JOURNAL OF INVESTMENT MANAGEMENT, Vol. 1, No. 1, (2003), pp. 135–151

JOIM
www.joim.com

SEGMENTATION, ILLIQUIDITY, AND RETURNS

Renato Staub [a,*] and *Jeffrey Diermeier* [a]

When investing in alternative assets, such as private equity or natural resources—which may be "locked-up" for prolonged periods of time—the question of compensation for illiquidity becomes important. No rational investor will choose the illiquid over the liquid asset unless he gets compensated for his loss of flexibility. We derive two approaches to model illiquidity compensation. In contrast to the ones most commonly seen in the literature, our methods do not analyze trading-based gains which cannot be realized as a result of illiquidity. Rather, we investigate the implications of illiquidity for a long-term investor.

1 Introduction

Pension plan sponsors, endowments, and other institutional investors have included non-traditional assets in their portfolios for many years. Larger, more aggressive funds, in particular those with longer investment horizons, have over time increased their allocations to venture capital, real estate, hedge funds, and other assets that fall outside the realm of regularly priced and traded securities.

At issue is the illiquidity of alternative investments and how it is compensated in the market. With "illiquidity" we do not refer to considerable bid–ask spreads, rather we mean that they cannot be traded at all for a significant amount of time. If, for instance, money is invested in venture capital, it is locked in until it becomes liquid, that is, until it can be traded in the market.

The main problem with illiquidity is that it keeps a portfolio from being continuously rebalanced. As a result, the real portfolio does not match its target allocation in most instances. Figures 1 and 2 documents this in a two-asset world.

In Figures 1 and 2, a mix containing 75% of the low-risk asset and 25% of the high-risk asset proves optimal in Sharpe ratio terms, and a mix containing 89% of the low-risk asset and 11% of the high-risk asset provides the lowest portfolio risk. If both assets can be continuously traded, it is feasible to achieve either one of these two targets. As soon as the portfolio drifts, the overweight is simply sold and the underweight is bought in turn. However, if one asset is illiquid, continuous rebalancing is not possible, and the portfolio is likely to drift away from its target

[a] UBS Global Asset Management, Chicago, IL, USA.
*Corresponding author. UBS Global Asset Management, One North Wacker Drive, Chicago, IL 60606, USA. Tel.: +1 312 525 7924; e-mail: renato.staub@ubs.com.

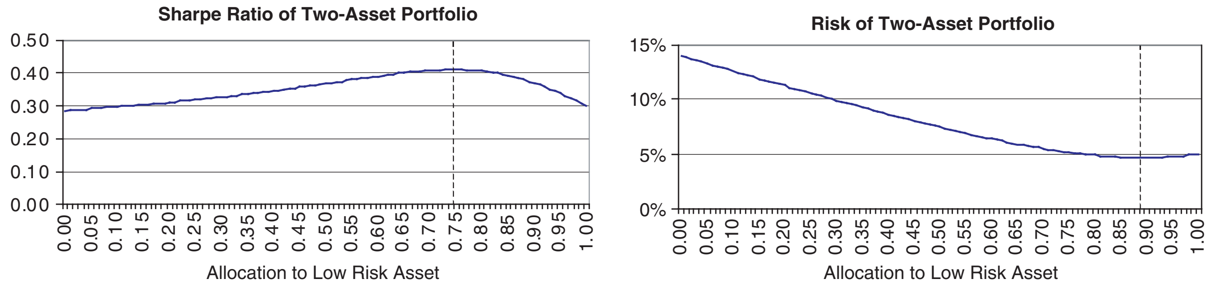

Figure 1 and 2 Sharpe ratio and risk of a two-asset portfolio where one asset has a 14% risk and a 4.5% excess return, and the second asset has a 5% risk and a 1.5% excess return. The correlation between the assets is 0.20.

allocation with the high-risk asset expected to gain in share. That is, there is a constant force increasing the share of the riskier asset. This drift away from the target allocation is the price of illiquidity.

Clearly, if two assets are identical apart from their degree of liquidity, no investor selects the less liquid one unless he gets an additional compensation for giving up flexibility. Here, it is our objective to examine how illiquidity should be compensated in the market. Thereby, it is paramount to examine this question in the context of the entire portfolio, since eventually the *overall* impact matters.

The rebalancing question is widely investigated in the literature. Perold and Sharpe (1988), for instance, conclude that constant rebalancing will

- outperform a comparable buy-and-hold strategy in *flat* markets;
- underperform a comparable buy-and-hold strategy when there are no *reversals*.[1]

Although important, these findings are irrelevant in our context, since we are interested in *policy* setting, that is, we deal with the long-term average composition of portfolios rather than responding to specific market circumstances. Further, when setting a policy, we cannot claim superior information, which is an issue for determining a temporal deviation from policy in order to exploit a given market environment by some strategy.

Unfortunately, there is no generally accepted compensation model for illiquid assets. There are some "ad hoc" approaches, but these suffer from drawbacks such as:

- operating with biased data;
- concentrating on single selected aspects;
- providing stand-alone investigations;
- being trading based (in 90% of all instances);
- making unrealistic assumptions.

We want to estimate equilibrium returns for illiquid assets as they can be expected in the long term, and we estimate them in a partially integrated world. The paper is designed as follows: First, we summarize some of the most relevant articles investigating illiquidity compensation. Second, we develop two new approaches. Third, we estimate the returns as implied by these approaches. Thereby, we do not assume superior information, since the market as a whole provides no active return. Fourth, we integrate these approaches into the existing integration/segmentation approach developed for liquid markets. Finally, we discuss the results.

2 Review of previous work

There are several articles dealing with illiquidity. Let us mention these which are referenced regularly or make strong statements.

Amihud and Mendelson (1986a) investigate the relationship between bid–ask spreads and gross returns of stocks, based on NYSE stock returns over the 1961–1980 period. Interestingly, they find that the statistical relationship between the bid–ask spreads and the gross returns was much stronger than between the betas and the gross returns. As a proxy for the betas, they use rolling historic estimates. Further, they claim that the monthly excess return of a stock with a 1.5% spread was 0.45% greater than that of a stock with a 0.5% spread, and hence, they conclude that the bid–ask spread is a very important return parameter.

Amihud and Mendelson (1986b), studying similar issues, find that: (i) market-observed average returns increase with their spread; (ii) asset returns net of trading costs increase with the spread; (iii) stocks with higher spreads are held by investors with longer holding periods; and (iv) returns of higher-spread stocks are less spread-sensitive.

Silber (1991) examines the price differences between two securities which are identical except for the fact that one is subject to SEC rule 144, which claims that holders of the so-called "restricted stock" are forced to hold the according stocks for at least two years, since they are not SEC registered. The incentive of public companies for placing private equity is *time*: getting an SEC registration takes time, and in many instances they cannot afford to wait.[2] Silber investigates some of the scarcely available data about private placements of common stock. He finds an average of 33% discount for a two-year illiquidity period. The problem with this investigation is the fact that many firms issuing SEC 144 stocks are in financial distress, which is why they have no time to issue regular stocks. Therefore, the large discount might not be for illiquidity per se, but information asymmetry.

Chaffe (1993) suggests that the value of a put with an exercise price at the asset's marketable price at the time of purchase is a fair compensation for non-marketability and thus should equal the discount: "When provided with an option to sell, otherwise non-marketable shares are given marketability. [. . .] Following this logic, the cost or price of the option represents all (or a major portion) of the discount to be taken from the marketable price to price the non-marketable shares." Unfortunately, his remark in parentheses weakens an otherwise strong statement. Based on put values derived by Black–Scholes, he concludes that discounts between 28 and 41% are reasonable estimates. As we shall demonstrate, the approach is fraught with some problems.

Longstaff (1995) asks how big the expected missed opportunity would be for an investor with perfect timing but with no opportunity to apply it, since he is locked in. He assumes an asset that follows a Brownian motion and derives, based on risk-neutral dynamics, a closed formula for the expected difference between market value at liquidation and the highest value between acquisition and liquidation.[3] Since there is no better information than perfect information, the according value is considered the greatest possible discount. The resulting discounts may be huge. Comparing his results with Silber's (1991) empirical research, he claims that the market prices are close to this upper bound. However, since we have reservations with regard to Silber's research and since Longstaff assumes superior information, there is not too much we can get out of this approach.

Longstaff (1999) assumes a two-asset market in which trades take place continuously. One asset is risk-free and the other follows a Brownian motion of variable, *instantaneous* volatility. He finds that an investor who maximizes his terminal wealth faces an optimal allocation which changes continuously. Further, he defines illiquidity by a trading cap per unit of time. But in the case of constrained trading, the investor does not have complete control over his allocation which should adapt permanently to the optimal allocation. This turns into a welfare loss. Longstaff's approach is an example of what a long-term investor is not looking for, for several reasons. First, the approach is mechanical in nature, and the world is considered a parameter-machine. While this might be helpful in educational terms, its practical use is limited. Second, the solution becomes so complicated that the author cannot provide a closed-form solution, not even for the two-asset world, and it is questionable whether it is feasible at all. In a multi-asset world, the problems encountered by this approach become insurmountable. Third, the objective function is driven by terminal wealth only, and the path to it, whether smooth or risky, does not matter; that is, there is no concept of risk aversion which is inconsistent with Modern Portfolio Theory. Fourth, the approach examines the asset on a stand-alone basis. Fifth, as a trading-based approach it should take into account the transaction costs which are essential in the case of a high trading frequency. And sixth, the approach is trading-based which does not fit into a long-term context.

Smith *et al.* (2000) document the SEC fair value standard recommending substantial discounts for non-marketable stocks. They observe that restricted-share transactions often take place when an issuer is under liquidity pressure and conclude that interpreting this as evidence for discounts is not valid. In addition, they claim that economic theory provides no basis for large return premiums for bearing illiquidity as provided by Silber (1991). A variety of institutional investors can invest in restricted shares without materially affecting their liquidity. Hence, they think that a fair discount for illiquidity is not very large.

For the reasons mentioned in Section 1, most of these examinations come up with illiquidity premia which seem unrealistically high. Our own experience with employee bonus deferral programs through restricted stocks imply significantly lower illiquidity premia.[4]

3 Existing approaches

The required returns for liquid assets can be modeled with the existing integration/segmentation approach. With regard to compensating illiquid assets, we intend to derive according approaches and reconcile them thereafter with the integration/segmentation approach. That is, we want to create a consistent set of approaches to be applied to the whole world of assets, regardless of whether liquid or illiquid. In this section, we describe briefly the integration/segmentation approach and then our approaches for modeling illiquidity. Finally, we combine the two.

3.1 Basis for risk premia: integration/segmentation approach[5]

A cornerstone in finance and generally accepted as an equilibrium model, the Capital Asset Pricing Model (CAPM) is usually the approach of choice for estimating risk premia. It claims that every asset is rewarded by a risk premium proportional to its beta versus the entire market plus a risk-free return, that is,

$$R_i = \beta_{i,M}(R_M - R_f) + R_f \tag{1}$$

Defining the entire market is an issue, however. Brinson *et al.* (1986) present criteria to reduce the list of potential asset classes that make up the "investable capital market".[6] According to their paper, "the investable capital market consists of primary wealth-generating assets where sufficient markets have developed and legal hurdles do not prohibit meaningful investment by tax-exempt investors".[7] Based on their criteria, our proxy for the entire market is our capital-weighted Global Investable Market (GIM) portfolio. In contrast to a common US-biased global balanced portfolio, it has less exposure to US assets (particularly US equity). On the other hand, GIM contains alternative assets, which are mainly but not exclusively represented by US real estate. Finally, the pricing of risk is based on the research provided by Singer and Terhaar (1997) who estimate the Sharpe ratio of GIM.[8]

CAPM, however, assumes perfect markets and is thus not an appropriate representation of the real world. While the US equity market approaches a state of high integration, it does not meet it perfectly. Barriers to international capital flows have come down, but they still exist to some extent, and many asset markets are still significantly segmented by national borders. These barriers are not necessarily erected by law; rather they may reflect investor preference.[9] Such restrictions on capital in- and outflows tend to create markets which are dominated by local investors, and, in the most extreme case, i.e. if a market is completely segmented, its own risk becomes its compensation reference, because there are no substitution opportunities. This means that the market's absolute rather than systematic risk is compensated[10]

$$R_i = \frac{\sigma_i}{\sigma_M}(R_M - R_f) + R_f \tag{2}$$

Of course, complete segmentation is as fictitious as perfect integration.

Therefore, in practice, it is important to estimate a market's degree of integration, and the according return is considered a weighted average of perfect integration and absolute segmentation. That is

$$R_{i,\text{weighted}} = w_{\text{int}} R_{i,\text{int}} + w_{\text{seg}} R_{i,\text{seg}} \tag{3}$$

where the two weights sum up to 100%.

3.2 Chaffe's put option approach

Chaffe thinks a fair discount for illiquidity is the value of a put option, the exercise price of which equals the asset's marketable price at the time of its purchase. In an equilibrium context, the marketable price corresponds to the asset's intrinsic value at the time of its purchase. He justifies the approach as follows: "When provided with an option to sell, otherwise non-marketable shares are given marketability. [. . .] Following this logic, the cost or price of the option represents all (or a major portion) of the discount to be taken from the marketable price to price the non-marketable shares".[11] Although the approach is based on option values, the investor is not intended to get the option; rather, he gets a discount of the same value theoretically enabling him to buy the option.[12]

Thus, in Chaffe's world, the investor has an asset/option portfolio (AOP) instead of the asset only. Its value at liquidation is

$$\max\,(S(T), S) \tag{4}$$

Figure 3 shows AOP's density function at liquidation, given that the asset's normally distributed continuous returns are reinvested.

All possible payoffs of AOP are on the vertical line marking the asset's value at times of purchase, S, or east of it. This is full shortfall insurance combined with unlimited upward potential. In other words, AOP is risk-free in shortfall terms. Figures 4 and 5 show the resulting discounts for selected risk levels

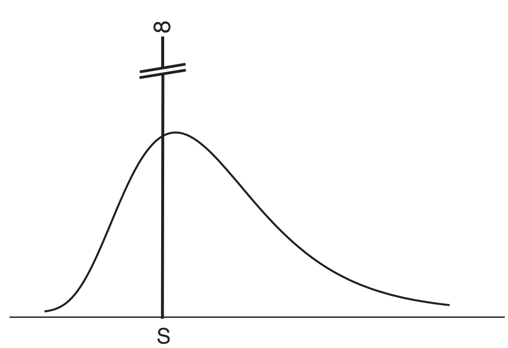

Figure 3 Payoff structure of Chaffe's Put Option approach.

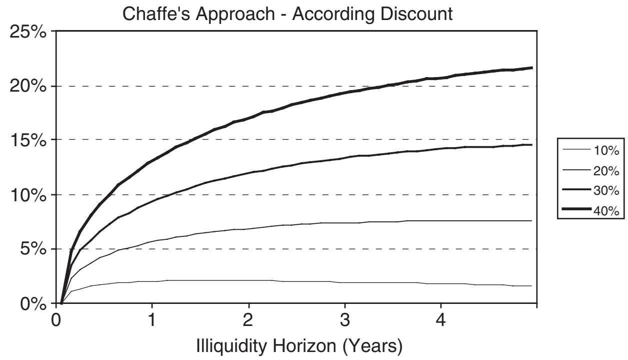

Figure 4 Discounts for four selected risk levels (10, 20, 30 and 40%). The x-axis marks the lock-in time and the y-axis the according discount. (Our equilibrium estimate for the continuous risk-free return is 4.597.)

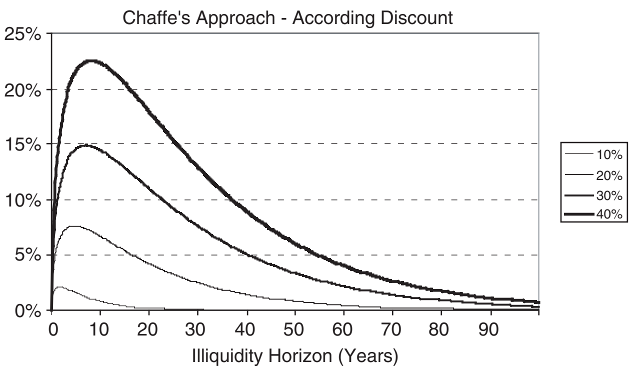

Figure 5 The same discount function given in Figure 4, but over a much longer horizon.

and illiquidity horizons as provided by the Black–Scholes formula for valuing European put options.

Apparently, the discount increases first with a rising illiquidity horizon, but after reaching a maximum it decreases and approaches zero asymptotically. The reason for this effect is straightforward: as the distribution effect of the asset's value only grows with the square root of time, it will be dominated by the return effect beyond a certain horizon, and hence the shortfall probability will decrease.

In the short run, free shortfall insurance combined with unlimited upward potential implies no real risk-return trade-off and hence looks too generous.[13] Therefore, Chaffe's additional remark in parentheses that the discount may only be "a major portion" of the put option value, weakens his overall statement and leaves the reader with the question of how big this portion should be. Further, for longer horizons, the implied discount and hence the illiquidity premium decreases with an increasing illiquidity horizon, whereas practical evidence implies the contrary. While longer horizons are probably less important for Chaffe, as he is tilted towards SEC 144 stocks which are banned from trading for a mere two years,[14] it is certainly relevant to us, since we are interested in assets with illiquidity horizons of up to 10 years.

In the end, although full shortfall insurance initially looks like an objective criterion, we think it is not. From a statistical perspective, it is fairly arbitrary, because it has no inherent relationship with both the asset's risk and the length of the lock-in. While we do not think that this is the right way to approach the problem, we consider the idea of a put option as one that is viable and can be improved.

3.3 Improved put option approach (POA)

In order to provide an improved POA, we consider more closely an asset's risk-return properties after some discrete amount of time. To that end, assume an asset providing normally distributed continuous returns that are reinvested

$$\tilde{\mu} \sim N(\mu, \sigma) \tag{5}$$

According to the theory of log-normal distribution, its expected value at time T is[15]

$$E[S(T)] = Se^{\mu^* T} \tag{6}$$

with

$$\mu^* = \left(\mu + \frac{\sigma^2}{2}\right) > \mu \tag{7}$$

and its median value, M, is

$$M[S(T)] = Se^{\mu T} \tag{8}$$

where μ is the *expected instantaneous* return and μ^* the *expected risk-corrected instantaneous* return. Note that μ^* is *instantaneous* as well, but it is corrected versus μ in such a way that it provides the asset's *expected value* at the end of the period when compounded. The difference between μ and μ^* happens to equal $\sigma^2/2$, and according to Eqs. (6) and (8), a risky asset's expected value after a discrete period is larger than its median value as reflected by compounding its continuous rate of return only. The longer the holding period, the greater the difference between the expected value and the median value. The reason for this is the asymmetric effect of risk through compounding (Figure 6).[16]

That is, if you buy and hold, you expect to have more by the end of the holding period than the median value, namely, the expected value. The difference between the two is shaded in Figure 6, and we call it the *holding* effect, implying that you need to hold the asset for some time until this effect materializes. On the other hand, if you buy and hold the asset, you still keep the option to resell it at any time, that is, the concept of the holding effect applies to any "normal" liquid asset. If illiquid assets are involved, illiquidity not only means that you buy and hold the asset, but that you buy and *commit* to hold it for some time.

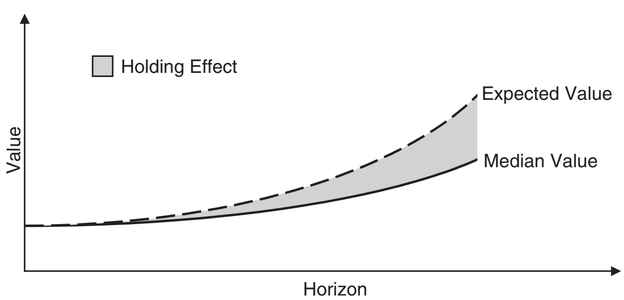

Figure 6 Expected value and median value of risky asset.

But what, ultimately, is the incentive to hold illiquid assets for a predetermined period, if the holding effect materializes for liquid and illiquid assets in a similar fashion? Like Chaffe, we do think that free insurance may be an incentive, but we think a fair incentive for committing to hold is free insurance for this portion to be gained through *holding*. This insurance of the holding effect is measured by the *difference* of two put options, p_1 and p_2,

$$\Delta = p_1 - p_2 \tag{9}$$

to be exercised at liquidation. The option p_1 has an exercise price equal to the asset's expected value

$$X_1 = E[S(T)] = Se^{\mu^* T} \tag{10}$$

And the option p_2 has an exercise price equal to its median value

$$X_2 = M[S(T)] = Se^{\mu T} \tag{11}$$

As in Chaffe's approach, the investor has an AOP, but this time, AOP has a short position as well; it is short p_2. At liquidation, it has the following value:

$$\max\,(S(T), E[S(T)],\ S(T)+E[S(T)]-M[S(T)]) \tag{12}$$

The resulting put option values are[17]

$$p_1 = Se^{(\mu^*-r)T} N((\mu - r + \sigma^2)\sqrt{T}/\sigma) - SN((\mu - r)\sqrt{T}/\sigma) \tag{13}$$

and

$$p_2 = Se^{(\mu-r)T} N((\mu - r + \sigma^2/2)\sqrt{T}/\sigma) - SN((\mu - r - \sigma^2/2)\sqrt{T}/\sigma) \tag{14}$$

Finally, the asset's resulting return, including compensation for a T-year lock-up, $\mu(T)$, is

$$\mu(T) = \log\left(\frac{e^{\mu^* T}}{1 - (p_1 - p_2)/S}\right) T^{-1} \tag{15}$$

A less formal and more intuitive explanation for this approach is that many investors are concerned about the fact that, over time, a risky asset's median performance is smaller than its expected performance, as the expected performance is driven by skewness.[18] Thus, in every single instance, the expected value is more likely to be underachieved than overachieved. Illiquidity in this context means the investor's agreement to exposure versus the potentially significant difference between the expected value and the median value with no chance to take any action until the asset becomes liquid. Hence, free insurance of the *expected performance* versus *the median performance* appears to be a fair compensation for *committing to stay* exposed to an asymmetric return pattern, and thus a fair compensation for accepting illiquidity. Figure 7 shows the density function for (12) at liquidation, given the asset's normally distributed continuous returns are reinvested.

Again, the approach provides some shortfall insurance, but the according insurance level, the asset's expected value at liquidation, obviously relates to the asset's properties and its illiquidity horizon. Further, there is no full insurance, since there is a distribution to the left of the insurance level. In the case of full insurance versus the asset's expected value, there would be no incentive at all to invest in liquid assets. But, as opposed to the asset only, the overall distribution is an improvement, as the left portion is shifted to the right. The smallest value possible is $E(T) - M(T)$.[19] Finally, the derived

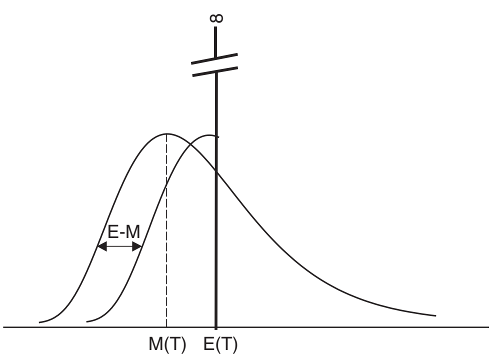

Figure 7 Payoff structure for POA.

discounts are not counterintuitive at the long end, as they consequently increase with an increasing illiquidity horizon. Overall, POA requires the following inputs:

- risk-free return (r);
- risk of the illiquid asset (σ);
- illiquidity horizon (T)[20];
- return of the entire market (μ_m^*);
- risk of the entire market (σ_m).

On the other hand, if μ is considered independent from the entire market, that is, in the case of absolute segmentation, the set of required inputs is reduced to r, σ, and T. Figure 8 shows the resulting discounts for selected risk levels.[21]

Apparently, the larger the asset's

- risk, the larger its discount;
- illiquidity horizon, the larger its discount and the steeper its discount function.

However, for certain combinations of risk and illiquidity horizon, the required discount exceeds 100%, which means that nobody will be willing to invest whatsoever.

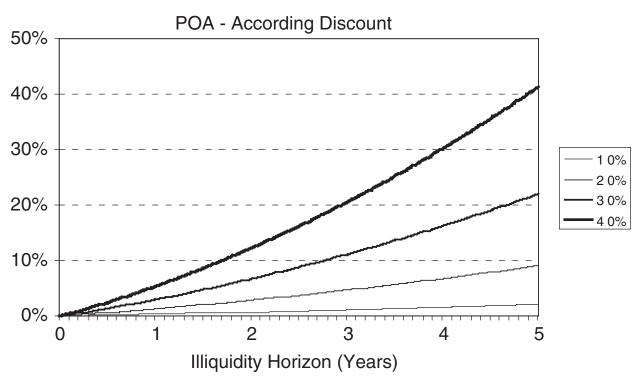

Figure 8 Discounts for four selected levels of compensated risk (10, 20, 30 and 40%). The x-axis marks the illiquidity horizon and the y-axis the according discount.

3.4 *Sharpe ratio approach (SRA)*

This approach is based on the key observation that the Sharpe ratio is a function of the holding period. Therefore, it is pointless to examine an asset's Sharpe ratio[22] for a holding period shorter than the illiquidity horizon. Rather, the expected wealth at liquidation and its distribution matter.

Assume a risky asset has a value of S at time 0, and its whole return is reinvested. Then, its expected value at time T is[23]

$$E[S(T)] = Se^{\mu^* T} \tag{6}$$

with

$$\mu^* = \left(\mu + \frac{\sigma^2}{2}\right) > \mu \tag{7}$$

Hence, the expected excess value of the risky asset over the value of the risk-free asset is

$$E[S(T) - S(T,r)] = Se^{\mu^* T} - Se^{rT} \tag{16}$$

where r is the risk-free return. Further, the expected variance of the expected excess value is[24]

$$\sigma^2[S(T) - S(T,r)] = \sigma^2[S(T)] = S^2 e^{2\mu^* T}(e^{\sigma^2 T} - 1) \tag{17}$$

Based on (16) and (17), we can derive the risky asset's Sharpe ratio, SR(T), over the entire holding period T. It equals the expected excess value at time T divided by its expected distribution; that is,[25]

$$\mathrm{SR}(T) = \frac{E[S(T) - S(T,r)]}{\sigma[S(T) - S(T,r)]} = \frac{e^{\mu^* T} - e^{rT}}{e^{\mu^* T}\sqrt{e^{\sigma^2 T} - 1}} \tag{18}$$

Figures 9 and 10 shows the Sharpe ratio function for selected parameters.

Apparently

- both assets have an optimal horizon in terms of SR(T);

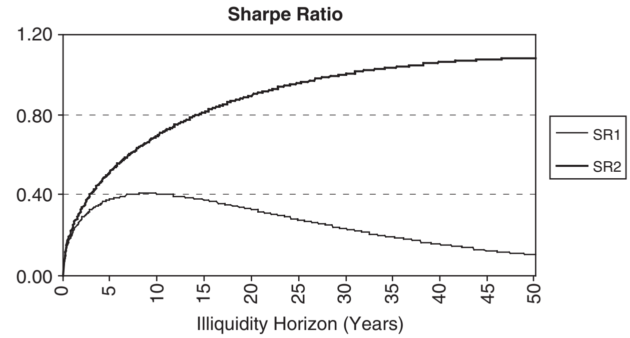

Figure 9 Sharpe ratio for an asset with 7% risk and 6.5% annual return (upper curve), and Sharpe ratio for an asset with 30% risk and 12% annual return (lower curve).

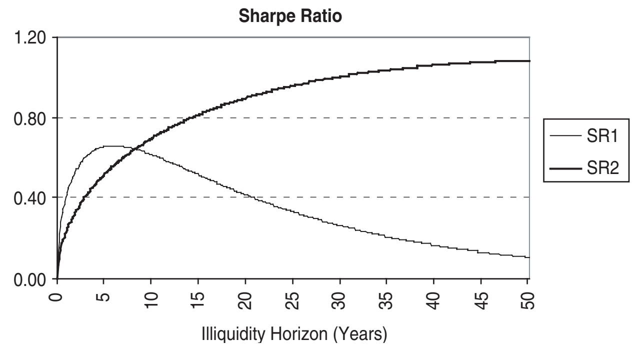

Figure 10 Sharpe ratio for an asset with 7% risk and 6.5% annual return (upper curve), and Sharpe ratio for an asset with 30% risk and 20% annual return (lower curve).

- the optimal horizon of the low-risk/low-return asset is further out;
- the optimal SR(T) of the low-risk/low-return asset class is larger.

Now, the key question is: what is the motivation to lock in a risky and non-marketable asset if its Sharpe ratio is smaller than *the entire market's* Sharpe ratio for the relevant horizon? Evidently none, as the investor requires a Sharpe ratio at least as big as the entire market's Sharpe ratio.

While an asset's risk is mainly *given* by risk factor exposure, its return can be influenced through the discount at which it is acquired. The larger the discount, the higher the expected return and, *ceteris paribus*, the more SR(T) in Figure 10 is shifted upward. As a consequence, the two curves' intersection moves up and to the right. In the given case, the illiquidity horizon is about 8 years, and the according Sharpe ratio for the illiquid asset and the entire market as proxied by GIM is about 0.67.[26]

In contrast to POA, SRA does not need to derive a discount rate in order to calculate the return rate of the examined asset. Rather, as the approach is built around the risk and return of a reference portfolio, the examined asset's return rate results directly. If the

- risk-free return (r)
- risk of the illiquid asset (σ)
- illiquidity horizon (T)
- return of the entire market (μ_m^*)
- risk of the entire market (σ_m)

are given, we can equal the Sharpe ratio for an alternative asset to the Sharpe ratio of the entire market and finally resolve for the asset's return which includes illiquidity compensation for a T-year lock-up, $\mu(T)$; it is[27]

$$\mu(T) = \log\left[\left(\left(1 - \sqrt{\frac{e^{\sigma^2 T} - 1}{e^{\sigma_m^2 T} - 1}}\right) \cdot \left(1 - \frac{1}{e^{(\mu_m^* - r)T}}\right)\right)^{-1/T} e^r\right] \quad (19)$$

Figure 11 shows the returns for selected risk levels and illiquidity horizons.

Obviously, the larger the asset's

- risk, the larger its return and the steeper the return function
- illiquidity horizon, the larger its return and the steeper the return function.

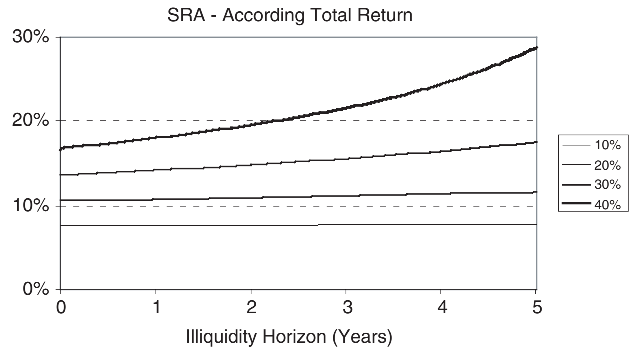

Figure 11 Implied annual returns for the four risk levels of 10 (lower curve), 20, 30, and 40% (upper curve) and illiquidity horizons between 0 and 5 years, given an entire market with 7% risk and 6.5% annual return.

Again, for certain combinations of risk and illiquidity horizon, the required return becomes infinite, which is the point where nobody is willing to invest. Practically speaking, this is the case when the spiky curve in Figure 10 becomes higher due to an increased asset return but not wider.[28]

3.5 *Comparison between the POA and the SRA*

As long as markets are considered integrated, both approaches require the same set of inputs: risk-free return, risk and return of the entire market, and risk and illiquidity horizon of the according asset. Further, both approaches imply an infinite return for a certain size of risk and illiquidity horizon, which means that nobody is willing to invest.

The key difference between the two approaches is the fact that SRA is always a *relative* approach, while POA may be an *absolute* approach. Namely, if μ of asset A is set independently from the entire market,[29] POA is absolute in that A's resulting illiquidity premium is exclusively a function of A's properties, regardless of how other assets or reference portfolios are characterized. On the other hand, SRA is always relative in that the illiquidity premium for asset A can only be determined if the properties of the reference market are known.

The most apparent difference between POA and SRA is SRA's ability to provide negative illiquidity premia. Although not too likely, this is the case if an asset's risk-free return plus risk premium provides Sharpe ratios for all horizons that are larger than the corresponding Sharpe ratios of the reference portfolio. Typical candidates would be low-risk assets with a relatively high degree of segmentation. Such an instance would require a negative illiquidity premium in order to satisfy condition (A.12). In contrast, a negative illiquidity premium can never result for POA, as p_1 is always larger than p_2 because of its higher strike price.

4 Combination of the illiquidity approach and the integration/segmentation approach

As mentioned in Section 3, we applied our illiquidity approaches so far to the *total risks* of the according markets which implies the full segmentation case, as no attention was paid to the *market context*. The opposite end of the spectrum is defined by absolute *integration* in which case only the systematic risk is compensated. In order to allow for perfect integration, we resolve Eqs. (15) and (19) with the systematic risk. Consequently, instead of talking of a holding period *Sharpe* ratio, we speak of a holding period *Treynor* ratio. Finally, as for liquid assets, we estimate a weighted mean reflecting the according degree of integration.

The premia for perfect integration are usually smaller than for absolute segmentation but never larger, since the systematic portion of an asset's risk can be no larger than its entire risk (Figure 12).

While the premia increase overproportionally with the illiquidity horizon, they increase linearly with

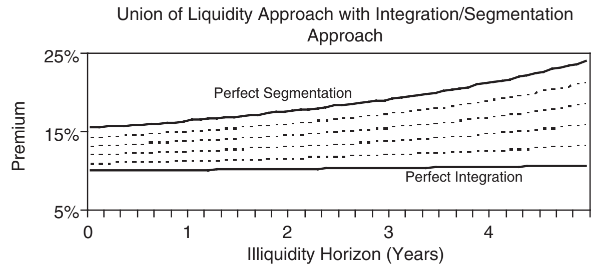

Figure 12 Liquidity combined with integration/segmentation. The upper bold line shows the total premia in case of full segmentation, and the lower bold line shows the premia in case of full integration, both as a function of the illiquidity horizon. The dashed lines represent various degrees of weighted means between integration and segmentation.

the degree of integration. By "illiquidity horizon" we mean the time span until the asset *can* be resold. This is not necessarily the time span after which the asset *will* be resold. This is an issue for real estate, in particular, which is usually held for a period considerably longer than our theoretical illiquidity horizon of 12 months.[30] Obviously, there is an incentive, such as a large bid–ask spread and/or a large risk in the bid/ask spread risk, to hold real estate beyond its illiquidity horizon. Further, this may be an issue for hedge funds as well, where the investor often pays a penalty if he redeems assets within three years.

Table 1 shows our inputs and the resulting returns. While the total risks and systematic risks are determined by the underlying covariance structure and the market capitalization weights, the determination of the illiquidity horizons is definitely more challenging. Although venture capital may require as long as 10 years to mature, the illiquidity horizon in a duration-sense is much shorter, since not all capital is invested at the very beginning and the investment starts to generate cash flows after a few years. In contrast to venture capital, private real estate seems to have a much shorter illiquidity horizon; technically, no longer than a year. In practice, however, the chance to resell real estate at

Table 1 Returns as implied by the Sharpe ratio approach (SRA) and the put option approach. Columns 7–12 are continuous returns while columns 13 and 14 are annualized returns as usually communicated in practice.

Asset	Illiquidity horizon	Beta vs. GIM	Risk (%)	Syst. risk (%)	Integration weight (%)	Segmentation weight (%)	Total return full integration		Total return full segmentation		Total return weighted mean		Total return weighted mean	
	1	2	3	4	5	6	7	8	9	10	11	12	13	14
Venture Early	4.00	3.77	44.9	25.0	50	50	12.8	13.5	26.6	27.4	19.7	20.5	21.8	22.7
Venture Late	2.50	3.17	35.2	21.0	50	50	10.8	11.4	16.3	18.3	13.6	14.9	14.6	16.0
LBO	4.00	3.76	36.0	24.9	50	50	12.8	13.5	18.8	20.1	15.8	16.8	17.1	18.3
Mezzanine	4.00	1.46	17.8	9.7	50	50	7.3	7.3	10.0	10.3	8.6	8.8	9.0	9.2
Distressed Debt	3.00	1.73	20.2	11.5	50	50	7.8	7.9	10.6	11.1	9.2	9.5	9.7	10.0
Distressed Securities	0.25	0.61	14.0	4.1	60	40	5.7	5.7	8.4	8.6	6.8	6.8	7.0	7.1
Event-Driven	0.25	0.76	16.0	5.0	60	40	6.0	5.9	9.0	9.2	7.2	7.3	7.5	7.5
Fund of Funds	0.25	0.44	11.0	2.9	70	30	5.4	5.3	7.6	7.6	6.1	6.0	6.2	6.2
Emerging	0.25	1.08	25.1	7.2	50	50	6.6	6.5	11.5	12.5	9.1	9.5	9.5	10.0
Growth	0.25	2.00	39.9	13.3	70	30	8.2	8.3	15.8	18.7	10.5	11.5	11.1	12.1
Value	0.25	1.58	22.2	10.4	70	30	7.5	7.5	10.7	11.4	8.4	8.7	8.8	9.0
Macro	0.25	0.00	20.0	0.0	70	30	4.6	4.6	10.1	10.6	6.2	6.4	6.4	6.6
Market-Neutral	0.25	0.16	11.1	1.1	75	25	4.9	4.9	7.6	7.7	5.6	5.6	5.7	5.7
Risk-Arbitrage	0.25	0.16	11.1	1.1	75	25	4.9	4.9	7.6	7.7	5.6	5.6	5.7	5.7
Convertible Arbitrage	0.25	0.16	11.1	1.1	60	40	4.9	4.9	7.6	7.7	6.0	6.0	6.2	6.2
Fixed-Income Arbitrage	0.25	−0.08	11.0	−0.5	60	40	4.7	4.7	7.6	7.6	5.9	5.9	6.1	6.1
Income	0.25	0.31	11.0	2.1	70	30	5.2	5.1	7.6	7.6	5.9	5.9	6.1	6.0
Sector Technology	0.25	1.88	35.7	12.4	70	30	8.0	8.1	14.6	16.8	10.0	10.7	10.5	11.3
Short Sellers	0.25	−1.23	22.1	−8.2	70	30	6.8	6.8	10.7	11.3	8.0	8.1	8.3	8.5
REITS (unleveraged)	0.25	0.69	8.9	4.6	70	30	5.8	5.8	7.0	7.0	6.2	6.2	6.4	6.3
RE apartment	2.00	0.56	9.3	3.7	60	40	5.6	5.6	7.2	7.2	6.2	6.2	6.4	6.4
RE industrial	2.00	0.77	10.8	5.1	60	40	6.0	6.0	7.6	7.7	6.6	6.6	6.9	6.9
RE office	2.00	0.85	11.6	5.6	60	40	6.1	6.1	7.8	7.9	6.8	6.8	7.0	7.1
RE retail	2.00	0.80	11.6	5.3	60	40	6.0	6.0	7.8	7.9	6.7	6.8	7.0	7.0
REITS apartment	0.25	0.76	14.2	5.1	70	30	6.0	5.9	8.5	8.6	6.7	6.7	7.0	7.0
REITS industrial	0.25	0.86	14.6	5.7	70	30	6.2	6.1	8.6	8.8	6.9	6.9	7.1	7.1
REITS office	0.25	0.93	16.1	6.2	70	30	6.3	6.2	9.0	9.3	7.1	7.1	7.4	7.4
REITS retail	0.25	0.86	16.0	5.7	70	30	6.1	6.1	9.0	9.2	7.0	7.0	7.3	7.3
Timber—U.S. South	10.00	0.67	14.3	4.4	50	50	5.8	5.8	9.3	9.3	7.5	7.5	7.8	7.8
Timber—U.S. West	10.00	0.76	17.1	5.0	50	50	5.9	6.0	10.7	10.6	8.3	8.3	8.7	8.6
Farm land (row crop)	2.00	0.55	13.8	3.6	50	50	5.6	5.5	8.5	8.6	7.0	7.1	7.3	7.4
Farm land (perm. crop)	2.00	0.89	19.6	5.9	50	50	6.2	6.2	10.3	10.8	8.2	8.5	8.6	8.8

fair conditions may require much longer, probably up to three years. Therefore, we set for real estate an "implied" illiquidity horizon significantly longer than only one year. Finally, timber investments are associated with the longest lock-up.

Generally, we find that the illiquidity premia increase overproportionally with the illiquidity horizon and the level of risk. This implies that high-risk investments with long lock-up times are compensated particularly well for illiquidity. This seems in line with common intuition, and all results are fairly straightforward from this perspective.

5 Conclusion

Most alternative assets are illiquid, thus reducing the investor's flexibility. In a rational world this should be compensated. The so-called illiquidity premium is a compensation for locking in investments without having the option to resell at any time until liquidation. Investors who can afford to wait should expect to be rewarded for assuming illiquidity. For purposes of asset allocation, we need to know how illiquidity is compensated; however, CAPM does not provide an answer, since it operates in a perfect world.

With the literature on this topic being fairly scarce, we derive our own approaches for illiquidity compensation. In this paper, we present Sharpe ratio and Put Option approaches providing fairly consistent results. However, illiquidity compensation is no free lunch, as an investor in illiquid assets does not have the choice to liquidate his entire portfolio immediately.

While illiquidity compensation turns out to be substantial for assets with long illiquidity horizons at high-risk levels, low-risk assets only get moderate liquidity premia, no matter how long the lock-in is. Furthermore, the illiquidity horizon is not the time between the first inflow and the last outflow; rather, it is a duration type of measure.

According to our approaches, LBOs and venture capital clearly get the highest compensation due to their high total and systematic risk as well combined with significant lock-up times. On the other hand, the relatively low compensation for real estate and timber is mainly due to both low total and systematic risk. On the other hand, real estate and timber investments are mainly driven by diversification purposes and their capability to hedge against inflation rather than boosting return.

Appendix

A.1 About continuous and simple returns

Assume a risk-free asset has a value of S at time 0 and provides a continuous return, r, which is fully reinvested. Further, it has a start value of 1 at time 0. Then its value at time T is

$$S(T) = e^{rT} \tag{A.1}$$

Apparently, the asset's value is defined by r at any time in the future. Its simple return, $R(T)$, between

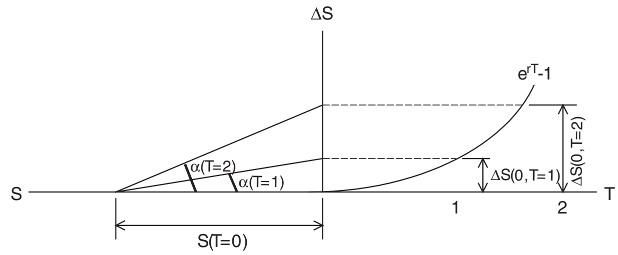

Figure A.1 Continuous and simple returns.

time 0 and T is

$$R(T) = \frac{S(T)}{S} - 1 = \frac{e^{rT}}{1} - 1 = e^{rT} - 1 \tag{A.2}$$

That is, the simple return, as opposed to the continuous return, is a function of the according time horizon which it is calculated for. In Figure A.1, we try to visualize simple returns.

$\Delta S(0, T)$ is the increase in value between time 0 and T. Obviously,

$$R(T) = \frac{\Delta S(0, T)}{S} = \tan(\alpha) \tag{A.3}$$

That is, the simple return is the *slope of the hypotenuse*. Unfortunately, the continuous return cannot be visualized; it is just a parameter defining the compounded curve's *shape*. Rather, the compounding *effect* can be visualized, which is, again, the simple return over a selected horizon, T.

Based on (A.2), we conclude that for

$$T > \frac{\log(1 + r)}{r} \tag{A.4}$$

the simple return is larger than the continuous return.[31] Given a continuous return of 10%, for instance, the simple return for 0.9531 years equals 10% as well, and the simple return for one year is 10.52%.

Note that we have not involved the concept of risk yet; the continuous and simple returns differ regardless of risk. To sum up this section, our main findings are:

- the geometric return *does not* depend on the horizon;
- the simple return *does* depend on the horizon;
- the asset has an *infinity* of simple returns;
- the simple return *depends on both*, the continuous return and the time horizon;
- beyond a certain horizon, T, the simple return is *always larger* than the continuous return.

Now, let us introduce the concept of risk by assuming an asset with the following *volatile* continuous return:

$$\tilde{\mu} \sim N(\mu, \sigma) \tag{5}$$

Then, according to the theory of the log-normal distribution, the asset's expected value at time T is[32]

$$E[S(T)] = e^{(\mu+\sigma^2/2)T} \tag{6}$$

and its median value is

$$M[S(T)] = e^{\mu T} \tag{8}$$

where σ is the continuous return's volatility. Comparing (A.1) and (7), it becomes apparent that the *asset's expected value at time T is larger in the case of a volatile rather than a constant continuous return.* The according reason is the asymmetric effect of risk through compounding.[33]

Further, as in (A.2), we calculate the asset's expected simple return between time 0 and T. It is

$$E[R(T)] = \frac{E[S(T)]}{S} - 1 = \frac{e^{(\mu+\sigma^2/2)T}}{1} - 1 = e^{(\mu+\sigma^2/2)T} - 1 \tag{A.5}$$

Hence, $\mu + \sigma^2/2$ is that continuous return which provides—when continuously compounded—the expected simple return over the according horizon. That is, an important conclusion is that we have to deal with three different types of returns, the

- *continuous* return;
- *risk-corrected continuous* return;
- expected *simple* return for a certain horizon.

Further, just as the continuous return, the *corrected continuous return does not depend on the horizon*. To make things more transparent, a numeric example: assume the asset has a continuous return of 10% and the annualized risk of the continuous return is 40%. Then, we get

- continuous return = 10%;
- risk-corrected continuous return = $0.1 + 0.4^2/2 = 18\%$;
- expected simple return over 1 year = $e^{0.18} - 1 = 19.72\%$.

Sometimes, people confuse the *risk-corrected continuous return* and the *simple return* over a selected horizon. However, *only* the difference between the continuous return and the risk-corrected continuous return equals $\sigma^2/2$, as the simple return over a selected horizon differs from the continuous return even if the continuous return is constant. Hence, $0.1972 > 0.1 + 0.4^2/2$. Note, particularly in the case of a moderate risk, the simple return exceeds the continuous return mainly because of compounding rather than because of the risk effect.

From this perspective, Hull's numerical example[34] is not too compelling. He intends to demonstrate the difference between the *continuous return* and the *risk-corrected continuous return* which is—exclusively—a risk effect. However, he computes simple *annual* returns which reflect a considerable compounding effect. Given his numbers, half of the estimated return difference is due to pure compounding of a continuous return and has nothing to do with risk. In order to empower his statement, he should take monthly or even weekly data. Hence, he clearly confuses the issues in question.

To sum up this section, our main findings are:

- the expected value at time T is larger in the case of a *volatile* rather than a constant continuous return;

- the corrected continuous return provides—when continuously compounded—the expected simple return over the according horizon;
- there are *three different types of returns*: continuous return, risk-corrected continuous return, and simple return for a certain horizon;
- simple returns are the result of both the *compounding and the risk effect*.

A.2 Option approach: derivation of formulas

Assume the asset's value at time 0 is S. For the exercise price

$$X_1 = Se^{\mu^* T} \tag{10}$$

Black–Scholes[35] provides

$$p_1 = Se^{(\mu^*-r)T}N(-d_{12}) - SN(-d_{11}) \tag{A.6}$$

with

$$d_{11} = \frac{(r-\mu)\sqrt{T}}{\sigma} \tag{A.7.1}$$

$$d_{12} = \frac{(r-\mu-\sigma^2)\sqrt{T}}{\sigma} \tag{A.7.2}$$

that is

$$p_1 = Se^{(\mu^*-r)T}N((\mu-r+\sigma^2)\sqrt{T}/\sigma) - SN((\mu-r)\sqrt{T}/\sigma) \tag{A.6$'$}$$

For the exercise price

$$X_2 = Se^{\mu T} \tag{11}$$

we get

$$p_2 = Se^{(\mu-r)T}N(-d_{22}) - SN(-d_{21}) \tag{A.8}$$

with

$$d_{21} = \frac{(r-\mu+\sigma^2/2)\sqrt{T}}{\sigma} \tag{A.9.1}$$

$$d_{22} = \frac{(r-\mu-\sigma^2/2)\sqrt{T}}{\sigma} \tag{A.9.2}$$

and hence

$$p_2 = Se^{(\mu-r)T}N((\mu-r+\sigma^2/2)\sqrt{T}/\sigma) - SN((\mu-r-\sigma^2/2)\sqrt{T}/\sigma) \tag{A.8$'$}$$

Further, because of the discount

$$\Delta = p_1 - p_2 \tag{9}$$

and the expected value at the end of the lock-in period of

$$E[S(T)] = Se^{\mu^* T} \tag{6}$$

we get

$$(S - (p_1 - p_2))e^{\mu(T)T} = Se^{\mu^* T} \tag{A.10}$$

and by rearranging

$$\mu(T) = \log\left(\frac{e^{\mu^* T}}{1-(p_1-p_2)/S}\right)T^{-1} \tag{A.10$'$}$$

where $\mu(T)$ is the total return including the illiquidity premium given a lock-in period of T. Finally, by substituting (A.7) and (A.9) into (A.10'), we get

$$\mu(T) = \left(\frac{e^{\mu^* T}}{1-e^{(\mu^*-r)T}N(-d_{12})+N(-d_{11})+e^{(\mu-r)T}N(-d_{22})-N(-d_{21})}\right)T^{-1} \tag{A.10$''$}$$

Ultimately, there is the point where the denominator is 0 and hence the resulting total return infinite. This is the case if

$$e^{(\mu^*-r)T}N(-d_{12}) - N(-d_{11}) - e^{(\mu-r)T}N(-d_{22}) + N(-d_{21}) = 1 \tag{A.11}$$

Beyond this point, that is, if (A.11) becomes negative, there is only an imaginary solution to the problem.

A.3 Sharpe ratio approach: derivation of formulas

All risks and instantaneous returns are given. While the market is liquid and hence its return rate, μ_m^*, does not depend on some particular horizon, the asset's return rate, $\mu(T)$, is a function of T. The asset's Sharpe ratio over its illiquidity horizon according to (18) is assumed to equal the market's Sharpe ratio for the same period:

$$\begin{aligned} \text{SR(entire market, } T) &= \frac{e^{\mu_m^* T} - e^{rT}}{e^{\mu_m^* T}\sqrt{e^{\sigma_m^2 T} - 1}} \\ &= \frac{e^{\mu(T)T} - e^{rT}}{e^{\mu(T)T}\sqrt{e^{\sigma^2 T} - 1}} \\ &= \text{SR(asset, } T) \qquad \text{(A.12)} \end{aligned}$$

By rearranging, we get

$$\mu(T) = \log\left[\left(1 - \sqrt{\frac{e^{\sigma^2 T} - 1}{e^{\sigma_m^2 T} - 1}} \cdot \left(1 - \frac{1}{e^{(\mu_m^* - r)T}}\right)\right)^{-1/T} e^r\right] \qquad \text{(A.12}')$$

The larger T, the steeper the curve. Finally, there is the point where the function becomes vertical and the return infinite. This point is achieved when the denominator is zero and hence meets the following condition:

$$\left(1 - \sqrt{\frac{e^{\sigma^2 T} - 1}{e^{\sigma_m^2 T} - 1}} \cdot \left(1 - \frac{1}{e^{(\mu_m^* - r)T}}\right)\right) = 0 \qquad \text{(A.13)}$$

Beyond this point, that is, if (A.13) is negative, there is only an imaginary solution to the problem. Referring to Figure 10, this is the case when the risk dominates so much that the spike is heightened but no more widened as a result of an increased required return.

Intuitively, it is not immediately clear why there is no solution to any combination of risk and illiquidity horizon. In this case, a look into Eq. (18) might help. It shows that an increase in the asset's expected return increases both the numerator and the denominator of SR(T), and hence it is not necessarily clear which one dominates.

Notes

1 Perold and Sharpe (1988), 21f.

2 The New York Times (2001) shares this view.

3 Longstaff (1995), p. 1770:

$$F(V, T) = V\left(2 + \frac{\sigma^2 T}{2}\right) N\left(\frac{\sqrt{\sigma^2 T}}{2}\right) + V\frac{\sqrt{\sigma^2 T}}{2}\exp\left(-\frac{\sigma^2 T}{8}\right) - V$$

4 For traded stocks of an average risk, it implies an illiquidity premium of about 1.0–1.5%.

5 The integration/segmentation approach is documented at length in Singer and Terhaar (1997). Further, Terhaar *et al.* (2003) make extensive use of this approach.

6 The criteria are: Analytical (adequate control and regulation, marketability and liquidity, meaningful impact, non-redundant, manageable estimation risk), Legal, and Talent Availability. For further details, see Brinson *et al.* (1986), p. 17.

7 Brinson *et al.* (1986), p. 17.

8 Singer and Terhaar (1997), 44ff.

9 In most instances, the investor's home bias is due to preference and hence voluntary.

10 We assume the same price for risk.

11 Chaffe (1993), p. 182.

12 Admittedly, the problem is more difficult in reality, as no options are traded for alternative asset. However, in practice, they might be substituted with options of their liquid counterparts.

13 Because in the short run, the risk is large relative to the average return.

14 Admittedly, this could be one justification for Chaffe's approach being generous in the short run, as many SEC 144 stocks are in financial distress.

15 Based on Aitchison and Brown (1966), Eq. (2.7), p. 8.

16 More generally, the reason is that $E[f(x)] \neq f(E[x])$, if the function is not linear.

17 For the derivation of the formulae, see Appendix A.2.

18 In particular, this is a crucial issue in the domain of private equity.

19 Which is still significantly below S.

20 Theoretically, the illiquidity horizon is an easy concept: it is the time span over which an investment is locked in and cannot be liquidated. However, in practice, it is more complicated, as illiquid investments such as private equity start to generate cash flows after a certain time. The illiquidity horizon is not the time between drawing the first

portion of the commitment and getting the last cash flow. Rather, the illiquidity horizon means the "'average" illiquidity horizon, and hence, it is a *duration* to be estimated based on the involved time spans and flow sizes. This is clearly where qualitative judgment becomes involved.

21 Note, for deriving the required total return of the illiquid asset based on a discount according to Eq. (A.10″), we need to know the total return of the same asset if it was liquid. For more details, see Appendix A.2.

22 Typically, it is a Sharpe ratio of annualized risks and returns.

23 For a detailed explanation of returns of risky assets, see Appendix A.1.

24 Derived from Aitchison and Brown (1966), 8, (2.8).

25 Hodges *et al.* (1997) make the same observation. They derive an according formula based on simple returns.

26 Again, as for the integration approach, the reference portfolio is the entire market and our proxy is GIM.

27 Note that a rate of return not being a function of T applies to any liquid asset, whereas a rate which is a function of T applies in the case of illiquidity with an horizon of T. For the formula and its derivation, see Appendix A.3.

28 For the mathematical properties of the function, see Appendix A.3.

29 That is, in the case of total segmentation.

30 According to our sources, real estate is held in the average for a little longer than 10 years.

31 $R(T) > r \Longrightarrow e^{rT} > 1 + r \Longrightarrow T > \log(1 + r)/r$.

32 Based on Aitchison and Brown, (1966), Eq. 2.7, p. 8.

33 More generally, the reason is $E[f(x)] \neq f(E[x])$, if the function is not linear.

34 Hull (1993), p. 213.

35 For the put option formula, see Wilmott *et al.* (1995), p. 48.

References

Aitchison, J. and Brown, J.A.C. (1966). *The Lognormal Distribution*. Cambridge: Cambridge University Press.

Amihud, Y. and Mendelson, H. (1986a). "Liquidity and Stock Returns." *Financial Analyst's Journal* May–June, 43–48.

Amihud, Y. and Mendelson, H. (1986b). "Asset Pricing and the Bid–Ask Spread." *Journal of Financial Economics* **17**, 223–249.

Brinson, G.P., Diermeier, J.J. and Schlarbaum, G.G. (1986). "A Composite Portfolio Benchmark for Pension Plans." *Financial Analysts Journal* March–April, 15–24.

Chaffe, D.B.H. (1993). "Option Pricing as a Proxy for Discount for Lack of Marketability in Private Company Valuations." *Business Valuation Review* December, 182–188.

Hodges, C.W., Taylor, W.R.L. and Yoder, J.A. (1997). "Stocks, Bonds, the Sharpe Ratio, and the Investment Horizon." *Financial Analysts Journal* November–December, 74–80.

Hull, J.C. (1993). *Options, Futures, and other Derivatives*, 2nd edn. Englewood Cliffs, NJ: Prentice Hall.

Longstaff, F.A. (1999). "Optimal Portfolio Choice and the Valuation of Illiquid Securities." Working paper.

Longstaff, F.A. (1996). "Placing No-Arbitrage Bounds on the Value of Nonmarketable and Thinly-Traded Securities." *Advances in Futures and Options Research* **8**.

Longstaff, F.A. (1995). "How much Can Marketability Affect Security Values?" *The Journal of Finance* **L**(5).

Mercer Capital (1998). Why Not Black-Scholes Rather than the Quantitative Marketability Discount Model? Internet-memo (www.bizval.com).

Perold, A.F. and Sharpe, W.F. (1988). "Dynamic Strategies for Asset Allocation." *Financial Analysts Journal* January–February, 16–27.

Silber, W.L. (1991). "Discounts on Restricted Stock: The Impact of Illiquidity on Stock Prices." *Financial Analysts Journal* July–August, 60–64.

Singer, B. and Terhaar, K. (1997). *Economic Foundations of Capital Market Returns*. The Research Foundation of the Institute of Chartered Financial Analysis.

Smith, J.K., Smith, R.L. and Williams, K. (2000). "The SEC's 'Fair Value' Standard for Mutual Fund Investment in Restricted Shares and Other Illiquid Securities," Working paper.

Terhaar, K., Staub, R. and Singer, B. (2003). "The Appropriate Policy Allocation for Alternative Investments." *Journal of Portfolio Mangement* (in press).

The New York Times (2001). "A Lifeline, With Conditions, C1 and C7." Thursday, May 10.

Wilmott, P., Howison, S. and Dewynne, J. (1995). *The Mathematics of Financial Derivatives*. Cambridge, UK: Cambridge University Press.