#18-99

Approximate Arbitrage

First Draft: August 1999
Revised: November 1999

Antonio Bernardo
Anderson Graduate School of Management
University of California, Los Angeles

and

Olivier Ledoit
Anderson Graduate School of Management
University of California, Los Angeles

Finance Working Paper
Sponsored by:
CAPITAL MANAGEMENT SCIENCES

Approximate Arbitrage

Antonio E. Bernardo and Olivier Ledoit*

November 1999

Abstract

We propose a formal definition of approximate arbitrage which can be used to extend the applicability of theories based on the absence of arbitrage. Our definition is based on the ratio of gain to loss, where gain (loss) is the expectation of the positive (negative) part of the excess payoff. Arbitrage is characterized by infinite gain-loss ratio, and approximate arbitrage by gain-loss ratio close to infinity. Our definition of approximate arbitrage has a useful *dual* interpretation in terms of pricing kernels. This allows us to compare the pricing kernel restriction implied by a limit on the maximum gain-loss ratio to other pricing kernel restrictions in the literature. We show theoretically that only the gain-loss restriction is consistent with the absence of arbitrage and approximate arbitrage opportunities. We demonstrate the practical differences of these alternative pricing kernel restrictions by examining their implications for the prices of call options on an asset that does not trade.

*Assistant Professors of Finance, Anderson Graduate School of Management, University of California, Los Angeles, 110 Westwood Plaza, Los Angeles, CA 90095-1481. E-mail: abernard@anderson.ucla.edu, ledoit@ucla.edu; phone - Bernardo: (310) 825-2198, Ledoit: (310) 206-6562; fax: (310) 206-5455. We thank Michael Brennan, John Cochrane, Mark Grinblatt, Stephen Leroy, Francis Longstaff, Richard Roll, Steve Ross, Pedro Santa-Clara, Jesús Saá-Requejo, Michael Stutzer, and participants to a session of the AFA meetings and seminars at Arizona, Chicago, Duke, LBS, UCLA, and UCSB for helpful discussions. Ilya Sharapov provided assistance with the numerical methods. Errors are ours.

1 Introduction

The existence of arbitrage opportunities (zero-cost investments that offer the possibility of gain with no possibility of loss) is incompatible with equilibrium in asset markets. This idea underpins some of the most important contributions to financial economics on topics ranging from corporate finance (Modigliani and Miller, 1958) to asset pricing (Ross, 1976, Black and Scholes, 1973). However, no-arbitrage arguments have an important limitation: while arbitrage opportunities are assumed not to exist, investments offering large potential gains with very small probabilities of very small losses are permitted to exist. Consequently, for many important problems, such as valuing options in incomplete markets, no-arbitrage arguments yield implications too weak to be practically useful.

To strengthen the implications of no-arbitrage arguments, it is also reasonable to assume that the existence of *approximate* arbitrage opportunities is incompatible with well-functioning capital markets. By defining what constitutes an approximate arbitrage opportunity, one could extend the implications of arbitrage-based theories by also precluding opportunities that are close to arbitrage. The goal of this paper is to propose a formal definition of approximate arbitrage. Our definition of approximate arbitrage includes the definition of arbitrage as a limiting case. And, like arbitrage, approximate arbitrage has an equivalent interpretation in terms of pricing kernels. This makes our definition of approximate arbitrage particularly useful for deriving asset pricing implications.[1] Furthermore, we show that within the extant literature, our definition is the only one that correctly measures closeness to arbitrage in an economically meaningful way.

Specifically, we show that investment opportunities with high gain-loss ratios constitute the set of arbitrage and approximate arbitrage opportunities. The gain (loss) of a portfolio is simply the expectation of the positive (negative) part of a portfolio's excess payoff. By definition, an arbitrage opportunity offers strictly positive gain with zero loss, thus it is characterized by an *infinite* gain-loss ratio. However, to go beyond pure arbitrage, one must use information about

[1] In frictionless markets, the price of an asset is determined by multiplying its payoff in any state by a state-contingent discount factor, or pricing kernel, and summing over all states according to their underlying probabilities.

a benchmark investor's preferences. This is accomplished by taking gain and loss expectations under an appropriate risk-adjusted probability measure. Using a notion of distance, or norm, that puts more weight on the states in which consumption is more valuable to the benchmark investor, we prove that every sequence that converges to pure arbitrage has a gain-loss ratio that tends to infinity. Conversely, every (converging) sequence whose gain-loss ratio tends to infinity converges to a pure arbitrage. Consequently, we can formally define an investment opportunity to be an approximate arbitrage if it has a high gain-loss ratio.

Our definition of approximate arbitrage has a useful *dual* interpretation in terms of pricing kernels. Ross (1978) showed that the absence of arbitrage imposes a positivity constraint on the set of admissible pricing kernels. Bernardo and Ledoit (1999) demonstrate that restricting the maximum gain-loss ratio is equivalent to a restriction on the *extreme* values of admissible pricing kernels. Thus, while absence of arbitrage implies that the pricing kernel cannot take negative values or values of zero or infinity in any state, the absence of approximate arbitrage implies that the pricing kernel cannot take values that are *too* low or *too* high. Such economically motivated pricing kernel restrictions are extremely valuable for asset pricing because they restrict the set of prices that can be assigned to any asset.

There have been numerous other results in the literature linking properties of asset returns to a dual restriction on pricing kernels. For example, Hansen and Jagannathan (1991) demonstrated that a restriction on the maximum Sharpe ratio in the economy is equivalent to a restriction on the second moment of the pricing kernel. They then showed that the pricing kernels implied by standard asset pricing models (with reasonable levels of risk aversion) do not exhibit sufficient variation to explain the historically large equity risk premium. Snow (1992) generalized the Hansen-Jagannathan (1991) result by deriving restrictions on the qth moment of the pricing kernel $\widetilde{m}$, i.e. $\mathrm{E}[\widetilde{m}^q]^{1/q}$, for $1 < q < \infty$. Another restriction is due to Stutzer (1993), who showed that restricting the maximum expected utility attainable by a CARA investor is equivalent to restricting the entropy of the pricing kernel $\mathrm{E}[\widetilde{m}\log(\widetilde{m})]$. Bansal and Lehmann (1997) showed that restricting the maximum expected utility that can be attained by an investor with logarithmic utility is equivalent to restricting $\mathrm{E}[-\log(\widetilde{m})]$. Finally, in Section 3.2 we extend the Bansal-

Lehmann result to a restriction on the maximum expected utility that can be attained by an investor with general constant relative risk aversion (CRRA) preferences. Although these pricing kernel restrictions were derived to provide diagnostic tests of asset pricing models, recently they have been used to restrict the set of admissible prices that can be assigned to an asset. For example, Cochrane and Saá-Requejo (1999) augment the Hansen-Jagannathan restriction with a positivity constraint to derive "no-good-deal" asset pricing implications in incomplete markets, while Cerný (1999) considers the pricing implications of the Stutzer and Bansal-Lehman bounds.

We show, however, that only the pricing kernel restriction implied by a maximum gain-loss ratio precludes the existence of arbitrage and approximate arbitrage opportunities. The reason is as follows. Arbitrage and approximate arbitrage opportunities exist when the prices of state claims (assets paying one in that state and zero elsewhere) are either *too* high or *too* low relative to a benchmark investor's willingness to pay for consumption in these states. The pricing kernel restrictions of Hansen-Jagannathan (1991), Snow (1992), and Stutzer (1993) allow state prices to be arbitrarily low (even negative). However, arbitrarily low state prices yield investments offering something tomorrow for nothing today - an arbitrage. Augmenting these restrictions with a positivity constraint on the admissible pricing kernels still allows investments that are arbitrarily close to arbitrage. Conversely, the Bansal-Lehman (1997) pricing kernel restriction, and our extension of this result to general CRRA preferences, allows state prices to be arbitrarily high. Arbitrarily high state prices yield investments offering something today for nothing tomorrow - also an arbitrage. Very high and very low values of the pricing kernel are symmetric in that they both imply arbitrage. To the best of our knowledge, the only restriction sensitive to values of the pricing kernel *both* close to zero and infinity is the one based on the gain-loss ratio. The classification of pricing kernel restrictions implied by this insight constitutes a secondary contribution of our paper.

We demonstrate the practical differences of these alternative pricing kernel restrictions by examining their implications for the prices of call options on an asset that does not trade. This is often the case, for example, when valuing real options. In this setting, the Black-Scholes (1973) dynamic replication argument fails. We find that the pricing kernels implied by a gain-loss re-

striction yield strictly positive lower bounds and finite upper bounds for call option prices. The other pricing kernel restrictions yield either a zero (or negative) lower bound or no upper bound on the price of some call options, implying the existence of arbitrage and approximate arbitrage opportunities.

2 Gain, Loss and Approximate Arbitrage

Consider a two-period model in which assets trade at a certain price today and deliver a random payoff tomorrow. There are S future states of the world with $p_j > 0$ denoting the probability of state j occurring ($j = 1, \ldots, S$). Portfolio payoffs are random variables $\widetilde{z} = (z_1, \ldots, z_S)$ in the space $Z \subset \mathbb{R}^S$ where z_j denotes the payoff in the jth state. We assume, for simplicity, that there exits a riskless asset with rate of return, r_F. Asset prices are given by a linear functional π defined on Z, i.e., the portfolio with payoff $\widetilde{z} \in Z$ has price $\pi(\widetilde{z})$. Given π, we can construct the space of excess payoffs $X = \{\widetilde{z} - (1 + r_F)\pi(\widetilde{z}) : \widetilde{z} \in Z\}$. Finally, we define the null payoff $\widetilde{0} = (0, \ldots, 0)$, the positive orthant $\mathbb{R}^S_+ = \{\widetilde{z} \in \mathbb{R}^S : \widetilde{z} \neq \widetilde{0} \text{ and } \widetilde{z} \geq 0 \ \forall j\}$, and the strict positive orthant $\mathbb{R}^S_{++} = \{\widetilde{z} \in \mathbb{R}^S : \widetilde{z} > 0 \ \forall j\}$.

2.1 The Gain-Loss Ratio

We are interested in characterizing the set of arbitrage and approximate arbitrage opportunities. Any reasonable notion of an approximate arbitrage opportunity must incorporate beliefs about investor preferences. While the absence of arbitrage requires only weak assumptions about investor preferences, i.e. monotonicity, characterizing approximate arbitrage requires stronger assumptions because any non-arbitrage portfolio will be fairly priced for some set of investor preferences.

Information about investor preferences is incorporated in the benchmark model $(u, \widetilde{c}^*)$ where u is a continuously differentiable von Neumann-Morgenstern utility function verifying $u' > 0$, and

$\widetilde{c}^* = (c_1^*, \ldots, c_S^*) \in \mathbb{R}^S$ is equilibrium consumption. Let $p_j^* = p_j u'(c_j^*)/\mathrm{E}[u'(\widetilde{c}^*)]$ for $j = 1, \ldots, S$ and let $\mathrm{E}^*[\cdot]$ denote the expectation under these risk-adjusted probabilities. The benchmark model correctly prices the assets in Z if and only if:

$$\forall \widetilde{x} \in X \text{ s.t. } \widetilde{x} \neq \widetilde{0} \quad \mathrm{E}[u'(\widetilde{c}^*)\widetilde{x}] = 0 \Leftrightarrow \mathrm{E}^*[\widetilde{x}] = 0 \Leftrightarrow \mathrm{E}^*[\widetilde{x}^+ - \widetilde{x}^-] = 0 \Leftrightarrow \frac{\mathrm{E}^*[\widetilde{x}^+]}{\mathrm{E}^*[\widetilde{x}^-]} = 1 \qquad (1)$$

where $\widetilde{x} = \widetilde{x}^+ - \widetilde{x}^-$ is the decomposition of a payoff into its positive part $\widetilde{x}^+ = \max(\widetilde{x}, 0)$ and negative part $\widetilde{x}^- = \max(-\widetilde{x}, 0)$.

Bernardo and Ledoit (1999) define $\mathrm{E}^*[\widetilde{x}^+]$ to be the *gain*, $\mathrm{E}^*[\widetilde{x}^-]$ the *loss*, and $\mathrm{E}^*[\widetilde{x}^+]/\mathrm{E}^*[\widetilde{x}^-]$ the *gain-loss ratio* and show that the gain-loss ratio is (i) mathematically defined on $\mathbb{R}^S$ (except for $\widetilde{0}$); (ii) is always nonnegative; (iii) is equal to $+\infty$ in the positive orthant $\mathbb{R}^S_+$ and finite elsewhere; (iv) is invariant to the multiplication of $\widetilde{x}$ by a positive scalar; and (v) the gain-loss ratio of a short position is the inverse of the ratio of the corresponding long position.

2.2 Dual Formulation in Terms of Pricing Kernels

For the benchmark investor described by $(u, \widetilde{c}^*)$ one can construct the benchmark pricing kernel:

$$\widetilde{m}_j^* = \frac{u'(\widetilde{c}_j^*)}{\mathrm{E}[u'(\widetilde{c}^*)](1 + r_F)}. \qquad (2)$$

The m_j^* represents the benchmark investor's willingness to pay, per unit of probability, for the state claim paying one in the jth state and zero elsewhere. If the investor is risk averse then u' is decreasing. Thus, a state claim that pays off when $\widetilde{c}^*$ is low (high) has relatively high (low) price. Intuitively, such an investment is desirable to risk-averse investors because it allows them to smooth consumption across future states of nature.

Bernardo and Ledoit (1999) demonstrate the following duality result:

Theorem 1

$$\max_{\substack{\widetilde{x}\in X\\ \widetilde{x}\neq 0}} \frac{\mathrm{E}^*[\widetilde{x}^+]}{\mathrm{E}^*[\widetilde{x}^-]} = \min_{\widetilde{m}\in M} \frac{\sup\limits_{j=1,\ldots,S}(m_j/m_j^*)}{\inf\limits_{j=1,\ldots,S}(m_j/m_j^*)} \tag{3}$$

where $M = \{\widetilde{m} \in \mathbb{R}^S_{++} : \forall \widetilde{z} \in Z \ \mathrm{E}[\widetilde{m}\widetilde{z}] = \pi(\widetilde{z})\}$ denotes the set of pricing kernels that correctly price all portfolio payoffs. If markets are complete, i.e. $Z = \mathbb{R}^S$, the set M has a unique element, otherwise M has many elements.

Proof of Theorem 1 *See Bernardo and Ledoit (1999).*

The duality result states that high gain-loss ratio exists when the pricing kernel exhibits extreme deviations from the benchmark pricing kernel. If *actual* state prices are equal to the benchmark investor's *willingness to pay* then all portfolios will be fairly priced and the maximum gain-loss ratio is one. Attractive investment opportunities exist when state prices differ from the benchmark investor's willingness to pay, in which case the benchmark investor can form attractive portfolios by buying (selling) cheap (dear) tradeable combinations of state claims.

2.3 Approximate Arbitrage

It is not surprising that the gain-loss ratio characterizes the set of arbitrage opportunities. Since arbitrage opportunities offer the possibility of gain with no probability of loss they must, by definition, have infinite gain-loss ratios. Formally:

Lemma 1 *The following three statements are equivalent:*

(i) The pricing functional $\pi(\cdot)$ admits arbitrage opportunities.

(ii) $\max\limits_{\substack{\widetilde{x}\in X\\ \widetilde{x}\neq 0}} \dfrac{\mathrm{E}^*[\widetilde{x}^+]}{\mathrm{E}^*[\widetilde{x}^-]} = +\infty.$

(iii) $\inf\limits_{\widetilde{m}\in M} \dfrac{\mathrm{esssup}(\widetilde{m})}{\mathrm{essinf}(\widetilde{m})} = +\infty.$

Proof of Lemma 1 (i) implies (ii) by definition. (ii) implies (iii) by Theorem 1. If (iii) holds then there does not exist a strictly positive pricing kernel which implies the existence of an arbitrage opportunity, therefore (i) holds.□

We now examine the relation between high gain-loss ratios and approximate arbitrage opportunities. We first introduce a notion of distance between excess payoffs and the positive orthant (the set of arbitrage opportunities) using the following norm on $\mathbb{R}^S$: $\|\widetilde{x}\| = \mathrm{E}^*[|\widetilde{x}|]$. Note that this norm is the L^1 (space of all random variables with finite expectations) norm except that it is defined with respect to the benchmark risk-adjusted probability measure, not the true probability measure: it puts more weight on the states where consumption is more valuable to the benchmark investor. This norm captures the fact that any definition of approximate arbitrage must make assumptions about investor preferences. For example, a state claim with a price ϵ close to zero is fairly priced for an investor whose marginal utility of consumption in that state is low. Although it may appear to be a near-arbitrage opportunity, it may not be if investors do not value consumption much in that state.

The next theorem shows that, in this metric, portfolios with high gain-loss ratio are close to being an arbitrage opportunity and vice-versa.

Theorem 2 *The normalized distance to the positive orthant decreases in the gain-loss ratio. Formally:*

$$\forall \widetilde{x} \in \mathbb{R}^S \text{ s.t. } \widetilde{x} \neq \widetilde{0} \qquad \frac{\min\limits_{\widetilde{a}\in\mathbb{R}^S_+} \|\widetilde{x}-\widetilde{a}\|}{\|\widetilde{x}\|} = \frac{1}{\dfrac{\mathrm{E}^*[\widetilde{x}^+]}{\mathrm{E}^*[\widetilde{x}^-]}+1}. \tag{4}$$

Proof of Theorem 2 If $\widetilde{x} \in \mathbb{R}^S_+$ then the normalized distance to the positive orthant is zero and the gain-loss ratio is $+\infty$, therefore Equation (4) holds. Now assume that $\widetilde{x} \notin \mathbb{R}^S_+$. In this case, $\forall \widetilde{a} \in \mathbb{R}^S_+ \quad \mathrm{E}^*[|\widetilde{x}-\widetilde{a}|] = \mathrm{E}^*[|\widetilde{x}^+ - \widetilde{a}\mathbf{1}_{\{\widetilde{x}\geq 0\}}|] + \mathrm{E}^*[|\widetilde{x}^- + \widetilde{a}\mathbf{1}_{\{\widetilde{x}<0\}}|]$, where $\mathbf{1}$ denotes the indicator function of an event. The way to minimize the first term of the decomposition is to impose that $\widetilde{a} = \widetilde{x}^+$ when $\widetilde{x} \geq 0$. The way to minimize the second term is to impose that $\widetilde{a} = 0$ when $\widetilde{x} < 0$. Therefore the minimum distance to the positive orthant is attained by $\widetilde{a} = \widetilde{x}^+$. And, at

the minimum, $\|\widetilde{x} - \widetilde{x}^+\| = \|\widetilde{x}^-\|$. It implies:

$$\frac{\min_{\widetilde{a}\in\mathbb{R}^S_+} \|\widetilde{x} - a\|}{\|\widetilde{x}\|} = \frac{\|\widetilde{x}^-\|}{\|\widetilde{x}\|} = \frac{\mathrm{E}^*[\widetilde{x}^-]}{\mathrm{E}^*[\widetilde{x}^+ + \widetilde{x}^-]} = \frac{1}{\frac{\mathrm{E}^*[\widetilde{x}^+]}{\mathrm{E}^*[\widetilde{x}^-]} + 1}. \qquad \square \tag{5}$$

This notion of normalized distance to arbitrage is invariant to scale as is the notion of arbitrage. The proof shows that we can describe the gain-loss ratio as the magnitude of the closest arbitrage opportunity divided by the distance to the closest arbitrage opportunity.

It is immediately apparent that any finite ceiling $\overline{L}$ on gain-loss ratios implies the absence of arbitrage opportunities. Furthermore, if you get close enough to an arbitrage opportunity eventually your gain-loss ratio will exceed $\overline{L}$.

Corollary 1 *If a sequence of excess payoffs converges to an arbitrage opportunity then its gain-loss ratio goes to infinity.*

Proof of Corollary 1 If $\widetilde{x}_n \to \widetilde{a} \in \mathbb{R}^S_+$ then $\min_{\widetilde{a}'\in\mathbb{R}^S_+} \|\widetilde{x}_n - \widetilde{a}'\| \leq \|\widetilde{x}_n - \widetilde{a}\| \to 0$. Furthermore, elements of the positive orthant have strictly positive norm, therefore $\|\widetilde{x}_n - \widetilde{a}\| \leq \|\widetilde{a}\|/2$ for large enough n, which implies $\|\widetilde{x}_n\| \geq \|\widetilde{a}\| - \|\widetilde{x}_n - \widetilde{a}\| \geq \|\widetilde{a}\|/2 > 0$. Since its numerator converges to zero and its denominator is bounded away from zero for large enough n, the normalized distance to arbitrage of $\widetilde{x}_n$ converges to zero. Corollary 1 then follows from Theorem 2. $\square$

The converse is also true.

Corollary 2 *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 Corollary 2 Assume that $\widetilde{x}_n \to \widetilde{x} \neq \widetilde{0}$ and that $\mathrm{E}^*[\widetilde{x}_n^+]/\mathrm{E}^*[\widetilde{x}_n^-] \to +\infty$. Then Theorem 2 implies $\min_{\widetilde{a}\in\mathbb{R}^S_+} \|\widetilde{x}_n - \widetilde{a}\|/\|\widetilde{x}_n\| \to 0$. Since $\|\widetilde{x}_n\| \to \|\widetilde{x}\| > 0$, we have $\min_{\widetilde{a}\in\mathbb{R}^S_+} \|\widetilde{x}_n - \widetilde{a}\| \to 0$. Therefore there exists $\widetilde{a}_n \in \mathbb{R}^S_+$ such that $\|\widetilde{x}_n - \widetilde{a}_n\| \leq \min_{\widetilde{a}\in\mathbb{R}^S_+} \|\widetilde{x}_n - \widetilde{a}\| + 1/n \to 0$. It implies

that $\|\widetilde{x} - \widetilde{a}_n\| \leq \|\widetilde{x} - \widetilde{x}_n\| + \|\widetilde{x}_n - \widetilde{a}_n\| \to 0$, i.e. $\min_{\widetilde{a} \in \mathbb{R}^S_+} \|\widetilde{x} - \widetilde{a}\| = 0$. Since $\|\widetilde{x}\| > 0$, the normalized distance to arbitrage of $\widetilde{x}$ is zero. By Theorem 2, it implies that the gain-loss ratio of $\widetilde{x}$ is infinite, therefore the limit $\widetilde{x}$ is an arbitrage opportunity. □

Intuitively these results mean that portfolios with gain-loss ratios above some threshold form a *neighborhood* of arbitrage. Formally, for every $\overline{L} < +\infty$ the set $\{\widetilde{x} \in \mathbb{R}^S : \widetilde{x} \neq \widetilde{0}$ and $\mathrm{E}^*[\widetilde{x}^+]/\mathrm{E}^*[\widetilde{x}^-] > \overline{L}\}$ is a topological neighborhood of the positive orthant $\mathbb{R}^S_+$. Thus, we say the excess payoff $\widetilde{x}$ is an approximate arbitrage opportunity if $\mathrm{E}^*[\widetilde{x}^+]/\mathrm{E}^*[\widetilde{x}^-] > \overline{L}$. By Theorem 1 we can say, equivalently, that approximate arbitrage opportunities exist when state prices are too *high* or too *low* relative to those implied by the benchmark investor's preferences.

3 Classification of Duality Results

We now examine the relation between other pricing kernel restrictions and the absence of arbitrage and approximate arbitrage opportunities based on whether a given restriction prevents state prices from being too *high* or too *low*.

3.1 Restricting High State Prices

The first category contains restrictions that, in effect, prevent state prices 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. Hansen and Jagannathan (1991) demonstrated that a bound on the variance of the pricing kernel is equivalent to a bound on the maximum Sharpe ratio (mean to standard deviation of the excess payoff). In the presence of a riskfree bond, the expectation of the pricing kernel is pinned down by $\mathrm{E}[\widetilde{m}(1 + r_F)] = 1$, therefore restricting the

variance is the same as restricting the second moment. Snow (1992) generalizes this result by deriving restrictions on the qth moment of the pricing kernel $\mathrm{E}[\widetilde{m}^q]^{1/q}$, for $1 < q < \infty$. By letting q go to infinity, we obtain a restriction on the supremum of the pricing kernel: $\sup(\widetilde{m})$. This is *one part* of the dual of the gain-loss ratio restriction stated in Theorem 1 (the other part being a restriction on the infimum), up to normalization by the benchmark pricing kernel $\widetilde{m}^*$. Another 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 $\mathrm{E}[\widetilde{m}\log(\widetilde{m})]$.

The problem with the restrictions in this category, for the purpose of deriving asset prices, is that they do not prevent state prices from being too close to zero. For example, Cochrane and Saá-Requejo (1999) derive price bounds on derivatives by ruling out the existence of "good deals" (investment opportunities with high Sharpe ratios) and arbitrage opportunities. Following Hansen and Jagannathan (1991), the Sharpe ratio bound is equivalent to a bound on the variance of the pricing kernel; and the no-arbitrage assumption implies the pricing kernel is positive. However, the variance bound does not rule out pricing kernels that take arbitrarily low (or even negative!) values in some states therefore it does not rule out arbitrage or approximate arbitrage opportunities. In other words, a Sharpe ratio constraint does not form a *neighborhood* of the set of arbitrage opportunities.[2] A striking example is the portfolio whose excess return follows the probability density function: $f(x) = 2/x^3$ if $x \geq 1$ and zero otherwise. The mean is two and the 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.[3] Adding a positivity constraint on the admissible pricing kernels still allows investments that are arbitrarily close to arbitrage because it permits state prices to be arbitrarily close to zero. Similarly, the restrictions of Snow (1992)

[2]A Sharpe ratio constraint does form a neighborhood of the set of arbitrage opportunities when the distribution of returns is Gaussian. In such cases, a gain-loss ratio restriction is equivalent to a Sharpe ratio restriction (Bernardo and Ledoit, 1999). The Cochrane-Saá-Requejo approach works well over short trading horizons, because at these horizons returns are almost normally distributed.

[3]This apparent paradox is related to the fact that mean-variance preferences, from which the Sharpe ratio is derived, requires a quadratic utility function (except for particular distributions of payoffs like the Gaussian), which displays satiation. Another facet of this problem is the proof by Dybvig and Ingersoll (1982) that the CAPM admits arbitrage opportunities if markets are complete. This suggests that mean-variance concepts are fundamentally incompatible with the no-arbitrage principle.

and Stutzer (1993) are incompatible with the no-arbitrage principle.

An alternative explanation for why such restrictions do not preclude arbitrage and approximate arbitrage opportunities is as follows. The Stutzer (1993) bound, for example, is derived by imposing a limit on the maximum attainable expected utility for an investor with CARA (constant absolute risk aversion) preferences. However, the maximum value that a negative exponential utility function can attain in any state is zero. Suppose that it is possible to get something in some states tomorrow for nothing today, i.e. an arbitrage. 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: if infinite wealth is attained with positive probability then expected utility is infinite. In that case, a restriction on the maximum expected utility attainable would have successfully ruled out the arbitrage opportunities that pay off only in some states. Thus, another way to express the problem with the restrictions in this category is that they correspond to utility functions that have a finite upper bound.[4]

3.2 Restricting Low State Prices

The second category contains the restrictions that, in effect, prevent state prices 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: Bansal and Lehmann (1997) show that restricting the maximum expected utility that can be attained by an agent with logarithmic utility is equivalent to restricting $\mathrm{E}[-\log(\widetilde{m})]$. We generalize their result to any CRRA utility function.

Theorem 3 *Let Z denote a space of payoffs, $\pi(\cdot)$ a pricing functional defined on Z, and M the*

[4]Clearly, this is also the case with quadratic utility.

they do not prevent state prices from being too close to infinity. In fact, they can allow some state prices to be *equal* to infinity. This is not economically plausible, since it allows arbitrage opportunities. This is the reverse of our earlier criticism. Again, the reason for the problem can be restated in terms of payoffs. Suppose, in the above example, that the call option is overpriced. In order to take advantage of the mispricing, one 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, implying 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 option prices that are too high. The problem with the restrictions in this category is that they correspond to utility functions that are only defined on the positive half-line.

3.3 Restricting High And Low State Prices Simultaneously

When we categorize pricing kernel restrictions in this way, it is easy to see that one ought to prevent state prices from getting too close to infinity *and* to zero. Consider the restriction stated in Theorem 1 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 (1997) logarithm bound for low state prices. The gain-loss restriction is the first one to restrict high and low state prices simultaneously.

4 Pricing an Option on a Non-Traded Asset

In this section, we derive pricing bounds for an option on an asset that does not trade using a gain-loss ratio restriction and then compare them to the price bounds implied by the other pricing

kernel restrictions in the literature.

We begin by using the pricing methodology in Bernardo and Ledoit (1999) to derive price bounds on the option consistent with a gain-loss ratio restriction. We are interested in the price a European call option with striking price K and a given maturity. We assume that the underlying does not trade, therefore, the Black-Scholes (1973) dynamic replication argument does not apply. We assume, however, that there exists a traded asset that is correlated with the underlying. We will use the information contained in the prices of a set of related basis assets to price the call option. The basis assets we choose include the riskless bond and a call 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.[5] To implement the gain-loss methodology, one must also choose a benchmark pricing kernel reflecting investor preferences. The natural choice is the Black-Schole pricing kernel. Rubinstein (1976) showed 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 portfolio underpriced relative to this benchmark equilibrium model. Even though the Black-Scholes formula is not the only valid one in this case, in our opinion it constitutes a reasonable benchmark.

Assume without loss of generality that the option expires at time one. Call $\widetilde{S}_n$ ($\widetilde{S}_t$) 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:

$$\begin{bmatrix} \log\left(\widetilde{S}_n/S_0\right) \\ \log\left(\widetilde{S}_t/S_0\right) \end{bmatrix} \overset{*}{\sim} \mathcal{N}\left(\begin{bmatrix} r-\dfrac{\sigma_n^2}{2} \\ r-\dfrac{\sigma_t^2}{2} \end{bmatrix}, \begin{bmatrix} \sigma_n^2 & \rho\sigma_n\sigma_t \\ \rho\sigma_n\sigma_t & \sigma_t^2 \end{bmatrix}\right) \tag{9}$$

where $\overset{*}{\sim}$ denotes distribution under the benchmark risk-adjusted probability measure, $\mathcal{N}(\cdot)$ denotes the bivariate normal distribution, r is the continuously compounded riskfree rate, σ_n (σ_t)

[5]This is a simplification: other derivative products of the traded asset could be included among basis assets to further tighten the resulting bounds. 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.

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:

$$\pi_t = S_0\Phi(d) - Ke^{-r}\Phi(d-\sigma_t) \quad \text{where } d = \frac{\log(S_0/K)+r}{\sigma_t} + \frac{1}{2}\sigma_t \tag{10}$$

and where Φ denotes the standard normal cumulative distribution function.

In this framework, the tightness of the bounds is controlled by the parameter $\overline{L}$, which is the maximum gain-loss ratio in the economy. For example, setting $\overline{L} = 1$ 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 $\overline{L}$ strictly above one. We choose the value $\overline{L} = 2$ meaning that the CRRA agent cannot receive twice as much gain as would be necessary for her to increase her holdings in the asset. Equivalently, we are defining an approximate arbitrage opportunity to be an excess payoff with gain-loss ratio above two. We will find all call option prices which are consistent with the absence of such opportunities.

The computational details are as follows. We generate $I = 10,000$ draws from the joint lognormal distribution of terminal prices for the traded and non-traded assets under the Black-Scholes risk-neutral probability measure. Let (y_1^i, y_2^i) denote two independent standard normal variates iid across $i = 1, \ldots, I$. The final value of the non-traded asset in the ith simulation is $S_n^i = S_0 \exp[r - (\sigma_n^2/2) + \sigma_n y_1^i]$, and the final value of the traded asset is: $S_t^i = S_0 \exp[r - (\sigma_t^2/2) + \sigma_t(\rho y_1^i + \sqrt{1-\rho^2} y_2^i)]$. Using these simulated payoffs, the Black-Scholes price of the call option on the traded asset is: $\pi_t = \frac{1}{I}\sum_{i=1}^{I} \max\left(S_t^i - K, 0\right) e^{-r}$. In the ith simulation the option on the non-traded asset has payoff: $z^i = \max(S_n^i - K, 0)$. Following Bernardo and Ledoit (1999), the bounds on the price π_n of this option implied by a limit $\overline{L}$ on the maximum gain-loss ratio are:

$$\begin{array}{ccccc} \max\limits_{\substack{w_0, w_1 \in \mathbb{R} \\ \frac{\frac{1}{I}\sum_{i=1}^{I}(b^i - z^i)^+}{\frac{1}{I}\sum_{i=1}^{I}(b^i - z^i)^-} \geq \overline{L} \\ b^i = w_0 + w_1 \max(S_t^i - K, 0)}} w_0 e^{-r} + w_1 \pi_t & \leq & \pi_n & \leq & \min\limits_{\substack{w_0, w_1 \in \mathbb{R} \\ \frac{\frac{1}{I}\sum_{i=1}^{I}(z^i - b^i)^+}{\frac{1}{I}\sum_{i=1}^{I}(z^i - b^i)^-} \geq \overline{L} \\ b^i = w_0 + w_1 \max(S_t^i - K, 0)}} w_0 e^{-r} + w_1 \pi_t \end{array} \tag{11}$$

where b^i is the payoff in the ith simulation of the replicating portfolio of basis assets with

weight w_0 on the riskfree bond and weight w_1 on the option on the traded asset. We computed these bounds using the Optimization Toolbox of the programming language MATLAB. The only numerical trick was to rewrite the constraint of the left-hand side maximization program as $\frac{1}{I}\sum_{i=1}^{I}(b_i - z_i)^+ \geq \frac{1}{I}\sum_{i=1}^{I}(b_i - z_i)^- \times \overline{L}$ and do the same thing for the right-hand side minimization program.

Figure 1 plots the resulting bounds for various values of the striking price K. The parameters are as follows: the continuously compounded 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 $\sigma_t = \sigma_n = 0.20$; and the correlation coefficient between the two assets is 0.70. The computed bounds lie strictly

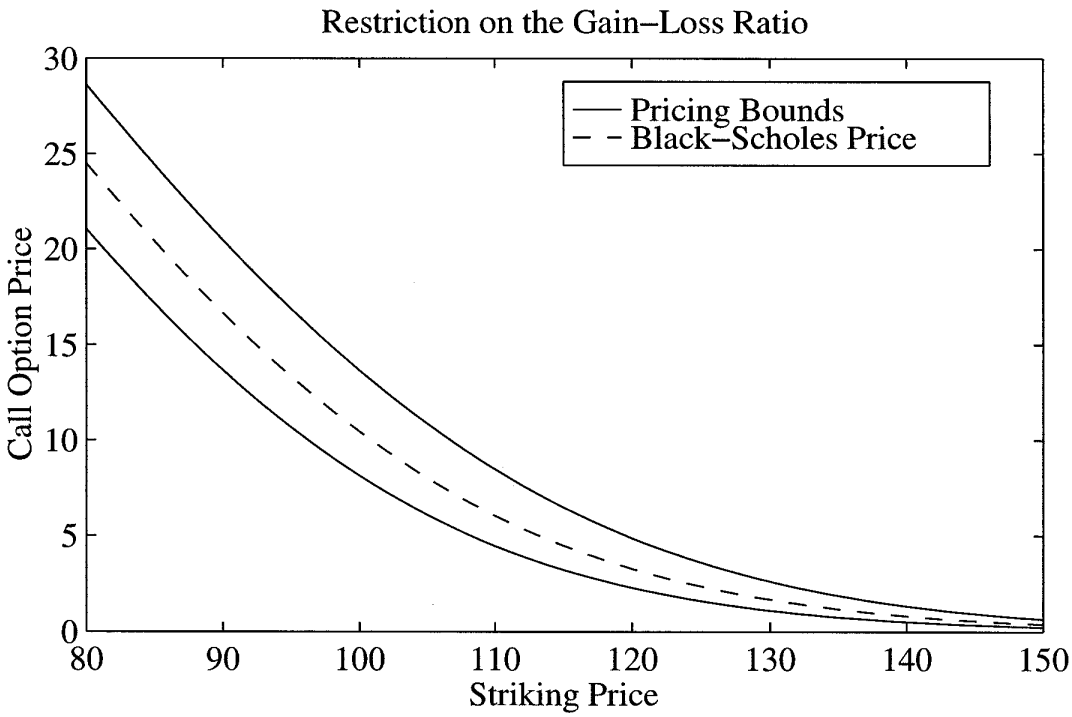

Figure 1: Bounds on the price of a call option on a non-traded stock.

between the Black-Scholes price and the no-arbitrage bounds (zero and infinity in this case). They get tighter as (i) the maximum gain-loss ratio goes down, or (ii) the correlation between traded and non-traded assets goes up. In the first case it is because the benchmark model is more reliable, and in second one it is because the basis assets allow more accurate replication. These bounds represent the option prices that rule out the existence of arbitrage and approximate

arbitrage opportunities.

4.1 Comparison with Other Restrictions

To illustrate the asset pricing implications of the other pricing kernel restrictions in the literature, we use the same call option example. The option pricing bounds implied by a restriction of the Stutzer (1993) type are plotted in Figure 2.[6] The setup is the same as in Figure 1. In addition, we assumed that both traded and non-traded 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. Notice that the price of the call can be

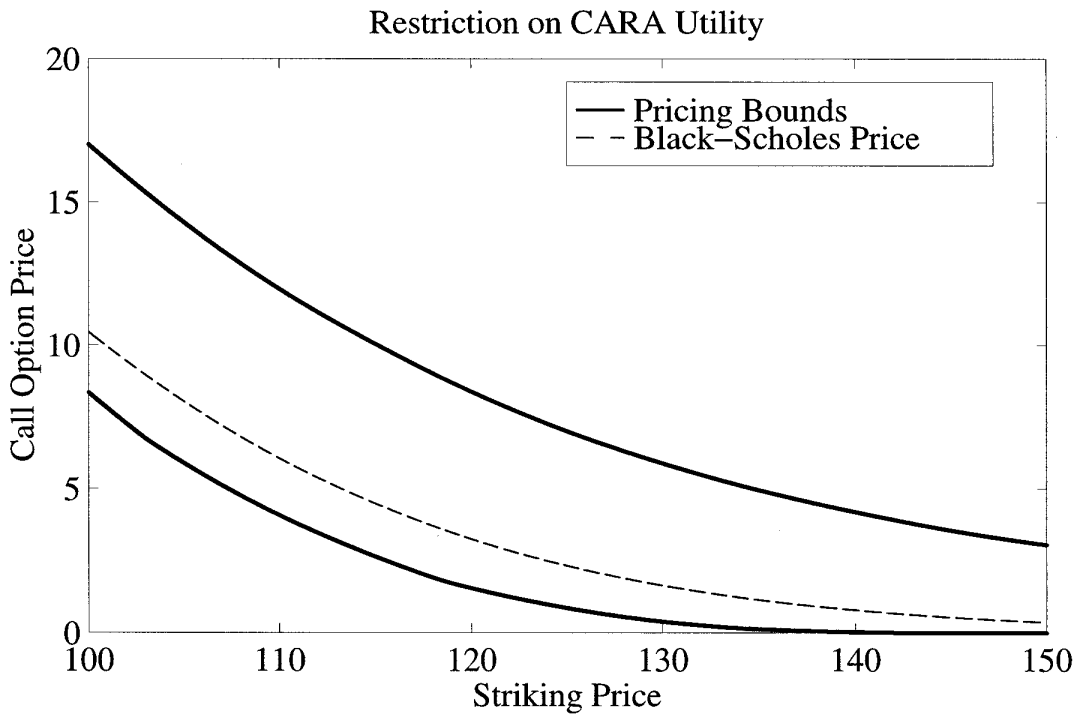

Figure 2: Bounds on the price of a call option on a non-traded stock. When the striking price is at 143 and above, the lower bound is *exactly* equal to zero: it is not due to numerical roundoff.

zero (or at least arbitrarily close to zero) in some cases, which implies an arbitrage opportunity.

[6]The computational details are omitted but are available from the authors.

The corresponding graphs for the restrictions of Hansen and Jagannathan (1991) and Snow (1992) are not shown because they look similar and in particular suffer from the same problem.[7]

We also use the call option example to derive option pricing bounds implied by a restriction of the Bansal-Lehmann (1997) type. The results are plotted in Figure 3. The setup is the same as in Figures 1 and 2. Notice that there is no upper bound: the price of the call can be infinite

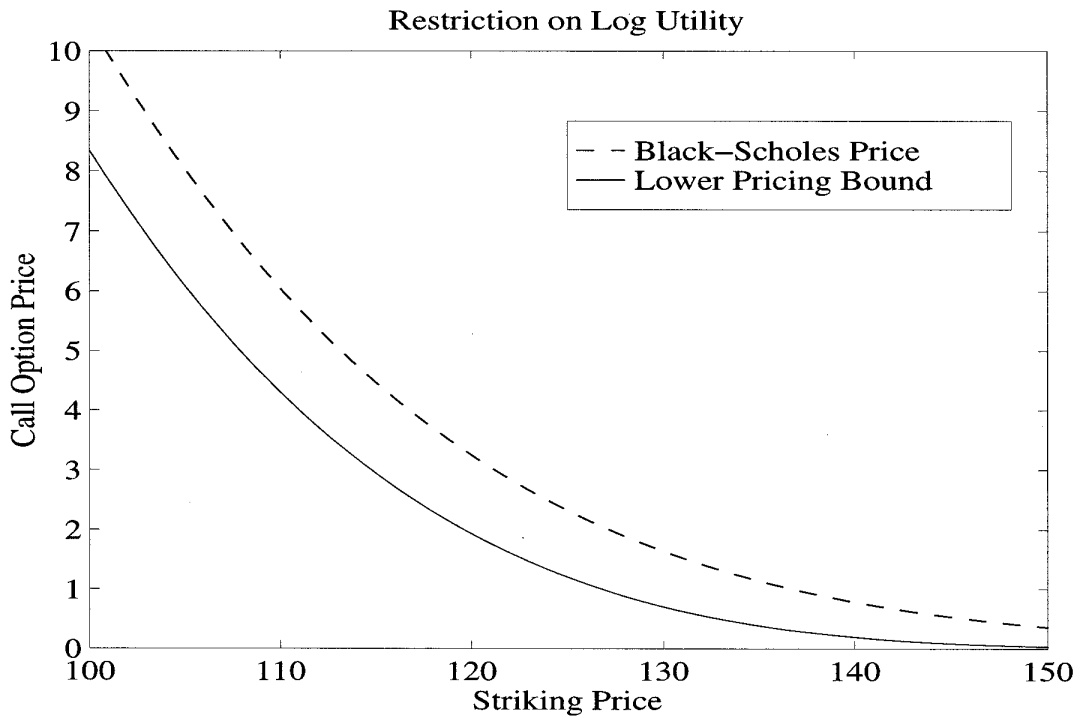

Figure 3: Lower bound on the price of a call option on a non-traded stock. The lower bound is always strictly positive, but the upper bound is infinity, i.e. the option price is allowed to be arbitrarily large.

(or at least arbitrarily close to infinity), which would also imply 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.[8]

In sum, we find that the pricing kernels implied by a gain-loss restriction yield strictly positive

[7]Assuming that we impose $\widetilde{m} > 0$ *in addition* to restricting the qth moment of $\widetilde{m}$. Without the additional positivity constraint, it is even worse: the pricing bounds can go negative.

[8]Strictly 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.

lower bounds and finite upper bounds for call option prices. By contrast, the other pricing kernel restrictions imply either a zero (or negative) lower bound or no upper bound on the price of some call options. Consequently, such restrictions yield price bounds so weak that they permit arbitrage and/or approximate arbitrage opportunities.

5 Conclusion

This paper defines rigorously the intuitive concept of approximate arbitrage. We show that the gain-loss ratio, and its dual formulation in terms of pricing kernel, correctly measures closeness to arbitrage. A new classification of duality results shows that no other one in the literature enjoys such a privileged relationship with the concept of arbitrage. The pricing of an option on a non-traded asset demonstrates the practical relevance of the theoretical developments.

The objective of this paper is to extend ideas from no-arbitrage theory beyond the realm of pure arbitrage and into the real world in which investors move quickly to exploit very attractive investment opportunities. Such an approach will be extremely useful for many asset pricing problems that inherently involve incomplete markets such as valuing real options, executive stock options, and fixed-income derivatives in the presence of default risk. Moreover, this approach provides another powerful test of intertemporal asset pricing models (see also Hansen and Jagannathan, 1991 and 1997).

References

Bansal, Ravi, and Bruce N. Lehmann, 1997, Growth optimal portfolio restrictions on stochastic discount factors, *Macroeconomic Dynamics* 1, 333-354.

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.

Bernardo, Antonio E., and Olivier Ledoit, 1999, Gain, loss and asset pricing, *Journal of Political Economy*, forthcoming.

Black, Fischer, and Myron Scholes, 1973, The pricing of options and corporate liabilities, *Journal of Political Economy* 81, 637-659.

Cerný, Ales, 1999, Generalized Sharpe ratios and consistent no-good-deal restrictions in a model of continuous trading, Unpublished Working Paper.

Cochrane, John, and Jesús Saá-Requejo, 1999, Beyond arbitrage: "Good-Deal" asset pricing bounds in incomplete markets, *Journal of Political Economy*, forthcoming.

Dybvig, Philip H., and Jonathan E. Ingersoll Jr., 1982, Mean-variance theory in complete markets, *Journal of Business* 55, 233-251.

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 P., and Ravi Jagannathan, 1997, Assessing specification errors in stochastic discount factor models, *Journal of Finance* 52, 557-590.

Hardy, G.H., J.E. Littlewood, and G. Pólya, 1964, *Inequalities* (Cambridge University Press, Cambridge).

Merton, Robert, 1973, Theory of rational option pricing, *Bell Journal of Economics and Management Science* 4, 141-183.

Modigliani, Franco, and Merton H. Miller, 1958, The cost of capital, corporation finance and the theory of investment, *American Economic Review* 48, 261-297.

Ross, Stephen A., 1976, The arbitrage theory of capital asset pricing, *Journal of Economic Theory* 13, 341-360.

Ross, Stephen A., 1978, A simple approach to the valuation of risky streams, *Journal of Business* 51, 453-475.

Rubinstein, Mark, 1976, The valuation of uncertain income streams and the pricing of options, *Bell Journal of Economics and Management Science* 7, 407-425.

Snow, Karl N., 1992, 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.