

A General Framework for the Derivation of Asset Price Bounds: An Application to Stochastic Volatility Option Models*

Oleg Bondarenko[†] and Iñaki R. Longarela[‡]

March 2009 (first version: April 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 demonstrates that 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

*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 Economics and Management, NFH, University of Tromsø, N-9037 Tromsø, Norway. E-mail: Inaki.Rodriguez@nfh.uit.no.

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 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. Unfortunately, in many realistic situations the resulting *no-arbitrage* bounds are too imprecise to be economically interesting.

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

In the first approach, Cochrane and Saá-Requejo concentrate on the Sharpe ratio as the relevant measure of attractiveness of investment opportunities. As demonstrated in Hansen and Jagannathan (1991), a bound on the maximum available Sharpe ratio is equivalent to a bound on the volatility of the admissible SDF. Building on this insight, Cochrane and Saá-Requejo suggest to rule out not only usual arbitrage opportunities but also *good deals*, or investments with *too* high Sharpe ratios, thus obtaining tighter price bounds when markets are incomplete.

In the second approach, Bernardo and Ledoit (2000) propose to measure the attractiveness of an investment by the *gain-loss* ratio, where gain (loss) is defined as the expectation of the investment's positive (negative) excess payoffs computed under a *benchmark* asset pricing model. Bernardo and Ledoit demonstrate that the existence of investments with a high gain-loss ratio is linked to SDF's exhibiting extreme deviations from the benchmark SDF. Therefore, by appropriately choosing a benchmark SDF and imposing an upper value on the above deviations one can again tighten the no-arbitrage interval.

An important advantage of Bernardo and Ledoit's approach is its ability to incorporate a benchmark or reference asset pricing model. In their approach, the benchmark represents an initial guess of an econometrician about the true model. Even though the econometrician might have legitimate concerns about potential model misspecification, the benchmark is nevertheless useful, as it helps in locating more accurately the set of candidate SDF's. Intuitively, this approach combines in a single framework the precision of the *model-based* pricing technique and the robustness of the *no-arbitrage* pricing technique. In contrast, the approach of Cochrane and Saá-Requejo does not incorporate a benchmark model, but it offers another important advantage – analytical tractability in the continuous-time setting.

Our first contribution and starting point is methodological. We demonstrate how the derivation of asset price bounds can be generally presented as an optimization over a certain set of admissible SDF's. This feasible set is determined by a distance measure, a specific point or benchmark SDF and a ceiling on the value of the distance from the point. It is then easily shown that the above contributions are particular cases of this general framework.

Next, we introduce within this setup a new type of asset price bounds, which blend the virtues of the these two seminal papers. Specifically, we generalize the good-deal approach by incorporating information about a given benchmark SDF. We introduce a new measure of model misspecification and define the attractiveness of an investment opportunity by the *benchmark* Sharpe ratio. The benchmark Sharpe ratio generalizes the standard Sharpe ratio, and the two concepts coincide when the benchmark is constant. We then establish a duality result which links the maximum available benchmark Sharpe ratio to the deviation of a candidate SDF from the benchmark SDF, as quantified by our new measure. Ruling out investments with high benchmark Sharpe ratios produces tighter price bounds, which are referred to as the *benchmark good-deal* (BGD) bounds.¹

This approach has a number of important features. First, it nests Cochrane and Saá-Requejo’ methodology as a special case of the constant (or, risk-neutral) benchmark SDF. More generally, it allows for a flexible and economically more sensible choice of a benchmark SDF. Second, our approach retains analytical tractability in continuous time, as opposed to the gain-loss methodology. Third, our approach has an intuitive property that it includes the model-based and no-arbitrage pricing techniques as two extreme cases of a single framework. When the maximum value of our model misspecification measure is set to zero, the model-based pricing obtains, while when the maximum value is set to infinity, one is back to the no-arbitrage pricing. Forth and most importantly, the performance of our BGD bounds is often much better than that of the standard GD bounds. By construction, the BGD bounds are never wider than the standard GD bounds, as one can always choose the constant benchmark SDF. In many situations, however, incorporating a more reasonable benchmark could improve the BGD bounds substantially.

It should be noted that as a first step in our derivations, we rely on a duality result which involves an elegant simplicity with far-reaching economic implications. Even though we regard these attributes as a virtue, the selection of the appropriate measure of model misspecification is not a trivial exercise. While many choices are possible in a static setting, our pick is, as far as we can tell, the only one that grants tractability in continuous time and retains the desired features listed above.

Finally, in order to illustrate the performance of our BGD bounds, we apply our methodology to a classical problem in finance of pricing an option in incomplete market. We consider Heston (1993) stochastic volatility model. In this model, one cannot perfectly replicate the option’s payoff by trading in the underlying only and the no-arbitrage pricing produces not a single price, but a wide range of prices. We consider a trader who does not know the exact form of the volatility risk premium function, but who nevertheless can place upper and lower limits on the function’s possible values. For this application, we derive a closed-form solution for the BGD bounds, which include the standard GD bounds as a particular case. To our best knowledge, this is the first

¹It is worth mentioning that there is an alternative way to incorporate a benchmark into the good-deal methodology. This extension relies on Hansen and Jagannathan (1997) (HJ) measure of model misspecification and it includes the good-deal bounds as a particular case. We analyze the economic aspects of this extension in a two-period model. However, the HJ measure does not allow for an easy extension to continuous time. For that reason, we introduce a different measure, which does the trick.

application which provides analytical expressions for the GD bounds in continuous time. A calibration exercise then allows us to conclude that under realistic conditions the BGD bounds can be considerably tighter than the standard GD bounds.

In our application, we focus on Heston (1993) because this popular model provides a particularly simple, clear-cut illustration of comparative advantages of our methodology. For this continuous-time problem, Bernardo and Ledoit's approach is not applicable, while the benefits of incorporating the benchmark model for constructing the BGD bounds are already very pronounced. It also helps that for this problem the bounds can be calculated analytically. Still, it is worth stressing that our approach is not limited to Heston (1993) and can be applied to more general models, although the bounds for those models might have to be computed numerically.

The remainder of the paper is organized as follows. Some methodological preliminaries are introduced in the next section. Section 3 discusses the economic properties of the bounds in a static economy. Section 4 deals with the continuous-time case and also presents the application to stochastic volatility option pricing. Finally, some conclusions and suggestions for further research are given in Section 5. All necessary mathematical proofs are confined to the 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

$$\mathbf{p} = E(m\mathbf{x}). \quad (1)$$

Let M^+ denote the set of all strictly positive SDF's.² These random variables give the price of any payoff $x \in X$ through their *pricing extensions* defined as

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

which assign the same price to all payoffs x in X^b but, in general, differ for payoffs $x \notin X^b$. For any $m \in M^+$, let Z_m denote the set of excess returns under the pricing function $\pi_m(\bullet)$, or

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

²Occasionally, we will use the term SDF even if condition (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.

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. Let m^* denote a strictly positive *benchmark* SDF, which will always be assumed to satisfy $E(m^*) = 1/R^f$.³ For any $x \in X$, consider its implied pricing function

$$\pi^*(x) \equiv E(m^*x).$$

If $E^*(\cdot)$ denotes the expectation under the risk-neutral probability measure implied by m^* , then we the implied pricing function can be represented as

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

In deriving asset price bounds the goal is always to tighten the set M^+ in accordance with some 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 or benchmark model and establishing how much discrepancy, measured in terms of a given distance,⁴ one is willing to allow. Unconstrained discrepancy will lead to the no-arbitrage bounds, while zero discrepancy will imply model-based pricing. It is in between these two extremes where our interest lies.

Formally, let $d(m, m^*)$ denote a distance between $m \in M^+$ and a given benchmark model m^* . Let \bar{d} be an appropriate ceiling on the maximum value of the above distance. To derive price bounds, one discards those $m \in M^+$ for which $d(m, m^*) > \bar{d}$. In particular, the lower bound on the price of a payoff $x^c \notin X^b$ can be found by solving the following constrained optimization problem:

$$\underline{C} = \min_{\substack{m \in M^+ \\ d(m, m^*) \leq \bar{d}}} E(mx^c) \quad (2)$$

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

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

because if the left-hand inequality does not hold the feasible set in (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 the distance $d(\cdot, \cdot)$ and the benchmark m^* , which are in the end the two key ingredients in

³In principle, the benchmark m^* can be assumed to be more general. In particular, it does not have to price the risk-free asset and, in many cases, it could even take negative values. However, we will not need such generality in our exposition.

⁴We somewhat abuse the term “distance” as the symmetry property does not necessarily hold.

the above derivation. The distance usually has a distinct economic interpretation which is obtained through a corresponding duality result. The benchmark m^* represents a reasonable initial guess of an econometrician about the true model and it helps one to better characterize the set of candidate SDF's, thus, yielding tighter bounds. In particular, when the ceiling \bar{d} goes to zero, the set of candidate SDF's shrinks to include only the benchmark m^* . When the ceiling \bar{d} goes to infinity, the set of candidate SDF's grows to include all SDF's in M^+ . Thus, the choice of the ceiling controls the trade-off between the precision of the model-based approach and the robustness of the no-arbitrage approach.

In the case when the benchmark m^* plays no explicit role in a derivation, the distance is determined by some moment of the candidate SDF's.

The above approach 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. In their approach, the distance is defined as

$$d(m, m^*) = d_{GL}(m, m^*) \equiv \frac{\sup \frac{m}{m^*}}{\inf \frac{m}{m^*}}. \quad (3)$$

On the other hand, Cochrane and Saá-Requejo (2000) derive their good-deal (GD) bounds by setting

$$d(m, m^*) = d_{GD}(m) \equiv \sigma(m) = [E(m^2) - E^2(m)]^{1/2}, \quad (4)$$

which can be viewed as a particular case of Hansen and Jagannathan's (HJ) distance with the constant benchmark SDF. The companion duality results go as follows.

Proposition 1 *The following relationships hold:*

1. *For the distance given in (3),*

$$\frac{E^*(x^+)}{E^*(x^-)} \leq d_{GL}(m, m^*), \quad \text{for all } x \in Z_m.$$

2. *For the distance given in (4),*

$$\frac{|E(x)|}{\sigma(x)} \leq d_{GD}(m) R^f, \quad \text{for all } x \in Z_m. \quad (5)$$

Proof. See Bernardo and Ledoit (2000) and Hansen and Jagannathan (1991). ■

In view of this proposition, the distance $d_{GL}(m, m^*)$ provides an upper bound on the maximum possible gain-loss ratio for any payoff 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 . To restrict the set of prices that can be assigned to assets, Bernardo and Ledoit propose to rule out investment opportunities with too high gain-loss ratios. In this approach, one

obtains bounds on asset prices by imposing a suitable ceiling \bar{d} and discarding admissible SDF's m whose distance $d_{GL}(m, m^*)$ to the benchmark m^* is greater than \bar{d} .

The gain-loss ratio is defined under the transformed measure implied by m^* and the role of the benchmark m^* is explicit. In contrast, the benchmark model plays no role in deriving the good-deal bounds, for which the distance is defined as the standard deviation of an admissible SDF m .

Cochrane and Saá-Requejo show how to calculate the good-deal bounds in one-period, multi-period, and continuous-time settings. In the one-period setting, the lower (upper) bound is generated by choosing a candidate SDF m which minimizes (maximizes) an asset's value, subject to constraints that

- i) m correctly prices the basis assets,
- ii) m is non-negative, and
- iii) the standard deviation $\sigma(m)$ does not exceed a prespecified threshold.

The multi-period problem can be solved recursively as a sequence of one-period problems. In the continuous-time limit, the dynamic optimization problem greatly simplifies. In particular, constraint ii) becomes unnecessary and constraint iii) reduces to the simple restriction on the instantaneous volatility of the discount factor process. This allows Cochrane and Saá-Requejo to derive a certain partial differential equation which characterizes the dynamics of GD bounds. Analytical tractability in continuous time is an important advantage of the GD methodology.

Cerny (2003) extends the good-deal machinery beyond the mean-variance framework. As a necessary first step, he establishes appropriate reward-for-risk measures for all utility functions within the HARA family. These constructs correspond to the equivalent of the Sharpe ratio for non-quadratic utility functions, which also display linear risk tolerance. Since these measures are also linked to restrictions on certain moments of the admissible SDF's, one can use this link to define the generalized good-deal bounds. For each utility function within the HARA family, the generalized good-deal bounds represent 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 (2005) and Jaschke and Küchler (2001) along these lines.

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

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

In defining BGD bounds, our primary motivation is to combine the main virtues of the approaches by Cochrane and Saá-Requejo and Bernardo and Ledoit: tractability in continuous time and ability to incorporate a benchmark model. Specifically, we would like our formulation to satisfy the following properties. The approach should

- (a) nest GD bounds as a particular case;
- (b) be able to incorporate a flexible benchmark m^* ;
- (c) allow for an intuitive duality characterization;
- (d) retain analytical tractability in continuous time.

In choosing an appropriate distance, some measure of discrepancy with the benchmark must be used. Natural choices along these lines are distances based on the *absolute* difference $(m - m^*)$ or the *relative* discrepancy m/m^* . Thus, one can consider the following two functionals:⁵

$$d(m, m^*) = [E(m - m^*)^2]^{\frac{1}{2}}, \quad (6)$$

and

$$d(m, m^*) = \frac{1}{R^f} \left[E \left(\frac{m}{m^*} - 1 \right)^2 \right]^{\frac{1}{2}}. \quad (7)$$

The distance in (6) has been introduced by Hansen and Jagannathan (1997) and it is used in defining their well-known measure of model misspecification. Note that we use the constant factor $1/R^f$ in the definition (7) to ensure that this distance coincides with the distance of Cochrane and Saá-Requejo (2000) in the special case of the constant benchmark $m^* = 1/R^f$.

Additionally, we consider two modifications of the distances in (6) and (9) which have not been previously studied in the literature. In these formulations, the expectation E is replaced with E^* :

$$d(m, m^*) = [E^*(m - m^*)^2]^{\frac{1}{2}}, \quad (8)$$

and

$$d(m, m^*) = \frac{1}{R^f} \left[E^* \left(\frac{m}{m^*} - 1 \right)^2 \right]^{\frac{1}{2}}. \quad (9)$$

It is easy to verify that the four alternatives in (6)-(9) all satisfy properties (a) and (b) above. In particular, for the constant benchmark $m^* = 1/R^f$, they all reduce to the standard deviation $\sigma(m)$. Furthermore, for a general benchmark m^* , the four

⁵The possibility of using distances (6) and (7) to incorporate a benchmark has been discussed in Cochrane and Saá-Requejo (2000).

specifications for $d(m, m^*)$ allow one to include the model-based and no-arbitrage pricing techniques as two extreme cases of a single framework. When the maximum value of $d(m, m^*)$ is set to zero, the model-based pricing obtains, and when the maximum value is set to infinity, the no-arbitrage pricing obtains.

However, only two of the four listed distances, (6) and (9), admit appealing economic interpretations, which can be found in the two propositions below. Before stating the propositions, it is convenient to rewrite the distances in (6) and (9) as

$$d(m, m^*) = \sigma(m - m^*), \quad (10)$$

and

$$d(m, m^*) = \frac{1}{R^f} \sigma^* \left(\frac{m}{m^*} \right), \quad (11)$$

where $\sigma^*(\cdot)$ denotes the standard deviation with respect to the expectation E^* . The above equations follow from the fact that $E(m) = E(m^*) = 1/R^f$ and that

$$E^* \left(\frac{m}{m^*} \right) = E(m)/E(m^*) = 1.$$

Proposition 2 *For the distance given in (6),*

$$d(m, m^*) = \max_{\substack{x \in X \\ E(x^2)=1}} |\pi_m(x) - \pi^*(x)|, \quad (12)$$

and

$$\frac{|E^*(x)|}{\sigma(x)} \leq d(m, m^*) R^f, \quad \text{for all } x \in Z_m. \quad (13)$$

Proposition 3 *For the distance given in (9),*

$$d(m, m^*) = \max_{\substack{x \in X \\ E^*(x^2) = 1}} |\pi_m(x) - \pi^*(x)|, \quad (14)$$

and

$$\frac{|E^*(x)|}{\sigma^*(x)} \leq d(m, m^*) R^f, \quad \text{for all } x \in Z_m. \quad (15)$$

In view of Proposition 2, the distance in (6) gives the maximum pricing discrepancy between π_m and π^* for payoffs in the span of all contingent claims whose second moments are equal to one. An alternative and more intuitive interpretation goes as follows. Suppose there are two different 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 m^* . 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^* , that is, as long as $d(m, m^*) > 0$. 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 the distance in (6) 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^* . Furthermore, a restriction on this 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, for which the expectation in the numerator is taken under the risk-adjusted probability measure that the benchmark implies.

By Proposition 3, the distance in (9) imposes a restriction on a ratio that can be naturally referred to as the *benchmark* Sharpe ratio. This relationship has clear similarities with the duality interpretation of Bernardo and Ledoit's bounds contained in Proposition 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. Note that the only difference between equalities (12) and (14) is that the normalization in (14) is modified by using the expectation $E^*(\cdot)$.

When m^* is constant, both the adjusted and benchmark Sharpe coincide with the standard Sharpe ratio and the inequalities in (13) and in (15) reduce to the inequality in (5). It is interesting that, although both choices (6) and (9) satisfy properties (a)-(c), only the latter one satisfies property (d), as the HJ distance in (6) does not allow for an easy extension of asset price bounds to continuous-time economies.⁶ This latter flaw is actually shared by many other alternative distances whose analysis may be pursued. Alternative distances may display attractive economic properties in a static setting, but, as far as we can tell, the distance in (9) is the only one that allows for an analytically tractable extension to the continuous-time setting. This is an unfortunate fact, since it precludes an analysis of the robustness of our continuous time results across a wider collection of distances. All these considerations justify our final choice which we turn to present now.

⁶In one-period or multi-period setting, one can define a new variable $y \equiv m - m^*$ and apply the same arguments as in Cochrane and Saá-Requejo (2000) to y instead of m . The derivation of the bounds can be obtained in this way with only minor adjustments. However, this approach does not lead to a tractable differential characterization in continuous time.

4 Benchmark good-deal bounds in continuous time

In this section we present the continuous time derivation of asset price bounds based on the distance in (9). We also illustrate our methodology in the context of Heston (1993) stochastic volatility model. For this model, the bounds on option prices can be derived in closed-form. We conclude with a numerical example where the GD bounds are compared to the BGD bounds. In order to simplify exposition, this section focuses on the case of a single asset and one state variable.⁷

4.1 Theoretical Derivations

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

$$\frac{dS_t}{S_t} = \mu_t dt + \sigma_t dW_t^s,$$

$$dV_t = \nu_t dt + \eta_t^s dW_t^s + \eta_t^v dW_t^v,$$

respectively. The stock 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 also exists a bond that pays the risk-free rate r_t .

Let $h_t = (h_t^s, h_t^v)$ be an adapted two-dimensional process which satisfies the Novikov condition. For a given process h , we can define a new measure $Q = \mathcal{Q}(h)$ and associated with it discount factor process $\Lambda = \mathcal{L}(h)$. Specifically, the measure Q is defined from the measure P and the process h via the Radon-Nikodim derivative

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

where for all $t \leq T$

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

The discount factor process Λ is defined as

$$\Lambda_t \equiv \xi_t B_t,$$

⁷The specific arguments for a multi-asset economy with several state variables are available from the authors upon request.

where $B_t \equiv \exp\left(-\int_0^t r_u du\right)$. By Itô's Lemma, the discount factor process satisfies⁸

$$\frac{d\Lambda_t}{\Lambda_t} = -r_t dt - h'_t dW_t.$$

By construction, ξ is a P -martingale with mean equal to one. Also, for any adapted process Y , it holds that

$$E_t^P\left(\frac{\Lambda_s}{\Lambda_t} Y_s\right) = E_t^Q\left(\frac{B_s}{B_t} Y_s\right) \quad t \leq s \leq T. \quad (16)$$

Along the same lines, we can define the benchmark model in terms of another vector process h^* with the corresponding martingale measure and SDF process denoted as $Q^* = \mathcal{Q}(h^*)$ and $\Lambda^* = \mathcal{L}(h^*)$.⁹

If we consider a security that pays a stream of dividends given by an adapted process X and has a terminal payoff X_T , then its price $C_t = C_t(h)$ under a candidate SDF process Λ is

$$C_t(h) = E_t^P\left(\int_t^T \frac{\Lambda_u}{\Lambda_t} X_u du\right) + E_t^P\left(\frac{\Lambda_T}{\Lambda_t} X_T\right) \quad (17)$$

Before formally stating the problem that defines the BGD bounds on the price of the security, it is convenient to first describe the set of feasible SDF processes. Since Λ is always strictly positive, only two restrictions are needed in continuous time: the pricing constraint and the volatility constraint.

First, we note that Λ prices the risk-free bond by definition. Moreover, Λ also prices the risky asset, if it satisfies

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

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 lemma characterizes the set of admissible SDF's and its proof is immediate.

Lemma 1 *The SDF process Λ prices the stock if and only if it has associated process satisfying $h^s = \lambda^s$.*¹⁰

⁸The 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 what follows, it will be assumed that all processes Λ under consideration have an associated process h which satisfies the Novikov condition.

⁹Note that this definition of the benchmark guarantees that Λ^* prices the risk-free asset.

¹⁰This 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.

Second, for a fixed benchmark SDF process Λ^* , the volatility constraint in continuous time is given as

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

where A is some adapted process. The following result establishes a formal link between the maximum benchmark Sharpe ratio and the upper bound on the volatility constraint.

Proposition 4 *For any security process U paying continuous dividends at rate δ_t^u and any process Λ pricing U , it holds that*

$$\frac{(\mu_t^u - r_t)^2}{\sigma_t^{u2}} dt \leq E_t^{Q^*} \left[\frac{d(\Lambda_t/\Lambda_t^*)}{\Lambda_t/\Lambda_t^*} \right]^2, \quad (18)$$

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

The inequality in (18) is the continuous-time equivalent of the inequality in (15). To see this more clearly, we can rewrite the inequality in (15) in the following equivalent form

$$\frac{|E^*(x)|}{\sigma^*(x)} \leq \frac{\sigma^* \left(\frac{m}{m^*} \right)}{E^* \left(\frac{m}{m^*} \right)}.$$

The right-hand side of (18) can be simplified as

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

where the first equality follows from the fact that the measure transformation leaves the volatility process unchanged. Lemma 1 implies that for an admissible SDF processes it holds 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 now be stated as

$$(h_t^v - h_t^{v*})^2 \leq A_t^2 - (\lambda_t^s - h_t^{s*})^2. \quad (19)$$

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

We are now ready to define the BGD bounds. Consider again the security in (17) whose price process is not spanned by the stock and the risk-free bond. Let H^v be the set of processes h^v which satisfy (19) on the time interval $[0, T]$. The lower bound is defined as

$$\underline{C}_0(h^*) \equiv \min_{h^v \in H^v} C_0(h^v; h^*, A), \quad (20)$$

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 (20) in differential terms describing a static optimization problem at every time t . Importantly, the static problem at each time has a linear objective and one quadratic constraint in (19). This implies that the constraint is always binding and it alone determines the optimal value of h_t^v . Formally,

Proposition 5 *The BGD bounds are given by*

$$\underline{C}_0(h^*) = C_0(\underline{h}^v), \quad \overline{C}_0(h^*) = C_0(\overline{h}^v),$$

where

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

Note that the GD bounds obtain as a particular case of the BGD bounds with the process $h_t^* \equiv (0, 0)$. It is also interesting to mention that it is possible to define the BGD bounds under an arbitrary objective probability measure P . We conclude this section by giving an alternative interpretation of the BGD bounds that exploits this fact.

Proposition 6 *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*

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

The corresponding equality for the upper bound is also satisfied.

In other words, the BGD bounds when the objective measure is P are equal to the 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 the BGD bounds for the case of European-style option prices in the context of a standard stochastic volatility model.¹¹ 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.¹² The state variable is the changing variance of the stock return, which follows a square-root mean-reverting process with long-run

¹¹Cochrane and Saá-Requejo (1999) present an application of the standard GD methodology to a model with stochastic volatility. To compute the GD bounds on option prices, they solve numerically the corresponding non-linear PDE.

¹²We specify the drift for dS_t/S_t in the same way as in, for example, Pan (2002) and Benzoni (2000).

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)). The above specification implies that $\text{Corr}_t(dS_t/S_t, dV_t)$ is equal to ρ .

The stock 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 V_t is nontradable, it is impossible to dynamically hedge the exposure to the shocks dW_t^v .

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

$$\frac{d\Lambda_t}{\Lambda_t} = -r dt - h_t^s dW_t^s - h_t^v dW_t^v.$$

Under Q , the price and variance processes are:

$$\frac{dS_t}{S_t} = \left(r + sV_t - h_t^s \sqrt{V_t} \right) dt + \sqrt{V_t} d\tilde{W}_t^s, \quad (21)$$

$$dV_t = (\alpha - \beta V_t - \lambda_t) dt + \sigma \sqrt{V_t} \left(\rho d\tilde{W}_t^s + \sqrt{1 - \rho^2} d\tilde{W}_t^v \right), \quad (22)$$

where \tilde{W}_t^s and \tilde{W}_t^v denote Brownian motions under the measure Q and

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

By Lemma 1, the SDF Λ_t is admissible if its associated h_t satisfies

$$h_t^s = \lambda_t^s = s \sqrt{V_t}. \quad (23)$$

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

$$\lambda_t = \lambda V_t \quad (24)$$

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 (23), the specification in (24) obtains if

$$h_t^v = v \sqrt{V_t},$$

where v is constant. Then,

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

Using (20) we write the problem for the lower bound of an European-style call option with strike K and maturity T as

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

where H^v is the set of processes $h_t = (s\sqrt{V_t}, h_t^v)$ for which the volatility constraint is satisfied. Our benchmark is assumed to be of the form

$$h^* = h_1^* \equiv \left(s\sqrt{V_t}, v^* \sqrt{V_t} \right),$$

for some constant v^* , which will be chosen below. This benchmark is admissible and it satisfies (24), so that the Heston's formula applies. For the special case of the standard 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. It is easy to verify, however, that their approach also works when the ceiling is chosen as a general positive adapted process. This fact is important for our analysis because we specify the ceiling process as

$$A_t = \bar{A} \sqrt{V_t}, \tag{25}$$

where \bar{A} is a positive constant.

The motivation for (25) is the following. In Heston (1993) model, the stock's instantaneous Sharpe ratio is proportional to $\sqrt{V_t}$ and a constant ceiling is not appropriate. This is because, for any constant \bar{A} , the bound on the conditional Sharpe ratio would be too loose and useless in states where V_t is low and the bound would be too tight in states where V_t is high. In fact, for the GD bounds, no feasible discount factor exists for *any* constant ceiling \bar{A} .

The specification in (25) has a 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 analytical expression for the BGD bounds 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 5. (For completeness, the Heston formula is provided in the Appendix.).

Proposition 7 *The lower and upper bounds for the European call can be written as*

$$\underline{C}_0(h^*) = C_0^H(\underline{\lambda}), \quad \bar{C}_0(h^*) = C_0^H(\bar{\lambda}),$$

with

$$\underline{\lambda} = \left(s\rho + \underline{v}\sqrt{1 - \rho^2} \right) \sigma, \quad \bar{\lambda} = \left(s\rho + \bar{v}\sqrt{1 - \rho^2} \right) \sigma.$$

For the special case of the BGD bounds with the process $h^* = h_1^* \equiv (s\sqrt{V_t}, v^* \sqrt{V_t})$, the constants \underline{v} and \bar{v} are given by

$$\underline{v} = v^* + \bar{A}, \quad \bar{v} = v^* - \bar{A}.$$

For the special case of the GD bounds with the process $h^* = h_2^* \equiv (0, 0)$, the constants \underline{v} and \bar{v} are given by

$$\underline{v} = \sqrt{\bar{A}^2 - s^2}, \quad \bar{v} = -\sqrt{\bar{A}^2 - s^2}.$$

To illustrate the practical relevance of our approach consider now a trader who needs to price a derivative security on the stock. The trader can compute the call price C_0 using the Heston formula, but he is concerned that the unobservable variance risk premium in (24) might be misspecified. For example, the actual 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 might not believe that the Heston model is completely correct, its discount factor is nevertheless a useful benchmark, because it gives a closed-form solution and it is certainly better than the GD benchmark h_2^* .

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

$$\lambda_l V_t \leq \sigma \left(\rho \lambda_t^s + \sqrt{1 - \rho^2} h_t^v \right) \sqrt{V_t} \leq \lambda_h V_t, \quad (26)$$

where λ_l and λ_h are constants. The Heston specification obtains by setting $\lambda_l = \lambda_h = \lambda$. When $\lambda_l < \lambda_h$, the condition in (26) defines the whole set of “plausible” candidate processes h , which produce a range of candidate call prices. Naturally, one wants to have the price bounds as tight as possible, subject to the constraint that the bounds should contain all candidate prices corresponding to the specification in (26). It is easy to verify that the optimal choice for \bar{A} is

$$\bar{A} = \max(|\lambda_l - v^*|, |\lambda_h - v^*|)$$

for h_1^* and

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

for h_2^* .

Suppose now that the parameter values of the Heston model are set as in the following table:

$\alpha = 0.097$	$\beta = 7.1$	$\sigma = 0.32$
$s = 8.6$	$v = -22.6$	$\rho = -0.53$

This choice of the parameters 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. In addition to the above benchmark case, we also consider two additional cases, where all the parameters are kept the same except for ρ . This variation allows us to investigate the effect of the correlation on the price bounds. The three cases for the correlations are:

- (1) $\rho = -0.10$;
- (2) $\rho = -0.53$ (benchmark);
- (3) $\rho = -0.90$;

Other parameters are set as $r = 0.05$, $S_0 = 100$, $V_0 = \alpha/\beta = 0.0137$, and $T = 0.25$. When constructing the price bounds, we assume that the true variance risk premium is bounded as in (26) with

$$\begin{aligned}\lambda_l &= \sigma(\rho s + \sqrt{1 - \rho^2} v_l), \\ \lambda_u &= \sigma(\rho s + \sqrt{1 - \rho^2} v_u),\end{aligned}$$

where $v_l = v - 0.5\Delta$ and $v_h = v + 0.5\Delta$. Intuitively, of the all parameters of the Heston model, the variance risk premium v is the hardest to estimate accurately because the variance is not traded. In contrast, the other parameters, including the stock risk premium s , can be easily estimated from the observable price process S_t . Therefore, the trader allows for the possibility that the unobservable variance risk premium might be misspecified.

We use Proposition 7 to analytically compute the bounds on call options when $\Delta = 0.0, 2.0$ and 4.0 . We compute the BGD bounds for the process $h^* = h_1^*$, where we set $v^* = -22.6$. The results are presented in Table 1 and in Figures 1-3. In all figures, the left panels correspond to the BGD bounds while the right panels correspond to the GD bounds.

Figure 1 presents the lower bound \underline{C}_0 and the upper bound \overline{C}_0 for a range of strikes K . For clarity, this figure shows the *difference* between the bounds and the Black-Scholes price C_0^{BS} , for which the volatility is set to $\sqrt{V_0}$. The bounds are shown for different Δ and the three cases of the correlation ρ .

Figure 2 plots the size of the bounds $\overline{C}_0 - \underline{C}_0$. As expected, the size of the bounds is the largest for near-the-money strikes. Finally, Figure 3 plots implied volatilities for the lower and upper bounds by inverting the Black-Scholes formula.

It is clear that the bounds for $h^* = h_1^*$ are considerably tighter than for the *risk-neutral* benchmark $h^* = h_2^*$. The difference between the two cases is particularly pronounced when Δ approaches zero, when the size of the bounds $\overline{C}_0 - \underline{C}_0$ is zero for h_1^* and it is strictly positive for h_2^* . (For example, the size of the GD bounds is about 0.68 for near-the-money call when $\rho = -0.53$ and $\Delta = 0$.)

The intuition for this result 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

determined by the ceiling \bar{A} . The ceiling must be set large enough so that the neighborhood contains all those candidate SDF's for which the variance risk premium function is bounded as in (26). 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

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

for some adapted process h_t^v , then another SDF, 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 the GD bounds must include all SDF's for which

$$(h_t^v)^2 \leq (\bar{A}^2 - s^2) V_t.$$

Many of the included SDF's may be economically implausible. (For example, SDF's that assign positive risk premium for the shocks dW_t^v are ruled out by extensive empirical evidence.) Consequently, the GD bounds are very wide, even when there is no uncertainty about the variance 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. As a result, the corresponding BGD bounds are narrow.

When comparing the results for the different correlations ρ , one observation is worth noting. As the correlation increases in absolute value from -0.1 to -0.9, the size of the price bounds shrinks. This is true for both the BGD and GD bounds. Intuitively, the larger the $|\rho|$, the smaller the relative contribution of unspanned shocks dW_t^v . Effectively, the market is less "incomplete" and the price bounds are tighter. In the extreme case of perfect correlation $|\rho| = 1$, the market is complete and the size of the bounds is zero for both the BGD and GD cases.

5 Conclusions

The literature on asset price bounds has matured at a quick pace since the seminal contributions of Cochrane and Saá-Requejo (2000) and Bernardo and Ledoit (2000). However, this prolific research has been so far mainly theoretical and it lacks a satisfactory line of companion empirical work. Our contribution is still of theoretical nature, 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*.

In this paper, we introduce a new distance, $\sigma^*(m/m^*)/R^f$, and find it useful for deriving asset price bounds. The same distance, however, could potentially be useful in other applications. In particular, Hansen and Jagannathan (1997) propose to quantify the degree of model misspecification by the distance $\sigma(m - m^*)$. The new distance could present an interesting alternative to the HJ measure. The new distance emphasizes the

relative discrepancy as opposed to the *absolute* discrepancy, offers an attractive economic interpretation, and is easy to work with in continuous time.

Empirical evidence confirms the presence of jumps in stock prices and possibly in their volatilities (see Andersen, Benzoni and Lund (2001), Eraker, Johannes and Polson (2001), among others). Therefore, inclusion of jumps in modeling price bounds might be an interesting theoretical development. Along these lines, the recent paper by Björk and Slinko (2006) extends the good-deal approach in continuous time by allowing price processes with jumps.

Appendix

Proof of Proposition 2

The equality is proved in Hansen and Jagannathan (1991). To prove the inequality, note that if $x \in Z_m$, then

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

which may be rewritten as

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

The fact that $d(m, m^*) = \sigma(m - m^*)$ now implies that

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

■

Proof of Proposition 3

From the Cauchy-Schwartz inequality, it follows that

$$\left| E^* \left(\left(\frac{m}{m^*} - 1 \right) x \right) \right| \leq \left[E^* \left(\frac{m}{m^*} - 1 \right)^2 \right]^{\frac{1}{2}} [E^*(x^2)]^{\frac{1}{2}} = \frac{1}{R^f} d(m, m^*) [E^*(x^2)]^{\frac{1}{2}}.$$

On the other hand,

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

This implies that

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

To prove that the bound is sharp, we note that this fact is obvious when $m = m^*$, in which case $d(m, m^*) = 0$. When $m \neq m^*$, the bound is achieved for the normalized payoff

$$\frac{x}{[E^*(x^2)]^{\frac{1}{2}}},$$

where

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

It only remains to prove the inequality in (15). For $x \in Z_m$,

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

Therefore,

$$|E^*(x)| \leq \sigma^* \left(\frac{m}{m^*} \right) \sigma^*(x).$$

■

In order to simplify notation, in what follows we will denote by $\Psi \equiv \Lambda/\Lambda^*$ for some fixed Λ^* . It will also be convenient to collect several auxiliary facts in the following lemma.

Lemma 2 *For any SDF processes Λ and Λ^* and any adapted process U , the following relationships hold:*

(a) *For all $t \leq s \leq T$,*

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

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

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

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

Proof. (a): Note that

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

where the second equality follows from (16).

(b): By definition Λ prices U if

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

and from part (a) it follows that

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

(c): From part (a),

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

because the process Λ/B is a P -martingale. ■

Proof of Proposition 4

By Lemma 2 (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 . Itô's Lemma implies that 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 2 (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 6

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

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

which together with Lemma 2 (b) gives that

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

Hence,

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

where the second equality follows from Lemma 2 Part (a). ■

The Heston formula

Given the risk-neutral model (21)-(22) and the market price of variance risk as in (24), Heston (1993) obtains the following closed-form formula for the European call option:

$$C^H(K, T; S_t, V_t, t, \theta, \lambda) = S_t P_1 + e^{-r\tau} K P_2,$$

where for $j = 1, 2$

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

and

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

References

- Andersen, T.G, L. Benzoni, and J. Lund (2001), "An Empirical Investigation of Continuous-Time Equity Returns Models," *Journal of Finance* 57, 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* 108, 144-172.
- Björk, T, and I. Slinko (2006), "Towards a General Theory of Good Deal Bounds," *Review of Finance* 10, 221-260.
- Cerny, A. (2003), "Generalized Sharpe Ratios and Asset Prices in Incomplete Markets," *European Finance Review* 7, 191-233.
- Cochrane, J. H., and J. Saá-Requejo (1999), "Good-Deal Option Price Bounds with Stochastic Volatility and Stochastic Interest Rate," working paper, University of Chicago.
- Cochrane, J. H., and J. Saá-Requejo (2000), "Beyond Arbitrage: Good-Deal Asset Price Bounds in Incomplete Markets," *Journal of Political Economy* 108, 79-119.
- Eraker B., M. S. Johannes, and N.G. Polson (2003), "The Impact of Jumps on Volatility and Returns," *Journal of Finance* 58, 1269-1300.
- Hansen, L.P., and R. Jagannathan, (1991), "Implications of Security Market Data for Models of Dynamic Economies," *Journal of Political Economy* 99, 225-262.
- Hansen L.P., and R. Jagannathan (1997), "Assessing Specification Errors in Stochastic Discount Factor Models," *Journal of Finance* 52, 557-590.
- Henderson, V., D. Hobson, S. Howison, and T. Kluge (2005), "A Comparison of Option Prices Under Different Pricing Measures in a Stochastic Volatility Model with Correlation," *Review of Derivatives Research* 8, 5-25.
- Jaschke S., and U. Küchler (2001), "Coherent Risk Measures and Good-Deal Bounds," *Finance and Stochastics* 5, 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.

Table 1: Price Bounds for Call Options

K	C	\underline{C}^{BGD}		\overline{C}^{BGD}		\underline{C}^{GD}			\overline{C}^{GD}		
		$\Delta=2$	$\Delta=4$	$\Delta=2$	$\Delta=4$	$\Delta=0$	$\Delta=2$	$\Delta=4$	$\Delta=0$	$\Delta=2$	$\Delta=4$
$\rho = -0.10$											
85	16.142	16.135	16.128	16.150	16.159	16.067	16.066	16.065	16.142	16.150	16.159
90	11.439	11.420	11.402	11.459	11.481	11.198	11.193	11.189	11.439	11.459	11.481
95	7.181	7.144	7.108	7.220	7.260	6.621	6.606	6.592	7.181	7.220	7.260
100	3.821	3.770	3.720	3.874	3.929	2.977	2.951	2.926	3.821	3.874	3.929
105	1.702	1.654	1.608	1.751	1.802	0.942	0.920	0.898	1.702	1.751	1.802
110	0.661	0.630	0.600	0.695	0.729	0.223	0.212	0.203	0.661	0.695	0.729
115	0.238	0.221	0.206	0.256	0.275	0.045	0.042	0.039	0.238	0.256	0.275
$\rho = -0.53$											
85	16.206	16.197	16.188	16.215	16.226	16.101	16.098	16.096	16.206	16.215	16.226
90	11.559	11.539	11.521	11.579	11.600	11.305	11.298	11.291	11.559	11.579	11.600
95	7.326	7.293	7.261	7.361	7.396	6.842	6.826	6.811	7.326	7.361	7.396
100	3.874	3.829	3.786	3.919	3.966	3.197	3.174	3.152	3.874	3.919	3.966
105	1.578	1.537	1.497	1.620	1.664	0.978	0.958	0.940	1.578	1.620	1.664
110	0.476	0.452	0.429	0.501	0.528	0.174	0.167	0.160	0.476	0.501	0.528
115	0.114	0.104	0.096	0.124	0.135	0.020	0.019	0.017	0.114	0.124	0.135
$\rho = -0.90$											
85	16.194	16.190	16.186	16.198	16.202	16.160	16.159	16.157	16.194	16.198	16.202
90	11.514	11.506	11.498	11.522	11.531	11.440	11.438	11.435	11.514	11.522	11.531
95	7.201	7.187	7.173	7.215	7.230	7.071	7.065	7.060	7.201	7.215	7.230
100	3.604	3.585	3.566	3.623	3.643	3.426	3.419	3.412	3.604	3.623	3.643
105	1.170	1.153	1.136	1.187	1.205	1.015	1.009	1.003	1.170	1.187	1.205
110	0.145	0.139	0.133	0.151	0.158	0.095	0.093	0.091	0.145	0.151	0.158
115	0.003	0.003	0.002	0.003	0.004	0.001	0.001	0.001	0.003	0.003	0.004

Notes: This table reports BGD and GD bounds for call options with different strikes K . The lower and upper bounds are computed using parameter values in Section 4.2, when $\Delta = v_h - v_l$ is equal to 0.0, 2.0, and 4.0. The second column is the theoretical call price C_0 . When $\Delta = 0.0$, the lower and upper BGD bounds coincide with C_0 and, thus, are not reported. The results are presented for the three cases of the correlation ρ .

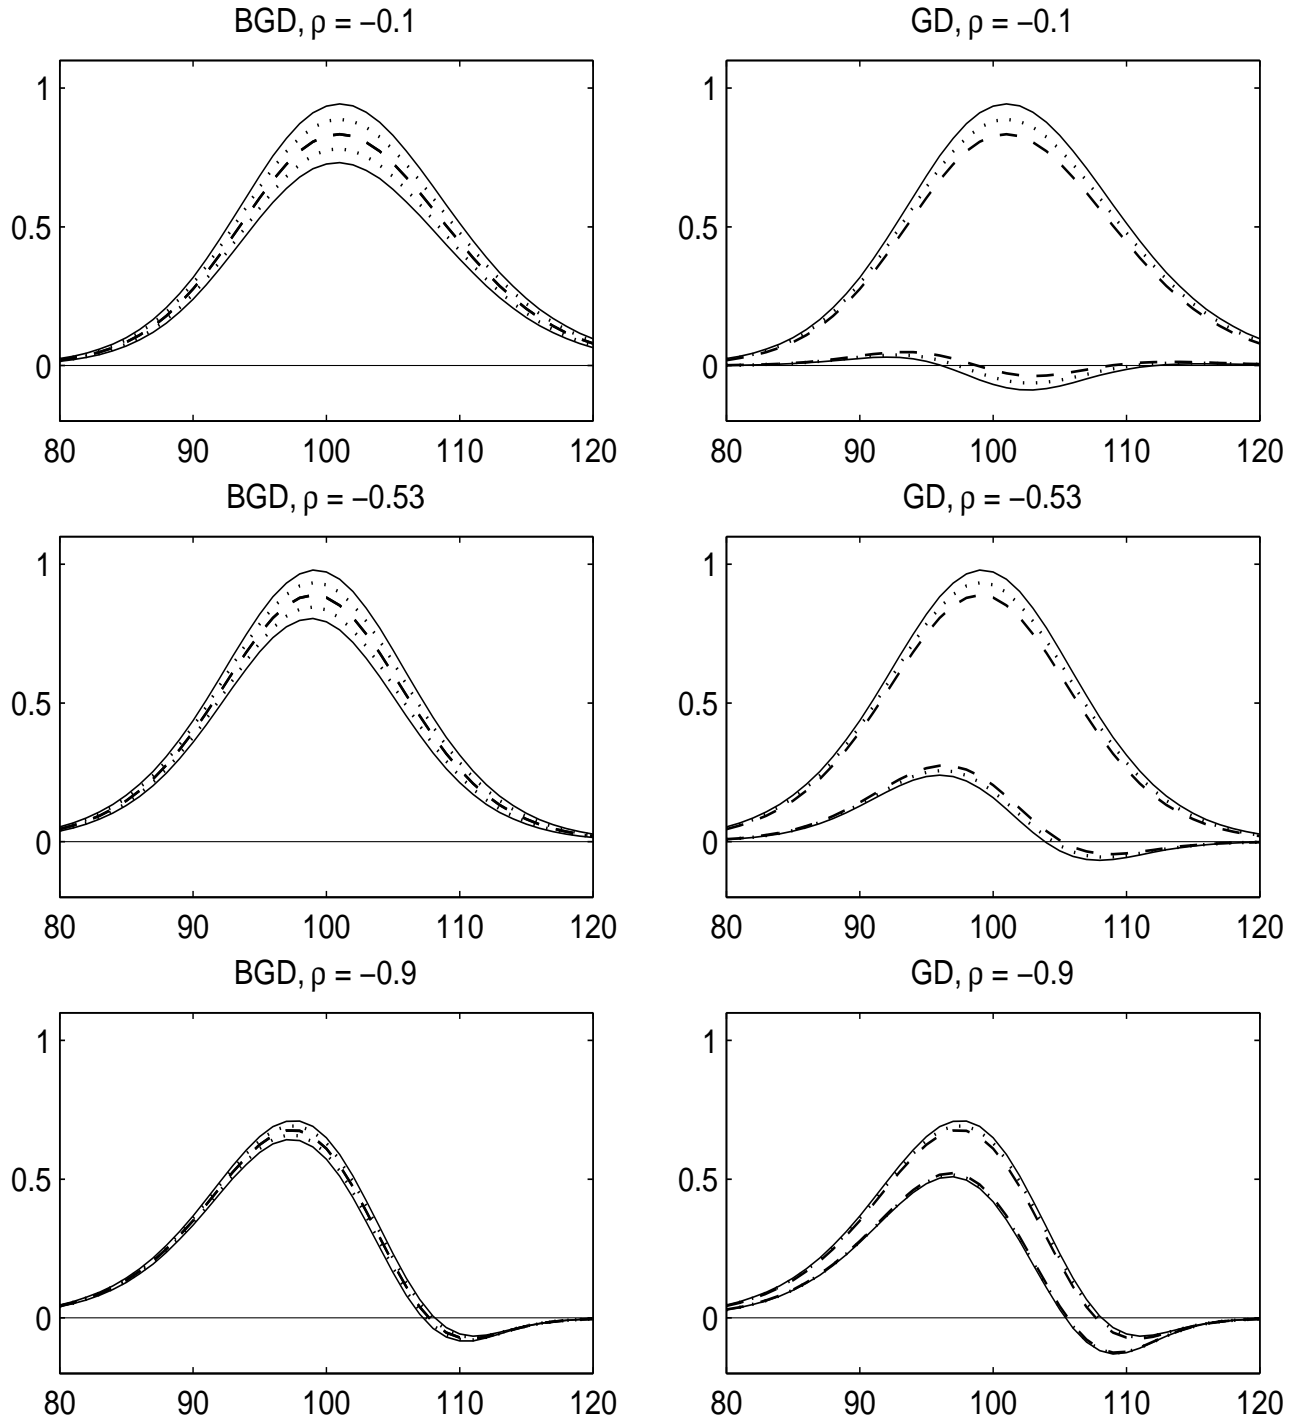


Figure 1: The difference between the option price bounds and the Black-Scholes price versus strike K , for h_1^* (the left panels) and h_2^* (the right panels). The differences $\underline{C}_0 - C_0^{BS}$ and $\overline{C}_0 - C_0^{BS}$ are computed using parameter values in Section 4.2, when $v_h - v_l$ is equal to 0.0 (the dashed lines), 2.0 (the dotted lines), and 4.0 (the solid lines). The results are presented for the three cases of the correlation ρ .

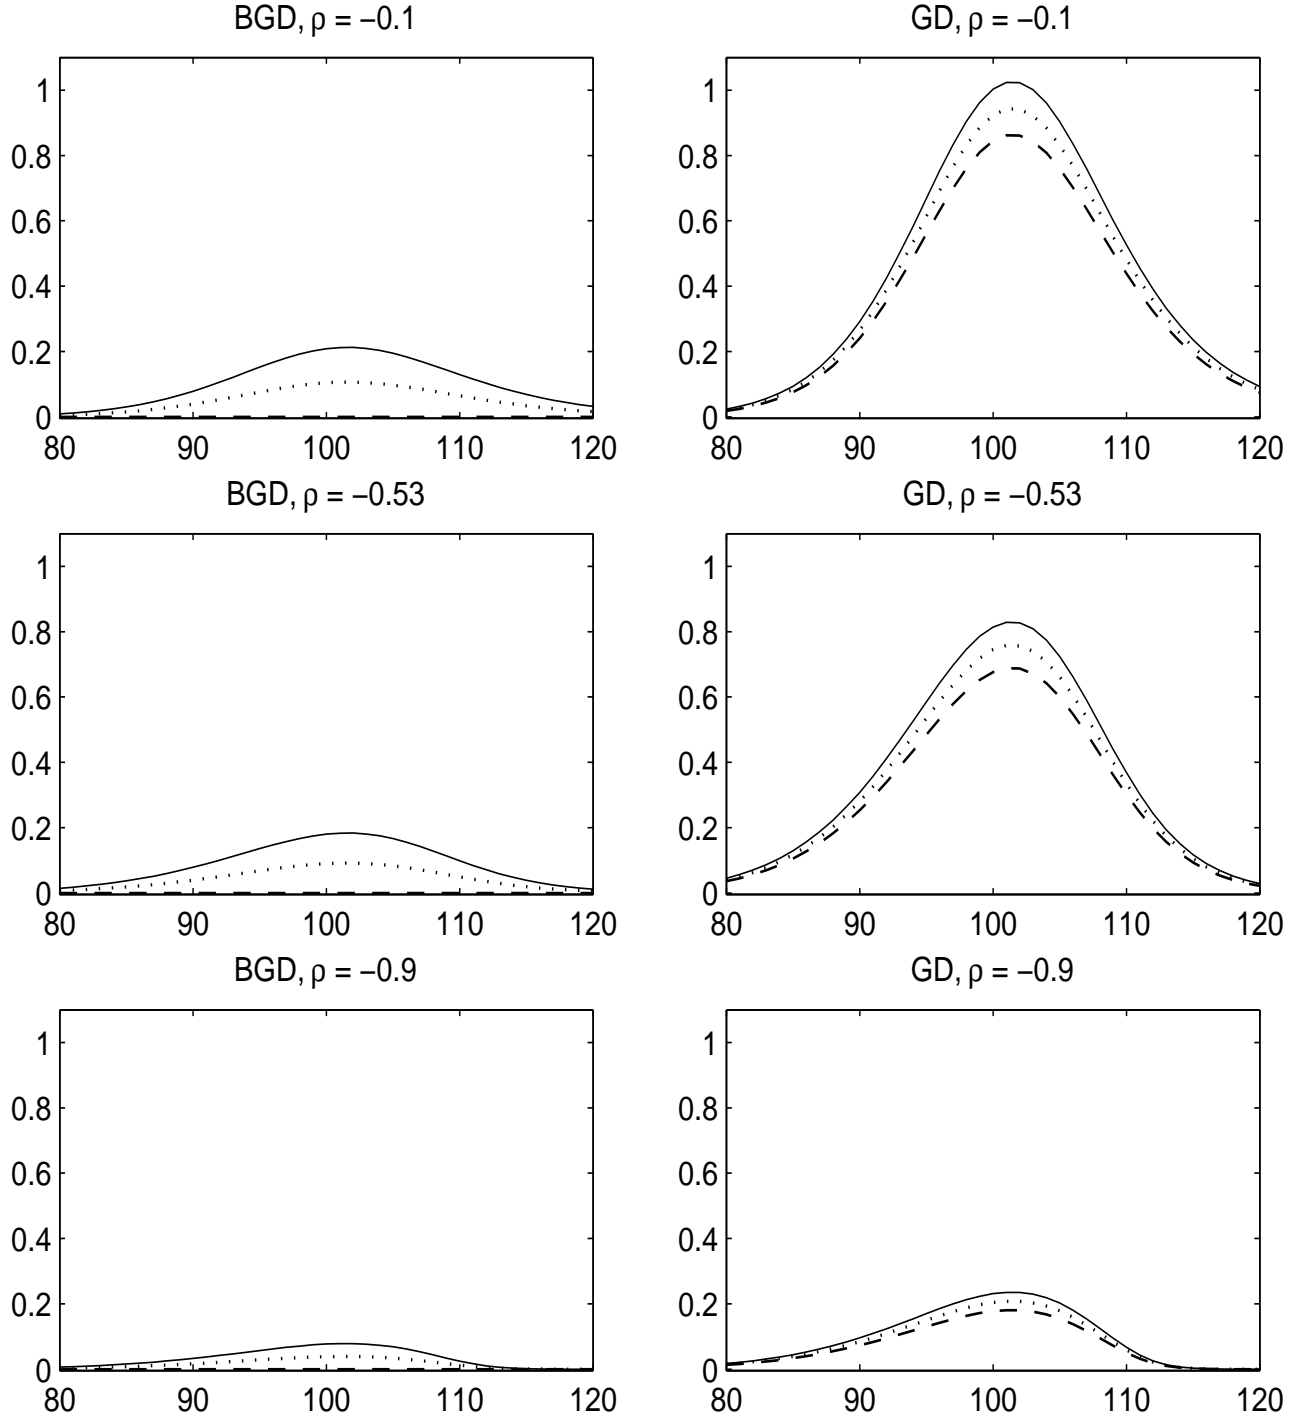


Figure 2: The size of the bounds $\overline{C}_0 - \underline{C}_0$ versus strike K , for h_1^* (the left panels) and h_2^* (the right panels). The bounds are computed using parameter values in Section 4.2, when $v_h - v_l$ is equal to 0.0 (the dashed lines), 2.0 (the dotted lines), and 4.0 (the solid lines). The results are presented for the three cases of the correlation ρ .




Figure 3: The Black-Scholes implied volatility versus strike K , for h_1^* (the left panels) and h_2^* (the right panels). 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 $v_h - v_l$ is equal to 0.0 (the dashed lines), 2.0 (the dotted lines), and 4.0 (the solid lines). The results are presented for the three cases of the correlation ρ .