

Gain, Loss, and Asset Pricing

Antonio E. Bernardo and Olivier Ledoit

University of California, Los Angeles

We develop an approach to asset pricing in incomplete markets that bridges the gap between the two fundamental approaches in finance: model-based pricing and pricing by no arbitrage. We strengthen the absence of arbitrage assumption by precluding investment opportunities whose attractiveness to a benchmark investor exceeds a specified threshold. In our framework, the attractiveness of an investment opportunity is measured by the gain-loss ratio. We show that a restriction on the maximum gain-loss ratio is equivalent to a restriction on the ratio of the maximum to minimum values of the pricing kernel. By limiting the maximum gain-loss ratio, we can restrict the admissible set of pricing kernels, which in turn allows us to restrict the set of prices that can be assigned to assets. We illustrate our methodology by computing price bounds for call options in a Black-Scholes economy without intermediate trading. When we vary the maximum permitted gain-loss ratio, these bounds can range from the exact prices implied by a model-based pricing approach to the loose price bounds implied by the no-arbitrage approach.

I. Introduction

There are two fundamental approaches for pricing assets. Each restricts the set of prices that can be assigned to an asset by restricting

We thank Michael Brennan, John Cochrane, Mark Grinblatt, Narayana Kocherlakota, Stephen LeRoy, Francis Longstaff, Richard Roll, Steve Ross, Pedro Santa-Clara, Jesús Saá-Requejo, Michael Stutzer, Eric Talley, Karen Van Nuy, Jean-Luc Vila, Ivo Welch, and participants to a session of the American Finance Association meetings and seminars at Arizona, Chicago, Duke, London Business School, UCLA, and California at Santa Barbara for helpful discussions. We are especially grateful to two anonymous referees and the editor (Lars Hansen) for many suggestions that greatly improved the paper. Ilya Sharapov provided assistance with the numerical methods. Errors are ours.

the set of admissible pricing kernels.¹ The first approach, known as *model-based pricing*, makes explicit assumptions about a benchmark investor's preferences, which in turn yield a specific pricing kernel embodying the investor's willingness to pay for consumption across states. By virtue of its strong assumptions, this approach yields pricing implications that are exact but sensitive to misspecification error. The second approach, known as *no-arbitrage pricing*, assumes only the existence of a set of basis assets (with known prices) and the absence of arbitrage opportunities to restrict the admissible set of pricing kernels to those that correctly price the basis assets and assign positive values to payoffs in every state. If the basis assets do not complete the market, the admissible set contains many pricing kernels. By virtue of its weak assumptions, this approach yields pricing implications in incomplete markets that are robust but often too imprecise to be economically interesting.

Much of the literature has heretofore presumed that these modeling approaches are mutually incompatible in incomplete markets, thereby leaving researchers and practitioners to make an uneasy choice between precision and robustness. Our paper, however, proposes a framework to unify them. Our analysis incorporates both information about investor preferences via a benchmark pricing kernel *and* information contained in the prices of basis assets and strengthens the no-arbitrage condition to also preclude investment opportunities whose attractiveness to the benchmark investor exceeds a specified threshold. The combination of these assumptions yields a restricted set of admissible pricing kernels to restrict asset prices in an economically meaningful way. Moreover, our analysis demonstrates that model-based and no-arbitrage pricing techniques represent extreme cases of a single framework.

In our framework, the attractiveness of an investment opportunity is measured by the "gain-loss" ratio, which is the expectation of the investment's positive excess payoffs divided by the expectation of its negative excess payoffs. By taking expectations under appropriately chosen risk-adjusted probabilities, we can incorporate information about investor preferences for consumption in different states. In general, investments with a high gain-loss ratio are very desirable for the benchmark investor, and, in the limit, investments with infinite gain-loss ratios constitute arbitrage opportunities.

Central to our approach is a new duality result linking the existence of investments with a high gain-loss ratio to pricing kernels

¹ In frictionless markets, the value 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 possible states according to their underlying probabilities.

exhibiting extreme deviations from the benchmark pricing kernel. By imposing a finite limit, \bar{L} , on the maximum gain-loss ratio, we restrict the admissible set of pricing kernels to those that do not exhibit such extreme deviations. If \bar{L} goes to one (its lower bound), the admissible set shrinks to contain only the benchmark pricing kernel. If \bar{L} goes to infinity, the admissible set grows to include all pricing kernels consistent with the absence of arbitrage among the basis assets.² Thus \bar{L} allows one to parameterize the trade-off between the precision of the model-based approach and the robustness of the no-arbitrage approach. Since many pricing problems arguably call for intermediate levels of precision and robustness, our framework lends considerable flexibility to existing pricing methodologies. To implement our framework, one must choose (i) a value for the parameter \bar{L} , (ii) a benchmark pricing kernel, and (iii) an appropriate set of basis assets. We suggest several ways to guide these choices in Section VC of the paper.

Our duality result is similar to the result in Hansen and Jagannathan (1991) linking the availability of attractive investment opportunities, measured by the Sharpe ratio (mean over standard deviation), to the standard deviation of the pricing kernel. The main advantage of our result for deriving asset pricing implications is that a restriction on the maximum gain-loss ratio, unlike a Sharpe ratio restriction, is equivalent to precluding the existence of arbitrage and approximate arbitrage opportunities.

We apply our methodology to a canonical problem in finance: pricing an option that cannot be replicated. We consider a Black-Scholes (1973) economy with no dynamic trading; thus the Black-Scholes dynamic hedging solution does not apply. However, Merton (1973) showed that in this setting the lower and upper no-arbitrage bounds are $\max(0, S - Ke^{-rt})$ and S , respectively, where S denotes the initial stock price, K the option strike price, r the continuously compounded risk-free rate, and t the option maturity. Moreover, Rubinstein (1976) showed that an appropriately chosen benchmark pricing kernel implies an option price equal to the Black-Scholes price. We examine the robustness of Rubinstein's utility-based result by adopting his benchmark pricing kernel and varying the maximum gain-loss ratio, \bar{L} . Our price bounds converge to the Black-Scholes price as \bar{L} goes to one and widen to the no-arbitrage bounds

² The benchmark pricing kernel must be positive, but it need not correctly price the set of basis assets. For example, our method allows a consumption-based benchmark pricing kernel or a risk-neutral benchmark, even if they misprice stocks included among the basis assets. This issue is discussed in detail in Secs. IVA and VC.

as \bar{L} goes to infinity. From a practical perspective, imposing reasonable restrictions on the maximum gain-loss ratio dramatically sharpens the call option price bounds from those implied by the no-arbitrage approach. We also find that the benchmark pricing kernel yields pricing implications that are relatively less robust for at-the-money options than for in-the-money or out-of-the-money options. This example illustrates that our methodology can be viewed either as sensitivity analysis around a specific asset pricing model or as strengthening the no-arbitrage principle.

The paper is organized as follows. Section II presents a simple numerical example to illustrate our methodology. Section III formalizes the gain-loss ratio concept and demonstrates our main duality result in a finite-state economy with a riskless asset. The result is extended to the case of an infinite-state economy in the Appendix. Section IV describes how our duality result can be used to derive pricing bounds that lie between those implied by a specific asset pricing model and the no-arbitrage bounds. Section V examines a detailed application of our method: pricing an option on a stock with no intermediate trading. We provide guidance on how to choose the key inputs of our pricing methodology—the maximum gain-loss ratio \bar{L} , the benchmark pricing kernel, and the set of basis assets—and also discuss related methods. Section VI concludes with potential applications and directions for future research.

II. Illustrative Numerical Example

Consider an incomplete markets economy in which a stock and a bond trade at some price today and deliver payoffs in three equally probable states next period. The stock payoff next period is given by $(2, 1, 0)$ in states 1, 2, and 3 and the bond payoff is given by $(1, 1, 1)$. Without loss of generality, today's stock and bond prices are normalized to one, implying a risk-free rate, r_F , of zero. Now suppose that we introduce a call option on the stock with exercise price $K = 1$. What can we say about its price, C ?

The *model-based pricing* approach assumes a benchmark pricing kernel, $\tilde{m}^* = (m_1^*, m_2^*, m_3^*)$, which assigns values to consumption in different states. Given \tilde{m}^* , the price of the call today is $E[\tilde{m}^* \tilde{z}]$, where $\tilde{z} = (1, 0, 0)$ is the option's contingent payoff next period. The values m_j^* can be interpreted as a benchmark investor's willingness to pay, per unit of probability, for the state claim paying one in the j th state and zero elsewhere. In the simplest case, the benchmark investor is risk-neutral, and $m_j^* = 1$ for all j . Thus one can assign a price $C = E[\tilde{m}^* \tilde{z}] = 1/3$ to the option.

The *no-arbitrage pricing* approach finds all the call option prices that do not permit arbitrage opportunities. This is equivalent to all the call option prices implied by the set of positive pricing kernels that correctly price the stock and bond. In this example, the admissible pricing kernels must satisfy $m_1, m_2, m_3 > 0$ and

$$\frac{1}{3}(m_1 \times 2) + \frac{1}{3}(m_2 \times 1) + \frac{1}{3}(m_3 \times 0) = 1 \quad (1)$$

and

$$\frac{1}{3}(m_1 \times 1) + \frac{1}{3}(m_2 \times 1) + \frac{1}{3}(m_3 \times 1) = 1, \quad (2)$$

which implies that $m_3 = m_1$, $m_2 = 3 - 2m_1$, and $0 < m_1 < \frac{3}{2}$. Consequently, the no-arbitrage bounds for the call option are $0 < C < \frac{1}{2}$.

Our paper proposes a third approach, which finds the set of option prices that precludes investment opportunities that exceed a threshold level of attractiveness for a benchmark investor. The measure of attractiveness we use is the *gain-loss ratio*. In the simplest case of a risk-neutral benchmark investor, the gain-loss ratio of any zero-price portfolio \tilde{x} is defined as $E[\tilde{x}^+]/E[\tilde{x}^-]$, where $\tilde{x}^+ = \max(0, \tilde{x})$ and $\tilde{x}^- = \max(0, -\tilde{x})$ represent the positive and negative parts of the payoff, respectively. The gain-loss ratio summarizes the attractiveness of any zero-price portfolio. A gain-loss ratio of one implies that the investment is fairly priced, and a gain-loss ratio above one implies the existence of an attractive investment opportunity.

The key to our approach is the following duality result, proved in Section III, relating the existence of high-gain-loss ratio investment opportunities to large ratios of the maximum to minimum values of the pricing kernel:

$$\max_{\{\tilde{x} \in X, \tilde{x} \neq 0\}} \frac{E[\tilde{x}^+]}{E[\tilde{x}^-]} = \min_{\{\tilde{m} > 0: E[\tilde{m}\tilde{x}] = 0 \forall \tilde{x} \in X\}} \frac{\sup_j(m_j)}{\inf_j(m_j)}, \quad (3)$$

where X denotes the space of zero-price payoffs. The maximand on the left-hand side is the gain-loss ratio, and its maximum value is found by searching over all possible zero-cost portfolios that can be constructed using the stock, the bond, and the option. The minimand on the right-hand side is the ratio of the highest to lowest values of the pricing kernel across states, and its minimum value is found by searching over all positive pricing kernels that price the assets. Thus a gain-loss restriction is equivalent to a restriction on admissible pricing kernels. For example, suppose that, in a well-functioning market, investments with gain-loss ratios above some value \bar{L} should not exist. With our duality result, this assumption

reduces the set of admissible pricing kernels to those that are positive and correctly price the basis assets, and whose ratio of highest to lowest values is below \bar{L} , that is,

$$m_1 > 0 \quad \text{subject to} \quad \frac{\sup_j(m_1, 3 - 2m_1)}{\inf_j(m_1, 3 - 2m_1)} \leq \bar{L}. \quad (4)$$

This yields the admissible pricing kernels $3/(\bar{L} + 2) < m_1 < 3\bar{L}/(2\bar{L} + 1)$, $m_2 = 3 - 2m_1$, and $m_3 = m_1$. For example, if we let $\bar{L} = 2$, this restricts the admissible pricing kernels to $3/4 < m_1 < 6/5$, implying the call price bounds $1/4 < C < 2/5$. These bounds are strictly narrower than the no-arbitrage bounds, and they contain the model-based price. The reader can verify that if \bar{L} is allowed to approach one (its lower bound), the call price bounds approach the unique price implied by the risk-neutral benchmark, $1/3$; and if \bar{L} is allowed to approach infinity, the call price bounds approach the no-arbitrage bounds, $(0, 1/2)$.

III. Theory

In this section, we derive our central duality result in a finite-state economy, which includes a riskless bond with known price. These assumptions are made for expositional clarity and are relaxed in the Appendix.

A. A Finite-State Framework for Asset Pricing

Consider a two-period model in which assets trade at a certain price today and deliver a random payoff next period. There are S future states of the world, with $p_j > 0$ denoting the probability that state j occurs ($j = 1, \dots, S$). The economy includes at most S linearly independent assets generating the space $Z \subset \mathbb{R}^S$ of portfolio payoffs. Portfolio payoffs are random variables $\tilde{z} = (z_1, \dots, z_S) \in Z$, where z_j denotes the payoff in the j th state. Of special interest are the null payoff $\tilde{0} = (0, \dots, 0)$, the positive orthant $\mathbb{R}_+^S = \{\tilde{z} \in \mathbb{R}^S: \tilde{z} \neq \tilde{0} \text{ and } z_j \geq 0 \forall j\}$, and the strict positive orthant $\mathbb{R}_{++}^S = \{\tilde{z} \in \mathbb{R}^S: z_j > 0 \forall j\}$.

Asset prices are given by a linear functional π defined on Z ; that is, the portfolio with payoff $\tilde{z} \in Z$ has price $\pi(\tilde{z})$. We assume that the pricing functional π does not allow arbitrage opportunities, that is, for all $\tilde{z} \in Z \cap \mathbb{R}_+^S$, $\pi(\tilde{z}) > 0$, and that there exists a riskless asset with risk-free rate of return r_f . Finally, given π , we can construct the space of excess payoffs $X = \{\tilde{z} - (1 + r_f)\pi(\tilde{z}): \tilde{z} \in Z\}$.

B. The Gain-Loss Ratio

We represent information about a benchmark investor with the pricing model (u, \tilde{c}^*) , where u is a continuously differentiable von Neumann–Morgenstern utility function verifying $u' > 0$, and $\tilde{c}^* = (c_1^*, \dots, c_S^*) \in \mathbb{R}^S$ is equilibrium consumption. This allows us to construct the benchmark pricing kernel:

$$m_j^* = \frac{u'(\tilde{c}_j^*)}{E[u'(\tilde{c}^*)]} \cdot \frac{1}{1 + r_f}. \quad (5)$$

The m_j^* represents the benchmark investor's willingness to pay, per unit of probability, for the state claim paying one in the j th state and zero elsewhere. If the investor is risk-averse, then u' is decreasing. Thus a state claim that pays off when \tilde{c}^* is low (high) has a relatively high (low) price. Such a state claim is more valuable for the benchmark investor because it allows her to smooth consumption across future states of nature.

In a frictionless market, the benchmark pricing kernel \tilde{m}^* correctly prices the assets in Z if and only if $E[\tilde{m}^* \tilde{z}] = \pi(\tilde{z})$ for all $\tilde{z} \in Z$ or, alternatively,

$$\begin{aligned} \forall \tilde{x} \in X \quad \text{subject to } \tilde{x} \neq \tilde{0}, E[\tilde{m}^* \tilde{x}] = 0 &\Leftrightarrow E^*[\tilde{x}] = 0 \\ &\Leftrightarrow E^*[\tilde{x}^+ - \tilde{x}^-] = 0 \Leftrightarrow \frac{E^*[\tilde{x}^+]}{E^*[\tilde{x}^-]} = 1, \end{aligned} \quad (6)$$

where $E^*[\cdot]$ denotes the expectation under the risk-adjusted probabilities $p_j^* = p_j u'(c_j^*) / E[u'(\tilde{c}^*)]$ for $j = 1, \dots, S$, and $\tilde{x} = \tilde{x}^+ - \tilde{x}^-$ is the decomposition of a payoff into its positive part $\tilde{x}^+ = \max(\tilde{x}, 0)$ and negative part $\tilde{x}^- = \max(-\tilde{x}, 0)$.

We call $E^*[\tilde{x}^+]$ the *gain*, $E^*[\tilde{x}^-]$ the *loss*, and $E^*[\tilde{x}^+] / E^*[\tilde{x}^-]$ the *gain-loss ratio*.³ The gain-loss ratio is mathematically defined on \mathbb{R}^S (except for $\tilde{0}$) and (i) is always nonnegative, (ii) is equal to $+\infty$ in the positive orthant \mathbb{R}_+^S and finite elsewhere, (iii) is invariant to the multiplication of \tilde{x} by a positive scalar, and (iv) is the inverse of the ratio of the corresponding short position.

The gain-loss ratio summarizes the attractiveness of a zero-cost investment for the benchmark investor. If the gain-loss ratio is equal to one, the investment is fairly priced for the benchmark investor.

³ The loss is related to the lower partial moment of order one suggested by Bawa and Lindenberg (1977) as an alternative measure of risk that could replace the variance in mean-variance analysis. The main differences are that we define loss by taking expectations under the benchmark risk-adjusted probability measure and we do not assume investors to be mean-loss optimizers.

If the gain-loss ratio is above (below) one, the investment offers a good buying (selling) opportunity. 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 increase her holdings in the asset. Equivalently, she risks only half the loss that she would be willing to accept to increase her holdings in the asset.

C. *Dual Formulation in Terms of Pricing Kernels*

Our main duality result relates the existence of high-gain-loss ratio investments to pricing kernels exhibiting extreme deviations from the benchmark pricing kernel.

THEOREM 1.

$$\max_{\substack{\tilde{x} \in X \\ \tilde{x} \neq 0}} \frac{E^*[\tilde{x}^+]}{E^*[\tilde{x}^-]} = \min_{\tilde{m} \in M} \frac{\sup_{j=1, \dots, S} (m_j / m_j^*)}{\inf_{j=1, \dots, S} (m_j / m_j^*)}, \tag{7}$$

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

We prove this equality by demonstrating that both inequalities \leq and \geq must be true. Proving the first inequality in equation (7) is simple. For all $\tilde{x} \in X$ and $\tilde{m} \in M$, we have $E[\tilde{m}\tilde{x}] = 0$, which is equivalent to $E^*[(\tilde{m}/\tilde{m}^*)\tilde{x}] = 0$. Therefore,

$$\begin{aligned} E^*[\tilde{x}^+] \times \inf_{j=1, \dots, S} \frac{m_j}{m_j^*} &\leq E^* \left[\left(\frac{\tilde{m}}{\tilde{m}^*} \right) \tilde{x}^+ \right] \\ &= E^* \left[\left(\frac{\tilde{m}}{\tilde{m}^*} \right) \tilde{x}^- \right] \leq E^*[\tilde{x}^-] \times \sup_{j=1, \dots, S} \frac{m_j}{m_j^*}. \end{aligned} \tag{8}$$

The second inequality in equation (7) is proved in the Appendix for the general case of infinite states. Furthermore, our result would also obtain if we replace the riskless bond with a limited liability asset.

D. *Discussion*

Theorem 1 is similar to the duality result in Hansen and Jagannathan (1991). These authors 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).⁴ The loss (gain) replaces the standard deviation (mean) of the excess payoff as a measure of risk and reward, and the extreme values replace the variance as a measure of dispersion of the pricing kernel. If 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) tradable combinations of state claims.

The main advantage of our duality result for deriving asset pricing implications is that the gain-loss ratio, unlike the Sharpe ratio, characterizes the set of arbitrage and approximate arbitrage opportunities. This is immediately apparent from theorem 1. Ross (1978) shows that the absence of arbitrage opportunities is equivalent to the existence of an admissible pricing kernel that is positive: $0 < \tilde{m} < +\infty$ or, equivalently, $0 < \tilde{m}/\tilde{m}^* < +\infty$. Thus, by theorem 1, a finite bound on the maximum gain-loss ratio precludes the existence of arbitrage opportunities. We strengthen the no-arbitrage restriction $0 < \tilde{m}/\tilde{m}^* < +\infty$ in the following way:

$$\alpha \leq \frac{\tilde{m}}{\tilde{m}^*} \leq \beta, \quad (9)$$

where $0 < \alpha \leq \beta < +\infty$. If we assume without loss of generality that the bounds in equation (9) are binding, the maximum gain-loss ratio, \bar{L} , is just the ratio β/α . Bernardo and Ledoit (1999) show that assuming \bar{L} is finite is equivalent to ruling out approximate arbitrage opportunities in the topological sense of being in the neighborhood of a pure arbitrage opportunity.

IV. Pricing Bounds

In this section, we demonstrate how to use our duality result to derive pricing implications that lie between those of a specific model and the no-arbitrage principle.

A. Compatibility

In our methodology, the benchmark pricing kernel must be positive but does not have to correctly price the basis assets. For example,

⁴ Technically, this result exploits the duality between the payoff space, L^2 , and the pricing kernel space, also L^2 . The infinite-dimensional extension of our result exploits the duality between the payoff space, L^1 , and the pricing kernel space, L^∞ .

we can use a consumption-based benchmark or a risk-neutral benchmark, even if they misprice some or all of the basis assets. To address this issue formally, let B denote the space of portfolio payoffs that can be formed from a set of basis assets with known prices. We assume that B includes the payoff on a limited liability asset. The distinction between the space B and Z is that the prices of payoffs in B are known and provide valuable information for our pricing procedure. This information can be summarized by the linear pricing functional π_B defined on B . Our methodology combines the information contained in the benchmark model with the information in basis asset prices; thus it is reasonable to ask to what extent they are compatible with one another. This is easily measured in our framework by the maximum gain-loss ratio among basis assets, denoted \bar{L}_B . As \bar{L}_B decreases (increases), the benchmark model does a better (poorer) job at pricing basis assets. In the limit $\bar{L}_B = 1$, the benchmark model is perfectly well specified in the sense that it prices all basis assets correctly. The ratio \bar{L}_B can never go below one because if some asset has a gain-loss ratio below one, its corresponding short position has a gain-loss ratio above one. In the other limit, as \bar{L}_B goes to infinity, the benchmark model becomes extremely misspecified in the sense that approximate arbitrage portfolios can be constructed from basis assets. Note that \bar{L}_B is similar in spirit to the measure of model misspecification developed by Hansen and Jagannathan (1997).

B. Economic Assumption

The economic assumption that defines our pricing methodology is given by the following assumption.

ASSUMPTION 1. Excess payoffs have a gain-loss ratio below \bar{L} :

$$\forall \tilde{z} \in Z \text{ subject to } \pi(\tilde{z}) = 0 \text{ and } \tilde{z} \neq \tilde{0}, \quad \frac{E^*[\tilde{z}^+]}{E^*[\tilde{z}^-]} \leq \bar{L}. \quad (10)$$

This assumption expresses the idea that if the benchmark model is reasonable, then high-gain-loss ratio investment opportunities are inconsistent with well-functioning capital markets: if high-gain-loss ratio investments existed, they would be (approximately) arbitrated away. When \bar{L} decreases, we express more confidence in the ability of the benchmark to price nonbasis assets; if \bar{L} increases, we express reluctance to assume anything stronger than the absence of arbitrage. If the benchmark model is misspecified, \bar{L} must be chosen to exceed \bar{L}_B (the maximum gain-loss ratio attainable from the set of basis assets alone).

C. Pricing Implications

Assumption 1 implies the following bounds on the price of a nonbasis asset:

$\forall \tilde{z} \in Z$ subject to $\tilde{z} \notin B$,

$$\max_{\substack{\tilde{b} \in B \\ \frac{E^*[(\tilde{z}-\tilde{b})^+]}{E^*[(\tilde{z}-\tilde{b})^-]} \geq \bar{L}}} \pi_B(\tilde{b}) \leq \pi(\tilde{z}) \leq \min_{\substack{\tilde{b} \in B \\ \frac{E^*[(\tilde{b}-\tilde{z})^+]}{E^*[(\tilde{b}-\tilde{z})^-]} \geq \bar{L}}} \pi_B(\tilde{b}). \quad (11)$$

The \tilde{b} in the bounds are the basis assets that come closest to replicating \tilde{z} from below and from above, respectively. These bounds are the tightest that can be formed for a single asset. In dual terms they can be expressed as

$$\forall \tilde{z} \in Z, \quad \min_{\substack{\tilde{m} \in \mathbb{R}_{++}^S \\ \frac{\sup(\tilde{m}/\tilde{m}^*)}{\inf(\tilde{m}/\tilde{m}^*)} \leq \bar{L} \\ \forall b \in BE[\tilde{m}\tilde{b}] = \pi_B(\tilde{b})}} E[\tilde{m}\tilde{z}] \leq \pi(\tilde{z}) \leq \max_{\substack{\tilde{m} \in \mathbb{R}_{++}^S \\ \frac{\sup(\tilde{m}/\tilde{m}^*)}{\inf(\tilde{m}/\tilde{m}^*)} \leq \bar{L} \\ \forall b \in BE[\tilde{m}\tilde{b}] = \pi_B(\tilde{b})}} E[\tilde{m}\tilde{z}]. \quad (12)$$

The bounds get wider (narrower) as \bar{L} increases (decreases). In the limit as \bar{L} goes to infinity, they converge to the no-arbitrage bounds, which can be expressed either as

$$\forall \tilde{z} \in Z, \quad \max_{\substack{\tilde{b} \in B \\ \tilde{b} \leq \tilde{z} \text{ a.s.}}} \pi_B(\tilde{b}) \leq \pi(\tilde{z}) \leq \min_{\substack{\tilde{b} \in B \\ \tilde{b} \geq \tilde{z} \text{ a.s.}}} \pi_B(\tilde{b}) \quad (13)$$

or as

$$\forall \tilde{z} \in Z, \quad \min_{\substack{\tilde{m} \in \mathbb{R}_{++}^S \\ \forall b \in BE[\tilde{m}\tilde{b}] = \pi_B(\tilde{b})}} E[\tilde{m}\tilde{z}] \leq \pi(\tilde{z}) \leq \max_{\substack{\tilde{m} \in \mathbb{R}_{++}^S \\ \forall b \in BE[\tilde{m}\tilde{b}] = \pi_B(\tilde{b})}} E[\tilde{m}\tilde{z}]. \quad (14)$$

As \bar{L} approaches \bar{L}_B , the bounds generically converge to each other. It is possible to build counterexamples in which the upper and lower bounds are different for $\bar{L} = \bar{L}_B$, but these examples are nongeneric. In general, there exists a unique $\tilde{m}' > 0$ almost surely that correctly prices basis assets and attains the bound

$$\frac{\sup(\tilde{m}'/\tilde{m}^*)}{\inf(\tilde{m}'/\tilde{m}^*)} = \bar{L}_B,$$

in which case both bounds converge to the value $E[\tilde{m}'\tilde{z}]$. The pricing kernel that attains the bound is very special because it is the smallest modification to the benchmark pricing kernel that correctly prices the basis assets.

V. Option Pricing without Intermediate Trading

In this section, we apply our methodology to derive price bounds for an option when there is no intermediate trading. Consequently, the Black-Scholes (1973) dynamic replication arguments do not apply here. Our economy consists of two basis assets: a riskless bond with continuously compounded rate of return $r = \log(1 + r_f)$ and a stock whose continuously compounded return is normally distributed with mean $\mu - (\sigma^2/2)$ and variance σ^2 . Thus the final stock price \tilde{S} is lognormally distributed. There are two dates (date 0 and date t) with no intermediate trading. Our goal is to price a call option with strike price K and time to expiration t .

Merton (1973) shows that we can obtain the Black-Scholes price even without continuous trading by specifying a utility function with constant relative risk aversion (see also Rubinstein 1976; Brennan 1979). Therefore, a natural benchmark is $u'(\tilde{c}^*) = A\tilde{S}^{-\gamma}$ for some positive constants A and γ . The value of the multiplier A does not matter, but there is only one value of the exponent γ that is consistent with the stock and bond parameters μ , σ , and r . It turns out that we do not even need to compute γ . As we showed in Section III, we need to know only the risk-neutral probability, defined as the product of $u'(\tilde{c}^*)/E[u'(\tilde{c}^*)]$ with the true probability. It is a standard result that, under this risk-neutral probability, the continuously compounded stock return is normally distributed with mean $r - (\sigma^2/2)$ and variance σ^2 . The only change is that μ is replaced by r . This property characterizes the risk-adjusted probability measure.

Let S denote the initial stock price, C the initial call option price, and $\tilde{C} = (\tilde{S} - K)^+$ the call option payoff. We shall be interested in the portfolio with w_S shares of stock and w_C call options, which has excess payoff $\tilde{x} = w_S(\tilde{S} - e^{rt}S) + w_C[(\tilde{S} - K)^+ - e^{rt}C]$.

A. Details on Numerical Computations

Recall that the gain is the expectation of a portfolio's excess payoff, computed over states in which the excess payoff is positive. Formally, the gain is $E^*[\tilde{x}\mathbf{1}_{\{\tilde{x}>0\}}]$, where $\mathbf{1}$ denotes the indicator function. The excess payoff \tilde{x} is positive when the stock price \tilde{S} belongs to a certain range. Consider an arbitrary interval for the stock price. If this interval does not contain the strike price K , \tilde{x} is a linear function of \tilde{S} over the whole interval. If it does contain K , this interval can be split into two subintervals at K so that \tilde{x} is linear in \tilde{S} over each subinterval. In summary, the range of values of \tilde{S} in which the excess payoff \tilde{x} is positive can be decomposed into three or fewer intervals so that \tilde{x} is linear in \tilde{S} over each one of them. Therefore, the gain can be

decomposed into the sum of three or fewer terms so that each term is the expectation of a linear function of \tilde{S} over an interval. Each term can be computed in closed form by the following formula:

$$E^*[(\alpha + \beta \tilde{S}) \mathbf{1}_{\{S_1 \leq \tilde{S} < S_2\}}] = \alpha[\Phi(d_1 - \sigma \sqrt{t}) - \Phi(d_2 - \sigma \sqrt{t})] + \beta S e^{rt}[\Phi(d_1) - \Phi(d_2)], \quad (15)$$

where

$$d_i = \frac{\log(S/S_i e^{-rt})}{\sigma \sqrt{t}} + \frac{1}{2} \sigma \sqrt{t}, \quad i = 1, 2. \quad (16)$$

In equations (15) and (16), α and β represent the coefficients of the linear function; S_1 and S_2 are the bounds of the interval, possibly zero or infinity; and $\Phi(\cdot)$ denotes the standard normal cumulative distribution function. It is worth noting that the familiar Black-Scholes price can be obtained directly from equations (15) and (16) by using the risk-neutral valuation formula $E^*[-(K + \tilde{S}) \mathbf{1}_{\{K \leq \tilde{S} < \infty\}}]/e^{rt}$.

This shows how to compute the gain (and by symmetry the loss) in closed form for any portfolio weights w_S and w_C , given initial prices S and C . The next step is to find the weights of the portfolio with the maximum gain-loss ratio, given S and C . Assume without loss of generality that $E^*[\tilde{C} - e^{rt}C] > 0$, and let $w_S = wE^*[\tilde{C} - e^{rt}C]$ and $w_C = 1 - wE^*[\tilde{S} - e^{rt}S]$, where w is a free parameter. As w varies, the expected excess payoff of the portfolio $E^*[\tilde{x}]$ remains constant. We then find the value of w that minimizes the first absolute moment (L_1 norm) of the excess payoff $E^*[\tilde{x}]$:

$$\begin{aligned} & \min_w E^*[\tilde{x}] \\ & \text{subject to } \tilde{x} = w_S(\tilde{S} - e^{rt}S) + w_C(\tilde{C} - e^{rt}C), \\ & \tilde{C} = (\tilde{S} - K)^+, \\ & w_S = wE^*[\tilde{C} - e^{rt}C], \\ & w_C = 1 - wE^*[\tilde{S} - e^{rt}S]. \end{aligned} \quad (17)$$

This is a simple univariate unconstrained convex optimization program. Fast and reliable algorithms abound. It is easy to verify that, at the optimum, the solution \tilde{x} has the maximum gain-loss ratio in the economy. Call it $L(S, C)$.

The final step is to impose a ceiling $\bar{L} > 1$ on the maximum gain-loss ratio in the economy. This ceiling implies upper and lower bounds on the call option price: $\bar{C}(S, \bar{L}) = \max\{C: L(S, C) \leq \bar{L}\}$ and $\underline{C}(S, \bar{L}) = \min\{C: L(S, C) \leq \bar{L}\}$. These bounds are obtained by inverting the function $L(\cdot, \cdot)$ in its second argument. A useful trick

is to work with the Black-Scholes implied volatility rather than the call price itself: although the two approaches are mathematically equivalent, the former approach is numerically better behaved.

An alternative method based on Monte Carlo simulations is as follows. We generate $I = 10,000$ draws from the lognormal distribution of the terminal stock price under the Black-Scholes risk-adjusted probability measure. Let y_i denote standard normal variates independently and identically distributed across $i = 1, \dots, I$. The final value of the stock in the i th simulation is $S_i = S \exp[rt - (\sigma^2/2)t + \sigma \sqrt{t} y_i]$. The corresponding payoff on the call option is $C_i = (S_i - K)^+$. Using these simulated payoffs, we can approximate the Black-Scholes price of the call option by

$$C = \frac{1}{I} \sum_{i=1}^I (S_i - K)^+ e^{-rt}.$$

The bounds implied by assumption 1 on the price of this option are derived from equation (11):

$$\begin{aligned} \max_{\substack{w_0, w_1 \in \mathbb{R} \\ b_i = w_0 + w_1 S_i}} w_0 e^{-rt} + w_1 S &\leq C \leq \min_{\substack{w_0, w_1 \in \mathbb{R} \\ b_i = w_0 + w_1 S_i}} w_0 e^{-rt} + w_1 S, \\ \frac{(1/I) \sum_{i=1}^I (b_i - C_i)^+}{(1/I) \sum_{i=1}^I (b_i - C_i)^-} \geq \bar{L} & \qquad \frac{(1/I) \sum_{i=1}^I (C_i - b_i)^+}{(1/I) \sum_{i=1}^I (C_i - b_i)^-} \geq \underline{L} \end{aligned} \tag{18}$$

where b_i is the payoff in the i th simulation of the replicating portfolio of basis assets with weight w_0 on the risk-free 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 on the left-hand-side maximization program as

$$\frac{1}{I} \sum_{i=1}^I (b_i - C_i)^+ \geq \frac{1}{I} \sum_{i=1}^I (b_i - C_i)^- \times \bar{L}$$

and do the same thing for the right-hand-side minimization program.

B. Resulting Bounds

The call price must lie between \bar{C} and \underline{C} , the right-hand and left-hand sides of equation (18), respectively, or else approximate arbitrage opportunities would exist. The location of the bounds is determined by the benchmark model, the tightness of the bounds is decreasing in the threshold \bar{L} , and the bounds lie strictly between

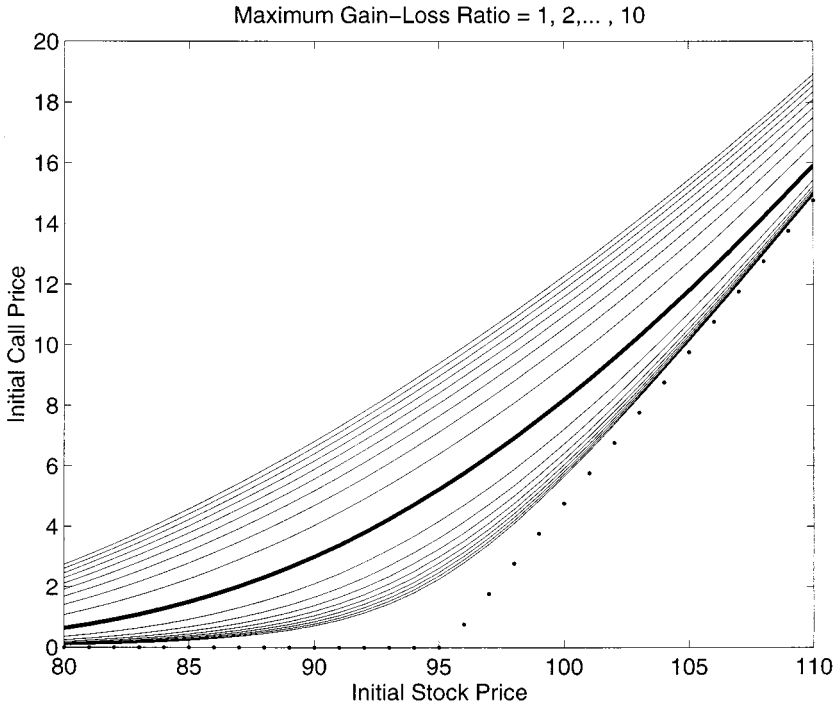


FIG. 1.—Bounds on call option price. The option has $t = 1$ year to expiration and strike price $K = 100$. The one-year risk-free rate of return is 5 percent and the standard deviation of the continuously compounded rate of return on the stock is $\sigma = 0.1409$ per year. These values are calibrated to mimic a broad-based stock market index. The benchmark model is chosen to yield the Black-Scholes price given by the thick line. The thin lines represent the upper and lower bounds obtained by ruling out approximate arbitrage. The dotted line represents the lower bound obtained by ruling out pure arbitrage (the corresponding upper bound is too high to be included in the picture).

the Black-Scholes price and the no-arbitrage bounds. As $\bar{L} \searrow 1$, \bar{C} and \underline{C} converge to the Black-Scholes price; as $\bar{L} \nearrow +\infty$, they converge to the no-arbitrage bounds.

The call option to be priced has time to maturity t of one year and strike price $K = 100$. We use parameter values calibrated to mimic a broad-based U.S. market index and bond data. The one-year risk-free interest rate is 5 percent, and the standard deviation of the stock is $\sigma = 0.1409$ per year. Figure 1 plots the gain-loss bounds on call option prices for the thresholds $\bar{L} = 1, 2, \dots, 10$. For $\bar{L} = 1$, the upper and lower bounds are equal to each other and to the Black-Scholes price. The bounds get wider as \bar{L} increases. By imposing reasonable restrictions on the maximum gain-loss ratio, we can dramatically sharpen the no-arbitrage bounds.




FIG. 2.—Approximate arbitrage bounds on the implied volatility. These bounds correspond to the option prices plotted in fig. 1. Implied volatilities were plotted using the Black-Scholes formula. The thick line represents the true volatility of the underlying asset, a constant 0.1409. The thin lines represent the volatilities implied by the upper and lower bounds obtained by ruling out approximate arbitrage.

Figure 2 plots the same data, except that call option prices have been mapped into implied volatilities by inverting the Black-Scholes formula.

In figures 1 and 2, it appears that the bounds are looser near the money. This can be explained by the well-known fact that near-the-money options are the least redundant ones. Intuitively, near-the-money options put a heavy probability weight on the nonlinearity (at K) that stocks and bonds, being linear, cannot reproduce. This is also the standard explanation for why near-the-money options have the highest traded volume.

Also, the upper bound \bar{C} is farther away from the Black-Scholes price than the lower bound \underline{C} . Intuitively, the reason is that the upper no-arbitrage bound \bar{S} is farther away than the lower no-arbitrage bound $\max(\bar{S} - e^{-rt}K, 0)$. Indeed, it is so far away that it does not even fit into figure 1. Since our bounds are, in some sense, interpolations between the no-arbitrage bounds and the Black-Scholes price,

it comes as no surprise that \bar{C} is farther away. There is a practical implication. If a trader does not know the true volatility—perhaps because it changes randomly or because of estimation error—then it may be better to overstate it systematically. The reason is that the penalty, in terms of creating approximate arbitrage opportunities that other traders can exploit, is not symmetric.

Another salient feature is that going from $\bar{L} = 1$ to $\bar{L} = 2$ moves the bounds much more than going from $\bar{L} = 9$ to $\bar{L} = 10$. To see why, remember that the Black-Scholes price minimizes $L(S, C)$. At the minimum, the first-order condition $\partial L / \partial C = 0$ implies that $\partial \bar{C}(S, 1) / \partial \bar{L} = +\infty$ and $\partial \underline{C}(S, 1) / \partial \bar{L} = -\infty$. In words, the bounds are infinitely sensitive to an increase of \bar{L} when $\bar{L} = 1$. This mathematical fact has an important economic interpretation: small errors in selecting the benchmark model can cause relatively large pricing deviations, even in the absence of approximate arbitrage opportunities.

The figures also show that bound sensitivity decreases sharply in \bar{L} . Mathematically, we have

$$\lim_{\bar{L} \rightarrow +\infty} \frac{\partial \bar{C}}{\partial \bar{L}}(S, \bar{L}) = \lim_{\bar{L} \rightarrow +\infty} \frac{\partial \underline{C}}{\partial \bar{L}}(S, \bar{L}) = 0.$$

Once we choose a ceiling \bar{L} high enough to be on the safe side, economically large increases in \bar{L} have little further effect on the bounds. This is important because the choice of the threshold \bar{L} is somewhat arbitrary and open to discussion.

To sum up, a bound on the gain-loss ratio generates pricing implications that are substantially different from model-based pricing and from no-arbitrage pricing. Furthermore, these implications can be relatively robust to differences of opinion over the level of the maximum gain-loss ratio.

C. *Modeler's Choices*

Our methodology involves several choices that the modeler must make ex ante in order to obtain pricing implications.

Ceiling on the Maximum Gain-Loss Ratio

The parameter \bar{L} controls the trade-off between the precision of a specific benchmark pricing model and the robustness of no-arbitrage bounds. Choosing the value of \bar{L} is difficult, but it is done implicitly by any modeler who assumes a specific model (implicitly choosing $\bar{L} = 1$) or insists on nothing stronger than a no-arbitrage

assumption (implicitly choosing \bar{L} arbitrarily large). In our opinion, neither one of these commonly made choices is optimal for deriving useful pricing implications in practice. Exact pricing models are so sensitive to misspecification that small values of \bar{L} are preferable to $\bar{L} = 1$. The no-arbitrage principle is so weak that it is always better to use a large but finite value of \bar{L} . Thus the problem of choosing \bar{L} is less daunting when compared against the alternatives.

One possibility is to let \bar{L} represent the maximum gain-loss ratio implied by a well-known mispricing puzzle, such as the historically high equity risk premium, and use it as a general indicator of worst-case error in economic models. For example, we compared the return on the value-weighted U.S. stock market index from the Center for Research in Security Prices to the risk-free rate on U.S. Treasury bills reported by the same source. When we use a benchmark model assuming logarithmic utility, the gain-loss ratio of the stock market index is $\bar{L} = 2.6$.⁵ This is representative of the magnitude of violation of the benchmark model in a classic asset pricing puzzle.

Alternatively, one could survey arbitrageurs to find out how large the prospect of gain must be relative to potential losses to persuade them to take the position.

Bawa and Lindenberg (1977) showed that in the special case of normally distributed returns, there exists a one-to-one mapping between the Sharpe ratio and the gain-loss ratio, under a risk-neutral benchmark. If the normally distributed excess payoff \tilde{x} has Sharpe ratio S , then

$$\frac{E[\tilde{x}^+]}{E[\tilde{x}^-]} = \frac{\phi(S) + S\Phi(S)}{\phi(-S) - S\Phi(-S)}, \quad (19)$$

where $\phi(\cdot)$ and $\Phi(\cdot)$ denote the standard normal probability density function and cumulative distribution function, respectively. This result is most useful when the asset being priced has returns that are close to normally distributed. This mapping has the advantage of building on our familiarity with the Sharpe ratio; for example, a Sharpe ratio of 0.3 maps into an \bar{L} of approximately 2. The mapping between the gain-loss ratio and the Sharpe ratio under normality is shown in figure 3. Note that the mapping between the gain-loss ratio has the same shape as the exponential function.

Finally, every model makes unrealistic assumptions, so it is important to understand which of a given model's pricing implications are sensitive to misspecification error and which are robust. A model's

⁵ To proxy for consumption, we used growth in personal consumption expenditures of nondurable goods reported by the Federal Reserve, using monthly data compounded annually from 1959 to 1997.

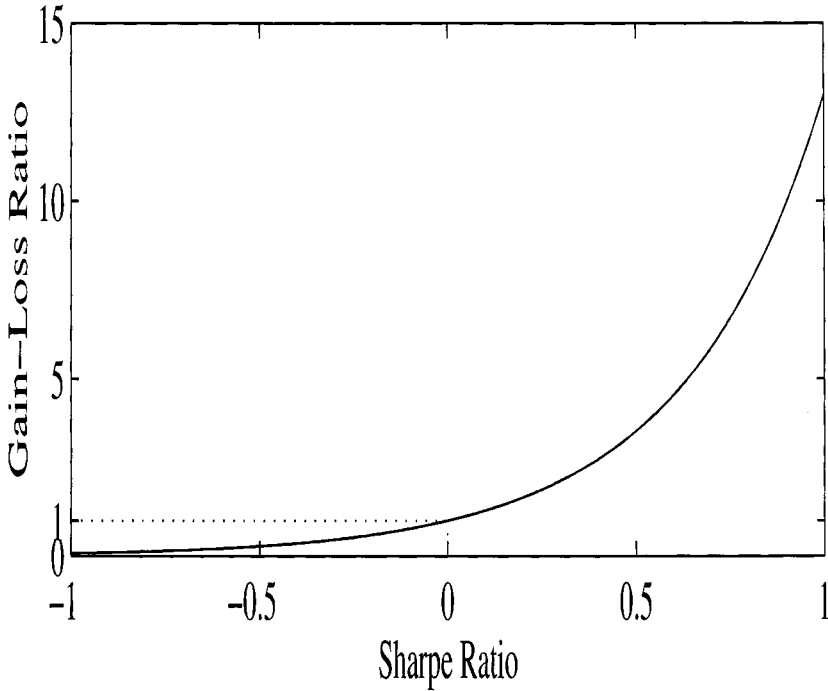


FIG. 3.—Mapping between the gain-loss ratio, using a risk-neutral benchmark, and the Sharpe ratio under normality.

pricing implications are robust when a violation of its assumptions creates an arbitrage or near-arbitrage opportunity. We can address the issue of robustness quantitatively by computing price bounds for different values of \bar{L} : price bounds that are narrow for large values of \bar{L} suggest that the benchmark model's predictions are robust.

Basis Assets

If there are assets with known prices and payoffs that nearly mimic the assets to be priced, including them among the basis assets will yield tighter pricing bounds. For example, suppose that there are two securities with identical payoffs except that they differ slightly in one state, which occurs with small probability. If the prices of these two securities are very different, then they create a near-arbitrage opportunity, and the price of one security will be very valuable information for determining reasonable prices for the other.

Also, the modeler should include only basis assets available to the investor in question. For example, suppose that we are interested

in obtaining bounds on the value of a stock option to a corporate executive. If the executive can sell short a market index to hedge the firm's market risk, an index option may be an appropriate asset to include in the analysis. However, if the executive is not permitted to buy and sell the firm's stock freely or is not permitted to short the firm's stock (e.g., because of insider trading restrictions), then it is inappropriate to include the firm's stock as a basis asset since the executive cannot freely form portfolios of the stock and option. By contrast, if we wanted to determine the value of such options to the firm, we would include the firm's stock as a basis asset.

Benchmark Pricing Kernel

To obtain a benchmark pricing kernel, one can specify a utility function u and a consumption plan \tilde{c}^* and use equation (5). But pricing kernels implied by representative agent models with time-additive utility do a poor job of explaining observed stock prices when aggregate consumption data are used (e.g., the equity premium puzzle of Mehra and Prescott [1985]). Alternatively, the modeler can choose a stock market index to proxy for aggregate consumption in order to fit observed data better. Recall, however, that there is no requirement that the benchmark pricing kernel correctly price all basis assets. For example, one could use a risk-neutral benchmark even though it does not price all basis assets. A gain-loss ratio restriction using this benchmark will yield robust pricing implications if investors' willingness to pay for consumption in different states does not vary considerably.

It is sometimes desirable to choose a benchmark model that prices basis assets correctly as in Rubinstein (1976). This is generally feasible when the benchmark is chosen from a family of models with enough free parameters to match the number of basis assets. The benchmark pricing kernel need not be inspired by theory; rather, it can be obtained from some parametric (e.g., Backus, Foresi, and Telmer 1996) or nonparametric (e.g., Bansal and Viswanathan 1993; Ait-Sahalia and Lo 1998) estimation technique. But the benchmark pricing kernel must be strictly positive to eliminate the possibility of arbitrage opportunities.

The choice of benchmark pricing kernel should account for the characteristics of the investor in question. Consider our earlier example of obtaining bounds on the value of an executive's portfolio of stock options. 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 that does not price idiosyncratic risk.

D. Varying the Benchmark

As mentioned earlier, our method does not require the benchmark model to price basis assets correctly. We now illustrate this by using a benchmark model that misprices most risky assets: the risk-neutral benchmark. In this case the benchmark risk-adjusted probability measure is equal to the true probability measure. Even when investors are risk-averse, this may still be a useful benchmark if we have little knowledge of what states of nature are “good” (low marginal utility) versus “bad” (high marginal utility).

Let us continue with the option pricing example. The first step is to specify the stock’s drift. We arbitrarily set it equal to $\mu = 0.1123$, calibrated to mimic a U.S. stock market index. We then redo the computations in Section VA with this drift instead. The maximum gain-loss ratio among basis assets is attained by a long position in the stock and is equal to $\bar{L}_B = 3.7927$. Setting $\bar{L} = \bar{L}_B$ gives a unique price for the call option. This price can be computed by taking the inner product of the call option payoff with the pricing kernel that takes the value \bar{m} when the excess return on the stock is negative and \underline{m} when the excess return is positive. The values \bar{m} and \underline{m} are derived from the condition that the bond and the stock be correctly priced, by solving a system of two equations and two unknowns. The resulting price can be compared to the Black-Scholes price.

The next step is to choose a ceiling on the maximum gain-loss ratio. We choose $\bar{L} = \bar{L}_B + 2$ and compare the resulting bounds against those obtained in subsection B for $\bar{L} = 1 + 2 = 3$ (in subsection B we had implicitly $\bar{L}_B = 1$). This is plotted in figure 4. Notice that varying the benchmark shifts the location of the bounds up or down but does not affect their width.

If we vary the benchmark over all arbitrage-free pricing models, then the bounds will sweep over all arbitrage-free prices between $\max(0, S - Ke^{-rT})$ and S , which is not particularly helpful. To narrow the bounds, one must choose a particular benchmark.

E. Comparison with Other Methods

Ledoit (1995) defines a “ δ -arbitrage opportunity” as a portfolio with a Sharpe ratio higher than a prespecified threshold δ and derives the implications of precluding δ -arbitrage opportunities for stock returns in the context of the arbitrage pricing theory (Ross 1976). The problem with this approach is that a Sharpe ratio constraint




FIG. 4.—Gain-loss bounds for different benchmarks. The dashed lines come from the risk-neutral benchmark, and the solid lines come from the power utility benchmark that correctly prices the stock and the bond. The inner lines correspond to $\bar{L} = \bar{L}_B$, where \bar{L}_B denotes the maximum gain-loss ratio attainable by constructing portfolios of the stock and bond. For the risk-neutral benchmark, $\bar{L}_B = 3.7927$, and for the power utility benchmark, $\bar{L}_B = 1$ by construction. If $\bar{L} = \bar{L}_B$, the upper and lower bounds coincide. The outer lines correspond to $\bar{L} = \bar{L}_B + 2$.

does not form a *neighborhood* of the set of arbitrage opportunities: in fact, it does not even rule out arbitrage opportunities since arbitrage opportunities can have a Sharpe ratio below any positive threshold δ . Intuitively, this occurs because *upside risk* increases the denominator of the Sharpe ratio. The problem is most pronounced when payoffs are heavily skewed, as with out-of-the-money options and lottery tickets.

A striking example is a lottery ticket costing one cent today with a payoff of \$50 billion next year with probability 10 percent, and nothing otherwise:

$$0.01 \begin{cases} \nearrow 50,000,000,000 & \text{with probability } .1 \\ \searrow \\ 0 & \text{with probability } .9. \end{cases} \quad (20)$$

This lottery has a very low Sharpe ratio, 0.33 (below a U.S. stock market index), yet it is obviously a very attractive investment opportunity for any reasonable individual. This apparent paradox is related to the fact that mean-variance preferences, from which the Sharpe ratio is derived, require a quadratic utility function (except for particular distributions of payoffs such as the Gaussian), which displays satiation. In fact, Dybvig and Ingersoll (1982) proved that the capital asset pricing model (CAPM) admits arbitrage opportunities if markets are complete.

Cochrane and Saá-Requejo (2000, this issue) derive pricing bounds on derivatives by ruling out the existence of “good deals” (investment opportunities with high Sharpe ratios) and arbitrage opportunities. As in 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 that the pricing kernel is nonnegative. Nonetheless, because the pricing kernel can take arbitrarily low positive values, this does not rule out some opportunities close to arbitrage.

The practical relevance of this theoretical point is seen by comparing figure 1 with the corresponding figure in the paper by Cochrane and Saá-Requejo. In the out-of-the-money region of figure 1 (low initial stock price), the gain-loss approach yields a strictly positive lower bound on the option price, whereas the “good-deal” bound yields a lower bound below zero. When the good-deal bound is augmented with the no-arbitrage restriction, it yields a lower bound of exactly zero. The gain-loss framework recognizes that out-of-the-money options can be attractive if they are sufficiently cheap; however, they will always have too much *upside risk* to be perceived as attractive *at any positive price* by the variance bound.

The method of Cochrane and Saá-Requejo works extremely well over short trading horizons, because at these horizons payoffs are almost normally distributed and are not shaped like out-of-the-money options. In addition, the authors show that their approach is remarkably tractable in dynamic problems. It is straightforward to add our gain-loss ratio restriction, under a risk-neutral benchmark, to their dynamic framework.

Hansen and Jagannathan (1997) define two general quadratic measures of distance: δ from a given benchmark pricing kernel to the nearest admissible pricing kernel and δ^+ to the nearest positive admissible pricing kernel. Both distances represent nonparametric measures of the degree of misspecification of the benchmark model. In practice, with nonskewed returns like the stocks and bonds that Hansen and Jagannathan use, the positivity constraint makes little difference, and δ^+ gives results qualitatively similar to δ . However,

with data involving heavily skewed returns, there could be a substantial difference: in such cases it might be worth looking at an alternative approach based on the gain-loss ratio for the reasons outlined above. This could be interesting, for example, when exploring the volatility smile.

Several other duality results in the literature could, in principle, be used to derive asset price bounds. It is important to note, however, that these results were not presented with this application in mind. Snow (1991) generalizes the Hansen-Jagannathan (1991) bounds by deriving restrictions on the q th moment of the pricing kernel $E[\tilde{m}^q]^{1/q}$, for $1 < q < \infty$. Another restriction is due to Stutzer (1995), who shows that restricting the maximum expected utility attainable by a constant absolute risk aversion investor is equivalent to restricting the entropy of the pricing kernel $E[\tilde{m} \log(\tilde{m})]$. Finally, Bansal and Lehmann (1997) show that restricting the maximum expected utility that can be attained by an investor with logarithmic utility is equivalent to restricting $E[-\log(\tilde{m})]$. Bernardo and Ledoit (1999) show that imposing these restrictions does not prevent state prices from being either arbitrarily close to zero or arbitrarily high; thus the implied pricing bounds would permit approximate arbitrage opportunities.

Constantinides (1994) uses second-order stochastic dominance arguments to derive option price bounds. The key assumption he makes is that the marginal utility of consumption is decreasing in the price of the underlying security. Thus, if a zero-cost portfolio, which may include options positions, has positive (negative) payoffs when the price of the underlying security is low (high), then it hedges consumption risk, and it must be true that, in equilibrium, the expectation of the portfolio payoffs is nonpositive. This condition holds for any increasing and concave utility function; thus the pricing bounds are robust to misspecification error. When prior knowledge about the specific form of the utility function is used, the price bounds could be tightened.

VI. Conclusion

This paper derives a new duality result relating the extreme values of the pricing kernel to the existence of approximate arbitrage opportunities, measured by the “gain-loss ratio.” By precluding the existence of approximate arbitrage opportunities, we impose restrictions on the set of admissible pricing kernels that can be used to price assets. Our notion of approximate arbitrage, high gain-loss ratios, depends on a set of benchmark preferences. While the absence

of arbitrage requires only weak assumptions about investor preferences (i.e., monotonicity), the notion of approximate arbitrage requires stronger assumptions because any nonarbitrage portfolio can be supported in equilibrium by some set of investor preferences.

We demonstrate the implications of a gain-loss ratio restriction by computing bounds on the price of options on a stock when there is no intermediate trading. 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. Thus our method offers a general way to chart the middle ground between a specific asset pricing model and no arbitrage. The optimal trade-off between the precision of a specific model and the robustness of the no-arbitrage principle often lies strictly between the two extremes.

Most of the interesting analysis that has been done in the mean-variance framework can be replicated in the gain-loss framework. For example, the standard quadratic portfolio selection problem of minimizing variance subject to attaining a certain level of expected return can be restated as a linear program to minimize loss subject to meeting a specified level of gain. Furthermore, a CAPM-like equilibrium pricing model can be derived (Bawa and Lindenberg 1977). Gain-loss is similar to mean-variance for close to normally distributed payoffs, but more consistent with no arbitrage for heavily skewed payoffs such as out-of-the-money options.

There are many other practical asset pricing problems that can be addressed with the gain-loss ratio restriction. For example, real options are difficult to value using arbitrage methods since 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 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. The dynamic replication argument implicit in the Black-Scholes approach also fails when valuing executive stock options because executives cannot freely trade in the underlying stock or options. To get useful price bounds the modeler could (i) account for the fact that the pricing kernel, which is relevant to the executive, prices idiosyncratic risks that he cannot diversify away; and (ii) include in the analysis the ability of the executive to trade some assets that are imperfect hedges for some firm risks, for example, index options. Finally, the gain-loss approach provides an economically meaningful way to evaluate mutual fund perfor-

mance: according to the attractiveness of the portfolio for some benchmark investor.⁶

Appendix

Proof of Theorem 1

We shall prove theorem 1 in the infinite-state case. We shall need to introduce some notation. Let $L^{1*} = \{\tilde{z}: \|\tilde{z}\| < \infty\}$ denote the space of all possible payoffs, where $\|\tilde{z}\| = E^*[\|\tilde{z}\|]$. The positive orthant of L^{1*} is given by $L_+^{1*} = \{\tilde{z} \in L^{1*}: \tilde{z} \geq 0 \text{ a.s. and } \tilde{z} \neq \tilde{0}\}$. Let $L^{\infty*} = \{\tilde{m}: \sup(\tilde{m}/\tilde{m}^*) < \infty\}$ denote the dual of L^{1*} and $L_{++}^{\infty*} = \{\tilde{m} \in L^{\infty*}: \tilde{m} > 0 \text{ a.s.}\}$ denote the strict positive orthant of $L^{\infty*}$ containing pricing kernels that do not permit arbitrage opportunities.

First, we show that we can assume, without loss of generality, that the benchmark investor is risk-neutral.

LEMMA 1. If equation (7) holds in the special case in which \tilde{m}^* is a constant, then it holds for any strictly positive \tilde{m}^* .

Proof of lemma 1. When \tilde{m}^* is a constant, equation (7) can be rewritten as

$$\max_{\substack{\tilde{x} \in X \\ \tilde{x} \neq \tilde{0}}} \frac{E[\tilde{x}^+]}{E[\tilde{x}^-]} = \min_{\tilde{m} \in M} \frac{\sup \tilde{m}}{\inf \tilde{m}}. \tag{A1}$$

Now let \tilde{m}^* denote any strictly positive benchmark pricing kernel. If we use equation (A1) under the $*$ probability, we obtain

$$\max_{\substack{\tilde{x} \in X \\ \tilde{x} \neq \tilde{0}}} \frac{E^*[\tilde{x}^+]}{E^*[\tilde{x}^-]} = \min_{\tilde{m}^\dagger \in M^\dagger} \frac{\sup \tilde{m}^\dagger}{\inf \tilde{m}^\dagger}, \tag{A2}$$

where $M^\dagger = \{\tilde{m}^\dagger \in L_{++}^{\infty*}: \forall \tilde{z} \in Z E^*[\tilde{m}^\dagger \tilde{z}] = \pi(\tilde{z})\}$. The difference between M^\dagger and M is that the elements of M^\dagger represent $\pi(\cdot)$ through expectation under the $*$ probability. Note that $\tilde{m} \in M$ if and only if $\tilde{m}^\dagger = E[\tilde{m}^*] \tilde{m} / \tilde{m}^* \in M^\dagger$. Therefore,

$$\min_{\tilde{m}^\dagger \in M^\dagger} \frac{\sup \tilde{m}^\dagger}{\inf \tilde{m}^\dagger} = \min_{\tilde{m} \in M} \frac{\sup(E[\tilde{m}^*] \tilde{m} / \tilde{m}^*)}{\inf(E[\tilde{m}^*] \tilde{m} / \tilde{m}^*)} = \min_{\tilde{m} \in M} \frac{\sup(\tilde{m} / \tilde{m}^*)}{\inf(\tilde{m} / \tilde{m}^*)}. \tag{A3}$$

Bringing equations (A2) and (A3) together proves lemma 1. Q.E.D.

To begin, we shall treat separately the case in which there are arbitrage opportunities.

⁶ In this sense, the gain-loss approach is related to the positive period weighting measure in Grinblatt and Titman (1990) when their weights are interpreted as marginal utilities.

LEMMA 2. The following three statements are equivalent: (i) The pricing functional $\pi(\cdot)$ admits arbitrage opportunities, (ii)

$$\max_{\substack{\tilde{x} \in X \\ \tilde{x} \neq \tilde{0}}} \frac{E^*[\tilde{x}^+]}{E^*[\tilde{x}^-]} = +\infty,$$

and (iii)

$$\inf_{\tilde{m} \in M} \frac{\sup(\tilde{m})}{\inf(\tilde{m})} = +\infty.$$

Proof of lemma 2. Part i implies part ii by definition. Part ii implies part iii by inequality (8). If part iii holds, then there does not exist a strictly positive pricing kernel, which implies the existence of an arbitrage opportunity; therefore, part i holds. Q.E.D.

Until the end of the Appendix, we shall assume that there are no arbitrage opportunities. To complete our proof we must show the inequality

$$\max_{\substack{\tilde{x} \in X \\ \tilde{x} \neq \tilde{0}}} \frac{E[\tilde{x}^+]}{E[\tilde{x}^-]} \geq \min_{\tilde{m} \in M} \frac{\sup(\tilde{m})}{\inf(\tilde{m})}. \quad (\text{A4})$$

Define

$$h = \max_{\substack{\tilde{x} \in X \\ \tilde{x} \neq \tilde{0}}} \frac{E[\tilde{x}^+]}{E[\tilde{x}^-]};$$

h is finite because we have ruled out the case $h = +\infty$ treated in lemma 2. Furthermore, $h \geq 1$ since any portfolio with a gain-loss ratio below one has a short position with a gain-loss ratio above one.

Consider the set $K = \{\tilde{y} \in L^1: \tilde{y} \neq \tilde{0}, E[\tilde{y}^+]/E[\tilde{y}^-] > h\}$ of random variables with a gain-loss ratio above h . We can rewrite it as $K = \{\tilde{y} \in L^1: \theta(\tilde{y}) < 0\}$, where the function

$$\theta: L_1 \rightarrow \mathbb{R}, \quad (\text{A5})$$

$$\tilde{y} \mapsto (h - 1)E[\tilde{y}^-] - E[\tilde{y}] \quad (\text{A6})$$

is convex and continuous. This implies that the set K is convex and open. In particular, the interior of K is K itself, which is nonempty. Note also that X is convex and does not intersect with K by the definition of h . Therefore, by Eidelheit's version of the separating hyperplane theorem (Luenberger 1969, p. 133), there exists a continuous linear functional $\psi: L_1 \rightarrow \mathbb{R}$ such that, for all $\tilde{x} \in X$, for all $\tilde{y} \in K$, $\psi(\tilde{x}) < \psi(\tilde{y})$. By the Riesz representation theorem, there exists a random variable $\tilde{m} \in L^\infty$ that represents ψ , that is, for all $\tilde{y} \in L^1$, $\psi(\tilde{y}) = E[\tilde{m}\tilde{y}]$. Thus \tilde{m} verifies

$$\forall \tilde{x} \in X \forall \tilde{y} \in K, \quad E[\tilde{m}\tilde{x}] < E[\tilde{m}\tilde{y}]. \quad (\text{A7})$$

First, substituting $\tilde{x} = \tilde{0}$ into equation (A7) yields, for all $\tilde{y} \in K$, $E[\tilde{m}\tilde{y}] > 0$. Since $L_+^1 \subset K$, we have, for all $\tilde{y} \in L_+^1$, $E[\tilde{m}\tilde{y}] > 0$. It implies that $\tilde{m} > 0$ almost surely. Second, suppose that there existed $\tilde{x} \in X$ with $E[\tilde{m}\tilde{x}] \neq 0$. By taking an $\alpha > 0$ sufficiently large, we would be able to construct $\tilde{x}' =$

$\alpha E[\tilde{m}\tilde{x}] \tilde{x} \in X$ that would verify $E[\tilde{m}\tilde{x}'] > E[\tilde{m}\tilde{y}]$, in violation of equation (A7). Therefore, for all $\tilde{x} \in X$, $E[\tilde{m}\tilde{x}] = 0$. Third, if \tilde{m} was replaced by $\tilde{m} \times E[\tilde{m}]/(1 + r_f)$, equation (A7) would still hold. Therefore, we can assume without loss of generality that \tilde{m} correctly prices the risk-free asset. When these three results are brought together, \tilde{m} is an admissible positive pricing kernel: $\tilde{m} \in M$.

Let $\bar{m} = \sup(\tilde{m})$ and $\underline{m} = \inf(\tilde{m})$. Then for all $\epsilon > 0$, we have $\Pr\{\tilde{m} \geq \bar{m} - \epsilon\} > 0$ and $\Pr\{\tilde{m} \leq \underline{m} + \epsilon\} > 0$. Fix $\epsilon > 0$ and define the random variable

$$\tilde{y}_\epsilon = \frac{(h + \epsilon) \times \mathbf{1}_{\{\underline{m} + \epsilon \geq \tilde{m}\}}}{\Pr\{\underline{m} + \epsilon \geq \tilde{m}\}} - \frac{\mathbf{1}_{\{\tilde{m} \geq \bar{m} - \epsilon\}}}{\Pr\{\tilde{m} \geq \bar{m} - \epsilon\}},$$

where $\mathbf{1}$ denotes the indicator function of an event. This random variable has gain $E[\tilde{y}_\epsilon^+] = h + \epsilon$, loss $E[\tilde{y}_\epsilon^-] = 1$, and gain-loss ratio $E[\tilde{y}_\epsilon^+] \div E[\tilde{y}_\epsilon^-] = h + \epsilon$; therefore, $\tilde{y}_\epsilon \in K$. Then equation (A7) yields

$$\begin{aligned} E[\tilde{m}\tilde{y}_\epsilon] &\geq 0, \\ E\left[\tilde{m}\left[\frac{(h + \epsilon) \times \mathbf{1}_{\{\underline{m} + \epsilon \geq \tilde{m}\}}}{\Pr\{\underline{m} + \epsilon \geq \tilde{m}\}} - \frac{\mathbf{1}_{\{\tilde{m} \geq \bar{m} - \epsilon\}}}{\Pr\{\tilde{m} \geq \bar{m} - \epsilon\}}\right]\right] &\geq 0, \\ (h + \epsilon) \times E[\tilde{m}|\underline{m} + \epsilon \geq \tilde{m}] - E[\tilde{m}|\tilde{m} \geq \bar{m} - \epsilon] &\geq 0, \\ (h + \epsilon) \times (\underline{m} + \epsilon) \geq (h + \epsilon) \times E[\tilde{m}|\underline{m} + \epsilon \geq \tilde{m}] \geq E[\tilde{m}|\tilde{m} \geq \bar{m} - \epsilon] \geq \bar{m} - \epsilon, \\ (h + \epsilon) \times (\underline{m} + \epsilon) &\geq \bar{m} - \epsilon, \\ h &\geq \frac{\bar{m} - \epsilon}{\underline{m} + \epsilon} - \epsilon. \end{aligned}$$

Since this is true for any $\epsilon > 0$, we have

$$h \geq \lim_{\epsilon \searrow 0} \left(\frac{\bar{m} - \epsilon}{\underline{m} + \epsilon} - \epsilon \right) = \frac{\bar{m}}{\underline{m}};$$

hence,

$$\max_{\substack{\tilde{x} \in X \\ \tilde{x} \neq 0}} \frac{E[\tilde{x}^+]}{E[\tilde{x}^-]} = \min_{\tilde{m} \in M} \frac{\sup \tilde{m}}{\inf \tilde{m}}. \tag{A8}$$

References

Ait-Sahalia, Yacine, and Lo, Andrew W. "Nonparametric Estimation of State-Price Densities Implicit in Financial Asset Prices." *J. Finance* 53 (April 1998): 499-547.

Backus, David; Foresi, Silverio; and Telmer, Chris. "Affine Models of Currency Pricing." Manuscript. New York: New York Univ., Stern School Bus., 1996.

Bansal, Ravi, and Lehmann, Bruce N. "Growth-Optimal Portfolio Restrictions on Asset Pricing Models." *Macroeconomic Dynamics* 1, no. 2 (1997): 333-54.

- Bansal, Ravi, and Viswanathan, S. "No Arbitrage and Arbitrage Pricing: A New Approach." *J. Finance* 48 (September 1993): 1231-62.
- Bawa, Vijay S., and Lindenberg, Eric B. "Capital Market Equilibrium in a Mean-Lower Partial Moment Framework." *J. Financial Econ.* 5 (November 1977): 189-200.
- Bernardo, Antonio E., and Ledoit, Olivier. "Approximate Arbitrage." Manuscript. Los Angeles: Univ. California, Anderson Grad. School Management, 1999.
- Black, Fischer, and Scholes, Myron S. "The Pricing of Options and Corporate Liabilities." *J.P.E.* 81 (May/June 1973): 637-54.
- Brennan, Michael J. "The Pricing of Contingent Claims in Discrete Time Models." *J. Finance* 34 (March 1979): 53-68.
- Cochrane, John, and Saá-Requejo, Jesús. "Beyond Arbitrage: Good-Deal Asset Price Bounds in Incomplete Markets." *J.P.E.* 108 (February 2000): 79-119.
- Constantinides, George. "Option Pricing Bounds Based on Stochastic Dominance." Manuscript. Chicago: Univ. Chicago, Grad. School Bus., 1994.
- Dybvig, Philip H., and Ingersoll, Jonathan E., Jr. "Mean-Variance Theory in Complete Markets." *J. Bus.* 55 (April 1982): 233-51.
- Grinblatt, Mark, and Titman, Sheridan D. "Portfolio Performance Evaluation: Old Issues and New Insights." *Rev. Financial Studies* 2, no. 3 (1990): 393-421.
- Hansen, Lars Peter, and Jagannathan, Ravi. "Implications of Security Market Data for Models of Dynamic Economies." *J.P.E.* 99 (April 1991): 225-62.
- . "Assessing Specification Errors in Stochastic Discount Factor Models." *J. Finance* 52 (June 1997): 557-90.
- Ledoit, Olivier. "Essays on Risk and Return in the Stock Market." Ph.D. dissertation, Massachusetts Inst. Tech., 1995.
- Luenberger, David G. *Optimization by Vector Space Methods*. New York: Wiley, 1969.
- Mehra, Rajnish, and Prescott, Edward C. "The Equity Premium: A Puzzle." *J. Monetary Econ.* 15 (March 1985): 145-61.
- Merton, Robert C. "Theory of Rational Optional Pricing." *Bell J. Econ. and Management Sci.* 4 (Spring 1973): 141-83.
- Ross, Stephen A. "The Arbitrage Theory of Capital Asset Pricing." *J. Econ. Theory* 13 (December 1976): 341-60.
- . "A Simple Approach to the Valuation of Risky Streams." *J. Bus.* 51 (July 1978): 453-75.
- Rubinstein, Mark. "The Valuation of Uncertain Income Streams and the Pricing of Options." *Bell J. Econ.* 7 (Autumn 1976): 407-25.
- Snow, Karl N. "Diagnosing Asset Pricing Models Using the Distribution of Asset Returns." *J. Finance* 46 (July 1991): 955-83.
- Stutzer, Michael. "A Bayesian Approach to Diagnosis of Asset Pricing Models." *J. Econometrics* 68 (August 1995): 367-97.