University of Warsaw
Faculty of Economic Sciences

Working Papers
No. 20/2026 (514)

Hedging Auction Volatility with Gap Call Options

Gilbert Mbara

Warsaw 2026
ISSN 2957-0506

WORKING PAPERS 20/2026 (514)

UNIVERSITY OF WARSAW
FACULTY OF ECONOMIC SCIENCES
WORKING PAPERS

Hedging Auction Volatility with Gap Call Options

Gilbert Mbara[1*]

[1] *University of Warsaw, Faculty of Economic Sciences*
[*] *Corresponding author: g.mbara@uw.edu.pl*

Abstract: We investigate the feasibility of hedging commodity price risk using gap call options within an auction-based market devoid of traditional derivatives. Using a novel, high-frequency dataset from the Nairobi Coffee Exchange (NCE), we model spot price dynamics by deriving a Geometric Brownian Motion process from the independent private values paradigm. The estimated model captures the unique microstructure of the NCE, where discrete weekly auctions generate prices characterized by extreme volatility (146.5% annualized). We utilize Monte Carlo simulation to price and evaluate the performance of gap call options for buyers seeking protection against catastrophic price spikes. Our results demonstrate that gap options – characterized by a trigger price higher than the strike – provide superior risk-adjusted returns compared to standard European calls. Our study offers a practical framework for developing tailored risk management instruments in emerging commodity exchanges, and provides empirical evidence for the viability of gap options as a cost-effective hedging tool in high-volatility, institutionally constrained markets.

Keywords: risk management, hedging, commodity price risk, gap options, Nairobi Coffee Exchange, geometric Brownian motion, Monte Carlo simulation

JEL codes: G11, G13, G32, Q14

1 Introduction

Coffee stands as one of the world's most commercialized tropical products and the second most traded commodity globally after petroleum (Mussatto, Machado, Martins and Teixeira, 2011; Pancsira, 2022). However, its market is characterized by pronounced price volatility, exceeding that of other soft commodities such as cocoa, corn, soybeans, and sugar (Hollstein, Prokopczuk and Würsig, 2020). For producers, traders, and buyers – particularly in developing economies – this volatility translates into significant financial risk, jeopardizing revenue stability, cost predictability, and long-term investment. Conventional risk management theory prescribes the use of derivative securities, such as futures and options, to hedge against adverse price movements (Myers, 1993; Hull and Basu, 2016). In well-developed markets, producers can sell part of their output forward, while buyers can lock in purchase prices, thereby mitigating spot price risk (Veld-Merkoulova and de Roon, 2003; Bush, 2012).

Yet, a critical gap exists in many emerging commodity markets: the absence of organized derivative exchanges where localized price risk can be traded. The Nairobi Coffee Exchange (NCE), a weekly auction market handling coffees from at least five African nations, epitomizes this challenge. While its volume is modest compared to large producers like Brazil, the NCE exerts a meaningful influence on global benchmarks; its "Colombian Milds" constitute 48–54% of the International Coffee Organization's price calculation. Despite this relevance, the NCE operates as a pure spot market with no affiliated futures or options contracts, leaving its participants fully exposed to the auction's inherent price swings. This deficit motivates an urgent question: can financially engineered derivatives, specifically tailored to the auction's microstructure, provide effective and viable hedging for its participants?

This paper investigates the feasibility of using *gap call options*, a form of exotic derivative, to manage price risk for buyers at the NCE. A gap option differs from a standard (vanilla) option through the inclusion of a trigger price H. For a gap call, the option only pays off if the underlying asset's price at expiration S_T exceeds H, at which point the payoff is $S_T - K$, where K is the strike price. This structure creates a discontinuous payoff profile, effectively providing *catastrophic insurance* against extreme price spikes while leaving the holder exposed to smaller, more frequent price increases. The core economic appeal is a significantly lower premium compared to a standard call, as the option is less likely to be exercised. We conceptualize the hedging strategy for a representative buyer, such as a large exporter or roaster, who must procure coffee weekly at the NCE. Facing uncertain auction prices, the buyer purchases a gap call option with a strike K near the current price level and a trigger $H > K$. The buyer pays a premium upfront. If the auction price concludes above H, the option pays out, offsetting the higher purchase cost. If the price finishes below H, the option expires worthless, and the buyer bears the loss of the premium but purchases at the lower market price. This structure allows the buyer to define a maximum effective purchase price ($K + Premium$) in all scenarios while participating in favorable moderate or low-price outcomes.

Using Monte Carlo simulation based on a structurally estimated price process, we evaluate the profitability and risk-adjusted performance of this hedge. Our results indicate that gap options can dominate standard calls. For instance, with a trigger set 50% above the strike ($H/K = 1.5$), the gap call achieves a higher expected profit and loss (P&L) and a superior Sharpe ratio. The reduction in premium more than compensates for the forgone payoffs in the region between K and H. This makes the gap call a cost-effective tool for buyers whose primary concern is insuring against severe, tail-risk price surges rather than smoothing all price variability.

Relation to the Literature Our study engages with and extends three interconnected strands of academic literature: commodity risk management in developing markets, exotic options for hedging, and auction theory for price process modeling. Prior research has extensively documented the challenges of price volatility in commodity-dependent economies and advocated for market-based risk management tools (Claessens and Duncan, 1994; Mohan, 2007). A common theme is the use of cross-hedging with internationally traded futures (e.g., ICE contracts) when local derivatives are absent (Ankirchner, Imkeller and Reis, 2010; Ankirchner, Dimitroff, Heyne and Pigorsch, 2012). Our study moves beyond the cross-hedging paradigm to explore the design of novel derivatives written directly on the local spot price, a solution more precise but contingent on the development of a local derivatives market.

The application of exotic options, including Asian (average price) and barrier options, for commodity risk management is well-established in financial engineering (Zhang, 1997; Hull and Basu, 2016). These instruments are prized for their ability to tailor payoffs to specific risk exposures, such as hedging an average price over a period. Our analysis specifically advances the literature on gap options, which have received less empirical attention in commodity contexts. We demonstrate how their binary-trigger feature is uniquely suited to the discrete, event-driven price jumps characteristic of auction markets, providing a more efficient payoff profile than standard options for hedging tail risk.

A rich literature models auction prices within the independent private values (IPV) paradigm (Paarsch, Hong and Haley, 2006; Donald and Paarsch, 1996). This paper bridges this auction-theoretic literature with financial econometrics. We derive the equilibrium price series from the IPV framework, showing it converges to a Geometric Brownian Motion (GBM). This provides a structural justification for using GBM – a workhorse of derivative pricing – in an auction context, grounding our simulations in the micro-foundations of bidder behavior rather than assuming GBM exogenously. This integration is a methodological contribution, offering a template for deriving tradable price processes from auction data in other nascent markets.

Novelty and Contribution The novelty of this study is threefold, anchored in its data, context, and methodological synthesis. First is the use of unique high-frequency auction data. We utilize a novel, granular dataset from the NCE, capturing the full sequence of bids, buyer identities, and transaction prices. This allows us to analyze heterogeneous buyer strategies and seller patterns, providing an empirical foundation rarely available for emerging market commodity auctions. Second, we address a real-world institutional void by focusing on a concrete, under-explored setting – a physically important but financially underserved auction market. Our study is not merely a theoretical exercise but a feasibility analysis with direct policy implications for exchange development and participant risk management in emerging economies. Finally, we integrate auction theoretic concepts to provide a novel microfoundation for stochastic processes. We uniquely combine the IPV approach to modeling bidder valuations with derivative pricing, demonstrating how financial instruments can be designed from the ground up based on the specific market microstructure of a commodity auction.

Structure of the Paper The remainder of this paper is structured as follows. Section 2 details the institutional workings of the NCE auction and presents descriptive statistics of our dataset, highlighting the price risk faced by major buyers and sellers. Section 3 derives the GBM price process from the IPV paradigm and presents parameter estimates. Section 4 introduces gap call options, details our Monte Carlo simulation framework, and presents a comprehensive analysis of hedging performance and sensitivity. Section 5 concludes, discussing the practical implications for the NCE, the limitations of our study, and avenues for future research.

2 Nairobi Coffee Exchange Auction and Data

The NCE is an auction where buyers and sellers meet to trade pre-submitted, graded and "tasted" lots of coffee. Both the buyers and sellers are registered members of the exchange. Auctions are held once every week, on Tuesdays or Wednesdays. In the next subsections, we give a brief summary of the auction mechanism and the organization of activities on a typical auction day. We then provide descriptive analysis of the data, form a both a buyer and seller perspective, demonstrating the price spikes generated by the auction mechanism.

2.1 The Auction Mechanism

The NCE operates a clock auction mechanism designed for the exchange by Aucsys,[1] a Belgian electronic trade systems developer. The NCE operates a "Dutch auction", where the seller selects a start price from which the price moves down until a buyer intervenes. A central display panel, as shown in Figure 1a below, shows the start price (usually augmented by a small random ± fraction of the start price). The display panel enables the buyers to identify the coffee on offer: lot number, description, production area, quality, proposed start price, the amount of each step by which the opening price will reduce (optional), buyer's name, plus information on upcoming lots, current auction status etc.

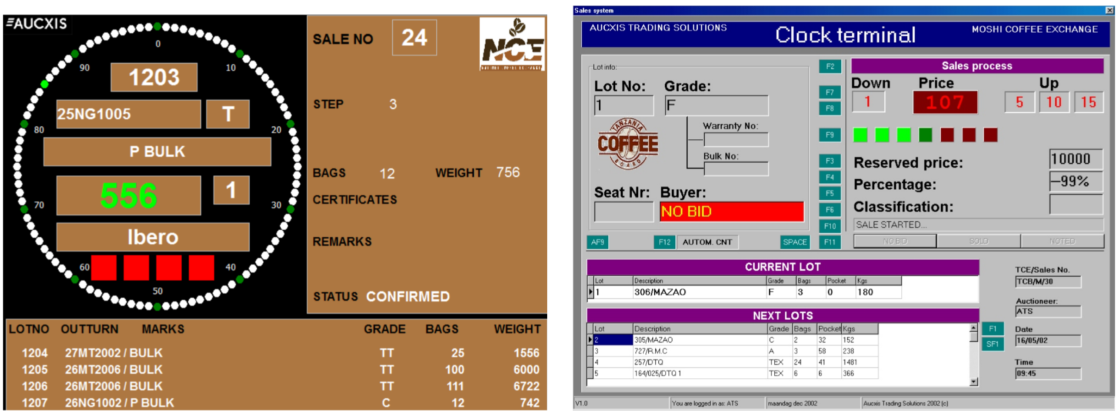

(a) Electronic Display Panel: NCE

(b) Seller Terminal: MCE

Figure 1: Auction Mechanism

[1] https://www.aucxis.com/en/news/aucxis-modernises-auction-system-nairobi-coffee-exchange

To begin the sale the auctioneer enters a price and a pre-determined clock step on a terminal as shown in Figure 1b. The price immediately starts to go down and, when a buyer is interested, s/he pushes one of three bid buttons, depending on the amount by which he wishes to increase the price. Then, as further bids are made, the value increases throughout the bidding time already determined by the auctioneer and displayed. When the time is up, the lot is assigned to the highest bidder. As soon as a lot is sold, the following lot moves from the "next lots" display to become the current lot. Before a typical auction day, "lots" of coffee are delivered to a common warehouse where "warrants" are issued by the "warehouse-men" in conjunction with the sellers or "marketing agents". After the a lot has been delivered to the warehouse with an accompanying sale catalogue, the sellers draw and present representative samples to the trade sample room. Sellers control the auction process, choosing their minimum prices, the speed of price changes and sequence of selling their inventory.

There are 80 different "buyer desks" where bids are submitted through a set of buttons available at each desk. The following information is collected at every sale: a transaction number, lot number, season, marks, grade, sale number (of the season), bags and weight bought, buyer code, price, seat number (of buyer in auction hall), agent code (marketing agent or seller identifier) and time. For each sale (of an individual lot from a single buyer), buyers are free to enter the bidding process or skip and wait for the next sale. Once the seller initiates the auctioning process by starting the clock timer, the price begins falling. By pressing a set of buttons available at each buyer's desk,[2] a buyer indicates interest in purchase, resulting into a reversal of the displayed price level which starts rising. If multiple buyers are pressing the buy button at the same time, the price rises faster, only stopping when there is one buyer left. The price at the last button press is the sale's price and the last remaining bidder is the winner of the lot. Once the auction has ended, sellers prepare invoices for the respective winning bidders and remit coffee warrants to the traders after payments have been made.

2.2 Data

The data used for this analysis are from Tuesdays; 13th August 2024, which represent any typical auction day(s). We give summary statistics showing the distribution prices and quantities sold/bought for different buyers and sellers. This forms the basis of our subsequent analysis of the price process and hedging.

2.3 Buyer Side Data

Table 1 shows the distribution of prices summarized by buyer identity. There are three major "winners" on this particular auction day: Ibero Kenya Ltd, Louis Dreyfus Company, and Taylor Winch (Coffee) Limited, each representing 29%, 19% and 21% of top bids, respectively. The distribution of buyer prices shows a market characterized by a dominant monopsonistic player, a diverse tier of strategic buyers, and significant price volatility that implies distinct sourcing strategies. Ibero Kenya emerges as the undisputed market leader, purchasing a commanding 29% of all lots at an average price of $211. This below-market-average price – when compared to the higher averages of other major players – suggests a high-volume, cost-conscious strategy, potentially focusing on commercial-grade coffee for blending or bulk export. However, its high standard deviation of 28.36 and wide price

[2]The NCE has 80 buyer desks. See: https://e-trade.aucxis.com/en/sector/other-sectors/

range (149-266) indicate a flexible mandate, acquiring both low-priced and mid-premium lots. The second and third largest buyers, Taylor Winch (21% share) and Louis Dreyfus Company (19% share), operated with notably different profiles. Taylor Winch achieved a substantially higher average price of \$ 242.31 with relatively low volatility (std dev 17.47), indicating a consistent, quality-focused acquisition strategy. In contrast, Louis Dreyfus exhibited the highest volatility among major buyers (std dev 38.93), with prices ranging from 130 to 301, reflecting a portfolio approach that spans the entire quality spectrum, from deeply discounted to premium lots.

Table 1: Price Distribution by Buyer, 13th August 2024

Buyer Name	Wins	% Win	Prices \$				Quantities		
			Mean	Min	Max	STD	Wgt. (tons)	%	Mean
BRIDGE EXPO & TRANS	1	0.2	261.0	261	261	0	1	0.1	733.0
Betco Coffee Co.	2	0.3	271.0	266	276	7	3	0.4	1430.5
C.Dorman	41	6.5	268.9	236	343	28	99	12.5	2408.9
Global Mark Foods	10	1.6	165.3	148	194	13	4	0.6	449.7
Ibero Kenya	184	29.1	211.7	149	266	28	184	23.1	997.6
Jabali The Coffee Company	7	1.1	253.0	245	262	6	7	0.9	1016.7
Jowam Coffee Traders	36	5.7	255.7	150	296	31	50	6.3	1392.6
Kenya Nut Company	2	0.3	278.5	275	282	5	1	0.2	680.5
Kenyacof	28	4.4	205.5	162	256	29	44	5.6	1574.5
Kyandu Trading Company	1	0.2	150.0	150	150	0	0	0.0	331.0
Louis Dreyfus Company	120	19.0	228.4	130	301	39	141	17.7	1172.2
Pazoori	4	0.6	260.2	244	292	22	4	0.5	935.5
Rockbern Coffee Group	2	0.3	214.5	213	216	2	1	0.1	541.5
Rosie Vam Dyke Company	21	3.3	162.1	121	211	19	11	1.3	506.3
Sanaark Investment	1	0.2	149.0	149	149	0	0	0.0	196.0
Sasini (K)	28	4.4	275.1	129	376	43	48	6.1	1715.8
Sondhi Trading Co.	9	1.4	204.2	180	228	17	5	0.6	535.2
Super Gibs	1	0.2	142.0	142	142	0	0	0.0	390.0
Taylor Winch (Coffee)	135	21.3	242.3	155	290	17	190	24.0	1409.4

Figure 2 represents the buyer annotated time series of the auction prices on the 13th Aug. 2024, with the prices paid by the three largest buyers marked. The price volatility evident in their purchase data – whether the targeted consistency of Taylor Winch or the extreme range of Louis Dreyfus – creates significant input cost risk for their downstream operations. The academic framework for managing such procurement risk favors derivatives that hedge against adverse average cost movements. For instance, an *Asian call* option would be appropriate for a buyer like Ibero Kenya or Louis Dreyfus, as it provides a payoff if the average price of coffee over the next week's auctions exceeds a strike level, effectively capping their average cost of goods sold (Hull and Basu, 2016). This would be suitable than hedging a single-day price for entities making repeated purchases. Alternatively, a *call gap* option could be strategically employed by a quality-focused buyer like Taylor Winch. This instrument would activate a hedge if prices rise above a certain threshold, protecting them from cost spikes on the premium lots they systematically target, while avoiding the cost of hedging for smaller, anticipated price movements (Zhang, 1997). The extreme volatility seen in buyers like Sasini (range 129-376, std dev 43.44) further underscores this need; such wide purchase prices translate directly into unpredictable production costs and compressed margins if not managed.

The bifurcation in buyer strategies – evidenced by the stark contrast between Ibero's volume-

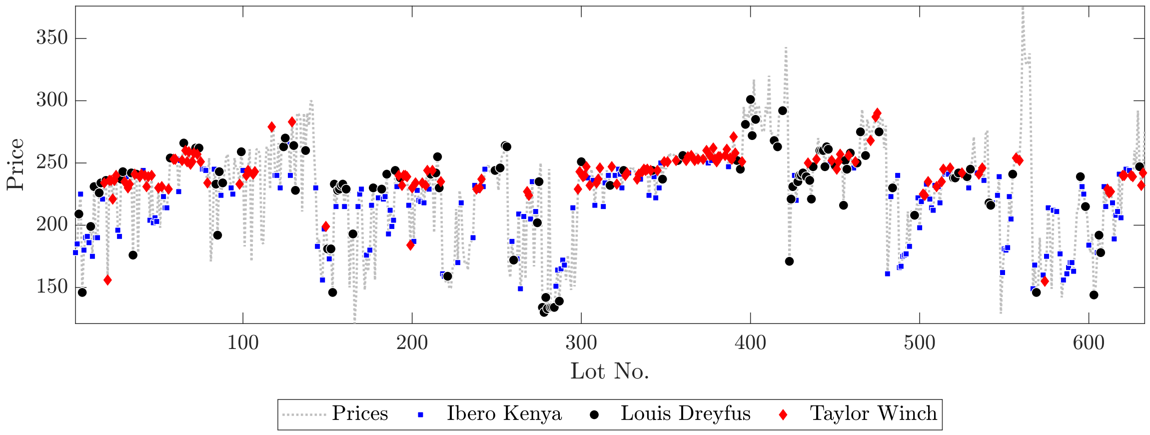

Figure 2: Buyer Prices

Notes: The solid gray line represents the price as the auction progresses through the day. The winning bids of the three largest buyers are each marked by a square, circle and diamond.

driven lower average and Taylor Winch's quality-driven higher average – creates a complex auction dynamic where basis risk is inherent. As Stulz (2022) notes, the primary objective of corporate risk management is to eliminate costly lower-tail outcomes that can jeopardize financial planning. Therefore, implementing a structured hedging program using average-price or gap options would allow these major buyers to secure a ceiling on their weekly average procurement costs. This transforms the auction's price volatility from a threat to operational stability into a managed variable, enabling firms to compete on blending, marketing, and supply chain efficiency rather than being exposed to the raw uncertainty of the auction floor.

2.4 Seller Side Data

Sellers information at the NCE is available through their catalogues. Every week, before the auction, sellers submit a detailed description of the lots of coffee they will have on sale. The catalogues includes standard identifying information such as seller name and date of intended auction/sale. Since the product is sold in lots, with each lot potentially being from a different region, harvest area and/or mill, among other characteristics, the sales are organized per lot, each of which presents product of different quality/grade and quantity. For each lot, a unique ID is issued by the "warehouse-men" and a descriptive "mark" added (and also recorded on 2.5kg sample bags). Finally grade, quantities on offer and packaging information is included.

Table 2 gives the seller(agent) distribution of prices and quantities on the same day. The distribution reveals a market structure with pronounced concentration and significant price dispersion. Three sellers – Alliance Berries (26% share), Meru County Coffee Marketing Agency (20%), and New Kenya Planters Co-operative Union PLC. (19%) – collectively dominate nearly two-thirds of the auction volume, positioning them as price-setters with considerable market influence. Notably, their price volatility profiles differ substantially. Alliance Berries, despite its commanding volume, exhibited a moderate standard deviation of 32.76, with prices ranging from \$146 to \$300, suggesting a diversified portfolio that averages out some risk. However, Meru County displayed greater consistency within a higher band (range 215-343, std dev 21.52), while New KPCU presented the most

Table 2: Price Distribution by Seller, 13th August 2024

Agent	Sales	%	Price per Sale, $ Mean	Min	Max	STD	Quantity per Sale Wgt (tons)	%	Mean
Alliance Berries	165	26.1	230.3	146	300	33	190	24.0	1154.1
Baringo Kawa Brokerage	6	0.9	185.2	121	229	45	2	0.2	251.0
KCCE Marketing Agency	64	10.1	206.7	148	255	31	46	5.9	726.3
Kiambu Coffee Marketing	21	3.3	240.3	190	264	16	19	2.3	881.4
Kipkelion Broker Company	41	6.5	175.6	130	250	34	26	3.3	630.9
Meru County Coffee M.A.	125	19.7	255.7	215	343	22	241	30.4	1931.0
Minnesota Coffee Marketers	65	10.3	249.9	161	293	25	126	15.8	1934.2
New Kenya Planters Co-op.	121	19.1	215.0	129	376	44	125	15.7	1031.9
United Eastern Kenya Coffee M.A	24	3.8	233.8	189	292	21	18	2.2	740.0

pronounced volatility in the entire dataset, with an extreme range from 129 to 376 and a high standard deviation of 43.58. This latter profile indicates a highly heterogeneous product mix, where the average price of $214.98 masks exposure to severe downside price swings alongside premium upside potential.

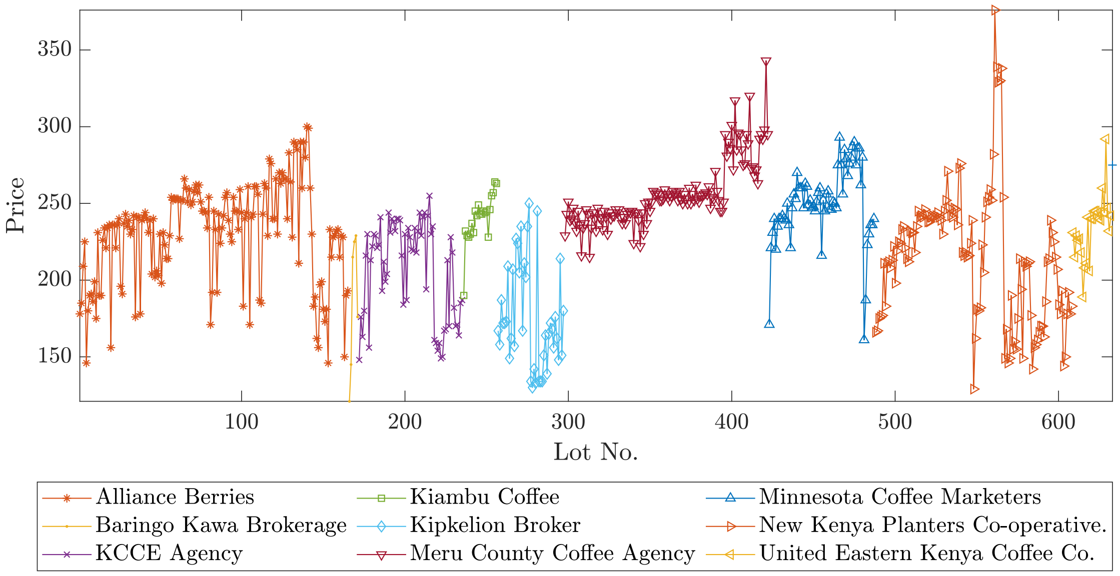

Figure 3: Seller Prices

Notes: The prices are sequential as the auction progresses through the day and each agent controls their selling process.

For these volume-dominant sellers, managing the inherent price volatility risk exposed in this auction data is critical for revenue stability. The academic literature on commodity risk management strongly advocates for the use of derivative instruments tailored to such auction-based, periodic cash flows. For instance, a *put gap options* could be strategically employed to establish a floor price while retaining upside participation above a higher strike. This would be suitable for a seller like New KPCU wishing to lock in a minimum acceptable price for its lower-quality lots while preserving the windfall potential from its occasional premium lots (Zhang, 1997). Figure 3 gives the time evolution of prices by seller through the auction on 24th August 2024. The need for such hedging is underscored by the stark price differentials observed between sellers, which can be attributed to unanticipated quality variations, weather impacts on bean grade, and shifting global demand – all classic sources of basis risk. As Clark (2014) argues, effective hedging mitigates these revenue uncertainties, allowing agricultural marketing agencies to focus on operational efficiency and quality improvement. Therefore, for the leading sellers in the Nairobi auction, deploying Asian put options to secure a floor for their weekly average price, or a carefully structured gap option strategy, would be a prudent financial strategy to transform the volatile price distributions evident on August 13th into predictable and stable revenue streams for the subsequent sales period.

3 Price Processes

We motivate how to model each buyer's private value for a lot of coffee as a log-normal random variable. The winning price is the maximum of these values. When we track these maximum values over successive auctions, their logarithmic changes behave like a random walk with drift – this is the GBM. The *drift* measures the average upward pressure from competition, while the *volatility* captures the uncertainty in how much the highest bidder is willing to pay. We elaborate this modelling framework of auction price time series from the buyer perspective. A similar analysis for the seller is included in Appendix A. We then discuss how we estimate the GBM parameters using maximum likelihood methods. We give results of the estimation as well as diagnostics at the end of this section.

3.1 Buyer Price Process

Following Paarsch, Hong and Haley (2006), we model the behaviour of individual bidders (buyers) within the independent private values paradigm. We assume that the logarithm of the valuation of a potential bidder $\log V$ is distributed normally with mean μ and variance σ^2. There is a sequence of auctions at times $t = 1,\ldots,T$, each with N_t potential bidders. We assume that the data recorded are the winning bid W_t and the number of participants N_t.

Let $V_{(1:N_t)} > V_{(2:N_t)} > \ldots V_{(N_t:N_t)}$ be the ordering of individual buyers' unobserved valuations. Assuming further that the valuations are independent and identically distributed, the cumulative distribution function of the winning(highest) bid $Y_t = V_{(1:N_t)}$ is

$$F_Y(y|N_t) = \Phi\left(\frac{\log y - \mu}{\sigma}\right)^{N_t}$$

where $\Phi(\cdot)$ is the cumulative normal distribution function. Since the Gaussian random variable $\log V_{(1:N_t)}$ lives within the location-scale family of random variables, we have that

$$\log V_{(i:N_t)} = \mu + \sigma Z_{(i:N_t)}$$

where $Z_{(i:N_t)}$ is the i^{th} largest order statistic from a sample of N_t independent identically distributed standard normal random variables. It then follows that for the highest order statistic Y_t, we have $E\left[\log Y_t|N_t\right] = \mu + \sigma E\left(Z_{(1:N_t)}\right)$, where

$$E\left(Z_{(1:N_t)}\right) = \mu + \sigma \int_{-\infty}^{+\infty} z N_t \Phi(z)^{N_t - 1}\phi(z)dz$$

and $\phi(z) = \frac{d}{dz}\Phi(z)$. Exploiting the location scale property we then have that $\log Y_t = \mu + \sigma E\left(Z_{(1:N_t)}\right) + U_t$, where $E(U_t|N_t) = 0$ and the variance of U_t depends on N_t. Over the small time interval Δt we have that

$$\log Y_{t+\Delta t} - \log Y_t = \left(\mu + \sigma E\left(Z_{(1:N_t)}\right)\right)\Delta t + \sigma_U\sqrt{\Delta t}U_{t+\Delta t}.$$

Taking limits as $\Delta t \downarrow 0$, we obtain the geometric Brownian motion

$$d\log Y_t = \bar{\mu}dt + \sigma_U dW_t \tag{1}$$

where $\bar{\mu} = \mu + \sigma E\left(Z_{(1:N_t)}\right)$ and $dW_t = \sqrt{dt}U_t$ is an increment of Brownian motion – a standard normal variable with mean 0 and variance dt.

There are several limitations to this approach. First, the assumption of constant drift and volatility may not fully capture auction dynamics, particularly jumps between lots and session effects. Second, market microstructure features such as discrete price movements, bid-ask bounce, and strategic interactions may induce deviations from continuous-time diffusion assumptions. Finally, we assume that the observation intervals are equal, but in actual inter-transaction times vary by several seconds.

3.2 Maximum Likelihood Estimation

We employ maximum likelihood estimation (MLE) to calibrate the GBM parameters for buyer price processes at the NCE. This approach provides statistically efficient parameter estimates under the assumption that observed auction prices follow the posited stochastic process. We discuss this estimation approach next.

Let $\{S_{t_0}, S_{t_1}, \ldots, S_{t_n}\}$ represent the sequence of winning bid prices for a particular buyer, observed at auction times $t_0 < t_1 < \cdots < t_n$. Logarithmic returns are given by $r_i = \Delta\log S_{t_i}$, $i = 1, \ldots, n$ which follow a normal distribution with linear drift. Under the assumption of constant drift μ and volatility σ, the discrete-time evolution of log prices follows:

$$\log S_{t_i} = \log S_{t_{i-1}} + \left(\mu - \frac{1}{2}\sigma^2\right)\Delta t_i + \sigma\sqrt{\Delta t_i}Z_i$$

where $Z_i \sim N(0,1)$ are independent standard normal innovations. The logarithmic returns are normally distributed as:

$$r_i \sim N\left[\left(\mu - \frac{1}{2}\sigma^2\right)\Delta t_i, \sigma^2\Delta t_i\right].$$

The negative log-likelihood function for the observed return sequence, conditional on the time intervals $\Delta t_i = t_i - t_{i-1}$, is:

$$\ell(\mu, \sigma^2) = -\frac{n}{2}\log(2\pi) - \frac{1}{2}\sum_{i=1}^{n}\log(\sigma^2\Delta t_i) - \frac{1}{2\sigma^2}\sum_{i=1}^{n}\frac{(r_i - (\mu - \frac{1}{2}\sigma^2)\Delta t_i)^2}{\Delta t_i}.$$

Minimizing the log-likelihood yields two interdependent equations which can be solved simultaneously for the maximum likelihood estimators:

$$\hat{\mu} = \frac{1}{T}\sum_{i=1}^{n}\frac{r_i}{\Delta t_i} + \frac{1}{2}\hat{\sigma}^2 \text{ and } \hat{\sigma}^2 = \frac{1}{T}\sum_{i=1}^{n}\frac{(r_i - (\hat{\mu} - \frac{1}{2}\hat{\sigma}^2)\Delta t_i)^2}{\Delta t_i} \tag{2}$$

where $T = \sum_{i=1}^{n}\Delta t_i$ represents the total observation period in years. These equations are interdependent and solved through numerical iteration. The estimation implemented using the following iterative algorithm: **initialization** using the method-of-moments starting values $\hat{\sigma}_0^2$, , $\hat{\mu}_0$, and **iteration** using the updating values $\hat{\sigma}_{k+1}^2$, $\hat{\mu}_{k+1}$ for $k = 0, 1, 2, \ldots$ until convergence with the criterion that stops when $\|\hat{\mu}_{k+1} - \hat{\mu}_k\| < \varepsilon$ and $\|\hat{\sigma}_{k+1}^2 - \hat{\sigma}_k^2\| < \varepsilon = 10^{-6}$. The initial values are given by eq. (3)

$$\hat{\sigma}_0^2 = \frac{1}{T}\sum_{i=1}^{n}\frac{(r_i - \bar{r})^2}{\Delta t_i}, \text{ where } \bar{r} = \frac{1}{n}\sum_{i=1}^{n} r_i\,, \quad \hat{\mu}_0 = \frac{1}{T}\sum_{i=1}^{n}\frac{r_i}{\Delta t_i} + \frac{1}{2}\hat{\sigma}_0^2 \tag{3}$$

and the iterations performed using the recursion eq. (4)

$$\hat{\sigma}_{k+1}^2 = \frac{1}{T}\sum_{i=1}^{n}\frac{(r_i - (\hat{\mu}_k - \frac{1}{2}\hat{\sigma}_k^2)\Delta t_i)^2}{\Delta t_i}, \quad \hat{\mu}_{k+1} = \frac{1}{T}\sum_{i=1}^{n}\frac{r_i}{\Delta t_i} + \frac{1}{2}\hat{\sigma}_{k+1}^2. \tag{4}$$

3.3 Parameter Estimates

We first estimate parameters of eq. (1) using the initial values eq. (3) and the iterative algorithm (4). We do this assuming a single time series with $N = 633$ prices. While the auction mechanism produces irregularly spaced observations, with time intervals $\Delta t_i = t_i - t_{i-1}$ varying based on bidding activity and lot characteristics. Based on previous timestamped data, we approximate that a single "*sale*" takes half a minute. There were 42 trading days per year during the 2023–2024 "season" of the NCE. Given 633 transactions in our data and half a minute per transaction means we have approximately 316 minutes per day multiplied by 42 trading days. The time increment for a single interval in units of years is the ratio of the interval length to the total annual trading time in minutes: $\Delta t = \frac{1}{42\times 316} \approx 7.5347e^{-05}$ and $T = \sum_{i=1}^{n}\Delta t_i = 0.0024$ years.

Table 3: Maximum Likelihood Estimates Buyer Price Process

Parameter	Estimate	Std.Error
μ	1.0804	0.0
σ	1.4652	0.0
Log-Likelihood	42.92	–
AIC	−81.85	
BIC	−72.95	
Pseudo R^2	0.5342	
Normality (JB - p-value)	0.0418	

Table 3 gives a summary the parameter estimates and a selection of goodness-of-fit measures. The two parameters are precisely estimated with standard errors ≈ 0. The pseudo $R^2 \approx 53\%$ meaning the model provides a good fit for the data. Figure 4 shows the MLE estimation diagnostics. We assess model fit through several diagnostic measures. We look at standardized residuals , perform an autocorrelation test, and check for distributional fit using the Kolmogorov-Smirnov statistic comparing $\{\hat{z}_i\}$ to the standard normal density.[3] The residuals are approximately normal according to the Jarque-Berra test and their empirical distribution. However, there is autocorrelation present, a feature that is typical of high frequency price time series. This could be induced by microstructure effects such as the bid-ask bounce or the the price impact of a trade. The confidence intervals are squeezed to the parameter estimates in the last panel of Figure 4 since the standard errors are ≈ 0.

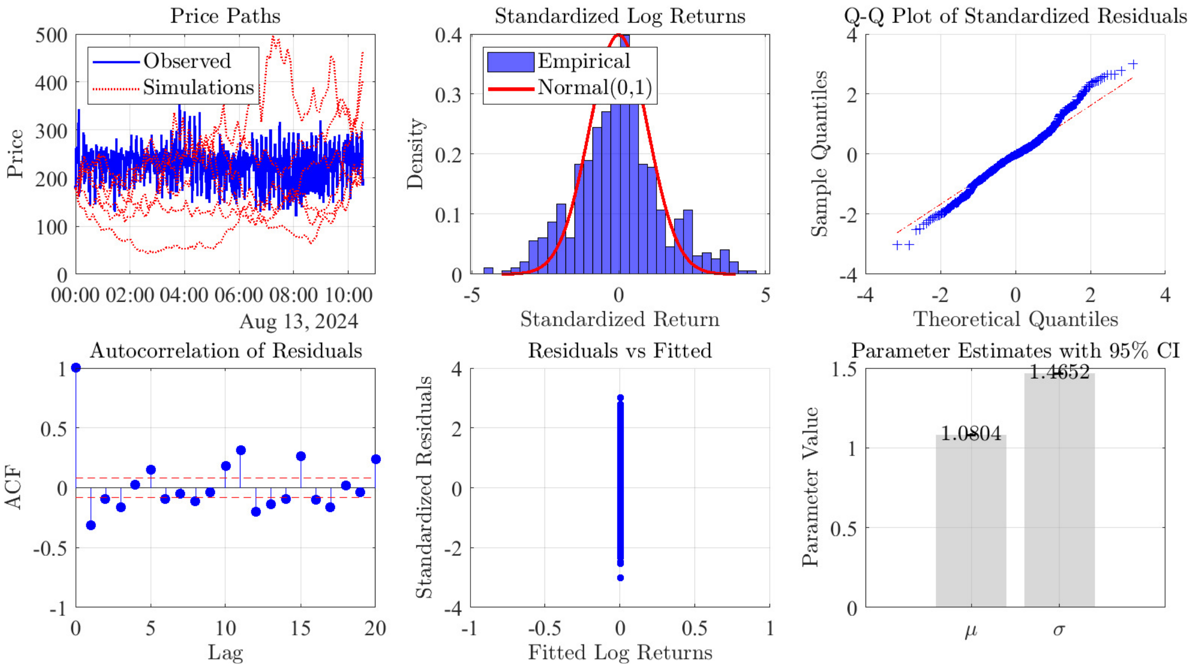

Figure 4: Maximum Likelihood Estimation Diagnostics

4 Hedging with Gap Call Options

This section discusses the risk management strategies that buyers/sellers can use to hedge against price variation. The analysis is purely conceptual as there are no actual derivative markets where options written on the NCE spot price are traded. We consider *gap call* options for buyers, and assume sellers at the auction take opposite sides of a bet. Given that buyers and sellers are on opposite sides of each transaction that generates the price processes, we do not include an analysis of the gains for sellers, since it would simply be a reversion of the buyer side gains/losses – with the exception that each seller would have a unique set of parameters driving the price process.[4] We consider the following strategy for buyers(sellers): we assume that there exists a gap put option that gives the right to buy(sell) output at the strike price K if the realized sale price in the next auction is higher(less) than

[3] The standardized residuals are computed as $\hat{z}_i = [r_i - (\hat{\mu} - \hat{\sigma}^2/2)\Delta t_i]/\hat{\sigma}\sqrt{\Delta t_i} \sim N(0,1)$. THe autocorrelation test uses the statistic $Q(m) = n(n+2)\sum_{k=1}^{m} \hat{\rho}_k^2/(n-k) \sim \chi^2(m)$, where $\hat{\rho}_k$ is the sample autocorrelation of $\{\hat{z}_i\}$ at lag k, and m is chosen as $\min(20, n/4)$.

[4] The assumption of unique parameter values for sellers can be motivated using arguments similar to to the buyer price process derivation in Section 3. See Appendix A for details.

H. We assume that buyers(sellers) take the opposite of these agreements, that is, sell call options with trigger price H.

We analyze gap call options that give the buyer the right, but not the obligation, to purchase coffee at strike price K if the terminal price S_T exceeds trigger price H. The payoff function is:

$$\text{Payoff} = \max(S_T - K, 0) \cdot \mathbf{1}_{\{S_T > H\}} \tag{5}$$

where $\mathbf{1}_{\{\cdot\}}$ is the indicator function. This structure creates a binary condition: the option provides value only if $S_T > \max(K, H)$. Option premiums are calculated using the Black-Scholes-Merton closed-form solution for gap options:

$$C = S_0 N(d_1) - Ke^{-rT} N(d_2) \tag{6}$$

with $d_1 = \frac{\ln(S_0/H)+(r+\sigma^2/2)T}{\sigma\sqrt{T}}$ and $d_2 = d_1 - \sigma\sqrt{T}$, where $N(\cdot)$ represents the cumulative standard normal distribution function. This framework assumes complete markets, no arbitrage, and continuous hedging possibilities.

4.1 Simulation and Performance Analysis

We use Monte Carlo simulations to model the stochastic evolution of coffee prices at the NCE. The simulation parameters are calibrated using transaction data from the NCE auction on August 13$^{\text{th}}$, 2024 as summarized in Table 3. The initial price $S_0 = \$275$ represents the closing price from that auction day. The annualized volatility $\sigma = 146.5\%$ is estimated from intraday price movements, capturing the extreme fluctuations of the action prices. A risk-free rate of $r = 4\%$, approximately equal to the annualized 90-day T-bill rate in Kenya, is used for the risk-neutral pricing. We generate $N =$ 10,000 independent price paths using the Euler-Maruyama discretization scheme of the stochastic differential equation (1):

$$S_{t+\Delta t} = S_t \exp\left[\left(\mu - \frac{1}{2}\sigma^2\right)\Delta t + \sigma\sqrt{\Delta t} Z_t\right] \tag{7}$$

where $Z_t \sim N(0,1)$ are independent standard normal random variables, $\Delta t = 1/252$ represents one trading day, and the total simulation horizon $T = 3/52$ years (three weeks). The simulation employs antithetic variates to reduce variance, generating pairs of paths with negatively correlated random shocks to improve statistical efficiency.

We evaluate option performance using multiple complementary metrics supplemented with downside risk assessment. We compute the expected profit & loss and its standard deviation, the Sharpe ratio, as well as the probabilities of profit and exercising the option. We evaluate riskiness by computing the value at risk and conditional value at risk. Table 4 gives summarizes these measures. We conduct systematic sensitivity analysis by varying the trigger-to-strike ratio H/K from 1.0 to 1.8 in increments of 0.1. This range encompasses scenarios from standard calls ($H = K$) to highly conditional gap calls requiring substantial price appreciation. For each H/K ratio, we compute the complete set of performance and risk metrics. We evaluate optimal trigger selection under three criteria: Sharpe ratio maximization to identify trigger levels providing optimal risk-adjusted returns, return maxi-

Table 4: Performance Measures Formulae

Description	Measure	Formulae
Performance		
Expected Profit and Loss (P&L)	$\mathrm{E}[P\&L]$	$N^{-1}\sum_{i=1}^{N}[\max(S_T^{(i)}-K,0)\cdot\mathbf{1}_{\{S_T^{(i)}>H\}}-\text{Premium}]$
Standard Deviation of P&L	$\sigma_{P\&L}$	$\sqrt{(N-1)^{-1}\sum_{i=1}^{N}(P\&L_i-\overline{P\&L})^2}$
Sharpe Ratio	Sharpe	$\mathrm{E}[P\&L]/\sigma_{P\&L}$
Profit Probability	P_{profit}	$N^{-1}\#\{P\&L_i>0\}\times 100\%$
Exercise Probability	P_{exercise}	$N^{-1}(\#\{S_T^{(i)}>H \text{ and } S_T^{(i)}>K\})\times 100\%$
Riskiness		
Value at Risk (95%)	$\mathrm{VaR}_{95\%}$	$-\inf\{x:P(P\&L\le x)>0.05\}$
Conditional Value at Risk (95%)	$\mathrm{CVaR}_{95\%}$	$-\mathrm{E}[P\&L \mid P\&L\le -\mathrm{VaR}_{95\%}]$

mization without explicit risk adjustment in terms of expected P&L, and target profit probability to determine trigger levels achieving specified success rates (e.g., 50% probability of profit). We compare gap option performance against standard European calls with identical strike prices, focusing on premium differentials, exercise probability trade-offs, and risk-return profile differences. This framework provides robust insights into the viability of gap options for risk management at the NCE, while documenting methodological choices and their implications for results interpretation.

4.2 Buyer-Side Hedging Performance

We first look at the performance of a plain vanilla call. The standard at-the-money call option, with a moneyness S_0/K of 110%, carries a premium of \$40.09. As an outright hedging instrument, it demonstrates the following characteristics. The expected profit and loss is a positive \$10.58, indicating that on average the option's payoff exceeds its cost. This expected gain comes with high uncertainty, given by the substantial P&L standard deviation of \$77.79, resulting in a low Sharpe ratio of 0.136. The option has a 54% probability of expiring in-the-money, yet the probability of the position being profitable is only 37.4%. This discrepancy highlights the cost of the premium. The gap call introduces a trigger $H > K$ where the option only pays off $S_T - K$ if $S_T > H$. This conditional payoff structure reduces the probability of exercise, thereby lowering the upfront premium. For $H = \$300$, the premium drops to \$33.16, which is 17.3% cheaper than the standard call. A key finding is that this gap call generates a higher expected P&L of \$12.62 and a superior Sharpe ratio of 0.158 compared to the standard call. This demonstrates that the premium savings can more than compensate for the forgone payoffs in scenarios where S_T is between K and H. The trade-off is a reduced probability of profit, at 33.7%, and exercise, also at 33.7%.

Figure 5 summarizes a sensitivity analysis across trigger levels, with H/K from 1.0 to 1.8, showing the spectrum of hedging strategies. Lower triggers, such as $H = \$275$, approximate standard call behavior with a near-50% profit probability. Analysis identifies optimal triggers for specific metrics: maximum expected P&L is achieved at $H = \$375$ with an expected P&L of \$15.79, while maximum risk-adjusted return, measured by the Sharpe ratio, is achieved at $H = \$400$ with a Sharpe ratio of

0.203. As H increases, the premium declines monotonically, but the exercise probability drops precipitously. The expected P&L first increases, peaks, and then decreases, illustrating an optimal region around $H/K = 1.5$. The distribution of returns becomes extremely positively skewed and leptokurtic at high triggers, indicating a strategy with frequent small losses, equal to the premium, and rare but extremely large gains.

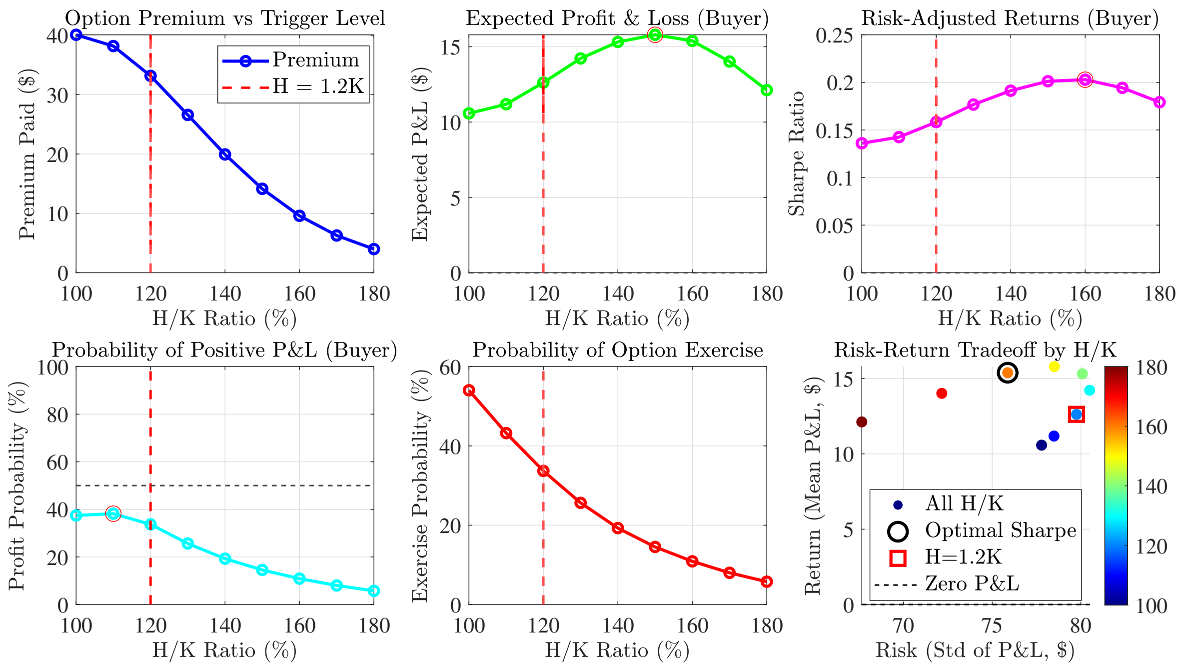

Figure 5: Gap call sensitivity analysis

The gap call significantly alters the risk profile. The maximum loss is capped at the premium paid, which decreases with H. The 95% Value-at-Risk and Conditional VaR are exactly equal to the premium, indicating that the worst 5% of outcomes constitute a total loss of the premium. The probability of losing more than 50% of the premium, which is equivalent to the probability of the option expiring worthless, increases dramatically with H, from 54.8% for $H = \$250$ to 94.2% for $H = \$450$. Conversely, the strategy offers a non-zero probability of very high percentage returns. From a commodity buyer's perspective, when employing the $H = \$300$ gap call, the buyer's average effective purchase price across all simulated paths is \$240.86, representing average savings of \$34.61 compared to the average unhedged market price of \$275.48. The hedge is effective in 33.7% of scenarios, where significant savings are realized. In all other scenarios, the buyer simply loses the premium and purchases at the market price. The break-even market price for the strategy is \$283.16; profitability requires the market price to exceed both this level and the \$300 trigger. The gap call is not a hedge against moderate price increases but a cost-effective tool for catastrophic insurance against extreme price spikes. It transforms the buyer's cost structure, securing a maximum price of $K + \text{Premium} = \$283.16$ in most scenarios, while providing substantial upside relief if prices surge beyond H.

For the volatile commodity environment specified, the gap call structure dominates the standard call on a risk-adjusted basis across a range of trigger prices. By accepting a lower probability of payoff, the hedger can reduce upfront hedging costs through a lower premium outlay, increase the expected economic profit of the hedging program, improve the Sharpe ratio of the hedging position, and tailor the hedge to specific risk tolerance. A buyer more concerned with extreme tail risk can opt for a higher H to maximize the Sharpe ratio, while one seeking a balance can choose a lower H for a higher probability of a positive result. In conclusion, the gap call option presents a superior and flexible hedging alternative for a commodity buyer in high-volatility regimes. It efficiently replaces expensive continuous protection with a contingent, leveraged payoff against severe adverse price movements, thereby optimizing the cost-benefit profile of the volatility risk management program.

4.3 Practical Implications

Buyers at the NCE can derive meaningful risk management benefits from gap call options with appropriate parameter selection. Strikes near current market prices ($K \approx \$250 - 275$) provide optimal balance between premium cost and protection value. Higher trigger levels ($H = 1.6 - 1.8K$) maximize risk-adjusted returns by reducing premiums while maintaining substantial upside potential. Gap calls function as "catastrophic insurance" against extreme price spikes rather than protection against normal market movements. However, such options should complement rather than replace other risk management strategies, given their binary payoff characteristics.

These findings suggest that gap call options could provide viable risk management tools for NCE buyers, particularly as protection against extreme price spikes. However, their effectiveness depends critically on appropriate parameter selection and should be evaluated in the context of the exchange's extreme volatility characteristics. For the NCE, developing a derivatives market based on gap options would require careful consideration of product design, participant education, and risk management infrastructure. The analysis provides a foundation for such development while highlighting the importance of continued research on auction market dynamics and derivative pricing in emerging commodity exchanges.

5 Conclusion

Sustainability and long term profitability are chronic challenges to commercial agriculture in many countries. In this study, we explored the possibility of using a simple risk management strategy to mitigate adverse price movements at the Nairobi Coffee Exchange. We modeled the once weekly auction prices from the NCE using a geometric Brownian motion and used it to price gap put options. The analysis showed that buyers would make meaningful gains in revenue and profits if there was a market for financial derivatives written on the NCE auction spot prices. However, these findings may be largely driven by the specific dates chosen and any other set of different auction days (weeks) could give other results. While we have not included analysis for the seller side, our framework shows how one could model the seller side parameters and engage in a similar risk–management exercise as we have demonstrated in the case of buyers. Our analysis contributes to the better understanding of the less studied commodity markets in the developing world. In many cases, most of these spot markets are neither well known nor studied. We hope that our analysis contributes to the eventual development of derivative securities with the NCE spot price as underlying security.

The analysis suggests that there is potential for developing viable options markets at the NCE.

There are several limitations of this study that warrant consideration. The 146.5% annual volatility estimate, while empirically derived from auction data, may incorporate microstructure noise, outliers, or period-specific effects. More sophisticated volatility estimation techniques and longer time series would improve reliability. Secondly, the Geometric Brownian Motion assumption may inadequately capture auction price dynamics, particularly jumps between lots and session effects. Alternative models incorporating auction-specific features could potentially provide more accurate pricing. Future research should address these limitations through multi-period analysis, alternative stochastic processes, incorporation of market frictions, and investigation of actual participant behavior and preferences at the NCE.

We have not included any evaluation of the effectiveness of using the futures contracts from the ICE to hedge spot price volatility for individual buyers/sellers at the NCE. Since there are no futures contracts specifically traded using the NCE spot price as the underlying, such an evaluation would follow the “cross hedging”[5] approaches elaborated by Ankirchner, Imkeller and Reis (2010); Ankirchner, Dimitroff, Heyne and Pigorsch (2012); Lioui, Trong and Poncet (1996) among others. We leave such an evaluation to future work.

[5]When an investor uses a liquid traded asset to hedge the risk of a correlated but non traded or illiquid underlying asset, they are cross–hedging (Basak and Chabakauri, 2012).

References

Ankirchner, S., Dimitroff, G., Heyne, G. and Pigorsch, C.: 2012, Futures cross-hedging with a stationary basis, *Journal of Financial and Quantitative Analysis* **47**(6), 1361–1395. (cited on page 3, 17)

Ankirchner, S., Imkeller, P. and Reis, G. D.: 2010, Pricing and hedging of derivatives based on nontradable underlyings, *Mathematical Finance* **20**(2), 289–312.
URL: *https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1467-9965.2010.00398.x* (cited on page 3, 17)

Basak, S. and Chabakauri, G.: 2012, Dynamic Hedging in Incomplete Markets: A Simple Solution, *The Review of Financial Studies* **25**(6), 1845–1896.
URL: *https://doi.org/10.1093/rfs/hhs050* (cited on page 17)

Bush, S. B.: 2012, *Coffee, Derivatives, and Income Security: Theory and Practice*, Palgrave Macmillan US, New York, pp. 51–100.
URL: *https://doi.org/10.1057/9781137062659_3* (cited on page 2)

Claessens, S. and Duncan, R. C.: 1994, *Managing Commodity Price Risk in Developing Countries*, The World Bank.
URL: *https://elibrary.worldbank.org/doi/abs/10.1596/0-8018-4662-5* (cited on page 3)

Clark, I. J.: 2014, *Commodity option pricing: a practitioner's guide*, John Wiley & Sons. (cited on page 9)

Donald, S. G. and Paarsch, H. J.: 1996, Identification, estimation, and testing in parametric empirical models of auctions within the independent private values paradigm, *Econometric Theory* **12**(3), 517–567. (cited on page 3)

Hollstein, F., Prokopczuk, M. and Würsig, C.: 2020, Volatility term structures in commodity markets, *Journal of Futures Markets* **40**(4), 527–555.
URL: *https://onlinelibrary.wiley.com/doi/abs/10.1002/fut.22083* (cited on page 2)

Hull, J. C. and Basu, S.: 2016, *Options, futures, and other derivatives*, Pearson Education India. (cited on page 2, 3, 6)

Lioui, A., Trong, P. N. D. and Poncet, P.: 1996, Optimal dynamic hedging in incomplete futures markets, *The Geneva Papers on Risk and Insurance Theory* **21**(1), 103–122.
URL: *http://www.jstor.org/stable/41953328* (cited on page 17)

Mohan, S.: 2007, Market-based Price-risk Management for Coffee Producers, *Development Policy Review* **25**(3), 333–354.
URL: *https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1467-7679.2007.00373.x* (cited on page 3)

Mussatto, S. I., Machado, E. M., Martins, S. and Teixeira, J. A.: 2011, Production, composition, and application of coffee and its industrial residues, *Food and Bioprocess Technology* **4**, 661–672. (cited on page 2)

Myers, R. J.: 1993, *Strategies for managing coffee price risks in Costa Rica.*, Johns Hopkins University Press, pp. 157–184. (cited on page 2)

Paarsch, H., Hong, H. and Haley, M.: 2006, *An Introduction to the Structural Econometrics of Auction Data*, Mit Press, MA.
URL: *https://books.google.co.ke/books?id=JusVAQAAMAAJ* (cited on page 3, 9)

Pancsira, J.: 2022, International coffee trade: a literature review, *Journal of Agricultural Informatics* **13**(1).
URL: *https://magisz.org/journal/index.php/jai/article/view/654* (cited on page 2)

Stulz, R. M.: 2022, Rethinking risk management, *Journal of Applied Corporate Finance* **34**(1), 32–46. (cited on page 7)

Veld-Merkoulova, Y. V. and de Roon, F. A.: 2003, Hedging long-term commodity risk, *Journal of Futures Markets* **23**(2), 109–133.
URL: *https://onlinelibrary.wiley.com/doi/abs/10.1002/fut.10060* (cited on page 2)

Zhang, P. G.: 1997, *Exotic options: a guide to second generation options*, World scientific. (cited on page 3, 6, 8)

A Seller Price Process

A.1 Seller Price Process

Figure 3 shows variation of prices across different sellers. The marks on the legend identify different sellers (product from a single marketer or producer). At each "sell session" prices seem to start low and progressively increase to a peak approximately \$300, then tumble and stay low through the session. While Figure 3 reveals distinctive price patterns across different sellers, we can formalize the seller's perspective using the same independent private values paradigm applied to buyers. Consider seller i offering a sequence of coffee lots throughout the auction day, with each lot potentially attracting a different set of buyers.

Let seller i face a sequence of auctions at times $t = 1,\ldots,T$, where each auction corresponds to a distinct lot sold to different buyers. In auction t, there are $N_{i,t}$ potential buyers with independent private valuations $V_{j,t}^{(i)}$ $(j = 1,\ldots,N_{i,t})$ for seller i's coffee. We assume buyer valuations for seller i's product follow a log-normal distribution:

$$\log V_{j,t}^{(i)} \sim N(\mu_i, \sigma_i^2), \quad \text{for } j = 1,\ldots,N_{i,t}$$

where μ_i captures the average quality perception of seller i's coffee, σ_i^2 represents heterogeneity in buyer valuations, and $N_{i,t}$ is the number of interested buyers at auction t. The seller observes the winning bid $W_{i,t}$ in each auction, which equals the maximum valuation among all potential buyers:

$$W_{i,t} = \max\{V_{1,t}^{(i)}, V_{2,t}^{(i)}, \ldots, V_{N_{i,t},t}^{(i)}\}$$

From order statistics theory, the cumulative distribution function of the winning bid is:

$$F_W(w|N_{i,t}) = \left[\Phi\left(\frac{\log w - \mu_i}{\sigma_i}\right)\right]^{N_{i,t}}$$

where $\Phi(\cdot)$ denotes the standard normal CDF. The expected log-winning bid follows from properties of normal order statistics:

$$\mathbb{E}[\log W_{i,t}|N_{i,t}] = \mu_i + \sigma_i \mathbb{E}[Z_{(1:N_{i,t})}]$$

where $Z_{(1:N_{i,t})}$ is the maximum of $N_{i,t}$ independent standard normal variables. The expected maximum of N i.i.d. standard normal variables can be approximated using Fisher-Tippett extreme value theory. For large N:

$$\mathbb{E}[Z_{(1:N)}] \approx a_N + b_N\gamma, \text{ where } a_N = \sqrt{2\log N} - \frac{\log\log N + \log(4\pi)}{2\sqrt{2\log N}}, b_N = \frac{1}{\sqrt{2\log N}},$$

and, $\gamma \approx 0.5772$ is the Euler-Mascheroni constant. For moderate N (typical in NCE auctions with 5–15 bidders per lot), we can use the exact expectation:

$$\mathbb{E}[Z_{(1:N)}] = \int_{-\infty}^{\infty} z \cdot N\Phi(z)^{N-1}\phi(z)dz \tag{A.1}$$

where $\phi(z) = \Phi'(z)$ is the standard normal density.

Letting $\{S_{i,t}\}$ represent the observed price sequence for seller i. Over a small time interval Δt between auctions, we have:

$$\log S_{i,t+\Delta t} - \log S_{i,t} = \left[\mu_i + \sigma_i \mathbb{E}[Z_{(1:N_{i,t})}]\right] \Delta t + \sigma_i \sqrt{\Delta t} \varepsilon_{i,t}$$

where $\varepsilon_{i,t} \sim N(0,1)$ captures the variation around the expected maximum. Assuming the number of bidders follows a Poisson process with intensity λ_i (reflecting the arrival rate of interested buyers for seller i), we obtain in the limit as $\Delta t \to 0$:

$$d \log S_{i,t} = [\mu_i + \sigma_i g(\lambda_i)]\, dt + \sigma_i dW_t \tag{A.2}$$

where $g(\lambda_i)$ represents the expected maximum for Poisson-arriving bidders and dW_t is a Wiener process. In the special case where there is constant competition such that each seller faces approximately N_i bidders per lot (a stable pool of interested buyers), then

$$d \log S_{i,t} = \bar{\mu}_i dt + \sigma_i dW_t$$

where the seller-specific drift is $\bar{\mu}_i = \mu_i + \sigma_i \mathbb{E}[Z_{(1:N_i)}]$. This yields the Geometric Brownian Motion form:

$$\frac{dS_{i,t}}{S_{i,t}} = \bar{\mu}_i dt + \sigma_i dW_t. \tag{A.3}$$

The structural parameters have clear economic interpretations: (i) *quality component* (μ_i) – Persistent quality differences across sellers' coffee, (ii) *perception heterogeneity* (σ_i) – dispersion in buyer valuations for seller i's product, (iii) *competition premium* ($\sigma_i \mathbb{E}[Z_{(1:N_i)}]$) – additional return earned through auction competition, and (iv) *effective bidders* (N_i) – size and engagement of seller i's buyer pool. The model explains the patterns in Figure 3: sellers with higher μ_i (superior quality) or larger N_i (more buyer competition) achieve steeper price trajectories.

A.2 Maximum Likelihood Estimation

Given transaction data for seller i: $\{S_{i,t_k}\}_{k=1}^{M_i}$ at times $t_1 < t_2 < \cdots < t_{M_i}$, we can estimate parameters via maximum likelihood. The log-likelihood for seller i under the GBM model (A.3) is:

$$\mathscr{L}_i(\bar{\mu}_i, \sigma_i) = -\frac{M_i - 1}{2} \log(2\pi\sigma_i^2) - \sum_{k=2}^{M_i} \frac{\left[\log\left(\frac{S_{i,t_k}}{S_{i,t_{k-1}}}\right) - (\bar{\mu}_i - \frac{1}{2}\sigma_i^2)\Delta t_k\right]^2}{2\sigma_i^2 \Delta t_k} \tag{A.4}$$

where $\Delta t_k = t_k - t_{k-1}$. Letting $T_i = \sum_{k=2}^{M_i} \Delta t_k$, the maximum likelihood estimators are:

$$\hat{\bar{\mu}}_i = \frac{1}{T_i} \sum_{k=2}^{M_i} \frac{\log(S_{i,t_k}/S_{i,t_{k-1}})}{\Delta t_k} + \frac{1}{2}\hat{\sigma}_i^2, \quad \hat{\sigma}_i^2 = \frac{1}{T_i} \sum_{k=2}^{M_i} \frac{\left[\log(S_{i,t_k}/S_{i,t_{k-1}}) - (\hat{\bar{\mu}}_i - \frac{1}{2}\hat{\sigma}_i^2)\Delta t_k\right]^2}{\Delta t_k}. \tag{A.5}$$

To recover the structural parameters μ_i and N_i from $\bar{\mu}_i$, we can exploit the relationship

$$\bar{\mu}_i = \mu_i + \sigma_i \mathbb{E}[Z_{(1:N_i)}] \tag{A.6}$$

and a two-step procedure as follows:

Step 1: Estimate μ_i from minimum bids – For auctions where only one bidder participates ($N_{i,t} = 1$), the winning bid equals the sole valuation. Using these observations:

$$\hat{\mu}_i = \frac{1}{|\mathscr{T}_1|} \sum_{t \in \mathscr{T}_1} \log S_{i,t} \tag{A.7}$$

where $\mathscr{T}_1 = \{t : N_{i,t} = 1\}$.

Step 2: Estimate N_i via method of moments – Using the estimated $\hat{\mu}_i$ and $\hat{\sigma}_i$, solve for N_i:

$$\hat{N}_i = \arg\min_{N \in \mathbb{N}} \left[\hat{\bar{\mu}}_i - \hat{\mu}_i - \hat{\sigma}_i \mathbb{E}[Z_{(1:N)}] \right]^2 \tag{A.8}$$

where $\mathbb{E}[Z_{(1:N)}]$ is computed numerically using (A.1).

University of Warsaw
Faculty of Economic Sciences
44/50 Długa St.
00-241 Warsaw
www.wne.uw.edu.pl
ISSN 2957-0506