Benchmark Good-Deal Bounds: an Application to Stochastic Volatility Models of Option Pricing*

Oleg Bondarenko† and Iñaki R. Longarela‡

April 23, 2004

Abstract

We present a generalization of Cochrane and Saá-Requejo's good-deal bounds which allows to include in a flexible way the implications of a given stochastic discount factor model. Furthermore, a useful application to stochastic volatility models of option pricing is provided where closed-form solutions for the bounds are obtained. A calibration exercise illustrates their good properties and how our *benchmark* good-deal pricing results in much tighter bounds. Finally, a discussion of methodological and economic issues is also provided.

1 Introduction

One of the most recurrent problems that the asset pricing literature has faced over the years may be presented in terms of the following question: given an initial set of asset payoffs with known price, what is the correct valuation to be assigned to a new postulated payoff? The former set is usually called the set of *basis payoffs* while the latter payoff is commonly referred to as the *focus payoff*. The above question has a clear-cut answer under the assumptions of market completeness and absence of arbitrage. In that case any new asset is redundant and its price is obtained by finding a portfolio of the basis assets whose payoff is identical to the focus payoff.

When the assumption of market completeness is dropped, a great deal of complexity is added to the problem. Several approaches can be followed

*Iñaki R. Longarela wants to thank the Wallander Foundation for financial support.

†Department of Finance (m/c 168), University of Illinois at Chicago, 601 S. Morgan St., Chicago, IL 60607-7124. E-mail: olegb@uic.edu.

‡Department of Finance, Stockholm School of Economics, Sveavägen 65, SE-113 83 Stockholm, Sweden. E-mail: Inaki.Rodriguez@hhs.se.

in this case. As a first approximation, the no-arbitrage assumption may be fully exploited to derive bounds which define the interval where the price of the focus asset should lie. However, in many realistic situations the size of this interval usually renders the *no-arbitrage* bounds uninformative.

With these considerations in mind, Cochrane and Saá-Requejo (2000) and Bernardo and Ledoit (2000) propose methods to shorten the size of the no-arbitrage interval by inserting stronger economic assumptions. Both contributions rely on a key concept of modern economic theory: the stochastic discount factor (SDF).

Cochrane and Saá-Requejo concentrate on the Sharpe ratio as the relevant measure of attractiveness of investment opportunities. Since it is possible to establish an equivalence between the maximum available Sharpe ratio and a restriction on the volatility of the admissible stochastic discount factors, this restriction delivers in a straightforward manner their *good-deal* bounds.

Bernardo and Ledoit consider the *gain-loss* ratio of excess payoffs, that is, the expectation of their positive part divided by the expectation of their negative part. For every admissible SDF, a duality result links the ratio with deviations from a given asset pricing model. Then by postulating a benchmark SDF and restricting the value of the above deviations one can again shorten the size of the no-arbitrage interval thereby deriving *gain-loss ratio* bounds.

An important difference between these two approaches is the key role that the benchmark model plays in Bernardo and Ledoit's bounds and that is absent in deriving good-deals bounds. This benchmark represents a reasonable initial guess of an econometrician of the true SDF or/and it is derived from a partial or general equilibrium model. In any case, there is a legitimate concern about potential misspecifications and the benchmark helps in locating more accurately the set of candidate SDF's, presumably resulting in tighter bounds.

Interestingly enough, a generalization of good-deal bounds along these lines can also be implemented and in fact, it is already suggested as a possibility in Cochrane and Saá-Requejo (2000). The key idea to perform this extension relies on an alternative interpretation of these objects. Suppose that one defines the level of attractiveness of an investment opportunity in terms of deviations from a misspecified benchmark SDF and quantifies those deviations with Hansen and Jagannathan's (HJ) measure of model misspecification. This idea delivers in a straightforward manner a general type of asset price bounds which happen to contain as a particular case good-deal bounds. We analyze the economic and computational aspects of these constructs in a two-period model.

However, in its pure form the HJ measure does not allow for a tractable extension to continuous time. For this reason, we propose another closely related measure. The new measure has the advantage that it permits an extension to continuous time and still nests good-deal bounds as a particular case. Furthermore, in the context of Heston's (1993) model we provide an application to option pricing with stochastic volatility where we obtain a closed-form solution for the bounds. A calibration exercise allows us to conclude the good performance of our *benchmark* good-deal (BGD) bounds as opposed to the standard good-deal (GD) bounds.

The remainder of the paper is organized as follows. Some methodological issues are discussed in the next section. Section 3 deals with the economic properties and computational aspects of the bounds in a static setup. A continuous time exposition is the content of Section 4 where we also present our application to stochastic volatility option pricing. Finally, some suggestions for further research are given in the last section. All necessary mathematical proofs are confined to an appendix.

2 Methodological preliminaries

Consider a two-period economy where N basis assets are traded today at a known price given by a vector $\mathbf{p}$ and they deliver a random payoff denoted by a vector $\mathbf{x}$. No frictions are allowed and hence any linear combination of the basis payoffs constitutes an attainable payoff x whose set will be denoted by X^b. Let $X \supset X^b$ be the space of payoffs in the span of all contingent claims. For any payoff $x \in X$ consider its positive and negative part $x^+ \equiv \max(x, 0)$ and $x^- \equiv \max(-x, 0)$, respectively. To simplify exposition, we focus on the finite-dimensional case with S states of nature which implies that $X = \Re^S$.

We assume from the beginning absence of arbitrage opportunities. This implies the existence of a strictly positive random variable, also known as SDF, which satisfies

$$\mathbf{p} = E(m\mathbf{x}). \tag{2.1}$$

Denote by M^+ the set of all strictly positive SDF's.[1] These random variables give the price of any payoff $x \in X$ through their *pricing extensions* defined as

$$\pi_m(x) \equiv E(mx)$$

[1] Ocassionally, we will use the term SDF even if condition (2.1) is not satisfied. Thus, whenever the context is not clear we will refer to those random variables belonging to M as *admissible* SDF's.

which assign the same price to all payoffs x in X^b but in general differ for payoffs $x \notin X^b$. Define also the set

$$Z_m = \{x \in X : \pi_m(x) = 0,\ x \neq \mathbf{0}\}$$

for any $m \in M^+$. Thus, Z_m is the set of excess returns under the pricing function $\pi_m(\bullet)$.

We also assume that one of the basis assets is risk-free with rate of return R^f and hence, it holds that $E(m) = 1/R^f$ for any $m \in M^+$.

The pricing implications of a given model are summarized by its implied SDF. Denote by m^* a strictly positive *benchmark* SDF which will always be assumed to satisfy $E(m^*) = 1/R^f$.[2] For any $x \in X$ consider also its implied pricing function $\pi^*(x) \equiv E(m^* x)$. Let also $E^*(\cdot)$ denote the expectation under the risk-neutral probability measure implied by m^* and thus, we have that

$$\pi^*(x) = E^*(x)\, E(m^*).$$

In deriving asset price bounds the goal is always to tighten the set M^+ in accordance to an economically meaningful criteria. Even though all elements of M^+ are consistent with the absence of arbitrage, many of them can be discarded by imposing stronger economic assumptions. At the same time, flexibility is also an important ingredient, since one wants to avoid falling into the rigidities of model-based pricing. All the above is accomplished by choosing a reference model and establishing how much discrepancy, measured in terms of a given distance,[3] one is willing to allow. Unconstrained discrepancy will lead to no-arbitrage bounds while zero discrepancy will imply model-based pricing. It is in between these two extremes where our interest lies.

Formally, for any $m \in M^+$ denote by $d(m, m^*)$ a distance between m and the benchmark model m^*. Let $\overline{d}$ be an appropriate ceiling on the maximum value of the above distance and discard those $m \in M^+$ such that $d(m, m^*) > \overline{d}$. The lower bound on the price of a payoff $x^c \notin X^b$ will be given by

$$\underline{C} = \min_{\substack{m \in M^+ \\ d(m,m^*) \leq \overline{d}}} E(m x^c) \tag{2.2}$$

[2] In principle, m^* can be assumed to be more general and in particular, it does not have to price the risk-free asset. In many cases, it could even take negative values. However, we will not need such generality in our exposition.

[3] We abuse the term "distance" in our framework since the symmetry property does not necessarily hold.

The upper bound $\overline{C}$ follows from replacing min with max in the above optimization. Clearly, the value of the ceiling must satisfy

$$\min_{m \in M^+} d(m, m^*) \leq \overline{d} \leq \max_{m \in M^+} d(m, m^*)$$

because if the left-hand inequality does not hold the feasible set in (2.2) is empty and if the right-hand inequality is violated the original no-arbitrage bounds are obtained.

It should be noted that all economic content can be introduced by means of d and/or m^* which are in the end the two key ingredients in the above derivation. The distance usually has a distinct economic interpretation which is obtained through a duality result. The benchmark m^* represents a reasonable initial guess of an econometrician about the true model. Since potential misspecification is always a concern, the benchmark helps when characterizing the set of candidate SDF's, resulting in tighter bounds. When m^* plays no explicit role, the distance is determined by a given moment of the admissible SDF's.

The above can be illustrated with the two seminal contributions in the asset price bounds literature. Bernardo and Ledoit (2000) present a derivation based on the gain-loss ratio. These bounds are obtained by setting

$$d(m, m^*) = \frac{\sup \frac{m}{m^*}}{\inf \frac{m}{m^*}} \tag{2.3}$$

and its interpretation is given by the duality result below.

Proposition 2.1 *For the distance in (2.3), the following equality holds*

$$d(m, m^*) = \max_{x \in Z_m} \frac{E^*(x^+)}{E^*(x^-)}.$$

Proof. A straightforward application of Theorem 1 in Bernardo and Ledoit (2000) gives the desired result.

So $d(m, m^*)$ can be written as the solution to an optimization where m and m^* are used to define the feasible set and the objective function, respectively. The distance gives the maximum gain-loss ratio defined under the transformed measure implied by m^* for any x in Z_m, that is, for any zero-price payoff in the span of all contingent claims under the pricing extension $\pi_m(\bullet)$ implied by m. Note that in this case the role of m^* is clearly explicit.

Cochrane and Saá-Requejo (2000) derive their good-deal (GD) bounds by setting

$$d(m, m^*) = \sigma(m) = \left[E\left(m^2\right) - E^2(m)\right]^{1/2} \tag{2.4}$$

which as discussed below can be looked at as a particular case of Hansen and Jagannathan's (HJ) distance with a constant benchmark SDF. The companion duality result whose proof is well-known goes as follows.

Proposition 2.2 *For the distance in (2.4), the following equality holds*

$$\frac{|E(x)|}{\sigma(x)} \leq d(m, m^*) R^f$$

for all $x \in Z_m$.

As we can see, good-deal bounds skip a benchmark model and define d as the standard deviation of the admissible SDF. This leads to a partial equilibrium interpretation based on the Sharpe ratio. Cerny (2003) extends the good-deal machinery beyond the mean-variance framework. As a necessary first step, he establishes the appropriate reward-for-risk measures that correspond to all utility functions within the HARA family. These constructs correspond to the equivalent of the Sharpe ratio for nonquadratic utility functions which also display linear risk tolerance. Since these measures are also linked to restrictions on certain moments of the admissible SDF's, the theoretical definition of generalized good-deal bounds is straightforward from here. Thus, for each utility function within the above family, these bounds define an interval of prices for a focus payoff that are compatible with a limited expansion of the investment opportunities (the equivalent of the mean-variance frontier for the given utility function), or in other words, a given ceiling on the maximum value of the appropriate reward-for-risk measure. One can also interpret some of the results in Henderson, Hobson, Howison and Kluge (2003) and Jaschke and Küchler (2001) along these lines.

However, as it has already been advanced, it is possible to generalize good-deal bounds by extending their definition in an alternative dimension which uncovers the role of the benchmark model. We will refer thus to these generalizations as *benchmark* good-deal bounds.

3 Benchmark good-deal bounds in a two-period economy

Hansen and Jagannathan (1997) introduce a measure of model misspecification which is based on the distance

$$d(m, m^*) = \left[E(m - m^*)^2\right]^{\frac{1}{2}} \tag{3.1}$$

and whose economic meaning is formalized in the proposition below.

Proposition 3.3 *For the distance given in (3.1), the following equality holds*

$$d(m, m^*) = \max_{\substack{x \in X \\ E(x^2)=1}} |\pi_m(x) - \pi^*(x)|. \tag{3.2}$$

Proof. See Hansen and Jagannathan (1997). ■

Hence, (3.1) gives the maximum pricing discrepancy between π_m and π^* for payoffs in the span of all contingent claims whose second moment are equal to one. An alternative and more intuitive interpretation goes as follows. Suppose there are two different (complete) financial markets where the whole set of contingent claims are traded. In one market, prices are set according to m and in the other one, prices are set according to the benchmark. Arbitrage opportunities do not exist within each market since both m and m^* are strictly positive. However, there are cross-market strategies that give infinite riskless benefits as long as there exist pricing discrepancies between π_m and π^* ($d(m, m^*) > 0$). The normalization $E(x^2) = 1$ guarantees boundedness thereby giving a measure of the size of the above benefits in relative terms. Hence, a restriction on the value of $d(m, m^*)$ is equivalent to a restriction on the optimal value of cross-market arbitrage strategies for those payoffs in X whose second moment is equal to one.

In other words, a ceiling on the value of (3.1) rules out investment opportunities that are too attractive where the level of attractiveness implied by a given admissible SDF is measured in terms of the size of the disintegration that creates with respect to the benchmark market given by m^*.

However, further economic intuition can be derived for (3.1) which we present once again in the form of a proposition.

Proposition 3.4 *For the distance defined in (3.1), it holds that*

$$\frac{|E^*(x)|}{\sigma(x)} \leq d(m, m^*) R^f$$

for all $x \in Z_m$.

Proof. See Appendix. ■

Therefore, a restriction on the above distance (which is equal to the volatility of the difference between the admissible SDF and the benchmark) implies a restriction on an *adjusted* Sharpe ratio, that is, a restriction on the Sharpe ratio where the expectation in the numerator is taken under the risk-adjusted probability measure that the benchmark implies. It should be noted that for a constant benchmark, Proposition 3.4 gives the well-known result that imposing a bound on the volatility of the SDF m implies a bound on the standard Sharpe ratio. Cochrane and Saá-Requejo (2000) use this restriction to derive good-deal (GD) bounds which therefore, can be considered a particular case of the derivation of asset price bounds based on the distance in (3.1).

The HJ distance has a very appealing interpretation. Unfortunately, in its pure form does not allow for an easy extension of asset price bounds to multi-period and continuous-time economies[4] because the positivity constraint could be binding. However, one can modify (3.1) in such a way that also nests the volatility restriction of the SDF (and hence GD bounds) as a special case and it inherits tractability in continuous time. Thus, we will also consider the following functional

$$d(m, m^*) = \left[E^* \left(\frac{m}{m^*} - 1 \right)^2 \right]^{\frac{1}{2}}. \tag{3.3}$$

Note that (3.3) emphasizes the *relative* discrepancy $\frac{m}{m^*}$ between SDF's as opposed to the *absolute* discrepancy $(m - m^*)$. This often leads to tighter asset pricing bounds. Again, the economic meaning of (3.3) is provided in the following two propositions.

Proposition 3.5 *For the distance given in (3.3), the following equality holds*

$$d(m, m^*) = \max_{\substack{x \in X \\ E^*(x^2) = 1}} |\pi_m(x) - \pi^*(x)| E(m^*). \tag{3.4}$$

Proof: See Appendix. ■

Hence, ignoring the constant term $1/R^f$, the only difference between (3.1) and (3.3) is that the normalization in Proposition 1.1 is modified by using the expectation $E^*(\cdot)$. Since the benchmark satisfies by assumption the condition $E(m^*) = 1/R^f$, then $E^*\left(\frac{m}{m^*}\right) = E(m)/E(m^*) = 1$ and we can represent the distance (3.3) as $d(m, m^*) = \sigma^*\left(\frac{m}{m^*}\right)$, where $\sigma^*(\cdot)$ is the standard deviation with respect to the *benchmark* measure.

Proposition 3.6 *For the distance given in (3.3), it holds that*

$$\frac{|E^*(x)|}{\sigma^*(x)} \leq d(m, m^*),$$

for all $x \in Z_m$.

[4] This is a limitation that gain-loss ratio bounds also share.

Proof. See Appendix. ■

So this second alternative imposes a restriction on a ratio that can be naturally referred to as the *benchmark* Sharpe ratio. It also coincides with the standard SR when m^* is constant. This result has clear similarities with the duality interpretation of Bernardo and Ledoit's bounds contained in Proposition 2.1. In both cases, a measure of attractiveness of investments is used (the gain-loss ratio and the Sharpe ratio) and in both cases the objective probability is replaced with the benchmark probability.

To summarize, our two choices in defining BGD bounds, (3.1) and (3.3), nest GD bounds when the benchmark is constant (risk-neutral). Also, their economic meaning has two dimensions: general equilibrium (Propositions 3.3 and 3.5) and partial equilibrium (Propositions 3.4 and 3.6). The first choice lacks tractability beyond the static setup. However, it is based on the HJ measure of misspecification. Hence, it has well-known economic properties which render its analysis useful even if only for illustrative purposes. The distance in (3.3) has not been studied in the literature but as shown above it has an attractive economic meaning and as we will see below, it easily admits an extension to continuous time economies.

We turn now to show how to obtain price bounds based on (3.1). We skip the corresponding derivations for (3.3) in the two-period framework since they can be derived by following similar arguments and we confine their discussion to the continuous time section.

Our formulations will parallel Cochrane and Saá-Requejo's (2000).[5] The problem to be solved can be written as

$$\underline{C} = \min_{m} E\left(mx^{c}\right) \tag{3.5}$$
$$\text{s.t.} \begin{cases} E\left(m\mathbf{x}\right) = \mathbf{p} \\ E\left(m - m^{*}\right)^{2} \leq \overline{d}^{2} \\ m \geq 0 \end{cases} .$$

where maximization gives the upper bound $\overline{C}$. (3.5) has two inequality constraints and the solution can be found by checking all possible combinations of binding and nonbinding constraints. When the second constraint is slack, no-arbitrage bounds are obtained so we will concentrate on the two remaining possibilities.

[5] We assume that $E\left(\mathbf{y}'\mathbf{y}\right) < \infty$ where $\mathbf{y} \equiv \left(\mathbf{x}', x^{c}\right)'$ and that without loss of generality $E\left(\mathbf{y}\mathbf{y}'\right)$ is nonsingular.

Assume that the second constraint in (3.5) is binding and that the positivity constraint is slack and define $\underline{C}'$ as

$$\underline{C}' \equiv \min_{y} E\left(yx^{c}\right) \tag{3.6}$$
$$\text{s.t.} \begin{cases} E\left(y\mathbf{x}\right)=\mathbf{q} \\ E\left(y^{2}\right) \leq \overline{d}^{2} \end{cases}.$$

where $\mathbf{q} \equiv \mathbf{p} - E\left(m^{*}\mathbf{x}\right)$ and $y \equiv m - m^{*}$. It is easy to see that

$$\underline{C} = \underline{C}' + E\left(m^{*}x^{c}\right) \qquad \text{and} \qquad \overline{C} = \overline{C}' + E\left(m^{*}x^{c}\right)$$

where $\overline{C}'$ is defined by replacing min with max in (3.6). Hence, a straightforward application of Proposition 1 in Cochrane and Saá-Requejo gives the result below.

Proposition 3.7 *The discount factor that generates the lower bound is*

$$\underline{m} = x^{*} - \underline{\upsilon} + m^{*}$$

and the bound is

$$\underline{C} = E\left(x^{*}x^{c}\right) - \underline{\upsilon}E\left(\omega^{2}\right) + E\left(m^{*}x^{c}\right)$$

where $x^{*} \equiv q'E\left(\mathbf{x}\mathbf{x}'\right)\mathbf{x}$,

$$\underline{\upsilon} = \sqrt{\frac{\overline{d}^{2} - E\left(x^{*2}\right)}{E\left(\omega^{2}\right)}}$$

and $\omega \equiv x^{c} - E\left(x^{c}\mathbf{x}'\right)E\left(\mathbf{x}\mathbf{x}'\right)^{-1}\mathbf{x}$*. The upper bound is given by* $\overline{\upsilon} = -\underline{\upsilon}$*.*

Note that the size of the bounds has an expression which is identical to equation (12) in Cochrane and Saá-Requejo (2000) with the exception that x^{*} is now the projection of y onto the space of payoffs X. Hence, the bounds are tighter if the value $\overline{d}$ is smaller, if the size of the residual $\sqrt{E\left(\omega^{2}\right)}$ is smaller or, equivalently, if the approximate hedge is better.

Assume now that both constraints in (3.5) bind. By introducing Lagrange multipliers we have

$$\underline{C} = \min_{m>0}\max_{\lambda,\delta>0} E\left(mx^{c}\right) + \lambda'\left[E\left(m\mathbf{x}\right)-\mathbf{p}\right] + \frac{\delta}{2}\left\{E\left[\left(m-m^{*}\right)^{2}\right] - \overline{d}^{2}\right\} \tag{3.7}$$

and the first-order conditions give that in the optimum

$$m = \max\left(-\frac{x^{c} + \lambda' x - m^{*}\delta}{\delta}\right)^{+}. \tag{3.8}$$

As in Hansen, Heaton and Luttmer (1995), we interchange min and max in (3.7) and use (3.8) to obtain after simplifying

$$\underline{C} = \max_{\lambda,\delta>0} E\left\{-\frac{\delta}{2}\left[-\frac{x^c+\lambda' x-\delta m^*}{\delta}\right]^{+2}\right\} - \lambda' p + \frac{\delta}{2}\left[E\left(m^{*2}\right)-\overline{\delta}^2\right].$$

This problem is solved by searching numerically over (λ, δ). Once again, the upper bound is found by replacing max with min but this time it must hold that $\delta < 0$.

4 Benchmark good-deal bounds in continuous time

In this section we present the continuous time derivation of asset price bounds based on the use of the distance in (3.3). An application in the context of Heston's (1993) model is also provided together with a close form expression for the bounds . We conclude with a numerical example where GD bounds are compared with BGD bounds. We focus on the single asset case with one state variable in order to simplify notation and make our exposition less obscure.[6]

4.1 Theoretical Derivations

Some notation is needed before extending our arguments to the continuous time set up. Consider a fixed time interval $[0, T]$ and a probability space $(\Omega, \mathcal{F}, P)$ with filtration denoted by $\{\mathcal{F}_t\}_{t\geq 0}$. We assume that there are one risky asset S_t and one additional (nontradable) state variable V_t. Thus, denote by $W_t = (W_t^s, W_t^v)$ a two-dimensional vector of standard and independent Brownian motions and let the dynamics of the security and the state variable be given by

$$\frac{dS_t}{S_t} = \mu_t dt + \sigma_t dW_t^s. \tag{4.1}$$

and

$$dV_t = \nu_t dt + \eta_t^s dW_t^s + \eta_t^v dW_t^v, \tag{4.2}$$

respectively. The asset pays continuous dividends at rate δ_t which is a deterministic scalar function of S_t and V_t and so are $\mu_t, \sigma_t, \nu_t, \eta_t^s$ and η_t^v. There is also a bond that pays the risk-free rate r_t.

[6] The specific arguments for a multi-asset economy with several state variables are available from the authors upon request.

Let $h_t = (h_t^s, h_t^v)$ be an adapted two-dimensional process which satisfies the Novikov condition. For the measure P, the process h defines a new measure $Q = \mathcal{Q}(h)$ via the Radon-Nikodim derivative

$$\frac{dQ}{dP} = \xi_T$$

where for all $t \leq T$

$$\xi_t = \exp\left[-\int_0^t h_u' dW_u - \frac{1}{2}\int_0^t \|h_u\|^2 du\right] \tag{4.3}$$

By construction ξ is a P-martingale with mean equal to one. Define also the process $\Lambda = \mathcal{L}(h)$ as

$$\Lambda_t \equiv \xi_t B_t \tag{4.4}$$

where $B_t \equiv \exp\left(-\int_0^t r_u du\right)$. By Itô's Lemma we have that

$$\frac{d\Lambda_t}{\Lambda_t} = -r_t dt - h_t' dW_t. \tag{4.5}$$

Intuitively, for a given h, Q represents a candidate martingale measure with associated stochastic discount factor process Λ.[7] The Cameron-Martin-Girsanov theorem guarantees that

$$dW_t + h_t' dt \tag{4.6}$$

is a standard Brownian motion under Q. Also, for any adapted process Y and a given h, it holds that

$$E_t^P\left(\frac{\Lambda_s}{\Lambda_t} Y_s\right) = E_t^Q\left(\frac{B_s}{B_t} Y_s\right) \qquad t \leq s \leq T. \tag{4.7}$$

Similarly, the benchmark model can thus be defined in terms of a vector process h^* with associated martingale measure and SDF processes $Q^* = \mathcal{Q}(h^*)$ and $\Lambda^* = \mathcal{L}(h^*)$, respectively.[8]

If we consider an asset that gives a stream of dividends given by an adapted process X and has a terminal payoff X_T, then its price, C_t, under a given process Λ, is

[7] The continuous time and discrete time discount factors are related by $m_{t,T} = \Lambda_T/\Lambda_t$, where $m_{t,T}$ is an SDF corresponding to the period $[t, T]$. In the sequel, it will be assumed that all processes Λ under consideration have an associated process h which satisfies the Novikov condition.

[8] Note that this definition of the benchmark guarantees that Λ^* prices the risk-free asset.

$$C_t(h) = E_t^P\left(\int_t^T \frac{\Lambda_u}{\Lambda_t} X_u du\right) + E_t^P\left(\frac{\Lambda_T}{\Lambda_t} X_T\right) \tag{4.8}$$

Before formally stating the problem that gives BGD bounds, it will be convenient to describe its feasible set. Since Λ will always be strictly positive in continuous time, we are left with only two restrictions: the pricing constraint and the volatility constraint.

First, note that by definition Λ prices the risk-free bond and if it satisfies

$$E_t^P\left[\frac{d(\Lambda_t S_t)}{\Lambda_t S_t}\right] = -\delta_t dt,$$

it prices the risky asset as well. Consider the market price per unit of risk of the stock given by

$$\lambda_t^s \equiv \frac{\mu_t + \delta_t - r_t}{\sigma_t}.$$

The following result characterizes the set of admissible SDF's and its proof is immediate.

Lemma 4.1 [9] *The process Λ prices the stock if and only if it has associated process satisfying $h^s = \lambda^s$.*

Second, the volatility constraint is given by

$$E_t^{Q^*}\left[\frac{d(\Lambda_t/\Lambda_t^*)}{\Lambda_t/\Lambda_t^*}\right]^2 \leq A_t^2 dt.$$

where A is an adapted process. Its interpretation is formalized in the result below which parallels Proposition 3.6.

Proposition 4.8 *For any security process U paying continuous dividends at rate δ_t^u and any process Λ pricing U, it holds that*

$$\frac{(\mu_t^u - r_t)^2}{\sigma_t^{u2}} dt \leq E_t^{Q^*}\left[\frac{d(\Lambda_t/\Lambda_t^*)}{\Lambda_t/\Lambda_t^*}\right]^2 \tag{4.9}$$

where $\mu_t^u \equiv E_t(dU_t/U_t)/dt + \delta_t^u$ and $\sigma_t^{u2} \equiv E_t\left(dU_t^2/U_t^2\right)/dt$.

[9] This is Lemma 2 in Cochrane and Saá-Requejo (2000). For the case of multiple risky assets the result is identical. In that case, the market price is a vector and it must be assumed that the variance-covariance matrix of the security prices is non-singular.

Proof. See Appendix. ■

Consequently, this inequality establishes that at each time, the maximum generalized Sharpe ratio is bounded by the value of the volatility constraint implied by the selected SDF. Also, the right-hand side in (4.9) can be rewritten as

$$
\begin{aligned}
E_t^{Q^*}\left[\frac{d\left(\Lambda_t/\Lambda_t^*\right)}{\Lambda_t/\Lambda_t^*}\right]^2 &= E_t^P\left[\frac{d\left(\Lambda_t/\Lambda_t^*\right)}{\Lambda_t/\Lambda_t^*}\right]^2 = E_t^P\left[\frac{d\Lambda_t}{\Lambda_t}-\frac{d\Lambda_t^*}{\Lambda_t^*}\right]^2 = \\
&= \|h_t - h_t^*\|^2\,dt
\end{aligned}
$$

where the first equality follows from the well-known fact that the measure transformation leaves the volatility process untouched. Hence, by Lemma 4.1 for admissible SDF's we have that

$$\|h_t - h_t^*\|^2 = (\lambda_t^s - h_t^{s*})^2 + (h_t^v - h_t^{v*})^2$$

and the volatility constraint can be stated as

$$(h_t^v - h_t^{v*})^2 \leq A_t^2 - (\lambda_t^s - h_t^{s*})^2 \tag{4.10}$$

In most partial equilibrium applications, like the one we present below, it makes sense to assume that the benchmark is an admissible SDF. In that case, the volatility constraint discards those SDF's for which $|h_t^v - h_t^{v*}| \geq A_t$.

We are now ready to define BGD bounds. Consider again the asset in (4.8) which is not assumed to belong to the span of the stock and the risk-free bond. Let H^v be the set of processes h^v which satisfy (4.10) on the time interval $[0,T]$, the BGD lower bound is defined as

$$\underline{C}_0\left(h^*\right) \equiv \min_{h^v \in H^v} C_0\left(h^v; h^*, A\right) \tag{4.11}$$

where the corresponding maximization gives the upper bound. By using identical arguments to those in Cochrane and Saá-Requejo (2000, p. 118), one can express (4.11) in differential terms giving a static problem at every time t with a linear objective and a quadratic constraint. Thus, (4.10) is binding and it gives alone the optimal value of h_t^v. Formally,

Proposition 4.9 *The BGD bounds are given by*

$$\underline{C}_0\left(h^*\right) = C_0\left(\underline{h}^v\right) \qquad \textit{and} \qquad \overline{C}_0\left(h^*\right) = C_0\left(\overline{h}^v\right)$$

where

$$\underline{h}_t^v = h_t^{v*} + \sqrt{A_t^2 - (\lambda_t^s - h_t^{s*})^2} \qquad \textit{and} \qquad \overline{h}_t^v = h_t^{v*} - \sqrt{A_t^2 - (\lambda_t^s - h_t^{s*})^2}.$$

Note that GD bounds are a particular case of BGD bounds with $h_t^* = (0,0)$. Obviously, it is possible to define BGD bounds under any objective probability measure P. We put an end to this section by giving an alternative interpretation of BGD bounds that exploits this fact.

Proposition 4.10 *Let $\underline{C}_0(h^*, P)$ denote the lower bound written as a function of the benchmark model and any measure P. Then, it holds that*

$$\underline{C}_0(h^*, P) = \underline{C}_0(\mathbf{0}, Q^*).$$

For the upper bound the equivalent equality is also satisfied.

In other words, BGD bounds when the objective measure is P are equal to GD bounds when the objective measure is Q^*.

4.2 An application: stochastic volatility

The goal of this section is to illustrate the derivation of BGD bounds for the case of European-style option prices in the context of a standard stochastic volatility model. The stock price evolves as in Heston (1993) which is a particular case of our model with

$$\begin{aligned} \frac{dS_t}{S_t} &= (r + sV_t)\,dt + \sqrt{V_t} dW_t^s, \\ dV_t &= (\alpha - \beta V_t)dt + \sigma\sqrt{V_t}\left(\rho dW_t^s + \sqrt{1-\rho^2} dW_t^v\right), \end{aligned}$$

where $r, s, \alpha, \beta, \sigma$, and ρ are constants.[10] Hence, the state variable is the changing variance of the stock return and it follows a square-root mean-reverting process with long-run mean α/β, speed of adjustment β, and variation coefficient σ. Imposing the restriction $\sigma^2 \leq 2\alpha$ is sufficient to guarantee that V_t stays in the open interval $(0, \infty)$ almost surely (see, for example, Cox, Ingersoll, and Ross (1985)). Also, we have that $Corr_t(dS_t/S_t, dV_t)$ is constant and equal to ρ.

We assume that the asset S_t can be traded continuously and frictionlessly. Therefore, for a derivative security written on S_t, one can completely eliminate the derivative's exposure with respect to the price shocks dW_t^s. However, because the asset volatility is nontradable, it is impossible to dynamically hedge the exposure to the volatility shocks dW_t^v.

[10] We specify the drift for dS_t/S_t in the same way as in, for example, Pan (2002) and Benzoni (2000).

Consider an adapted process $h_t = (h_t^s, h_t^\upsilon)$ and its associated martingale measure $Q = \mathcal{Q}(h)$ and SDF process $\Lambda = \mathcal{L}(h)$, where

$$\frac{d\Lambda_t}{\Lambda_t} = -rdt - h_t^s dW_t^s - h_t^\upsilon dW_s^\upsilon.$$

Under Q, the price and variance follows the dynamics[11]

$$\begin{aligned} \frac{dS_t}{S_t} &= \left(r + sV_t - h_t^s\sqrt{V_t}\right)dt + \sqrt{V_t}dW_t^s, & (4.12)\\ dV_t &= (\alpha - \beta V_t - \lambda_t)dt + \sigma\sqrt{V_t}\left(\rho dW_t^s + \sqrt{1-\rho^2}dW_t^\upsilon\right), & (4.13)\end{aligned}$$

where

$$\lambda_t = \sigma\sqrt{V_t}\left(\rho h_t^s + \sqrt{1-\rho^2}h_t^\upsilon\right)$$

By Lemma 4.1 the SDF Λ_t is admissible if its associated h_t satisfies

$$h_t^s = \lambda_t^s = s\sqrt{V_t}. \tag{4.14}$$

In order to achieve tractability, Heston (1993) assumes that the volatility risk premium λ_t is proportional to the variance of the stock, V_t, that is,

$$\lambda_t = \lambda V_t \tag{4.15}$$

for some constant λ. This specification implies that under Q the process for the variance V_t is also mean-reverting, with the long-run mean $\alpha/(\beta+\lambda)$ and speed of adjustment $(\beta+\lambda)$. Heston derives a closed-form formula for prices of standard European-style options via Fourier inversion of the conditional characteristic function. In view of (4.14), the specification in (4.15) obtains if

$$h_t^\upsilon = \upsilon\sqrt{V_t}$$

where υ is constant. Then,

$$\lambda = \sigma\left(s\rho + \upsilon\sqrt{1-\rho^2}\right).$$

Using (4.11) we write the problem for the lower bound of an European-style call option with strike K and maturity T as

$$\underline{C}_0 = \min_{h\in H^\upsilon} C_0(h) = \min_{h\in H^\upsilon} E_0^P\left[\frac{\Lambda_T}{\Lambda_0}\max(S_T - K, 0)\right]$$

[11] For simplicity, we use the same notations W_t^s and W_t^υ for Brownian motions under both the objective measure P and the equivalent martingale measure Q.

where H^{υ} is the set of processes $h_t = \left(s\sqrt{V_t}, h_t^{\upsilon}\right)$ for which the volatility constraint is satisfied. Our benchmark is assumed to be of the form $h^* = h_1^* \equiv \left(s\sqrt{V_t}, \upsilon^*\sqrt{V_t}\right)$, for some constant υ^*, which will be chosen below. This benchmark is admissible and it satisfies (4.15), so that the Heston's formula applies. For GD bounds we set $h_t^* = h_2^* \equiv (0,0)$.

Cochrane and Saá-Requejo consider a case where the ceiling process A in the volatility constraint is a positive constant. However, their approach also works when the ceiling is chosen as a general positive adapted process. This fact is important for our analysis as we set

$$A_t = \overline{A}\sqrt{V_t} \tag{4.16}$$

where $\overline{A}$ is a positive constant.

The motivation for (4.16) is the following. In our economy, the stock's instantaneous Sharpe ratio is proportional to $\sqrt{V_t}$. In this case, a constant ceiling is inappropriate because, for any constant $\overline{A}$, the bound on the conditional Sharpe ratio is too loose and useless in states where V_t is low and it is too tight in states where V_t is high. In fact, for GD bounds, no feasible discount factor exists for *any* constant ceiling $\overline{A}$.

The specification in (4.16) has the good economic property that at each moment the ceiling is proportional to the stock's instantaneous Sharpe ratio. Moreover, this specification allows us to derive the BGD bounds analytically via the Heston formula. Specifically, let $C_0^H(\lambda) = C^H(K,T;S_0,V_0,0,\theta,\lambda)$ denote the Heston call price, where $\theta = (\alpha,\beta,\sigma,\rho)$ is the vector of model parameters which can be estimated from the stock price process S_t. The result below follows from Proposition 4.9. (The Heston formula is provided in the Appendix.).

Proposition 4.11 *For the call price we have that*

$$\underline{C}_0\left(h^*\right) = C_0^H(\underline{\lambda}) \qquad \text{and} \qquad \overline{C}_0\left(h^*\right) = C_0^H(\overline{\lambda})$$

with

$$\underline{\lambda} = \left(s\rho + \underline{\upsilon}\sqrt{1-\rho^2}\right)\sigma \qquad \overline{\lambda} = \left(s\rho + \overline{\upsilon}\sqrt{1-\rho^2}\right)\sigma$$

where for h_1^*

$$\underline{\upsilon} = \upsilon^* + \overline{A} \qquad \overline{\upsilon} = \upsilon^* - \overline{A},$$

and for h_2^*

$$\underline{\upsilon} = \sqrt{\overline{A}^2 - s^2} \qquad \overline{\upsilon} = -\sqrt{\overline{A}^2 - s^2}.$$

To illustrate the practical relevance of our approach consider now a trader (or an option market maker) who needs to price a derivative security on the asset. The trader can compute the call price C_0 using Heston's formula but he is concerned that the volatility risk premium in (4.15) might be misspecified. For example, the actual and unobservable premium might be a general nonlinear function of state variables S_t and V_t. Note that the availability of an explicit formula is a good reason to choose the benchmark h_1^*. That is, even if the trader does not believe that Heston's model is correct, its benchmark is useful, because it gives a closed-form solution and it is certainly better than the GD benchmark h_2^*.

The trader wishes to compute bounds on the call price which allow for some uncertainty about the exact form of the above function. Specifically, suppose that the trader believes that, for all t, S_t, and V_t, it holds that the function lies within the band given by

$$\lambda_l V_t \leq \sigma \left(\rho \lambda_t^s + \sqrt{1-\rho^2} h_t^{\upsilon} \right) \sqrt{V_t} \leq \lambda_h V_t, \tag{4.17}$$

where λ_l and λ_h are constants. The Heston specification in (4.15) obtains by setting $\lambda_l = \lambda_h = \lambda$. When $\lambda_l < \lambda_h$, the condition in (4.17) defines the whole set of "plausible" candidate h processes, which produce a range of candidate call prices. Obviously, we want to have the pricing bounds as tight as possible subject to the constraint that the bounds contain all candidate prices corresponding to the specification in (4.17). It is easy to verify that the optimal setting for $\overline{A}$ is

$$\overline{A} = \max\left(|\lambda_l - \upsilon^*|, |\lambda_h - \upsilon^*|\right)$$

for h_1^* and

$$\overline{A} = \sqrt{s^2 + \max\left({\lambda_l}^2, {\lambda_h}^2\right)}$$

for h_2^*.

Now, suppose that the parameter values of our model are set as

$$\alpha = 0.097, \qquad \beta = 7.1, \qquad \sigma = 0.32, \qquad \rho = -0.53.$$

This choice corresponds to Pan (2002), who fits the stochastic volatility model using the S&P 500 Index options data over the period from 1989 to 1996. She also estimates that $\lambda = -7.6$, $s = 8.6$ and $\upsilon = -22.6$.

Therefore, we assume that the true volatility risk-premium is bounded as in (4.17) with $\lambda_l = \lambda - \Delta$ and $\lambda_h = \lambda + \Delta$, where $2\Delta = 0.0$, 1.0 and 2.0. Other parameters are $r = 0.05$, $S_0 = 100$, $V_0 = \alpha/\beta = 0.0137$, and $T = 0.25$.

We use Proposition 4.11 to compute the bounds where we set $\upsilon^* = \lambda = -7.6$ for $h^* = h_1^*$. The results are shown in Figures 1-3.

Figure 1 presents the lower bound $\underline{C}_0$ and the upper bound $\overline{C}_0$ for a range of strikes K from 80 to 120. For clarity, this figure shows the *difference* between the bounds and the Black-Scholes price C_0^{BS}, for which the volatility parameter is set to $\sqrt{V_0}$. Figure 2 plots the size of the bounds $\overline{C}_0 - \underline{C}_0$. As expected, the size of the bounds are the largest for near-the-money strikes. Finally, Figure 3 plots the Black-Scholes implied volatility for the lower and upper bounds.

It is clear that the bounds for $h^* = h_1^*$ are considerably tighter than for the *risk-neutral* benchmark h_2^*. This is particularly true when Δ approaches zero, in which case the size of the bounds $\overline{C}_0 - \underline{C}_0$ is zero for h_1^* and it is strictly positive for h_2^* (the size of the bounds is about 1.23 for near-the-money call).

The intuition for this is the following. The BGD bounds $\underline{C}_0$ and $\overline{C}_0$ determine the range of possible prices prescribed by a set of admissible discount factors which all lie in the "neighborhood" of a given benchmark, where the "radius" of the neighborhood is given by the ceiling $\overline{A}$. The ceiling must be set large enough so that the neighborhood contains candidate discount factors for which the volatility risk premium function is bounded as in (4.17). For h_2^*, the neighborhood is "symmetric" with respect to the shock dW_t^v. That is, if an admissible Λ_t is in the neighborhood with the representation

$$\frac{d\Lambda_t}{\Lambda_t} = -rdt - s\sqrt{V_t}dW_t^s - h_t^v dW_t^v$$

for some adapted process h_t^v, then another discount factor, which has the same representation except that h_t^v is replaced with -h_t^v, is also admissible and belongs to the neighborhood. This means that, even when $\Delta = 0$, the neighborhood for GD bounds must include all SDF's for which $(h_t^v)^2 \leq \left(\overline{A}^2 - s^2\right) V_t$. Many of the included discount factors may be economically implausible (for example, those that have positive market price for the shocks dW_t^v, which is inconsistent with the empirical findings). Consequently, the pricing bounds are very wide, even when there is no uncertainty about the volatility risk premium.

In contrast, the benchmark discount factor for h_1^* is selected in such a way that the neighborhood only includes economically sensible candidate discount factors, and the corresponding BGD bounds are narrow.

5 Conclusions

The literature on asset pricing bounds has matured at a quick pace since its seminal contributions. However, this prolific research has been so far mainly theoretical and it lacks a satisfactory line of companion empirical work. Our contribution is still theoretical but it includes an attempt to pave the way for future efforts aimed at filling that gap. Perhaps, final steps in this direction should focus on deriving more sophisticated procedures to set the value of the ceiling in computing the bounds. Using the jargon of nonparametric statistics, there seems to be a need for an ex-ante way of fixing this pricing *bandwidth.*

Empirical evidence (e.g., Andersen, Benzoni and Lund (2001) and Eraker, Johannes and Polson (2001)) confirms the presence of jumps in returns and possibly in their volatility. Even though their effect in option pricing is still an issue, their inclusion in modelling asset price bounds might be an interesting theoretical development.

Finally, our calibrations seem to point out an aspect which may deserve further research. They report a surprisingly low sensibility of option prices to misspecifications in the market price of volatility.

Appendix

Proof of Proposition 3.4

Since $E(m^*) = E(m) = 1/R^f$, it holds that

$$d(m, m^*) = \sigma(m - m^*). \tag{A.1}$$

We have that if $x \in Z_m$, then

$$E[(m - m^*)x] = -E(m^* x).$$

which by the definition of covariance may be rewritten as

$$cov(m - m^*, x) = -\frac{E^*(x)}{R^f}.$$

Finally, the definition of correlation and its properties together with (A.1) give the desired result.

□

Proof of Proposition 3.5

Consider the following inner product and norm associated with the expectation operator $E^*(\cdot)$:

$$< h_1, h_2 >_* \equiv E^*(h_1 h_2), \qquad \text{and} \qquad \|h\|_*^2 \equiv < h, h >_*, \qquad h_1, h_2, h \in X.$$

From the Cauchy-Schwartz inequality, it follows that

$$\left| E^* \left(\left(\frac{m}{m^*} - 1 \right) x \right) \right| \leq \left\| \frac{m}{m^*} - 1 \right\|_* \cdot \|x\|_* = d(m, m^*) \cdot \|x\|_*.$$

On the other hand,

$$E^* \left(\left(\frac{m}{m^*} - 1 \right) x \right) = E((m - m^*)x)E(m^*) = (\pi_m(x) - \pi^*(x))E(m^*).$$

This implies that

$$d(m, m^*) \geq \max_{\substack{x \in X \\ \|x\|_* = 1}} |\pi_m(x) - \pi^*(x)| \cdot E(m^*).$$

To establish the fact that the bound is sharp, we note that the fact is obvious when $m = m^*$, in which case $d(m, m^*) = 0$. When $m \neq m^*$, the bound is achieved for a normalized payoff $x/\|x\|_*$ where

$$x = \frac{m}{m^*} - 1.$$

□

Proof of Proposition 3.6

Since $E(m) = E(m^*) = 1/R^f$ it holds that

$$cov^* \left(\frac{m}{m^*}, x \right) = E^* \left(\frac{m}{m^*} x \right) - E^* \left(\frac{m}{m^*} \right) E^*(x) = R^f E(mx) - E^*(x).$$

Therefore,

$$|E^*(x)| \leq \sigma^* \left(\frac{m}{m^*} \right) \sigma^*(x), \qquad \text{for } x \in Z_m,$$

□

In order to simplify notation, in the sequel we will denote by Ψ_t the process Λ/Λ^* for some given Λ^*. The following auxiliary lemma will be used below.

Lemma A.1 *For any SDF processes* Λ *and* Λ^* *and any adapted process* U*, the following holds*

a) *For all* s*, satisfying* $t \le s \le T$

$$E_t^P\left(\frac{\Lambda_s U_s}{\Lambda_t U_t}\right) = E_t^{Q^*}\left(\frac{\Psi_s B_s U_s}{\Psi_t B_t U_t}\right).$$

b) *If* Y *is a security process paying continuous dividends at rate* δ_t^u*. The process* Λ *prices* U *if and only if*

$$E^{Q^*}\left[\frac{d\left(\Psi_t B_t U_t\right)}{\Psi_t B_t U_t}\right] = -\delta_t^u dt.$$

c) *The process* Ψ *is a* Q^**-martingale.*

Proof. Part a: Note that

$$E_t^P\left(\frac{\Lambda_s U_s}{\Lambda_t U_t}\right) = E_t^P\left(\frac{\Lambda_s^* \Psi_s U_s}{\Lambda_t^* \Psi_t U_t}\right) = E_t^{Q^*}\left(\frac{\Psi_s B_s U_s}{\Psi_t B_t U_t}\right) \qquad t \le s \le T$$

where the second equality follows from (4.7).

Part b: By definition Λ prices U if

$$E_t^P\left[\frac{d\left(\Lambda_t U_t\right)}{\Lambda_t U_t}\right] = -\delta_t^u dt$$

and from part a it easily follows that

$$E_t^P\left[\frac{d\left(\Lambda_t U_t\right)}{\Lambda_t U_t}\right] = E_t^{Q^*}\left[\frac{d\left(\Psi_t B_t U_t\right)}{\Psi_t B_t U_t}\right].$$

Part c: By definition, the process Λ/B is a P-martingale and hence,

$$E_t^P\left[\frac{d\left(\Lambda_t/B_t\right)}{\Lambda_t/B_t}\right] = 0$$

which together with part a gives

$$E_t^{Q^*}\left[\frac{d\Psi_t}{\Psi_t}\right] = E_t^P\left[\frac{d\left(\Lambda_t/B_t\right)}{\Lambda_t/B_t}\right] = 0.$$

□

Proof of Proposition 4.8

By Lemma A.1 b) we have that

$$E_t^{Q^*}\left(\frac{d(\Psi_t B_t U_t)}{\Psi_t B_t U_t}\right) = -\delta_t^u dt$$

because by assumption Λ prices U. By using Itô's Lemma the above equation can be rewritten as

$$E_t^{Q^*}\left(\frac{d\Psi_t}{\Psi_t} + \frac{dB_t}{B_t} + \frac{dU_t}{U_t} + \frac{d\Psi_t}{\Psi_t}\frac{dU_t}{U_t}\right) = -\delta_t^u dt$$

which by Lemma A.1 c) implies that

$$E_t^{Q^*}\left(\frac{dU_t}{U_t}\right) - r_t dt + \delta_t^u dt = -Cov_t^{Q^*}\left(\frac{d\Psi_t}{\Psi_t}, \frac{dU_t}{U_t}\right),$$

and by the definition of correlation and its properties we get the desired result.

□

Proof of Proposition 4.10

Let H^1 and H^2 be the feasible sets in obtaining $\underline{C}_0\left(h^*, P\right)$ and $\underline{C}_0\left(\mathbf{0}, Q^*\right)$, respectively. We have that

$$E_t^{Q^*}\left[\frac{d\left(\Lambda_t/\Lambda_t^*\right)}{\Lambda_t/\Lambda_t^*}\right]^2 = E_t^{Q^*}\left[\frac{d\left(\Psi_t\right)}{\Psi_t}\right]^2 = E_t^{Q^*}\left[\frac{d\left(\Psi_t B_t\right)}{\Psi_t B_t}\right]^2$$

which together with Lemma A.1 b) gives that

$$\Lambda \in H^1 \Leftrightarrow \Psi B \in H^2.$$

Hence,

$$\begin{aligned}
\underline{C}_0\left(h^*, P\right) &= \min_{h\in H^1} E_0^P\left(\int_0^T \frac{\Lambda_u}{\Lambda_0} X_u du\right) + E_0^P\left(\frac{\Lambda_T}{\Lambda_0} X_T\right) = \\
&= \min_{h\in H^1} E_0^{Q^*}\left(\int_0^T \frac{\Psi_u B_u}{\Psi_0 B_0} X_u du\right) + E_0^{Q^*}\left(\frac{\Psi_T B_T}{\Psi_0 B_0} X_T\right) = \\
&= \min_{h\in H^2} E_0^{Q^*}\left(\int_0^T \frac{\Lambda_u}{\Lambda_0} X_u du\right) + E_0^{Q^*}\left(\frac{\Lambda_T}{\Lambda_0} X_T\right) = \underline{C}_0\left(\mathbf{0}, Q^*\right),
\end{aligned}$$

where the second equality follows from Lemma A.1.

□

The Heston formula

Given the risk-neutral model (4.12)-(4.13) and the market price of volatility risk as in (4.15), Heston (1993) obtains the following closed-from formula for the European call option:

$$C^H(K,T;S_t,V_t,t,\theta,\lambda) = S_tP_1 + e^{-r\tau}KP_2,$$

where for $j = 1,2$

$$\begin{aligned}
P_j(S_t,V_t,\tau;K) &= \frac{1}{2} + \frac{1}{\pi}\int_0^\infty Re\left[\frac{e^{-i\phi\ln(K)}F_j(S_t,V_t,\tau;\phi)}{i\phi}\right]d\phi,\\
F_j(S_t,V_t,\tau;\phi) &= \exp[C(\tau;\phi) + D(\tau;\phi)V_t + i\phi\ln(S_t)],\\
C(\tau;\phi) &= r\tau\phi i + \frac{\alpha}{\sigma^2}\left[(\beta_j - \rho\sigma\phi i + h)\tau - 2\ln\left(\frac{1-ge^{h\tau}}{1-g}\right)\right],\\
D(\tau;\phi) &= \frac{\beta_j - \rho\sigma\phi i + h}{\sigma^2}\left(\frac{1-e^{h\tau}}{1-ge^{h\tau}}\right),\\
g &= \frac{\beta_j - \rho\sigma\phi i + h}{\beta_j - \rho\sigma\phi i - h},\\
h &= \sqrt{(\rho\sigma\phi i - \beta_j)^2 - \sigma^2(2u_j\phi i - \phi^2)},
\end{aligned}$$

and

$$\tau = T - t, \qquad u_1 = 1/2, \qquad u_2 = -1/2, \qquad \beta_1 = \beta + \lambda - \rho\sigma, \qquad \beta_2 = \beta + \lambda.$$

References

Andersen, T.G, L. Benzoni, and J. Lund (2001), "An Empirical Investigation of Continuous-Time Equity Returns Models", *Journal of Finance*, vol 57, no. 2, 1239-1284.

Benzoni, L. (2000), "Pricing Options under Stochastic Volatility: An Empirical Investigation," working paper, University of Minnesota.

Bernardo, A.E., and O. Ledoit (2000), "Gain, Loss, and Asset Pricing", *Journal of Political Economy*, vol. 108, no. 1, 144-172.

Cerny, A. (2003), "Generalized Sharpe Ratios and Asset Prices in Incomplete Markets", *European Finance Review*, vol 7, no. 2, 191-233.

Cochrane, J. H., and J. Saá-Requejo (2000), "Beyond Arbitrage: Good-Deal Asset Price Bounds in Incomplete Markets", *Journal of Political Economy,* vol. 108, no. 1, 79-119.

Eraker B., M. S. Johannes, and N.G. Polson (2003), "The Impact of Jumps on Volatility and Returns", *Journal of Finance*, vol. 58, no. 3, 1269-1300.

Hansen L.P., and R. Jagannathan (1997), "Assessing Specification Errors in Stochastic Discount Factor Models", *Journal of Finance*, vol 52, no.2, 557-590.

Hansen, L.P, J. Heaton, and Erzo Luttmer (1995), "Econometric Evaluation of Asset Pricing Models", *Review of Financial Studies*, vol 8, no. 2, 237-274.

Henderson, V., D. Hobson, S. Howison, and T. Kluge (2003), "A Comparison of Option Prices Under Different Pricing Measures in a Stochastic Volatility Model with Correlation", working paper, University of Oxford.

Jaschke S., and U. Küchler (2001), "Coherent Risk Measures and Good-Deal Bounds", *Finance and Stochastics*, vol. 5, no. 2, 181-200.

Pan, J. (2002), "The Jump-Risk Premia Implicit in Options: Evidence from an Integrated Time-Series Study", *Journal of Financial Economics,* 63, 3-50.

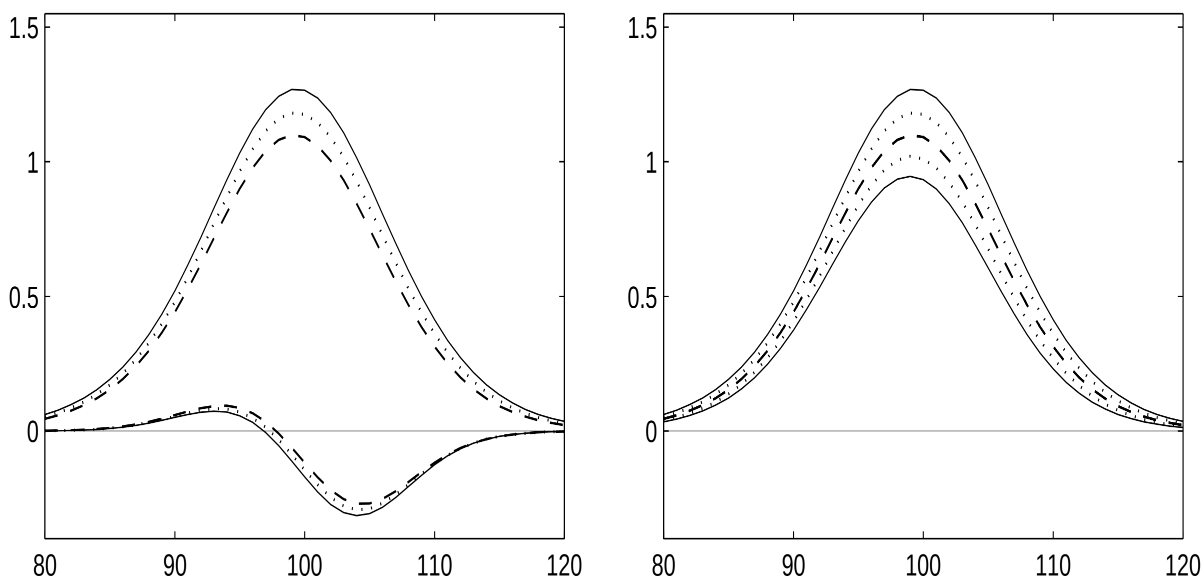

Figure 1: The difference between option price bounds and the Black-Scholes price versus strike K, for h_1^* (the right panel) and case h_2^* (the left panel). The differences $\underline{C}_0 - C_0^{BS}$ and $\overline{C}_0 - C_0^{BS}$ are computed using parameter values in Section 4.2, when $\lambda_h - \lambda_l$ is equal to 0.0 (the dashed lines), 1.0 (the dotted lines), and 2.0 (the solid lines).

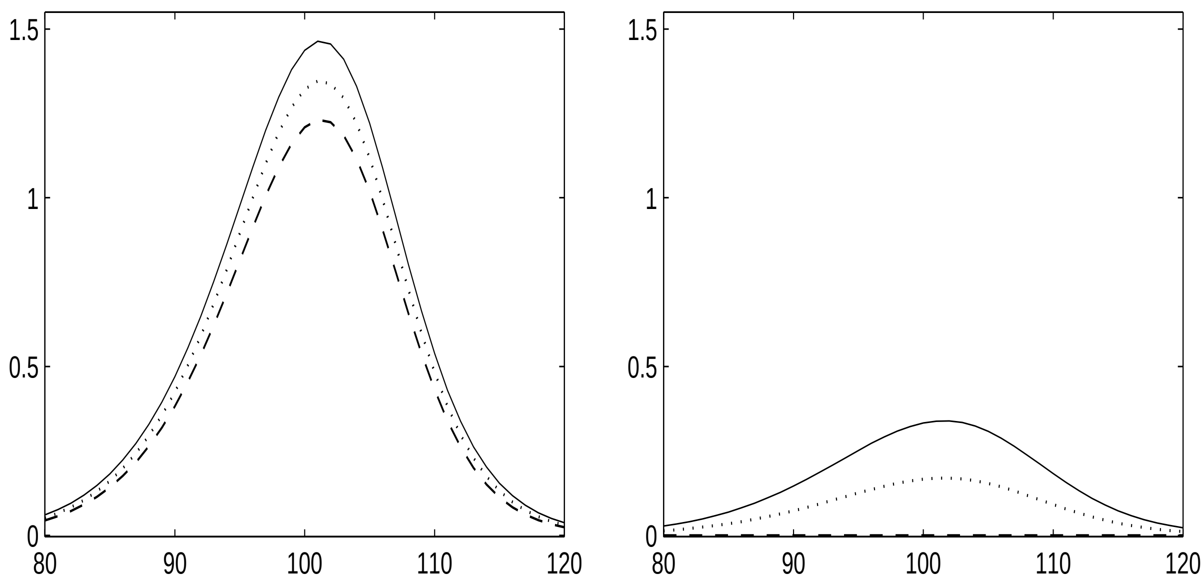

Figure 2: The size of the bounds $\overline{C}_0 - \underline{C}_0$ versus strike K, for case h_1^* (the right panel) and case h_2^* (the left panel). The bounds are computed using parameter values in Section 4.2, when $\lambda_h - \lambda_l$ is equal to 0.0 (the dashed lines), 1.0 (the dotted lines), and 2.0 (the solid lines).

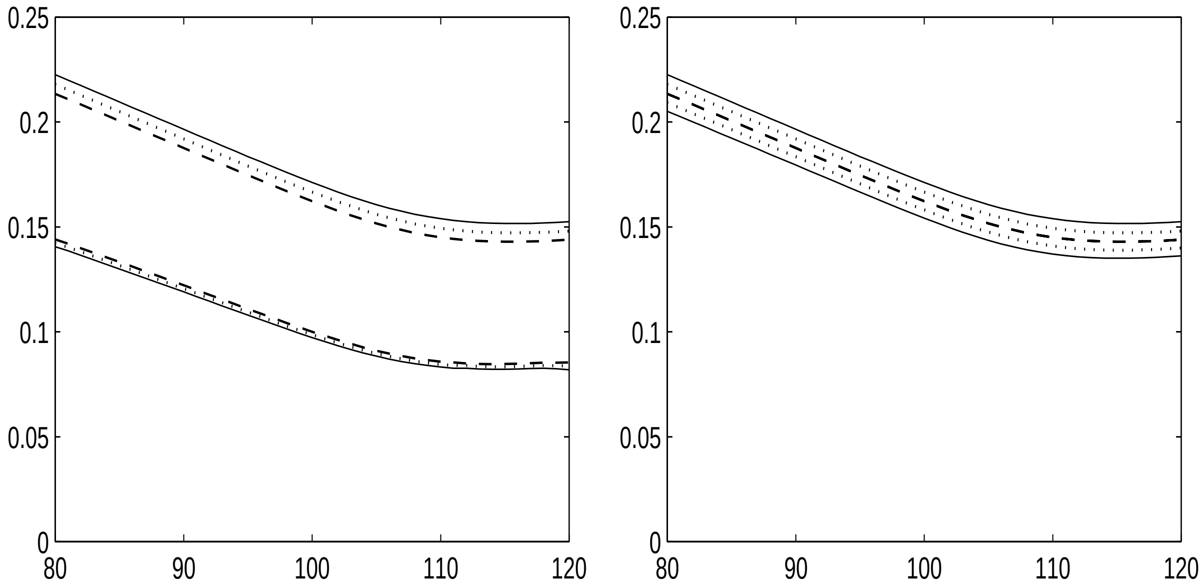

Figure 3: The Black-Scholes implied volatility versus strike K, for h_1^* (the left panel) and h_2^* (the right panel). The implied volatilities are computed for the lower and upper bounds $\underline{C}_0$ and $\overline{C}_0$ using parameter values in Section 4.2, when $\lambda_h - \lambda_l$ is equal to 0.0 (the dashed lines), 1.0 (the dotted lines), and 2.0 (the solid lines).