Hannan and Blackwell meet Black and Scholes: Approachability and Robust Option Pricing

October, 2005

This Revision: 4/20/09

ABSTRACT.

We study the link between the game theoretic notion of approachability or "regret minimization" and robust option pricing. This concept was introduced by Hannan and Blackwell and is the basis behind a growing literature in game theory and microeconomics. We demonstrate how trading strategies, based on approachability, that minimize regret also imply robust, and empirically relevant, upper bounds for the prices of European call options. We then argue that the gradient strategy proposed by Hannan-Blackwell is path dependent and therefore suboptimal when finite horizons are considered. Based on path independence we solve for the optimal strategy and bound.

1. Introduction

There is a growing literature in game theory that is based on "approachability" or "regret minimization" for games under uncertainty.[1] Regret is defined as the difference between the outcome of a strategy and that of the ex-post optimal strategy (within a given class). This literature is based on earlier work by Hannan (1957) and Blackwell (1956) who studied robust dynamic optimization, and is the basis for the more recent work on calibration and the dynamic foundations of correlated equilibria; see Hart (2005), Foster, Levine, and Vohra (1999), and Fudenberg and Levine (1998) for excellent surveys. In this paper we consider a financial application of these ideas and demonstrate a link to the robust pricing of financial assets. In particular we focus on standard European-style options, which we can think of as contracts that allow investors to minimize their regret when choosing an investment portfolio.

Using the link between regret minimization and option pricing, we then derive robust pricing bounds for financial options. The classic, structural approach to option pricing developed by Black and Scholes (1973) and Merton (1973), posits a specific stock price process (geometric Brownian motion), and then shows that the payoff of an option can be replicated using a dynamic trading strategy for the stock and a risk-free bond. No arbitrage then implies that price of the option must equal the cost of this trading strategy. But because empirical stock prices do not follow the process assumed by Black-Scholes-Merton, their argument is not a true arbitrage: the replication is perfect only for a very restricted set of price paths. While our results are weaker -- we provide bounds, rather than exact prices -- they are robust in that we do not assume a specific price process.

In sum, the goal of this paper is two-fold. First, we develop a finance-based interpretation for the notion of regret minimization by showing the link to robust (distribution-free) bounds for the value of financial options. This interpretation is interesting as monetary payoffs provide a tangible way to measure the performance of regret-minimizing strategies (compared with standard results based on the asymptotic average performance of these strategies). Second, we look for the optimal such bounds. These bounds provide some measure of the performance of

[1] We use the terms regret minimization and approachability interchangeably. Formally, approachability is the more general concept while regret minimization is a classic example of it.

known, heuristic regret minimizing strategies (which despite being asymptotically efficient need not be optimal).

The roots of approachability and regret minimization in game theory can be traced to Hannan (1957) and Blackwell's (1956) work on dynamic optimization when the decision maker has very little information about the environment. They considered an infinitely repeated decision problem in which in each period the agent chooses an action from some fixed finite set. Although the set of actions is fixed, the payoffs to these actions vary in a potentially non-stationary manner, so that learning is not possible. They show that in the limit, there is a dynamic strategy that guarantees the agent an average payoff that is at least as high as that from the ex-post optimal static strategy in which the same action is taken repeatedly. Thus, in terms of the long run average payoff, the agent suffers no regret with respect to any static strategy.

More generally, approachability and regret minimization can be viewed as an alternative objective when optimizing under uncertainty to the traditional approach in economics that considers an absolute objective function for the decision-maker (e.g. Gilboa and Schmeidler (1989)).[2] A regret-minimizing decision-maker is not concerned about the absolute performance of her strategy, but rather how well it performs compared to a defined set of alternative strategies. Of course, our purpose is not to evaluate whether regret minimization is an appropriate objective, but rather to consider the properties of the strategies that achieve it in a financial context.

In an uncertain financial market, we can define regret as the ratio between the investor's wealth and the wealth he could have obtained had he followed an alternative investment strategy. By comparing the investor's payoff to that which could be attained from a buy and hold investment of a stock or a bond, we can interpret regret as the difference in payoff between a dynamic trading strategy and a call option, allowing us to link regret minimization to no arbitrage upper bounds for option prices. To see why, suppose we have a strategy that can guarantee a return no worse than 80% of the return on investment in a stock or in a bond (i.e., has a "maximum regret"

[2] Several papers in economics consider regret minimization as an objective. Bergmann and Schlag (2005) examine a monopolist who minimizes regret. Boze, Ozdorzen and Pape (2004) examine optimal auctions in this context. Milnor (1954) and later Hayashi (2005) provide axiomatic foundations for such preferences. A general computer science setting where the performance of online algorithms is compared to the optimal one is competitive analysis (see Borodin and El-Yaniv (1998)).

of 20%). Let the interest rate on the bond be zero for simplicity. Then by borrowing \$100 and investing \$125 in this strategy, we can attain a payoff that is no worse than the payoff of an at the money call option on \$100 worth of stock.[3] Thus, the value of the call option cannot exceed the initial investment of \$25. Thus a regret guarantee of no more than 20% is equivalent to an upper bound of \$25 for an at-the-money call option.

In Section 2 of the paper, we first review Hannan-Blackwell's approachability. While Hannan-Blackwell focus on limiting results (similar to the traditional work on regret minimization), we consider minimizing regret over a finite horizon. We also construct a simple but important generalization of their original "gradient" strategy. This generalization will allow us to improve the finite-horizon performance of the Hannan-Blackwell strategy, and will also prove extremely useful when applying the results to options with different strike prices.

In Section 3 we apply the results to the pricing of call options. To do so, we need to adjust for the fact that in an investment context, payoffs are multiplicative and not additive. We then apply the Hannan-Blackwell gradient strategy, with mixed strategies replaced by portfolio weights, to derive robust upper bounds for the prices of call options. The bounds we derive are based on no arbitrage and are robust in that they do not depend on specific distributional assumptions for the stock price path. For example, we can allow both jumps in the stock price process and trading halts. The bounds we derive depend solely upon the total quadratic variation of the stock price path. Because quadratic variation is an equivalent measure to volatility in the Black-Scholes-Merton (BSM) framework, we can directly compare our price bounds with those of the Black-Scholes formula.[4] Unlike Black-Scholes, however, our bounds do not require the stock price paths to be continuous.

The trading strategies we develop are simple trend-following or momentum strategies, and unlike standard option hedging strategies, these strategies are history dependent. While they are asymptotically optimal, they are not necessarily optimal for a finite investment horizon and a given quadratic variation. In Section 3 we also show that we can improve upon the Hannan-Blackwell strategy using our generalized gradient approach. In Section 4, we further argue that the optimal regret minimizing strategy must be path independent; and hence the Hannan-

[3] The future payoff of the option is the max{\$100,S_1}-100 where S_1 is the future value of \$100 invested in the stock.
[4] Our results in this regard are related to work by Cover (1991, 1996) on the "universal portfolio," a dynamic trading strategy designed to perform well compared to any alternative fixed-weight portfolio.

Blackwell strategy is suboptimal. This property suggests using a recursive application of our generalized gradient method. We use this idea to solve numerically for the optimal bound using dynamic programming and derive the optimal robust trading strategy. This strategy is the lowest cost strategy with a payoff that exceeds the option payoff for any stock price path with a quadratic variation below a given bound. This strategy is also the the optimal strategy for minimizing regret in our setting.

In Section 4.5, we compare our optimal price bounds to the BSM model. Again. our optimal bounds necessarily exceed the BSM price – thus, we can interpret the bounds as the BSM price corresponding to a higher implied volatility. We show that our bounds are sufficiently "tight" to be empirically relevant. Indeed, the pattern of implied volatility determined by our bound resembles the volatility "smile" that has been documented empirically in options markets. We also compare our trading strategy to the delta hedging strategy of BSM. We show that it is similar in nature but that the stock position is more sensitive overall to movements in the underlying stock price. This strategy insures against jumps in the stock price that are not considered by BSM, which we demonstrate makes it much more robust than the BSM hedging strategy when the assumption of continuous price paths is dropped.

1.1. Literature Review

The classic work of Foster and Vohra (1998) renewed the interest of game theorists in the importance of the approachability theorem.[5] While much of this literature is focused on abstract settings a natural question is whether these results can be applied to financial markets. Our paper tries to build such a link by examining the implication to the classic Black-Scholes model.

While the Black-Scholes formula is one of the most useful formulas developed in economics, in recent years extensive empirical research has identified several anomalies in the data. In general the formula seems to generate prices for stock index options that are too low. Said another way, the implied volatility of the stock index computed based on the Black-Scholes formula is significantly higher on average than the ex-post realized volatility (see Bollen and Whaley (2004)). In addition, this effect is more pronounced for call options whose strike price is low.

[5] Recent contributions include Dekel, E. and Y., Feinberg (2006), Al-Najjar, and Weinstein. (2008) and Olszewski and Sandroni (2008).

This effect is often referred to as the volatility "smirk" or "smile." As a response to these findings, there has been an active research trying to modify the Black and Scholes formula to account for these discrepancies.[6] These papers examine different stochastic processes for the index, with modifications that include jump processes and stochastic volatility models. The result of our study will complement this analysis by offering a new perspective. Rather than focusing on a specific formulation for the stochastic process we rely on a generic trading strategy that works with any evolution for the risky asset as long as it satisfies some bounds on quadratic variation.

As a result of both academic and practical interest there are several papers similar in spirit to our work as the goal is to provide a robust bound by relaxing the specific assumption on the price process made by Black and Scholes. For example, it has been shown that the Black-Scholes formula provides a price bound if a maximal instantaneous volatility or maximal quadratic variation is used in place of a constant volatility.[7] These results all maintain the assumption of a continuous price path and continuous trading, however. These results therefore do not hold in models with discrete trading or price jumps (which are obviously relevant in practice), unlike the approach considered here.

An alternative approach to that taken here is developed by Bernardo and Ledoit (2000) and Cochrane and Saa-Requejo (2000), who strengthen the no-arbitrage condition by using an equilibrium argument. They assume a specific stochastic process for the stock and put bounds for the risk-reward ratio that should be achievable in the market.[8]

Our research is also related to research in Computer Science and Statistics. In particular there is a literature that applies regret minimizing algorithms (called competitive algorithms in this literature) in the context of investments.[9] Most of the literature follows the seminal work by Cover (1991), who consider strategies to optimize the long-run asymptotic performance of a portfolio relative to fixed-weight strategies. Cover and Ordentlich (1998) consider a finite horizon setting, and look at strategies that minimize regret against all fixed-weight strategies

---

[6] E.g. Pan (2002), Eraker, Johannes and Polson (2003), and Eraker (2004).

[7] See Shreve, El Karoui, and Jeanblanc-Picque (1998), Grundy and Wiener (1999), Mykland (2000), and Shafer and Vovk (2001).

[8] See also Lo (1987) and Bertsimas, Kogan and Lo (2001) who examine strategies that "almost" replicate the payoff of an option given a stochastic process for the underline stock.

[9] The competitive ratio of an algorithm is the maximum, over all realizations, of the ratio of the performance of the best ex-post algorithm to that of the given algorithm (see, e.g., Sleator and Tarjan (1985)).

over this horizon. They then interpret the result in terms of the price of the exotic derivative that pays ex-post the best constant rebalanced portfolio. While their results are mathematically very elegant, they are not useful for standard options such as call options that are traded in the market, and the bounds are much too weak – even weaker than standard, static bounds – to be of practical or theoretical importance for standard options.

2. Hannan and Blackwell's Gradient Strategy

We first examine the original strategy proposed by Blackwell (1956) for adversarial games. While Blackwell focuses on the infinite horizon properties of this strategy, we provide a characterization of the strategy's finite horizon performance. We also construct a simple generalization of the strategy considered by Blackwell that will be useful in our analysis.

2.1. The Multi-Action Game

Blackwell (1956) considers a setting in which an agent repeatedly chooses a single action among I possible alternatives. The realized payoffs of the actions at time n are given by $\pi_n \in \mathbb{R}^{1\times I}$, with $\pi_{j,i}$ the payoff of alternative i at time n. At this point we make no assumptions regarding the underlying distribution of these payoffs. In particular the payoffs need not be stationary. The agent's action choice at time n is given by the random variable $\xi_n \in \{1..I\}$. The agent's payoff at time n is π_{n,ξ_n}, and the agent's aggregate payoff up to time N is $\sum_{n=1}^{N} \pi_{n,\xi_n}$.

At each date, the agent observes the payoff of *every* action choice. Given the history up to date $n-1$ of prior payoffs and action choices, the agent's strategy at date n specifies a distribution for the action choice ξ_n, which we denote by the vector $\rho_n \in \mathbb{R}^{I\times 1}$, where $\rho_{n,i} = \Pr(\xi_n = i)$.

Given the few assumptions we have made, the agent does not seek to maximize absolute performance. Instead, Blackwell considered a relative benchmark: How does the agent's aggregate payoff compare to the payoff of a simple static strategy of choosing the same alternative each period? Specifically, we compute the regret with the best static strategy ex-post as:

$$\max_i \left\{ \sum_{n=1}^{N} \pi_{n,i} \right\} - \sum_{n=1}^{N} \pi_{n,\xi_n} \tag{1}$$

2.2. A Gradient Strategy

In his seminal paper Blackwell (1956) constructs a simple randomized strategy that minimizes the asymptotic regret in (1). Let $A_{n,i}$ denote the agent's relative performance compared to alternative i after n periods, $A_{n,i} \equiv \sum_{n'=1}^{n} \pi_{n',i} - \pi_{n',\xi_{n'}}$, let $A_{n,i}^{+} = \max\{A_{n,i}, 0\}$ and let A_n denote the vector of relative performance after n periods where we suppress i. The randomized strategy is given by:

$$\Pr(\xi_n = i) = \rho_{n,i} = \frac{A_{n-1,i}^{+}}{\sum_{i'} A_{n-1,i'}^{+}}, \tag{2}$$

with $\Pr(\xi_n = i) = 1/I$ for all i if $\sum_i A_{n-1,i}^{+} = 0$.

This strategy has a nice geometric interpretation. The vector A_n^{+} represents the loss of the agent at time j relative to the different alternatives. The agent starts at the origin and tries to slow his drift away from the "no-regret region," which is the negative orthant shown in **Figure 1**. He does so by choosing a strategy that moves him orthogonally relative to the closest point of the no regret region. Specifically, define the vector $\Delta A_n \equiv A_n - A_{n-1}$ to be the change in the aggregate loss on date j. Then given the strategy defined in (2), the expected change ΔA_n is conditionally orthogonal to A_{n-1}^{+}:[10]

$$E\left[\Delta A_n \cdot A_{n-1}^{+} \middle| \pi_n, A_{n-1}\right] = \pi_n \cdot A_{n-1}^{+} - \pi_n \cdot \rho_n \sum_i A_{n-1,i}^{+} = 0 \tag{3}$$

[10] We use the conditional expectation loosely here, as we have not defined a probability distribution for the payoff vectors. The expectation in (3) is over the agent's action choice.

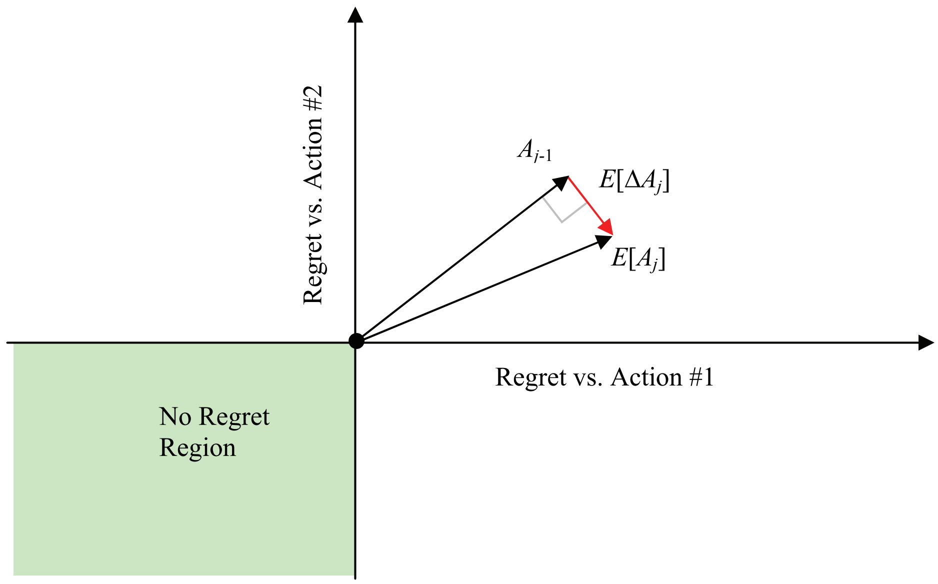

Figure 1 Updating the Loss Vector using the Blackwell Gradient Strategy

Figure 1 illustrates the updating when there are two action choices. Because of the orthogonal updating of the loss vector, the squared distance from the no regret region accumulates in an additive fashion. This fact leads to the following key result:

PROPOSITION 1. For any payoff history $\pi^n = (\pi_1, ..., \pi_n)$, the gradient strategy described above satisfies:

$$E\left[\max_i A^+_{n,i} \middle| \pi^n\right] \le \sqrt{\sum_{n'=1}^{n} E\left[\left\|\Delta A_{n'}\right\|^2 \middle| \pi^n\right]} \qquad (4)$$

Proof: Note that

$$\left\|A^+_n\right\|^2 \le \left\|A^+_{n-1} + \Delta A_n\right\|^2 = \left\|A^+_{n-1}\right\|^2 + \left\|\Delta A_n\right\|^2 + 2\Delta A_n \cdot A^+_{n-1}$$

Therefore, we can conclude using (3) that:

$$E\left[\left\|A^+_n\right\|^2 \middle| \pi^n\right] \le \sum_{n'=1}^{n} E\left[\left\|\Delta A_{n'}\right\|^2 \middle| \pi^n\right]$$

The result then follows from the fact that $E\left(\max_i A^+_{n,i}\right) \le E\left(\left\|A^+_n\right\|\right) \le \sqrt{E\left(\left\|A^+_n\right\|^2\right)}$. ♦

For the special case of two actions, because $\|\Delta A_n\|^2 \le (\pi_{n,1} - \pi_{n,2})^2$ we have as an immediate corollary:

Corollary A. For $I = 2$, Eq. (4) is equivalent to

$$E\left[\max_i A^+_{n,i} \middle| \pi^n \right] \le \sqrt{\sum_{j=1}^{n} (\pi_{n,1} - \pi_{n,2})^2} \qquad (5)$$

Hence, $\sqrt{\sum_{j=1}^{n} (\pi_{n,1} - \pi_{n,2})^2}$ is the maximum expected loss of this strategy relative to the best static alternative ex-post. Note that this expectation is with regard to the agent's strategy choice and does not reflect any randomness in the payoff structure as it holds for <u>any</u> payoff history. Nevertheless there is uncertainty regarding the actual realization of the loss.

A Generalized Gradient strategy

Consider a simple yet useful generalization of the Blackwell gradient strategy. As we shall later see this generalization is useful for (i) improving the bound for at the money call options, and (ii) bounds for options with different strike prices

We describe it for the special case of just two alternatives $I = 2$; this is the relevant case for our application. Suppose that in the strategy we have just described, instead of starting with a loss vector $A_0 = 0$, we start at some arbitrary point $A_0 = (x, y) > 0$. Let $A_n(x, y)$ denote our location after *j* rounds. In this case, following the same logic as above, we have:

$$E\left[\max_i A^+_{N,i}(x,y) \middle| \pi^N \right] \le \sqrt{\sum_{n=1}^{N} (\pi_{n,1} - \pi_{n,2})^2 + x^2 + y^2} \qquad (6)$$

2.3. Approachability and Calibration

While in this paper we focus on finite horizons, we note that **Proposition 1** implies the famous "Approachability" result of Blackwell (1956). Specifically, if the payoffs are appropriately bounded, then the gradient strategy implies that the agent will have no asymptotic regret:

Corollary B. (Approachability) If $|\pi_{j,i} - \pi_{j,i'}|$ is uniformly bounded for all (*i, i', j*), then with the generalized gradient strategy the agent has no long run average regret:

$$\liminf_{N\to\infty}\left[\max_i\left\{\frac{1}{N}\sum_{n=1}^{N}\pi_{j,i}\right\}-\frac{1}{N}\sum_{n=1}^{N}\pi_{j,\xi_j}\right]\leq_{a.s.} 0 \qquad (7)$$

The proof for this follows the fact that Eq. (4) implies that the average regret goes to zero at a rate of $n^{-1/2}$ and using strong law of large numbers. To appreciate this remarkable result, consider the following example:

EXAMPLE (Approachability and No Regret): *Suppose an agent must repeatedly predict the outcome of a coin toss, where the coin need not be a fair one. We are interested in the success rate he can guarantee in the long run, that is in the limit as the number of coin tosses goes to infinity.*

If the agent repeatedly predicts heads or tails randomly with equal probability, then in the long run his success rate is 50% with probability one. He could do better if he knew the parameter of the coin. If the probability of heads, p, is more than 0.5, then he could always announce heads and achieve a success rate of p. Conversely if $p < 0.5$ he could achieve a success rate of $1-p$ by always predicting tails. Note that if the same coin is used over and over again the fact that p is unknown is not an issue in the long run. The agent can first focus on estimating p, as failure in the first rounds does not affect his long term success rate. Once he learns p he can guarantee a success rate of $\max\{p,1-p\}$. We can describe this in terms of regret minimization by considering two simple strategies: (i) always predict heads, and (ii) always predict tails. Thus, we can obtain zero asymptotic regret provided that the sequence is stationary; we use the same coin over and over again.

Corollary B implies that the agent can achieve a similar performance even if the coin can change over time and learning is not possible. That is, the generalized gradient strategy achieves an asymptotic success rate of $\max\{p(n),1-p(n)\}$ where $p(n)$ denotes the fraction of heads in the first n rounds.[11]

This example demonstrates the link to calibration. It implies that it may be difficult to tell whether an agent indeed knows the distribution of the coin, as the agent achieves the same performance even when he knows nothing about the coin. This surprising conclusion led to the

[11] Note that $p(n)$ need not even converge as n goes to infinity. Hence, the formal statement of the result is that for any choice of coins $\liminf_{n\to\infty}\{X\{n\}-\max(p(n),1-p(n))\}\geq 0$ almost surely, where $X(n)$ denotes the success rate in the first n tosses.

growing literature on calibration that further investigates the problem of asymptotically testing whether an "expert" is indeed knowledgeable.[12]

3. Regret-Based Bounds for Call Options

In this section we first describe a simple trading model. This model is different from the setup we described before in that it is multiplicative in nature and not additive. We show to adapt the original Blackwell strategy and construct gradient trading strategies. We then show how this can translated to an upper bound for the value of a European call option.

3.1. A Simple Financial Trading Model

Consider a discrete-time n-period model where time is denoted by $n \in \{0,1,...,N\}$. There is a risky asset (e.g., stock) whose value (price) at time n is given by S_n. We normalize the initial value to one, $S_0 = 1$, and assume that the asset does not pay any dividends. We denote by r_n the return between $n-1$ and n so that $S_n = S_{n-1}(1+r_n)$. We call $r = r_1,...,r_N$ the *price path*. In addition to the risky asset we have a risk-free asset (e.g., a bond). Unless otherwise stated, we assume that the risk-free rate is zero.

A dynamic trading strategy specifies a portfolio to hold at each date. This portfolio has initial value $G_0 = 1$. Each period, the current value of the portfolio G_n is invested in the assets, with some fraction w_n in the risky asset and $1-w_n$ in the risk-free asset. The portfolio weight w_n is specified as a function of the price path of the stock up to date n. Since we assume a zero interest rate, at time $n+1$ its value is $G_{n+1} = (w_n G_n)(1+r_{n+1}) + (1-w_n)G_n = G_n(1+w_n r_{n+1})$; its final value is G_n. Thus, a trading strategy determines a mapping $G_n(r)$ from price paths to payoffs at each date n.

In section 2.1 we introduced a regret measure in an additive framework with multiple alternatives. Here we have just two alternatives: investing in the risk free asset or the stock. Also, because the model of returns is multiplicative, it is natural to measure regret in percentage terms, so that the analog of the regret measure in (1) is:

---

[12] See e.g. Foster and Vohra (1997, 1998), Foster (1999), Dekel and Feinberg (2006), Al-Najjar and Weinstein (2008), Wojciech and Sandroni (2008).

$$\frac{\max\{1,S_N\}-G_N}{\max\{1,S_N\}}=1-\frac{G_N}{\max\{1,S_N\}}. \quad (8)$$

3.2. The Gradient Trading Strategy

To adapt the construction described in Section 2 to our setting, one needs to make a few adjustments. First, since returns are compounded, our setup is multiplicative rather than additive. Second, we seek a deterministic strategy rather than a randomized one; we use portfolio weights to replace randomization.

First, we consider two alternative benchmarks based on the two financial assets by letting the payoffs $\pi_{n,i}$ be the log returns of each asset: $\pi_{n,1}=\ln(1+r_n)$, $\pi_{n,2}=0$. Second, we construct a deterministic strategy by investing a fraction of $w_n=\Pr\left(\xi_n=1\middle|A_{n-1}\right)$ in the risky asset and $1-w_n$ in the risk-free asset at time *j*.

Our return at time j is given by $1+w_nr_n$, and our final payoff is given by $\Pi_{n=1}^{N}(1+w_nr_n)$. Since $w_n\in[0,1]$ and $r_n>-1$ we have that:[13]

$$\ln(1+w_nr_n)\geq w_n\ln(1+r_n)=E\left(\pi_{n,\xi_n}\middle|\pi_n,A_{n-1}\right)$$

Hence we conclude that:[14]

$$\sum_{n=1}^{N}\ln(1+w_nr_n)\geq E\left(\sum_{n=1}^{N}\pi_{n,\xi_n}\middle|\pi^N\right)$$

Because $(\pi_{n,1}-\pi_{n,2})^2=(\ln(1+r_n))^2$, we can apply **Corollary A** to conclude that:

---

[13]For a given $w\in[0,1]$ let $f(r)=1+wr$ and $g(r)=(1+r)^w$. Note that $f(0)=g(0)=1$, $f'(0)=g'(0)=w$. Since *g* is concave while *f* is linear in *r* we have that $f(r)\geq g(r)$ for $r>-1$.

[14] Note that the use of log returns is crucial here. Consider a two period model in which the stock price doubles in both periods with certainty. Suppose that an investor first chooses with equal probabilities whether to invest his entire wealth in the stock or the bond. He does not change his decision in the second period so in each period the expected fraction invested in the stock equals a half. This random strategy yields 1 with probability 0.5 and 4 with probability 0.5 so on average 2.5. Using the procedure outlined in the text we transform this strategy to a deterministic one by investing half of our wealth in the stock in both periods; this strategy yields 2.25 with certainty. However, once we look at log returns the randomized strategy yields on average $0.5\ln(4)=\ln(2)$ while the deterministic one yields $2\ln(1.5)=\ln(2.25)$.

PROPOSITION 2. The gradient trading strategy implies $G_N \geq \exp(-q(r))\max\{1, S_N\}$ where $q^2(r)$ is the quadratic variation of the log returns, or:

$$q(r) \equiv \sqrt{\sum_{n=1}^{N}(\ln(1+r_n))^2}\ . \tag{9}$$

Consider now the generalized gradient strategy with starting point (x, y). Let $A_n(x, y)$ denote our location after n rounds. Recall that in this case we have:

$$E\left[\max_i A^+_{n,i}(x,y) \middle| \pi^n\right] \leq \sqrt{q^2(r)+x^2+y^2}$$

In characterizing the performance of our trading strategy we must credit back the initial regret of x with respect to the stock and y with respect to the bond. Hence, we conclude that:

PROPOSITION 3. The generalized gradient trading strategy implies

$$G_N \geq \max\left\{\exp\left(y - \sqrt{q^2(r)+x^2+y^2}\right), \exp\left(x - \sqrt{q^2(r)+x^2+y^2}\right)S_N\right\},$$

3.3. Price Bounds for at-the-Money Options

In the previous section we described a family of trading strategies and characterized their regret conditional on the realized price path. We now show how we can use this result to provide upper bounds for the value of options.

Let Φ be a set of possible price paths for the stock. Conditional on this set we assume that there is *no arbitrage* in prices. Namely, for any trading strategy G such that $G_n(r) > 1$ for some $r \in \Phi$, there exists another price path $r' \in \Phi$ such that $G_n(r') < 1$. Otherwise, investing in G and shorting the bond would lead to a profit at date j given path r_1 with no possibility of a loss.[15] A European call option with strike price K that matures at time N has a final payoff of $\max\{0, S_N - K\}$. Let Φ be the set of feasible stock price paths, and let $C(K|\Phi)$ be the highest value of the call option at time 0 that is consistent with no arbitrage.

Before proceeding we should discuss the importance of imposing restrictions on the price path given by Φ. The first part of Merton's (1973) paper addresses this question by asking what can

[15] An arbitrage opportunity is any trading strategy that generates a profit in some state without the possibility of a loss. The condition given here is sufficient to rule out arbitrage opportunities using the stock and the bond.

be said about the value of a call option without making any additional assumptions about the price path. The answer is that

$$C(K \mid \Phi) \le S_0$$

Hence, the option is not more valuable than the underlying asset. This is a very weak bound but cannot be improved if arbitrary price paths are allowed. Our goal is to find bounds for the option value $C(K|\Phi)$ using the gradient trading strategies; in this section we begin by considering the case of an at-the-money option (K=1). We begin by formalizing the link between regret limiting strategies and option prices discussed (as an example) in the introduction:

PROPOSITION 4. Suppose we have a dynamic trading strategy that satisfies $G_N \ge \beta(r) \max\{1, S_N\}, \forall r \in \Phi$. Let

$$\beta^* = \inf_{r\in\Phi}\{\beta(r)\} \tag{10}$$

Then no arbitrage implies the following upper bound for the value of an at-the-money call option:

$$C(1 \mid \Phi) \le \frac{1}{\beta^*} - 1 \tag{11}$$

Proof: Investing $1/\beta^*$ in the trading strategy and borrowing \$1 leads to a payoff

$$\frac{G_N(r)}{\beta^*} - 1 \ge \frac{\beta(r)\max\{1, S_N\}}{\beta^*} - 1 \ge \max\{0, S_N - 1\} \text{ QED}$$

To gain some intuition, consider a very simple trading strategy. Suppose we decide to use a buy and hold strategy in which we invest a fraction $\beta = 1/2$ in both assets. The future payoff of this static portfolio is

$$G_n = 0.5 + 0.5S_N \ge \max\{0.5, 0.5S_N\}$$

Note that the above holds for all possible price paths. Using the above proposition we conclude that:

$$C(1 \mid \Phi) \le \frac{1}{0.5} - 1 = 1 = S_0$$

As mentioned before, S_0 is the simple known upper bound from Merton (1973). To improve upon this bound, we must consider strategies with better performance guarantees that lead to a higher β^*.

Consider first the simple gradient trading strategy and define $q(\Phi) = \sup_{r\in\Phi} q(r)$, the highest possible realized quadratic variation of the log returns for the price paths in Φ. Combining **PROPOSITION 2** and **PROPOSITION 4**, we have that

$$C(1\,|\,\Phi) \le \exp(q(\Phi)) - 1 \tag{12}$$

Now consider generalized gradient trading strategy with x=y; in this case,

$$\beta(r) = \exp\left(x - \sqrt{q^2(r) + 2x^2}\right)$$

Taking $q(\Phi)$ and maximizing over x we get $x^*(q) = q(\Phi)/\sqrt{2}$. Hence, we obtain the following stronger result:

> **PROPOSITION 5.** Based on the generalized gradient trading strategy, the no arbitrage price of an at-the-money call option satisfies
>
> $$C(1\,|\,\Phi) \le \exp\left(\tfrac{1}{\sqrt{2}} q(\Phi)\right) - 1 \tag{13}$$

By focusing on small q we can quickly compare it to the Black-Scholes option pricing formula. Using the original gradient trading strategy we have:

$$C(1\,|\,\Phi) \le \exp\big(q(\Phi)\big) - 1 \approx q(\Phi) \text{ for small } q. \tag{14}$$

Using the optimal generalized gradient trading strategy we have:

$$C(1\,|\,\Phi) \le \exp\left(q(\Phi)\big/\sqrt{2}\right) - 1 \approx q(\Phi)\big/\sqrt{2} \text{ for small } q. \tag{15}$$

When contrasting these bounds with the Black-Scholes option pricing formula, note that the Black-Scholes formula also assumes the stock price path is continuous, and the volatility remains constant over time. Given a volatility of σ for the stock price, and the opportunity to trade continuously until the option's expiration on date T, the Black-Scholes assumptions on the stock price path implies $q(r) = \sigma\sqrt{T}$ for all $r \in \Phi_{BS}$. We have not made such assumptions (namely the stock price can jump and trading may not be continuous) and so the bounds we derive are necessarily weaker than that given by Black-Scholes. In particular, we have[16]

$$C(1 \mid \Phi_{BS}) \approx q(\Phi_{BS}) \big/ \sqrt{2\pi} \text{ for small } q. \tag{16}$$

Comparing equations (15) and (16), we can see that option price bound provided by the generalized gradient trading strategy exceeds the Black-Scholes price by an amount equivalent to an increase in the stock's volatility by a factor of $\sqrt{\pi}$, or approximately 77%. This increase bounds the impact on the option price from relaxing the Black-Scholes assumption of continuous price paths and constant volatility.

3.4. Different Strike Prices

We can generalize the methods of Section 3.3 to options with different strike prices using the following result:

PROPOSITION 6. Suppose we have a dynamic trading strategy that satisfies $G_N \geq \max\{\alpha(r), \beta(r) S_N\}, \forall r \in \Phi$ and let

$$\beta^* = \inf_{r \in \Phi} \{\beta(r), \alpha(r)/K\} \tag{17}$$

Then no arbitrage implies the following upper bound for the value of a call option with strike price K:

$$C(K \mid \Phi) \leq \frac{1}{\beta^*} - K \tag{18}$$

Proof: Investing $1/\beta^*$ in the trading strategy and borrowing \$$K$ leads to a payoff

$$\frac{G_N(r)}{\beta^*} - K \geq \frac{\max\{\alpha(r), \beta(r) S_N\}}{\beta^*} - K \geq \max\{K, S_N\} - K = \max\{0, S_N - K\} \text{ QED}$$

In **PROPOSITION 3** we demonstrated that the generalized gradient trading strategy with initial regret (x, y) provides a performance guarantee of

---

[16] To derive this result, let $f(\sigma)$ be the Black-Scholes value of an at-the-money call option with expiration T when the stock's price is 1, its volatility is σ, and the interest rate is zero. Then $f(\sigma)$ is the corresponding "vega" (volatility sensitivity) of the option, and for σ small, $f(\sigma) \approx f(0) + f'(0)\sigma = 0 + \left(\frac{1}{\sqrt{2\pi}} \sqrt{T}\right)\sigma$.

$$\alpha(r) = \exp\left(y - \sqrt{q^2(r) + x^2 + y^2}\right), \qquad \beta(r) = \exp\left(x - \sqrt{q^2(r) + x^2 + y^2}\right)$$

Note that in this case, $\alpha(r) = \exp(y - x)\beta(r)$. If we define $q(\Phi) = \sup_{r\in\Phi} q(r)$ and $k = \ln K$, then from (17),

$$\ln \beta^* = \min\{0, y - x - k\} + x - \sqrt{q^2(\Phi) + x^2 + y^2} \tag{19}$$

Eq. (19) implies that to maximize β^*, we should choose $y - x = k$. For if $y - x > k$, we can increase β^* by increasing x, and if $y - x < k$, we can increase β^* by increasing y. Thus, we choose $y = x + k$, and (19) reduces to

$$\ln \beta^* = x - \sqrt{q^2(\Phi) + x^2 + (x + k)^2} \tag{20}$$

Thus, the quality of the option price bound provided by the generalized gradient strategy depends upon the choice of x in Eq. (20). Recall that x determines the initial regret point, $A_0 = (x, y = x + k)$. Because the initial regret must be non-negative, we require

$$x \geq \max(0, -k) \tag{21}$$

The best bound possible using the generalized gradient strategy can thus be found by choosing x to maximize (20) subject to (21). The optimal choice of x depends upon the maximal quadratic variation q^2 as follows:

$$x^*(q) = \tfrac{1}{2}\sqrt{2q^2 + k^2} - \tfrac{1}{2}k \tag{22}$$

Replacing x in (20) with $x^*(q(\Phi))$ leads to best possible performance guarantee:

$$\ln \beta^* = -\tfrac{1}{2}\sqrt{2q^2(\Phi) + k^2} - \tfrac{1}{2}k \tag{23}$$

Combining (23) and **PROPOSITION 6**, the gradient trading strategies imply the following no arbitrage bound for the price of a call option:

PROPOSITION 7. The no arbitrage price of a call option with strike price K satisfies

$$C(K \mid \Phi) \leq \exp\left(\tfrac{1}{2}\sqrt{2q^2(\Phi) + (\ln K)^2} + \tfrac{1}{2}\ln(K)\right) - K \tag{24}$$

This bound is achieved by borrowing K and investing the total in the generalized gradient trading strategy with initial regret $x = x^*(q(\Phi))$ and $y = x + \ln K$.

Again, we can compare the option price bound from the generalized gradient strategy with that of Black-Scholes. **Figure 2** shows the generalized gradient bound, expressed as a Black-Scholes implied volatility, when q = 10%. We can interpret the figure as showing the potential impact on the value of the option from dropping the Black-Scholes assumptions of a continuous price path with a constant volatility, where that impact is expressed in terms of the equivalent increase in the volatility of the stock.

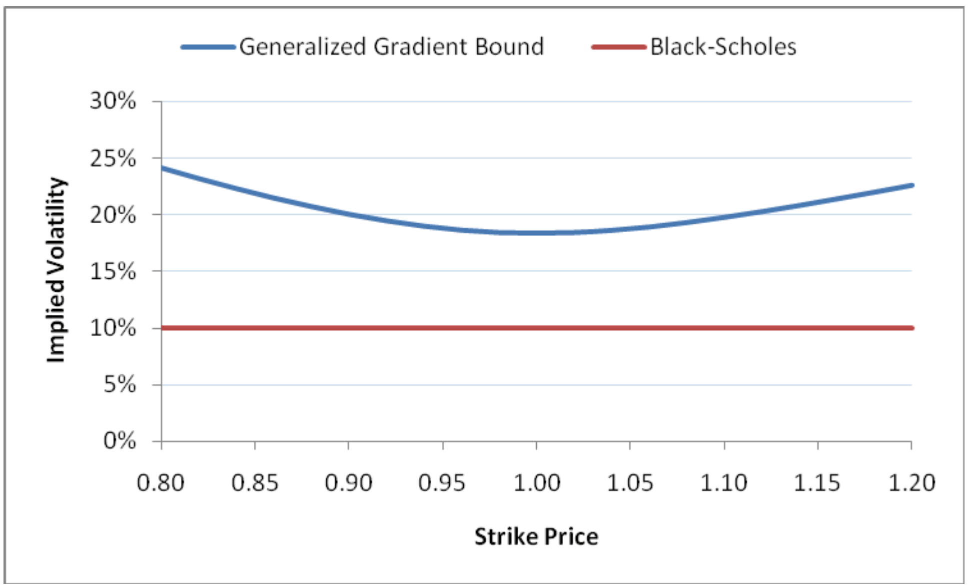

Figure 2 Generalized Gradient Bounds versus Black-Scholes

4. Optimal *q*-based Bounds

Thus far, we have demonstrated that we can obtain meaningful option price bounds by extending and generalizing Blackwell's gradient strategy to an investment trading environment. These bounds depend upon the maximal quadratic variation of the possible stock price paths, $q^2(\Phi)$. In this section we take a different approach, and define the set of allowable stock price paths as all paths with quadratic variation less than some bound:

$$\Phi(\overline{q}) = \{r \mid q(r) \le \overline{q}\}$$

We then ask, what is the best bound available for the set $\Phi(\overline{q})$? That is, what is $C(K \mid \Phi(\overline{q}))$? We demonstrate how to compute this bound recursively in this section. It will improve on the bounds that we derived based on gradient strategies in Section 3.

4.1. Optimality and Path Independence

The Hannan and Blackwell gradient strategy that we examined in subsection 3.2 is an example of a strategy that is path dependant. That is, the portfolio chosen by the strategy depends on the strategy's past performance. In contrast, the trading strategy in Black and Scholes is path independent, in that it only depends on the current stock price and the volatility of the stock over the remainder of the option's life. In this section we argue that the optimal regret minimizing strategy should also share this property of path independence:

> Definition: We say that a dynamic trading strategy is path independent if the portfolio weights at time *n* depend only on the stock price at time *n*, S_n, and the possible future paths Φ*n*.

To see why an optimal strategy should be path independent, consider a scenario where at time *n* the stock price is given by S_n and the value of our portfolio is G_n. Let $G_{n,N}^{\vec{w}}$ denote the payoff at time *N* of a strategy that starts with a wealth of \$1 at *n* and follows a dynamic portfolio strategy $\vec{w}$. Then the regret minimizing strategy must solve

$$\max_{\vec{w}} \min_{r \in \Phi_n} \frac{G_n G_{n,N}^{\vec{w}}}{\max\{1, S_N\}} = G_n \left(\max_{\vec{w}} \min_{r \in \Phi_n} \frac{G_{n,N}^{\vec{w}}}{\max\{1, S_N\}} \right)$$

Thus, the optimal strategy is independent of G_n.

An immediate implication is that the Hannan-Blackwell gradient trading strategy is suboptimal. The claim extends to original additive setup; also there the Hannan and Blackwell does not minimize regret over finite horizon. This is clearly not a criticism of their result as when the horizon is infinite the average regret of their strategy is optimal.

One way of making the gradient strategy to be path independent can be based on our generalized gradient strategy. Recall from Section 2.2 that we can interpret the generalized gradient strategy as though we started the game with an arbitrary initial relative performance vector A_0. Then in

Section 33.3 we computed the optimal generalized gradient trading strategy, and showed that the optimal starting position depended only on the remaining quadratic variation of the price paths, and the ratio of the current stock price to the option strike price. That is, it is optimal at each point in time to "forget" the actual past performance of the strategy, and choose the current portfolio in a path independent manner. In the remainder of this section, we implement this idea by using backward induction to solve for the optimal regret minimizing trading strategy.

4.2. The Super-Replicating Portfolio

We solve for the optimal price bound by computing the lowest cost of a super-replicating portfolio, which is a portfolio of the stock and the bond whose value exceeds the value of the call option after any move of the stock price. Let $V(S, q^2)$ be the cost of the super-replicating portfolio, given a current stock price of S, a strike price of 1, and a future return path with a quadratic variation of at most q^2. Obviously, at the boundary $q^2 = 0$, we have

$$V(S,0) = \max\{0, S - 1\}. \tag{25}$$

For $q^2 > 0$, the function V is an upper bound for the value of the option if in addition to (25), V satisfies

$$\begin{aligned} V(S,q^2) = \quad & \min\nolimits_{\Delta,B}\ \Delta \times S + B \\ & s.t. \quad \Delta \times S e^{\tilde{r}} + B \geq V\left(S e^{\tilde{r}}, q^2 - \tilde{r}^2\right) \\ & \qquad\ \ \text{for all } \tilde{r} \text{ such that } \tilde{r}^2 \leq q^2 \end{aligned} \tag{26}$$

That is, V is the minimum cost of a portfolio, which holds Δ shares of the stock and B units of the bond, such that given any feasible return for the stock, the new value of the portfolio will be sufficient to purchase the option given the new stock price and the remaining quadratic variation of the returns.

The optimal bound is the smallest function V^* that satisfies (25) and (26). To see that V^* is well-defined, note that $V(S, q^2) = S$ for $q^2 > 0$ is a solution, and that if V^1 and V^2 are solutions, so is $\min(V^1, V^2)$. Finally V is bounded below by the intrinsic value of the option.

4.3. A Lower Bound

While we cannot solve for V^* analytically, we can compute it numerically. But to gain further insight into the form of V^*, as well as the behavior of the worst case price paths, in this section we derive an analytic lower bound for V^* that turns out to be a very good approximation of the true bound.

The limit on the quadratic variation of the stock price path bounds the size of any jumps in the stock price. Specifically, the constraint $\tilde{r}^2 \le q^2$ implies that the percentage jump in the stock price must be within the range $[-q_-, q_+]$, where

$$q_+ = e^q - 1 \text{ and } q_- = 1 - e^{-q} \tag{27}$$

In this section we argue that a lower bound for V^* is equal to the function $\hat{V}$ defined as follows:

$$\hat{V}(S, q^2) \equiv \begin{cases} \dfrac{q_+ q_-}{q_+ + q_-} S^{1/q_-} & \text{for } S \le 1 \\ \dfrac{q_+ q_-}{q_+ + q_-} S^{-1/q_+} + S - 1 & \text{for } S \ge 1 \end{cases} \tag{28}$$

To motivate the function $\hat{V}$, suppose $S = 1$ and $n = 1$. Then, the highest value for the option is obtained if the stock price moves maximally, to $(1 + q_+)$ or $(1 - q_-)$. The risk-neutral probability[17] that the stock price rises is then $q_- / (q_+ + q_-)$, so that the value of the option is $[q_- / (q_+ + q_-)][(1 + q_+) - 1] = q_+ q_- / (q_+ + q_-) = \hat{V}(1, q^2)$. Thus, this must be a lower bound for V when $S = 1$.

Next suppose $S < 1$. Suppose the stock price drifts up slowly at some rate $\hat{r}\,dt$ until the stock price equals 1, but with some chance stock price drops q_- and stays there. If there is no drop, the stock price will reach 1 at time t such that $Se^{\hat{r}t} = 1$. Because the risk-neutral hazard rate that the stock price will drop is $(\hat{r}/q_-)dt$,[18] the risk-neutral probability that there is no drop before time t is $e^{-(\hat{r}/q_-)t} = S^{1/q_-}$. Given the value of the option when $S = 1$ calculated earlier, this leads to the

[17] That is, the probability that makes the expected return of the stock equal to the risk-free rate of zero.

[18] If the arrival rate of the drop is λdt, then for the expected return to be zero, $\lambda\, q_- dt = \hat{r}\, dt$.

lower bound for V when $S < 1$ shown in (28).[19] A symmetric calculation (assuming the stock price drifts down to 1 or jumps up maximally) leads to the lower bound for V when $S > 1$.[20]

Thus, we have shown by construction the following result:

PROPOSITION 8. If V satisfies (26), then $V(S,q^2) \geq \hat{V}(S,q^2)$.

For an at-the-money call option, this lower bound implies that

$$C(1\,|\,\Phi) \geq \frac{\exp(q)+\exp(-q)-2}{\exp(q)-\exp(-q)} \approx q(\Phi)/2 \text{ for small } q. \tag{29}$$

Comparing (29) with (15) and (16), we see that there is room for substantial improvement beyond the generalized gradient strategy, but that (not surprisingly) there is still a significant increase in the potential value of the option from relaxing the Black-Scholes assumptions.

4.4. Numerical Computation of V^*

We can compute the optimal bound V^* via backward induction on the number n of remaining price movements. The limiting case as n becomes large is equivalent to the solution when there is no limit to the number of price movements. Specifically, let $V(S,q^2,n)$ be the value with n periods remaining, and note that

$$V(S,q^2,0) = \max\{0, S-1\} \tag{30}$$

Then we can compute the bound for $n > 0$ inductively as follows:

$$\begin{aligned} V(S,q^2,n) \equiv \quad & \min_{\Delta,B} \Delta \times S + B \\ & s.t. \quad \Delta \times Se^{\tilde{r}} + B \geq V\left(Se^{\tilde{r}}, q^2 - \tilde{r}^2, n-1\right) \\ & \qquad\quad \text{for all } \tilde{r} \text{ such that } \tilde{r}^2 \leq q^2 \end{aligned} \tag{31}$$

It is easy to show by induction that for all n, V is weakly increasing in n and $V \leq V^*$. Thus, V converges monotonically to V^*.

Note also that to facilitate computation, we can use the result of section 4.3 and start the numerical calculation with

[19] Note that if the stock price drifts smoothly, the remaining quadratic variation is unchanged at q^2, as a smooth continuous drift has no quadratic variation.

[20] Note that we can compute the value of the call option as the value of a put option plus the stock plus \$1 of debt.

$$V(S,q^2,0)=\hat{V}(S,q^2) \quad (32)$$

in place of (30).

We find numerically that while the optimal bound does indeed exceed $\hat{V}$, the difference is rather negligible. For example, for q = 50%, the optimal bound V^* exceeds $\hat{V}$ by leas than 0.03%, and the gap disappears as q approaches 0. Thus $\hat{V}$ provides a very useful approximation to the optimal bound.

4.5. Comparison with Black-Scholes

Again, we can compare the optimal q-based bound with the option price that would be obtained if the Black-Scholes assumptions were satisfied, as well as the generalized gradient strategy of Section 3. **Figure 3** shows the comparison in terms of the corresponding Black-Scholes implied volatility.

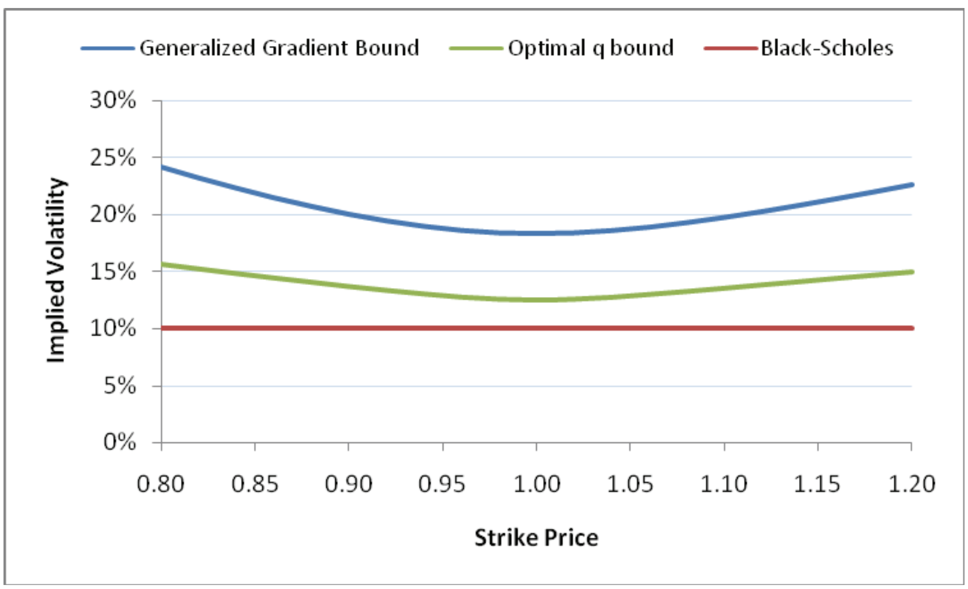

Figure 3 Option Price Bounds versus Black-Scholes

In **Figure 4**, we show the optimal hedging strategy for the optimal q-based bound. Specifically, given the bound V^*, we can compute the number of shares of the stock Δ to hold at each point in time in order to super-replicate the option at the lowest possible cost. Indeed, from Eq. (26), we have $\Delta = V^*_1(S,q^2)$, shown in **Figure 4** for the case q = 0.10. As a comparison, we also show the delta from the Black-Scholes trading strategy given a volatility of q, as well as the implied volatility of the option if it were trading at the price implied by the bound.

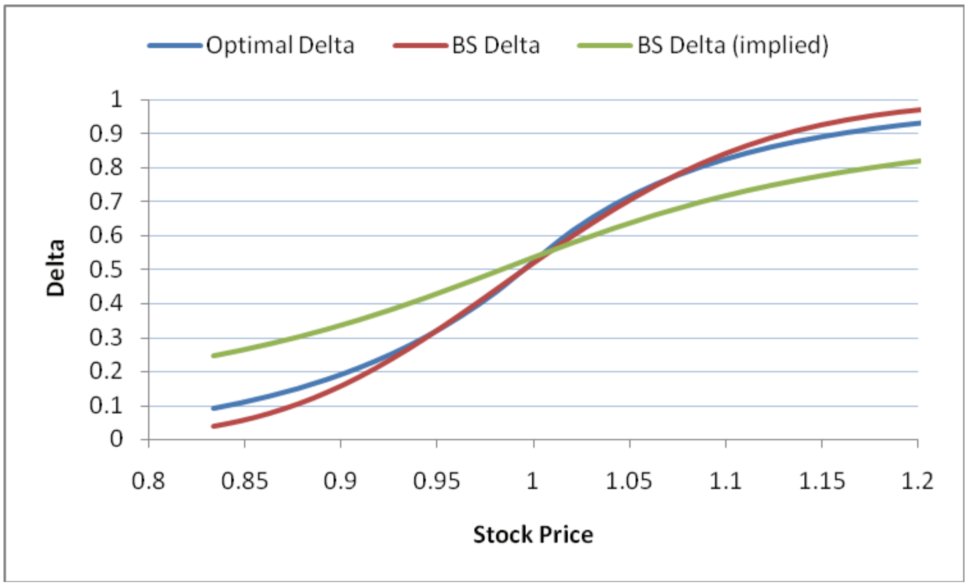

Figure 4 Delta Hedging Strategy for the Optimal q-based Bound versus Black-Scholes

Finally, we compute the cost of hedging a call option using the trading strategy from the optimal q-based bound. By definition, the cost of hedging the option is at most $V^* (S,q^2)$. But because the trading strategy only super-replicates, rather than replicates, the option, the cost will be less than this amount for almost all price paths. In **Figure 5** we compute the range of hedging cost for an at the money option ($S = 1$) with an initial quadratic variation of $q = 0.50$. We also show the same range when we restrict the price paths to be continuous. As a benchmark we compare this with the range for the Black-Scholes strategy. Note that with a continuous price path and a fixed quadratic variation, the cost using Black-Scholes is identically equal to the Black-Scholes option value for all such price paths. But once we allow for jumps, the worst case under Black-Scholes is indeed worse than that under the optimal strategy.

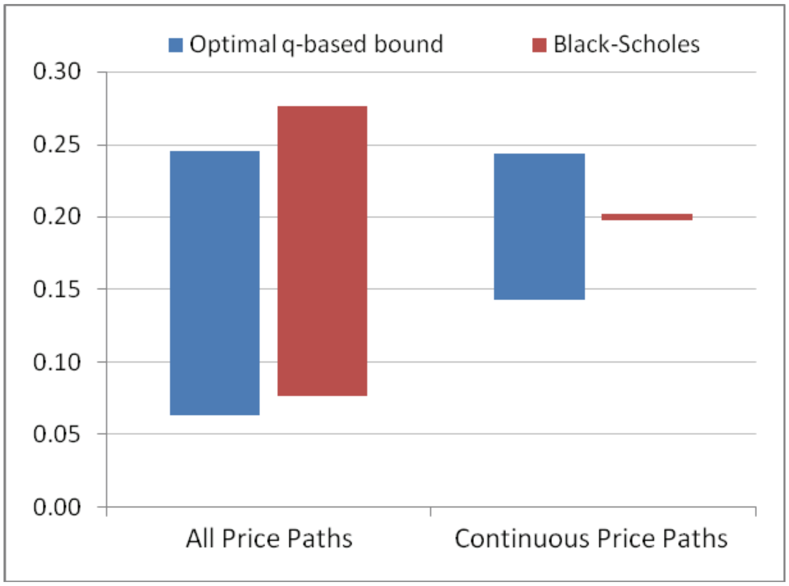

Figure 5 Hedging Cost using the Optimal q-based Bound versus Black-Scholes

5. Conclusion

We have examined the relation between concepts of approachability and regret minimization in games and no arbitrage price bounds in financial markets. We demonstrate that the simple strategies of Hannan and Blackwell can be used to derive dynamic trading strategies that imply robust upper bounds for the value of a call option under different assumptions. We then showed that these strategies are path dependent and thus suboptimal. We then solved for the optimal upper bound for the option price, which simultaneously implies the regret minimizing strategy in the dynamic trading game.

The option price bounds that we derive are for stock price paths whose quadratic variation is bounded, as in the Black-Scholes framework, but where we have dropped the assumption of continuity and continuous trading. One striking result of our analysis is that the optimal quadratic variation based bounds that we derive are "close" to the option prices derived in a Black-Scholes framework. For at-the-money call options, our bounds lead to Black-Scholes implied volatilities that are roughly 25% higher than what one would calculate based on the quadratic variation of the price path. In other words, relaxing the Black-Scholes assumption that price paths are continuous leads to relatively small change in the price of the option. Moreover, the implied volatilities predicted by our model are much closer to those found empirically. This result suggests that our bounds may be relevant for option pricing in practice.

6. References

Al-Najjar, Nabil, and Jonathan Lewis Weinstein. 2008. "Comparative Testing of Experts," Econometrica., forthcoming

Bergmann D. and K. Schlag (2005) "Robust Monopoly Pricing: The Case of Regret," working paper, Yale University.

Bertsimas D., L. Kogan and A.W. Lo (2001) "Hedging derivative securities and incomplete markets: an ε-arbitrage approach,*" Operation Research* 49(3) 372-397.

Bollen N.P and R. Whaley (2004) "Does Net buying Pressure Affect the Shape of Implied Volatility Functions?," *Journal of Finance* 59(2) 711-754

Bernardo A. E. and O. Ledoit (2000) "Gain, loss and asset pricing," *Journal of Political Economy,* 108:144—172.

Black F. and M. Scholes (1973) "The pricing of options and corporate liabilities," *Journal of Political Economy*, 81:637—654.

Blackwell D. (1956), "An analog of the MiniMax theorem for vector payoffs," *Pacific Journal of Mathematics*, 6:1—8.

Borodin A. and R. El-Yaniv. (1998) "Online Computation and Competitive Analysis," Cambridge University Press.

Bose, S., E. Ozdorzen and A. Pape (2004) "Optimal Auctions with Ambiguity", working paper, University of Texas, Austin and University of Michigan.

Cochrane J.H. and J. Saa-Requejo (2000) "Beyond arbitrage: Good-deal asset price bounds in incomplete markets," *Journal of Political Economy*, 108:79—119.

Cover T. (1991), "Universal portfolios", *Mathematical Finance*, 1:1-29.

Cover T. (1996), "Behavior of sequential predictors of binary sequences," *Transactions of the Fourth Prague Conference on Information Theory.*.

Cover T. and E. Ordentlich.(1996) "Universal portfolios with side information," *IEEE Transactions on Information Theory*, 42: 348--368.

Cover T. and E. Ordentlich.(1998) "The cost of achieving the best portfolio in hindsight," Mathematics of Operations Research, 960--982,

Darrell Duffie D. (2001). "Dynamic Asset Pricing Theory," Princeton University Press, 2001.

Dekel, E. and Y., Feinberg 2006, "Non-Bayesian Testing of a Stochastic Prediction" Review of Economic Studies, 73, 4, pp. 893-906.

Eraker B. (2004) "Do stock prices and volatility jump? reconciling evidence from spot and option prices," *Journal of Finance*, 59:[1367-1403.

Eraker B., M. Johannes, and N. G. Polson. (2003) "The impact of jumps in returns and volatility," *Journal of Finance*, 53:1269-1300.

Dean Foster (1999) , "A proof of Calibration via Blackwell's Approchability Theorem," Games and Economic Behavior, 73 - 79.

Foster D., and R. Vohra. (1993) "A randomized rule for selecting forecasts," *Operations Research*, 41:704-709.

Foster D., and R. Vohra. (1997) "Regret in the on-line decision problem," *Games and Economic Behavior*, 21:40-55.

Foster D., and R. Vohra. (1998) "Asymptotic calibration," *Biometrika*, 85:379-390.

Foster D., D. Levine and R. Vohra. (1999) "Introduction to the Special issue," , *Games and Economics Behavior,* 29:1-6.

Fudenberg D., and D. Levine (1998) "Theory of Learning in Games", *Cambridge, MA: MIT Press*

Gilboa I. and D. Schmeidler (1989) "Maxmin expected utility with non-unique prior," *Journal of Mathematical Economics*, 18:141-153.

Grundy B. and Z. Wiener (1999) "The analysis of Deltas, State Prices and Var: A New Approach," *Working paper*

Hannan J. (1957) “Approximation to bayes risk in repeated plays,” *In M. Dresher, A. Tucker, and P. Wolfe, editors, Contributions to the Theory of Games,* 3: 97--139. Princeton University Press, 1957.

Hart S. (2005) “Adaptive Heuristics,” *Econometrica* 73:5:1401-1431

Hart S., and A. Mas-Colell (2000) “A simple adaptive procedure leading to correlated equilibrium,” *Econometrica*, 68:1127-1150.

Hayashi T. (2005) “Regret aversion and opportunity dependence,” working paper University of Texas, Austin.

Merton, R. C. (1973) “Theory of rational option pricing,” *Bell Journal of Economics and Management Science*, 4 (1):141-183.

Milnor J. (1954) ”Games Against Nature”, In Decision Processes ed R.M. Thrall, C.H. Coombs & R.L. Davis. New York Riley.

Lo A.W. (1987) “Semi-parametric upper bounds for option prices and expected payoffs,” *Journal of Financial Economics* 19:373-387.

Mykland P. (2000) “Conservative delta hedging,” *The Annals of Applied Probability*, 664—683.

Pan J. (2002) “The jump-risk premia implicit in options: Evidence from an integrated time-series study,” *Journal of Financial Economics,* 63: 3-50.

Shreve, S., N. El Karoui, and M. Jeanblanc-Picque (1998) “Robustness of the Black and Scholes formula,” *Mathematical Finance*, 8:93-126.

Sion M. (1958), “On general minmax theorems,” Pacific *Journal of Mathematics*, 8, 171-176.

Shafer G. and V. Vovk (2001) “Probability and Finance: It's Only a Game!” *Wiley*

Sleator D. and R. E. Tarjan (1985) “Amortized efficiency of list update and paging rules,” *Communications of the ACM*, 28:202-208.

Wojciech O. and A. Sandroni (2008) “Manipulability of Future-Independent Tests”, *Econometrica*, forthcoming.