, the strike price of the call is observed in order to obtain weekly observations about the calls' moneyness.

3.1 Option data

All option data has been downloaded via Thomson Reuters Datastream from the Option Price Reporting Authority (OPRA), which stores data from all major U.S. option exchanges, including the Chicago Mercantile Exchange (CME) and the Chicago Board Options Exchange (CBOE). Compared to other sources, the OPRA provides data for a wider range of strike prices,

[1] The underlying being the value of the S&P 500 at that specific date.

allowing a higher degree of precision in the matching process[2]. Historical prices and implied volatility have been downloaded for European-style dead put and call options expiring at the end of every month for each year in the time span under analysis.

4. Findings

Over the time window under analysis, our model revealed that hedging all drops bigger than 15% in the index value would have required to turn down all upward moves greater than 5.49%, on average. This result indicates that put protection on the S&P 500 is, in fact, expensive: an investor pursuing this strategy would have still been exposed to a significant slice of the downside, while only being able to benefit from a much smaller percentage of the upside – participating in all index moves from –15.00% to +5.49% (see **Table 1** below). Thus, the specific form of cashless collar that we set up is far from being costless as these results carry broader implications: over the 7-year period analyzed, put and call options that were being sold on the market for the same price would average a 9.51% difference in their moneyness.

The figures for the calls' moneyness seem steady over the 7 years (standard deviation of only 0.55% among the 7 annual averages). Remarkably, the put protection seems to become more and more expensive in most recent years.

Possible intuitions as to why puts appear more expensive than calls, *ceteris paribus*, are discussed in paragraph 4.1: "Possible Explanations".

[2] The minimum jump in the option strike price typically consists of $5 for the most commonly traded moneyness. However, the more OTM the options are, the more sporadic the data becomes – with data available only for every $50 of strike price, thus making the price matching process more inaccurate for deep-below-the money options.

Annual average moneyness (OTM)			
	PUT	CALL	Avg. Diff.
2012	15.06%	6.19%	8.87%
2013	14.99%	6.19%	8.81%
2014	14.96%	5.45%	9.51%
2015	15.00%	5.18%	9.82%
2016	15.00%	5.24%	9.76%
2017	14.99%	4.69%	10.30%
2018	14.98%	5.47%	9.51%
Total period	**15.00%**	**5.49%**	**9.51%**

Table 1 - Annual average option moneyness.

Furthermore, not only this protection would have been costly, but also not convenient: from 2012, the index underwent a steady growth, averaging a positive annual return of 15.92% (arithmetic return) associated with an average volatility of 11.58% per year. Most importantly, over the whole period the index never dropped by more than 8.54% from one week to the next one and therefore the 15%-OTM protective put would have never been exercised. (see **Table 6, Appendix**).

On the other hand, the index value surpassed the profit limit set by the calls' strike prices in several occasions during these seven years, as can be seen from **Table 7, Appendix**: depending on whether the strategy would have assumed option roll-over at either the first or the last week of every 4 months, call options would have been exercised either 11 or 14 times overall, respectively.

Figure 1 below shows the corridor created by our option strategy: the upper line represents the profit cap imposed by writing the call options, while the lower line plots the maximum loss limit set by the protective puts. Every upward or downward index move that is inside the corridor is fully experienced by the investor. On the other hand, every profit above the upper bound is given up and every loss greater than 15% is systematically hedged.

Note that setting how often options are rolled over defines how closely the corridor mimics the underlying: rolling over options with shorter maturities – i.e., one or two months – would imply both more frequent transaction costs and a less rigid corridor – that is, because the

15%-OTM put strike price that is based on the monthly value of the index is being updated more often. On the contrary, transaction costs become less frequent and the corridor becomes more rigid for options with longer maturities.

In our case, setting new strike prices for the put options every 4 months allows for the corridor to have some flexibility while still protecting against sudden falls. On the contrary, a strategy that sets the put's moneyness at the beginning of the investment period and then rolls the option over at expiration always with the same initial strike price would result in a completely inelastic corridor represented by two parallel lines, thus completely capping off all upcoming profits.

The tradeoff between higher and lower frequencies of options rollover ultimately depends on transactions costs as well as the investor's expectations about the future performance of the index.

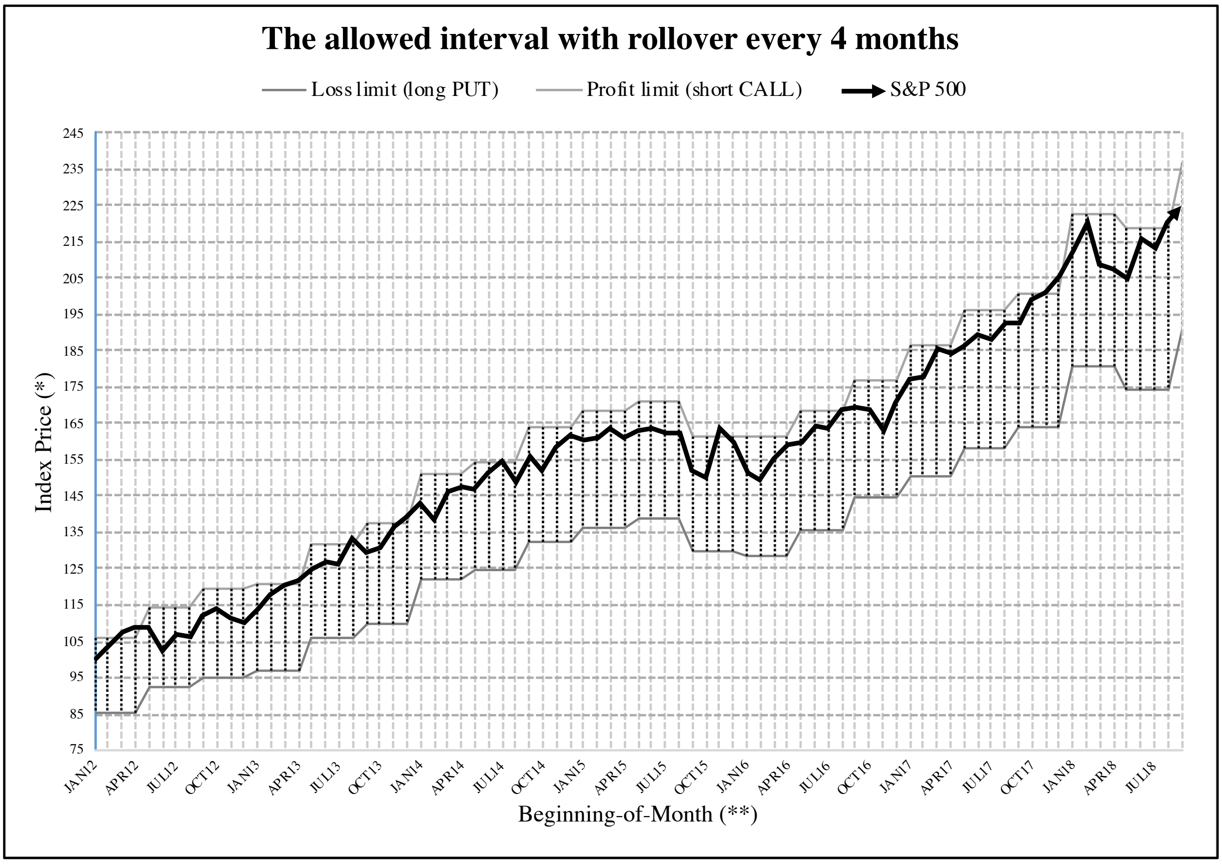

Figure 1 - Our investment strategy with rollover every 4 months. (*)(**)[3]

4.1 Extensions of Results

The results presented in **Table 1** above are derived using real-world option data and represent the outcome of a strategy that makes use of options that were actually traded on the market. However, for the sake of a broader significance of our conclusions, we can obtain the interpolated moneyness for a call's price that exactly matches the price of the 15%-OTM put. By doing this, it is possible to eliminate any noise due to imperfect price matches and obtain a theoretical strike price for the calls[4]. **Table 3, Appendix** shows that the interpolated values do

[3] (*) Indexed at 100 at Jan. 2012. (**) Arithmetic returns and option moneyness are computed at the 1st weekly observation available for each month; for months whose 1st weekly moneyness was not available, the moneyness from the last week of the previous month was plotted.

[4] Note that these newly obtained strike prices are fictitious and not actually being traded on the market.

not differ significantly from the previously obtained values: for a put moneyness that is now exactly 15% OTM for every weekly observation, the average interpolated call moneyness over the same 7-year period is 5.63% instead of 5.49%. The average difference in call and put moneyness decreases from 9.51% to 9.37%.

What if our strategy was to be applied to a broader horizon? Assuming that losses were capped at exactly 15% and that call options would always be exactly 5.49% OTM, we can extend the horizon from January 1964 to November 2018. During this time window, even though we are including the 2007 financial crisis and other times of significant market downturns, the put would have been exercised only 6 times against 93 times for the call[5].

Figure 2, Appendix offer a visual representation of the corridor created by our strategy from September 1984 to November 2018.

4.2 Possible Explanations

What factor could possibly justify the consistent difference observed in the moneyness of put and call options that share the same underlying and maturity and that are priced at the same level by the market? It would seem as market participants value the obligation to sell the index during good times less than the right to sell the index during bad times[6].

Indeed, if we investigate calls that have the same moneyness (and expiration) of the puts we are using for protection, we find out that these are extremely cheap and much more illiquid:

---

[5] That is, if the 4-month option rollover is done on the first week of the month; with a rollover done on the last week of each month, the put would have been exercised 9 times and the call 101 times.

[6] Once again, think of our strategy this way: an investor would still be exposed to losses as big as 15%, but he could only benefit from up to 5.49% profits, on average. In times of high market returns, the short position on the call would force him to sell the (profitable) underlying every time the index value would exceed the call's strike price. *Vice versa*, the protective put would only trigger during bad times, and the investor would be able cap his loss exactly at the put's strike price.

over the whole period considered, these 15%-OTM calls on the S&P 500 were always trading for less than a dollar. Contrarily to put options with the same moneyness, data available for these calls is much more sporadic, suggesting that calls for such moneyness are less frequently traded than puts with similar characteristics. From the first week of January 2012 to the last week of September 2018, a 15%-OTM call traded for an average $0.43. However, for the 352 weeks in our sample, only data for 229 calls was available (see **Table 8, Appendix**).

Could perhaps this be linked to the fact that the S&P 500 index had been performing incredibly well and markets participants were pricing options according to the expectations of a future sustained growth? Intuitively, investors would be less willing to write call options that are around 6% OTM if they anticipate weekly index returns to be well above 6% with high probabilities. If so, the fact that 15%-OTM calls are not as liquid as puts with similar characteristics would not constitute a surprise. To compensate investors for writing these roughly 6%-OTM calls, such options would also have to offer high returns and thus, they would be cheaply available on the market.

Instead, maybe the option pricing process is incorporating some factor that behaves differently for call and for put options. According to traditional asset-pricing theory, options should be priced based on their systematic risk – similarly to any other asset in the economy; more specifically, the six factors directly affecting option prices on which academic research has reached a consensus are the following: the current underlying's price, the option's strike price, its maturity, the volatility of the underlying, the relevant risk-free rate and the present value of the dividends paid by the underlying (in case of a stocks or equity indices).

However, Coval and Shumway (2001) suggested that something besides such factors affects option pricing: in their 2001 paper, the authors claim the importance of an additional volatility factor, which reflects the systematic stochastic volatility of options, in explaining why these assets tend to exhibit anomalous low returns compared to their degrees of systematic risk.

Could indeed the observed difference in option pricing incorporate market's expectations about the future performance of the underlying index?

4.3 Time-series Analysis of Results

To investigate possible time-dependencies that could help explain our results, we defined three sets of variables, shown in **Table 2** below, in order to run five sets of time-series analyses that include up to five lags for each variable. All variables are time-dependent and expressed in percentages.

Regression variables:		
		Unit:
$IV_{PUT,t}$	Implied Volatility of the Put	%
$IV_{CALL,t}$	Implied Volatility of the Call	%
VIX_t	VIX (CBOE S&P 500 Volatility)	%
RV_t	Realised Volatility of the S&P 500 (20-week rolling st. dev.)	%
SPX_t	Arithmetic return of the S&P 500	%
MON_t	Difference between the moneyness of the call and that of the put	%

Table 2 - Regression variables.

The first group of variables keeps track of implied volatility (IV): either that of the call and put that are matched by our model at each week, or that expressed by the VIX index[7]. Returns of the S&P 500 are included either directly in the regressions, or indirectly through their 20-week realized volatility; the choice of which of the two variables to use has been based on which of the two performs the best as an explanatory variable (i.e., which variable returns the higher adjusted R^2). Finally, the observed option moneyness was incorporated by defining

[7] Source: Thomson Reuter Datastream.

a synthetic measure (MON_t) that expresses the difference between the moneyness of the put and that of the call with the same price at each point in time[8].

Using once again weekly data from 5 January 2012 to 30 August 2018, we run five different sets of regressions: each regression set is defined by choosing a dependent variable from one of the three groups defined in **Table 2** above; each set studies the explanatory capabilities of variables from the remaining two groups on the dependent variable. Due to the inclusion of five weekly lags for each variable as well as the use of a 20-week rolling standard deviation, 25 observations from the original dataset were lost: all regression sets include exactly 323 observations at matching dates.

Regression Set n. 1 investigates the effects of past implied volatility and past index returns on the observed moneyness (dependent variable: MON_t; see **Table 9, Appendix**). Noticeably, the explanatory power of the model seems to increase with the number of lags included, but only marginally and up to a point where including an additional lag would not benefit the model's coefficient of determination: for all three regressions in the set, the adjusted R^2 starts decreasing when moving from four to five lags. Because a similar trend is also observed in all other regression sets for a similar number of lags, we limited our analysis to five lags per variable[9].

While past observations do pretty well at explaining the variance of the moneyness, which seems to incorporate information up to four weeks prior and therefore suggesting some predictability, the inclusion of VIX lags and index returns lags does not appear to contribute

[8] This synthetic measure has the advantage of accounting for the *actual* moneyness of the put-call pairs selected by our model: remember that the put moneyness is the closest to – but not necessarily exactly corresponding to being 15% OTM.

[9] The inclusion of additional lags would also be costly in terms of the number of observations: any extra lag included requires sacrificing one observation (the earliest one in the dataset). Note that in case of missing observations for the IV of put and call options, the last previously available data point is assumed.

significantly to the efficacy of the model: the highest possible R^2 value for any number of lags obtained using index returns and VIX values as regressors, 64.84%, is lower than the 64.90% value observed by regressing the moneyness variable on its lags alone.

Thus, although there seems to be evidence of high predictability in past moneyness values, it appears as both past index returns and past expectations about future volatility do not have a significant impact on the observed moneyness factor and their impact as explanatory variables is so marginal that it can be ignored altogether.

Regression Set n. 2 explores possible dependencies between realized index return and the factors from the two remaining groups: while there seems to be very little predictability in S&P 500 returns alone (highest adjusted R^2 achievable of 3.70% for one lag – see **Table 10, Appendix**), moneyness lags do not add significant explanatory power to the model[10]. The highest possible coefficient of determination, 4.78% – still very low – is obtained by regressing index returns on its own lags, on VIX lags and on MON lags for one lag.

On the other hand, Regression Set n. 3 tells a different story: while VIX observed values seem to be highly dependent on their past history (maximum adjusted R^2 of 49.98% for four lags – see **Table 11, Appendix**), they do not display any meaningfully relationship with past observations of the index or the moneyness factor. In this case, including more than three or four lags decreases the predictive capability of the model.

Finally, Regression Sets n. 4 and n. 5 show how the variance of implied option volatility (IV) allegedly display significant predictability features: for calls, observed IV is explained up to 68.50% by its lags alone (for three lags) and the adjusted R^2 remains incredibly high even up

[10] When VIX lags are added to the right-hand side of the equation, the explanatory power of the model increases by roughly 1% – but only for the first weekly lag. The topic of return predictability has long been studied in the literature and it is probably still one of the most debated subjects today. For the purpose of this paper, we are only interested in studying intertemporal dependencies among the factors we defined, and in how outputs change when adding or removing variables to the regression.

to five lags (**Table 12, Appendix**). While the model becomes less efficient when incorporating past index returns as regressors, moneyness lags together with calls' IV own lags still demonstrate exceptionally high explanatory capabilities, with adjusted R^2 values that remain very similar to the ones observed by regressing calls' IV on its own lags alone.

Notably, the trend is reversed in case of puts: IV regressed on its lags still shows very high adjusted coefficient of determination (around 62% for up to five lags – see **Table 13, Appendix**), yet the inclusion of index returns yields very similar adjusted R^2 values and it does not decrease the explanatory power of the model in a significant way. Contrarily to what is observed for calls, however, the moneyness factor generally reduces the adjusted R^2 by approximately 15%, on average.

Could the different sensibility to these two factors provide any clue in explaining the different pricing process that we observed for call and put options sharing the same moneyness levels? All in all, calls' IV appear significantly less susceptible to the index's past performance compared to puts' IV (7.68% on average, for one to five lags). Perhaps this difference indeed solidifies the intuition that calls are less sought-after than puts because investors are expecting the tremendous positive growth experienced by the index in the past to continue indefinitely.

We believe these results constitute an important starting point for any further research aimed at shedding light on the different pricing dynamics that are affecting index options.

5. Model Limitations

Transaction costs. To account for transaction costs, bid prices can be used for all assets on which our strategy holds long positions (i.e., the S&P 500 and the put options on the index), whereas ask prices can be downloaded for the assets that are being short-sold (namely, the call options on the index).

However, weekly data on bid-ask option prices is more prone to inaccuracy compared to closing prices, as put and call prices are available for fewer strike prices. This affects the validity of the model as the matching process between put prices and call prices results in wider gaps between the two; therefore, the protection on the strategy is being financed by selling calls that have a significantly higher or lower price. A replication of our model on a small subsample ranging from 5 January 2012 to 26 September 2013 using bid-ask prices shows that, for puts' moneyness levels that are almost equal to the ones obtained using closing prices (15.01% compared to 15.02%, respectively. See **Table 14, Appendix**), the average call moneyness decreases from 6.19% (using closing prices) to 6.14%.

However, of the 91 weeks in this subsample, 24 are missing, making it hard to tell how much of this difference is actually due to transaction and how much is due to a more imprecise price match, instead.

The matching process. The results derived from our model heavily depend on both the investor preferences and the technical assumptions made; while the former involve the frequency of the option roll-over (4 months in our case) and the percentage of protection desired, the latter include the rounding of strike prices and the matching criteria.

Although some degree of discretionality is unavoidable, we favored results that are as real-world oriented as possible; for this reason, all the option prices used in the model are those of actual end-of-day transactions.

6. Appendix

Annual average moneyness	* = *interpolated value*		
	PUT*	**CALL***	**Avg. Diff.**
2012	15.00%	7.11%	7.89%
2013	15.00%	6.23%	8.77%
2014	15.00%	5.47%	9.53%
2015	15.00%	5.18%	9.82%
2016	15.00%	5.30%	9.70%
2017	15.00%	4.68%	10.32%
2018	15.00%	5.46%	9.54%
Total period	**15.00%**	**5.63%**	**9.37%**

Table 3 - Annual average option moneyness (interpolated).

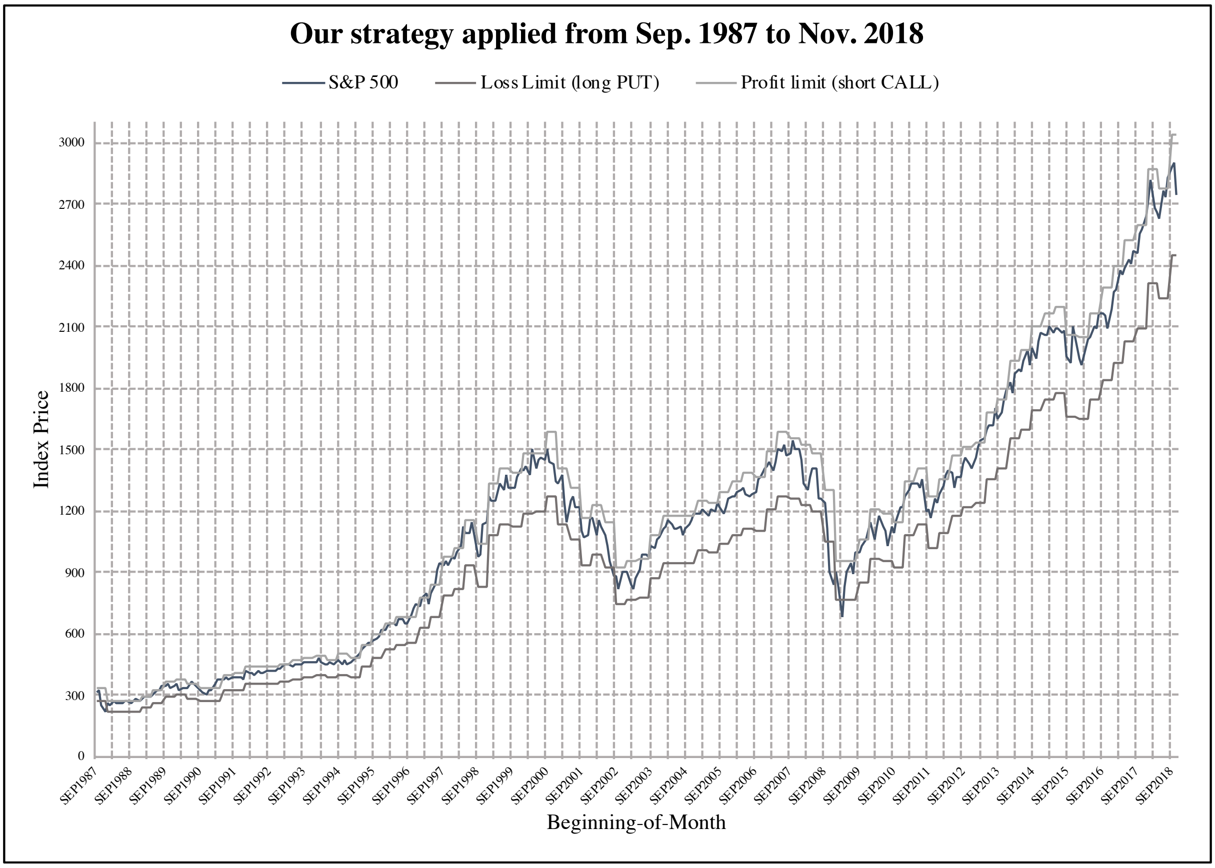

Figure 2 - Our strategy applied from Sep. 1987 to Nov 2018.

S&P 500 Volatility		
	StDev	**Var**
Monthly	4.57%	0.21%
Annualized	15.82%	2.50%
Start month		JAN1964
End month		NOV2018
N. of months in sample		659

Table 4 - S&P 500 volatility.

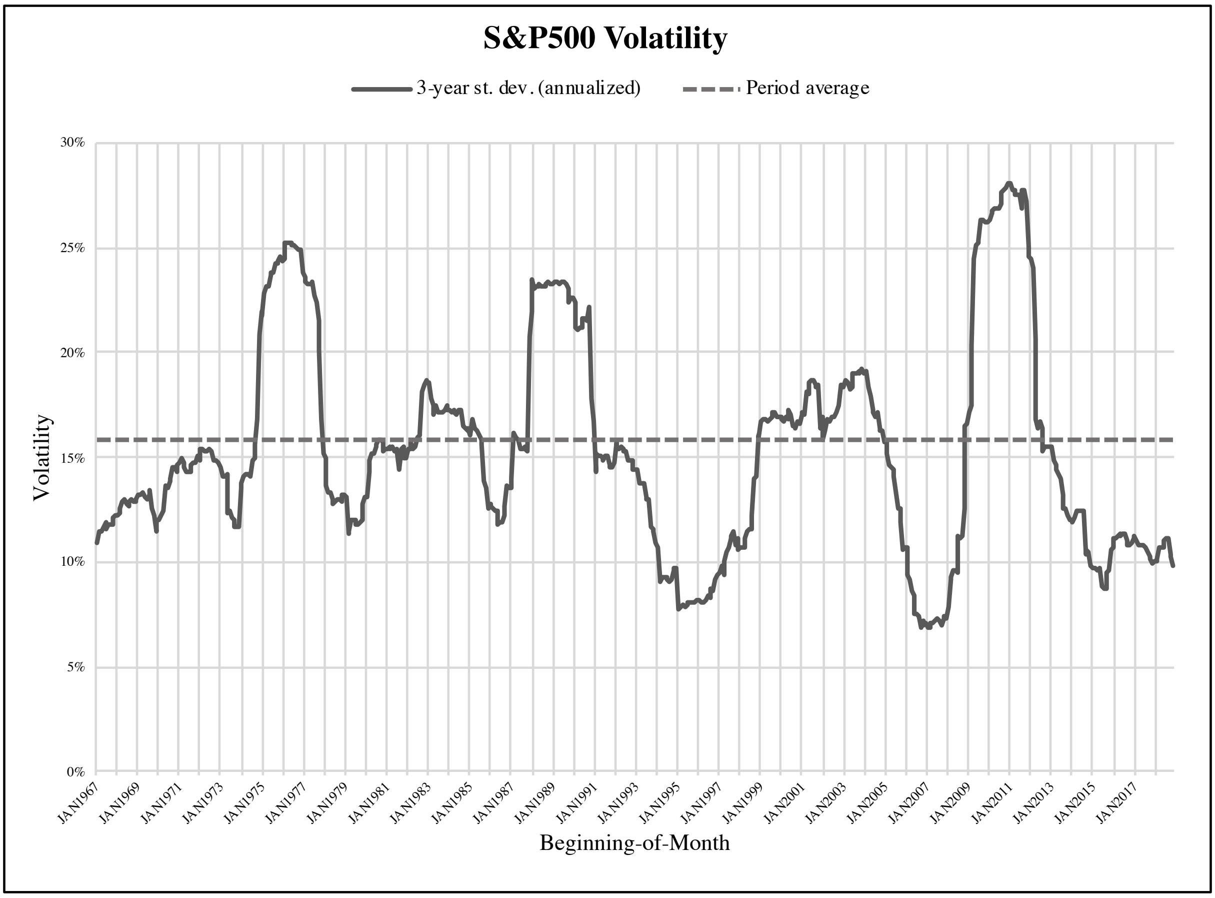

Figure 4 – S&P 500 volatility plotted from January 1967 to November 2018 (using monthly returns).

Call-Put price match		
	Avg. Diff.	**Sum**
2012	0.25	12.95
2013	0.07	3.4
2014	0.03	1.45
2015	-0.02	-1.13
2016	0.08	3.93
2017	-0.15	-6.29
2018	-0.12	-4.41
Total period	**0.04**	**14.31**

Table 5 - Call-Put price match.

Years 2012 - 2018		
		S&P 500
Start date	15-mar-12	1402.60
End date	27-set-18	2914.00
N. of weeks in sample		352
Base		52
Maximum weekly profit		5.82%
Minimum weekly loss		-8.54%
Total period return		107.76%
Annual arithmetic return		**15.92%**
Annual geometric return		**11.41%**
Annualized volatility		**11.58%**

Table 6 - S&P 500 performance statistics.

How many times were the options exercised per year?				
	Beginning-of-month		*End-of-month*	
	PUT	**CALL**	**PUT**	**CALL**
2012	0	2	0	2
2013	0	3	0	2
2014	0	1	0	2
2015	0	1	0	2
2016	0	1	0	2
2017	0	2	0	2
2018	0	1	0	2
Total period	**0**	**11**	**0**	**14**

Table 7 - Option exercise statistics.

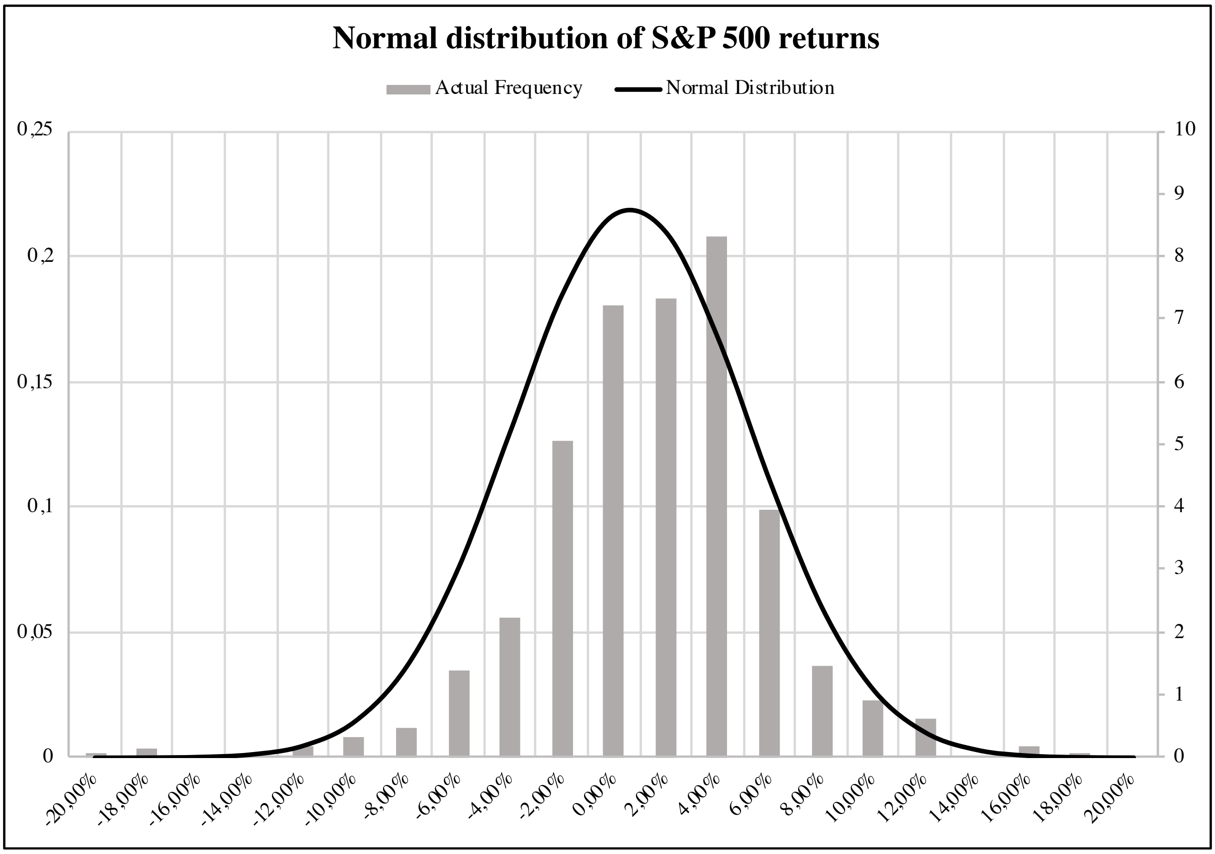

Figure 4 - Normal distribution of S&P 500 monthly returns from January 1964 to November 2018.

15%-OTM Call prices				
	Avg. Price	**N. of missing obs.**	**N. of valid obs.**	**Tot. N. of obs.**
2012	0.45	17	35	52
2013	0.33	0	52	52
2014	0.24	0	52	52
2015	0.28	7	46	53
2016	0.51	30	22	52
2017	0.35	43	9	52
2018	0.85	26	13	39
Total period	**0.43**	**123**	**229**	**352**

Table 8 - 15%-OTM Call prices.

Regression Set 1					
Y:	*MON*				
X:	*MON lags*				
Lags	**1**	**2**	**3**	**4**	**5**
R Square	0.6071137	0.63655754	0.63655931	0.65337028	0.65407801
Adjusted R Square	0.60588976	0.63428602	0.63314137	0.64901016	0.64862183
Standard Error	0.00508353	0.00489697	0.00490463	0.00479738	0.00480003
Observations	323	323	323	323	323
F	496.030278	280.234748	186.240765	149.851369	119.878317
Significance F	4.32198E-67	4.6742E-71	8.845E-70	7.234E-72	6.8417E-71
Y:	*MON*				
X:	*MON lags, RV lags, VIX lags*				
Lags	**1**	**2**	**3**	**4**	**5**
R Square	0.61087016	0.64029741	0.64039184	0.6582935	0.66091338
Adjusted R Square	0.60721064	0.63346761	0.63005167	0.64506616	0.64434563
Standard Error	0.00507501	0.00490245	0.00492524	0.00482426	0.00482915
Observations	323	323	323	323	323
F	166.925932	93.7505889	61.932424	49.7676096	39.8915594
Significance F	4.6641E-65	3.6258E-67	2.8323E-64	5.3657E-65	5.605E-63
Y:	*MON*				
X:	*MON lags, SPX lags, VIX lags*				
Lags	**1**	**2**	**3**	**4**	**5**
R Square	0.61346448	0.63955149	0.64077513	0.66149189	0.66414323
Adjusted R Square	0.60982935	0.63270753	0.63044598	0.64838836	0.64773329
Standard Error	0.00505806	0.00490753	0.00492261	0.00480163	0.0048061
Observations	323	323	323	323	323
F	168.759968	93.447593	62.0356128	50.4819245	40.4720088
Significance F	1.6081E-65	5.0182E-67	2.4019E-64	1.2794E-65	1.331E-63

Table 9 - Summary output for Regression Set n. 1.

Regression Set 2					
Y:	*SPX*				
X:	*SPX lags*				
Lags	**1**	**2**	**3**	**4**	**5**
R Square	0.03999073	0.04081239	0.04166662	0.04166705	0.04375479
Adjusted R Square	0.037000050	0.03481747	0.03265408	0.02961255	0.02867205
Standard Error	0.01597583	0.01599392	0.01601183	0.01603699	0.01604476
Observations	323	323	323	323	323
F	13.3717724	6.80782638	4.62318266	3.45655502	2.90098533
Significance F	0.0002984	0.00127224	0.00350723	0.00877716	0.01409928
Y:	*SPX*				
X:	*SPX lags, VIX lags, MON lags*				
Lags	**1**	**2**	**3**	**4**	**5**
R Square	0.05670054	0.05757352	0.06314184	0.0673448	0.06819927
Adjusted R Square	0.04782938	0.03967935	0.03620343	0.03124202	0.02267154
Standard Error	0.01588574	0.01595359	0.01598243	0.01602352	0.01609424
Observations	323	323	323	323	323
F	6.39156234	3.21744528	2.34393311	1.86536315	1.49797222
Significance F	0.00032302	0.00441435	0.01434193	0.03798881	0.10419284
Y:	*SPX*				
X:	*SPX lags, MON lags*				
Lags	**1**	**2**	**3**	**4**	**5**
R Square	0.04314780	0.04525314	0.04764117	0.04839952	0.05046160
Adjusted R Square	0.03716748	0.03324375	0.02955841	0.02415492	0.02002767
Standard Error	0.01597444	0.01600695	0.01603743	0.01608202	0.01611599
Observations	323	323	323	323	323
F	7.21495783	3.76814502	2.63461783	1.99630115	1.65807073
Significance F	0.00086135	0.00519747	0.01656266	0.04648402	0.08976176

Table 10 - Summary output for Regression Set n. 2.

Regression Set 3					
Y:	*VIX*				
X:	*VIX lags*				
Lags	**1**	**2**	**3**	**4**	**5**
R Square	0.48240696	0.49550883	0.504162680	0.50462468	0.50463203
Adjusted R Square	0.48079452	0.49235576	0.49949963	0.49839355	0.49681865
Standard Error	0.02497988	0.0247002	0.02452579	0.02455287	0.02459139
Observations	323	323	323	323	323
F	299.178349	157.151239	108.118724	80.9843787	64.5856667
Significance F	7.94978E-48	2.8614E-48	2.5946E-48	2.5369E-47	2.4326E-46
Y:	*VIX*				
X:	*VIX lags, RV lags, MON*				
Lags	**1**	**2**	**3**	**4**	**5**
R Square	0.49234018	0.50729241	0.51510618	0.51603402	0.52009839
Adjusted R Square	0.48756595	0.4979372	0.50116355	0.49729985	0.49665042
Standard Error	0.02481645	0.02456404	0.02448498	0.02457962	0.0245955
Observations	323	323	323	323	323
F	103.124516	54.2256725	36.944683	27.545074	22.1809637
Significance F	1.0996E-46	8.9004E-46	2.7624E-44	4.3339E-42	1.714E-40
Y:	*VIX*				
X:	*VIX lags, SPX lags, MON*				
Lags	**1**	**2**	**3**	**4**	**5**
R Square	0.49704332	0.50313884	0.51425197	0.52025917	0.52112522
Adjusted R Square	0.49231332	0.49370477	0.50028478	0.50168855	0.49772743
Standard Error	0.02470123	0.02466736	0.02450654	0.0244721	0.02456917
Observations	323	323	323	323	323
F	105.083154	53.3320938	36.8185556	28.0151857	22.2724118
Significance F	2.5035E-47	3.2991E-45	3.6176E-44	1.159E-42	1.2494E-40

Table 11 - Summary output for Regression Set n. 3.

Regression Set 4					
Y:	*IV CALL*				
X:	*IV CALL lags*				
Lags	**1**	**2**	**3**	**4**	**5**
R Square	0.6793868	0.68520165	0.687984249	0.6883146	0.688676
Adjusted R Square	0.678388	0.68323417	0.68504993	0.68439403	0.68376553
Standard Error	0.01216316	0.01207117	0.01203652	0.01204905	0.01206104
Observations	323	323	323	323	323
F	680.20643	348.261884	234.461428	175.564883	140.246362
Significance F	2.76486E-81	4.8438E-81	2.4855E-80	3.4999E-79	4.1171E-78
Y:	*IV CALL*				
X:	*IV CALL lags, SPX lags*				
Lags	**1**	**2**	**3**	**4**	**5**
R Square	0.52876872	0.5475756	0.56794136	0.57066587	0.57330743
Adjusted R Square	0.52582352	0.54188473	0.55973771	0.55972742	0.55963138
Standard Error	0.01476897	0.01451669	0.01423102	0.01423119	0.01423274
Observations	323	323	323	323	323
F	179.536031	96.2199659	69.2303658	52.1706385	41.9205604
Significance F	5.2175E-53	1.5018E-53	1.5018E-53	2.8201E-53	5.7496E-52
Y:	*IV CALL*				
X:	*IV CALL lags, MON lags*				
Lags	**1**	**2**	**3**	**4**	**5**
R Square	0.68054356	0.68647823	0.68886109	0.690751	0.69152327
Adjusted R Square	0.67854695	0.68253456	0.68295339	0.68287204	0.68163619
Standard Error	0.01216015	0.01208449	0.01207652	0.01207807	0.01210158
Observations	323	323	323	323	323
F	347.34626	177.442438	119.033949	88.8847613	71.7757093
Significance F	6.4637E-81	1.094E-79	4.9821E-78	4.8113E-76	7.9277E-75

Table 12 - Summary output for Regression Set n. 4.

Regression Set 5					
Y:	*IV PUT*				
X:	*IV PUT lags*				
Lags	**1**	**2**	**3**	**4**	**5**
R Square	0.61189411	0.6263323	0.630866929	0.6310956	0.63267178
Adjusted R Square	0.61068506	0.62399688	0.62739546	0.62645529	0.62687796
Standard Error	0.01495958	0.0147016	0.01463501	0.01465346	0.01464517
Observations	323	323	323	323	323
F	506.093865	268.187933	181.728999	136.002986	109.197683
Significance F	6.03497E-68	3.9599E-69	1.0503E-68	1.3967E-67	8.8464E-67
Y:	*IV PUT*				
X:	*IV PUT lags, SPX lags*				
Lags	**1**	**2**	**3**	**4**	**5**
R Square	0.61970264	0.62799755	0.63366296	0.64184497	0.64203771
Adjusted R Square	0.61732578	0.62331828	0.6267072	0.63272	0.63056456
Standard Error	0.01483145	0.01471486	0.01464852	0.01453006	0.01457264
Observations	323	323	323	323	323
F	260.723401	134.208273	91.0989405	70.3394155	55.9600194
Significance F	6.6028E-68	5.2537E-67	6.3758E-66	1.7461E-65	1.1348E-63
Y:	*IV PUT*				
X:	*IV PUT lags, MON lags*				
Lags	**1**	**2**	**3**	**4**	**5**
R Square	0.46350387	0.47876815	0.48572458	0.49653986	0.506703580
Adjusted R Square	0.46015077	0.47221177	0.47595986	0.48371285	0.4908928
Standard Error	0.01761591	0.01741802	0.01735606	0.01722719	0.01710699
Observations	323	323	323	323	323
F	138.231416	73.0232951	49.7427905	38.7104921	32.0479757
Significance F	5.3784E-44	7.8544E-44	7.1118E-43	1.3514E-42	2.369E-42

Table 13 - Summary output for Regression Set n. 5.

Average moneyness (subsample)			
		Bid-Ask Prices	**Last Prices**
Start date	5-gen-12		
End date	26-set-13		
N. of weeks in sample		91	91
N. of missing observations		24	0
Average moneyness Put		15.01%	15.02%
Average moneyness Call		6.14%	6.19%

Table 14 - Average moneyness using bid-ask prices (subsample).

7. References

Coval, Joshua D. and Shumway, Tyler. 2001. "Expected Option Returns". *The Journal of Finance*, Vol. 56 (No.3): pp. 983-1009.

Tufano, Peter and Kyrillos, Barbara B. 1995. "Leland O'Brien Rubinstein Associates Incorporated: Portfolio Insurance". *Harvard Business Review*, Rev. September 15, 1995: HBS Case No. 294-061.

Welch, Ivo. 2016. "The (Time-Varying) Importance of Disaster Risk". *Financial Analysts Journal*, Vol. 72 (N. 5): pp. 14-30.

7.1 Online Resources

Chicago Board Options Exchange. 2000, April. *Margin Manual*, p. 14.
Available at: https://www.cboe.com/learncenter/pdf/margin2-00.pdf
(Link visited on 26 November 2018).