The Cost of Capital for Alternative Investments

Jakub W. Jurek and Erik Stafford*

Abstract

This paper studies the cost of capital for alternative investments. We document that the risk profile of the aggregate hedge fund universe can be accurately matched by a simple index put option writing strategy that offers monthly liquidity and complete transparency over its state-contingent payoffs. The contractual nature of the put options in the benchmark portfolio allows us to evaluate appropriate required rates of return as a function of investor risk preferences and the underlying distribution of market returns. This simple framework produces a number of distinct predictions about the cost of capital for alternatives relative to traditional mean-variance analysis.

First draft: June 2011

This draft: August 2011

*Jurek: Bendheim Center for Finance, Princeton University; *jjurek@princeton.edu*. Stafford: Harvard Business School; *estafford@hbs.edu*.

This paper studies the required rate of return for a risk averse investor allocating capital to alternative investments. There are two key aspects to this asset allocation decision that are inconsistent with (present a challenge for) standard decision-making tools. First, the alternative investment exposure is nonlinear with respect to the remainder of the portfolio at the horizon where the investor is able to rebalance the portfolio, making static mean-variance analysis inappropriate. Second, the typical allocation is large relative to the aggregate outstanding share in the economy, such that charging for a small marginal allocation to this risk is a poor approximation to the actual contribution to the portfolio's overall risk. These two features interact to produce very large required rates of return relative to commonly used models.

Investments made by sophisticated individual and institutional investors in private investment companies like hedge funds and private equity funds are referred to as alternative investments. These investments are frequently combined with financial leverage to bear risks that may be unappealing to the typical investor or that require flexibility that public investment funds may not provide. The economic nature of these investments is commonly described as "picking up pennies in front of a steamroller" indicating that it is understood that there is a real possibility of a complete loss of invested capital. Moreover, the aggregate performance of these unappealing positive net supply risks tend to be correlated with aggregate economic conditions, such that losses are more likely when other positive net supply assets are also experiencing losses.

As emphasized by Shleifer and Vishny (1997), investors in hedge funds executing risk arbitrage strategies typically do not have the technology and information required to engage in these activities directly and must infer the risks from realized returns. One point of this paper is that even with direct knowledge of the underlying risks the commonly used tools for asset allocation and determining required rates of return are inappropriate for these types of risk.

A practical challenge in measuring the risks from realized returns is that the portfolios are not directly observable at any point in time and the risks of any fund are likely to be changing through time because of the manager's discretion over portfolio composition, financial leverage, and hedging rules. To estimate the aggregate risks of the hedge fund universe, we exploit the 2008 stock market decline to identify an index put writing strategy that matches the aggregate hedge fund drawdown. This put writing strategy matches the risks of the hedge fund universe throughout the sample period (1996-2010) in terms of other drawdown patterns, stock market beta, and time series return volatility. The implied fee on the put writing strategy required to match the cumulative returns on the hedge fund index is 3.6% per year, which corresponds closely with the 2% fixed fee plus 20% profit sharing compensation structure that is common in alternative investments.

With a complete state-contingent description of an investable risk-matched alternative to the aggregate hedge fund universe, we can determine the rate of return that an investor would require as a function of his risk aversion and the underlying return distributions of other asset classes, all of which are necessary for any asset allocation decision. The comparative statics of the generalized asset allocation framework suggest that the nonlinearity of alternative investments creates several meaningful differences relative to the mean-variance framework. These differences become larger as the allocation to alternatives increases. In traditional mean-variance analysis, as an investor allocates capital between riskfree and risky investments the required rate of return on the risky investment increases linearly in the risky share. We show that when the risky portfolio consists of downside risks the required rate of return is convex in the risky share. In addition to various effects related to portfolio composition, we examine situations that systematically affect required rates of return, but where traditional mean-variance or CAPM-style analyses predict no variation in returns. For example, increasing skewness in the underlying return distribution while holding variance constant has a large effect on the cost of capital for investments with downside exposure. Similarly, increasing leverage while holding CAPM β constant has a large effect on required rates of return in the generalized framework, but not in the traditional framework. An interesting implication is that investors relying on traditional analyses for benchmarking alternatives are likely to be attracted *ex ante* to strategies and historical return series that are highly levered investments in safe assets that will be disappointing in the event of a market crash.

The remainder of the paper is organized as follows. Section 1 describes the risk profile of hedge funds. Section 2 presents a simple recipe for replicating the aggregate hedge fund risk exposure. Section 3 develops a generalized asset allocation rule appropriate for combining securities with nonlinear payoff functions. Section 4 discusses implications of the framework relative to traditional mean-variance analysis. Finally, Section 5 concludes the paper.

1 Describing the Risk Profile of Hedge Funds

To compute the required rate of return – or cost of capital – for an allocation to hedge funds, one must first characterize the risk profile of a typical investment. Rather than examine risk exposures of individual funds (Lo and Hasanhodzic (2007)) or strategies (Fung and Hsieh (2001), Mitchell and Pulvino (2001), Agarwal and Naik (2004)), we focus on the the aggregate risk properties of the asset class. Consequently, the cost of capital we derive can be thought of as applying to an investor in a diversified hedge fund portfolio (e.g. a fund-of-funds, or an endowment holding a portfolio of alternative investments).

We proxy the performance of the hedge fund universe using two indices: the Dow Jones Credit Suisse

Broad Hedge Fund Index, and the HFRI Fund Weighted Composite Index. Such indices typically provide an upward biased assessment of hedge fund performance due to the presence of backfill and survivorship bias. For example, Malkiel (2005) reports that the difference between the mean annual fund return in the backfilled and non-backfilled TASS database was 7.34% per year in the 1994-2003 sample. Moreover, once defunct funds are added in the computation of the mean annual returns to correct for survivorship bias, the mean annual fund return declines by 4.42% (1996-2003). To the extent that the survivorship bias also affects the measured risks, it is unlikely that the true risks are lower than those estimated from the realized returns over this period. We discuss the implications of higher underlying risks and how alternative economic outcomes are likely to affect the risks of alternative investments in Section 4.

Table 1 reports summary statistics for the two indices computed using quarterly returns from 1996:Q1-2010:Q3 ($N = 59$ quarters), and compares them to the S&P 500 index.[1] The attraction of hedge funds over this time period is clear: mean returns on alternatives exceeded that of the S&P 500 index, while incurring lower volatility. Moreover, the estimated linear systematic risk exposures (or CAPM β values) indicate that hedge fund performance was largely unrelated to the performance of the public equity index, and suggests that relative to this risk model they have outperformed. The realized Sharpe ratios on alternatives were three times higher than that of the S&P 500 index. Under all of the standard risk metrics inspired by the mean-variance portfolio selection criterion, hedge funds represented a highly attractive investment. This is further illustrated in Figure 1, which plots the value of $1 invested in the various assets through time. By September 2010, the hedge fund investor had amassed a wealth roughly 50% larger than the wealth of the investor in public equity markets, and roughly twice the wealth of an investor rolling over investments in short-term T-bills.

Another risk metric popular among practitioners is the drawdown, which measures the magnitude of the strategy loss relative to its highest historical value (or high watermark). Hedge funds perform relatively well on this measure over the sample period with a drawdown of approximately -20%, which is less than half of the -50% drawdown sustained by investors in public equity markets.

Figure 1 also demonstrates that the performance of hedge funds as an asset class is *not* market-neutral. For example, hedge funds experience severe declines during extreme market events, such as the credit crisis during the fall of 2008 and the LTCM crisis in August 1998. During the two-year decline following the bursting of the Internet bubble, hedge fund performance is flat. And, finally, in the "boom" years hedge funds perform well. Empirically, the downside risk exposure of hedge funds as an asset class is reminiscent of writing out-of-the-money put options on the aggregate index. Severe index declines cause the option to

[1] Although the indices are available at daily and monthly frequencies, we focus on quarterly returns to ameliorate the effects of stale prices and return smoothing (Asness, et al. (2001), Getmansky, et al. (2004)).

expire in-the-money, generating losses that exceed the put premium. Mild market declines are associated with losses comparable to the the put premium, and therefore flat performance. Finally, in rising markets the put option expires out-of-the-money, delivering a profit to the option-writer.

There are structural reasons to view the aggregated hedge fund exposure as being similar to short index put option exposure. Many strategies explicitly bear risks that tend to realize when economic conditions are poor and when the stock market is performing poorly. For example, Mitchell and Pulvino (2001) document that the aggregate merger arbitrage strategy is like writing short-dated out-of-the money index put options because the underlying probability of deal failure increases as the stock market drops. Hedge fund strategies that are net long credit risk are effectively short put options on firm assets – in the spirit of Merton's (1974) structural credit risk model – such that their aggregate exposure is similar to writing index puts. Other strategies (e.g. distressed investing, leveraged buyouts) are essentially betting on business turnarounds at firms that have serious operating or financial problems. In the aggregate these assets are likely to perform well when purchased cheaply so long as market conditions do not get too bad. However, in a rapidly deteriorating economy these are likely to be the first firms to fail.

The downside exposure of hedge funds is induced not only by the nature of the economic risks they are bearing, but also by the features of the institutional environment in which they operate. In particular, almost all of the above strategies make use of outside investor capital and financial leverage. Following negative price shocks outside investors make additional capital more expensive, reducing the arbitrageur's financial slack, and increasing the fund's exposure to further adverse shocks (Shleifer and Vishny (1997)). Brunnermeier and Pedersen (2008) provide a complementary perspective highlighting the fact that, in extreme circumstances, the withdrawal of funding liquidity (i.e. leverage) to arbitrageurs can interact with declines in market liquidity to produce severe asset price declines.

2 Replicating Aggregate Hedge Fund Risk Exposure: A simple recipe

In order to replicate the aggregate risk exposure of hedge funds, we examine the returns to simple strategies that write naked (unhedged) put options on the S&P 500 index.[2] Each strategy writes a single, short-dated put option, and is rebalanced monthly. We consider a range of strategies with different downside risk exposures, as measured by how far the put option is out-of-the-money and how much leverage is applied

[2] Academic approaches to hedge fund return replication fall into three broad categories: factor-based, rule-based, and distributional. The factor-based (APT) methods use regression analysis to identify replicating portfolios of tradable indices (Fung and Hsieh (2002, 2004), Lo and Hasanhodzic (2007)); in some cases, the factors include option-based strategies. The rule-based methods use mechanical algorithms to assemble portfolios mimicking basic hedge fund strategies (Mitchell and Pulvino (2001), Duarte, et al. (2007)). The distributional methods focus on matching the distributional properties of returns through dynamic trading of futures (Kat and Palaro (2005)).

to the portfolio. We place primary focus on matching the drawdowns experienced by the hedge fund indices during the fall of 2008, and the LTCM crisis. Our secondary focus is on matching the realized return volatility and CAPM betas. The emphasis on matching drawdowns is motivated by the potential non-linearity in the underlying economic risk exposures, and the diminished scope for managers to smooth returns during extreme market moves. Once we have selected the strategy that most closely matches the hedge fund drawdowns, we then examine its risk and return properties over the full sample period (1996:1-2010:10).

2.1 Strike Selection through Time

Unlike previous studies, which have focused on strategies with fixed option moneyness – measured as the strike-to-spot ratio, K/S, or strike-to-forward ratio – we construct strategies that write options at fixed strike Z-scores. The option strike corresponding to a Z-score, Z, is given by:

$$K(Z) \quad = \quad S \cdot \exp\left\{\left(r_f(\tau) - q(\tau) + \frac{\sigma^2}{2}\right) \cdot \tau + \sigma(\tau) \cdot \sqrt{\tau} \cdot Z\right\} \tag{1}$$

where $r_f(\tau)$, $q(\tau)$, and $\sigma(\tau)$ are the risk-free rate, dividend yield, the stock index volatility corresponding to the *trade* maturity, τ. Since the option maturity date will generally not coincide with the trade maturity (i.e. roll date), we select the contract with nearest expiration date, T, following the trade roll date. Consequently, we distinguish between the trade initiation date, t_o, the trade closure date, t_c, and the option expiration date, T. In our empirical implementation, we set the trade maturity, $\tau = t_c - t_o$, equal to one month, rolling the positions on the last business day of each month. To measure volatility at the one-month horizon we use the CBOE VIX implied volatility index; dividend yields and risk-free rates are from OptionMetrics.

Selecting strikes on the basis of their corresponding Z-scores ensures that the systematic risk exposure of the options at the roll dates is fixed, when measured using the option delta. For example, recall that the delta of the put with strike $K(Z)$ is:

$$\delta_{put}\left(Z\right) \quad = \quad -e^{-q(T-t_o)\cdot(T-t_o)} \cdot \Phi\left[Z\right] \tag{2}$$

Since the dividend yield over the life of the option, $q\left(T - t_o\right)$, does not vary substantially in the data, the option deltas will be comparable at the roll dates. This contrasts with applications which involve fixing the strike moneyness (Glosten and Jagannathan (1994), Bakshi and Kapadia (2003), Agarwal and Naik (2004)). In particular, options selected by fixing moneyness have higher systematic risk – as measured by delta or market beta – when implied volatility is high, and lower risk when implied volatility is low.

2.2 The Mix of Investor Capital and Leverage

Option writing strategies require the posting of capital (or, margin). The capital represents the investor's equity in the position, and bears the risk of losses due to changes in the mark-to-market value of the liability. In the case of put writing strategies the maximum loss per option contract is given by the option's strike value. Consequently, a put writing strategy is fully-funded or unlevered – in the sense of being able to guarantee the terminal payoff – if and only if, the investor posts the discounted value of the exercise price less the proceeds of the option sale, κ_A:

$$\kappa_A \quad = \quad e^{-r_f(T-t_o)\cdot(T-t_o)} \cdot K - \mathcal{P}(K,\, S,\, T;\, t_o) \tag{3}$$

In practice, it is uncommon for the investor to post the entire asset capital, κ_A. Instead, the investor contributes equity of, κ_E, with the broker contributing the balance, κ_D. Although the broker's contribution is conceptually equivalent to debt, the transfer of the principal never takes place, and the interest rate on the loan is paid in the form of haircut on the risk-free interest rate paid on the investor's capital contribution. The ratio of the asset capital to the investor's capital contribution (equity), represents the leverage of the position, $L = \frac{\kappa_A}{\kappa_E}$. Allowable leverage magnitudes are controlled by broker and exchange limits, with values up to approximately 10 being consistent with existing CBOE regulations.[3]

The investor's capital – comprised of his contribution κ_E and the put premium proceeds – is assumed be invested in securities earning the risk-free rate less the broker haircut, h. If the broker haircut exceeds the prevailing risk-free interest rate the deposit earns an interest rate of zero. This produces an terminal *accrued interest* payment of:

$$\mathrm{AI}\,(t_o,\, t_c) \quad = \quad \left(\frac{\kappa_A}{L} + \mathcal{P}(K,\, S,\, T; t_o)\right) \cdot \left(e^{\max\left(0,\, r_f(\tau)-h\right)\cdot\tau} - 1\right) \tag{4}$$

The investor's return on capital is comprised of the change in the value of the put option and the accrued interest divided by his capital contribution (or equity):

$$r\,(t_o,\, t_c) \quad = \quad \frac{\mathcal{P}(K,\, S,\, T;\, t_c) - \mathcal{P}(K,\, S,\, T;\, t_o) + \mathrm{AI}\,(t_o,\, t_c)}{\kappa_E} \tag{5}$$

[3]The CBOE requires that writers of uncovered (i.e. unhedged) puts "deposit/maintain 100% of the option proceeds plus 15% of the aggregate contract value (current index level) minus the amount by which the option is out-of-the-money, if any, subject to a minimum of [...] option proceeds plus 10% of the aggregate exercise amount:"

$$\min\, \kappa_E^{CBOE} \quad = \quad \max\left(0.10 \cdot K,\, \mathcal{P}^{bid}(K,\, S,\, T;\, t_o) + 0.15 \cdot S - \max(0,\, S-K)\right)$$

Unless otherwise noted, we assume that the investor buys (sells) puts at the ask (bid) prevailing at the close of the trade date.

2.3 The Returns to Naked Put Writing

On the last trading day of each month from January 1996 through September 2010, we invest the investor's capital in a one-month U.S. Treasury bill and write an option on the S&P 500 index corresponding to the strike at Z-score, Z, receiving the bid price. The quantity of posted capital, κ_E, relative to the total exposure, κ_A, determines the leverage, L, of the strategy. The portfolio is rebalanced monthly by buying back last month's option at the prevailing ask price, and writing a new index put option at the strike, $K(Z)$, receiving the bid. We consider strategies $[Z, L]$ with $Z \in \{-1, -2, -3\}$ and $L \in \{1, 2, 4, 6, 8, 10\}$.

2.3.1 An Example

To illustrate the portfolio construction mechanics consider the first trade of the $[Z = -1, L = 2]$ strategy. The initial positions are established at the closing prices on January 31, 1996, and are held until the last business day of the following month (February 29, 1996), when the portfolio is rebalanced. At the inception of the trade the closing level of the S&P 500 index was 636.02, and the implied volatility index (VIX) was at 12.53%. Together with the risk-free rate and dividend yield corresponding to the one-month holding period, these values pin down a proposal strike price, $K(Z)$, for the option to be written via (1). To obtain the risk-free rate and dividend yield we use the Option Metrics index dividend yield and zero-coupon yield curve files to find $r_f(\tau) = 5.50\%$ and $q(\tau) = 2.82\%$, where τ equals one month. We then select an option maturing *after* the next rebalance date, whose strike is closest to the proposal value, $K(Z)$. In this case, the selected option is the index put with a strike of 615 maturing on March 16, 1996. The $[Z = -1, L = 2]$ strategy writes the put, bringing in a premium of \$3.00, corresponding to the option's *bid* price at the market close. The required asset capital, κ_A, for that option is \$607.99, and since the investor deploys a leverage, $L = 2$, he posts capital of exactly half that, or equivalently, \$303.99. The investor's capital is invested at the risk-free rate, with the positions held until February 29, 1996. At that time, the option position is closed by repurchasing the index put at the close-of-business *ask* price of \$2.25. This generates a profit of \$0.75 on the option and \$1.32 of accrued interest, representing a 68 basis point return on investor capital. Finally, a new strike proposal value, which reflects the prevailing market parameters is computed, and the entire procedure repeats.

2.3.2 Risk Profiles

We report the risk exposures – volatility, (linear) CAPM beta, and drawdown – of the various naked put writing strategies in Table 2. To make comparisons with the *after-fee* statistics in Table 1 meaningful, we equalize the compounded in-sample performance of each put writing strategy and the HFRI Fund Weighted Composite Index by taking out a flat management fee, f, prior to computing the risk statistics. The implied, annualized fee is reported along with the risk statistics in Table 2.

Table 2 illustrates a wide range of risk exposures available in naked put writing strategies. Consistent with intuition, holding the strike Z value (or option delta) fixed, the three risk measures increase with leverage, L; holding leverage fixed, strategies become "safer" as the strike Z value becomes more negative, since the options being written are further out-of-the-money. Despite the explicit leverage of the strategy and the implicit leverage of the options, most of the strategies experience less volatility that the S&P 500 index. In fact, in the absence of marking-to-market, all of the strategies – even those with leverage values of 10 (the CBOE limit) – would have survived despite: (a) equity market performance being below its long-run historical average in this sample; and, (b) the extreme events during the fall of 2008.

The put writing strategy, which provides the closest match to the aggregate risk properties of the hedge fund indices, writes index puts at strikes corresponding to $Z = -1$ and deploys leverage of $L = 2$. The strategy sustains a maximum drawdown of -22.6%, which is in line with the -20% drawdown realized by hedge funds, and its annualized volatility and (linear) CAPM β match the corresponding values of the hedge fund indices (Table 1). Finally, the implied annual management fee of 3.60%, necessary to match the performance of the HFRI Fund Weighted Composite Index is consistent with estimates of the all-in hedge fund fees over this sample. For example, using cross-sectional data from the TASS database for the period 1995-2009, Ibbotson, et al. (2010) find that the average fund collected an annual fee of 3.43%.

To further explore the quality of the empirical fit provided by the [$Z = -1$, $L = 2$] put writing strategy, we plot its *after-fee* performance alongside the HFRI Fund Weighted Composite Index in Figure 2. The overall fit is remarkable: the strategy matches the losses during the fall of 2008 and the LTCM crisis, the flat performance during the bursting of the Internet bubble, as well as, the strong returns during boom periods. The panels on the bottom, which plot the times series of drawdowns, illustrate the superior performance of hedge funds relative to public equities, and highlight their similarity to naked put writing. The close fit of our replicating strategy indicates that in spite of variation in the popularity of individual hedge fund strategies and institutional changes in the industry, the underlying economic risk exposure of hedge funds has remained essentially unchanged over the 15-year sample. This is consistent with the notion that hedge funds specialize in the bearing of a particular class of scary, positive *net* supply risks, that may be highly unappealing to less sophisticated investors. Consequently, in the ensuing analysis we use the [$Z = -1$, $L = 2$] strategy to characterize the economic risk exposure of the aggregate hedge fund universe.

3 Portfolio Selection with Alternatives

The two fundamental ingredients of any portfolio selection framework are: (1) a specification of investor preferences (utility); and, (2) a description of the joint payoff profiles (or return distributions) of the assets under consideration. Under a special set of circumstances, the two combine to make mean-variance portfolio rules optimal (Ingersoll (1987)). The three canonical scenarios under which this is the case are: (a) dynamic trading with asset prices following diffusions; (b) static portfolio selection with mean-variance utility and arbitrary return distributions; and (c) static portfolio selection with elliptical distributions (e.g. Gaussian) and some additional restrictions on preferences.[4] Unfortunately, none of these situations apply to the typical investor in alternatives.

3.1 Why Is Mean-Variance Analysis Not Appropriate?

First, investors in alternatives (e.g. pension plans, endowments) rebalance their portfolios infrequently, and are typically subject to lockups. While they could hedge their investments dynamically, the *instantaneous* risk exposures necessary to do so are not observable to them due to the limited information dissemination typical in alternatives. In principle, this creates a role for a third-party auditor to collect information about hedge fund positions at intermediate frequencies, say daily, and compute benchmarks with instantaneously matched risk exposures for use in performance evaluation and hedging. Absent such an institution, an investment in alternatives is best viewed as owning an unspecified derivative contract (e.g. perhaps reflecting the dynamic trading of the hedge fund manger), whose terminal payoff has to be fit into their portfolio at their rebalancing frequency.[5]

Second, given investors' concerns about portfolio drawdowns, expected shortfalls, and other (left) tail measures, it is clear that investor preferences are not of the mean-variance type. A reasonable utility function should therefore describe preferences over all moments of the return distribution. Third, since the returns to hedge fund investing are similar to naked put writing strategies, which exhibit negative skewness,

[4]A classical example of the third case is the combination of exponential (CARA) utility with multivariate Gaussian returns. Power (CRRA) utility can be combined with multivariate Gaussian *log* returns to make mean-variance portfolio choice approximately optimal (Campbell and Viceira (2001)).

[5]Merton (1990) distinguishes between three economically important investor horizons (planning, decision, and trading) and emphasizes that these are likely to vary across investors and specific situations. While the managers of alternatives are likely to have very short trading horizons, the relevant trading horizon for the investors supplying capital to these funds is on the order of monthly or quarterly.

they cannot be described using an elliptical distribution. In particular, even if market index returns were elliptical, such that mean-variance portfolio choice applied well to public investments, the non-linearity of the hedge fund replicating portfolio, would make hedge fund returns fall outside the class of elliptical distributions, rendering them incompatible with mean-variance portfolio selection. Finally, there is strong evidence of stochastic volatility and market crashes at the level of the aggregate stock market index, such that the index returns themselves are not well described by the class of symmetric, elliptical distributions.

3.2 Incorporating Alternatives

To accommodate these concerns, our static portfolio selection framework combines power utility (CRRA) preferences, with a state-contingent asset payoff representation. Under power utility the investor prefers more positive values for the odd moments of the terminal portfolio return distribution (mean, skewness), and penalizes for large values of even moments (variance, kurtosis).[6] The second ingredient, the state-contingent payoff representation, originates in Arrow (1964) and Debreu (1959). To specify the joint structure of asset payoffs, we describe each security's payoff as a function of the aggregate equity index (here, the S&P 500).[7] This applies trivially to index options, since their payoffs are already specified contractually as a function of the index value. More generally, the framework requires deriving the mapping between a security's payoff and the market state space. Coval, et al. (2008) illustrate how this can be done for portfolios of corporate bonds, credit default swaps, and derivatives thereon. Importantly, by specifying the joint distribution of returns using state-contingent payoff functions, we can allow security-level exposures to depend on the market state non-linearly, generalizing the linear correlation structure implicit in mean-variance analysis. Finally, to operationalize the framework we need to specify the investor's risk aversion, γ, and the distribution of the state variable, which we parameterize using the log market index return, r_m.

The portfolio choice problem we study involves selecting the optimal mix of a risk-free security, the equity index, and hedge funds. The terminal distribution of the log index return at the investment horizon – assumed to be one month – is given by, $\phi(r_m)$. To illustrate, we parameterize this distribution using a normal inverse Gaussian (NIG) probability density, which allows the user to flexibly specify the first four moments. While this can be used to match the presence of skewness and kurtosis in returns, the probability

[6] To see this, note that if initial wealth is normalized to one, and r_w is the log return on the investors aggregate wealth, the expected utility can be expressed as:

$$E\left[\frac{W^{1-\gamma}}{1-\gamma}\right] = E\left[\frac{\exp\left((1-\gamma)\cdot r_w\right)}{1-\gamma}\right] = \frac{1}{1-\gamma}\cdot\sum_{n=0}^{\infty}\frac{(1-\gamma)^n}{n!}\cdot E\left[r_w^n\right]$$

[7] The same state-contingent payoff model is used in Coval, et al. (2009) to value tranches of collateralized debt obligations relative to equity index options, and in Jurek and Stafford (2011) – to elucidate the time series and cross-section of repo market spreads and haircuts.

and severity of tail events may be even more severe than implied by this parametrization.[8] For every \$1 invested, the state-contingent payoffs of the three assets are as follows: the risk-free asset pays $\exp(r_f \cdot \tau)$ in all states, the equity index payoff is, by definition, $\exp(r_m)$, and the payoff to the hedge fund investment is $f(r_m)$. The investor's problem is then to maximize his utility of terminal wealth, by varying his allocation to the equity market, ω_m, and the alternative investment, ω_a:

$$\max_{\omega_m, \omega_a} \quad \frac{1}{1-\gamma} \cdot E\left[\Big((1-\omega_m-\omega_a)\cdot \exp\left(r_f \cdot \tau\right) + \omega_m \cdot \exp\left(\tilde{r}_m\right) + \omega_a \cdot f(\tilde{r}_m)\Big)^{1-\gamma}\right] \tag{6}$$

where we have normalized total investor wealth to \$1, and the expectation is evaluated over the distribution of realizations for the log index return, $\tilde{r}_m$.

The payoff of the alternative investment is represented using a levered, naked put writing portfolio, as in the empirical analysis in Section 2. Specifically, we assume that the investor places his capital, ω_a, in a limited liability company (LLC) to eliminate the possibility of losing more than his initial contribution. Limited liability structures are standard in essentially all alternative investments, private equity and hedge funds alike, effectively converting their payoffs into put spreads. This has important implications for the investor's cost of capital, which we return to in the next section. Given a leverage of L, the quantity of puts that can be supported per \$1 of investor capital is given by:

$$q \quad = \quad \frac{L}{\exp\left(-r_f \cdot \tau\right) \cdot K(Z) - \mathcal{P}(K(Z),\, 1,\, \tau)} \tag{7}$$

where $K(Z)$ is the strike corresponding to a Z-score, Z. The put premium and the agent's capital grow at the risk free rate over the life of the trade, and are offset at maturity by any losses on the index puts to produce a terminal state-contingent payoff:

$$f\left(\tilde{r}_m\right) \quad = \quad \max\Big(0,\, \exp\left(r_f \cdot \tau\right) \cdot \left(1 + q \cdot \mathcal{P}(K(Z),\, 1,\, \tau)\right) - q \cdot \max\left(K(Z) - \exp\left(\tilde{r}_m\right),\, 0\right)\Big) \tag{8}$$

Using this payoff function, we also deduce that the limited liability legal structure corresponds to owning q puts at the strike, $K(LLC)$:

$$K(LLC) \quad = \quad K(Z) - \exp\left(r_f \cdot \tau\right) \cdot \frac{1 + q \cdot \mathcal{P}(K(Z),\, 1,\, \tau)}{q} \tag{9}$$

Having specified (1) the investor's utility function and (2) a description of the asset payoff profiles –

[8] Jurek and Stafford (2011) describe a procedure for constructing a crash return distribution using daily data. This distribution can then be combined with a stochastic volatility model for the diffusive component of stock return dynamics to produce a monthly return distribution.

which are the fundamental ingredients of any portfolio choice framework – we can now either solve for optimal allocations taking put prices, $\mathcal{P}$, as given; or solve for the investor's required rate of return on a hedge fund with parameters, $[Z,\, L]$, as a function of his portfolio allocation, $\{\omega_m,\, \omega_a\}$.

3.3 The Investor's Cost of Capital

In order to compute the investor's cost of capital at an ω_a allocation to a $[Z,\, L]$ hedge fund, we must first compute his shadow value for the put option written within the LLC structure. Given the investor's portfolio allocation, $\{\omega_m,\, \omega_a\}$, let the marginal utility of wealth in the state in which the return on public equities is r_m, be denoted by $U'\left(r_m \mid \omega_m,\, \omega_a\right)$. Note that the marginal utility is itself a function of the shadow put value, $\mathcal{P}^*(K(Z),\, 1,\, \tau)$, through its influence on the hedge fund payoff. The shadow value of the put option satisfies the following individual Euler equation:

$$E\left[U'\left(\tilde{r}_m \mid \omega_m,\, \omega_a\right) \cdot \max\left(K(Z) - \exp\left(\tilde{r}_m\right),\, 0\right)\right] \quad = \quad \mathcal{P}^*(K(Z),\, 1,\, \tau) \tag{10}$$

Once the shadow value of the index put with strike $K(Z)$ has been pinned down, the shadow value of the alternative investment can be obtained from the individual Euler equation for the LLC payoff:

$$p_a^*\left(\omega_m,\, \omega_a\right) \quad = \quad E\left[U'\left(\tilde{r}_m \mid \omega_m,\, \omega_a\right) \cdot f\left(\tilde{r}_m;\, \mathcal{P}^*\right)\right] \tag{11}$$

which corresponds to an annualized required rate of return of:

$$r_a^*\left(\omega_m,\, \omega_a\right) \quad = \quad \frac{1}{\tau} \cdot \log \frac{E\left[f\left(\tilde{r}_m;\, \mathcal{P}^*\right)\right]}{p_a^*} \tag{12}$$

3.3.1 The CAPM cost of capital

In the ensuing analysis, we compare the proper cost of capital, (12), with the value obtained under a CAPM budgeting rule. To re-derive this rule recall that under mean-variance analysis, the optimal portfolio allocation is $\omega = \frac{1}{\gamma} \cdot \Sigma^{-1} \cdot (\mu - r_f)$, where γ is the investor's risk aversion, Σ is the covariance matrix of asset returns, and μ is a vector of mean asset returns. The solution to the investor's optimal portfolio choice problem, ω, can be viewed as specifying weights under the which the agent's required rate of return on each asset matches its observed mean rate of return. When this is the case, further portfolio rebalancing cannot make the agent better off and the proposed portfolio is indeed optimal. We can also apply this result "in reverse" to derive the agent's required rate of return on an asset as a function of his allocation, $r^*(\omega) = r_f + \gamma \cdot \Sigma \cdot \omega$.

Recall that mean-variance analysis applies under a restrictive set of preferences and distributional assumptions that generally will not be applicable to alternative assets due to their underlying payoff non-linearities. Nonetheless, suppose an investor were to characterize the joint payoffs to equities and alternatives using exclusively their mean returns and covariances, as opposed to their joint state-contingent payoff profiles. Let the variances of the two risky assets – equity and the alternative – be given by σ_m^2 and $\sigma_a^2 = \beta \cdot \sigma_m^2 + \sigma_\epsilon^2$, where $\beta = \frac{Cov[r_a, r_m]}{Var[r_m]}$ is the CAPM β of the alternative on the equity index and σ_ϵ is the idiosyncratic volatility of the alternative. We can then write the investor's required rate of return on the alternative asset as:

$$r_a^{*,CAPM}(\omega_m, \omega_a) \quad = \quad r_f + \omega_m \cdot \beta \cdot \left(\gamma \cdot \sigma_m^2\right) + \omega_a \cdot \left(\gamma \cdot \sigma_a^2\right) \tag{13}$$

The two notable features of this formula are that, an investor with a non-infinitesimal allocation to the hedge fund charges for its total risk, σ_a, rather than simply the component that covaries with the market returns, and that the required rate of return is linear in the portfolio weights.

In practice, this rule is applied assuming that: (a) prior to adding the new security (the alternative), the agent is at his optimal mix of cash and the market portfolio; and, (b) the investment in the new asset – the alternative – is infinitesimal ($\omega_a \approx 0$). If we denote the market risk premium by λ, the agent's optimal cash-equity mix has, $\omega_m^* = \frac{\lambda}{\gamma \cdot \sigma_m^2}$, which taken together with a marginal allocation to the new securities, yields the following cost of capital for alternatives:

$$\begin{aligned} r_a^{*,CAPM}(\omega_m^*, 0) \quad &= \quad r_f + \left(\frac{\lambda}{\gamma \cdot \sigma_m^2}\right) \cdot \beta \cdot \left(\gamma \cdot \sigma_m^2\right) \\ &= \quad r_f + \beta \cdot \lambda \end{aligned} \tag{14}$$

The CAPM equilibrium logic identifies the market risk premium, λ, as the rate of return, under which the representative (marginal) investor is fully invested in the risky portfolio. Given a risk aversion, $\tilde{\gamma}$, for the representative agent, the equilibrium risk premium is given by, $\lambda = \tilde{\gamma} \cdot \sigma_m^2$.

3.3.2 Baseline model parameters

The investor's cost of capital is a function of model parameters describing the distribution of the (log) market return, investor's risk tolerance, investor's allocation to other assets, and the structure of the alternative investment (e.g. option strike price and leverage). Before turning to a discussion of the comparative statics of the investor's cost of capital, we describe the baseline model parameters:

- **Risk aversion, γ:** We consider investors with two different levels of risk aversion: marginal ($\gamma = 2$)

and endowment ($\gamma = 3.3$). These levels of risk aversion are chosen such that, in the absence of alternatives, the marginal investor would be fully invested in equities, and the endowment investor would hold a portfolio of roughly 40% cash and 60% equities, corresponding to an allocation commonly used as a benchmark by endowments and pension plans.

- **Distribution, $\phi(r_m)$:** To illustrate the key features of the framework, we rely on the normal inverse Gaussian (NIG) probability density. The parameters of the monthly return distribution are chosen to roughly match historical features of monthly S&P 500 Z-scores. Following Jurek and Stafford (2011), we select skewness and kurtosis values that combine to produce a left-tail "event" once every 5 years that results in a mean monthly Z-score of -3.5.[9] These values are skewness of -1 and kurtosis of 7. This pins down a conditional Z-score distribution from which we simulate log monthly index returns:

$$r_m = \left(r_f + \lambda - \frac{1}{\tau} \cdot \ln E\left[\exp\left(\sigma \cdot \sqrt{\tau} \cdot \tilde{Z} \right) \right] \right) \cdot \tau + \sigma \cdot \sqrt{\tau} \cdot \tilde{Z}, \quad \tilde{Z} \sim \text{NIG}\,(0,\, 1, -1, 7)$$

  We set σ to 18%, which corresponds to the annualized volatility of monthly returns in sample; and the equity risk premium, λ, to 6.5%, which is approximately equal to the required return for the marginal investor. Additionally, we set the riskfree rate, r_f , equal to 3% and the dividend yield, q, to 2%.[10]

- **Alternative investment, $[Z,\, L]$:** Given the empirical results in Sections 1 and 2, the aggregate risk exposure of the alternative investment universe is described using the naked S&P 500 put writing strategy, which writes puts with strikes corresponding to $Z = -1$, a deploys a leverage of two, $L = 2$.

3.3.3 Comparative statics: portfolio composition

A prediction of mean-variance analysis is that the investor's required rate of return is a linear function of his allocation, (12). We explore the practical deviations between this rule and the model-implied cost of capital for investments in equities and alternatives. Figure 3 compares the investor's cost of capital as he shifts weight from the risk-free security into either equities (left panel) or alternatives (right panel). The linear mean-variance cost of capital is computed using (13), setting $\omega_a = 0$ and varying $\omega_m \in (0, 1)$ in the

[9] For comparison, the mean value of the Z-score under the standard normal (Gaussian) distribution – conditional on being in the left 1/60 percent of the distribution – is -2.5.

[10] The risk premium required by an investor with risk aversion γ, who holds exclusively the equity index is given by:

$$\lambda \quad = \quad \gamma \cdot \sigma^2 + \frac{1}{\tau} \cdot \left(\sum_{n=3}^{\infty} \frac{\kappa_n}{n!} \cdot \left(\sigma \cdot \sqrt{\tau} \right)^n \cdot \left(1 + (-\gamma)^n - (1-\gamma)^n \right) \right)$$

where the κ_i are the cumulants of the distribution of the Z innovation. For a Gaussian distribution, all cumulants $n > 2$ are equal to zero, and the equity risk premium is equal to $\gamma \cdot \sigma^2$. Under the baseline model parameters, the Gaussian component of the equity risk premium equals 6.48%, with the higher order cumulants contributing an additional 0.19%.

left panel; and by setting $\omega_m = 0$ and varying $\omega_a \in (0, 1)$ in the right panel. The proper (model) cost of capital is computed using (12).

The left panel indicates that the model and linear (mean-variance) costs of capital are essentially identical for the equity investment, in spite of the fact that the equity return distribution is not elliptical and therefore, formally incompatible with mean-variance analysis. The practical deviation between the proper and mean-variance costs of capital for traditional assets is negligible. By contrast, the required rate of return for an investor in cash and alternatives is meaningfully convex in the risky share, resulting in large deviations relative to the linear mean-variance rule. The rapid growth of the required rate of return on alternatives highlights the strong interaction between the portfolio allocation and the nonlinearity in the payoff profile of alternatives. This suggests that mean-variance cost of capital computations are likely to be useful for traditional assets, but misleading for investments in alternatives.

The required costs of capital for the two risky assets interact when the securities are combined in a single portfolio. To examine this interaction we compute the investor's cost of capital for the alternative as a function of its share in the risky portfolio of each investor (Figure 4). We contrast this cost of capital with the value obtained under the CAPM β logic, (14), which is predicated on an infinitesimal allocation to the alternative. The β of the alternative investment is given by $q \cdot \Delta$, where Δ is the delta of the option portfolio that describes the systematic exposure of the alternative investment (short q options at $K(Z)$ and long q options at $K(LLC)$).[11] Under the baseline model parameters, the beta of the alternative with option strike, $Z = -1$, and leverage, $L = 2$, is equal to 0.33, essentially matching the empirical beta of the hedge fund indices examined earlier. We consider the portfolios of two investors: an equity only investor ($\gamma = 2$), and an endowment investor ($\gamma = 3.3$). The endowment investor is assumed to hold 80% of his portfolio in risky securities (equities + alternatives) and 20% of his portfolio in cash, corresponding to the typical allocation of a sophisticated endowment (Lerner and Schoar (2008)).[12]

Figure 4 illustrates two important points. First, the cost of capital for alternatives meaningfully departs from the CAPM β rule even for infinitesimal allocations for both investors. For example, the marginal investor demands an additional 2% in excess return relative to the CAPM benchmark, for an infinitesimal allocation. The same wedge is roughly 3.5% for the endowment investor, reflecting his greater risk-aversion, and also the somewhat aggressive risk posture of his baseline portfolio allocation (i.e. the 20% cash holding

[11] Option deltas are computed using the Black-Scholes (1973) formula with implied volatilities set according to, $\sigma(K) = \sigma \cdot \exp(-\eta \cdot \log K)$. The elasticity parameter, η, is set equal to -3, to roughly match the historical CAPM β's of the cross-section of put writing strategies in Table 2.

[12] Lerner and Schoar (2008) highlight a three-fold increase in alternative investment allocation at endowments over the period from 1992-2005. For example, at the end of 2005 the median Ivy League endowment held 37% in alternatives, representing a 50% share of their risky asset portfolio (alternatives + equities). For the purposes of our risk analysis, we conservatively treat fixed income investments as risk free.

is below his benchmark allocation of 60% equities and 40% cash in the absence of alternatives). In other words, even at infinitesimal allocations, investors reliant on CAPM cost of capital benchmarks will "observe" meaningful α's, even though a proper cost of capital would indicate the assets are priced correctly. Second, the magnitude of the wedge between the proper cost of capital and CAPM benchmark is increasing in the share of alternatives in the risky asset portfolio. In practice, the fixed costs of investing in alternatives, imply that the share of alternatives in investor portfolios will not be infinitesimal. As mentioned earlier, sophisticated endowments hold between 25% and 50% of their *risky* portfolio in alternatives. At these allocations, endowment investor's would have to observe CAPM α's of 4% to 5% per year just to cover their properly computed cost of capital.

3.3.4 Comparative statics: return distribution

Next we examine the investor's cost of capital as a function of the riskiness of the equity market return distribution, while holding investor portfolios fixed. In particular, each investor is assumed to allocate 35% of their risky portfolio to alternatives: the marginal investor holds 0% cash, 65% equities, and 35% alternatives; and the endowment investor holds 20% cash, 52% equities, and 28% alternatives. We then consider two dimensions of risk: (1) the volatility of the market return, σ_m; and (2) the severity of extreme, left-tail events. We measure tail risk as the average Z-score, conditional on being below the p^{th} percentile of the distribution. We set $p = \frac{1}{60}$, such that the mean tail Z value is interpretable as corresponding to a "once-in-five-year" event. To vary the severity of extreme events we change the skewness, s, of the return distribution, holding volatility fixed at its baseline value. We make use of a restriction of the normal inverse Gaussian distribution that the minimum feasible kurtosis can be described as a function of the skewness ($k_{min}(s) = 3 + \frac{5}{3} \cdot s^2$), allowing us to uniquely map the comparative static in skewness into tail risk. As before, we compare the proper required rates of return with the CAPM β cost of capital, (14).

The top panel of Figure 5 displays the proper and CAPM required rates of return for the marginal investor and the endowment investor for the alternative investment as volatility increases. The CAPM rule always underestimates the required rate of return, and this wedge is increasing in volatility. This wedge becomes severe at modest levels of volatility, as the linear beta-approximation fails to keep pace with the true risks. The bottom panel compares the proper and CAPM costs of capital as tail risk increases. Here, since volatility is being held constant, the CAPM required rates of return are constant. As before, the wedge between the proper required rates of return and the approximation are steeply increasing in the mean severity of the tail events.[13]

[13] The CAPM cost of capital is based on an equilibrium market risk premium equal to $\tilde{\gamma} \cdot \sigma_m^2$. At low tail risk magnitudes, this overestimates the equity market risk premium that would be charged by the marginal investor, when 35% of his risky portfolio

3.3.5 Comparative statics: leverage

An interesting feature of the framework described in this paper is that there are specific dimensions of risk that systematically affect the investor's required rate of return, but that are completely missed by the CAPM rule. Strategies that shift risk into the left tail, while holding their CAPM betas constant, increase the investor's required rate of return. This effect is presented in Figure 6 for the marginal investor and the endowment investor.

To illustrate this effect, we first select a target CAPM beta of 0.33, which coincides with the exposure of the $[Z = -1, L = 2]$ strategy. Then, as we vary the leverage factor, L, we adjust the strike price, $K(Z^*)$, and quantity of the options written within the LLC to keep the CAPM beta of the alternative portfolio – inclusive of the LLC put option – fixed at the target value. Intuitively, to keep the portfolio β fixed, the higher leverage strategies require writing options that are further out-of-the-money, and thus have lower deltas. For example, the $[Z = -1, L = 2]$, $[Z = -1.8, L = 4]$, and $[Z = -2.6, L = 10]$ strategies all have the same CAPM β of 0.33. The CAPM rule therefore predicts that the required cost of capital in excess of the risk free rate is a constant 2.15% ($= 0.33 \cdot 6.5\%$) across these strategies. The marginal investor ($\gamma = 2$), who is fully invested in risky assets, with 35% in alternatives and 65% in the equity index, requires 4.50% for the $[Z = -1, L = 2]$ strategy, 7.36% for the $[Z = -1.8, L = 4]$ strategy, and 14.66% for the $[Z = -2.6, L = 10]$ strategy. Similarly, the endowment investor requires higher rates of return as leverage increases, even though the CAPM β is fixed. The extreme variation in investor required rates of return highlights that evaluating the performance of alternative investments is likely to be challenging in practice. We discuss some further implications of levering safe assets to match the CAPM risks of inherently riskier assets in the next section.

4 Implications

The historical investor experience in alternatives has been very attractive from the perspective of the CAPM. Table 1 reports that the aggregate hedge fund universe realized roughly a 6% excess return from 1996 through 2010, a period where the S&P 500 stock index averaged only 4.4% above Treasury bills. The aggregate hedge fund indices realized a CAPM beta around 0.35, implying an annual CAPM alpha of 4.5%. To the extent that investors are uncomfortable with the CAPM benchmark, they often take comfort in the fact that absolute returns for alternatives have been higher than those of traditional risky asset portfolios

in held in alternatives. Intuitively, when tail risks are not severe, the 35% allocation to the $\beta = 0.33$ alternative, makes the marginal investor's portfolio less risky than with a 100% allocation to equities. Consequently, the investor's cost of capital for the alternative, can actually drop below that implied by the CAPM.

with half the return volatility of the stock market.

What are the true risks of alternatives? Properly evaluating the risks of alternative investments is challenging. At the individual fund level, this will be especially difficult. The nature of a private investment company is such that the manager typically has discretion over the composition of the underlying risks, financial leverage, and hedging rules, all of which can vary at high frequency (e.g. intraday). Periodic reporting occurs monthly or quarterly, and direct monitoring of portfolio risks by limited partners is typically not feasible. The investor must infer risks from realized returns, which are typically unrevealing about the specific nature of their downside exposure. Measuring the risks at the asset class level is easier because the idiosyncrasies of individual funds tend to diversify, but this too is a complex exercise. The empirical evidence reported in Section 2 suggests that the realized risks of the aggregate hedge fund indices are better matched by a simple put writing strategy than by the linear CAPM benchmark.

What are the required rates of return, given the risks? Investors relying on the CAPM rule will charge for the risks contributed to their portfolios assuming small linear exposures, both of which are inaccurate descriptions of the real world investment for a typical investor in alternatives. The comparative statics of the generalized model indicate that the required rates of return for downside exposures can deviate meaningfully from those implied by a CAPM rule.

The observation that the risks of alternative investments can be accurately mimicked with a simple portfolio of index put options and the simple framework described in Section 3 combine to produce a number of interesting implications relative to the standard mean-variance framework. We now explore some implications of the investor experience from the perspective of the CAPM and the generalized model by looking across investors with different risk preferences, as well as across strategies with different downside exposures.

4.1 Who Should Own Alternatives?

The decision to allocate capital to alternative investments is typically made by sophisticated investors with a long investment horizon who are in a financial position to bear the risks of these often illiquid investments. Given the specialized investment expertise required to properly evaluate and monitor these investments, it is common for allocations to be relatively large to amortize the fixed costs associated with expanding the investment universe to include alternatives (Merton (1987)). Additionally, we would expect that the investors in alternatives are relatively risk tolerant given the tendency for these investments to fail to pay off in poor economic states where the marginal value of wealth is high. Highly risk averse investors will require very large risk premia for downside risks.

The primary investors in alternatives are wealthy individuals, endowments, and pension funds. Wealthy individuals are principal investors (often with the help of advisors) who are likely to be relatively risk tolerant. Pension funds are intermediaries who are investing on behalf of individuals, who in the aggregate must have average risk aversion. Finally, there is a wide cross section of endowment funds. Many large university endowments use a benchmark portfolio consisting of 60% stock and 40% bonds, which implies a relatively high level of risk aversion ($\gamma = 3.33$ in our baseline setup).

The so-called "endowment model" is based on the premise that illiquid investments should earn an additional risk premium that long-term institutional investors will have a comparative advantage in bearing (Swensen (2000)). This is the motivation for relatively large allocations to alternatives. The comparative statics demonstrate that a typical endowment requires 4%-5% CAPM alpha to cover the true required rate of return of the individuals on whose behalf they act, which coincides roughly with what has been realized by the aggregate hedge fund indices. Table 3 reports the annual realized returns in excess of the one-month Treasury bill return for the HFRI Composite Index and the put-writing portfolio (Panel A). The table also reports the *ex ante* required rates of return computed using the linear CAPM benchmark with a constant beta ($\beta = 0.35$), and the proper rates accounting for the non-linear downside payoff exposure, computed for both the marginal investor and the endowment investor described above. To calculate the *ex ante* required excess returns, each month we determine the monthly required rate of return assuming that volatility will be 0.8 times the prevailing VIX level (Figure 7) and then compound these monthly returns in excess of the riskfree rate of interest to produce an annualized excess return. Panel B of Table 3 shows that over the period from 1996-2010, the CAPM alpha for the aggregate hedge fund universe as proxied by the HFRI Composite Index was 3.5% per year (t-stat: 1.8). However, to the extent that the downside exposure of this index is better explained by the [$Z = -1$, $L = 2$] put-writing strategy, the average annual required rates of return for the marginal and endowment investor were considerably higher than suggested by the CAPM. For example, the marginal investor required 5.6% on average over this period, implying an alpha of only 0.5% (t-stat: 0.3) per year, while the endowment investor required an average excess return of 7.9%, implying an average annual alpha of -1.8% (t-stat: -0.8). The annualized alphas of the mechanical put writing strategy exceed those of the HFRI index by 3.8% per year for both investors, indicating that a passive exposure to downside risk may be preferable to direct investment in hedge funds.

Finally, it is reasonable to expect that the first endowments to allocate to alternatives have actually earned the returns associated with the published indices, suggesting that they have nearly covered their cost of capital, but have not earned any additional risk premium for bearing the illiquidity of alternatives. It is also likely that more recent investors in alternatives have earned something closer to the average fund-of-fund return (Table 1), which is 300 basis points lower than the reported average aggregate hedge fund return, suggesting that these investors have not covered their cost of capital.

4.2 Evaluating Levered Strategies Matched on CAPM β

When viewed from the perspective of the linear CAPM β risk adjustment, levering safe assets frequently empirically dominates bearing inherently risky exposures, both among traditional assets and alternatives. For example, a levered portfolio of investment grade bonds has historically outperformed an unlevered portfolio of non-investment grade bonds. Black, Jensen, and Scholes (1972) observed that a levered portfolio of low beta stocks seemed to outperform portfolios of high beta stocks. Frazzini and Pedersen (2010) link the risk-adjusted returns on strategies that short high-beta instruments and are long, levered low-beta instruments to funding constraints. Similarly, the implied fees in Table 2 illustrate that – holding the realized CAPM β's constant – levered short positions in further out-of-the-money (i.e. safer) put options outperform less leveraged portfolios of put options that are closer to at-the-money.

However, the analysis in Section 3 indicates that investor required rates of return vary significantly depending on the strategy's exposure to losses in the left tail of economic outcomes, even after matching CAPM β's. In particular, strategies that choose relatively safer economic risk exposures (i.e. exposures less likely to sustain a loss), but apply higher leverage, require much higher rates of return than predicted by the linear CAPM β risk adjustment. Proper evaluation of the performance of these strategies therefore hinges crucially on understanding their downside risk exposure. The statistical analysis of realized returns, particularly when originating from within *ex post* successful economies and/or periods of economic expansion, requires the use of risk models that allow for the possibility of non-linear tail risk exposures.

Our analysis predicts that investors using the CAPM rule-of-thumb are likely to conclude in favor of large *ex ante* α's for strategies that apply high leverage to safe assets, and therefore overallocate their portfolios to these investments. In Figure 8, we illustrate the proper portfolio allocations to three alternative strategies that group losses in progressively worse economic states, by increasing the leverage factors (higher L values), while making the underlying assets safer (terminal payoffs produce losses beginning at more negative Z values). The strategies are designed to have constant CAPM β's equal to 0.33, and offer an *increasing* sequence of CAPM α's of 3% (left panel), 4% (center panel), and 5% (right panel). While the standard mean-variance logic predicts that all investors – irrespective of risk aversion – would increase their allocations as α's increase, in fact, the opposite is observed. As the leverage factor increases, the required rate of return grows faster than the linear CAPM α, such that it becomes optimal for all investor to *decrease* their allocation to the alternative. This example demonstrates that without a proper understanding of the underlying downside risk exposure, errors in required rates of return can translate into inappropriate portfolio allocations.

The same failure to appreciate the underlying downside risk exposures, which gives rise to spurious CAPM α estimates, translates *ex post* into observations of "black swans" and increased downside correlations during significant economic declines.[14] From the perspective of the framework described in this paper "black swans" can be seen as arising from a combination of two investor errors. First, the distribution of economic conditions has a fat left-tail, such that large systematic shocks occur more frequently and are more severe than the normal distribution underlying mean-variance analysis predicts. Second, the presence of a non-linear downside exposure in many strategies means that the losses from these events are significantly more severe than investors expected them to be, conditional on the magnitude of the underlying shock. To the extent that a large number of investors make these errors, their short-run equilibrium impact may be magnified even further as there is an aggregate shortage of risk bearing capital.

Among alternative investments – distressed investing and leveraged buyouts – are two examples where extreme downside exposure may be created by levering safe assets. These are situations were common intuition on required rates of return is particularly susceptible to errors, and therefore where opportunistic agents can be exploitative. To illustrate investor mistakes arising from using linear CAPM benchmarking rules, we examine the returns to these strategies. As in the case of the hedge fund indices, we compare the asset class performance to that of an investable, risk-matched alternative implemented in equity index options. We show that while generating high absolute returns by levering safe assets during the 1996-2010 period was not difficult, endowment investors who were attracted to such "superstar" strategies, may well have underperformed their proper *ex ante* required rates of return.

4.2.1 Distressed investing

On the surface, distressed investing may not appear consistent with the description of a levered safe asset. However, because these assets are typically purchased at a considerable discount (margin of safety), they have a high probability of returning more than their cost. These investments often require time and management expertise or discipline to realize a business turnaround that is largely independent of aggregate economic conditions. However, in the event of a severe deterioration of aggregate market conditions, these investments will have very low recovery values. Their state-contingent payoff is therefore similar to a deep out-of-the-money digital call on a broad market index (e.g. a highly levered bet against economic

[14]Because hedge fund and put-writing returns are not elliptically distributed, the linear (Pearson's) correlation coefficient is an inadequate measure of dependence. To see this we can compute two alternative measures of dependence, Kendall's τ and Spearman's ρ. Since the value function of any put writing strategy is a strictly increasing function of r_m, in our model we have $\tau = \rho = 1$ (i.e. the two random variables are perfectly dependent and co-monotonic).

catastrophe).

The nature of this risk exposure is consistent with the fact that in 2008 the HFRI Distressed/Restructuring subindex sustained a drawdown that was 30% larger than that of the aggregate hedge fund universe (Table 1). As in the case of hedge funds, we construct a put writing strategy that matches the risk metrics (volatility, CAPM β, and drawdown) of the distressed index. Using the data in Table 2, one quickly sees that the $[Z = -2, L = 6]$ strategy, which applies higher leverage and writes deeper out-of-the-money (safer) options relative to strategy replicating the aggregate hedge fund risk exposure, represents the best match. The strategy drawdown of -27% exactly matches the drawdown on the HFRI Distressed/Restructuring subindex, the CAPM β of the strategy is 0.35 vs. 0.36 for the index, and the volatilities are 9.48% for the strategy and 9.36% for the index. Crucially, the implied fee indicates that the put writing strategy outperforms the HFRI Composite index by 7.52% per year, and the Distressed/Restructuring subindex by slightly under 7% per year. This suggests that the margin of safety in most distressed investments may not have been large enough, as they underperform this risk-matched, investable alternative.

4.2.2 Superstar investors

Managers whose strategies deliver high average returns with few negative realizations are frequently considered to be "superstars." We show that such return series are consistent with applying high leverage to safe assets, and may simply be earning a risk premium appropriate for bearing the risk of loss in a severe aggregate recession. Moreover, these strategies can be constructed to have moderate linear CAPM β's and return volatilities. As a result, investors who evaluate "superstar" investors using conventional linear models risk concluding in favor of manager skill, where in fact there is none.

To illustrate this, we compute the returns to a strategy that applies a leverage of $8x$ to very safe assets (short positions in monthly index put options that are on average 15% out-of-the-money). While the leverage factor of eight is extreme from the perspective of a strategy implemented in liquid securities subject to intra-daily mark-to-market, managers of alternatives are frequently given considerable discretion with regards to how and when portfolios are marked. We report the performance of the $[Z = -2, L = 8]$ strategy in Table 4 alongside the annual returns on public equities (S&P 500), hedge funds (HFRI Composite), and the hedge-fund replicating strategy, $[Z = -1, L = 2]$. The table reports annual returns for the put writing strategies before fees.

The $[Z = -2, L = 8]$ strategy earns has the properties of a "superstar" investment: (1) it features a pre-fee average annual return of 23.2%; (2) its 2008 drawdown is half as severe as the drawdown in the S&P 500; (3) it experiences only one year with a negative return (2008), while delivering double-digit returns in all but one of the remaining years; and (4) has an annualized volatility smaller than the S&P 500 index. In fact, many other strategies with similarly (or more) attractive return series can be constructed within this sample. This highlights that nothing disastrous has occurred in U.S. markets during the last 15 years to truly reveal the hidden, downside risks exposure of strategies applying high leverage to safe assets.

The risk-return characteristics of the $[Z = -2,\ L = 8]$ strategy, when assessed either in an absolute sense, relative to peers, or relative to an index, would lead many to conclude that it represents a highly compelling investment. To evaluate this claim, we adopt the perspective of an endowment and compute its *ex ante* required rate of return for such an investment. As in Section 4, we assume the endowment investor has a risk aversion of $\gamma = 3.3$, and holds a baseline portfolio of 20% cash and 80% risky assets. He is then allowed to adjust the share of the risky portfolio invested in alternatives, which are now modeled as an equally-weighted mix of hedge fund exposure ($[Z = -1,\ L = 2]$) and highly levered safe assets ($[Z = -2,\ L = 8]$). Figure 9 plots the required rate of return on each of the alternatives and their equally-weighted mix. Fixing all of the model parameters at their baseline values, and assuming a 35% share of alternatives in the risky portfolio, the endowment will be willing to make a marginal investment in the $[Z = -2,\ L = 8]$ strategy if it offers an *ex ante* risk premium of 17.8%. Consequently, if the strategies employed by "superstar" managers are well described by the payoff to levering safe assets ($[Z = -2,\ L = 8]$), endowment investors are likely to find themselves underperforming their required rate of return despite achieving impressive absolute returns.

Figure 9 also highlights that an increased portfolio allocation to highly-levered safe assets meaningfully increases the required rate of return on other alternatives, such as hedge funds. Recall from Section 4, that when the alternative portfolio was comprised exclusively of the hedge fund replicating strategy, the required risk premium on hedge funds was between 6.0-6.5% as the share of alternatives in the risky portfolio varied between 25-50% (Figure 4). However, once half of the alternative portfolio is comprised of highly-levered safe assets, the corresponding range of the required rates of return on the marginal hedge fund investment increases to 7.6%-11.0%. From the perspective of our state-contingent portfolio allocation framework, the dramatic rise in the required rates of return reflects the fact that both assets are predicted to decline in tandem as economic conditions worsen. Any additional allocation to alternatives represents a significant marginal contribution to the risk of the entire portfolio, when viewed from the perspective of the investor's marginal utility. In practice, investors frequently seem surprised by increases in return correlations between alternatives and traditional assets (or between alternatives themselves) as economic conditions deteriorate, suggesting they may not fully appreciate their portfolio-level downside risk exposure. This *ex post* surprise likely coincides with meaningful *ex ante* errors in estimates of required rates of return, and therefore inappropriate capital allocations.

5 Conclusion

This paper argues that the risks born by hedge fund investors are likely to be positive net supply risks that are unappealing to average investors, such that they may earn a premium relative to traditional assets; and that over the rebalancing horizon of the typical investor in these strategies the payoff profile has a distinct possibility of being nonlinear with respect to a broad portfolio of traditional assets. We document that a simple put writing strategy closely matches the risks observed in the time series of the aggregate hedge fund universe. We then study the required rates of return for a variety of investors allocating capital to the risk-matched put writing portfolio to develop estimates of the cost of capital for alternative investments.

The setup studied in this paper, one where an investor holds a nonlinear exposure over a substantial period of time, is one that is infrequently examined. Typically, it is assumed that investors hold assets whose risks can be well described by their covariance with each other over their rebalancing horizon. However, the analysis in this paper suggests that the risks of alternative investments do not comply with this assumption. One of the attractive features of this simple generalized framework is that it conceptually requires no information beyond the traditional analysis, although in practice it will require more sophisticated judgment over the state-contingent risk profile of alternative investments.

An accurate assessment of the cost of capital is fundamental to the efficient allocation of capital throughout the economy. Investment managers should select risks that are expected to deliver returns at least as large as those required by their capital providers. The investors in alternatives should require returns for each investment that compensate them for the marginal contribution of risk to their overall portfolio. In the case of investments with downside exposure, the magnitude of these required returns is large relative to those implied by linear risk models. As the allocation to downside risks gets large, the marginal contribution of risk to the overall portfolio expands quickly, requiring further compensation. The calibrations in this paper suggest that despite the seemingly appealing return history of alternative investments, many investors have not covered their cost of capital.

References

[1] Agarwal, V., and N. Naik, 2004, "Risks and Portfolio Decisions Involving Hedge Funds," *Review of Financial Studies*, 17, p. 63-98.

[2] Arrow, K. J., 1964, The Role of Securities in the Optimal Allocation of Risk Bearing, *Review of Economic Studies*, 31, p. 91-96.

[3] Asness, C., R. Krail, and J. Liew, 2001, "Do Hedge Funds Hedge?," *Journal of Portfolio Management*, 28, p. 1-19.

[4] Bakshi, G., and N. Kapadia, 2003, "Delta-Hedged Gains and the Negative Market Volatility Risk Premium," *Review of Financial Studies*, 16, p. 527-566.

[5] Black, F., 1986, "Noise," *Journal of Finance*, 41, p. 529-543.

[6] Black, F., and M. Scholes, 1973, "The pricing of options and corporate liabilities," *Journal of Political Economy* 81, p. 617-654.

[7] Black, F., M. Jensen, and M. Scholes, 1972, "The Capital Asset Pricing Model: Some Empirical Tests," in *Studies in the Theory of Capital Markets.* Michael C. Jensen, ed. New York: Praeger, p. 79-121.

[8] Brunnermeier, M., and L. Pedersen, 2009, "Market Liquidity and Funding Liquidity," *Review of Financial Studies*, 2009, 22, p. 2201-2238.

[9] Campbell, J. Y., and L. Viceira, 2001, *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors*, Oxford University Press.

[10] Coval, J., J. W. Jurek, and E. Stafford, 2009, "Economic Catastrophe Bonds," *American Economic Review*, 99, p. 628-666.

[11] Debreu, G., 1959, *Theory of Value*, Wiley.

[12] Duarte, J., F. Longstaff, and F. Yu, 2007, "Risk and Return in Fixed-Income Arbitrage: Nickels in Front of a Steamroller?," *Review of Financial Studies*, 20, p. 769-811.

[13] Frazzini, A. and L. Pedersen, 2010, "Betting Against Beta," NYU Stern School of Business, working paper.

[14] Fung, W., and D. Hsieh, 2001, "The Risk in Hedge Fund Strategies: Theory and Evidence from Trend Followers," *Review of Financial Studies*, 14, p. 313-341.

[15] Fung, D. and W. Hsieh, 2002, "Asset-Based Hedge-Fund Styles Factors for Hedge Funds," *Financial Analysts Journal*, 58, p. 16–27.

[16] Fung, D. and W. Hsieh, 2004, "Hedge Fund Benchmarks: A Risk Based Approach," *Financial Analysts Journal*, 60, p. 65–80.

[17] Getmansky, M., A. Lo, I. Makarov, 2004, "An Econometric Model of Serial Correlation and Illiquidity in Hedge Fund Returns," *Journal of Financial Economics*, 74, p. 529-609.

[18] Glosten, L., and R. Jagannathan, 1994, "A Contingent Claim Approach to Performance Evaluation," *Journal of Empirical Finance*, 1, p. 133-160.

[19] Hasanhodzic, J., and A. Lo, 2007, "Can Hedge-Fund Returns Be Replicated?: The Linear Case," *Journal of Investment Management*, 5, p. 5-45.

[20] Ibbotson, R. G., P. Chen, and K. Zhu, 2010, "The ABCs of Hedge Funds: Alphas, Betas, and Costs," Yale School of Management, working paper.

[21] Ingersoll, J. E., 1987, *Theory of Financial Decision Making*, Rowman and Littlefield Publishers.

[22] Jurek, J. W., and E. Stafford, 2011, "Crashes and Collateralized Lending," Princeton University, working paper.

[23] Kat, H. and H. Palaro, 2005, "Who Needs Hedge Funds? A Copula-Based Approach to Hedge Fund Return Replication," Cass Business School, working paper.

[24] Lerner, J., A. Shoar, and J. Wang, 2008, "Secrets of the Academy: The Drivers of University Endowment Success," *Journal of Economic Perspectives*, 22, p. 207-222.

[25] Merton, R. C., 1974, "On the Pricing of Corporate Debt: The Risk Structure of Interest Rates," *Journal of Finance*, 29, p. 449-470.

[26] Merton, R. C., 1987, "A Simple Model of Capital Market Equilibrium with Incomplete Information," *Journal of Finance*, 42, p. 483-510.

[27] Merton, R. C., 1990, *Continuous-Time Finance*, Blackwell.

[28] Mitchell, M., and T. Pulvino, 2001, "Characteristics of Risk in Risk Arbitrage," *Journal of Finance*, 56, p. 2135-2175.

[29] Shleifer, A. and R. Vishny, 1997, "The Limits of Arbitrage," *Journal of Finance*, 52, p. 35-55.

[30] Swensen, D., 2000, *Pioneering Portfolio Management: An Unconventional Approach to Institutional Investment*, Free Press.

Table I
Historical Hedge Fund Performance

This table reports the performance of investments in risk-free bills, public equities and hedge funds between January 1996 and October 2010. *T-bill* is the return on the one-month U.S. Treasury T-bill obtained from Ken French's website. *S&P 500* is the total return on the S&P 500 index obtained from the CRSP database. *HFRI Composite Index* is the Hedge Fund Research Inc. Fund Weighted Composite Index. *DJ/CS Broad Index* is the Dow Jones Credit Suisse Broad Hedge Fund Index. *HFRI Fund-of-funds Index* is the Hedge Fund Research Inc. Fund of Funds Composite Index. *HFRI Distressed Index* is the Hedge Fund Research Inc. Distressed/Restructing Index. The reported means and volatilities are annualized. *SR* is the Sharpe ratio, β is the slope from the regression of the excess strategy return on the excess return of the S&P500 index, and *Drawdown* measures the magnitude of the strategy loss relative to its highest historical value. Means, volatilities and Sharpe ratios are computed using quarterly returns; β estimates are from regressions based on quarterly returns to account for potential asynchronicity due to return smoothing. Drawdown values are based on monthly return series.

Asset	Mean	Volatility	SR	CAPM β	Drawdown
T-bill	3.16%	0.99%	-	-	0.00%
S&P 500	7.53%	18.09%	0.24	1.00	-50.21%
HFRI Composite Index	9.20%	9.14%	0.66	0.42	-21.42%
DJ/CS Broad Index	9.04%	8.73%	0.68	0.33	-19.67%
HFRI Fund-of-funds Index	6.13%	8.01%	0.38	0.30	-22.20%
HFRI Distressed Index	9.78%	9.36%	0.70	0.36	-27.41%

Table II
Risk Properties of Naked Put-writing Strategies

This table summarizes the maximum drawdowns, annualized volatilities and CAPM βs of naked put writing strategies using S&P500 index options. Strategies are defined by the option strike Z-score, Z, and leverage, L, and are rebalanced monthly. Drawdowns are computed using monthly returns after subtracting off a flat management fee (*Implied fee*) that equates the after-fee performance of the strategy and the Hedge Fund Research Inc. Fund Weighted Composite Index. Volatility and CAPM β are computed using quarterly returns.

	Z = -1			
Leverage (L)	Drawdown	Volatility (σ)	CAPM β	Implied fee
1	-10.18%	3.87%	0.18	-0.65%
2	-22.56%	7.72%	0.37	3.60%
4	-44.19%	15.92%	0.77	11.51%
6	-61.85%	24.65%	1.18	18.50%
8	-75.48%	33.94%	1.63	24.37%
10	-85.66%	43.79%	2.09	28.68%

	Z = -2			
Leverage (L)	Drawdown	Volatility (σ)	CAPM β	Implied fee
1	-3.33%	1.74%	0.06	-2.91%
2	-7.96%	3.15%	0.11	-0.76%
4	-17.61%	6.26%	0.23	3.44%
6	-26.97%	9.48%	0.35	7.52%
8	-35.60%	12.77%	0.48	11.46%
10	-43.54%	16.14%	0.60	15.26%

	Z = -3			
Leverage (L)	Drawdown	Volatility (σ)	CAPM β	Implied fee
1	-0.47%	1.21%	0.02	-4.22%
2	-2.35%	1.68%	0.05	-3.36%
4	-6.06%	2.92%	0.10	-1.64%
6	-9.68%	4.27%	0.15	0.05%
8	-13.35%	5.65%	0.20	1.73%
10	-17.23%	7.06%	0.25	3.38%

Table III
Required Excess Rates of Return (1996-2010)

Panel A of this table compares the *ex post* realized excess rates of return for the HFRI Composite Hedge Fund Index, and the naked put-writing strategy, $[-1, 2]$, with *ex ante* required risk premia. The *ex ante* required risk premia are computed using a linear CAPM benchmark with constant beta ($\beta = 0.35$), and under the non-linear model for the marginal investor ($\gamma = 2$) and the endowment investor ($\gamma = 3.3$). The annual excess returns are computed by compounding monthly values. The realized returns to the put writing strategy are reported before fees. The returns for 1996 include February-December; returns for 2010 are through the end of October. Panel B reports the annualized values of the arithmetic mean monthly (excess) returns, and computes investor alphas with respect to the linear CAPM benchmark and the model implied excess return (t-statistics in brackets).

Panel A: Annual excess returns

	Realized (ex post)		**Required (ex ante)**		
Year	HFRI Composite	Put-writing [-1, 2]	CAPM ($\beta = 0.35$)	Model (marginal)	Model (endowment)
1996	11.4%	9.2%	1.2%	2.4%	3.3%
1997	11.5%	14.6%	2.5%	5.1%	7.3%
1998	-4.0%	10.1%	3.4%	7.5%	10.7%
1999	19.8%	21.1%	2.8%	5.7%	8.1%
2000	4.6%	6.5%	2.5%	5.1%	7.3%
2001	0.7%	2.2%	3.1%	6.6%	9.4%
2002	-1.2%	0.0%	3.3%	7.2%	10.3%
2003	15.9%	22.5%	2.5%	5.1%	7.3%
2004	8.1%	14.3%	1.1%	2.1%	2.9%
2005	6.1%	9.1%	0.8%	1.4%	2.0%
2006	8.2%	10.6%	0.7%	1.3%	1.8%
2007	5.9%	9.9%	1.3%	2.6%	3.7%
2008	-20.2%	-11.2%	5.0%	12.5%	18.0%
2009	18.5%	21.3%	5.4%	12.8%	18.7%
2010	5.9%	8.8%	2.8%	5.9%	8.4%

Panel B: Investor alphas

	HFRI Composite	Put writing [-1, 2]
Realized excess return, R*	6.0%	9.7%
CAPM R* ($\beta = 0.35$)	2.5%	2.5%
alpha	3.5%	**7.2%**
	[1.8]	[3.5]
Model R* (equity investor)	5.4%	5.4%
alpha	0.6%	**4.3%**
	[0.3]	[2.1]
Model R* (endowment)	7.6%	7.6%
alpha	-1.6%	2.1%
	[-0.8]	[1.0]

Table IV
Generating "Superstar" Returns by Levering Safe Assets (1996-2010)

This table summarizes the annual returns to the S&P 500 index, the HFRI Composite Hedge Fund Index, and two naked put writing strategies, $[Z, L]$. HFRI returns are reported net of fees; the returns to the put writing strategies are reported before fees. The returns for 1996 include February-December; returns for 2010 are through the end of October. Annual returns are obtained by compounding monthly strategy returns.

		HFRI	**Before fees**	
Year	S&P 500	Composite	[-1, 2]	[-2, 8]
1996	21.4%	16.2%	13.9%	22.2%
1997	28.7%	17.3%	20.6%	21.3%
1998	23.9%	0.7%	15.6%	3.1%
1999	20.9%	25.4%	26.8%	50.4%
2000	-3.9%	10.7%	12.7%	25.5%
2001	-12.2%	5.00%	6.5%	18.9%
2002	-16.5%	0.5%	1.6%	19.7%
2003	15.2%	17.1%	23.8%	40.9%
2004	12.9%	9.3%	15.6%	28.0%
2005	8.7%	9.1%	12.2%	20.8%
2006	14.1%	13.3%	15.9%	24.7%
2007	7.9%	11.0%	15.1%	29.2%
2008	-37.8%	-18.7%	-9.6%	-18.2%
2009	25.5%	18.7%	21.5%	39.6%
2010	6.0%	6.0%	8.8%	22.7%

Figure 1. Asset Class Performance Comparison. This figure plots the total return indices for two hedge fund indices – the Hedge Fund Research Inc. Fund Weighted Composite Index, and the Dow Jones Credit Suisse Broad Hedge Fund Index – the S&P 500 Index, and a strategy that rolls over one-month U.S. Treasury bills over the period from January 1996 to October 2010 ($N = 177$ months).

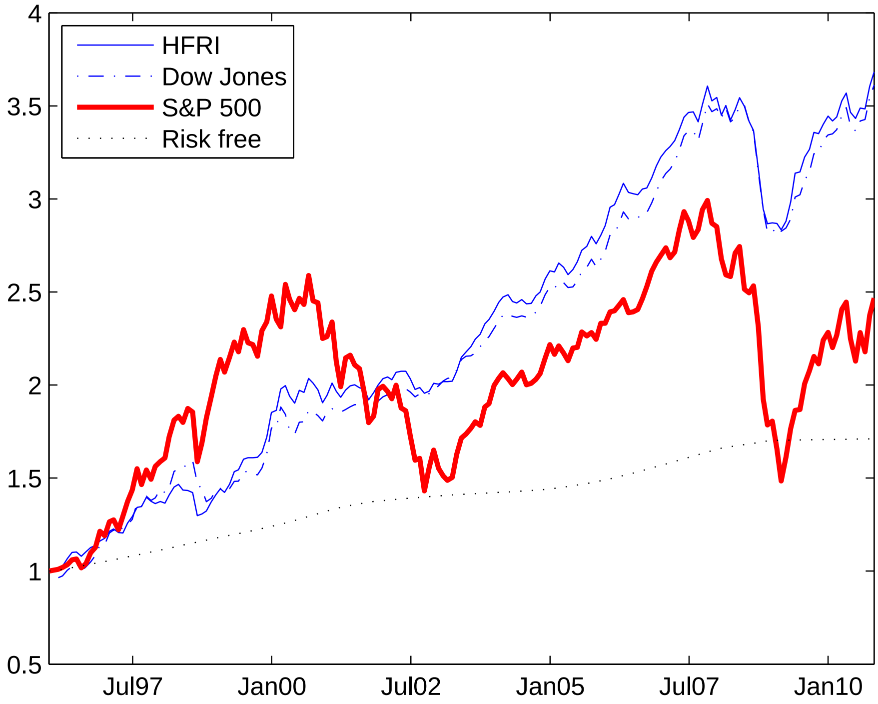

Figure 2. The [Z = -1, L = 2] S&P 500 Put-Writing Strategy and Hedge Fund Returns. The left panel plots the total return indices for the $[Z = -1, L = 2]$ put-writing strategy with a 3.6% annual fee, the Hedge Fund Research Inc. Fund Weighted Composite Index, and the S&P 500 Index, over the period from January 1996 to October 2010. The bottom panels plot the drawdowns – measured as the magnitude of the strategy loss relative to its highest historical value – for the three strategies over the corresponding period.

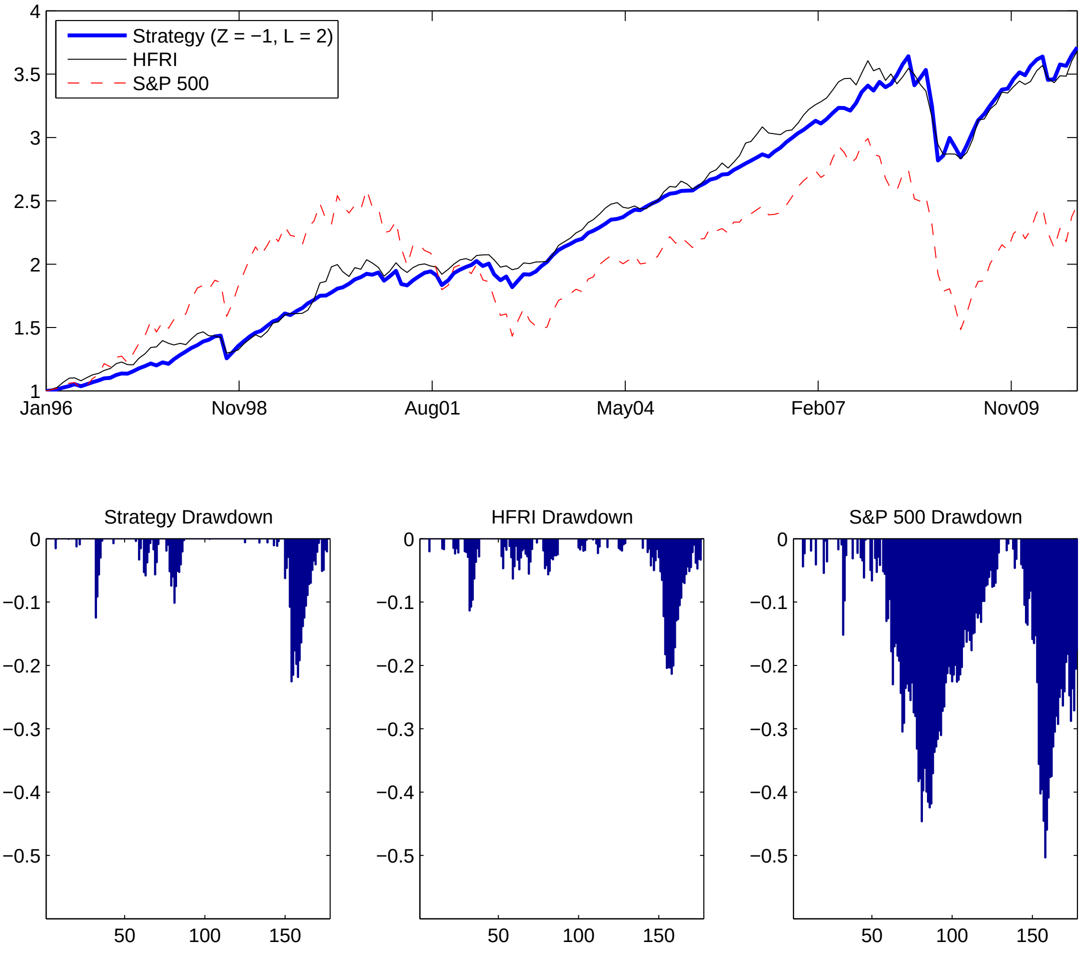

Figure 3. The Cost of Capital for Risky Assets. The following figure displays the investor's cost of capital as a function of his allocation to the the risky asset. The left (right) panel illustrates the case of a shift in the allocation from the risk free asset to equities (alternatives) for the marginal investor ($\gamma = 2$) and an endowment investor ($\gamma = 3.3$). *Proper* is the model-implied cost of capital at the given allocation; *Linear* is the cost-of-capital computed under a linear mean-variance rule.

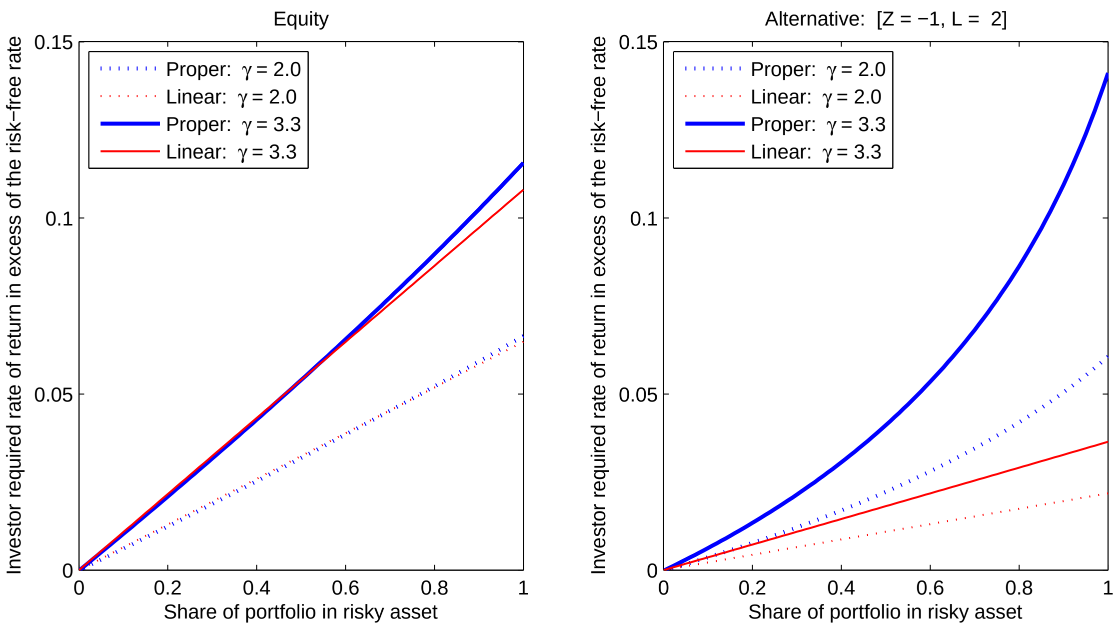

Figure 4. The Effect of Portfolio Composition on the Cost of Capital. The following figure plots the investor's cost of capital for the alternative investment, as a function of its weight in his risky portfolio. The marginal investor is assumed to hold a portfolio comprised entirely of risk assets; the endowment investor holds 20% of his portfolio in the risk-free asset. *Proper* is the model-implied cost of capital at the given allocation; *CAPM* is the cost-of-capital computed under the CAPM β rule.

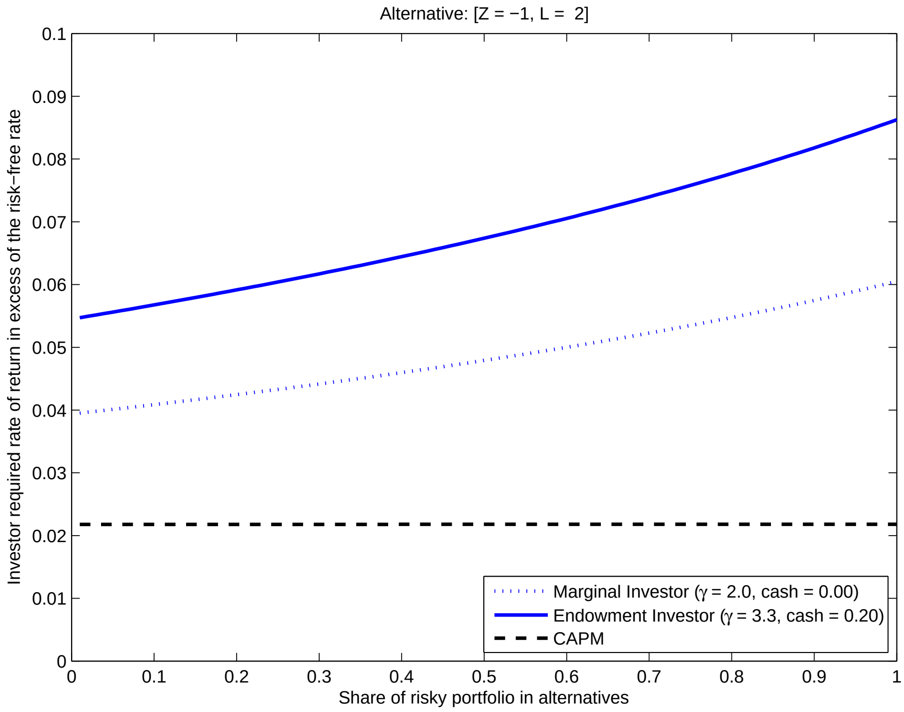

Figure 5. The Effect of Risk on the Cost of Capital. The following figure plots the investor's required rate of return on the alternative as a function of the riskiness of the equity market return distribution, while holding investor allocations fixed. Each investor is assumed to allocate 35% of their risky portfolio to alternatives: the marginal investor holds 0% cash, 65% equities, and 35% alternatives; and the endowment investor holds 20% cash, 52% equities, and 28% alternatives. The top (bottom) panel examines variation in the cost of capital as a function of volatility (mean tail Z-score). We measure tail risk as the average Z-score, conditional on being below the p^{th} percentile of the distribution. We set $p = \frac{1}{60}$, such that the mean tail Z value is interpretable as corresponding to a "once-in-five-year" event.

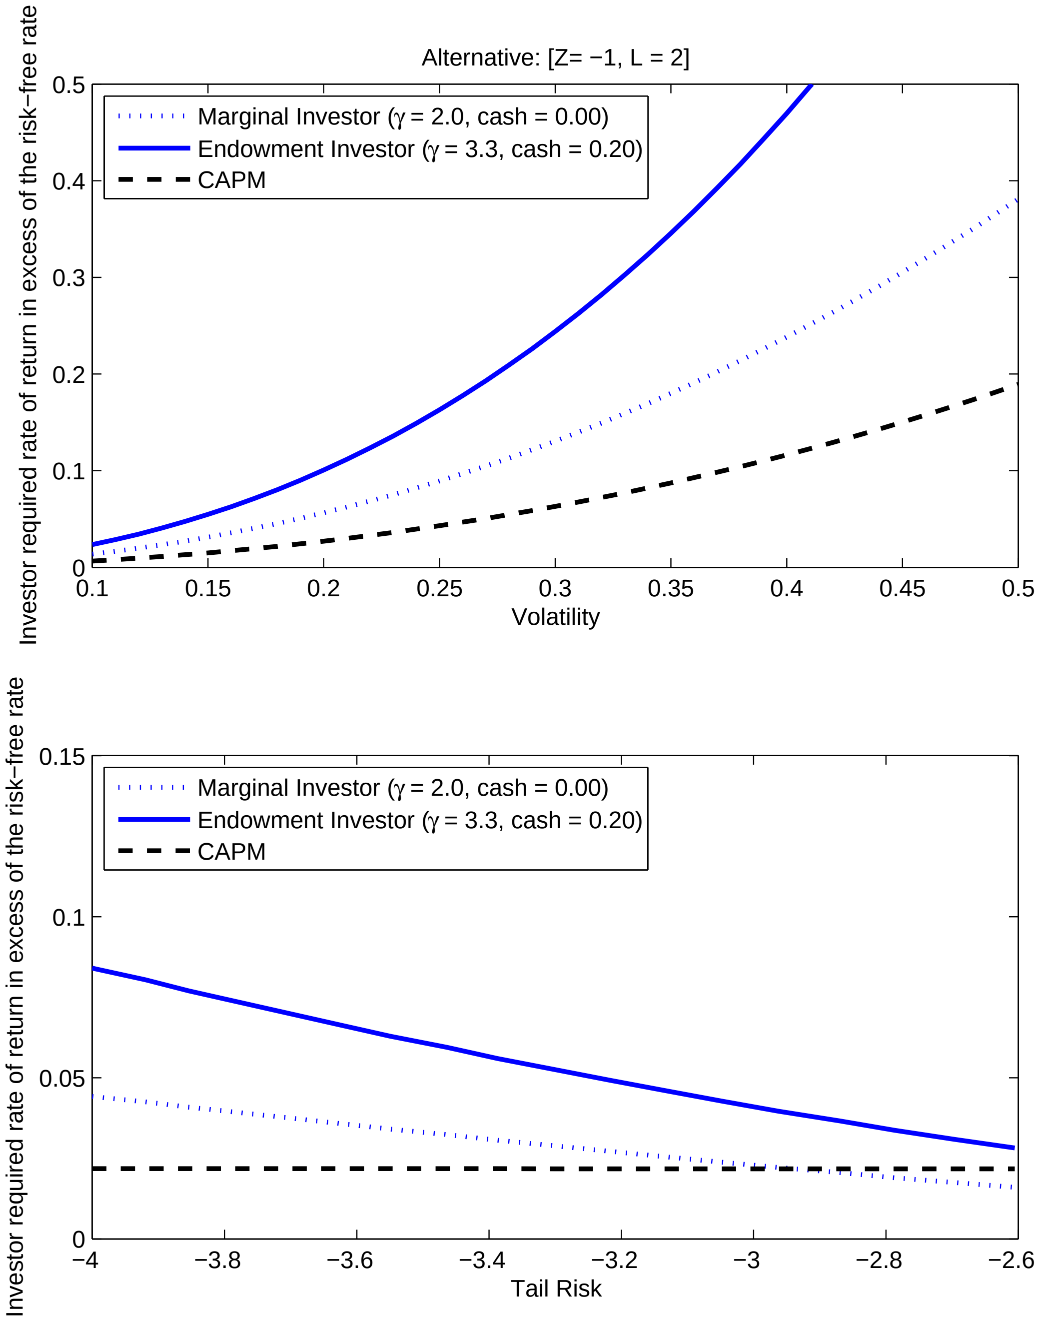

Figure 6. The Effect of Portfolio Leverage on the Cost of Capital. The following figure plots the investor's required rate of return on the alternative as a function of the leverage, L, of the replicating portfolio, while holding investor allocations fixed. As L is varied, the strike price of the options in the replicating portfolio for the alternative are adjusted to keep the instantaneous CAPM β fixed at the value corresponding to the baseline replicating portfolio, $[Z = -1, L = 2]$. Each investor is assumed to allocate 35% of their risky portfolio to alternatives: the marginal investor holds 0% cash, 65% equities, and 35% alternatives; and the endowment investor holds 20% cash, 52% equities, and 28% alternatives. *Proper* is the model-implied cost of capital at the given allocation; *CAPM* is the cost-of-capital computed under the CAPM β rule.

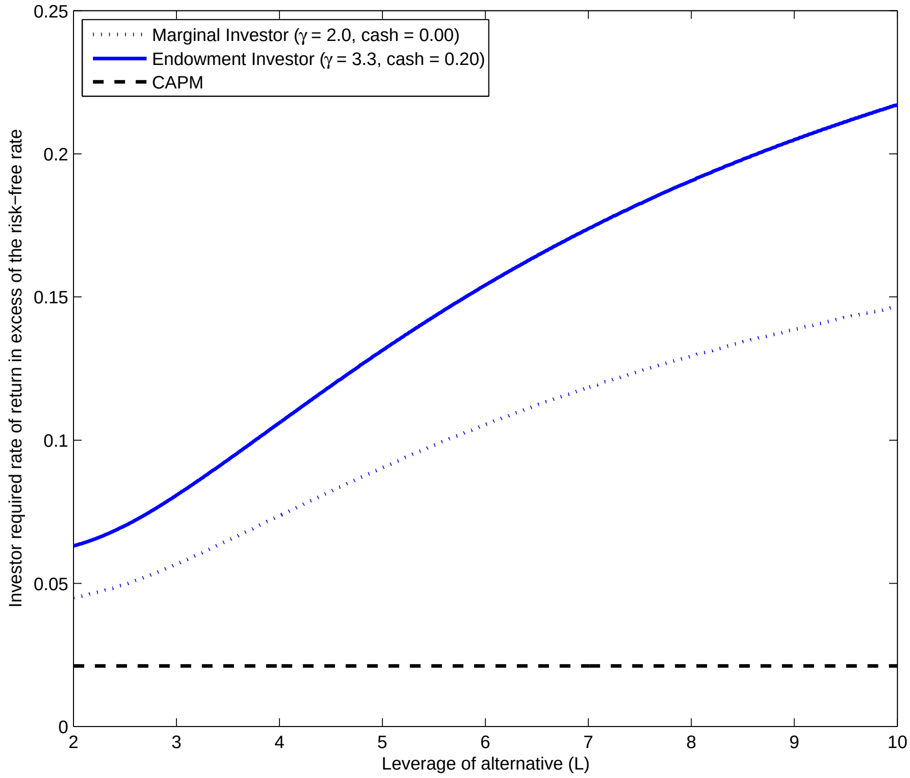

Figure 7. Required Rate of Return for Hedge Funds Through Time. The following figure plots the required rate of return for the hedge fund replicating strategy ($[Z = -1, L = 2]$) through time for the marginal investor ($\gamma = 2$) and the endowment investor ($\gamma = 3.3$). Each investor is assumed to allocate 35% of their risky portfolio to alternatives: the marginal investor holds 0% cash, 65% equities, and 35% alternatives; and the endowment investor holds 20% cash, 52% equities, and 28% alternatives. The equity return volatility, σ, used to compute the required rate of return for month t is set equal to 0.8 of the implied volatility (VIX) index on the last business day of the preceding month. The skewness and kurtosis of the parameterized normal inverse Gaussian (NIG) distribution are held fixed at the baseline values.

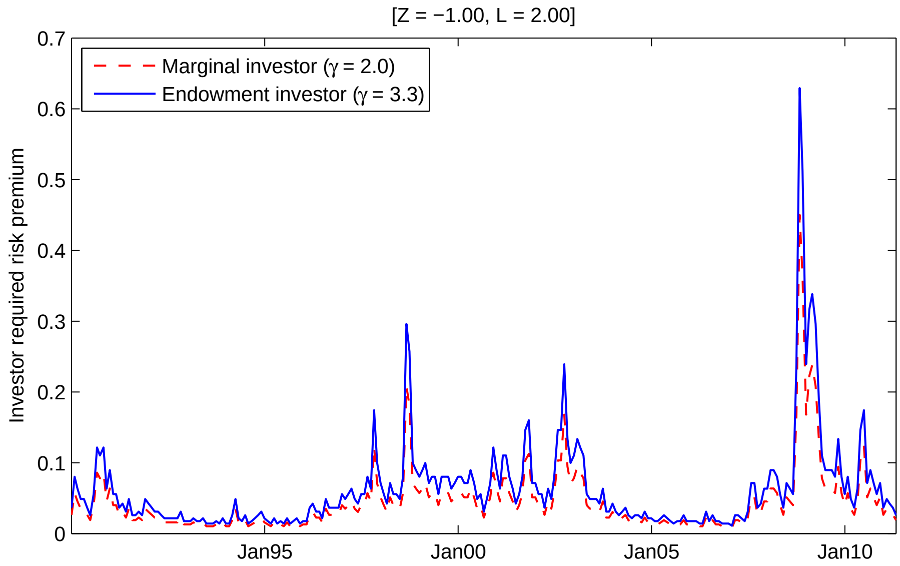

Figure 8. Declining Allocations to Assets with Increasing CAPM α's. This figure illustrate the proper portfolio allocations – as a function of investor risk aversion, γ – to three alternative strategies that group losses in progressively worse economic states by increasing the leverage factors (higher L values), while making the underlying assets safer (more negative Z values). The strategies are designed to have constant CAPM β's equal to 0.33, and offer an *increasing* sequence of CAPM α's of 3% (left panel), 4% (center panel), and 5% (right panel).

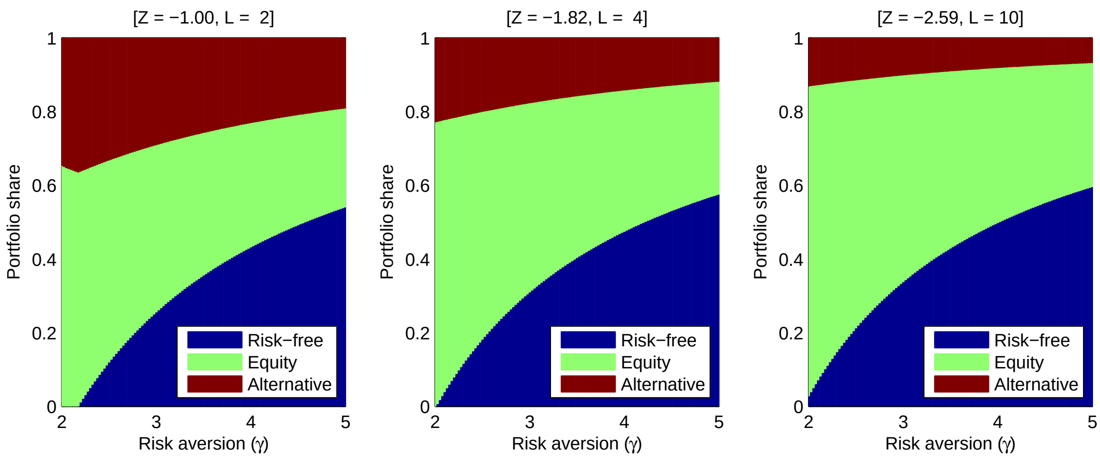

Figure 9. Required Rates of Return with a Mix of Alternatives. The following figure plots the required rate of return for two alternative investments – the hedge fund replicating strategy ($[Z = -1, L = 2]$) and a strategy that applies high leverage to safe assets ($[Z = -2, L = 8]$) – and their equally-weighted mix, as a share of alternatives in the risky portfolio of an endowment ($\gamma = 3.3$) investor. The endowment investor holds 20% of his portfolio in the risk-free asset with the remainder in risky assets (equities and alternatives).