

Minimax Asset Price Bounds in Incomplete Markets

Antonio S. Mello

School of Business

University of Wisconsin, Madison

Unyong Pyo

Faculty of Business

Brock University, Canada

Minimax Asset Price Bounds in Incomplete Markets

Antonio S. Mello¹

School of Business

University of Wisconsin, Madison

Unyong Pyo²

Faculty of Business

Brock University, Canada

April 2006

Abstract

This paper develops an approach to tighten the bounds on asset pricing in an incomplete market that combines no-arbitrage pricing and preference-based pricing, and the approach is applied to call options without dynamic rebalancing. With the no-arbitrage pricing, it is straightforward to obtain the initial bounds, which are too wide to be of practical uses. By accepting that investors exhibit risk aversion from benchmark pricing kernels, it is possible to narrow the bounds considerably. Using the minimax deviation implicit in the parameters, one can restrict further the set of plausible values for call options on a stock.

JEL Classification: G12; D52

Keywords: Incomplete Markets; Real Option; Minimax Deviation

¹Corresponding author: Antonio S. Mello, School of Business, University of Wisconsin, Madison, Wisconsin. Tel: (608) 263-3423, E-mail: amello@bus.wisc.edu. Antonio S. Mello acknowledges financial support from the Aschenbrener Fellowship at the School of Business, University of Wisconsin, Madison.

²Unyong Pyo, Faculty of Business, Brock University, St. Catharines, Ontario, L2S 3A1, Canada. Tel: (905) 688-5550 ext. 3147, Fax: (905) 688-9779, E-mail: upyo@brocku.ca.

1 Introduction

In a complete market, the value of an option is derived from the prices of other assets. Continuous trading in the underlying assets makes it possible to mimic the option's payoffs, and the value of the option is equal to the price of the replicating portfolio in the absence of arbitrage opportunities. The value of the option does not require specifying investors' attitudes toward risk, and is obtained by discounting at the risk free rates the option's expected payoff computed with state-contingent probabilities, which correctly price assets in a risk-neutral economy.

In many instances, attempting to replicate option values by trading in a portfolio of basis assets is both impracticable and impossible. It is impracticable because trading cannot be done continuously and is costly. It is impossible because the risk factors that affect an option payoffs are not represented by assets that are traded in the economy. The no-arbitrage approach then provides only bounds on the value of the option. In technical terms, there are multiple elements in the set of admissible martingale measures that correctly price the basis assets. Since the portfolio of basis assets does not hedge the option perfectly, holding the option becomes inherently risky.

There are several approaches to option pricing in incomplete markets. One approach focuses on different optimization criteria applied to the set of equivalent martingale measures, but usually offers little practical use, [Follmer and Schweizer (1991), and Choulli and Stricker (2005)]. Another imposes constraints on the volatility of the pricing kernel or on the gain-loss ratio to construct bounds, outside of which deals are considered to be too good [Cochrane and Saa-Requejo (2000), Staum (2004), and Bernardo and Ledoit (2000)]. A third approach makes limited assumptions about the risk preferences of an investor and combines these assumptions with no-arbitrage pricing conditions. In this paper, we follow this approach. This choice is determined by the finding that in the presence of risk, asset pricing must not ignore investors' risk preferences and be based solely on the assets' payoffs. Setting bounds based on assets' payoffs alone leads in many cases to inconsistencies about investors preferences: an investor shows both risk aversion and risk loving within the bounds set for the price.

Resorting to investors' preferences yields exact pricing relationships, which are an obvious improvement over the bounds obtained by the no-arbitrage approach. The problem is that, in many circumstances, it is very difficult to assert precision with which risk preferences characterize investors. Pricing based on incorrectly characterized utility functions, or the use of wrong parameter values creates problems and errors that make researchers reluctant to build approaches that use

investors' risk attitudes. The challenge is then to impose the fewest possible set of restrictions on investors' preferences, which are supported by the economy, and use them in combination with the assets payoffs in ways that achieve meaningful valuations, while avoiding arbitrage opportunities.

In our approach, bounds on asset prices depend on how much information there exists about investors' preferences. When an investor has little knowledge of its attitudes toward risk, the information used to price an option is that incorporated in the basis assets, and the option can have many values within the bounds set by the no-arbitrage approach. However, when an investor has some knowledge of its risk preferences, it rationally uses this information to narrow the bounds on the value of the option.

We show that by making the simple assumption that investors are risk-averse – a standard assumption in finance – it is possible to narrow the bounds on the values of the option considerably. Further tightening of the option bounds requires additional restrictions on investors' preferences. We propose a method to narrow the bounds that uses information from the initially estimated parameters, and consider the possibility of misspecification as well as of estimation errors to determine the neighborhood within which the correct value of the option will most likely lie. Since it is possible that the utility is misspecified or that the parameters characterizing attitudes towards risk are misestimated, the method relies on a minimax deviation to set a threshold in terms of the investor's confidence range. The range depends on maximum deviations: upper and lower maximum deviation. The lower maximum deviation measures the distance between the option value computed using the investor's benchmark risk-adjusted probabilities and the option value computed under the maximum level of risk aversion for the investor. The upper maximum deviation measures the distance between the option value computed using the objective probabilities (as if the investor were risk neutral) and the option value computed with the investor's benchmark risk-adjusted probabilities. The minimum of these two maximum deviations measures how much deviation the investor can confidently accept relative to its benchmark valuation. Option values outside this minimax deviation are then either too high or too low to be considered acceptable by the investor.

The article is organized as follows. First, we present the minimax deviation approach to narrow the bounds on the value of options in an incomplete market. Section 3 applies the methodology to tighten the bounds on the values of a call option on a stock. In Section 4 we compare our approach with other recent methods developed by Cochrane and Saa-Requejo (2000) and by Bernardo and Ledoit (2000). We show that our approach dominates these two others: first, it does not impose

additional restrictions to narrow the bounds beyond the information obtained from the initially estimated parameters, such as restrictions on maximum Sharpe ratios or maximum gain-loss ratios. Second, and more important, our approach avoids the inconsistencies in Cochrane and Saa-Requejo (2000) and in Bernardo and Ledoit (2000), that lead to pricing kernels representing investors displaying both risk-aversion and risk-loving attitudes. Third, our approach is more efficient in that it provides narrower bounds. In Section 5, we conclude.

2 A Minimax Deviation Approach

We present the minimax deviation approach to narrow the bounds on option prices in a finite-state economy with a risk-free asset. The financial market has a number of linearly independent securities smaller than the number of states in the economy. Investors are risk-averse and pricing kernels decline monotonically with a state factor. This implies that if a state of nature has higher/lower probability of lower (higher) payoff all other states also have higher/lower probability of lower (higher) payoff.

2.1 A Setup for Asset Pricing

Consider a two period economy where assets trade at a known price today and have random payoffs in the next period. There are S future states of the world, with $q_j > 0$ denoting the objective probability that state j occurs ($j = 1, \dots, S$). The economy offers at most S linearly independent assets generating the space $Z \subset \mathcal{R}^S$ of portfolio payoffs. The 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. Define three types of payoffs in the non-negative area: the null payoff $\tilde{0} = (0, \dots, 0)$, the positive orthant $\mathcal{R}_+^S = \{\tilde{z} \in \mathcal{R}^S : z_j \neq \tilde{0} \text{ and } z_j \geq 0 \forall j\}$, and the strictly positive orthant $\mathcal{R}_{++}^S = \{\tilde{z} \in \mathcal{R}^S : z_j > 0 \forall j\}$.

A linear functional π defined in Z spans asset prices in a way that the portfolio with payoff $\tilde{z} \in Z$ has price $\pi(\tilde{z})$. The pricing functional $\pi(\tilde{z})$ is strictly positive to remove arbitrage opportunities: $\pi(\tilde{z}) > 0 \forall \tilde{z} \in Z \cap \mathcal{R}_+^S$. With a risk-free asset offering the risk-free rate of return r_f , we present the space of excess payoffs X given π as $X = \{\tilde{z} - (1 + r_f)\pi(\tilde{z}) : \tilde{z} \in Z\}$.

2.2 Deviations from the Benchmark Price

An investor characterized by the vector (u, \tilde{c}^*) , where u is a continuously differentiable and concave von Neumann-Morgenstern utility function with marginal utility of consumption, $u'(\tilde{c}_j^*) > 0$, and

$\tilde{c}^* = (c_1^*, \dots, c_S^*) \in \mathcal{R}^S$ is an equilibrium state contingent consumption package. A pricing kernel m_j^* is given by

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

where $E[\cdot]$ denotes the expectation under the objective probabilities q_j . This pricing kernel represents how much the benchmark investor is willing to pay, per unit of probability, for a state contingent claim that pays 1 unit of consumption in state j and 0 in all other states. For a risk-averse investor $u''(\tilde{c}_j^*) < 0$. We assume that the pricing kernel declines monotonically with the state variable, that is, a state claim that pays off when \tilde{c}^* is low (high) has a relatively high (low) value, since such claim allows the benchmark investor to smooth consumption across future states of nature.

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

$$\forall \tilde{x} \in X \text{ subject to } \tilde{x} \neq \tilde{0}, E[\tilde{m}^* \tilde{x}] = 0 \Leftrightarrow E^*[\tilde{x}] = 0 \quad (2)$$

where $E^*[\cdot]$ denotes the expectation under the correct risk-adjusted probabilities $p_j^* = \frac{q_j u'(c_j^*)}{E[u'(\tilde{c}^*)]}$ for $j = 1, \dots, S$.

Since the correct risk-adjusted probabilities depend on the risk-aversion implied by the utility function u , any deviations resulting from misspecification of u or any misestimation of the parameters characterizing the investor's risk preferences will make the estimated risk-adjusted probabilities, \hat{p}_j , differ from the true yet unobserved risk-adjusted probabilities, p_j^* . Naturally, the risk-adjusted probabilities differ from the objective probabilities, q_j . Consider that \hat{p}_j is a point estimate in a distribution of possible values for p_j^* , characterized by an estimation error of $+/- e_j$, or e_j^+ and e_j^- . With e_j added to or subtracted from the estimated risk-adjusted probabilities, the correct risk-adjusted probability can be expressed as $p_j^* = \hat{p}_j \pm e_j$ with $0 \leq p_j^*, \hat{p}_j, e_j \leq q_j$. Then, equation (2) becomes

$$\forall \tilde{x} \in X \text{ subject to } \tilde{x} \neq \tilde{0}, E[\tilde{m}^* \tilde{x}] = 0 \Leftrightarrow E^*[\tilde{x}] = 0 \quad (3)$$

$$\Leftrightarrow \hat{E}[\tilde{x}] = 0 \Leftrightarrow e_j = 0 \quad (4)$$

where $\hat{E}[\cdot]$ denotes the expectation under the estimated risk-adjusted probabilities $\hat{p}_j = p_j^* \pm e_j$ for $j = 1, \dots, S$.

From the monotonicity of the pricing kernels and risk aversion, \hat{p}_j must be in the same range as p_j^* for all j : either both are in the interval $(0, q_j)$ or in the interval $(q_j, 1)$. For example, if we

have $p_j^* \in (0, q_j)$, we must have $\hat{p}_j \in (0, q_j)$ because the other range of values $(q_j, 1)$ would imply that the investor displayed risk-loving preferences for state j and then it would be a contradiction to contain \hat{p}_j . The deviation, $\pm e_j = p_j^* - \hat{p}_j$, of the estimated risk-adjusted probability from the correct risk-adjusted probability is mathematically defined in R^S (except for $\tilde{0}$) and its absolute value, $|e_j|$, (i) takes on values between zero and q_j for $0 \leq \hat{p}_j \leq q_j$, (ii) changes to values between 0 and $1 - q_j$ for $q_j \leq \hat{p}_j \leq 1$. This deviation translates into precision in the set of pricing kernels. If the deviation is zero, the pricing kernels are correct to the benchmark investor. If the deviation is positive (negative) for $0 \leq \hat{p}_j \leq q_j$, the investor will exhibit less (more) risk-aversion than implied by the estimated risk-adjusted probabilities, and the pricing kernels are under(over)-estimated. An investment with positive (negative) deviation for $0 \leq \hat{p}_j \leq q_j$ offers a good buying (selling) opportunity. For example, a deviation that raises the risk-adjusted probability to the level of the objective probability for $0 \leq \hat{p}_j \leq q_j$ implies that a risk-averse investor values the investment as if it were risk neutral, and pays more for the investment than the value implied by the correct pricing kernel.

$$0 \leq p_j^* \leq q_j \text{ and } 0 \leq |e_j| = |p_j^* - \hat{p}_j| \leq q_j \quad \text{if } 0 \leq \hat{p}_j \leq q_j \quad (5)$$

$$q_j \leq p_j^* \leq 1 \text{ and } 0 \leq |e_j| = |p_j^* - \hat{p}_j| \leq 1 - q_j \quad \text{if } q_j \leq \hat{p}_j \leq 1 \quad (6)$$

Large errors in estimated risk preferences generate pricing kernels significantly different from the correct pricing kernel. From the benchmark pricing kernel and using the estimated risk-adjusted probabilities \hat{p}_j ,

$$\hat{m}_j = \frac{\hat{p}_j}{q_j(1+r_f)} \quad (7)$$

it is possible to obtain the deviations in pricing kernels:

$$m_j^* = \frac{p_j^*}{q_j(1+r_f)} = \frac{\hat{p}_j \pm e_j}{q_j(1+r_f)} = \hat{m}_j + \frac{e_j^+ - e_j^-}{q_j(1+r_f)} \quad (8)$$

2.3 A Minimax Deviation

The total deviation $e_j = e_j^+ - e_j^-$ can be decomposed into a positive part $e_j^+ = \max(e_j, 0)$ and a negative part $e_j^- = \max(-e_j, 0)$. Focusing on the range $(0, q_j)$ for \hat{p}_j ¹, we use these positive and negative parts of the deviation to develop a method that refines the range of plausible pricing kernels. The method works as follows: on the one hand, the maximum deviation that is consistent with the lowest possible level of risk aversion is given by the value of $e_j^+ = q_j - \hat{p}_j$; on the other

¹The case of a range $(q_j, 1)$ for \hat{p}_j can be treated similarly.

hand, the maximum deviation that is consistent with the highest possible level of risk aversion is given by the value of $e_j^- = \hat{p}_j - 0 = \hat{p}_j$. Consider the minimum of these maximum deviations. This *minimax deviation*, defined by e_j^* , is the deviation

$$e_j^* = \min \left(e_j^+, e_j^- \right) = \begin{cases} \min(q_j - \hat{p}_j, \hat{p}_j) & \text{if } 0 \leq \hat{p}_j \leq q_j \\ \min(1 - \hat{p}_j, \hat{p}_j - q_j) & \text{if } q_j < \hat{p}_j \leq 1 \end{cases} \quad (9)$$

2.4 Minimax Bounds

This minimax deviation is then used to determine the deviation that is consistent with a change of the same magnitude in utility on the other side of the risk preference. Since the utility is strictly monotone in the risk-adjusted probability, it is possible to obtain the inverse of the utility as an explicit function of the risk-adjusted probabilities. To do this we first compute the change in utility level implied by the minimax deviation from the estimated risk-adjusted probability. By allowing the same change in utility but with a different sign on the other side of the risk preference, we then derive the maximum reasonable deviation on that side. The deviation derived in this way is now used to construct the other bound on the pricing kernels, e_j^{**} , which is computed by solving the following equations with the superscript of -1 denoting the inverse function:

$$u^{-1}(\hat{p}_j + e_j^*) - u^{-1}(\hat{p}_j) = u^{-1}(\hat{p}_j) - u^{-1}(\hat{p}_j - e_j^{**}) \quad \text{if } q_j - \hat{p}_j \leq \hat{p}_j \quad (10)$$

$$u^{-1}(\hat{p}_j) - u^{-1}(\hat{p}_j - e_j^*) = u^{-1}(\hat{p}_j - e_j^{**}) - u^{-1}(\hat{p}_j) \quad \text{if } q_j - \hat{p}_j > \hat{p}_j \quad (11)$$

For example, if we have that $q_j - \hat{p}_j < \hat{p}_j$, then $e_j^* = q_j - \hat{p}_j$. The important thing to note is that now the pricing kernels with minimax deviation have narrow bounds around the unobserved (true) benchmark pricing kernel

$$m_j^* \in \left(\frac{\hat{p}_j + e_j^*}{q_j(1+r_F)}, \frac{\hat{p}_j - e_j^{**}}{q_j(1+r_F)} \right) = \left(\frac{1}{1+r_F}, \frac{\hat{p}_j - e_j^{**}}{q_j(1+r_F)} \right) \quad (12)$$

The economic intuition for the minimax deviation is straightforward. An investor that exhibits a smaller maximum deviation on one side of the benchmark risk preference expresses more confidence on that side than on the other side of the risk preference, because the minimax deviation limits the maximum possible level of utility changes that correspond to the estimated risk-adjusted probability. Based on the uncertainty of deviation, it is reasonable to apply the same level of change in the utility to the other side of the estimated risk-adjusted probability. Therefore, the investor is able to use the minimum of the maximum deviations to set the deviation on the other side; the final result is narrow bounds around the true yet unobserved benchmark pricing kernel.

3 Option Price Bounds without Rebalancing

While the formula by Black and Scholes (1973) provides us an exact price for an European option, it is valid only in a complete market. When intermediate trading is not allowed, we cannot use the formula because the market becomes incomplete without dynamic rebalancing. We illustrate the minimax deviation approach to narrow the price bounds for an option in an incomplete market in an economy with two basis assets: a riskless bond and a risky stock. The bond offers a continuously compounded rate of return of $r = \log(1 + r_f)$, while the continuously compounded return on the stock is normally distributed with mean $\mu - (\sigma^2/2)$ and variance σ^2 . The stock return makes the stock price \tilde{S} lognormally distributed. We assume that it is not possible to dynamically replicate the option payoffs between two dates (date 0 and date t). We apply the approach to the case of an European call option with exercise price K and time to expiration t .

Minimax bounds can be constructed for different classes of utility functions. Here we employ a power utility function $u(\tilde{c}^*) = \tilde{S}^{(1-b)}/(1-b)$, which exhibits constant relative risk aversion (CRRA). With this utility function, we can obtain the Black-Scholes price in the absence of dynamic rebalancing as shown in Merton (1973), Rubinstein (1976), and Brennan (1979). We first compute the risk-neutral probability \hat{p} , which is the product of $u'(\tilde{c}^*)/E[u'(\tilde{c}^*)]$ by the objective probability q . With the risk neutral probability we derive the risk-adjusted probability, which is determined by the continuously compounded return of the stock being normally distributed with $r - (\sigma^2/2)$ and variance σ^2 . Denoting S by the initial stock price, C the initial call option price, and the call option payoff as $\tilde{C} = (\tilde{S} - K)^+$, we obtain the benchmark Black-Scholes price from the risk-neutral valuation $C = E^* \left[(\tilde{S} - K)^+ 1_{[K \leq \tilde{S} < \infty]} \right] / e^{rt}$.

3.1 The Price of a Call Option

In this section we choose parameter values to reflect the U.S. stock market index. The one-year continuously compounded rate of return on the risk-free bond is $r = \log(1 + r_f) = \log(1.05) = 4.88\%$. The one-year continuously compounded rate of return on the stock is $\mu = \log(1.13) = 12.22\%$ and its annual volatility is $\sigma = 14.09\%$.

In this economy with two basis assets, a risk-free bond with $r = 4.88\%$ and a stock with its return distributed normally with mean $\mu - \sigma^2/2 = 0.1123$ and variance $\sigma^2 = 0.019853$, we price a call option with strike price $K = 100$ and time to expiration $t = 1$, in the absence of dynamic rebalancing.

The benchmark Black-Scholes price consistent with a power utility function $u(\tilde{c}^*) = \tilde{S}^{(1-b)}/(1-b)$ is obtained by Monte Carlo simulation as in Bernardo and Ledoit (2000). Under the Black-Scholes risk-adjusted probability measure, the lognormal distribution of the final stock price is used to take on 10,000 random numbers, which are then used to generate independently and identically distributed standard normal random variables y_i . The terminal stock price S_i in the i th simulation with the initial stock price S is given as

$$S_i = S \exp \left[rt - (\sigma^2/2)t + \sigma\sqrt{t}y_i \right] = S \exp \left[0.0488 - (0.1409^2/2) + 0.1409y_i \right] \quad (13)$$

The call option has the terminal payoff $C_i = (S_i - K)^+ = (S_i - 100)^+$, which is employed to approximate the benchmark Black-Scholes call option price

$$C = \frac{1}{I} \sum_{i=1}^I (S_i - K)^+ e^{-rt} \quad (14)$$

where I is the number of runs performed in the simulation.

The power utility function implied by the Black-Scholes price assumes that investors are risk-averse. When investors are risk-neutral, they compute the unique price of the call option with the objective probability measure as the benchmark risk-adjusted probability measure. This price is obtained from the inner product of the terminal payoff of the call option with the pricing kernel, which takes the value \underline{m} when the excess return on the stock is positive and \overline{m} when the excess return is negative. We must solve a system of two equations with two unknowns to determine the values \underline{m} and \overline{m} , using the condition that the stock and the bond are accurately priced. The benchmark price and the risk-neutral price are shown in Figure 1.

3.2 Minimax Price Bounds

The minimax deviation approach presented above is able to construct bounds that are consistent with risk aversion, even without imposing any additional restrictions on the pricing kernel. To see this, consider the utility function consistent with the Black-Scholes call option price in an incomplete market. With a power utility function exhibiting constant relative risk-aversion (CRRA), Rubinstein (1976) shows that there is a relation between the parameters of the utility function and the parameters of the distribution of stock returns. Following Rubinstein, we can specify the parameter of a utility function consistent with the Black-Scholes option price in an economy with

values described in section 3.1, as

$$U(\tilde{c}^*) = \frac{1}{1-b}c^{1-b} \tag{15}$$

$$b = \frac{(\mu - \sigma^2/2) - r}{\sigma^2} + \frac{1}{2} = \frac{0.1123 - 0.0488}{0.019853} + \frac{1}{2} = 3.6986 \tag{16}$$

Given the current stock price, S_0 , the expected utility maximization (with the parameter b) that generates the Black-Scholes option value exhibits some degree of risk-aversion. Around the Black-Scholes option value, there is a risk-neutral price that can be determined which is consistent with risk-neutrality. In an economy that exhibits risk aversion from looking at the market value of the option, this risk neutral price is an extreme value (bound). The other extreme value (bound) is the no-arbitrage price that is consistent with infinite risk-aversion. For example, given the current stock price of 90, the Black-Scholes price, the risk-neutral price, and the no-arbitrage low bound consistent with infinite risk-aversion are 2.9078, 3.2386, and 0, respectively. Thus, the risk-averse bounds around the benchmark Black-Scholes value become (0, 3.2386). Since $3.2386 - 2.9078 = 0.3308 < 2.9078 (= 2.9078 - 0)$, the minimax upper bound is chosen as 3.2386. Computing the change in utility from the Black-Scholes benchmark to the upper bound given by the risk-neutral value, and then applying the same change in utility to the other side of the benchmark value, we obtain the minimax lower bound as 2.6751. On the other hand, in cases where the Black-Scholes option value is greater than the risk-neutral option value, the risk-neutral option value becomes the minimax lower bound and the minimax upper bound is computed similarly. That is, we compute the change in utility from the benchmark Black-Scholes price to the minimax lower bound (equal to the risk-neutral price), and then apply the same change in utility to determine the upper bound.

Figure 2 shows the minimax bounds constructed around the benchmark Black-Scholes option market price. As observed in the graph, the minimax bounds generate values consistent with a risk-averse economy, preferences which are reflected in the Black-Scholes price. When the benchmark price is lower (higher) than the risk neutral price, any risk-adjusted prices from risk aversion must also be lower (higher) than the risk neutral price. Therefore, the minimax bounds lie below (above) the risk neutral bounds. The minimax bounds are able to narrow the bounds around the Black-Scholes price, while not requiring arbitrary restrictions such as maximum values on the Sharpe ratio or on the gain-loss ratio.

4 Comparison with Other Approaches

We use a numerical example to illustrate how our approach compares with two well known methods for the case of a European call option when there is no dynamic rebalancing. The two approaches are those the “good-deal bounds” presented in Cochrane and Saa-Requejo (2000) and “the gain-loss ration” proposed in Bernardo and Ledoit (2000).

4.1 Good-Deal Bounds

In calibrating the discount factor volatility constraint, we use the assumption in Cochrane and Saa-Requejo (2000) that the investor takes any opportunity (sees the opportunity as a good deal) with a Sharpe ratio twice that of a U.S. stock market index. Using the values in section 3.1:

$$E(R) = \exp(\mu) = 1.13 : R_f = \exp(r) = 1.05 \quad (17)$$

$$\sigma(R) = \exp(\mu) \times (\exp(\sigma^2) - 1)^{1/2} = 1.13 \times (\exp(0.1409^2) - 1)^{1/2} = 0.16 \quad (18)$$

$$h = 2 \times \frac{E(R - R_f)}{\sigma(R)} = 2 \times \frac{1.13 - 1.05}{0.16} = 1 \quad (19)$$

Since the good-deal bounds approach by Cochrane and Saa-Requejo rests on imposing a constraint on the volatility of the pricing kernels, we compute pricing kernels corresponding to the upper bound and the lower bound for the call option. The pricing kernel of the lower bound and its volatility, $\sigma(m)$, are:

$$\underline{C} = \min_{\{m\}} E(mx^c) \quad \text{subject to } \mathbf{p} = E(m\mathbf{x}); \quad m \geq 0; \quad \sigma(m) \leq \frac{h}{R_f} = \frac{1}{1.05} = 0.9524 \quad (20)$$

and the upper good-deal bound \bar{C} solves the corresponding maximum.

Figure 3 shows the resulting good-deal bounds by Cochrane and Saa-Requejo around the Black-Scholes price, and the our minimax price bounds. The minimax bounds are much tighter than the good-deal bounds. Recall that in both instances the investor takes any opportunity with Sharpe ratio twice that of a U.S. stock market index. Why are the minimax bounds much tighter then? Because the “good-deal bounds” approach of Cochrane and Saa-Requejo does not use as a reference the benchmark price and therefore generate much more loose bounds with just the threshold of twice the Sharpe ratio. The problem with the Cochrane and Saa-Requejo approach is that their bounds conflict with the risk preferences of any risk-averse investor in the economy, because the Sharpe ratio constraint builds bounds around a risk-neutral price, which is in most cases not the price prevailing in the economy. For an initial stock price of 90, the good-deal bounds

by Cochrane and Saa-Requejo for the call option are (0.2537, 4.6571). The risk-neutral price at 3.2386 is higher than the Black-Scholes price equal to 2.9078. Thus, the interval comprised by the values (0.2537, 3.2386) represents risk-averse preferences, and the interval comprised by the values (3.2386, 4.6571) exhibits risk-loving preferences. On the other hand, when the initial stock price is 100, the good-deal bounds are (5.3428, 9.3104) and the risk-neutral call option price price is 7.7014, which is lower than the Black-Scholes price of 8.0642. In this case, the lower range (5.3428, 7.7014) of the risk-neutral price represents risk-loving preferences, and the upper range of (7.7014, 9.3104) exhibits risk-averse preference. Both cases show that the good-deal bounds (TBR: of Cochrane-Saa Requejo) by Cochrane and Saa-Requejo lead to the impossible situation of an investor that sets those bounds behaving simultaneously as a risk averse and a risk loving investor.

4.2 Gain-Loss Ratio Bounds

Bernardo and Ledoit (2000) choose the same parameter values as do Cochrane and Saa-Requejo except $\mu = 0.1123$. However, it is understood that they mean $\mu = \log(1.13) = 0.1222$ to have $\mu - \sigma^2/2 = \log(1.13) - 0.1409^2/2 = 0.1123$. Thus, we choose $\mu = 0.1222$ for consistency. Although Bernardo and Ledoit (2000) present the method of constructing the narrow bounds in detail, the procedure to implement the numerical computation of the gain-loss bounds is not trivial. Longarela (2002) develops a simple procedure for computation of the bounds, which requires solving a linear optimization program. We employ the Longarela approach to construct the gain-loss bounds by solving

$$\begin{aligned} \min_{\mu_1, \dots, \mu_S, \theta_1, \theta_2} \sum_{j=1}^S \mu_j z_j^* & \quad (21) \\ \text{s.t.} \quad \left\{ \begin{array}{l} \sum_{j=1}^S \mu_j z_j^i = \pi(\tilde{z}^i), \quad i = 1, \dots, N \\ \theta_1 \leq \frac{\mu_j}{\mu_j^*} \leq \theta_2, \quad j = 1, \dots, S \\ \theta_2 = \theta_1 \bar{L} \\ \theta_1 \geq 0 \end{array} \right. \end{aligned}$$

for the lower bound. The upper bound is found by solving the corresponding maximization problem. A vector of state prices, denoted by $[\mu_1, \dots, \mu_S]$, and asset prices are given by a linear function π on Z , the space of portfolio payoffs.

Figure 4 illustrates the gain-loss bounds resulting from a gain-loss ratio of 2 with the Black-Scholes call value used as benchmark. The minimax bounds improve significantly on the result

of applying both the good-deal bounds and the gain-loss bounds. As in the previous case of the Cochrane-Saa Requejo good-deal bounds, the gain-loss ratio approach also generates bounds that are inherently inconsistent with the risk preferences of the decision maker. This is because the gain-loss ratio constraint ignores any information regarding investors' preferences embedded in the benchmark price, and sets the bounds exclusively from the asset payoffs. Setting bounds based on assets' payoffs alone is very likely to include values that reveal inconsistent attitudes toward risk from investors that trade the assets in the incomplete market. Indeed, inside the bounds set by a ratio $\bar{L} = 2$, it is possible to find a region of values that conform with investors displaying risk aversion, while in a neighboring region also inside the bounds set by the same ratio $\bar{L} = 2$, the values are consistent with investors displaying a risk-loving preferences. Such anomalies are less desirable, and more so because in an incomplete market the gain-loss ratio is an individual-specific measure and therefore impossible to separate from the decision-maker. The same investor cannot be simultaneously risk-averse and risk-loving with respect to the same project.

Figure 5 combines the three alternative approaches – the good-deals, the gain-loss ratio and the minimax bounds – to illustrate the relative merits of each. In addition to providing consistency in investors' risk-preferences, the minimax bounds are the most efficient bounds, tighter than those constructed from the other two approaches.

4.3 More Efficient Bounds

A natural question that emerges from the previous comparison is why the minimax approach provides narrower bounds than the other two approaches? Are the narrow bounds that result from applying the minimax deviation just a coincidence of the particular example? Or can one safely consider that such bounds when applied to other typical pricing problems in an incomplete market provide better results? Since the merit of the minimax bounds depends on the coefficient of risk aversion that determines risk-adjusted probabilities relative to objective probabilities, these questions are directly related to another interesting question, that of whether risk-adjusted probabilities are sufficiently close to objective probabilities.

Since the gain-loss ratio bounds appear to perform better than the good-deal bounds, we focus our discussion about efficiency to the case of the maximum gain-loss ratio. Bernardo and Ledoit (2000) approximate the ratio to represent the U.S. stock market index. Assuming logarithmic utility, Bernardo and Ledoit (2000) compute the gain-loss ratio of the stock market index as $\bar{L} = 2.6$, which represents the extent of violation of the benchmark model in a classic equity premium puzzle.

They then choose the gain-loss ratio $\bar{L} = 2$, which excludes all ratios above that level, including $\bar{L} = 2.6$. Note, however, that typical investment opportunities in an incomplete market can easily dominate the gain-loss ratio of the U.S. stock market index, $\bar{L} = 2.6$. The exclusion of values above $\bar{L} = 2$ allows Bernardo and Ledoit to narrow the bounds considerably. Yet, as we show next, even for an economically very low gain-loss ratio, the bounds still are significantly wider than the minimax bounds derived in our paper.

To evaluate the two approaches, we modify them to provide equivalent results. Consider either (1) narrowing the gain-loss bounds with the smaller maximum gain-loss ratio of $\bar{L} = 1.5$, and (2) widening the minimax bounds with twice the risk-averse parameter $b' = 2b = 2 \times 3.6986 = 7.3972$. Modifying the gain-loss bounds is straightforward. We recalculate the bounds with $\bar{L} = 1.5$ in Figure 6. To modify the minimax bounds with new risk-averse parameter, we need to recalculate the maximized utility. Alternatively, we can also reduce in half the variance of the stock market return $\sigma^2/2 = 0.019853/2 = 0.0099265$ in Figure 7. Only then, the two sets of bounds (the gain-loss ratio bound and the minimax bounds) are equivalent.

A problem with the first modification is that a gain-loss ratio of $\bar{L} = 1.5$ is too restrictive to be justified by the stock market data, when the U. S. stock market index has a gain-loss ratio of $\bar{L} = 2.6$. Bernardo and Ledoit also compare the gain-loss ratio with the Sharpe ratio under the assumption of normal distribution. They find that a gain-loss ratio of $\bar{L} = 2$ is equivalent to a Sharpe ratio of 0.3. This connection implies that the gain-loss ratio $\bar{L} = 1.5$ must be equivalent to a Sharpe ratio lower than 0.3. Using a Sharpe ratio lower than 0.3 appears to show that the investor has no tool reflecting real investment opportunities on hand other than just constructing imaginary narrow bounds on a benchmark price. Thus a gain-loss ratio $\bar{L} = 1.5$ simply removes room for misspecification without plausible economic justification.

A lesson from the second modification is that doubling the risk-averse parameter does not give up the efficiency of the minimax bounds to the gain-loss bounds with $\bar{L} = 2$, while it significantly relaxes its assumption on the stock market return. For the minimax bounds to lose comparative advantage to the gain-loss bounds, investors must be much more risk-averse than twice that implied by the stock market returns. Unless we allow risk-adjusted probabilities to be unrealistically different from risk-neutral probabilities, the comparative advantage of the minimax bounds over the gain-loss bounds does not change.

4.4 Properties of the Minimax Bounds

The minimax bounds can be characterized as simple to construct, consistent with the implications of benchmark prices, and efficient in deriving tight bounds on asset prices.

Simplicity Asset price bounds in the literature require investors to set arbitrary ratios to derive somewhat narrow bounds on a asset price in an incomplete market. Those ratios are not easy to choose and often force investors to compromise between the benchmark price and the no-arbitrage price. However, the minimax bounds are simple because they are constructed without arbitrary ratios. As shown in Equation (12), $m_j^* \in \left(\frac{1}{1+r_F}, \frac{\hat{p}_j - e_j^{**}}{q_j(1+r_F)} \right)$ depends only on the utility difference implied by the difference between an objective probability and a risk-adjusted probability in addition to the risk-free rate of return.

Consistency Although asset price bounds obtained with any approaches must be consistent with their benchmark prices, they easily violate the implications inherent in the benchmark prices. When the benchmark price is derived from a concave utility function, bounds around the price are supposed to exhibit risk-aversion. The minimax bounds are guaranteed to be consistent with risk-aversion implied in the benchmark price by construction as Equation (10) and (11)

Efficiency For any asset price bounds to be of practical use, the pricing implication from the bounds must be sufficiently narrow. Although price bounds in the literature have been somewhat successful to obtain narrow bounds, the minimax bounds are more efficient than other bounds in deriving narrow bounds. Unlike the simplicity and the consistency of the minimax bounds, the efficiency of the minimax bounds must be determined by comparing them to the bounds suggested by other approaches in the literature. The results above show conclusively that the minimax bounds are more efficient.

5 Conclusion

In an incomplete market, it is not possible to replicate the values of certain assets using other assets. In many instances, such relative pricing formulas provide pricing implications that are too wide to be economically meaningful. To obtain more precision, it is necessary to specify investors' risk preferences and rely on arguments about the marginal rate of substitution for consumption across time and states. The disadvantage of a preference-based approach is the restriction imposed

on the agents' attitudes toward risk, which makes it susceptible to misspecification. The challenge is to achieve a good balance between the level of precision in prices and the restrictions imposed on the economy.

We develop an approach to preclude arbitrage opportunities and to improve pricing implications in an incomplete market. We use weak restrictions on the investor's attitude toward risk to restrict the range of possible pricing kernels that can be used to price risky assets. In its simplest form, the approach narrows the arbitrage bounds by using the information from the benchmark price (in the case of a call option, the Black Scholes market price) that investors are risk-averse. Further tightening is possible, since the no-arbitrage bounds imply extreme risk preferences, which are implausible in many instances. To obtain narrow bounds under the assumption of risk aversion, we develop a minimax deviation approach. We revise one of the bounds by relating to each other the deviations around the price computed by a benchmark investor, which might be misspecified. A minimax deviation represents the minimum deviation from the value computed with the estimated risk-adjusted probabilities to the value computed with the objective probabilities, or to the value with the extreme level of risk-aversion. The narrow bounds constructed by the minimum deviation are consistent with no-arbitrage pricing and with the risk preferences embedded in the observed benchmark pricing kernels.

Call options on stocks, perhaps more than other risky assets, have features that test the preference-free relative pricing approach when dynamic rebalancing is not feasible. Often, the event on which the option payoff depends is not a traded asset. Thin trading of real assets may also make it impossible to hedge. The key question is then how to deal with risk factors that cannot be hedged but can significantly impact the value of the option. While current approaches in the literature, notably those that have attracted more attention, by Cochrane and Saa-Requejo and by Bernardo and Ledoit, arbitrarily and exogenously impose restrictions on the values to obtain narrower bounds, our approach extracts additional information from the investor's confidence on the initial parameter estimation to compute narrower bounds in an incomplete market.

References

- Bernardo, Antonio E. and Olivier Ledoit, 2000, Gain, Loss, and Asset Pricing, *Journal of Political Economy* 108, 144–172.
- Black, Fischer and Myron S. Scholes, 1973, The Pricing of Options and Corporate Liabilities, *Journal of Political Economy* 81, 637–654.
- Brennan, Michael J., 1979, The Pricing of Contingent Claims in Discrete Time Models, *Journal of Finance* 34, 53–68.
- Choulli, T. and C. Stricker, 2005, Minimal Entropy-Hellinger Martingale Measure in Incomplete Markets, *Mathematical Finance* 15, 465–490.
- Cochrane, John H. and Jesus Saa-Requejo, 2000, Beyond Arbitrage: Good-Deal Asset Price Bounds in Incomplete Markets, *Journal of Political Economy* 108, 79–119.
- Follmer, H. and M. Schweizer, 1991, Hedging of Contingent Claims under Incomplete Information, in M. H. A. Davis and Eds. New York R. J. Elliott, eds.: *Applied Stochastics Monographs* (Gordon and Breach), 389–414.
- Longarela, Inaki Rodriguez, 2002, A Simple Linear Programming Approach to Gain, Loss and Asset Pricing, *Topics in Theoretical Economics* 2, 1–8.
- Merton, Robert C., 1973, Theory of Rational Option Pricing, *Bell Journal of Economics and Management Science* 4, 341–360.
- Rubinstein, Mark, 1976, The Valuation of Uncertain Income Streams and the Pricing of Options, *Bell Journal of Economics* 7, 407–425.
- Staum, J., 2004, Fundamental Theorems of Asset Pricing for Good Deal Bounds, *Mathematical Finance* 14, 141–161.

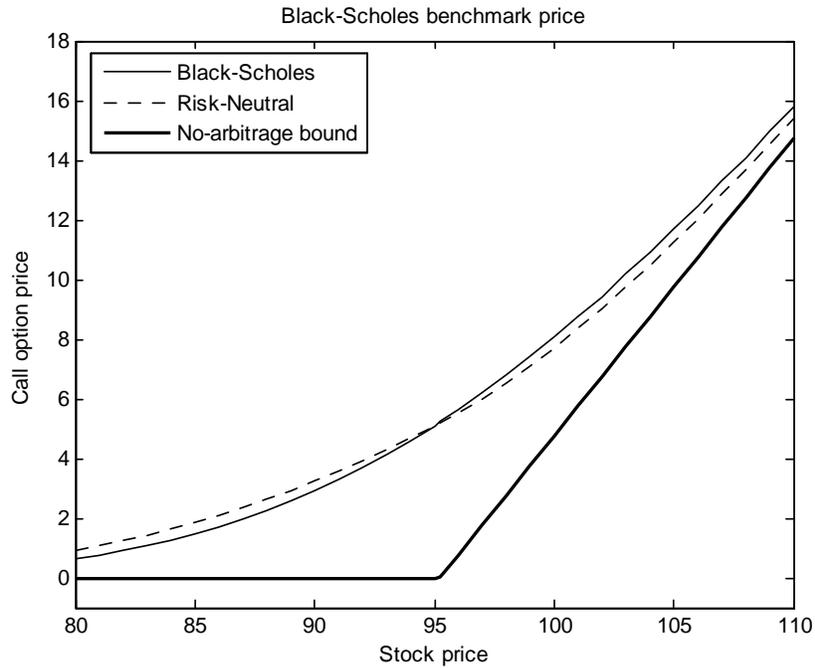


Figure 1: The Black-Scholes benchmark price and the risk-neutral price. The call options have one year to expiration and strike price $K = \$100$. The stock is lognormally distributed with parameters $\exp(\mu) = 1.13$ and $\exp(\mu) \times (\exp(\sigma^2) - 1)^{1/2} = 0.16$. The parameter values are chosen to imitate a stock market index. The benchmark price is represented by the Black-Scholes price. The risk-neutral price is computed with the objective probability. The no-arbitrage bound represents the lower bound derived from excluding pure arbitrage. The corresponding upper bound is too high to be drawn in the graph.

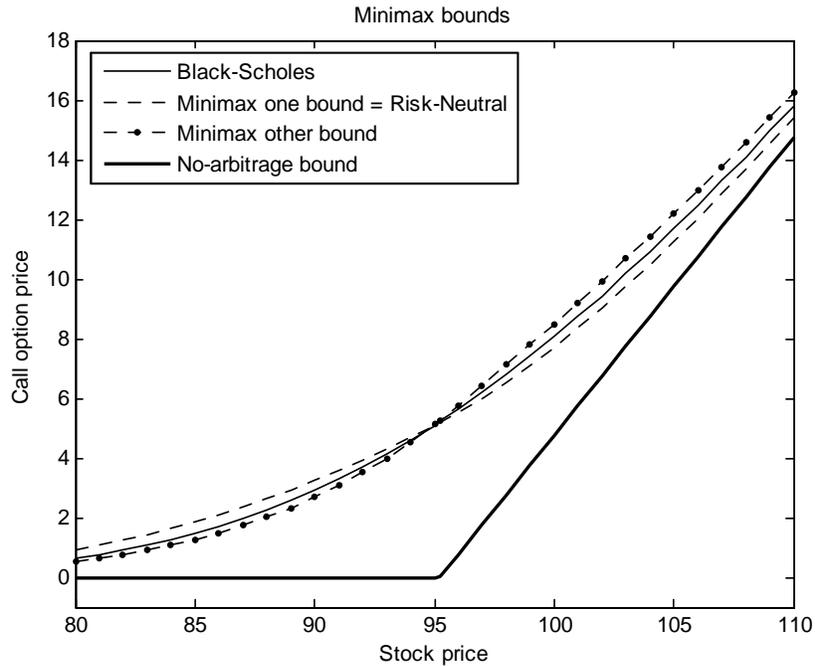


Figure 2: Minimax bounds around the Black-Scholes benchmark price. The call options have one year to expiration and strike price $K = \$100$. The stock is lognormally distributed with parameters $\exp(\mu) = 1.13$ and $\exp(\mu) \times (\exp(\sigma^2) - 1)^{1/2} = 0.16$. The parameter values are chosen to imitate a stock market index. The benchmark price is represented by the Black-Scholes price. The minimax bounds are constructed from employing a power utility function with a parameter $b = 3.6986$. The risk-neutral price, which is computed with the objective probability, serves one bound and the other bound is computed with the same change in utility from the benchmark and the risk-neutral price. The no-arbitrage bound represents the lower bound derived from excluding pure arbitrage. The corresponding upper bound is too high to be drawn in the graph.

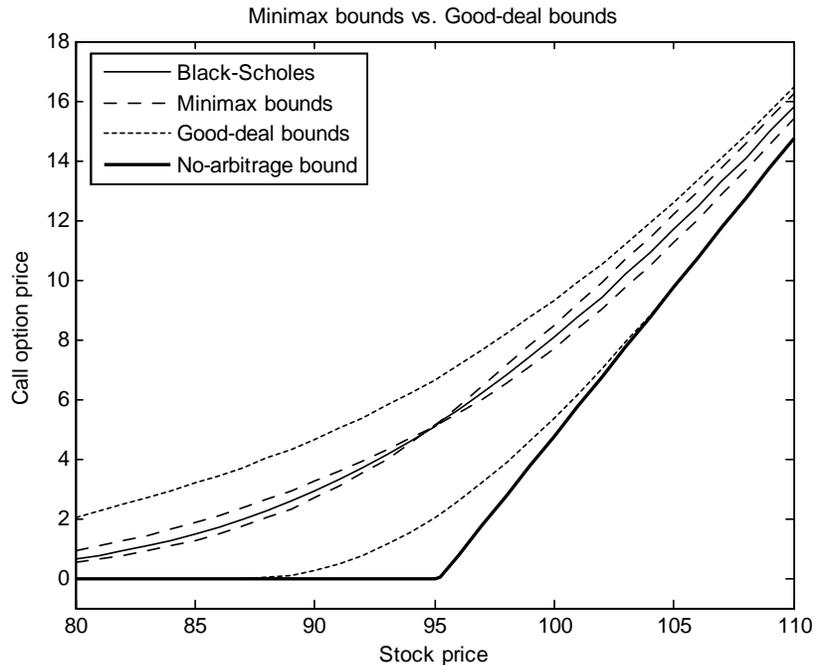


Figure 3: The minimax bounds with the good-deal bounds around the call option values. The call options have one year to expiration and strike price $K = \$100$. The bounds assume no intermediate trading and a discount factor volatility bound $h = 1$ associated with twice the market Sharpe ratio. The stock is lognormally distributed with parameters $\exp(\mu) = 1.13$ and $\exp(\mu) \times (\exp(\sigma^2) - 1)^{1/2} = 0.16$. The parameter values are chosen to imitate a stock market index. The Black-Scholes price represents the benchmark price for the minimax bounds and the reference for the good-deal bounds. The no-arbitrage bound represents the lower bound derived from excluding pure arbitrage. The corresponding upper bound is too high to be drawn in the graph.

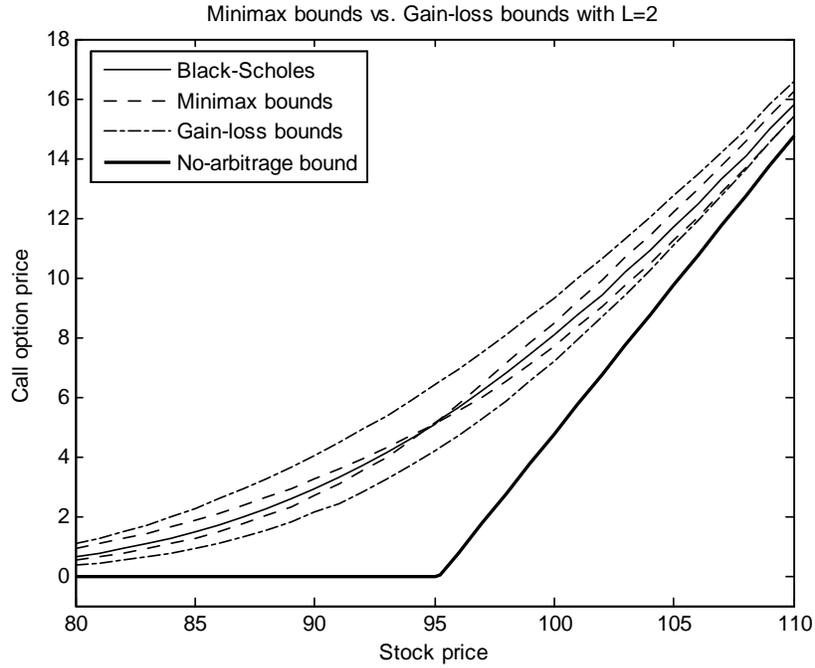


Figure 4: The minimax bounds with the gain-loss bounds around the call option values. The call options have one year to expiration and strike price $K = \$100$. The bounds assume no intermediate trading and the gain-loss ratio $L = 2$. The stock is lognormally distributed with parameters $\mu - \sigma^2/2 = 0.1123$ and $\sigma = 0.1409$. The parameter values are chosen to imitate a stock market index. The benchmark price is represented by the Black-Scholes price. The no-arbitrage bound represents the lower bound derived from excluding pure arbitrage. The corresponding upper bound is too high to be drawn in the graph.

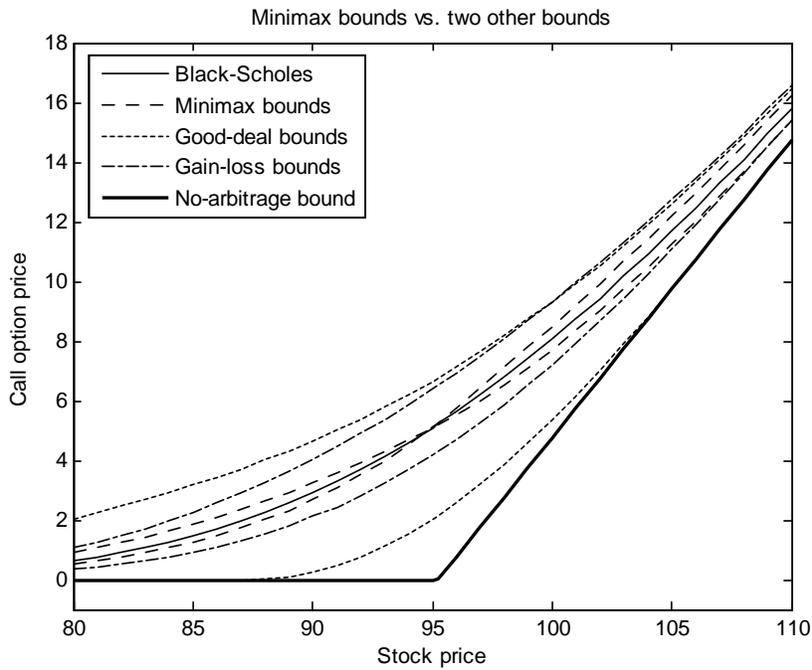


Figure 5: Alternative bounds around the call option values from the three approaches. The call option have one year to expiration and strike price $K = \$100$. The stock is lognormally distributed with parameters $\exp(\mu) = 1.13$ and $\exp(\mu) \times (\exp(\sigma^2) - 1)^{1/2} = 0.16$. The parameter values are chosen to imitate a stock market index. The Black-Scholes price represents the benchmark price for the minimax bounds and the gain-loss bounds, and the reference for the good-deal bounds. The minimax bounds are constructed from employing a power utility function with a parameter $b = 3.6986$. The risk-neutral price, which is computed with the objective probability, serves one bound and the other bound is computed with the same change in utility from the benchmark and the risk-neutral price. The good-deal bounds have a Sharpe ratio 2 and the gain-loss bounds have a gain-loss ratio 2. The no-arbitrage bound represents the lower bound derived from excluding pure arbitrage. The corresponding upper bound is too high to be drawn in the graph.

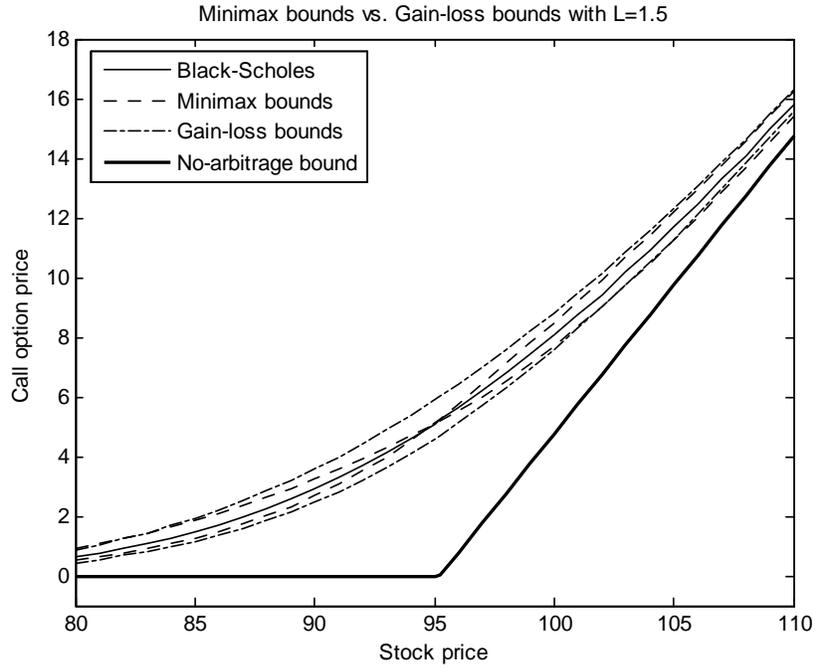


Figure 6: The minimax bounds with the gain-loss bounds around the call option values. The call options have one year to expiration and strike price $K = \$100$. The bounds assume no intermediate trading and the gain-loss ratio $L = 1.5$. The stock is lognormally distributed with parameters $\mu - \sigma^2/2 = 0.1123$ and $\sigma = 0.1409$. The parameter values are chosen to imitate a stock market index. The benchmark price is represented by the Black-Scholes price. The no-arbitrage bound represents the lower bound derived from excluding pure arbitrage. The corresponding upper bound is too high to be drawn in the graph.

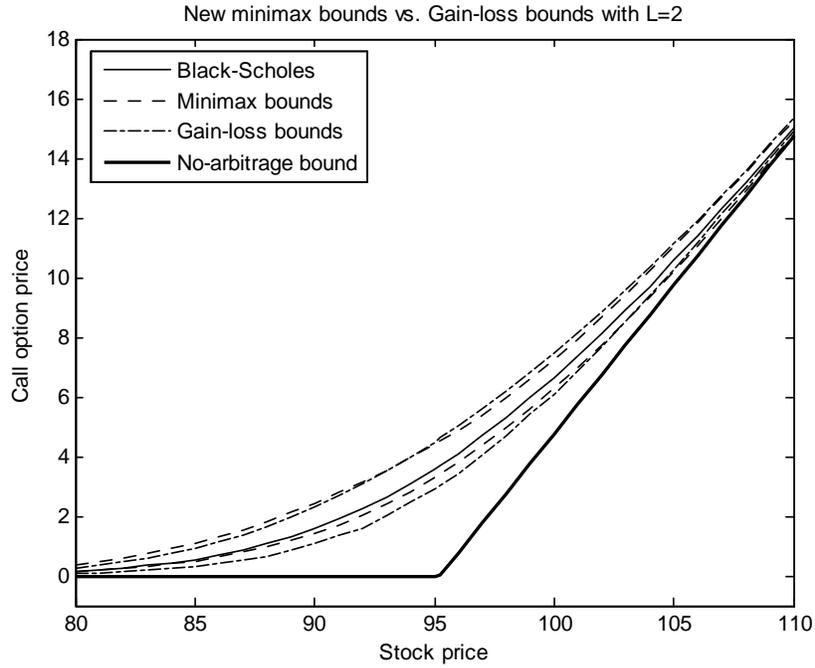


Figure 7: The minimax bounds with the gain-loss bounds around the call option values. The call options have one year to expiration and strike price $K = \$100$. The bounds assume no intermediate trading and the gain-loss ratio $L = 2$. The stock is lognormally distributed with parameters $\mu - \sigma^2/2 = 0.1172$ and $\sigma = 0.0996$. The parameter values are chosen to imitate a stock market index. The benchmark price is represented by the Black-Scholes price. The no-arbitrage bound represents the lower bound derived from excluding pure arbitrage. The corresponding upper bound is too high to be drawn in the graph.