

Gain, Loss and Asset Pricing

Antonio Bernardo and Olivier Ledoit*

December 11, 1996

Abstract

In this paper we develop an approach to asset pricing in incomplete markets that gives the modeller the flexibility to control the tradeoff between the precision of equilibrium models and the credibility of no-arbitrage methods. We rule out the existence of investment opportunities that are very attractive to a benchmark investor. The key feature of our approach is the measure of attractiveness employed: the gain-loss ratio. The gain (loss) of a portfolio is the expectation, under a benchmark risk-adjusted probability measure, of the positive (negative) part of the portfolio's excess payoff. The benchmark risk-adjusted probability measure incorporates valuable prior information about investor preferences and portfolio holdings. A restriction on the maximum gain-loss ratio in the economy has a dual representation in terms of admissible pricing kernels: it is equivalent to a bound on the ratio of extreme deviations from the benchmark pricing kernel.

Price bounds are derived by computing all prices which do not permit the formation of portfolios with gain-loss ratios in excess of some prespecified level. We give an example where we bound the price of an option on a non-traded asset that is correlated with a traded asset. The resulting bounds lie strictly between the Black-Scholes price and the no-arbitrage bounds, and they are sharper when (i) the maximum allowable gain-loss ratio is lower, (ii) the correlation between the non-traded and traded asset is higher, and (iii) the volatility of the non-traded asset is lower. This has implications for pricing real options and executive stock options, and for performance evaluation of portfolio managers who use derivatives.

*Assistant Professors of Finance, Anderson Graduate School of Management, University of California Los Angeles, Box 951481, 110 Westwood Plaza, Los Angeles, CA 90095-1481. E-mail: antonio.bernardo@anderson.ucla.edu, ledoit@ucla.edu; phone - Bernardo: (310) 825 2198, Ledoit: (310) 206 6562; fax: (310) 206 5455. We thank Michael Brennan, Mark Grinblatt, Richard Roll, Pedro Santa-Clara, Jean-Luc Vila, and participants to seminars at Arizona and UCLA for helpful discussions. Ilya Sharapov provided assistance with the numerical methods. Errors are ours.

1 Introduction

Much of finance is concerned with pricing new securities. The most credible approach to valuing new securities is to use no-arbitrage arguments. In fact, if markets are (dynamically) complete no-arbitrage also yields precise prices: every new security is redundant and thus its price must be equal to the price of a portfolio which replicates its payoffs. The main limitation of the no-arbitrage approach is that the pricing implications are not very precise in the more realistic case when markets are (dynamically) incomplete. Another approach to valuing new securities is to specify an equilibrium model in which investors have a *specific* system of preferences and endowments (e.g. Lucas, 1978). Equilibrium models yield precise prices for new securities. However, such specific assumptions can be too strong and the pricing implications many not be credible. Perhaps we do not know exactly what preferences are or what the aggregate endowment is (cf. Roll's (1977) critique).

In this paper we develop an approach to pricing new securities in incomplete markets which allows us to control the trade-off between the precise pricing implications of a specific equilibrium model and the credible pricing implications of no-arbitrage arguments. We do so by ruling out more than arbitrage opportunities: we also rule out investment opportunities that are very attractive to a benchmark investor. The key feature of our approach is the measure of attractiveness used: the gain-loss ratio. The gain (loss) of a portfolio is simply the expectation, under a benchmark risk-adjusted probability measure, of the positive (negative) part of the portfolio's excess payoff.¹ The benchmark risk-adjusted probability measure incorporates valuable prior information about investor preferences and portfolio holdings. Consider, for example, a frictionless economy in which prices are determined as if there exists a representative investor with utility function u and aggregate endowment e . In this case the benchmark risk-adjusted probability measure is obtained by multiplying the true probability of a state ω occurring, $p(\omega)$, with the marginal utility of consumption in that state, $u'(e(\omega))$, and normalizing. It is straightforward to show that a necessary condition for equilibrium in such an economy is that the gain-loss ratio equals unity. The gain-loss ratio summarizes the attractiveness of an investment opportunity to an investor

¹If x is the excess payoff then the gain is $E^*[\max\{x, 0\}]$ and the loss is $E^*[\max\{-x, 0\}]$, where the expectation is taken under the benchmark risk-adjusted probability measure.

described by (u, e) . If the ratio is one the investment is fairly priced, below one it is unattractive, and above one it is attractive. For example, a gain-loss ratio of two means that the benchmark investor receives twice as much gain as is necessary for her to hold the asset or, equivalently, that she risks only half the loss that she would be willing to accept to hold the asset.

If we are certain that u and e characterize the representative investor then we can impose the restriction that the maximum gain-loss ratio available in the economy is unity and derive exact prices for the new security. If, however, we are uncertain about the characteristics of the representative investor we may want to rule out investment opportunities that are too attractive; that is, investment opportunities with gain-loss ratios which exceed some specified level \bar{L} . Thus, our notion of attractiveness depends not only on the ceiling imposed, \bar{L} , but on the beliefs about the characteristics of the representative investor. The choice of \bar{L} allows the modeller the flexibility to control the trade-off between the precise, but unreliable, predictions of a benchmark equilibrium model and the reliable, but imprecise, predictions of no-arbitrage models. We obtain price bounds on new securities by finding all prices which do not permit gain-loss ratios above \bar{L} . The bounds can be made sharper by including other assets, called basis assets, in the analysis: price bounds are then obtained by finding all prices of the new security which do not allow the benchmark investor to construct high gain-loss ratio portfolios from the new security and the basis assets.

The idea of obtaining pricing implications by imposing a restriction on the attractiveness of investment opportunities available in an economy has been proposed by Ledoit (1995) and Cochrane and Saá-Requejo (1996). In both papers, the Sharpe ratio is used as the measure of attractiveness. This is reasonable in their settings because they consider securities with returns that are almost normally distributed and the Sharpe ratio summarizes the attractiveness of investments in such cases: Ledoit (1995) considers stocks, for which normality is a good first-order approximation, and Cochrane and Saá-Requejo (1996) focus on derivatives that traded frequently, so that prices almost (or exactly) follow a Wiener diffusion process with normally distributed increments. If returns are not normally distributed, however, the Sharpe ratio is not a good measure of the attractiveness of an investment opportunity. In fact, arbitrage opportunities can have arbitrarily low Sharpe ratios!

The gain-loss ratio, on the other hand, is an economically reasonable measure of attractiveness

because it accounts explicitly for investor preferences. Furthermore, it is compatible with no-arbitrage: an investment opportunity is an arbitrage if and only if it has infinite gain-loss ratio. We also demonstrate that an investment opportunity with a high gain-loss ratio is *closer* to arbitrage, in an economically reasonable norm, than an investment with low gain-loss ratio. And, like a Sharpe ratio bound, a gain-loss ratio bound gives the modeller (i) the flexibility to let the degree of tightness vary over a wide range; (ii) nonparametric pricing bounds; and (iii) an intuitive economic interpretation. Furthermore, if returns are normally distributed then a gain-loss restriction is equivalent to a Sharpe ratio restriction.

The maximum gain-loss ratio in an economy has a useful, and equivalent, representation in terms of admissible stochastic discount factors. A stochastic discount factor, or pricing kernel, is a random variable m that prices all excess payoffs x by $E[mx] = 0$. The value $m(\omega)$ has the interpretation of the price per unit of probability of an Arrow-Debreu security that pays one in state ω and zero elsewhere; or what is commonly known as the state price per unit of probability. If the representative agent is described by (u, e) then the value of the pricing kernel at each ω is proportional to the marginal utility of consumption evaluated at the aggregate endowment in that state: $m^*(\omega) = \frac{u'(e(\omega))}{E[u'(e)](1+r)}$. We can think about $m^*(\omega)$ as representing the benchmark investor's *willingness to pay* per unit of probability for an ω -state claim. We demonstrate that a bound on the maximum gain-loss ratio in an economy is equivalent to ruling out pricing kernels that deviate too much from the benchmark pricing kernel. Intuitively, state claim prices that are too far away from the benchmark investor's willingness to pay would be arbitrated away by buying (selling) relatively cheap (dear) state claims.

Our restriction on the gain-loss ratios available in the economy, and its equivalent interpretation in terms of pricing kernels, is closely related to a large literature on pricing kernel restrictions. Hansen and Jagannathan (1991) showed that if a riskless asset exists then a restriction on the maximum Sharpe ratio in the economy is equivalent to a bound on the second moment of admissible pricing kernels. Snow (1991) generalized this result to the q^{th} moment of the stochastic discount factor $E[m^q]^{1/q}$ for $1 < q < \infty$. Stutzer (1993) showed that restricting the maximum expected utility attainable by an agent with constant absolute risk aversion preferences is equivalent to a restriction on the value of $E[m \log(m)]$. The problem with these restrictions, for deriving pricing implications on new securities, is that they do not prevent state prices from being too

close, or even equal to, zero. Consequently, such restrictions do not preclude zero prices for many realistic securities such as out-of-the-money calls and puts. Bansal and Lehmann (1996) showed that restricting the maximum expected utility attainable by an agent with log utility is equivalent to a restriction on $E[-\log(m)]$. Such a restriction, however, fails to prevent state prices from being too close to infinity! With such a bound, we cannot preclude very high prices for many reasonable securities. In contrast, our gain-loss restriction imposes a restriction on pricing kernels that prevents state prices (per unit of probability) from being too high or too low.

We demonstrate the power of a gain-loss restriction by computing pricing bounds for options on an asset that does not trade. The dynamic replication arguments of Black-Scholes (1973) do not apply here, thus the option cannot be priced exactly by arbitrage. Furthermore, equilibrium models such as Rubinstein's (1976) or Brennan's (1979) require exact knowledge of the utility function, which is difficult to obtain. We assume that there exists another traded asset which is correlated with the non-traded asset and whose price is known. Price bounds on the option are computed by finding all prices which do not permit the existence of attractive portfolios created from the option on the non-traded asset and a similar option on the non-traded asset. We obtain credible pricing bounds which get tighter as (i) the maximum gain-loss ratio goes down; (ii) the correlation between the traded and non-traded asset goes up; and (iii) the volatility of the non-traded asset decreases. We believe that the gain-loss approach is extremely valuable in applications where the security returns are not normally distributed, including (i) valuing real options on non-traded assets; (ii) valuing executive stock options when the executive cannot trade the options or the underlying due to insider trading restrictions; (iii) evaluating the performance of portfolio managers who invest in derivatives; (iv) pricing options on a security whose price follows a jump-diffusion or a fat-tailed Pareto-Levy diffusion process; and (v) pricing fixed-income derivatives in the presence of default risk.

The organization of this paper is as follows. Section 2 demonstrates the problems with a bound on the maximum Sharpe ratio when returns are not normally distributed. Section 3 introduces the gain-loss ratio for a benchmark investor, the relation between a restriction on the gain-loss ratio and arbitrage, and demonstrates the equivalence of a restriction on the gain-loss ratio and a restriction on admissible pricing kernels. In section 4 we demonstrate the implications of a bound on the gain-loss ratios available in the economy for asset prices. We discuss some issues related

to the modeller's choice of (i) benchmark pricing kernels; (ii) ceiling on the gain-loss ratio; and (iii) the set of basis assets to include in the analysis. We then demonstrate the effectiveness of a gain-loss restriction in obtaining price bounds for options on non-traded assets and discuss some important applications. In Section 5 we link our gain-loss approach to earlier literature on alternatives to mean-variance such as mean- p^{th} moment preferences and prospect theory, and to other restrictions on stochastic discount factors in the literature and their relation to our approach. Section 6 concludes.

2 Restricting the Maximum Sharpe Ratio in the Economy

In this section, we review an approach to asset pricing introduced earlier by Ledoit (1995) and Cochrane and Saá-Requejo (1996), paying special attention to the case where returns are not normal.

2.1 Literature Review

Ledoit (1995) adapts Ross's (1976) Arbitrage Pricing Theory (APT) to a finite universe of stocks by assuming that the maximum Sharpe ratio is below a ceiling \bar{S} . The tightness of his pricing implication depends explicitly on the value of \bar{S} , which must be specified *ex-ante*. Varying \bar{S} yields a family of models ranging from an equilibrium model as tight as the CAPM to a no-arbitrage model as loose as the APT. Thus, restricting the maximum Sharpe ratio can be seen as a flexible way to chart the middle ground between the two paradigms of asset pricing: equilibrium and no-arbitrage.

Independently, Cochrane and Saá-Requejo (1996) apply the same method to derivatives. They obtain option pricing bounds that converge to the Black-Scholes (1973) price as the time between consecutive trading dates converges to zero. They also obtain option pricing bounds when trading is continuous but volatility and interest rates are stochastic. As above, the tightness of these bounds decreases in the ceiling \bar{S} imposed *ex-ante* on the maximum Sharpe ratio in the economy.

A key feature in both these papers is that returns are almost normally distributed. Ledoit (1995) considers stocks, for which normality is a good first-order approximation. Cochrane and

Saá-Requejo (1996) focus mainly on short (or infinitesimal) trading intervals, where prices almost (or exactly) follow a Wiener diffusion process with normally distributed increments. The next section shows what happens when returns are not normal.

2.2 Non-Normality

There are many important problems where returns are far from normally distributed. Here we show how the approach of Ledoit (1995) and Cochrane and Saá-Requejo (1996) fares outside normality. Counter-examples show that low Sharpe ratio portfolios may be *very* attractive, and also that high Sharpe ratio portfolios may *not* be attractive

2.2.1 Arbitrage

Ledoit (1995) calls portfolios with Sharpe ratios above \bar{S} “approximate arbitrage opportunities”, and Cochrane and Saá-Requejo (1996) call them “good deals”. In both cases the basic intuition is to rule out more than arbitrage opportunities. Indeed, a normally distributed security is an arbitrage opportunity if and only if it has infinite Sharpe ratio. But outside normality this linkage breaks down: arbitrage opportunities can have arbitrarily low Sharpe ratios. An extreme example is the portfolio whose excess return follows the probability distribution function:

$$f(x) = \begin{cases} \frac{2}{x^3} & \text{if } x \geq 1 \\ 0 & \text{if } x < 1. \end{cases} \quad (1)$$

The mean excess return is two and variance is infinite, therefore the Sharpe ratio is zero, even though it is an arbitrage opportunity, since you double your money even in the worst case. There exists a deep incompatibility between the Sharpe ratio and no-arbitrage outside normality, and more generally between any mean-variance concept and no-arbitrage. Intuitively, it is because quadratic utility displays satiation.²

Cochrane and Saá-Requejo (1996) attempt to resolve the incompatibility by ruling out arbitrage opportunities *in addition* to Sharpe ratios above \bar{S} . To see why this falls short of the mark,

²Dybvig and Ingersoll (1989) show that arbitrage opportunities exist in equilibrium if agents are mean-variance optimizers and markets are complete.

consider the lottery whose ticket costs one cent today, and where winners pocket fifty billion dollars next year with probability 0.1, and nothing otherwise:

$$\begin{array}{rcl}
 & & 50,000,000,000 \text{ with probability } 0.1 \\
 0.01 & \nearrow & \\
 & \searrow & 0 \text{ with probability } 0.9.
 \end{array} \tag{2}$$

This lottery is not a pure arbitrage opportunity, and its Sharpe ratio is very low: 0.33 (below a U.S. stock market index), so it is not ruled out by the combined no-arbitrage / Sharpe ratio restriction.³ Yet it gives you one chance in ten of becoming the richest person in the world in exchange for a coin that you may not even bother picking up from the pavement. It corresponds to the intuitive meaning of “approximate arbitrage” or “good deal”. A proper definition of these words should include everything that comes close to an arbitrage opportunity without actually being one. It should form an airtight envelope around the set of arbitrage opportunities – in topological terms: a *neighborhood*. But there is no way to construct that with the Sharpe ratio (cf. Footnote 4). Therefore we need to develop an alternative to the Sharpe ratio that is compatible with the notion of arbitrage when returns are not normal.

2.2.2 Preferences

Under normality, the Sharpe ratio completely summarizes the attractiveness of an investment opportunity. But outside normality it is impossible to make general statements that are preference-free other than no-arbitrage (and stochastic dominance).

It may seem that, for any set of preferences, securities with high enough Sharpe ratios will be too attractive to exist in equilibrium. This is simply not true. For example, consider a two-period Lucas (1978) model where markets are complete, agents have constant relative risk aversion γ , and aggregate endowment is distributed according to a χ^2 with n degrees of freedom. When $n/2 > \gamma$ the expectation of marginal utility is finite, therefore the model is not pathological. Yet when $\gamma \geq n/4$ the maximum Sharpe ratio in the economy is infinite. Proposition 1 shows that this is a general result.

³The way Cochrane and Saá-Requejo combine no-arbitrage with the Sharpe ratio restriction is slightly more complicated, but this does not affect the point we are making.

Proposition 1 Consider a two-period equilibrium model where markets are complete, agents have utility function u , and aggregate endowment is distributed according to a cumulative distribution function F . If the marginal utility u' is unbounded then there exists distributions F such that the expectation of marginal utility is finite and the maximum Sharpe ratio in the economy is infinite.

Proof of Proposition 1 Let e denote aggregate endowment. Since u' is unbounded, there exists distributions F for e such that $E[u'(e)] < +\infty$ and $E[u'(e)^2] = +\infty$. When $E[u'(e)] < +\infty$ the model is not pathological. By the Hansen-Jagannathan (1991) duality result, $E[u'(e)^2] = +\infty$ means that the maximum Sharpe ratio in the economy is infinite. \square

Intuitively, providing insurance against bad states of nature remains unattractive despite a high Sharpe ratio. Proposition 1 shows that high Sharpe ratios cannot always be ruled out. In some sense, there is an assumption on preferences hidden inside the Sharpe ratio restriction. We believe that this is unavoidable: any restriction stronger than no-arbitrage must make *some* assumption on preferences. But we also believe that this assumption should be stated in a more explicit way. Furthermore, the economist should have the freedom to specify any type of assumption that he or she deems appropriate in any particular context. Therefore we need to develop an alternative to the Sharpe ratio that takes preferences into account explicitly.

2.2.3 Parametric vs. Nonparametric

As Hansen and Jagannathan (1991) point out, Sharpe ratio bounds are *nonparametric*: if we wanted to parametrize the set of utility functions that satisfy the bound, we would need an infinite number of parameters. Under normality, the nonparametric character of this restriction is not exploited, since a parametric restriction on the market price of independent sources of risk has the same effect. For example, when Cochrane and Saá-Requejo (1996) price options in continuous time with stochastic volatility, there is a single unknown parameter: the market price of volatility risk. It is always possible to specify a reasonable range for this parameter, and derive option pricing bounds from that range. There is little reason to invoke a sophisticated nonparametric method when a simple parametric method does the same job.

Since the Sharpe ratio constraint runs into problems outside normality, we may be tempted to try a parametric alternative. As it turns out, parametric methods run into a severe problem of their own. Assume for simplicity that there is a single normally distributed source of risk whose market price is unknown. Under normality, it is sufficient to know one parameter: the price of a security perfectly correlated with that source of risk. Outside normality, we need to know an *infinite* number of parameters: the price of every contingent claim derived from that security. And a parametric method with an infinite number of parameters truly is *nonparametric*. Thus, we need to impose a nonparametric restriction similar to the Sharpe ratio bound.

3 Gain-Loss Restriction

In this section we try to extend the work of Ledoit (1995) and Cochrane and Saá-Requejo (1996) to assets with payoffs that are not normally distributed. Given the specific problems highlighted in Section 2.2, the wish list for an alternative to the Sharpe ratio restriction is that: (1) it is compatible with no-arbitrage, (2) it takes preferences into account explicitly, and (3) it is nonparametric. Ideally, it should also present the same advantages as the Sharpe ratio constraint, namely: (4) be flexible enough to let the degree of tightness vary over a wide range, (5) have an intuitive economic interpretation, and (6) have a simple expression both in terms of returns and in terms of stochastic discount factors.

We consider a very general two-period model where assets trade at a certain price today and promise a random payoff in the future. Let Ω denote the set of future states of the world with generic element ω , and let $p(\omega)$ denote the probability density function. We assume that there exists a default-free bond with interest rate r . Let Z denote the space of asset payoffs with generic element z and pricing functional $\pi(z)$. Excess payoffs are constructed as $x = z - (1 + r)\pi(z)$.

3.1 Benchmark

We start from an equilibrium model where prices are determined as if a representative agent with utility function u held the aggregate endowment e . In general, we do not know the true model: we do not have direct access to utility functions, and we only have proxies for the aggregate

endowment. However, we can often choose a benchmark equilibrium model (u, e) that we find “reasonable” by some criterion. This benchmark will embody the assumption of preferences that we need to make in order to have a restriction stronger than no-arbitrage. It will appear explicitly in our alternative to the Sharpe ratio restriction. One important point is that the economist has the freedom to specify any benchmark model that he or she deems appropriate.

If (u, e) was the true model, then the first-order condition of the utility maximization program under budget constraint would imply the following relation for any security with payoff $z \in Z$ and price $\pi(z)$:

$$\pi(z) = \frac{\mathbb{E}[u'(e)z]}{\mathbb{E}[u'(e)](1+r)}. \quad (3)$$

If the representative investor is risk averse then u' is decreasing, thus an investment that has relatively high payoff in states of the world in which e is low has relatively high price. Intuitively, such an investment is desirable to risk-averse investors because it allows them to smooth consumption across future states of nature. In general, the price an investor is willing to pay for an asset does not depend solely on the statistical properties of the payoffs (i.e. mean and variance) but also depends on preferences u , and the aggregate endowment e .

3.2 Gain and Loss

Next, we define what constitutes an attractive investment opportunity relative to the benchmark. Decompose the excess payoff $x = z - (1+r)\pi(z)$ into positive and negative part: $x = x^+ - x^-$, where $x^+ = \max(x, 0)$ and $x^- = \max(-x, 0)$. Equation (3) can be rewritten as:

$$\begin{aligned} \mathbb{E}[u'(e)x] &= 0 \\ \mathbb{E}[u'(e)x^+] - \mathbb{E}[u'(e)x^-] &= 0 \\ \frac{\mathbb{E}[u'(e)x^+]}{\mathbb{E}[u'(e)x^-]} &= 1 \\ \frac{\mathbb{E}^*[x^+]}{\mathbb{E}^*[x^-]} &= 1, \end{aligned} \quad (4)$$

where $\mathbb{E}^*[\cdot]$ denote expectation under the benchmark risk-adjusted probability measure $p^*(\omega) = p(\omega)u'(e(\omega))/\mathbb{E}[u'(e)]$.

We call $E^*[x^+]$ the *gain*, since it measures how much you have a chance to win; and $E^*[x^-]$ the *loss*, since it measures how much you risk to lose. Thus $E^*[x^+]/E^*[x^-]$ is the *gain-loss ratio*. The basic properties of the gain-loss ratio are analogous to those of the Sharpe ratio: it is invariant to the multiplication of x by a positive scalar, and the gain-loss ratio of a short position is the inverse of the ratio of the long position. The gain-loss ratio completely summarizes the attractiveness of a zero-cost portfolio relative to the benchmark: if the gain-loss ratio is equal to one then the portfolio is fairly priced, below one it is unattractive, and above one it is attractive. For example, a gain-loss ratio of two means that the benchmark investor receives twice as much gain as would be necessary for her to hold the asset. Or, equivalently, that she risks only half the loss that she would be willing to accept to hold the asset.

3.3 The Restriction

If the benchmark equilibrium model is close to the truth then we should not observe portfolios that are very attractive relative to that benchmark. This suggests that the gain-loss ratio should never be far from one.

Restriction 1 *The maximum gain-loss ratio in the economy is below \bar{L} :*

$$\frac{E^*[x^+]}{E^*[x^-]} \leq \bar{L}, \quad (5)$$

for every excess payoff x .

The value of the ceiling \bar{L} must be chosen *ex-ante* in the range $1 \leq \bar{L} \leq +\infty$. Restriction 1 is equivalent to imposing the floor $1/\bar{L}$ on the minimum gain-loss ratio. It is tight if \bar{L} is close to one, loose if \bar{L} is close to infinity. If we are certain that u and e characterize the representative agent in the economy then we can impose the ceiling $\bar{L} = 1$ to obtain precise equilibrium pricing implications. If, however, we are uncertain about the characteristics of the representative agent, we may want to rule out opportunities that are too attractive relative to the benchmark. Thus, our notion of attractiveness depends not only on the ceiling imposed, \bar{L} , but also on the beliefs about the characteristics of the representative agent.

3.4 Arbitrage

There are many different ways to measure deviation from the benchmark, and the gain-loss ratio is only one of them. We chose it because it is intrinsically compatible with the notion of arbitrage. Intuitively, an arbitrage opportunity is potential gain without potential loss.

Proposition 2 *An asset is an arbitrage opportunity if and only if it has infinite gain-loss ratio.*

Proof of Proposition 2 Let x denote an excess payoff.

$$\begin{aligned}
 x \text{ is an arbitrage opportunity} &\iff \Pr\{x^+ > 0\} > 0 \text{ and } \Pr\{x^- > 0\} = 0 \\
 &\iff \mathbb{E}^*[x^+] > 0 \text{ and } \mathbb{E}^*[x^-] = 0 \\
 &\iff \frac{\mathbb{E}^*[x^+]}{\mathbb{E}^*[x^-]} = +\infty. \quad \square
 \end{aligned}$$

Let us define the following metric on excess payoffs: $\|x\| = \mathbb{E}^*[|x|]$. This is known as the L^1 norm (under the risk-adjusted probability measure). According to this metric, when we get closer to an arbitrage opportunity the gain-loss ratio goes up.

Theorem 1 *The normalized distance to the set of arbitrage opportunities decreases in the gain-loss ratio. For every excess payoff x :*

$$\frac{\min_{a \in A} \|x - a\|}{\|x\|} = \frac{1}{\frac{\mathbb{E}^*[x^+]}{\mathbb{E}^*[x^-]} + 1} \tag{6}$$

where $A = \{x : x \geq 0\}$ denotes the set of excess payoffs on arbitrage opportunities.

Proof of Theorem 1 Notice that $\forall a \in A \quad \mathbb{E}^*[|x - a|] = \mathbb{E}^*[x^+ - a\mathbf{1}_{\{x \geq 0\}}] + \mathbb{E}^*[|x^- + a\mathbf{1}_{\{x < 0\}}|]$, where $\mathbf{1}_{\{\cdot\}}$ denotes the indicator function of an event. The way to minimize the first term of the decomposition is to impose that $a = x^+$ when $x \geq 0$. The way to minimize the second term is to impose that $a = 0$ when $x < 0$. Therefore the minimum distance to the arbitrage set is attained by $a = x^+$. And, at the minimum, $\|x - x^+\| = \|x^-\|$. It implies:

$$\frac{\min_{a \in A} \|x - a\|}{\|x\|} = \frac{\|x^-\|}{\|x\|} = \frac{\mathbb{E}^*[x^-]}{\mathbb{E}^*[x^+ + x^-]} = \frac{1}{\frac{\mathbb{E}^*[x^+]}{\mathbb{E}^*[x^-]} + 1}. \quad \square \tag{7}$$

As a consequence, if you get close enough to an arbitrage opportunity, eventually your gain-loss ratio will exceed \bar{L} . This is the sense in which gain-loss ratios above \bar{L} form a *neighborhood* of the set of arbitrage opportunities A .⁴

Proposition 3 *If a sequence of excess payoffs converges to an arbitrage opportunity then its gain-loss ratio goes to infinity.*

Proof of Proposition 3 If $x_n \rightarrow a \in A$ then $\min_{a' \in A} \|x_n - a'\| \leq \|x_n - a\| \rightarrow 0$. Furthermore, arbitrage opportunities have strictly positive norm, therefore $\|x_n - a\| \leq \|a\|/2$ for large enough n , which implies $\|x_n\| \geq \|a\| - \|x_n - a\| \geq \|a\|/2 > 0$. Since its numerator converges to zero and its denominator is bounded away from zero for large enough n , the normalized distance to arbitrage of x_n converges to zero. Proposition 3 then follows from Theorem 1. \square

The converse is also true.

Proposition 4 *If a sequence of excess payoffs converges to a nonzero limit and its gain-loss ratio goes to infinity, then the limit is an arbitrage opportunity.*

Proof of Proposition 4 Assume that $x_n \rightarrow x \neq 0$ and that $E^*[x_n^+]/E^*[x_n^-] \rightarrow +\infty$. Then Theorem 1 implies $\min_{a \in A} \|x_n - a\|/\|x_n\| \rightarrow 0$. Since $\|x_n\| \rightarrow \|x\| > 0$, we have $\min_{a \in A} \|x_n - a\| \rightarrow 0$. Therefore there exists $a_n \in A$ such that $\|x_n - a_n\| \leq \min_{a \in A} \|x_n - a\| + 1/n \rightarrow 0$. It implies that $\|x - a_n\| \leq \|x - x_n\| + \|x_n - a_n\| \rightarrow 0$, i.e. $\min_{a \in A} \|x - a\| = 0$. Since $\|x\| > 0$, the normalized distance to arbitrage of x is zero. By Theorem 1, it implies that the gain-loss ratio of x is infinite. By Proposition 2, the limit x is an arbitrage opportunity. \square

These results show that the gain-loss ratio restriction is fully compatible with the no-arbitrage principle. Of course, the results are dependent on the metric, but that is always the case. One advantage of our metric is that it is induced by benchmark preferences chosen to be “reasonable”. Given any other metric than $\|\cdot\|$, we could construct a measure of attractiveness from the

⁴By contrast, Sharpe ratios above some ceiling \bar{S} form a neighborhood of the set of *zero-variance* arbitrage opportunities $A_0 = \{x : x \in A \text{ and } \text{Var}[x] = 0\}$. Using the quadratic norm $\|x\|_2 = \sqrt{E[x^2]}$ (under the true probability measure), the Sharpe ratio is inversely related to the normalized distance to the set of zero-variance arbitrage opportunities by: $\min_{a \in A_0} \|x - a\|_2/\|x\|_2 = (1 + (E[x]/\sqrt{\text{Var}[x]})^2)^{-1/2}$. Since A_0 is a strict subset of A , high Sharpe ratio portfolios do not form a neighborhood of A in general. Under normality, however, it works out fine, since the only arbitrage opportunities that are traded are the ones with zero variance.

normalized distance to arbitrage defined with that metric. Any such measure of attractiveness would have the same fundamental properties as the gain-loss ratio, although perhaps a less intuitive interpretation.

3.5 Dual Restriction on the Stochastic Discount Factor

In this section we demonstrate that a restriction on the maximum gain-loss ratio in the economy is equivalent to a dual restriction on the set of admissible stochastic discount factors. A stochastic discount factor (or pricing kernel) is a random variable $m(\omega)$ that prices all excess payoffs x by $E[mx] = 0$. The value $m(\omega)$ has the interpretation of the price per unit of probability of an Arrow-Debreu security that pays one in state ω and zero elsewhere; or what is commonly known as the state price per unit of probability. If the benchmark model (u, e) was true then the pricing kernel would be proportional to the marginal utility of the benchmark agent evaluated at his wealth in that state:

$$m^*(\omega) = \frac{u'(e(\omega))}{E[u'(e)](1+r)}. \quad (8)$$

A useful way to think about $m^*(\omega)$ is as the benchmark investor's *willingness to pay* per unit of probability for an ω -state claim. The willingness to pay is greater for claims that pay off in bad states (assuming that the benchmark investor is risk-averse).

If every payoff $z \in Z$ has price $\pi(z)$, then the space of admissible pricing kernels is defined as: $M = \{m : \forall z \in Z E[mz] = \pi(z)\}$. We can also define the set of non-negative admissible pricing kernels $M^+ = \{m \in M : m \geq 0 \text{ a.s.}\}$ and the set of positive admissible pricing kernels $M^{++} = \{m \in M : m > 0 \text{ a.s.}\}$.

It is well-known that there are relations between the investment opportunities available in an economy and the set of admissible pricing kernels. For example, a necessary and sufficient condition for absence of arbitrage opportunities is the existence of a positive admissible pricing kernel, i.e. $M^{++} \neq \emptyset$. Intuitively, if $m(\omega) = 0$ for some ω then investors can get something for nothing by purchasing claims that pay off in state ω . Hansen and Jagannathan (1991) also demonstrate that a bound on the maximum Sharpe ratio available in an economy is equivalent to

a bound on the variance of admissible pricing kernels:

$$\max_{\substack{x \in Z \\ \pi(x)=0}} \frac{E[x]}{\sqrt{\text{Var}[x]}} = \min_{m \in M} \frac{\sqrt{\text{Var}[m]}}{E[m]}. \quad (9)$$

Another important relation is that the expectation of any admissible discount factor is equal to $1/(1+r)$.

We will now show the relation between the gain-loss ratios available in an economy and the set of admissible stochastic discount factors. To provide intuition, first consider an economy with two future states of the world and assume that the benchmark investor is risk-neutral. In this case, the benchmark investor's willingness to pay is the same in both states: $m_1^* = m_2^*$. Suppose, however, that the true state prices verify $m_1 < m_2$. Then the benchmark investor would do well to buy (sell) claims that pay off in state one (two) because such claims are relatively cheap (expensive). This suggests that attractive investment opportunities exist for the benchmark investor. We can measure attractiveness by the gain-loss ratio. If (p_1, p_2) represents the probabilities of state 1 and state 2 occurring then the price of an excess payoff (x_1, x_2) is given by $E[mx] = p_1 m_1 x_1 + p_2 m_2 x_2 = 0$. Since the investor "buys low and sells high", $x_1 > 0$ and $x_2 < 0$ and the gain-loss ratio is given by:

$$\frac{E^*[x^+]}{E^*[x^-]} = \frac{E[x^+]}{E[x^-]} = \frac{p_1 x_1}{-p_2 x_2} = \frac{p_1 x_1}{p_1 m_1 x_1 / m_2} = \frac{m_2}{m_1}. \quad (10)$$

The greater is the ratio of the highest to lowest state price per unit probability in the economy, the greater is the gain-loss ratio available to the benchmark investor. Intuitively, when *actual* state prices differ from the benchmark investor's *willingness to pay* attractive investment opportunities are available.

This intuition generalizes to cases where the representative investor is risk-averse and the number of future states is infinite. The following theorem demonstrates that a restriction on the maximum gain-loss ratio in the economy is equivalent to a restriction on the ratio of extreme deviations from the benchmark pricing kernel.

Theorem 2 *Let Z denote a space of payoffs, $\pi(\cdot)$ the pricing functional, and M^{++} the set of*

positive admissible pricing kernels. Then we have:

$$\max_{\substack{x \in Z \\ \pi(x)=0}} \frac{E^*[x^+]}{E^*[x^-]} = \min_{m \in M^{++}} \frac{\text{esssup}(m/m^*)}{\text{essinf}(m/m^*)} \quad (11)$$

where esssup (essinf) denotes the essential supremum (infimum) of a random variable with respect to its underlying probability measure.⁵

Proof of Theorem 2 The following three propositions are equivalent: a) there is a zero-cost arbitrage opportunity in Z ; b) the maximum gain-loss ratio is infinite; c) the set M^{++} is empty. In this case, Equation (11) holds, since the minimum of an empty set is $+\infty$ by convention.

Now rule out arbitrage opportunities. Assume for the moment that m^* is a constant, i.e. $E^*[\cdot] = E[\cdot]$. Let $X = \{x \in Z : \pi(z) = 0\}$ denote the set of excess payoffs. Consider any $x \in X$ and $m \in M^{++}$. We have $E[mx] = 0$, hence $E[mx^+] = E[mx^-]$. But $E[mx^+] \geq \text{essinf}(m)E[x^+]$ and $E[mx^-] \leq \text{esssup}(m)E[x^-]$, therefore $E[x^+]/E[x^-] \leq \text{esssup}(m)/\text{essinf}(m)$. Since this holds for any $x \in X$ and $m \in M^{++}$, we have:

$$\max_{x \in X} \frac{E[x^+]}{E[x^-]} \leq \min_{m \in M^{++}} \frac{\text{esssup}(m)}{\text{essinf}(m)}. \quad (12)$$

To prove equality, let us consider a gain-loss efficient portfolio x° . It is a portfolio in X whose gain-loss ratio is maximal. Thus: $E[x^{\circ+}]/E[x^{\circ-}] = \max_{x \in X} E[x^+]/E[x^-]$. The existence of this asset and uniqueness of this portfolio (up to positive scalar multiplication, when no asset is redundant) are proven by Bawa and Lindenberg (1977), along with the ‘‘CAPM’’-like beta pricing equation:

$$E[z] - (1+r)\pi(z) = \frac{\text{Cov}[z, \text{sign}(x^\circ)]}{\text{Cov}[x^\circ, \text{sign}(x^\circ)]} E[x^\circ], \quad (13)$$

where z denotes a payoff with price $\pi(z)$ (see Section 5.2 for more details). As in the CAPM, only systematic risk matters, except that here it is measured by covariance with the *sign* of x° ,

⁵The essential infimum of a random variable m is equal to $\sup\{a \mid \Pr(m \leq a) = 0\}$, and the essential supremum is $\inf\{b \mid \Pr(m \leq b) = 1\}$. They work like the usual sup and inf, except that they ignore zero-probability events.

instead of x° itself. Equation (13) can be rewritten as:

$$\pi(z) = \frac{1}{1+r} \left\{ E[z] - \frac{E[z \operatorname{sign}(x^\circ)] - E[z]E[\operatorname{sign}(x^\circ)]}{\operatorname{Cov}[x^\circ, \operatorname{sign}(x^\circ)]} E[x^\circ] \right\} \quad (14)$$

$$= \alpha E[z] + \beta E[z \operatorname{sign}(x^\circ)] \quad (15)$$

for some scalars α and β that depend on r and x° but *not* on z . Rewriting Equation (15), there exists scalars \bar{m}° and \underline{m}° such that:

$$\pi(z) = E \left[z \left(\bar{m}^\circ \mathbf{1}_{\{x^\circ < 0\}} + \underline{m}^\circ \mathbf{1}_{\{x^\circ \geq 0\}} \right) \right], \quad (16)$$

where $\mathbf{1}_{\{\cdot\}}$ denotes the indicator function of an event. It implies that the discount factor

$$m^\circ = \begin{cases} \bar{m}^\circ & \text{if } x^\circ < 0 \\ \underline{m}^\circ & \text{if } x^\circ \geq 0 \end{cases} \quad (17)$$

correctly prices all the assets in the economy, i.e. $m^\circ \in M^{++}$. In particular, m° correctly prices x° , thus:

$$\bar{m}^\circ E[x^{\circ+}] - \underline{m}^\circ E[x^{\circ-}] = 0. \quad (18)$$

It implies that \bar{m}° and \underline{m}° have the same sign, which must be positive, since $E[m^\circ] = 1/(1+r) > 0$. This guarantees that $m^\circ > 0$. Note that the gain-loss ratio of $-x^\circ$ cannot exceed the gain-loss ratio of x° , therefore $E[x^{\circ+}] \geq E[x^{\circ-}]$, and by Equation (18): $\bar{m}^\circ \geq \underline{m}^\circ$. The identity

$$\frac{\operatorname{esssup}(m^\circ)}{\operatorname{essinf}(m^\circ)} = \frac{\bar{m}^\circ}{\underline{m}^\circ} = \frac{E[x^{\circ+}]}{E[x^{\circ-}]} \quad (19)$$

shows that there has to be an equal sign in Equation (12).

Now let us generalize the result to any non-constant m^* . If we apply the above reasoning to the risk-adjusted probability measure, we obtain:

$$\max_{\substack{x \in Z \\ \pi(x)=0}} \frac{E^*[x^+]}{E^*[x^-]} = \min_{\hat{m} \in \hat{M}^{++}} \frac{\operatorname{esssup}(\hat{m})}{\operatorname{essinf}(\hat{m})} \quad (20)$$

where $\hat{M}^{++} = \{\hat{m} : \hat{m} > 0 \text{ and } \forall z \in Z E^*[\hat{m}z] = \pi(z)\}$. Note that $m \in M^{++}$ if and only if

$\widehat{m} = E[m^*] m/m^* \in \widehat{M}^{++}$. Therefore

$$\min_{\widehat{m} \in \widehat{M}^{++}} \frac{\text{esssup}(\widehat{m})}{\text{essinf}(\widehat{m})} = \min_{m \in M^{++}} \frac{\text{esssup}(E[m^*] m/m^*)}{\text{essinf}(E[m^*] m/m^*)} = \min_{m \in M^{++}} \frac{\text{esssup}(m/m^*)}{\text{essinf}(m/m^*)}. \quad (21)$$

Bringing Equations (20) and (21) together proves Theorem 2. \square

Restriction 1 is thus equivalent to the existence of an admissible pricing kernel m such that:

$$\frac{\text{esssup}(m/m^*)}{\text{essinf}(m/m^*)} \leq \bar{L}. \quad (22)$$

It rules out state prices that are too far away from the benchmark investor's willingness to pay for such state claims.⁶ Intuitively, large deviations would be arbitrated away by buying (selling) relatively cheap (expensive) state claims. The following proposition shows another way to express this.

Proposition 5 *Restriction 1 is equivalent to assuming that the benchmark marginal rate of substitution across states is wrong at worst by a factor \bar{L} :*

$$\forall \omega_1, \omega_2 \in \Omega \quad \frac{1}{\bar{L}} \leq \frac{m(\omega_1)/m(\omega_2)}{m^*(\omega_1)/m^*(\omega_2)} \leq \bar{L}, \quad (23)$$

where m denotes an admissible pricing kernel.

Proof of Proposition 5 Ignoring zero-probability events:

$$\frac{\text{esssup}(m/m^*)}{\text{essinf}(m/m^*)} \leq \bar{L} \iff \forall \omega_1, \omega_2 \in \Omega \quad \frac{m(\omega_1)/m^*(\omega_1)}{m(\omega_2)/m^*(\omega_2)} \leq \bar{L} \quad (24)$$

$$\iff \forall \omega_1, \omega_2 \in \Omega \quad \frac{m(\omega_1)/m(\omega_2)}{m^*(\omega_1)/m^*(\omega_2)} \leq \bar{L} \quad (25)$$

$$\iff \forall \omega_1, \omega_2 \in \Omega \quad \frac{1}{\bar{L}} \leq \frac{m(\omega_1)/m(\omega_2)}{m^*(\omega_1)/m^*(\omega_2)} \leq \bar{L}. \quad \square \quad (26)$$

⁶This is an alternative to Hansen and Jagannathan's (1996) measure of distance between the set of non-negative admissible pricing kernel and the benchmark pricing kernel. One important difference is that our measure penalizes heavily pricing kernels with zero state prices (because they allow arbitrage opportunities), whereas theirs does not. On the other hand, their measure is easier to estimate by conventional econometric methods.

This shows that the gain-loss ratio restriction has an intuitive interpretation in terms of the stochastic discount factor.

4 Implications for Asset Prices

In this section we demonstrate how to obtain pricing bounds on a new security. We assume that markets are incomplete so that, in general, the new security cannot be priced exactly by absence of arbitrage. Instead, we obtain price bounds on the new security by computing all prices that are consistent with the absence of very attractive investment opportunities as measured by the gain-loss ratio.

4.1 Modeller's Choices

Our methodology involves several choices that must be made by the modeller *ex-ante* in order to derive pricing implications.

4.1.1 Benchmark Pricing Kernel

In order to obtain pricing bounds on a new security which are consistent with the absence of high gain-loss ratios the modeller must choose an appropriate benchmark pricing kernel. One approach is to specify a utility function u and a consumption plan c , in which case the benchmark pricing kernel is simply proportional to the marginal utility of consumption evaluated at $c(\omega)$, i.e. $u'(c(\omega))$. For example, we could assume that $u(c) = c^{1-\gamma}/(1-\gamma)$ so that $m^*(\omega)$ is proportional to $c^{-\gamma}$ where c is the representative agent's consumption. However, it is well-known that such pricing kernels implied by representative agent models with time additive utility do not do a good job of explaining observed stock prices when using aggregate consumption data (e.g. the equity premium puzzle of Mehra and Prescott, 1985). The modeller can choose some market index to proxy for the aggregate endowment if it fits observed data better. It is important to remember that the benchmark pricing kernel need not price all assets exactly: the ability to choose $\bar{L} > 1$ gives the modeller the flexibility to accommodate pricing errors of the benchmark pricing kernel.

Furthermore, the benchmark pricing kernel does not necessarily have to be inspired by theory; rather, it can also be obtained from some parametric (e.g. Backus, Foresi and Telmer, 1996) or non-parametric (e.g. Bansal and Viswanathan, 1993; Aït-Sahalia and Lo, 1995) estimation technique. To ensure the absence of arbitrage opportunities the modeller should restrict the benchmark pricing kernel to be strictly positive.

An important consideration for the choice of benchmark pricing kernel is its applicability to the investor in question. Consider, for example, the problem of obtaining bounds on the value to an executive of a stock option on the firm's stock. The executive may not hold a diversified portfolio because of insider trading restrictions, in particular short-sale constraints, and may have considerable human capital tied up in the firm; thus it would be inappropriate to use a benchmark pricing kernel which does not price idiosyncratic risk.

4.1.2 Ceiling on the Maximum Gain-Loss Ratio

If the modeller imposes the ceiling $\bar{L} = 1$ on the gain-loss ratio then the new security will be priced as if there was equilibrium with a representative agent whose marginal utility of consumption is described by the benchmark pricing kernel. It will yield an exact, but perhaps not credible, price for the new security. On the other hand, as the ceiling on the gain-loss ratio tends to infinity the price bounds for the new security will converge to the no-arbitrage bounds. These bounds will be credible, but imprecise. Consequently, the modeller must choose \bar{L} to control the trade-off between credibility and precision or, alternatively, the middle ground between the pricing implications of the equilibrium model and no-arbitrage.

4.1.3 Basis Assets

If we are interested in obtaining price bounds on a new security our gain-loss methodology does not require the use of any information about the price of other assets in the economy. Once the benchmark pricing kernel is chosen, we can compute all prices for the security which do not make investing in this security too attractive in the sense of permitting high gain-loss ratios. This can be called *absolute* pricing.

However, our methodology is more powerful for *relative* pricing. The price bounds can be

made much tighter by recording the prices of other assets, called basis assets, and imposing the restriction that high gain-loss ratio portfolios cannot be created from these basis assets and the new security. Thus, if there exists a security with known price and highly correlated payoffs with the new security then sharp pricing bounds can be achieved by including this security in the set of basis assets. For example, suppose that there are two securities with identical payoffs except that they differ slightly in one state which occurs with small probability. Clearly, if the prices of these two securities are too different then there will exist a great investment opportunity, so the price of one security will be very valuable information for determining reasonable prices of the other.

An important issue is that the benchmark pricing kernel used should do reasonably well at pricing the basis assets chosen. If not, it may be possible to create high gain-loss ratio portfolios using only the basis assets. If the gain-loss ratios attainable from the basis assets exceeds the ceiling \bar{L} then the restriction is not credible for the new security. Consequently, the modeller must choose an upper bound that exceeds the gain-loss ratios already attainable with the basis assets.

Another important issue is that the modeller should only include basis assets which are available to the investor in question. For example, suppose we are interested in obtaining bounds on the value of a stock option to a corporate executive. If the executive is not permitted to buy and sell the firm's stock freely or is not permitted to sell the firm's stock short at all, then it is inappropriate to include the firm's stock as a basis asset because the executive cannot freely form portfolios of the stock and option. On the other hand, the executive may be able to sell a market index short to hedge the firm's market risk, and so an index option may be an appropriate asset to include in the analysis. For another example, consider the problem of pricing a real option on an asset which is not traded. There may exist securities which are traded and are correlated, albeit not perfectly, to the value of the untraded underlying security. To get sharp pricing bounds, the modeller would want to include those securities in the set of basis assets.

4.2 Example: Pricing Options on Non-Traded Assets

In this section, we show that the gain-loss ratio restriction implies pricing bounds for options on an asset that does not trade. This is a problem of practical importance which is not addressed satisfactorily by existing asset pricing methods. The Black-Scholes (1973) dynamic replication argument does not work since the underlying is not traded, thus the option cannot be priced by arbitrage. Equilibrium models such as Rubinstein's (1976) require exact knowledge of the utility function, which is difficult to obtain. Sharpe ratio restrictions suffer from the non-normality of the payoff distribution in ways that were enumerated in Section 2.2, and will be further clarified in Section 5.1.1. In summary, this is a good place to start to demonstrate the power of a restriction on the maximum gain-loss ratio.

Suppose that we want to price a European call option with striking price K and a given maturity. Even though the underlying is not traded, we assume there exists a traded asset that is correlated with the underlying. It is reasonable to assume that the prices of European call options on this traded asset are known, either because they are themselves traded, or because they can be dynamically replicated with the traded asset. The key idea is to find out what the prices of options on the traded asset imply for the prices of options on the non-traded asset. Without loss of generality, we can rescale both assets so that they have the same value today, making it meaningful to compare striking prices. With this convention, we take as the collection of basis assets the riskless bond and the option on the traded asset with the same striking price K and the same maturity as the option on the non-traded asset that we wish to price.

We assume that the two assets are jointly lognormally distributed. There exists a natural choice for the benchmark pricing kernel: the Black-Scholes pricing kernel. Rubinstein (1976) shows that it can be obtained as an equilibrium pricing kernel by assuming Constant Relative Risk Aversion (CRRA) utility. Thus, an attractive investment opportunity will be defined as a security underpriced relatively to this benchmark equilibrium model. Even though the Black-Scholes formula is not always applicable, it always provides a useful benchmark.

The tightness of the pricing bound is controlled by the parameter \bar{L} , which is the maximum gain-loss ratio in the economy. For example, setting $\bar{L} = 1$ would reduce to the Rubinstein (1976) equilibrium pricing model, therefore it would allow only the Black-Scholes price. But we have

no reason to assume that agents have exactly CRRA utility, therefore we choose a value of \bar{L} strictly above one. The higher the value of \bar{L} , the more reliable and the looser the bounds. As \bar{L} increases, we get closer to the no-arbitrage bounds, which only say that the option price must lie somewhere between zero and infinity. We choose the value $\bar{L} = 2$ to capture our degree of confidence in the benchmark. It means that the CRRA agent cannot receive twice as much gain as would be necessary for her to hold the asset. Or, that the marginal rate of substitution *across states* implied by the Black-Scholes pricing kernel is never more than twice (or less than half) the true one. The intuition is that the Black-Scholes formula must be taken with a grain of salt when its assumptions are not verified.

Without loss of generality, assume that the option expires at time one. Call S_n (S_t) the final value of the non-traded (traded) asset, and S_0 their common initial value. The continuously compounded rates of return on these assets are jointly distributed as:

$$\begin{bmatrix} \log(S_n/S_0) \\ \log(S_t/S_0) \end{bmatrix} \overset{*}{\sim} N \left(\begin{bmatrix} \log(1+r) - \frac{\sigma_n^2}{2} \\ \log(1+r) - \frac{\sigma_t^2}{2} \end{bmatrix}, \begin{bmatrix} \sigma_n^2 & \rho\sigma_n\sigma_t \\ \rho\sigma_n\sigma_t & \sigma_t^2 \end{bmatrix} \right) \quad (27)$$

where $\overset{*}{\sim}$ denotes distribution under the benchmark risk-adjusted probability measure, $N(\cdot)$ denotes the bivariate normal distribution, σ_n (σ_t) denotes the volatility of the non-traded (traded) asset, and ρ is the correlation coefficient. The price π_t of the call with striking price K on the traded asset is given by the Black-Scholes formula:

$$\pi_t = S_0 \Phi(d) - \frac{K}{1+r} \Phi(d - \sigma_t) \quad \text{where } d = \frac{\log(S_0/K) + \log(1+r)}{\sigma_t} + \frac{1}{2}\sigma_t \quad (28)$$

and where Φ denotes the standard normal cumulative distribution function.

Now consider the economy consisting of the riskless bond with payoff $1+r$ and price 1, the option on the traded asset with payoff $\max(S_t - K, 0)$ and price π_t , and the option on the non-traded asset with payoff $\max(S_n - K, 0)$ and price π_n . The last quantity, π_n , is the only unknown. For each value of π_n , let $X(\pi_n)$ denote the space of excess payoffs on portfolios of

these three assets:

$$X(\pi_n) = \left\{ w_1 + w_2 \max(S_t - K, 0) + w_3 \max(S_n - K, 0) : \frac{w_1}{1+r} + w_2 \pi_t + w_3 \pi_n = 0 \right\}. \quad (29)$$

We can compute the portfolio with highest gain-loss ratio by solving $\max_{x \in X(\pi_n)} E^*[x^+]/E^*[x^-]$. We call it the gain-loss efficient portfolio (see Section 5.2 for more details on gain-loss efficiency). The composition of this portfolio obviously depends on the price of the option on the non-traded asset. In particular, if π_n was extremely low (high) then the efficient portfolio would be heavily tilted towards holding (writing) this option. Either way, there would exist an attractive investment opportunity in the sense that the gain-loss efficient portfolio would have high gain-loss ratio. Therefore there exists only a narrow domain for the option price π_n that does not create investment opportunities that are too attractive. This domain defines the pricing bounds for the option on the non-traded asset. Formally, Restriction 1 implies that:

$$\min \left\{ \pi_n : \max_{x \in X(\pi_n)} \frac{E^*[x^+]}{E^*[x^-]} \leq \bar{L} \right\} \leq \pi_n \leq \max \left\{ \pi_n : \max_{x \in X(\pi_n)} \frac{E^*[x^+]}{E^*[x^-]} \leq \bar{L} \right\}. \quad (30)$$

There exists a dual formulation in terms of pricing kernels. Let m^* denote the benchmark pricing kernel. Consider the set of the pricing kernels that are positive, price the basis asset correctly, and do not differ from the benchmark by more than a factor \bar{L} :

$$M_{\bar{L}}^{++} = \left\{ m > 0 : E[m(1+r)] = 1, E[m \max(S_t - K, 0)] = \pi_t \text{ and } \frac{\text{esssup}(m/m^*)}{\text{essinf}(m/m^*)} \leq \bar{L} \right\}. \quad (31)$$

If we restrict our attention to such pricing kernels, we obtain pricing bounds for the option on the non-traded asset as:

$$\min_{m \in M_{\bar{L}}^{++}} E[m \max(S_n - K, 0)] \leq \pi_n \leq \max_{m \in M_{\bar{L}}^{++}} E[m \max(S_n - K, 0)]. \quad (32)$$

Using Theorem 2, it is easy to show that the bounds in Equations (30) and (32) are the same.

Figure 1 plots these bounds as a function of the striking price K for particular values of the parameters: the riskless interest rate is $r = 0.05$; both traded and non-traded assets have the same initial price $S_0 = 100$, and the same volatility $\sigma_n = \sigma_u = 0.20$, which is calibrated to imitate

a stock or stock index at the one-year horizon; and the correlation coefficient between the two assets is 0.70. These bounds lie strictly between the Black-Scholes price and the no-arbitrage

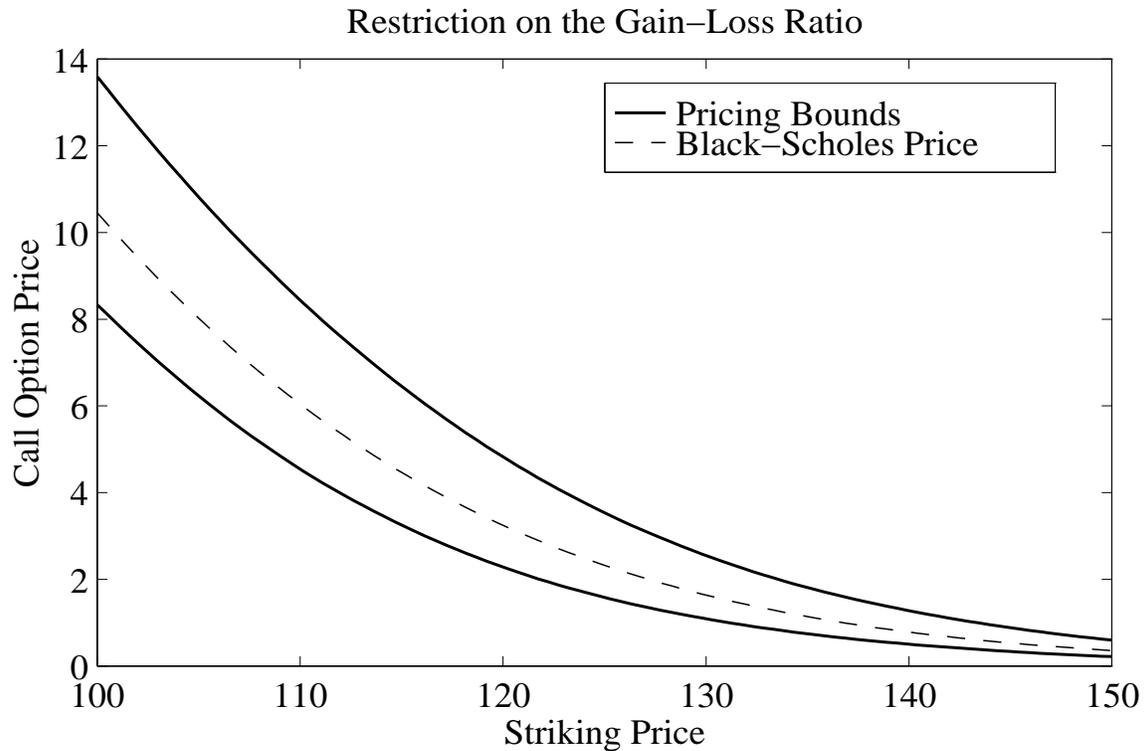


Figure 1: Bounds on the price of a call option on a non-traded stock.

bounds (zero and infinity in this case). They get tighter as: 1) the maximum gain-loss ratio goes down, 2) the correlation between traded and non-traded assets goes up, or 3) the volatility of the non-traded asset decreases. In the first case it is because the benchmark model is more reliable, and in the other two it is because the basis assets allow better hedging. These bounds represent what we know if the Black-Scholes formula is not too false and if attractive investment opportunities are ruled out.

4.3 Applications

In this section we briefly review some practical asset pricing problems that can be solved using the gain-loss ratio restriction.

4.3.1 Valuing Real Options

An important application of the gain-loss approach is to value real options. Real options capture the flexibility that managers often have when making capital budgeting decisions such as (i) the option to wait to invest rather than invest now; (ii) the option to abandon a failing project; and (iii) the option to make follow-on investments. Arbitrage-based option pricing models, such as Black-Scholes, are often used to value real options. This approach is valid if the stochastic component of the options return can be replicated over time by trading in some set of basis assets. An important feature of real options, however, is that the stochastic component of the options return often cannot be replicated because the underlying asset does not exist, does not trade, trades in an illiquid market, or is not spanned by a portfolio of traded assets. If the risk associated with some state variable cannot be traded the modeller can introduce an estimate of the market value of this non-traded risk and then value the option as if by arbitrage. This risk premium, however, should come from an equilibrium model (see Schwartz, 1994).

The gain-loss approach allows the modeller to specify a particular equilibrium model and compute price bounds on the value of the real option consistent with the absence of very attractive investment opportunities. Furthermore, the modeller has the flexibility to introduce basis assets which yield sharper pricing bounds for the real option. Thus, if one can construct an imperfect hedging strategy by using some combination of existing assets then our gain-loss restriction yields bounds consistent with the inability to construct extremely attractive portfolios using these basis assets. To apply our example in Section 4.2 to a real option example, we can interpret the basis asset as an imperfect hedging instrument for some non-traded risk and compute pricing bounds accordingly.

4.3.2 Valuing Executive Stock Options

Another example for which the dynamic replication argument implicit in the Black-Scholes approach fails is the valuation of executive stock options. Executives cannot freely trade in the underlying stock, due to insider trading restrictions, nor can they trade the options due to contractual restrictions; consequently, dynamic replication arguments are unrealistic in this case. Thus, the gain-loss approach can be used to obtain price bounds. To get useful price bounds the mod-

eller should (i) account for the fact that the pricing kernel which is relevant to the executive prices idiosyncratic risks which he cannot diversify away; and (ii) include in the analysis, as we did in the above example, the ability of the executive to trade some assets which are imperfect hedges for some firm risks, e.g. index options.

4.3.3 Performance Evaluation

Benchmarking for the risk of a portfolio is a crucial aspect of evaluating mutual fund performance. Many studies employ models such as the CAPM or APT to determine the riskiness of a portfolio.⁷ For example, Jensen's alpha, the difference between the average return on a portfolio and the return implied by the CAPM for a portfolio with the same beta, is often used as a measure of a portfolio manager's ability to achieve a superior risk-return tradeoff. If portfolio returns are normally distributed then Jensen's alpha is a reasonable measure of a portfolio manager's ability to select individual stocks if they are not market timers (Grinblatt and Titman, 1989).⁸ However, many mutual funds use derivatives and have return distributions that are not normally distributed (Koski and Pontiff, 1996). In such cases, measures such as Jensen's alpha can be misleading.

The gain-loss approach can be used to evaluate performance when portfolio returns are not normally distributed. Moreover, gain-loss yields an economically meaningful way to evaluate performance: according to the attractiveness of investing in the portfolio for some benchmark investor. Suppose, for example, that an investor holds an endowment which is proxied by some market index, e . If we assume that the investor has a utility function in the constant relative risk aversion class, $u(e) = \frac{1}{1-\gamma}e^{1-\gamma}$ then there exists a relative risk aversion parameter, γ , which correctly prices the index. We can now evaluate portfolio performance by measuring its gain and loss under the benchmark pricing kernel implied by (u_γ, e) : if the gain exceeds the loss then the portfolio is an attractive investment for the benchmark investor.⁹ A potential problem with this

⁷See Grinblatt and Titman (1995) for a survey.

⁸Portfolio managers are defined to be market timers if they change the beta of their portfolio when they believe the expected return on the market portfolio will change.

⁹The gain-loss approach is related to the positive period weighting measure in Grinblatt and Titman (1989) when their weights are interpreted as marginal utilities. The gain-loss ratio provides a useful metric for evaluating whether portfolio managers have valuable stock picking information when returns are not normally distributed. However, if portfolio managers are also market timers then the gain-loss measure may not be appropriate outside of normality. These issues will be left to future research.

approach is that the benchmark pricing kernel may price the market index but may not correctly price various portfolio styles such as small size and high book-to-market portfolios. This can be remedied by including such portfolios in the set of basis assets and imposing the condition that the benchmark pricing kernel also price these portfolios correctly.

5 Links to Earlier Literature

The gain-loss ratio restriction was developed in Section 3 in an effort to find an alternative to the Sharpe ratio restriction that would take preferences into account and be consistent with the no-arbitrage principle outside normality. It would be surprising if something with such a simple expression was not related to some earlier work. In this section, we show that it is related to *three* different strands of literature (whose relation to one another had not been previously noted). This is reassuring: we take it as an additional indication that the gain-loss ratio makes economic sense. For simplicity we ignore the difference between the true probability and the benchmark risk-adjusted probability throughout this section.

5.1 Other Restrictions on the Stochastic Discount Factor

Since Hansen and Jagannathan (1991) obtained a duality result based on mean-variance, several other duality results have appeared in the literature. Each one of them corresponds to a different restriction on the stochastic discount factor. Even though these restrictions are all presented as “intuitive”, it is sometimes hard to figure out what they *really* mean, how they are related to one another, and what differentiates them. We propose an original classification of these restrictions that sheds light on these issues. Our classification is based on whether a given restriction prevents state prices (per unit of probability) from being too *high* or from being too *low*.

5.1.1 Restricting High State Prices

The first category contains the restrictions that, in effect, prevent state prices (per unit of probability) from getting too close to infinity. Intuitively, an infinite state price means that you can get

something today for nothing tomorrow: it is an arbitrage opportunity. Therefore it makes sense to rule out state prices that are infinite, and by continuity those that are very high too.

Several restrictions fall into this category, first among which the Hansen-Jagannathan (1991) variance bound. In the presence of a riskfree bond, the expectation of the stochastic discount factor is pinned down by $E[m(1+r)] = 1$, therefore restricting the variance is the same as restricting the second moment. Snow (1991) generalizes this approach by deriving restrictions on the q^{th} moment of the stochastic discount factor $E[m^q]^{1/q}$, for $1 < q < \infty$. By letting q go to infinity, we obtain a restriction on the essential supremum of the stochastic discount factor: $\text{esssup}(m)$. This is *one part* of the dual of the gain-loss ratio restriction derived in Theorem 2 (the other part being a restriction on the essential infimum), up to normalization by the benchmark pricing kernel m^* . The last restriction in this category is due to Stutzer (1993), who shows that restricting the maximum expected utility attainable by a CARA agent is equivalent to restricting the entropy of the pricing kernel $E[m \log(m)]$.

The problem of the restrictions in this category is that they do not prevent state prices from being too close to zero. As a matter of fact, they can allow some state prices to be exactly *equal* to zero. This is not economically meaningful, since it allows arbitrage opportunities. This is a generalization of the criticism of Section 2.2.1, where we argued that the Sharpe ratio restriction is not compatible with the no-arbitrage principle: indeed, neither are the restrictions of Snow (1991) and Stutzer (1993). To illustrate the consequences of this problem, we use the call option example from Section 4.2. The option pricing bounds implied by a restriction of the Stutzer (1993) type are plotted in Figure 2. The setup is the same as in Figure 1. In addition, we assumed that both assets have continuously compounded returns with the same expectation $\mu = 0.10$ under the true probability measure. The restriction is calibrated to give the same lower bound at-the-money as in Figure 1. We see that the price of the call can be zero (or at least arbitrarily close to zero) in some cases, which would open up an arbitrage opportunity. The corresponding graphs for the restrictions of Hansen and Jagannathan (1991) and Snow (1991) are not shown because they look similar and in particular suffer from the same problem.¹⁰

The reason for the problem can be restated in terms of payoffs, which might clarify where

¹⁰Assuming that we impose $m \geq 0$ in addition to restricting the q^{th} moment of m . Without the additional positivity constraint, it is even worse: the pricing bounds can go negative.

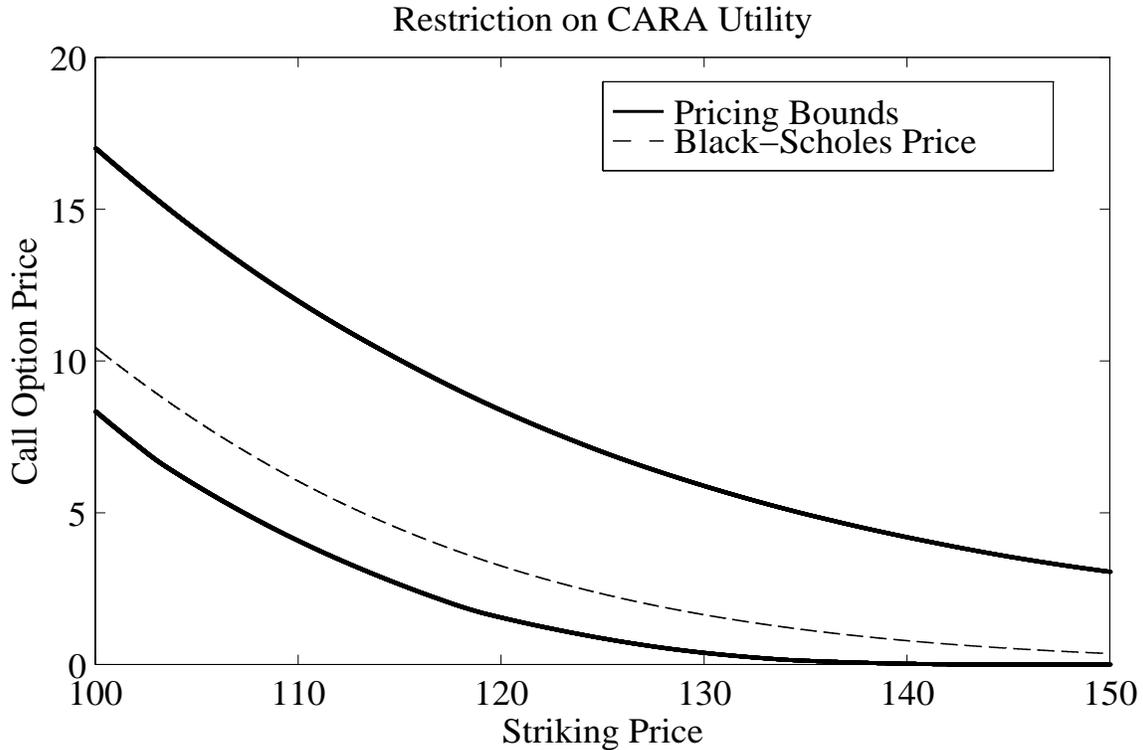


Figure 2: Bounds on the price of a call option on a non-traded stock. When the striking price is at 143 and above, the lower bound is *exactly* equal to zero: it is not due to numerical roundoff.

it comes from. Take, for example, Stutzer's (1993) bound. Remember that the maximum value that a CARA utility function can take is zero. Suppose that it is possible to get something in some states tomorrow for nothing today. Then, in those states, the bliss point is attained, which corresponds to zero utility. However, the expected utility is not zero because in the other states the bliss point is not attained. Thus, restricting the expected utility of a CARA agent allows arbitrage opportunities that pay off only in some states.

Contrast this with the situation when the utility function has no upper bound. In that case, if infinite wealth is attained with positive probability, then expected utility is infinite. Thus, a restriction on expected utility would have successfully ruled out the arbitrage opportunities that pay off only some states. The problem with the restrictions in this category is that they correspond to utility functions that have a finite upper bound.¹¹

¹¹Clearly, this is also the case with quadratic utility.

5.1.2 Restricting Low State Prices

The second category contains the restrictions that, in effect, prevent state prices (per unit of probability) from getting too close to zero. Intuitively, a zero state price means that you can get something tomorrow for nothing today: it is an arbitrage opportunity. Therefore it makes sense to rule out state prices that are zero, and by continuity those that are very low too.

Only one published restriction falls into this category: the Bansal-Lehmann (1996) logarithm bound. Bansal and Lehmann show that restricting the maximum expected utility that can be attained by an agent with logarithmic utility is equivalent to restricting $E[-\log(m)]$. We generalize their result to any CRRA utility function.

Theorem 3 *Let Z denote a space of payoffs, $\pi(\cdot)$ the pricing functional, and M^+ the set of non-negative admissible pricing kernels. Then we have:*

$$\max_{\substack{z \in Z \\ \pi(z)=1}} E \left[\frac{1}{1-\gamma} z^{1-\gamma} \right] \leq \min_{m \in M^+} \frac{1}{1-\gamma} E \left[m^{1-(1/\gamma)} \right]^\gamma \quad (33)$$

when $0 < \gamma < 1$ or $1 < \gamma < \infty$.

Proof of Theorem 3 The result relies on an extension of the Hölder inequality. If $z \geq 0$, $m \geq 0$ and $1/p + 1/q = 1$ then we have:

$$E[z^p]^{1/p} E[m^q]^{1/q} \leq E[mz], \quad (34)$$

when $-\infty < p < 0$ or $0 < p < 1$. Notice that the inequality gets reversed when the exponents p and q fall outside their usual range (see e.g. Hardy et al. (1964, Equation 2.8.4) for a reference). Take $p = 1 - \gamma$ and $q = 1 - (1/\gamma)$. For all $z \in Z$ such that $\pi(z) = 1$ and for all $m \in M^+$, we have $E[mz] = 1$, therefore:

$$\begin{aligned} E \left[z^{1-\gamma} \right]^{1/(1-\gamma)} E \left[m^{1-(1/\gamma)} \right]^{-\gamma/(1-\gamma)} &\leq 1 \\ E \left[z^{1-\gamma} \right]^{1/(1-\gamma)} &\leq E \left[m^{1-(1/\gamma)} \right]^{\gamma/(1-\gamma)} \\ E \left[\frac{1}{1-\gamma} z^{1-\gamma} \right] &\leq \frac{1}{1-\gamma} E \left[m^{1-(1/\gamma)} \right]^\gamma. \end{aligned} \quad (35)$$

Strictly speaking, Equation (35) requires $z \geq 0$, because for negative values the left hand side would be a complex number. But it makes more economic sense to define the CRRA utility of a negative payoff as minus infinity instead, and with this convention Equation (35) always holds. Since it holds for any unit-price payoff and any non-negative admissible pricing kernel, Theorem 3 follows. \square

By letting γ go to zero in Equation (33), we obtain a restriction on the essential infimum of the stochastic discount factor: $\text{essinf}(m)$. This is the *other part* of the dual of the gain-loss ratio restriction derived in Theorem 2, again up to normalization by the benchmark pricing kernel m^* .

The problem of the restrictions in this category is that they do not prevent state prices from being too close to infinity. As a matter of fact, they can allow some state prices to be exactly *equal* to infinity. This is not economically meaningful, since it allows arbitrage opportunities. This is the reverse of the criticism we had above. To illustrate the problem, we use again the call option example developed in Section 4. The option pricing bounds implied by a restriction of the Bansal-Lehmann (1996) type are plotted in Figure 3. The setup is the same as in Figures 1 and 2. We see that there is no upper bound: the price of the call can be infinite (or at least arbitrarily close to infinity), which would open up an arbitrage opportunity. The corresponding graphs for the other CRRA restrictions are not shown because they look globally the same and in particular suffer from the same problem.¹²

Again, the reason for the problem can be restated in terms of payoffs, which might clarify where it comes from. In the above example, suppose that the call option is overpriced. In order to take advantage of the mispricing, we would have to sell it short, thus accepting unlimited downside potential. No CRRA agent would accept that, because there would be some probability of having negative wealth, which would imply an expected utility of minus infinity. Therefore, restricting the expected utility of a CRRA agent allows options to be arbitrarily overpriced.

Contrast this with the situation where the utility function is defined on the *entire* real line. In that case, even if negative wealth is attained with positive probability, expected utility need not be equal to minus infinity. Thus, a restriction on expected utility can successfully rule out

¹²Strictly speaking, this statement is only correct for $0 < \gamma < 1$. When $1 < \gamma < \infty$, we get the worst of both worlds: the upper bound is infinite *and* the lower bound can be zero.

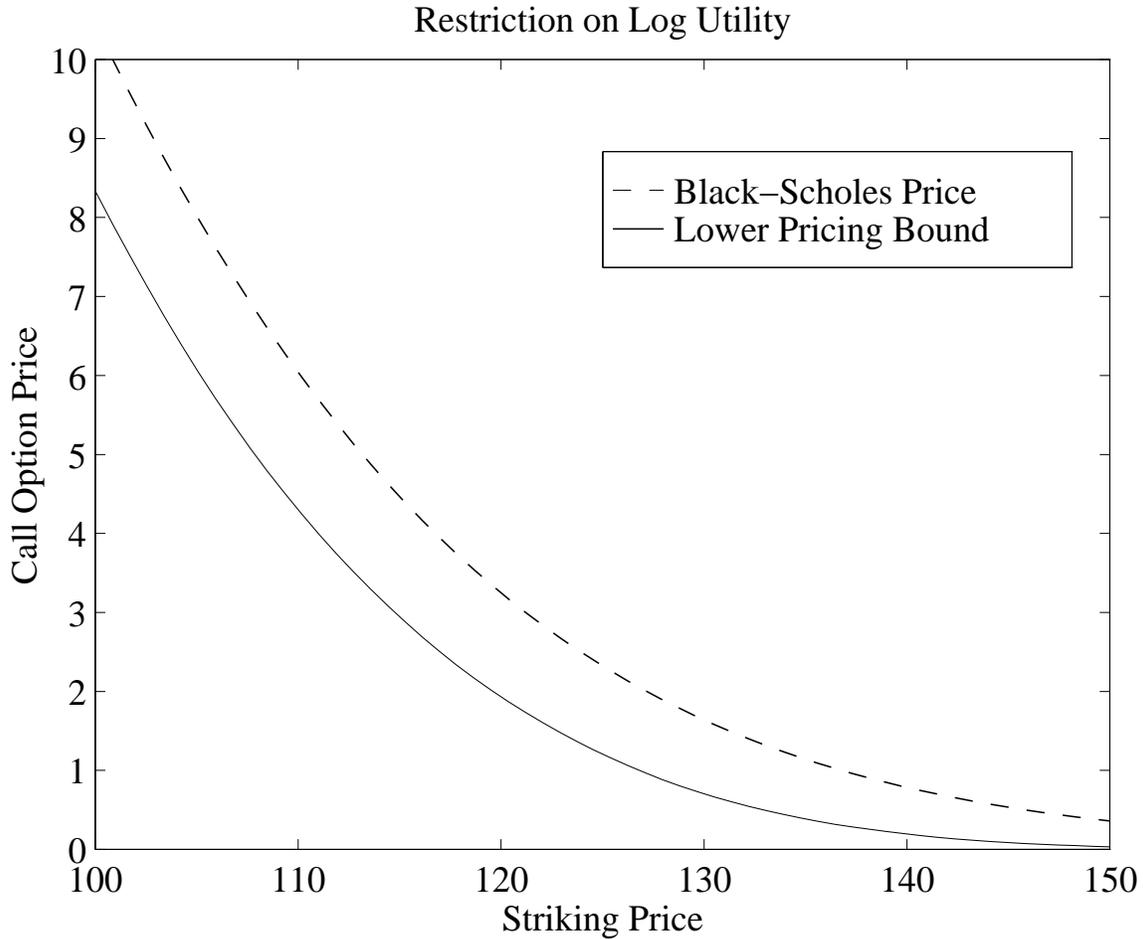


Figure 3: Lower bound on the price of a call option on a non-traded stock. The lower bound is always strictly positive, but the upper bound is infinity, i.e. the option price is allowed to be arbitrarily large.

option prices that are too high. The problem with the restrictions in this category is that they correspond to utility functions are only defined on the positive half-line.

5.1.3 Restricting High And Low State Prices Simultaneously

When we categorize the pricing kernel restrictions in the existing literature in this way, it is easy to see that we ought to prevent state prices (per unit of probability) from getting too close to infinity *and* to zero. Consider the restriction that we obtained in Theorem 2 as the dual of the gain-loss ratio restriction: one part of it is the limiting case of the generalization of the Hansen-Jagannathan (1991) variance bound for high state prices, and the other part is the limiting case of

the generalization of the Bansal-Lehmann (1993) logarithm bound for low state prices. Therefore our restriction is the first one to restrict high and low state prices simultaneously.

It is not the only one. Any such restriction could be constructed by considering a utility function that is defined on the entire real line and has no upper bound.¹³ Another possibility would be to combine a restriction of the first category with a restriction of the second category. Our restriction does not have better properties than these combinations. The only reason why we chose the gain-loss approach is because it is intuitive and has a simple explicit formulation both in terms of pricing kernel and in terms of payoffs, which cannot be said of the others. But in all generality our method for obtaining pricing bounds would work just as well using any restriction that prevents state prices from getting too high *and* too low (relative to some benchmark).

5.2 Alternatives to Mean-Variance

Mean-variance preferences can be seen as a special case of the mean- p^{th} moment preferences defined by:

$$\min_{\substack{x \\ E[x]=\mu}} E[|x|^p]. \quad (36)$$

Setting $p = 2$ yields mean-variance preferences, and $p = 1$ yields mean-first absolute moment preferences, which are equivalent to gain-loss preferences (if we ignore the change of probability measure). This shows that gain-loss analysis (under the true probability measure) is in the same general family as mean-variance analysis.

Another alternative to mean-variance analysis was suggested by Markowitz (1952) in his seminal paper. He defined the semivariance of excess payoff x as $E[(x^-)^2]$, and proposed mean-semivariance as an alternative to mean-variance. Motivated by stochastic dominance arguments, Bawa (1975) introduced a generalization called mean-lower partial moment of order p :

$$\min_{\substack{x \\ E[x]=\mu}} E[(x^-)^p]. \quad (37)$$

Again for $p = 1$ we obtain mean-loss preferences, which are equivalent to gain-loss preferences.

¹³Surprisingly, it is hard to find a familiar utility function with both properties.

It is interesting to note that $p = 1$ is the only exponent for which the families defined by Equations (36) and (37) intersect.

Bawa and Lindenberger (1977) derive all the properties of mean-lower partial moment of order p preferences, of which gain-loss is a special case. In particular, they show that it is equivalent to mean-variance preferences if returns are normally distributed. For convenience, we prove again their result below.

Proposition 6 *If the normally distributed excess payoff x has Sharpe ratio S , then its gain-loss ratio is:*

$$\frac{E[x^+]}{E[x^-]} = \frac{\phi(S) + S \Phi(S)}{\phi(-S) - S \Phi(-S)} \quad (38)$$

where ϕ (respectively Φ) denotes the standard normal probability density function (respectively cumulative distribution function).

Proof of Proposition 6 Let x be normally distributed with mean μ and variance σ^2 . We have:

$$\begin{aligned} E[x^+] &= \int_0^\infty \frac{1}{\sqrt{2\pi\sigma^2}} x e^{-(x-\mu)^2/(2\sigma^2)} dx \\ &= \int_{-\mu/\sigma}^\infty \frac{1}{\sqrt{2\pi}} (\mu + \sigma z) e^{-z^2/2} dz \\ &= \mu \int_{-\mu/\sigma}^\infty \frac{1}{\sqrt{2\pi}} e^{-z^2/2} dz + \sigma \int_{-\mu/\sigma}^\infty \frac{1}{\sqrt{2\pi}} z e^{-z^2/2} dz \\ &= \mu \int_{-\mu/\sigma}^\infty \frac{1}{\sqrt{2\pi}} e^{-z^2/2} dz + \sigma \left[-\frac{1}{\sqrt{2\pi}} e^{-z^2/2} \right]_{-\mu/\sigma}^\infty \\ &= \mu \left[1 - \Phi\left(-\frac{\mu}{\sigma}\right) \right] + \sigma \phi\left(-\frac{\mu}{\sigma}\right) \\ &= \mu \Phi\left(\frac{\mu}{\sigma}\right) + \sigma \phi\left(\frac{\mu}{\sigma}\right). \end{aligned}$$

Furthermore, the loss of x is equal to the gain of the short position $-x$, which is normally distributed with mean $-\mu$ and variance σ^2 , therefore: $E[x^-] = E[(-x)^+] = -\mu \Phi(-\mu/\sigma) + \sigma \phi(-\mu/\sigma)$. Bringing both formulas together shows that the gain-loss ratio of x is given by Equation (38). \square

Proposition 6 shows that, under normality, there is a one-to-one relationship between the Sharpe ratio and the gain-loss ratio. Therefore gain-loss preferences are exactly the same as mean-variance preferences in this case. This is to be expected, since any reasonable set of preferences

leads to mean-variance under normality. In particular, Restriction 1 on the maximum gain-loss ratio is equivalent to a restriction on the Sharpe ratio if returns are normal. We can use Equation (38) to map the value of the gain-loss ratio ceiling \bar{L} in Restriction 1 into the value of the Sharpe ratio ceiling \bar{S} imposed by Ledoit (1995) and Cochrane and Saá-Requejo (1996). Figure 4 graphs the functional relationship between them. Another consequence of Proposition 6 is that

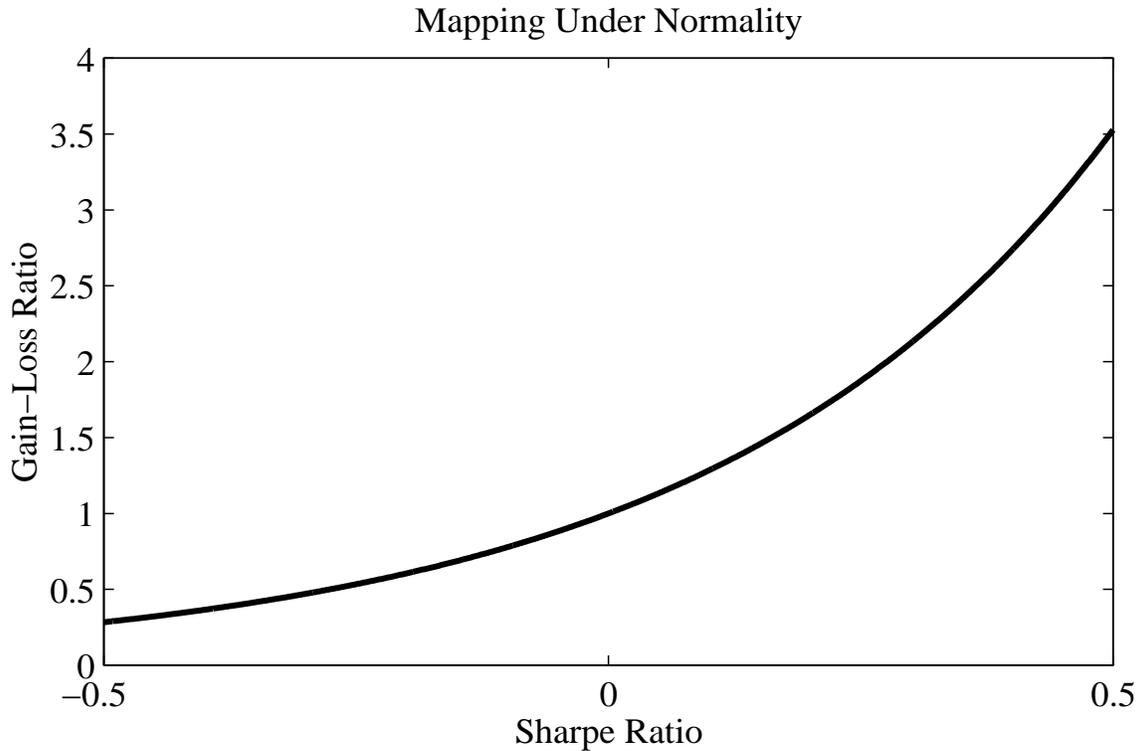


Figure 4: Mapping between the gain-loss ratio and the Sharpe ratio under normality. It has the same overall shape as the exponential function. For example, a stock market index with Sharpe ratio near 0.3 has gain-loss ratio near 2 (on an annual basis).

the gain-loss approach does not contribute anything new under normality. This is why, for stocks or for short trading intervals, the Sharpe ratio restriction is still perfectly adequate. We see the gain-loss ratio restriction as a way to generalize the Sharpe ratio restriction outside normality.

Bawa and Lindenberger (1977) also show that the set of attainable gains and losses forms a convex set in (gain, loss) space. Thus, there exists a well-defined gain-loss efficient frontier. In the presence of a riskless bond, the frontier is a half-line and two-fund separation holds, with the bond being one of the two funds. In this case, all gain-loss efficient portfolios have the same excess return up to positive scalar multiplication. Figure 5 shows a typical gain-loss efficient

frontier.

Let x^* denote the excess return on a gain-loss efficient portfolio. Bawa and Lindenberg (1977) prove a ‘‘CAPM’’-like beta pricing equation that we used in the proof of Theorem 2:

$$E[z] - (1 + r)\pi(z) = \frac{\text{Cov}[z, \text{sign}(x^*)]}{\text{Cov}[x^*, \text{sign}(x^*)]}E[x^*], \quad (39)$$

where z denotes a payoff with price $\pi(z)$. As in the CAPM, only systematic risk matters, except that here it is measured by covariance with the *sign* of x^* , instead of x^* itself. Therefore gain-loss analysis forms a fully developed alternative to mean-variance analysis.

It is interesting to rewrite Equation (39) as:

$$\frac{E[x\mathbf{1}_{\{x^* > 0\}}]}{E[x\mathbf{1}_{\{x^* < 0\}}]} = \frac{E[x^{*+}]}{E[x^{*-}]}, \quad (40)$$

where $\mathbf{1}_{\{\cdot\}}$ denotes the indicator function of an event. $E[x\mathbf{1}_{\{x^* > 0\}}]$ can be interpreted as systematic gain, and $E[x\mathbf{1}_{\{x^* < 0\}}]$ as systematic loss. Thus, systematic gain is proportional to systematic loss, and the coefficient of proportionality is the gain-loss ratio of a gain-loss efficient portfolio. The idea that only systematic risk matters is familiar, but the idea that only *systematic reward* matters is more original. It did not appear in mean-variance analysis because, there, the measure of reward (i.e. the mean) is a linear functional.

5.3 Prospect Theory

Kahnemann and Tversky (1979) report that the axioms of expected utility are violated in controlled experiments. They summarize their experimental evidence into what they call ‘‘prospect theory’’. The main point of prospect theory is that gains relative to a break-even reference point are treated differently from losses. This is very close in spirit to gain-loss preferences.

More formally, Benartzi and Thaler (1995) argue that the key features of prospect theory are captured by a piecewise linear utility function $u(x) = ax^+ + bx^-$, where $0 < a < b$ are parameters, and x denotes excess payoff relative to the break-even reference point (see Figure 6). It is easy to see that this simplified version of prospect theory is equivalent to gain-loss preferences (once again, if we ignore the change of probability measure). Thus, the gain-loss

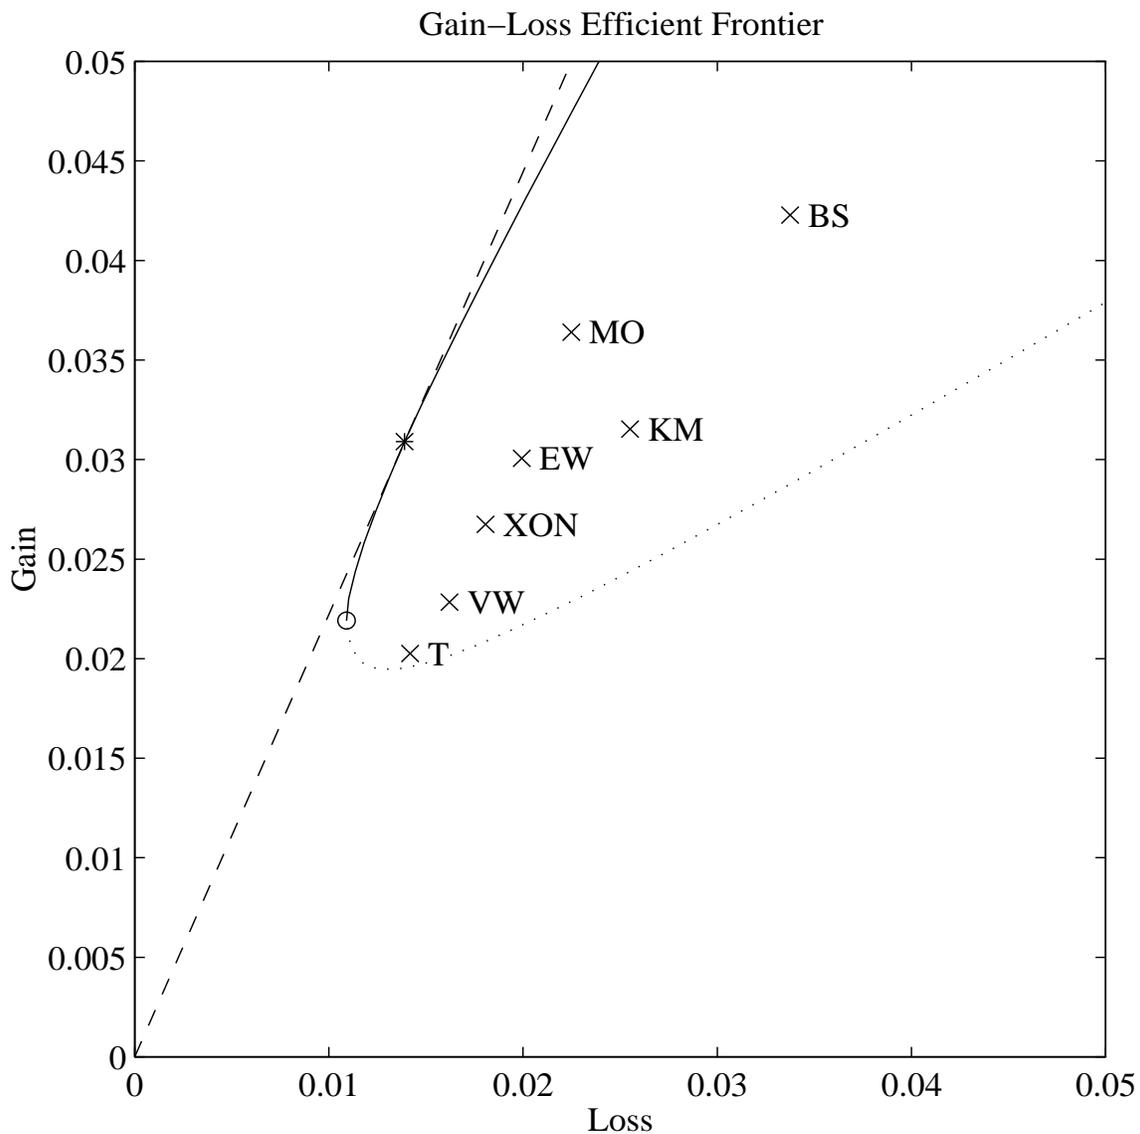


Figure 5: Gain-Loss Efficient Frontier. Data come from the CRSP database and cover 70 years from 1926 to 1995. The assets are the CRSP value-weighted (VW) and equal-weighted (EW) indices, and five stocks: AT&T (T), Bethlehem Steel (BS), K-Mart (KM), Philip Morris (MO), and Exxon (XON). Gain and loss are computed from monthly returns in excess of the CRSP riskfree rate. The solid (dashed) line is the ex-post gain-loss efficient frontier without (with) a riskless asset. The dotted line is the inefficient part of the frontier with no riskless asset. The portfolio of stocks with the highest gain-loss ratio is represented by a star (*), and the one with the lowest loss by a circle (o).

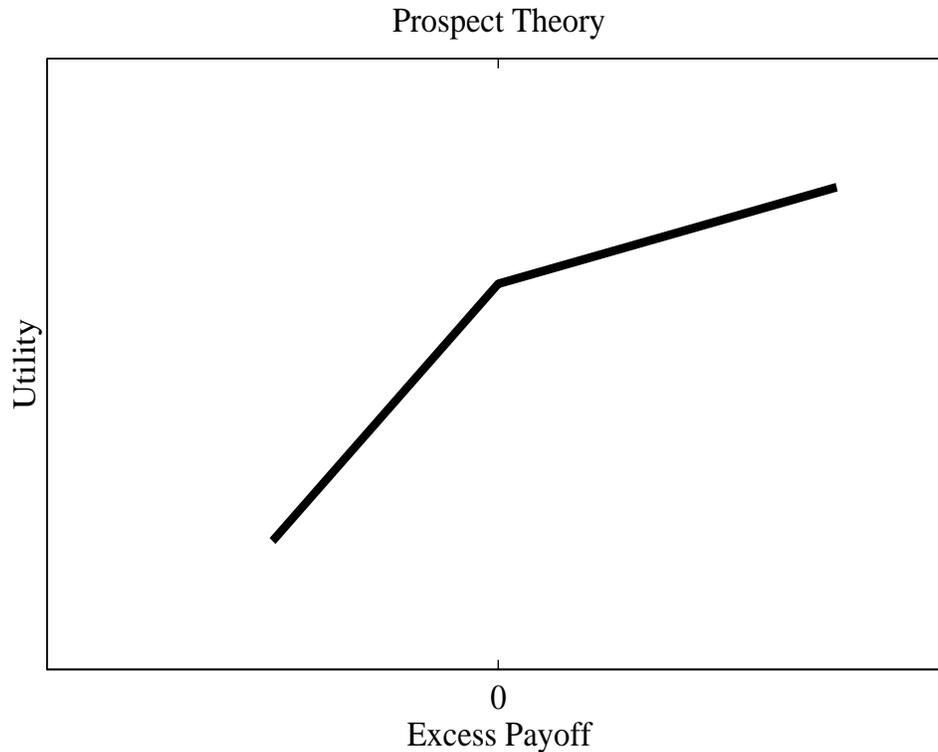


Figure 6: Profile of the utility function that summarizes prospect theory.

approach is supported by evidence gathered in experimental economics.

6 Conclusion

Ledoit (1995) and Cochrane and Saá-Requejo (1996) develop a methodology that charts the middle ground between the two paradigms of asset pricing: equilibrium and no-arbitrage. Their idea is to restrict the maximum Sharpe ratio in the economy. It is primarily designed to work well under normality. There exists an infinite number of ways to generalize their approach to non-normal returns, and restricting the Sharpe ratio itself is but one of them. We show that it is not a satisfactory one because high Sharpe ratio portfolios do not form a neighborhood of the set of arbitrage opportunities, and because the maximum Sharpe ratio is unbounded in some (non-pathological) equilibrium models.

Thus we propose a different way to generalize the Sharpe ratio restriction outside normality. It is to restrict the maximum *gain-loss* ratio in the economy, defined as the ratio of the expectation

of the positive part to the expectation of the negative part of the excess payoff, taking expectations under a benchmark risk-adjusted probability measure. On top of its intuitive appeal and links to various strands of earlier literature, the gain-loss ratio has three economic justifications: it measures violations of the first-order condition for maximizing utility, it is an inverse function of normalized distance to the set of arbitrage opportunities, and it is the dual of the ratio of extreme deviations from the benchmark pricing kernel.

We demonstrate the implications of the gain-loss ratio restriction by computing bounds on the price of options on an asset that does not trade, but is imperfectly correlated with a traded asset. This is useful for example if we want to price real options or value executive stock options. By construction, our bounds lie strictly between the Black-Scholes price (obtained here as an equilibrium price since dynamic replication is impossible) and the no-arbitrage bounds (zero and infinity in this case). Thus, our method offers a general way to chart the middle ground between equilibrium and no-arbitrage. The point is that the optimal trade-off between the precision of an equilibrium model and the credibility of the no-arbitrage principle often lies strictly between the two extremes. This can be seen either as sensitivity analysis for an equilibrium model, or as a strengthening of the no-arbitrage principle.

References

- Aït-Sahalia, Yacine, and Andrew Lo, 1995, Nonparametric estimation of state price densities implied in financial asset prices. Working paper, University of Chicago.
- Backus, David, Silverio Foresi, and Chris Telmer, 1996, Affine models of currency pricing, *New York University Working Paper*.
- Bansal, Ravi, and Bruce N. Lehmann, 1995, Growth optimal portfolio restrictions on stochastic discount factors. Fuqua School of Business, Duke University.
- Bansal, Ravi, and S. Viswanathan, 1993, No arbitrage and arbitrage pricing: A new approach, *Journal of Finance* 48, 1231-1262.
- Bawa, Vijay S., and Eric B. Lindenberg, 1977, Capital market equilibrium in a mean-lower partial moment framework, *Journal of Financial Economics* 5, 189-200.
- Benartzi, Shlomo, and Richard Thaler, 1995, Myopic loss aversion and the equity premium puzzle, *Quarterly Journal of Economics*: 73-92.
- Black, Fischer, and Myron Scholes, 1973, The pricing of options and corporate liabilities, *Journal of Political Economy* 81, 637-659.
- Brennan, Michael J., 1979, The pricing of contingent claims in discrete time models, *Journal of Finance* 24, 53-68.
- Cochrane, John, and Jesús Saá-Requejo, 1995, Beyond arbitrage: Derivatives pricing in incomplete markets. Graduate School of Business, University of Chicago.
- Dybvig, Philip H., and Jonathan E. Ingersoll, Jr., 1982, Mean-Variance Theory in Complete Markets, *Journal of Business* 55, 233-251.
- Grinblatt, Mark, and Sheridan Titman, 1989, Portfolio performance evaluation: Old issues and new insights, *Review of Financial Studies* 2: 393-421.
- Grinblatt, Mark, and Sheridan Titman, 1995, Portfolio evaluation, in R.A. Jarrow, V. Maksimovic, and W. T. Ziemba, *Finance*, North-Holland.
- Hansen, Lars P. and Ravi Jagannathan, 1991, Implications of security market data for models of dynamic economies. *Journal of Political Economy* 99, 225–262.
- Hansen, Lars, and Ravi Jagannathan, 1996, Assessing specification errors in stochastic discount factor models, forthcoming *Journal of Finance*.
- Hardy, G.H., J.E. Littlewood, and G. Pólya, 1964, *Inequalities*, Cambridge University Press.
- Harrison, J. M. and David M. Kreps, 1979, Martingales and arbitrage in multiperiod securities markets. *Journal of Economic Theory* 20, 381–408.
- Ingersoll, Jonathan E. Jr., 1987, *Theory of financial decision making*, Rowman and Littlefield.

- Kahnemann, Daniel, and Amos Tversky, 1979, Prospect theory: An analysis of decision under risk, *Econometrica* 47 (2): 263-291.
- Koski, Jennifer Lynch, and Jeff Pontiff, 1996, How mutual funds use derivatives, *University of Washington Working Paper*.
- Ledoit, Olivier, 1995a, Essays on risk and return in the stock market, 2nd Chapter. Finance Ph.D. Thesis, Sloan School of Management, Massachusetts Institute of Technology.
- Lucas, Robert E., 1978, Asset prices in an exchange economy, *Econometrica* 46: 1429-1445.
- Mandelbrot, Benoît, 1963, The variation of certain speculative prices. *Journal of Business* 36, 394-419.
- Markowitz, Harry, 1952, Portfolio selection, *Journal of Finance* 7: 77-91.
- Mehra, R., and E. Prescott, 1985, The equity premium: A puzzle, *Journal of Monetary Economics* 15: 145-62.
- Merton, Robert, 1973, Theory of rational option pricing, *Bell Journal of Economics and Management Science* Spring, 141-183.
- Roll, Richard, 1977, A critique of the asset pricing theory's tests, part 1: on past and potential testability of the theory, *Journal of Financial Economics* 4, 129-176.
- Ross, S. A., 1976, The arbitrage theory of capital asset pricing. *Journal of Economic Theory* 13, 341-360.
- Rubinstein, Mark, 1976, The valuation of uncertain income streams and the pricing of options, *Bell Journal of Economics and Management Science* 7.
- Snow, K. N., 1991, Diagnosing asset pricing models using the distribution of returns, *Journal of Finance* 46, 955-983.
- Stutzer, Michael, 1993, A Bayesian approach to diagnosis of asset pricing models, *Journal of Econometrics* 68: 367-397.