

Hedge fund transparency: quantifying valuation bias for illiquid assets

Risk measures, such as the Sharpe ratio, used by investment professionals are only as good as the accuracy of the asset price data used to derive them. Nowhere is this issue more relevant than for hedge funds, which often invest in less liquid assets such as convertible bonds. Here, Eric Weinstein and Adil Abdulali devise a 'phantom price' framework for illiquid assets and show how to generalise the Sharpe ratio to incorporate liquidity risk

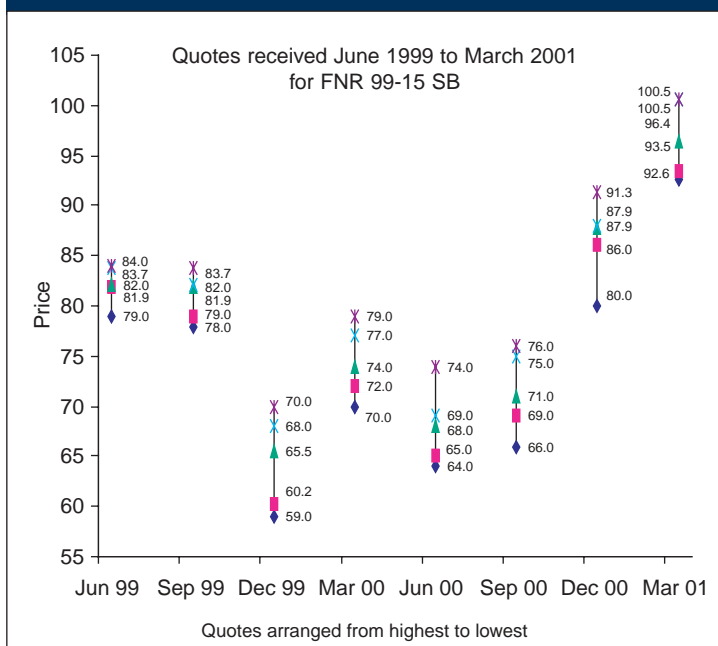
In this article, we introduce the 'phantom price' framework for measuring valuation risk. The techniques revealed relate to assets that fall in the middle of the liquidity continuum between liquid assets with transparent pricing such as IBM common stock, and illiquid assets with opaque pricing such as real estate. These assets, which we label 'translucent', include mortgage-backed securities, high-yield bonds, certain convertible bonds and others. Translucent assets are traded so privately or with such infrequency that transaction prices are not readily available for pricing. Instead, the market convention is to gather a collection of 'indicative prices' from several sources (eg, broker-dealers). A representative asset from our hedge fund¹, an inverse-floater called the FNR 99-15 SB, provides an example that shows the range of such indicative prices (see figure 1).

In the absence of consensus pricing, it is difficult to calculate the benchmarking statistics that are applied to other liquid asset classes. In particular, the Sharpe ratio, which gives a measure of historical return adjusted for volatility, has been unavailable to traders who operate in these markets. This frustrates the analysis of portfolios that mix liquid and illiquid assets, as the statistics in one sector of the portfolio fail to extend to other sectors.

We argue here that for this important asset category, the valuation problem can be resolved when such illiquidity is considered an additional source of risk, for which the hedge fund investor should rightly expect additional reward. In particular, we develop a generalisation of the Sharpe ratio for the representative hedge fund investor relative to instruments that may fail to have a single agreed-upon market price.

We propose a two-step process to find the best representative price for

1. Time series for broker-dealer quotes on an MBS derivative



a security for which multiple indicative prices are received. The first step is to decide, for each indicative price, what the probability is of it being the trade price. For example, if five indicative prices received for a security are 99, 99.5, 100, 100.5 and 50, one might well believe that 50 has a lower probability of being the trade price than 99.5. The assignment of such probabilities represents a belief about the market. We represent market beliefs as a probability distribution constructed from indicative prices.

The second step is to decide a single price that best represents the distribution. That is, once the probability of each price occurring as a transaction price is established, the person ascribing an instantaneous value to the asset has to answer the question: what is the lowest single definite price I would accept for this asset in exchange for the ambiguous situation of the multiple indicative prices if my intention were to immediately liquidate the security in question? To make this clear, one might imagine that many sellers of the asset in the above example would be likely to accept a definite price of 98, while only the most risk-fearing would exchange the uncertainty for a price of 80. It is a feature of our framework that this step of deciding the representative price involves the risk tolerance of the person ascribing a value to the asset. In the above example, a risk-loving investor might be comfortable with the asset valued at 100, and would not sell at 95, as opposed to a risk-averse investor who would be happy with 95.² In this paper, we use a Von-Neumann-Morgenstern risk preference framework to set up the computational details of mathematically representing such investor risk preferences.

We will demonstrate that this approach can be made concrete in the case of a 'representative investor'. Any statistic that can be computed from historical prices can be generalised meaningfully to translucent assets, and collapses to the usual one when a single market clearing price is in evidence.

Our approach

We begin by assuming that for each security in our portfolio, there is a collection of m broker-dealers who are willing and able in every period t to give indicative prices for the various instruments in question. We let $v_t \in \mathbb{R}_+^m$ represent the non-negative indicative prices. It is then natural to try to replace the fuzziness and uncertainty contained in the range of broker-dealer prices represented by v_t with a single 'phantom price'.

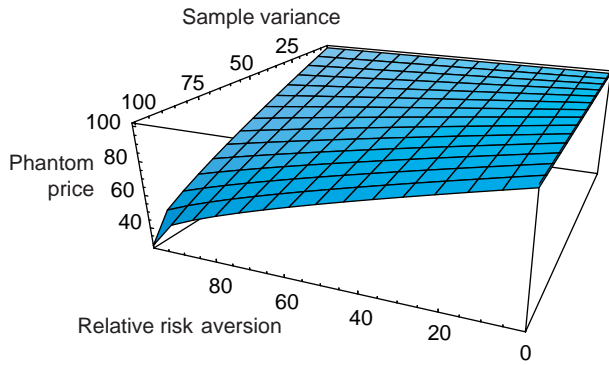
We can immediately see that it is possible to implement this approach at a crude level: using the mean price of the sample will formally accomplish the task, albeit at the unacceptable cost of masking the additional risk presented by price uncertainties. If mean prices were to be used as the phantom, a risk-averse investor would be unable to distinguish between an illiquid portfolio consisting of translucent assets and a related portfolio of common stocks whose prices followed the mean time-series of its translucent 'twin'.

To incorporate this missing risk component naturally, we shall seek our phantom price as a 'certainty equivalent' of the uncertain situation confronting

¹ The specific fund from which this time series was drawn is Mortgage Back Opportunity, managed by the second author at AdKap

² Although, for clarity, the previous discussion contemplates an investor valuing a long position, one can easily make an analogous argument that applies to both short and long positions by substituting 'position value' for 'price'

2. Phantom price as a function of both variance and risk aversion



the representative investor. We assume that the representative investor views the indicative prices as constituting a random sample from a probability distribution supported on the non-negative real numbers. If the actual price that would result from asking for binding quotes is drawn from the same distribution, the investor can attempt to work backwards at guessing the nature of the distribution from the nature of the sample of indicative prices.

Let v_t be the vector of indicative prices in period t and assume that the representative investor has Von-Neumann-Morgenstern risk preferences given by the sub-utility function u_t in period t . If given v_t the representative investor believes that B_{v_t} is the probability distribution generating the sample, the result will be a set of time-dependent phantom prices:

$$\mathbb{P}(v_t) = u_t^{-1} \left(\int_{\mathbb{R}_+} u_t(x) B_{v_t}(x) dx \right) \quad (1)$$

determined from the expected value to the investor.

In contrast to the case of liquid assets, formula (1) shows that the market beliefs and investor preferences are essential in the construction of performance statistics which recognise the additional risk from uncertain pricing.

Representing investor market beliefs

The simplest approach to modelling investor beliefs about the nature of translucent prices is to imagine that prices are symmetrically distributed about the sample mean as random variables following a familiar bell-shaped curve. Unfortunately, since negative prices are not generally possibilities for the securities under study, choosing normal distributions will associate unacceptable positive probabilities for the security reaching negative valuations. Therefore, we seek a similar two-parameter family of distributions that are supported on the non-negative real numbers. The most natural such family is the family of gamma distributions

$$\text{Gamma distribution PDF} = \Psi_{[\alpha, \lambda]}^\Gamma(x) = \frac{\lambda^{-\alpha} e^{-\frac{x}{\lambda}} x^{\alpha-1}}{\Gamma[\alpha]} \quad (2)$$

which are parameterised not by mean and variance but by a shape parameter α and a scale parameter λ . We can, however, make a change of variables by defining α and λ as functions of the sample mean μ and variance σ^2 according to:

$$\alpha = \frac{\mu^2}{\sigma^2} \quad \lambda = \frac{\sigma^2}{\mu} \quad (3)$$

It is now possible to check to see that calculating the shape and scale parameters α and λ from a choice of μ and σ^2 in this manner returns the gamma probability density function with precisely the desired sample mean and variance with which one began.

Representing investor preferences

Within the risk calculus, a representative investor can generally be expected to possess risk preferences given by a sub-utility function $u(x)$ ex-

hibiting decreasing absolute risk aversion³ at any given moment in time. The simplest way of ensuring this is to make the simplifying assumption that our representative investor's relative Arrow-Pratt risk aversion is both constant and time-independent:

$$a_t(x) = - \frac{x u''[x]}{u'[x]} = a \in \mathbb{R} \quad (4)$$

with a strictly positive⁴ so that his/her Von-Neumann Morgenstern sub-utility function may be given according to the rule⁵

$$CRRR_a(x) = u_a(x) = \begin{cases} \frac{x^{1-a}}{1-a} & a \neq 1 \\ \ln[x] & a = 1 \end{cases} \quad (5)$$

Given a vector $v(t) = \{v_i(t)\}_{i=1}^m$ of sample indicative prices, the sample mean μ and sample variance σ^2 are given by the usual formulas

$$\mu(v_t) = \frac{\sum_{i=1}^m v_i(t)}{m} \quad \sigma^2(v_t) = \frac{\sum_{i=1}^m (v_i(t) - \mu(v(t)))^2}{m-1} \quad (6)$$

Thus, in the generic case where $a \neq 1$ we replace the uncertain price spread represented by the vector of sample prices v_t by the phantom price:

$$\mathbb{P}_a^\Gamma(v_t) = \left((a-1) \int_0^\infty \frac{x^{1-a}}{a-1} \lambda^{-\alpha} e^{-\frac{x}{\lambda}} \frac{x^{\alpha-1}}{\Gamma[\alpha]} dx \right)^{\frac{1}{1-a}} \quad (7)$$

$$= \left(\frac{\sigma_t^2}{\mu_t} \right) \left(\frac{\Gamma \left[\frac{\mu_t^2 - (a-1)\sigma_t^2}{\sigma_t^2} \right]}{\Gamma \left[\frac{\mu_t^2}{\sigma_t^2} \right]} \right)^{\frac{1}{1-a}} \quad (8)$$

calculated relative to the coefficient of relative risk aversion a .

In the case that $a = 1$ the result is that:

$$\mathbb{P}_1^\Gamma(v_t) = \text{Log} \left(\frac{\sigma_t^2}{\mu_t} \right) + \frac{\Gamma' \left[\frac{\mu_t^2}{\sigma_t^2} \right]}{\Gamma \left[\frac{\mu_t^2}{\sigma_t^2} \right]} \quad (9)$$

Having established explicit formulas⁶ for phantom prices, we examine how these formulas behave in various limiting cases.

1Variance $\sigma^2 \rightarrow 0$. For a collection of samples with a shared positive sample mean of μ_0 and a fixed level of risk-aversion, as the variance approaches 0, the phantom price approaches μ_0 .

1Constant of relative risk aversion $\alpha \rightarrow 0$. For a given non-zero variance, as the coefficient of relative risk aversion a approaches 0, the phantom price approaches μ_0 , which corresponds to an idealised risk-neutral investor.

1Constant of relative risk aversion $a \rightarrow \pm\infty$. For fixed variance, as a becomes a large positive (high risk-aversion), the phantom price drops well

³ For a discussion of absolute and relative risk aversion, see for example Arrow (1965) and Pratt (1964)

⁴ If $a = 0$, the individual will be risk neutral and if a is strictly negative, the individual will exhibit gambling or 'risk-loving' behaviour

⁵ Note that with all such choices of 'functional forms', the use of power utility functions is not without its difficulties. Nevertheless, from among the common choices of simplified utility functions, those in equation (5) would appear to compare favourably

⁶ The recipe we have given produces a concrete formula for phantom prices as a so-called 'certainty equivalent' in standard terminology. The nomenclature, however, is slightly unfortunate; a better term within the model might be 'riskless equivalent' as 'uncertainty' is properly reserved for the element of chance in the model which defies objective quantification and is therefore irremovable (Knight, 1921). Any practical recipe such as ours cannot pretend to eliminate all such uncertainty from the model at hand

below μ_0 , and as a becomes increasingly negative (gambling behaviour) the phantom price rises well above μ_0 .

Figure 2, representing samples with $\mu_0 = 100$, shows that when either the risk aversion is set to $a = 0$ (ie, pure risk neutrality), or consensus pricing is available (ie, the variance of the sample is negligible), the value of the non-linear phantom price function is given by the familiar arithmetic mean μ_0 .

Bounded assets

In instances where security prices are bounded from above (eg, zero-coupon bonds) or below away from zero (eg, securities with embedded put options), the above approach must be modified slightly. For a security bounded on two sides, the minimal change can be affected by positing that the indicative prices of the bounded security should no longer be thought of as drawn from a member of the family of Gamma distributions but from the family of generalised Beta distributions of the form:

$$\begin{aligned} \text{Generalised beta distribution PDF} &= \Psi_{[\alpha, \beta]}^{\text{Beta}}(x) \\ &= \frac{\left(\frac{x-L}{U-L}\right)^{\alpha-1} \left(1 - \frac{x-L}{U-L}\right)^{\beta-1} \Gamma[\alpha + \beta]}{(U-L)\Gamma[\alpha]\Gamma[\beta]} \end{aligned} \quad (10)$$

supported on the interval $[L, U]$ so as not to associate finite probabilities with impossible prices. At time t , α_t and β_t are calculated from the security's upper bound U (eg, par), lower bound L (eg, 0), and the relevant mean and variance according to the formulas:⁷

$$\alpha_t = \frac{(L - \mu_t)((L - \mu_t)(\mu_t - U) - \sigma_t^2)}{(L - U)\sigma_t^2} \quad (11)$$

$$\beta_t = \frac{(U - \mu_t)((L - \mu_t)(U - \mu_t) + \sigma_t^2)}{(L - U)\sigma_t^2} \quad (12)$$

Sample application

In this section, we will apply our method for valuing illiquid securities to a simple three security portfolio. The translucent assets in the example are a subset of holdings of our hedge fund and the prices used are the historical indicative prices from a collection of real broker-dealers.

When a portfolio consists solely of liquid assets with point prices, the value of the portfolio:

$$P_{\text{Portfolio}} = \sum_{i=1}^n m_i P_i \quad (13)$$

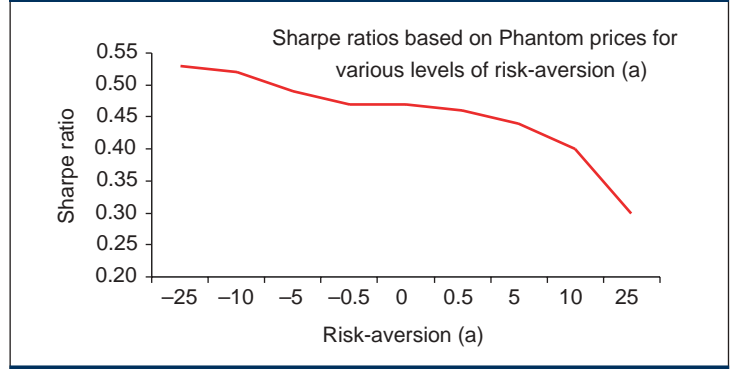
is equal to the sum of the individual security prices P_i weighted by the number of shares m_i held. With phantom prices, the above formula actually ceases to be a tautology and gives only an approximation. If one wants an exact formula, a simple theoretical adjustment is required, which renders the exact answer somewhat more cumbersome in closed form (eg, requiring use of hypergeometric functions). For brevity, in this section we will instead calculate the portfolio performance according to (13) and (10) using phantom prices in place of point prices for every translucent asset in the portfolio.

With regard to transparency, we note that for all the statistics contemplated in the examples, the portfolio manager need not disclose specific assets, just the dealer quotes received for them and the quantities of each. In fact, managers would not be required to change or disclose their algorithms or methods for arriving at and reporting prices that they use for net asset value (NAV) calculations.

Phantom price Sharpe ratio versus manager Sharpe ratio

Our first application deals with measuring bias in reporting the volatility of returns. A common concern for potential investors in hedge funds with translucent assets is the discretion that managers have in the calculation of NAV. The result of a manager valuing assets at the low end of the range of quotes received when the portfolio performs well and at the high end when it performs poorly is a smoothing of the volatility of returns. This necessarily inflates

3. Risk-sensitive Sharpe ratio for a sample translucent portfolio



risk-adjusted returns as measured by statistics such as the Sharpe ratio.

Investors would like to get insight into the manager's behaviour with respect to the range of prices received for translucent assets. One would like to accomplish this without forcing the manager to necessarily reveal or change valuation algorithms or require him to disclose specific positions.

We present a basic example of the uses of phantom prices in measuring the performance of a mortgage derivative portfolio. We posit a hypothetical manager with a strategy of buying mortgage derivatives and hedging their price volatility with a five-year maturity US Treasury note (5yT).

The prices are real but the quantities and trading strategy of the 5yT is imaginary. The imaginary trading strategy results in a Sharpe ratio of 2.37 for the two-year period of our example. In comparison with the usual Sharpe ratio, we now calculate the Sharpe ratio based on the phantom prices calculated from the broker-dealer quotes for various values of risk-aversion. The results are presented in figure 3.

As the phantom price framework we are presenting depends on a risk aversion parameter 'a', the Phantom price Sharpe ratio based on indicative price series picks up the same parameterisation; there are at least two aspects of this phenomena that bear mentioning. First, the Phantom price Sharpe ratio would converge to the usual Sharpe ratio based on manager marks as the spread between broker-dealer quotes narrowed and converged to the manager's marks (ie, in the zero variance limit, the function would become a function of 'a' that was identically constant).

Second, the Phantom Sharpe ratio (now naturally function-valued as in figure 3) remains a mechanical consequence of the time-series data that is investor independent. The scalar-valued hedge fund Sharpe ratio has not disappeared in this framework, however. As we shall see in the next section, it ceases to be a purely utility independent gauge of fund performance and becomes instead a tool for detecting smoothing irregularities and potential gaming in manager marking of translucent assets; a new tool specifically suited to a problem that does not arise in the presence of well-defined spot prices.

Manager NAV/Phantom price NAV

Phantom prices can provide some insight regarding the absolute level of valuation relative to the range of quotes received for translucent assets. We look at the portfolio constructed above and examine the ratio (Manager NAV/Phantom price NAV) for various values of risk aversion. At any moment in time this ratio gives a measure of the aggressiveness of pricing used by the manager relative to the quotes received. The sensitivity of the ratio for various values of risk aversion gives an indication of the presence of translucent assets in that portfolio. But if the ratio is greater than one for a risk aversion parameter of zero, then one can conclude that the manager has on average priced translucent assets at a price greater than the mean price received. Likewise, if the ratio is less than one for $a = 0$, this implies that the manager is conservative relative to the mean prices.

A useful statistic is the manager's average risk aversion over time. We

⁷ For further details and phantom price formulas, see Weinstein & Abdulali (2001)

construct this by examining the average of the time series of the (Manager NAV/Phantom price NAV) and solving for the risk-aversion parameter which makes this average = 1. This level of a is the answer to the question: what has been our manager's average risk aversion for the period in question? To illustrate this we present this time series for our hypothetical manager's portfolio (see table A). It turns out that in our example, the hypothetical manager's average risk-aversion is represented by $a = -5$.

Based on these results one could conclude facts such as the following: compared with the historical average pricing behaviour of the manager, he/she started out being conservative until Sep 1999, was aggressive until Sep 2000, and has been conservative since then. The maximum deviation over the average was 7.8% and the maximum deviation under the average was 5.5%.

These statistics for two different managers with the same strategy and assets would reveal differences in valuation practices. Style change relative to the managers' own historical biases in valuing translucent assets will be revealed. Measuring the change in Phantom price valuations for high levels of risk aversion would alert the investor if a manager suddenly changed the quantity of translucent assets in a portfolio.

Just like the usual Sharpe ratio, investing based purely on such considerations would turn out to be misguided. However, the added knowledge of the manager's biases relative to a risk-neutral investor in evaluating translucent assets should provide a yardstick with which to compare managers as well as give investors a level of comfort regarding the degree of discretion their portfolio managers are using to price translucent assets.

In figure 4, the manager claims a Sharpe ratio of 2.34 while the phantom price evaluation reveals a significantly lower value of 0.45. The value for 'a' has been chosen here to match the manager's average risk aversion in NAV reporting over the period in question.

Extensions of the approach

Here, we indicate some slight modifications to our earlier approach.

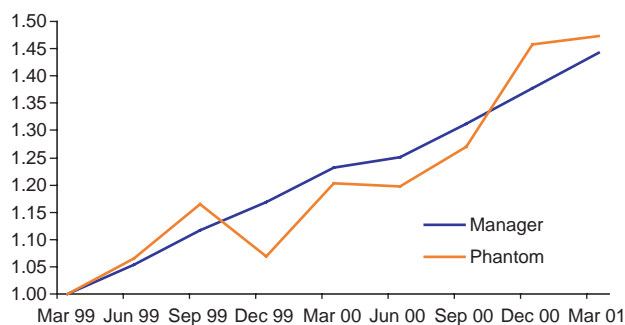
1The asset integration hypothesis. An individual with fixed risk preferences given by an unchanging sub-utility function $u(x)$ may, nevertheless, exhibit changing risk-taking behaviour. If a given investment proposition comes with a moderate probability of leaving the investor destitute, most rational investors are unlikely to accept, even if the expected returns are strongly positive. Should such investors become wealthier, investments, which at a lower wealth level threatened ruin, may become newly attractive without any change in their risk preferences given by the sub-utility function $u(x)$.⁸

1Portfolio aggregation. In a fully liquid portfolio, the total price of the portfolio at any given time is the sum of the prices of the individual stocks weighted by the number of shares held as in (13). In the phantom price framework, one has to use tools from the calculus of probability density functions in order to aggregate the values of individual assets. At the portfolio level, the full aggregation procedure requires convolving the respective probability density functions, but has been de-emphasized here for purposes of exposition.

A. Time series for hypothetical manager's

Value of 'a'	-5
Start	1.000
June 99	0.989
September 99	0.962
December 99	1.078
March 2000	1.010
June 2000	1.027
September 2000	1.018
December 2000	0.945
March 2001	0.972
Average	1
Max	1.078
Min	0.945

4. Manager marks vs phantom prices: growth of \$1



1Testing for the coefficient of relative risk aversion a. We have shown how to produce Sharpe ratios through phantom prices depending on the average risk aversion exhibited by a manager's marks. In effect, the indicative prices from the various time periods face the manager with a series of questions about values for which the manager's marks form the answers.

If an investor finds himself with a different risk profile from the manager, he may wish to calculate phantom price statistics using his own value for 'a'. As for managers, the best way to determine this information is to give the investor a series of questions designed to elicit the level of risk aversion a which best models the investor's personal risk preferences.⁹

Conclusion

We have argued here that the absence of consensus pricing common among hedge fund instruments is best treated as an additional form of risk for which the investor should expect to be rewarded. As such, the risk tolerance of the investor becomes critical for accurate portfolio evaluation once the uncertain situation represented by a collection of indicative prices is made concrete. Assuming constant relative risk aversion leads to explicit prescriptions for portfolio-level phantom prices, which can in turn be used to extend Sharpe ratios or other performance statistics familiar from the evaluation of portfolios based upon more liquid instruments. 1

Adil Abdulali is principal and Eric Weinstein is head of research at AdKap. Research sponsored jointly by Capital Market Risk Advisors and AdKap. Visit <http://www.transparentrisk.com> for technical discussions in greater depth. The authors wish to thank Leslie Rahl for valuable input and feedback on this paper. Please address all correspondence to weinstein@post.harvard.edu

⁸ For a more complete discussion see Weinstein & Abdulali (2001)

⁹ See, for example, Weinstein & Abdulali (2001)

References

- Arrow K, 1965
Aspects of the theory of risk bearing
Yrjo Hahnsson Foundation, Helsinki
- Asness C, R Krail and J Liew, 2001
Do hedge funds hedge
Journal of Portfolio Management 28(1), pages 6-19
- Knight F, 1921
Risk, uncertainty, and profit
Boston, Houghton-Mifflin
- Pratt J, 1964
Risk aversion in the small and in the large
Econometrica 32, pages 122-136
- Weinstein E and A Abdulali, 2001
Phantom prices: valuation transparency for illiquid assets
Working paper, available at www.transparentrisk.com
- Weisman A, 2002
Dangerous attractions: informationless investing and hedge fund performance measurement bias
To appear in Journal of Portfolio Management, spring