

FQ Perspective

by Max Darnell, Ghene Faulcon and Chuck Fannin, CFA

Foundations, Endowments and other institutional investors commonly have significant capital commitments in place with private equity and venture capital managers who are waiting for the right circumstances to arise before they call that capital. While there are a wide range of factors that may determine when private equity managers will call in committed capital, there is a tendency, particularly for buy-out funds, for those capital calls to come in at approximately the same time, i.e., they tend to cluster. We witnessed exactly that as the market reacted to the credit crisis in 2008 and early 2009. The common thread that leads buyout funds, in particular, to cluster is a decline in equity market prices which quite simply makes the price of potential deals more attractive. Low interest rates also increase the probability of capital calls (particularly for those private equity investments involving leveraged buy-outs where low interest rates mean lower borrowing costs for financing buy-outs), but because low interest rate environments typically arise gradually – they don't tend to arise through as discrete a change in the markets as low equity prices do - low interest rates don't lead to a short-term clustering of capital calls.

How should the committed capital be managed while it lies in wait for capital calls (or, for that matter, after it has been returned and is awaiting a new placement)? One answer is to hold the capital in short-term cash instruments. The liquidity is attractive since investors may have to fulfill capital calls within a matter of days even though the commitment may have existed for as long as three to five years. While this preserves the principle value of the capital set aside for future capital calls and generates a small return along the way, cash is, unfortunately, poorly matched with the risk and return profile intended for that capital. Such an approach to managing

the capital creates a significant performance drag on the total allocation to private equity.

Given the high correlations of private equity with public equity, public equities seem, on the surface, to be a better place to park committed, but as yet uncalled capital. The problem with parking this unemployed capital in public equities, however, is that if a decline in equities is what will tend to trigger a call for this capital, then one is, in effect, holding an asset to serve a future need in a form that will have its lowest value at the time the need arises. Equities will have to be sold right after they fall. "Sell low" is not a rule we should set ourselves up to follow. With this approach, if one has a \$1b capital commitment, then one would need to park some amount of capital greater than \$1b in public equities – say an additional 10% if a 10% decline in public equities is what will trigger a cluster of capital calls – to fully service that capital call when it occurs (this allocation can be reduced if and as equities deliver an attractive return during the period, of course). In short, the problem with holding public equities is, quite simply, that capital is most likely to be called when the value of public equities has fallen.

This presents a dual goal, therefore, in managing committed capital in advance of its drawdown. We'd like to generate a return that best approximates the return on private equity, while at the same time protecting the value of the capital in the context of an equity market decline. These goals are somewhat at odds with each other, so our search should be for the optimal compromise between these two goals. We are, in effect, in search of a hybrid asset that generally delivers private equity-like returns, but which will preserve most or all of its value in a declining equity market.



Hedged Equities

Overlaying equity exposure with a portfolio of options designed to hedge against downside risk would seem to be the natural starting point for addressing this need (most likely with the equities and the hedge being leveraged to match the leveraged nature of, and the return expectation for, private equity, but that's a separable issue). Some form of a self-financing collar would be attractive because such a hedge is financed out of opportunity cost rather than current assets or cash flow. By writing out of the money (OTM) calls to finance the purchasing of OTM puts, one gives up a portion of the best returns in tradeoff for avoiding the worst.

Because there's skew in the pricing of equity volatility (OTM puts cost more than calls with strikes that are equally distant in opposite direction from the current market level) simple collars give up more upside compared with the downside they protect against, i.e., the opportunity cost exceeds the value of what's protected. To give the reader a general feel, commonly skew will be such that the downside exposure will be 40-60% greater than the upside capture (e.g., a 7% OTM put might be financed with a call sold at only 4.5% OTM). Of course, the skew can become much worse than that during periods of heightened market stress. Given the negative skew in equity returns, the skew in options pricing may be entirely appropriate, but that doesn't mean that it is attractive to an investor looking for the most cost effective protection.

A better alternative is to use a put spread collar. In a put spread collar, one buys an OTM put, and finances it with a combination of writing a call and writing a further OTM put. This allows one to either/both reduce the opportunity cost associated with writing the call (by writing calls that are further OTM), and/or cause the protection that the purchased put provides by allowing one to buy a put that isn't as far OTM. These benefits come with the cost of the protection being less in the rarer, more extreme market declines, as the further OTM put that has been written re-exposes the portfolio to downside risk if the market falls beyond the strike of this put. Under such extreme outcomes, what has been protected against is only that portion of the price decline that lies between the strike on the nearer to the money put that's been bought, and the strike on the further OTM put that's been sold.

For example, we might buy a 7% OTM put, write a 14% OTM put, and then sell an OTM call to finance the remainder. The call we sell might be expected to fall in the neighborhood of 4.5% to 7% OTM depending on the skew. If the market was to fall during a quarter by as much as 30% and we had bought a put at 7% OTM and sold a put

at 14% OTM, then we'd have suffered a 23% reduction in portfolio value – a full 7% improvement in the outcome. This still represents a significant amount of protection, reducing a disastrous decline in this example by a quarter.

Just how much protection should one demand? That preference will vary according to each institution's risk tolerance, of course, so a simple starting point would be to assume we want to guard against a one standard deviation move in equities. The S&P 500¹ has an annualized long-term volatility of roughly 14%; the options we'd use for hedging are three month options, so we want to hedge against a one standard deviation move over a calendar quarter; that translates into a 7% move, so we'd want to buy quarterly puts that are 7% OTM¹.

Furthermore, we should be less interested in paying for protection after the equity market has fallen, and/or when the cost of insurance is on the high side. Given that the cost of insurance quite predictably rises with market declines, insurance will be most expensive after a market decline, right when our need for further insurance should be lower (since the damage is likely to have already been done). It is also precisely at these times when, if you assume mean-reversion in market returns, you want to maintain a higher upside exposure to the market. The moneyness on the put we buy should, therefore, vary with the cost of insurance, i.e., with options price implied volatility. When implied volatility is high, we should look to buy puts that are farther out of the money, e.g., greater than 7% OTM, and when implied volatility is low, we should look to buy puts that are closer to ITM, e.g., less than 7% OTM².

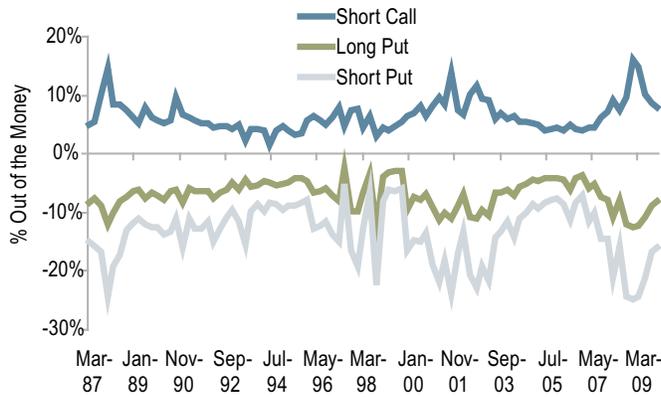
For simplicity, we'll use the following rule: the strike on the put we buy, instead of always being roughly 7% OTM, will instead be equal to 70% of implied quarterly volatility. Why 70%? We've chosen 70% because that gets us to an average expectation of buying 7% OTM puts. Implied volatility averages 20% historically, so again converting that to a quarterly volatility number – since we're buying quarterly options contracts – gets us to 7%. If implied volatility is at 15%, we'll look to buy puts that are only 5.25% OTM, whereas if implied volatility is at 25%, we'll look to buy puts that are 8.75% OTM.

The options market was thin in the early 1980's, and skew in options pricing didn't exist until after the market decline in October of 1987. We should start our analysis in 1988, therefore, but because 1987 is such an important event to include, we've reached back one year earlier and have started the analysis in 1987. Using this approach applied to historical options contracts since 1987, we would have seen the following strikes chosen for the options:



FIGURE 1: STRIKES FOR OPTIONS STRUCTURES

March 1987 – December 2009

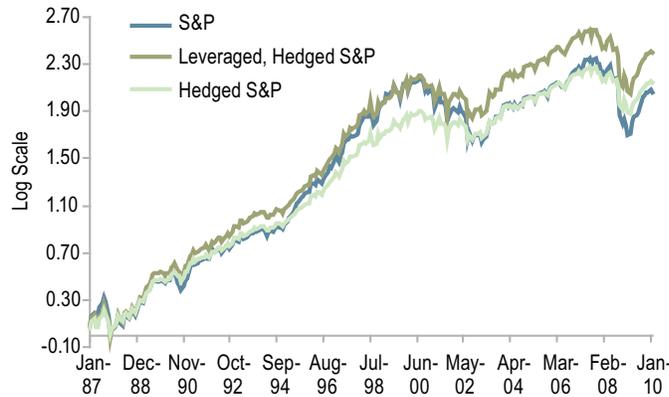


Source: First Quadrant, L.P.

Since this approach will, of course, reduce the volatility of the return, we should leverage the returns (scaling through the use of futures) to the same volatility level of the S&P. Both the unleveraged and leveraged versions of the hedged portfolio would have outperformed the S&P 500 over this period³. Notably the unlevered portfolio would have kept up with the S&P until roughly 1995, then lagged behind during the strong bull market of the late 1990's. The leveraged portfolio keeps up with the S&P during the whole of the 1990's, and then moves ahead of the S&P during

FIGURE 2: CUMULATIVE GROWTH \$1

January 1987 – February 2010



Source: First Quadrant, L.P.

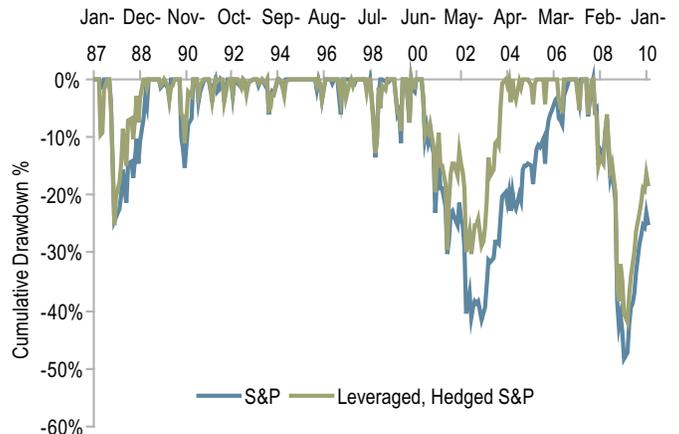
the first market decline of the 2000's, and moves farther ahead during the second decline. That's exactly what we're looking for it to do.

Given that what we're most interested in reducing is our exposure to drawdowns, the historical comparison of cumulative drawdowns demonstrates that we've achieved this end. Nowhere does the leveraged, and hedged S&P portfo-

lio realize any meaningfully worse drawdown, and in all three major downdraft episodes for the S&P, this portfolio provides the protection we're looking for.

FIGURE 3: DRAWDOWN ANALYSIS (STANDARD DEVIATIONS)

January 1987 – February 2010



Source: First Quadrant, L.P.

Notably, this approach significantly outperforms the use of either put-spread collars that are unresponsive to the price of implied volatility (by always looking for a 7% OTM strike on the put that's purchased), or simple collars (that do use wider strikes when implied volatility is higher, but that finance all of the cost of the put that's purchased with a call, thereby cutting off the upside more severely). In the table below, the "leveraged, hedged S&P" version we're focused on outperforms the others in terms of sharpe ratio, and it reduces both the negative skew and the excess kurtosis (kurtosis is reported relative to a zero normal rather than a normal of three) relative to all but the simple collar. The simple collar results in even lower skew and lower kurtosis, but because it has such a dramatic, negative impact on risk, it suffers from higher leverage costs to reach the same risk level and ends up delivering less return as a result.

Generalized Measure of Downdraft Exposure

Investment managers often cite their performance in terms of the annualized return relative to the maximum cumulative downdraft which is intended to represent risk. The problem with this is that "risk" is defined by a single event. What investors want is to avoid shortfalls generally, but particularly shortfalls that are deep and shortfalls that are prolonged. Said otherwise, our private equity investors want to have the greatest probability weighted portfolio value available when the capital calls arrive. This is an issue of both how much below the peak value the portfolio falls, and how long we remain there. The portfolio value has everything to do with how deep the downdraft is. Since



Monthly Returns Sorted Into S&P Return Quintiles					
January 1987 – February 2010					
	Btm Quintile	4th Quintile	3rd Quintile	2nd Quintile	Top Quintile
Quintile Means					
S&P	-5.67%	-0.60%	1.17%	2.95%	6.26%
Leveraged, Fixed Put-Spread (7%)	-5.44%	-0.70%	1.28%	2.97%	6.42%
Leveraged, Costless Collar	-4.97%	-0.80%	1.48%	2.23%	5.89%
Leveraged, Hedged S&P	-5.17%	-0.69%	1.63%	2.90%	6.05%

Quintile Medians					
S&P	-4.79%	-0.53%	1.04%	2.92%	5.64%
Leveraged, Fixed Put-Spread (7%)	-4.14%	-0.52%	1.26%	2.96%	6.13%
Leveraged, Costless Collar	-4.56%	-0.53%	1.30%	2.47%	5.32%
Leveraged, Hedged S&P	-4.03%	-0.53%	1.50%	2.85%	5.79%

Performance Summary Statistics				
	S&P	Leveraged, Fixed Put-Spread (7%)	Leveraged, Costless Collar	Leveraged, Hedged S&P
Ann Rtn %	9.33%	10.42%	8.55%	10.89%
Ann Risk %	15.24%	15.24%	15.21%	15.23%
Sharpe Ratio	0.61	0.68	0.56	0.72
Skew	-0.97	-0.68	-0.48	-0.79
Kurtosis	3.27	2.39	1.17	2.39

Source: First Quadrant, L.P.

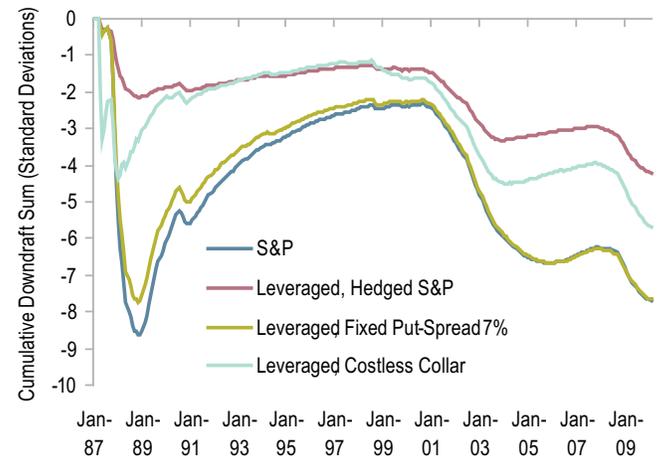
it is unknown as to when the call on capital will come, we should treat every point in time where the market has fallen as equally likely to be subject to the capital calls. The disutility of underperformance is, therefore, a function of both magnitude and time, and when evaluating the historical drawdraft experience, investors should consider not just the worst case, but the cumulation of all cases.

In effect, the sum total of “unpleasantness” is represented by the area under zero in each of the two lines plotted in the cumulative drawdraft in Figure 4. (Mathematically, we’re measuring and comparing the integrals of each the two lines in the following chart that displays cumulative drawdraft episodes.) A month of being underwater by 5% would be considered a 5% “unit” of drawdraft exposure. If that drawdraft was maintained into a second month, that would be measured at a total of 10% units of drawdraft exposure. This treats time as equally important to depth, which may or may not be correct, but time certainly matters, so time-weighting is certainly important. A 5% drawdraft from which investors fully recover in a month is certainly less painful than a 5% drawdraft from which it takes six months to recover.

The chart in Figure 4 shows the cumulative sum of the time-weighted drawdrafts of the S&P relative to our three versions of the options overlay or hedge. So that this measure is generalizable, we cumulate shortfall returns in units

of standard deviations. (Since each of the series we’re examining has the same historical standard deviation, this has no impact here.) If you look to the “leveraged, hedged S&P” line in the Figure 4 chart, you would observe, for example, that as of January, 2000, that portfolio had historically delivered an average of 1.4 annualized standard deviations of shortfall per year. Given the larger drawdrafts experienced since that time, that average annual shortfall

FIGURE 4: DOWNDRAFT INDEX
January 1987 – February 2010



Source: First Quadrant, L.P.



grew to 4.2 by February of 2010. The fact that these lines trend downward is due to the increased degree of shortfall experienced late in the period. As we move forward into the future, we would expect the index to rise, assuming we experience less shortfall in the future than we have over the past (we've seen what we all hope is an extreme).

What we're looking for in this chart is a steady rather than episodic growth in cumulative shortfall, and a shallower shortfall ideally. Steady growth means that we minimize the risk that our capital is called right after a large, more episodic shortfall. That's why we like a steady decline. The clear implication is that the leveraged, hedged portfolio shows itself to have delivered less downside pain than the S&P on its own or the other two approaches to hedging.

Overhedging

Hedging the equities is a sensible starting point, but if we think more abstractly about our goals and about the nature of the options portfolio, we can accomplish much more. Our goal is not really to specifically participate in equity returns, but rather to hold an asset that will deliver equity-like returns with less risk of finding ourselves in a downdraft situation at the time when committed capital is likely to be called. Because committed capital is likely to be called at a time when equities, specifically, have fallen, we most definitely do want to perform well - in fact we'd like to perform *especially* well, at the very time that equities decline. The rest of the time, we shouldn't be concerned about how closely the asset relates to equity performance so long as the long-term expected return is similar to equities (again, we might want to be thinking in terms of leveraged equity exposure, but that doesn't change the analysis here). In fact, if the return potential is the same as equities with equal or less risk, but the correlation is otherwise low, then we have the benefit of carrying a diversifying asset in our portfolio rather than just more of the same equity, meaning we're better off if the asset doesn't correlate to equities generally. In short, we should be willing to bear significant tracking error relative to the equity market.

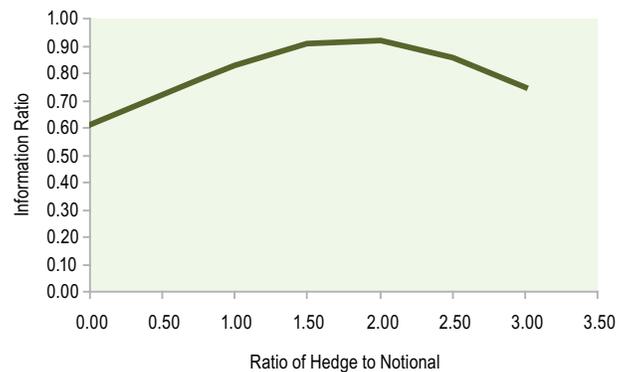
Because the options portfolio stands ready to deliver its best performance when equities fall, the more options exposure we have at such times, the better. We'd prefer to be "over hedged," in other words, at such times. Since we can't predict when market drops will occur, we have to optimize the combination of equity exposure and the amount we over hedge based upon the long-term risk-return characteristics. We find that the optimal combination falls roughly in the neighborhood of running the hedge at double the size of the underlying equity investment. This is roughly consistent with an information-ratio based risk

allocation which is the right theoretical approach to optimizing the combination.

What makes this optimal? First, we can look to the impact this scale of hedging has on the return and information ratio. A hedge/notional of 2, for example, indicates that the hedge over twice the portfolio notional is optimal over the past 25 years.

FIGURE 5: IMPACT OF HEDGE AT DIFFERENT RATIOS OF HEDGE TO NOTIONAL January 1987 – February 2010

Impact of Hedge at Different Ratios of Hedge to Notional							
Hedge/Notional	0	0.75	1	1.5	2	2.5	3
Ann. Return %	9.33%	9.64%	9.69%	9.73%	9.67%	9.51%	9.26%
Ann. Risk %	15.24%	12.42%	11.70%	10.74%	10.52%	11.11%	12.39%
Sharpe Ratio	0.61	0.78	0.83	0.91	0.92	0.86	0.75
Skew	-0.97	-0.89	-0.79	-0.51	-0.22	-0.11	-0.16
Kurtosis	3.27	2.67	2.38	2.41	4.35	7.54	9.70



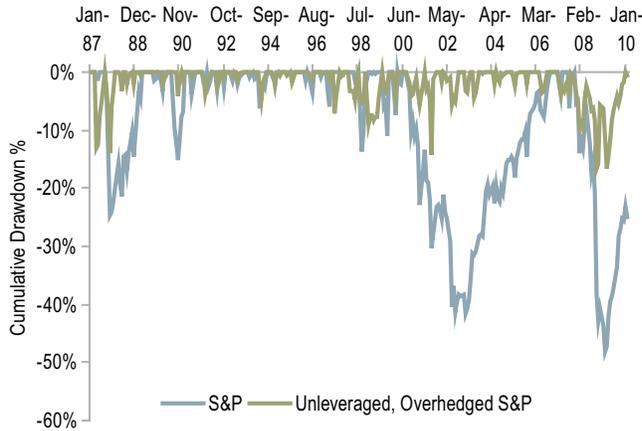
Source: First Quadrant, L.P.

Similar to the problem we incurred with the simple collar, by overhedging, the risk is reduced more dramatically, so more leverage cost is involved in reaching the same risk as equities. In this case, however, leverage is not required to obtain the same level of return as equities. We can enjoy both higher return and lower risk in this case. Where the greatest advantage lies in how the downdraft risk is reduced relative to the case above where the notional scale of the options overlay was the same as the notional scale of the equity exposure. We can see this in both of the charts in Figure 6 and 7 which reflect dramatic improvement on our goals.



FIGURE 6: DRAWDOWN ANALYSIS

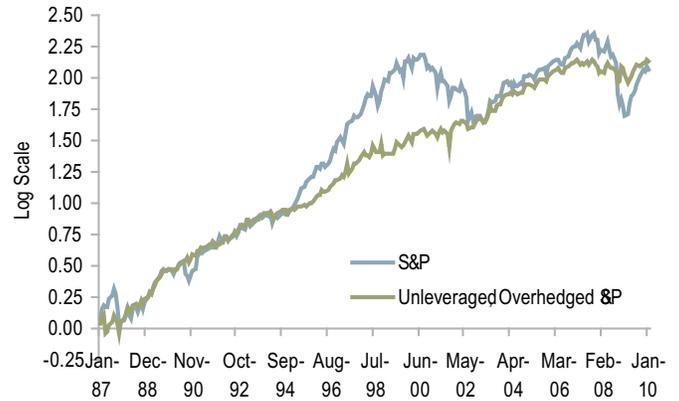
January 1987 – February 2010



Source: First Quadrant, L.P.

FIGURE 8: CUMULATIVE GROWTH \$1

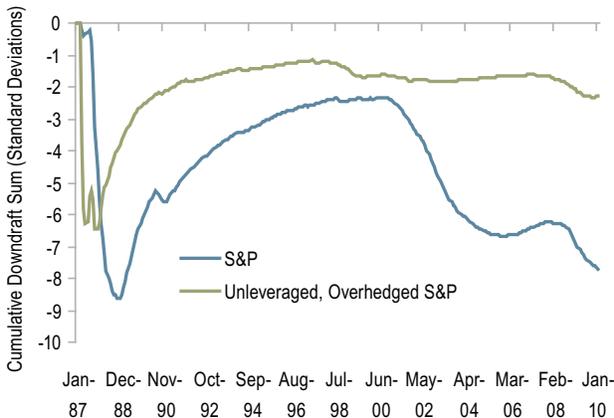
January 1987 – February 2010



Source: First Quadrant, L.P.

FIGURE 7: DRAWDOWN INDEX

January 1987 – February 2010



Source: First Quadrant, L.P.

Conclusion

When waiting for capital commitments to be called, particularly buyout related capital commitments, investors need a source of return that will best approximate the returns they would obtain if the capital commitments were already filled, while at the same time avoiding the obvious problem that what triggers some of those capital calls is a decline in that very instrument, equities, that best approximates the return to private equity. A solution is to invest in equities and to utilize a hedge. Which hedging method you use makes a significant difference. Simple collars give up too much upside, and being unresponsive to risk in pricing put-spread collars causes one to pay too much to protect the portfolio when implied volatilities are high, and to give up too much upside right after market declines have occurred. The use of a put-spread collar that is sensitive to market volatility provides a superior approach to protecting the portfolio against periods of drawdowns – periods when private equity investors are most likely to receive calls on their committed, but as of yet, uncalled capital.

This approach can be further enhanced if one is willing to take on increased tracking error relative to equities, and is willing to increase the investment in the “insurance” that the options overlay provides by increasing the notional associated with the hedge above and beyond the notional of the equity exposure. Such an approach puts more emphasis on protecting capital, and historically would have better returns than equities at lower risk. This comes only at the cost of increased tracking error relative to equities – a cost that we think should not be of issue when used as a means of investing for future capital calls.





Endnotes

¹ For the duration of this paper the S&P total return is calculated as the S&P futures return plus the return (1 month LIBOR) on the cash needed to fully equitize the investment. The pricing of futures considers expected dividends and current interest rates. This has very good tracking characteristics with respect to the S&P total return index, but allows us to extend our analysis further back in time in order to cover the very important events of October 1987. Since the S&P return forms part of all of the comparisons being made, this formulation of the S&P total return is not giving any advantage or disadvantage to any of the candidates being compared. An additional result of this construction is that leveraging this portfolio consists solely of scaling the futures investment, but requires no additional finance costs as futures have the financing cost built into them.

² Exploring such an approach historically requires dealing with the practical realities of the options available to us. You'll rarely find options that are struck at exactly 7% out of the money, of course, and in the earliest days for S&P 500 options, strikes were often in the neighborhood of 5% or more apart. The further back in time we go, we also find an increasing number of days upon which existing options didn't trade, and for which we don't have a market price for the options. We calculate all returns from the 17th of the prior month to the 18th of the subsequent month since options tend to expire roughly around the 18th. In cases where the options prices were stale on the 18th, we rolled the return forward into the next month rather than making an assumption about likely prices. At each quarter-end, this problem resolved itself at expiration.

³ The hypothetical performance results of the paper portfolios presented in this study are for illustrative purposes only and are not representative of any actual strategies deployed by First Quadrant, L.P. No actual trading has taken place. Strategy guidelines are constant through the life of the portfolios. Transaction costs were set at 7.5 basis points per transaction (30 basis points per year). Option structures were held until maturity at which point it is rolled into the next 3 month option position with new strikes as described in the main text. All trading takes place once a month (just prior to expiration of the options – generally around the 18th of each month). Performance returns for one year or longer are annualized. With respect to the put spread collar, when options could not be found close to the strikes that we were looking for, we chose the farthest out of the money put option for our short put option and the put option closest to half as much out of the money for our long put option. Then our long call option was chosen to make the entire set nearly costless. This option structure is held until maturity at which point it is rolled into the next 3 month option position with new strikes as described.