Energy Economics 131 (2024) 107351

Contents lists available at ScienceDirect

Energy Economics

14.7 CiteScore 12.8 Impact Factor

journal homepage: www.elsevier.com/locate/eneeco

Valuation of carbon emission allowance options under an open trading phase

Mingyu Fang [a], Ken Seng Tan [b], Tony S. Wirjanto [c,d,*]

[a] RBC Insurance, Toronto, Ontario, Canada
[b] Division of Banking & Finance, Nanyang Technological University, Singapore
[c] Department of Statistics & Actuarial Science, University of Waterloo, Ontario, Canada
[d] School of Accounting & Finance, University of Waterloo, Ontario, Canada

ARTICLE INFO

JEL classification:
C15
G13

Keywords:
Climate change
Carbon emission allowances
Emission derivatives
Carbon pricing
Sustainable finance

ABSTRACT

This paper presents valuation models of emission allowance options under an emission trading scheme, operating in an open trading phase, where unused allowances are banked to subsequent phases without any limit. Empirical studies are carried out to show that allowance option prices exhibit similar volatility smiles to those of the stock market. Three reduced-form econometrics models, namely a Lognormal allowance price model, a Skewness–Kurtosis-Modified Lognormal allowance price model, and a Mixture Lognormal allowance price model are introduced, with each being accorded with a rich interpretation of its own. Numerical illustration of the models is performed through calibration to the European-Union Emission-Trading-Scheme's allowance futures option prices collected for EU ETS Phase 3 and Phase 4 respectively, where statistical fitness of the models is assessed comparatively within each sample and across the two samples collected to ensure robustness of the conclusions.

1. Introduction

This paper presents valuation models of allowance options under an ETS operating in an open trading phase. A phase is defined to be open if the allowances can be banked without any limit into subsequent trading phases, while intra-phase banking and borrowing across compliance years are also permitted. Therefore, under an open trading phase, an emission allowance is a non-expiring investible contract with no upper bound on its price. As a result, for an open trading phase, the valuation of allowance derivatives can appropriately be based on continuous-time reduced-form econometrics models.

Reduced-form econometric models for emission allowance price have been proposed only in a few existing studies in the literature. However, the models presented in most of these studies usually contain a certain structural emission component as, for instance, in Çetin and Verschuere (2009), or they rely on discrete-time approximations as, for instance, in Mnif (2012). Moreover, they are often seen to be lacking formalization on the scope of application, where model implementations are mostly based on EU ETS phase 1, which is a closed trading phase. A closed trading imposes expiration of unused allowances at the end of phase, and thereby induces a dichotomous terminal condition on allowance price that is best captured by structural valuation models as argued in Fang et al. (2022). To the best of our knowledge, the formalized use of continuous-time reduced-form models for allowance option valuation under open trading phases is still somewhat scanty in the literature.

The present study makes an attempt to fill this gap. Our contributions to the literature are threefold. First, we perform an empirical assessment of the stylized facts of allowance prices and allowance option prices using relatively recent market data, with a focus on option volatility smiles. Observations from this exercise reveal several stylized behaviors, such as reverse volatility skews similar to those observed from stock index options. Second, we present three valuation models for allowance options based on continuous-time reduced-form allowance price processes adapted from modern stock option price models. The models have different characteristics and trade-offs in complexity, statistical fitness, as well as the ability to capture stylized facts in market prices. Finally, we calibrate the models to real market option data from two distinct samples collected for EU ETS Phase 3 and Phase 4, as well as perform comparative assessments of their statistical fitness within each sample and across the two samples for the purpose of a robustness check. The two-component Mixture Lognormal model (MLN-2) performs the best in terms of overall fitness across the two samples. While all three models show signs of mispricing the out-of-money options, the Skewness–Kurtosis-Modified Lognormal model (SKM) and MLN-2 models outperform the Lognormal model (LN) and exhibit somewhat opposite mispricing behavior toward the out-of-money region. This leads to the conclusion that there is apparently not

* Correspondence to: University of Waterloo, 200 University Avenue West, Waterloo, Ontario, N2L 3G1, Canada.
E-mail addresses: bmfang@uwaterloo.ca (M. Fang), kenseng.tan@ntu.edu.sg (K.S. Tan), twirjanto@uwaterloo.ca (T.S. Wirjanto).

https://doi.org/10.1016/j.eneco.2024.107351
Received 23 May 2023; Received in revised form 10 January 2024; Accepted 26 January 2024
Available online 1 February 2024

a perfect-for-all model, whose selection should instead be based on the particular application it intends to serve.

2. Basics of an emission trading scheme

Rising awareness of the impacts of climate change and global warming has triggered initiatives designed to control the output of greenhouse gases, commonly measured as carbon dioxide equivalents. Among these efforts, the establishment of an emission trading scheme (ETS) is perhaps the most consequential one. A typical ETS adopts a cap-and-trade strategy by using emission allowances. At the beginning of each compliance year, participating firms receive from the regulators an initial grant of allowance contracts, each of which allows outputting a set unit of carbon dioxide equivalent. At the end of a given year, firms must undergo compliance filing and submit sufficient allowances to cover their emissions over the year, or a non-compliance penalty must be paid on an excess emission. The allowances can be traded freely on the secondary market along with their corresponding futures and option contracts. An ETS comprises several trading phases. At the beginning of each phase, the regulator adjusts the ETS parameters such as the annual grant and penalty levels to reflect lessons from past experience and plans to meet future emission control target. Within each phase, allowances can be banked and borrowed between years.

A formal emission allowance market did not exist until the establishment of an European Union Emission Trading Scheme (EU ETS) in 2005. Phase 1 of EU ETS spanned the years 2005–2007. As a pilot period to train and learn from the market, it was designed to be a closed phase where unused allowances at the end of phase expired with no value. Such a structure forced a dichotomy of terminal allowance at either zero or the penalty level as verified by actual price data. In contrast, the subsequent three EU ETS phases are all open phases characterized by non-expiration of the allowances, under which the allowance price reflects the market's view of the aggregate allowance balance (surplus or shortages vs. emission generated) as well as the expected penalty levels in subsequent trading phases.

Over the past decade or so, this market has grown substantially in size, stability, and liquidity, inspiring the establishment of various global ETS in both developed countries and emerging markets. The majority of the in-force emission markets operate in open trading phases, which include the current EU ETS, Swiss ETS, New Zealand ETS, Korean ETS, and the Chinese national ETS launched by the world's largest emission producing economy. As a result, emission allowance has become an attractive asset class on its own for many institutional investors with a thematic preference toward sustainable carbon-neutral investment portfolios. This is because no other asset classes, so far, is able to provide such a direct carbon offset. Therefore, research on an efficient valuation model for allowance derivatives is called for. To make this study concrete yet manageable in scope, the analyses reference the directives of Phase 3 (2013–2020) and 4 (2021–2030) of the EU ETS.

It is also important to note that for most of the affected industries, the EU ETS has been migrating toward an auction market allocation instead of a free distribution of allowances at the beginning of each compliance year. Although not a defining characteristic of an open trading phase, such a feature is assumed to apply in this paper, since it is a mechanism that most ETS is expected to adopt in the near future to fund the emission-mitigating projects. Without loss generality, an analogy can thus be drawn between an open-phase ETS and the stock market on market tiers; i.e., the auction-based allocation of allowances can be viewed as a primary market dealing or seasoned offering, while the subsequent emission trading activities constitute the secondary market, which is the subject of our study. The proposed valuation approach using reduced-form models has a distinct advantage of being able to conveniently avoid the complexity of capturing the auction market price formation.

Fig. 1. Daily closing EUA spot price from January 2014 to December 2020.

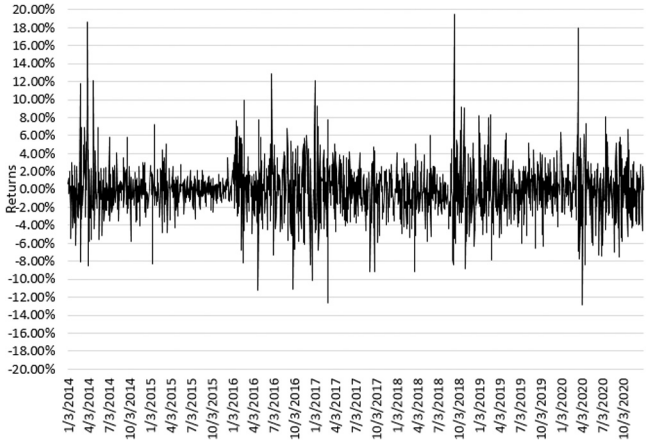

Fig. 2. Daily EUA returns from January 2014 to December 2020.

3. Stylized facts of allowance prices under an open trading phase

Stylized facts of emission allowance prices are examined from EU ETS Phase 3 data. Phase 3 is the most recently completed trading phase spanning the period from 2013 to 2020 in our study. Compared to the previous two trading phases, the phase 3 market is more mature and stable as evidenced from elevated trading volumes and a harmonized set of market disciplines that extends into phase 4, which includes the introduction of a EU-wide emission cap and a wider coverage in sectors and emission gas types. The market price behavior of Phase 3 is rarely studied in the literature, to which our study hopefully is able to provide a fresh contribution.

Fig. 1 shows the daily closing EUA spot prices during the period from January 2014 to December 2020. The daily prices during the period are relatively stable. The lack of price convergence to either zero or the penalty level (100 Euro/tonne) toward the end of the period resonate with the defining characteristic of an open trading phase. The lack of clearly observable discrete jumps indicates the suitability of continuous-time reduced-form econometrics models for the price dynamics in our paper.

Fig. 2 shows the daily log returns of the EUA prices for the same period. Similar to a typical stock price return series, the EUA daily returns display a ubiquitous presence of both volatility clustering and distinct volatility regimes, which calls for consideration of more advanced models to capture these volatility variations.

Summary statistics of the daily allowance returns are presented in Table 1. Although the return data displays approximately zero

Table 1
Summary statistics of daily EUA returns January 2014 to December 2020.

Mean	Median	Variance	Skewness	Excess kurtosis
−0.001	−0.001	0.001	0.491	4.492

sample mean and median, it exhibits a small sample skewness and a large sample excess-kurtosis, which are expected from the observed return volatility clustering behavior over the period. These descriptive statistics provide insights into our efforts to specify the plausible price processes in this paper. They will also be used as qualitative criteria to assess each model presented in this section, in addition to the statistical fitness of the resultant option valuation models to real market data. Stylized facts of allowance option prices, particularly the volatility smile, will be discussed in a later section.

4. Model specifications and assumptions

We denote a filtered probability space by a triplet $(\Omega, \mathcal{F}_t, \mathcal{P})$, where $t \in [0, T]$ in units of years is the time index and T is the terminal time of the current trading phase.[1] The following assumptions are made for all of the models presented in this section:

(1) The market is complete for the trading phase(s) of interest.
(2) Borrowing and lending activities are allowed to take place at a constant continuously compounded risk-free rate r.
(3) All tradings take place continuously over time with negligible transaction costs.

Externality such as undue influences of political forces are not considered in our models.

Under the open trading phase, the actual emission information has reduced effect on allowance prices. Releases of the market emission status may induce shifts in market sentiments toward an allowance's compliance value and subsequently price movements, but do not necessarily lead to discrete price jumps, the absence of which is evidenced in Fig. 1. Therefore, replicating portfolios can be readily constructed for traded contingent claims on allowance and allowance futures contracts, including the options market which is of particular interest in this paper. The improved market stability and volume in the open trading phase has substantially incentivized the participation of global market makers, leading to cost-efficient hedging for allowance derivatives. Hence, we view the complete market assumption as reasonable.

5. Lognormal allowance price model

For the first reduced-form econometrics model, we assume that the allowance price follows a GBM with constant drift and volatility parameters. This specification coincides with that in a classic Black–Scholes model for a stock option valuation proposed in the seminal work of Black and Scholes (1973). Assume that under the real-world probability space $(\Omega, \mathcal{F}_t, \mathcal{P})$, the spot allowance price $S(t)$ follows a time-homogeneous GBM given by:

$$dS(t) = \mu S(t)dt + \sigma S(t)dZ(t) \tag{1}$$

with an initial condition $S(0) = S_0$, where μ and σ are the drift and volatility parameters of the allowance price return, and $Z(t)$ is a standard Brownian motion under the real world measure $\mathcal{P}$. The risk-neutral allowance price process is obtained in the usual manner by a standard application of Girsanov's Theorem; so that under the risk-neutral measure, $\mathcal{Q}$, we have:

$$dS(t) = rS(t)dt + \sigma S(t)dW(t), \tag{2}$$

[1] The phase end time T plays a limited role in allowance option valuation for open trading phases, but it is included here for the sake of completeness only.

where $W(t)$ is a standard Brownian motion under $\mathcal{Q}$ with a likelihood ratio:

$$\frac{d\mathcal{Q}}{d\mathcal{P}} = \exp\left(\frac{\mu - r}{\sigma} Z(T^*) - \frac{1}{2}\left(\frac{\mu - r}{\sigma}\right)^2 T^*\right), \quad \forall T^* > 0.$$

From now on, we work under the risk-neutral measure and hence the probability space $(\Omega, \mathcal{F}_t, \mathcal{Q})$ unless specified otherwise. The arbitrage-free price of an allowance futures contract settling at τ' is given by:

$$F(t, \tau') = e^{r(\tau' - t)} S(t). \tag{3}$$

By construction, an allowance futures contract always settles before the end of the current trading phase. This cap to life of a futures contract represents an important difference between the emission allowance and the stock market, as well as practical implications on rolling futures portfolios over trading phases, for which prudent strategies should be devised. The allowance futures price can be shown to follow a martingale. These results directly translate to conditional lognormality (LN) in price distribution and hence direct adaption of the Black–Scholes results.

Adopt the following notations where superscripts indicate the valuation model used:

(1) $C^{LN}(t, K, \tau)$ and $P^{LN}(t, K, \tau)$ denote respectively the time t values of a call and put option on the allowance contract with strike price K and maturity τ under the LN model.
(2) $C_F^{LN}(t, K, \tau, \tau')$ and $P_F^{LN}(t, K, \tau, \tau')$ denote respectively the time t values of a call and put option with maturity τ and strike price K written on the allowance futures to be settled at τ' under the LN model, where $0 \le t < \tau \le \tau' \le T$.

In addition, whenever the context is clear, the time index may also be moved from the bracket to the subscript to emphasize that the information is known as of the valuation time point.

Proposition 1. *Consider an open trading phase spanning the period* $[0, T]$. *Assume that the spot allowance price follows the GBM given by Eq.* (1). *At time* t, *the risk-neutral prices of the European allowance call and put option with strike price* K *maturing at time* $\tau \in (t, T)$, *respectively, are given by:*

$$C^{LN}(t, K, \tau) = S_t \Phi(d_1) + K e^{-r(\tau - t)} \Phi(d_2),$$
$$P^{LN}(t, K, \tau) = K e^{-r(\tau - t)} \Phi(-d_2) - S_t \Phi(-d_1),$$

where

$$d_1 = \frac{\ln\left(\frac{S_t}{K}\right) + (r + 0.5\sigma^2)(\tau - t)}{\sigma\sqrt{\tau - t}}, \qquad d_2 = d_1 - \sigma\sqrt{\tau - t}.$$

The risk-neutral prices of the European allowance futures call and put option with strike price K *maturing at time* τ, *written on an allowance futures contract to be settled at* τ' *where* $0 \le t < \tau \le \tau' \le T$, *respectively, are given by:*

$$C_F^{LN}(t, K, \tau, \tau') = e^{-r(\tau - t)}\left(F(t, \tau')\Phi(d_1') + K\Phi(d_2')\right),$$
$$P_F^{LN}(t, K, \tau, \tau') = e^{-r(\tau - t)}\left(K\Phi(-d_2') - F(t, \tau')\Phi(-d_1')\right)$$

where

$$d_1' = \frac{\ln\left(\frac{F(t, \tau')}{K}\right) + 0.5\sigma^2(\tau - t)}{\sigma\sqrt{\tau - t}}, \qquad d_2' = d_1 - \sigma\sqrt{\tau - t}.$$

The results in the proposition are obtained by using risk-neutral valuation as well as the expected payoff relationship in Eq. (4) below.

$$E\left[\left(F(\tau, \tau') - K\right)^+\right] = e^{r(\tau' - \tau)} E\left[\left(S(\tau) - K e^{-r(\tau' - \tau)}\right)^+\right]. \tag{4}$$

Notice that the allowance futures settlement time τ' does not enter the allowance futures option value expression, as the option payoff is

fully realized at maturity and any resultant marked-to-market cashflows from the futures position is not attributable to the option.

The LN model serves as a benchmark model in our analysis, based on which richer specifications of econometrics models are proposed in subsequent sections. Numerical results from the implementation of allowance futures option valuations under alternative models are compared to those from Proposition 1. Due to the heavy tails observed in actual allowance price returns in Section 3 which, by construction, is not captured by the GBM price process, it is expected that models which accommodate the excess kurtosis will bring tangible improvements in statistical fitness. This will be demonstrated later on by calibrating them to market option data.

5.1. Volatility smile in allowance options

Volatility smile refers to the phenomenon that the option-implied volatility varies with respect to the strike price where the relationship graphically displays a convex curve. It is a well-known stylized behavior for stock options, but rarely studied in the emission allowance market. This section contributes to filling this gap.

Heuristically, implied volatility is derived from the spot option price by solving for the volatility value that equates the market and model price of an option. In practice, allowance options are traded over-the-counter (OTC), forming a small market with comparatively low liquidity and trading volumes. Hence, the implied volatility are more reliably derived from the allowance futures option prices. In contrast to allowance options, allowance futures options are standardized contracts actively traded in major commodity exchanges such as the ICE Europe. A formal definition for implied volatility is given below accordingly.

Definition 1. The implied volatility of an allowance futures option, denoted by σ_{im}, is the volatility that equates the option's price under the LN model to its market price. For a European allowance futures call with strike price K and maturity τ written on an allowance futures contract to be settled at τ', its time-t implied volatility is defined by the following relationship:

$$C_F^m(t,K,\tau,\tau') = C_F^{LN}(t,K,\sigma_{im},\tau,\tau'), \tag{5}$$

where $C^m(t,K,\tau,\tau')$ is the time-t market price of the allowance futures call option.

Result of Proposition 1 indicate that solving for the implied volatility from market prices requires numerical procedures except for the ideal case where the futures option is at the money (ATM), in which analytical expression for the implied volatility can be obtained.

According to the specification of the LN model, the constant volatility parameter σ is invariant to the option parameters. However, our empirical analysis reveals the existence of volatility smiles where (holding all else constant) the implied volatility varies with the strike price to form the shape of a "smile". Fig. 3 illustrate the implied volatilities at selected dates derived from allowance futures call options maturing in December 2023 (CKZ23) and December 2024 (CKZ24) respectively where the corresponding futures prices are also shown.

From Fig. 3, the existence of volatility smile is evident in the allowance market. Similar to those seen in the financial market, the curvature is more pronounced for the short-term option compared to the long-term. One important yet interesting observation is that the implied volatilities exhibit a reverse skew, which is often seen for index options. Since emission allowance may be viewed as a production input for industries under the emission control scheme, one may expect it to share similar characteristics of the commodity market in which volatility smiles with forward skews are more common. Leveraging various plausible explanations and proposed theories, a primary cause for this phenomenon could be the so-called crash-o-phobia effect, which is captured through a density cluster of a series of substantial losses. The fear of suffering such losses increases demand of out-of-money put options for downside protection. This effect also contributes to the excess kurtosis observed in the allowance returns identified in Section 3, which is not captured by the LN model. A secondary cause for the reverse smile could be the strong demand for deeply in-the-money call options by investors pursuing the leverage effect of options. Such behavior is indicative of market momentum that contribute to non-zero skewness in the short term.

On the solution side, efforts aimed at capturing volatility smiles in stock options include stochastic volatility models studied in Hull and White (1987) and Heston (1993), the local volatility pricing method proposed in Dupire (1994), and econometric models by incorporating GARCH-type conditional volatility processes introduced in Engle and Mustafa (1992). To the best of our knowledge, no studies have been published to this date on volatility smiles observed in allowance option markets. This inspires us to present the models in the subsequent sections which explicitly account for this stylized behavior in allowance option valuation.

5.2. Accommodation for skewness and kurtosis

One immediate refinement to the original specification of the LN model draws inspiration from the work of Corrado and Su (1996) who introduced skewness and kurtosis adjustments to the stock return model in analyzing the observed deviations from normality in S&P500 index returns. In principle, a similar adjustment can be applied to the LN model to achieve an improved statistical fitness of the model to account for the excess kurtosis in the allowance price returns as discussed in Section 3. First, consider the allowance return over an arbitrary period $[u,t]$ given all price information up to u. Under the LN model specification, the standardized allowance price return $R(u,t)$ conditioning on $\mathcal{F}_u$ follows the Standard Normal Distribution.

$$R(u,t) = \frac{\ln\left(\frac{S(t)}{S_u}\right) - (r-0.5\sigma^2)(t-u)}{\sigma\sqrt{t-u}} \qquad \forall u < t. \tag{6}$$

To accommodate skewness and excess kurtosis in the return distribution, Corrado and Su (1996) apply a Gram–Charlier expansion to the density function with truncated Hermite polynomial terms, so that the standardized return $R(u,t)$ now has the following form of the conditional density function:

$$g_R(x) = \phi(x)\left[1 + \frac{\mu_3(x^3-3x)}{3!} + \frac{(\mu_4-3)(x^4-6x^2+3)}{4!}\right], \tag{7}$$

where $\phi(x)$ is the Standard Normal density function and μ_3 and μ_4 are respectively the explicit parameters for the skewness and kurtosis of the allowance price returns: under this transformed density function, we have $E(R^3)=\mu_3$ and $E(R^4)=\mu_4$ with first and second moment being 0 and 1. Therefore, this modified model can be viewed as a generalized version of LN model, which allows for nonzero skewness and excess kurtosis.

Now, consider the time-t risk-neutral price of an allowance call option with strike price K and maturity τ under this skewness–kurtosis-modified LN model (abbreviated as a SKM model), denoted by $C^{SKM}(t,K,\tau)$. Under $\mathcal{Q}$, risk-neutral valuation gives:

$$\begin{aligned} C^{SKM}(t,K,\tau) &= e^{-r(\tau-t)}[(S(\tau)-K)^+] \\ &= e^{-r(\tau-t)}\int_k^\infty \left(S_t\exp(x\sigma\sqrt{\tau-t} + (r-0.5\sigma^2)(\tau-t)) - K\right) g_R(x)dx, \end{aligned} \tag{8}$$

where

$$k = \frac{\ln\left(\frac{K}{S_t}\right) - (r-0.5\sigma^2)(\tau-t)}{\sigma\sqrt{\tau-t}}.$$

This expression is evaluated to yield the following:

$$C^{SKM}(t,K,\tau) = C^{LN}(t,K,\tau) + \mu_3 Y_3 + (\mu_4-3)Y_4, \tag{9}$$

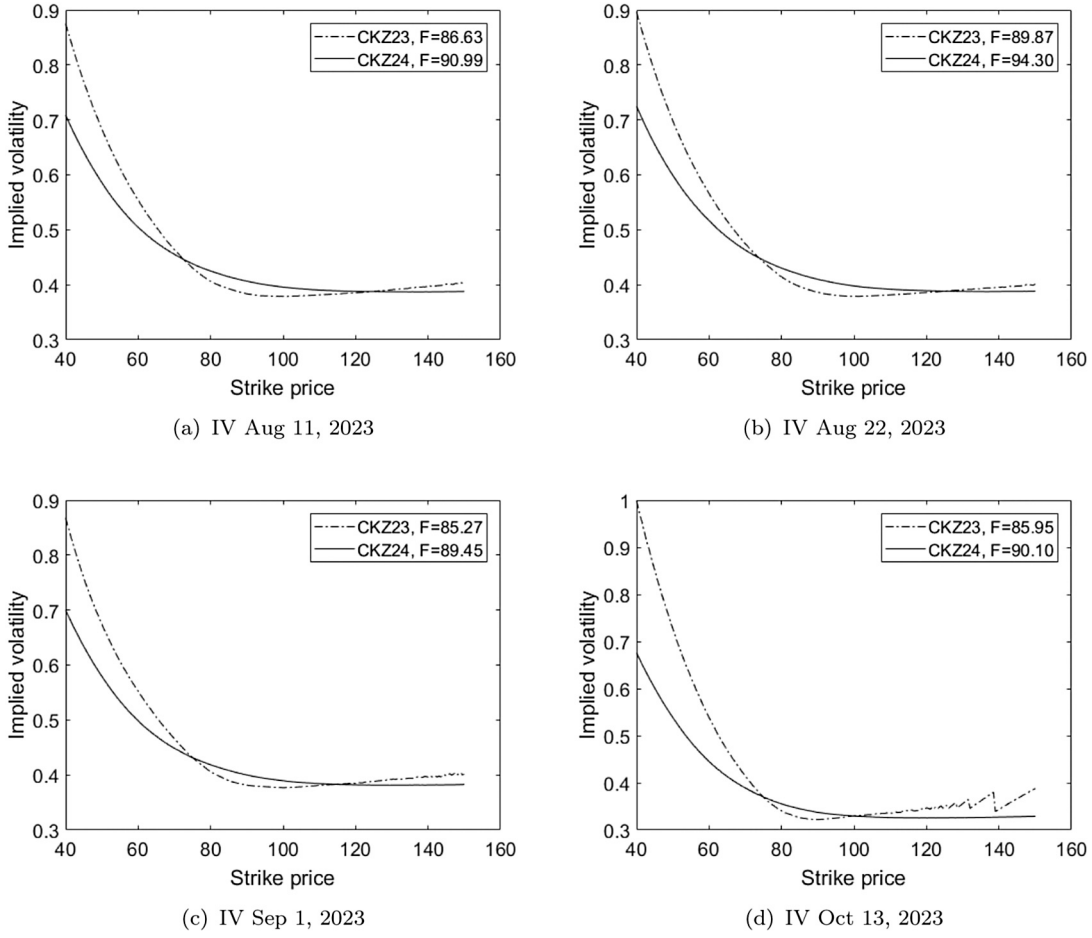

Fig. 3. Implied volatilities (IV) vs. strike price at selected dates.

where $C^{LN}(t,K,\tau)$ is the corresponding call option price under the LN model, and

$$Y_3 = \frac{S_t\sigma\sqrt{\tau-t}}{3!}\left((2\sigma\sqrt{(\tau-t)}-d_1)\phi(d_1)-\sigma^2(\tau-t)\Phi(d_1)\right),$$

$$Y_4 = \frac{S_t\sigma\sqrt{\tau-t}}{4!}\left((d_1^2-1-3\sigma\sqrt{(\tau-t)}d_2)\phi(d_1)+\sigma^3(\tau-t)^{1.5}\Phi(d_1)\right),$$

with d_1 and d_2 as defined in Proposition 1.

The Y_3 and Y_4 components in this option price expression represent, respectively, the marginal effect of nonnormal skewness and kurtosis in allowance returns to the call option value under the risk-neutral measure. Using the results in Eqs. (4) and (8), we make the incremental contribution by deriving the valuation expression of allowance futures options under the SKM model, under which the time-t risk-neutral price of a European allowance futures call option with strike price K and maturity $\tau \in (t,T)$ written on an allowance futures contract to be settled at τ', denoted by $C_F^{SKM}(t,K,\tau,\tau')$, $0 \le t < \tau \le \tau' \le T$, can be shown to follow:

$$C_F^{SKM}(t,K,\tau,\tau') = C_F^{LN}(t,K,\tau,\tau') + \mu_3 Y_3' + (\mu_4-3)Y_4', \tag{10}$$

where $C_F^{LN}(t,K,\tau,\tau')$ is the corresponding allowance futures call option price under the LN model given in Proposition 2, and

$$Y_3' = \frac{F(t,\tau')\sigma\sqrt{\tau-t}}{3!e^{r(\tau-t)}}\left(\left(2\sigma\sqrt{(\tau-t)}-d_1'\right)\phi(d_1')-\sigma^2(\tau-t)\Phi(d_1')\right),$$

$$Y_4' = \frac{F(t,\tau')\sigma\sqrt{\tau-t}}{4!e^{r(\tau-t)}}\left(\left({d_1'}^2-1-3\sigma\sqrt{(\tau-t)}d_2'\right)\phi(d_1')+\sigma^3(\tau-t)^{1.5}\Phi(d_1')\right),$$

with d_1' and d_2' defined in Proposition 1.

The Y_3' and Y_4' components share similar interpretations as Y_3 and Y_4 but in the context of the allowance futures options. The put option prices under the SKM model can be found by using the put–call parity relationships, whose validity is model-invariant.

6. Mixture lognormal allowance price model

In this section, we present a mixture lognormal model for allowance prices (abbreviated as an MLN model) that is expected to deliver improved statistical fitness over the benchmark model by attempting to capture the excess kurtosis and volatility smile.

Assume that under the real-world probability space $(\Omega, \mathcal{F}_t, \mathcal{P})$, the spot allowance price $S(t)$ follows a mixture of N diffusion processes:

$$S(t) = \sum_{i=1}^{N} 1_{C=i} S_i(t), \tag{11}$$

$$dS_i(t) = S_i(t)\mu_i dt + \sigma_i S_i(t) dZ(t), \quad i \in \{1,2...N\}, \tag{12}$$

with common initial condition $S_i(0) = S_0$, where $Z(t)$ is a standard Brownian motion under $\mathcal{P}$. In addition:

(1) $1_{C=i}$ is a binary indicator variable that equals 1 when $C = i$ and 0 otherwise.
(2) C is an underlying market state random variable taking values on $1,2,3...N$ with $Pr(C = i) = \lambda_i$ subject to the adding-up restriction that $\sum_{i=1}^{N} \lambda_i = 1$.
(3) μ_i and σ_i, respectively, are the drift and volatility parameters of the allowance price return given that the market is in state i.

Mathematically, this mixture diffusion specification of allowance price represents a super-imposition of N components price processes

denoted by triplets (S_i, μ_i, σ_i), with corresponding mixing weights denoted as $\lambda_i, i \in \{1, 2, \ldots, N\}$. These weights also serve as probability masses of the underlying market state variable C that dictates the component process in force, and hence make up the joint distribution of the parameter pairs (μ_i, σ_i). This additional source of randomness introduces flexibility to the model in terms of being able to capture potential excess kurtosis in the allowance price returns.

Intuitively, the mixed diffusion specification emphasizes that the allowance price dynamics is subject to uncertainty due to the agents' different behaviors and reactions to information under different market states. This is consistent with the view espoused by modern economics theory as pointed out by Neumann (2002). More specifically, market participants trade in response to the arrival of new information. Each component diffusion process can be perceived as random results of information events under the corresponding market state. However, the price responds differently to a given information event based on the market state. For instance, news of a large emitter downsizing its business should not result in much price reactions in an optimistic market that confidently expects a steady economic growth, compared to a pessimistic market with high informational asymmetry. The market state thus represents the aggregate effect of economic outlooks and market sentiments, which respectively determine the drift and volatility in the component price processes in Eq. (12). Since the market state is not observable in the market, it is modeled as a latent categorical variable C with discrete probability masses. Note that the state is assumed for convenience to remain time-invariant in the specification. As a result, inter-temporal changes in market state are not explicitly captured in this specification. We can now formally introduce the mixture lognormal distribution in Definition 2 below:

Definition 2. Under a filtered probability space $(\Omega, \mathcal{F}_t, \mathcal{P})$, let C denote the underlying discrete state variable taking values on $\{i \mid i \in 1, 2, \ldots, N\}$ with $P(C = i) = \lambda_i$ and $1_{C=i}$ denote a binary indicator variable that takes the value of 1 if $C = i$ and 0 otherwise. A random variable Y follows a mixture of lognormal distributions iff

$$Y = \sum_{i=1}^{N} 1_{C=i} X_i,$$

where the X_i's are independent lognormal random variables with the following density functions:

$$f_{X_i}(x; m_i, \sigma_i) = \frac{1}{\sqrt{2\pi} x \sigma_i} \exp\left\{-\frac{(\ln x - m_i)^2}{2\sigma_i^2}\right\}, \quad i \in \{1, 2 \ldots N\}.$$

For notation convenience, let Λ, M, and $\Sigma \in R^N$ be vectors of the weight parameters λ_i, the mean parameters m_i, and the standard deviation parameters σ_i, $i \in 1, 2 \ldots N$ in Definition 2. Then the distributional relationship of Y can be expressed as:

$$Y \sim MLN(\Lambda, M, \Sigma)$$

While it is tempting to apply risk-neutral valuation to derive the option value expression under the MLN model, the MLN model by construction is an incomplete-market model due to the additional source of randomness introduced by the market state variable, overturning the uniqueness of risk-neutral measure. Approaches to valuation of contingent claims under incomplete-market models are abundant in the literature. This include:

(1) Assuming a particular form of the market's utility function. This, in effect, is equivalent to specifying a market price of risk to be used for an application of the change of measure. Examples of this first approach include Pellizzari and Gamba (2002) and Zhang and Han (2013), both of which are studied under the stock market settings.
(2) Assuming that the risk-neutral measure is chosen to be the one closest to the real-world measure under some metric. Examples of this second approach include He and Zhu (2015) on European stock option valuation by using the minimal entropy martingale measure, as well as Çetin and Verschuere (2009) on allowance valuation under closed trading phases by using the minimal martingale measure.
(3) Directly specifying the risk-neutral dynamics of the allowance price process and calibrating the resultant model option prices to market data. Examples of this third approach include Leisen (2004) on stock option valuation by using mixture models.

From our perspectives, the third approach is especially appealing for its objectivity and is therefore chosen to be implemented in our analysis. There is no convincing evidence to establish that any particular utility functions are adopted by market agents; this challenges the validity and robustness of the first approach. A similar argument can be made for the second approach in that it lacks a solid theoretical foundation. There is no economic rationale as to why the market must price a derivative under a risk-neutral measure close to $\mathcal{P}$. In contrast, the third approach does not require any assumptions about the utility functions or a particular risk measure being used by the market. Instead, The risk-neutral price dynamics is directly specified and the model option prices are calibrated to their market quotes, which reflect all the information from the true risk-neutral measure $\mathcal{Q}$ chosen by the market.

For the rest of this section, we work under the risk-neutral measure and therefore the probability space of interest is given by $(\Omega, \mathcal{F}_t, \mathcal{Q})$. Assume that the spot allowance price under $\mathcal{Q}$ follows a mixture of diffusion processes given by Eq. (11) and component processes:

$$dS_i(t) = rS_i(t)dt + \sigma_i S_i(t) dW(t), \quad i \in \{1, 2 \ldots N\}, \tag{13}$$

with common initial condition $S_i(0) = S_0$, where $W(t)$ is a standard Brownian motion under $\mathcal{Q}$ and r is the continuously compounded risk-free rate. Since each component process follows a GBM that translates readily to conditional lognormality as in the result of Section 5, based on Definition 2, the allowance price follows a conditional mixture lognormal distribution that can be expressed as:

$$S_t \mid S_u \sim MLN\left(\Lambda, 1^N lnS_u + 1^N r(t-u) - \frac{1}{2}\Sigma^2(t-u), \Sigma\sqrt{t-u}\right), \quad \forall u < t \tag{14}$$

where 1^N denotes a vector of ones of size N. We assume that the market is arbitrage-free despite the MLN allowance price specification being an incomplete-market model, so that the relationship in Eq. (4) holds.[2] The allowance futures price is a martingale with dynamics given by a mixture of the following diffusion processes:

$$F(t, \tau') = \sum_{i=1}^{N} 1_{C=i} F_i(t, \tau'), \tag{15}$$

$$dF_i(t, \tau') = \sigma_i F_i(t, \tau') dW(t), \quad i \in \{1, 2 \ldots N\}, \tag{16}$$

with a common initial condition $F_i(0, \tau') = F_0$ and terminal conditions $F_i(\tau', \tau') = S_i(\tau')$. These results are sufficient for the purpose of valuation of allowance options in this section. The most strikingly convenient property of the MLN model is that, for a given option, its value is simply a linear combination of the corresponding N option values from the LN model calculated with the component drift and volatility parameters using the mixture weights as coefficients. This is formally summarized in Proposition 2 below.

Proposition 2. *Consider an open trading phase spanning the period of* $[0, T]$. *Assume that the spot allowance price follows a mixture diffusion process given by Eqs.* (11) *and* (13) *under* $\mathcal{Q}$. *The time-t price of a European*

[2] Detailed discussions on the relationship between arbitrage-free and complete market can be found, for instance, in Bjork (2004).

allowance option with strike price K maturing at time τ is equal to the linear combination of those under the LN model calculated by using component volatilities σ_i weighted by the mixing weights λ_i, $i \in \{1, 2, \dots, N\}$:

$$C^{MLN}(t, K, \tau) = \sum_{i=1}^{N} \lambda_i C^{LN}(t, K, \tau, \sigma_i),$$

$$P^{MLN}(t, K, \tau) = \sum_{i=1}^{N} \lambda_i P^{LN}(t, K, \tau, \sigma_i).$$

Similarly, The time-t price of a European option with strike price K maturing at time τ, written on an allowance futures contract to be settled at τ' where $0 \le t < \tau \le \tau' \le T$, is equal to the linear combination of those under the LN model calculated by using component volatilities σ_i weighted by the mixing weights λ_i, $i \in \{1, 2, \dots, N\}$:

$$C_F^{MLN}(t, K, \tau, \tau') = \sum_{i=1}^{N} \lambda_i C_F^{LN}(t, K, \tau, \tau', \sigma_i),$$

$$P_F^{MLN}(t, K, \tau, \tau') = \sum_{i=1}^{N} \lambda_i P_F^{LN}(t, K, \tau, \tau', \sigma_i),$$

where $C^{LN}(t, K, \tau, \sigma_i)$, $P^{LN}(t, K, \tau, \sigma_i)$, $C_F^{LN}(t, K, \tau, \sigma_i)$ and $P_F^{LN}(t, K, \tau, \sigma_i)$ denote the European allowance and allowance futures option values under the LN model in Proposition 1 calculated with volatility σ_i.

Proof. See Appendix.

Compared to other advanced models built on state-contingent price processes, such as the Regime-switching Lognormal (RSLN) Model proposed by Hardy (2001), the MLN allowance price model has the key advantage of having closed-form results, where model parameters can be easily calibrated to market prices using error-minimizing criteria. Overall, the MLN model is more parsimonious and numerically efficient in practice, while capturing both the volatility smile and heavy tails in the allowance return.

6.1. A two-component mixture lognormal model

For a practical implementation of the MLN model, the total number of mixture components (which is N) must be specified. In principle, introducing additional components is expected to lead to an improved statistical fitness of the model to market option prices but at the cost of numerical efficiency. In this section, we examined the two-component MLN model (abbreviated as an MLN-2 model), which is the simplest instance of the MLN model to incorporate the crash-o-phobia effect and volatility smiles.[3] The model relies on a mixture of two component variables with interesting interpretations that are not particularly transparent in a general MLN model. Under $\mathcal{P}$, assume that the spot allowance price follows a mixture of two diffusion processes:

$$S(t) = \sum_{i=1}^{2} 1_{C=i} S_i(t), \tag{17}$$

$$dS_i(t) = \mu_i S_i(t) dt + \sigma_i S_i(t) dZ(t), \quad S_1(0) = S_2(0) = S_0, \quad i \in \{1, 2\}, \tag{18}$$

where $Pr(C = 1) = \lambda \in [0, 1]$ and $Z(t)$ is a standard Brownian motion under $\mathcal{P}$. Following the option valuation approach discussed for the MLN model, the allowance price under $\mathcal{Q}$ is assumed to follow a mixture of diffusion given by (17) and component process:

$$dS_i(t) = rS_i(t) dt + \sigma_i S_i(t) dW(t), \quad S_1(0) = S_2(0) = S_0, \quad i \in \{1, 2\}, \tag{19}$$

where $W(t)$ is a standard Brownian motion under $\mathcal{Q}$. The time-t risk-neutral option prices follow

$$C^{MLN2}(t, K, \tau, \sigma_1, \sigma_2) = \lambda C^{LN}(t, K, \tau, \sigma_1) + (1 - \lambda) C^{LN}(t, K, \tau, \sigma_2) \tag{20}$$

$$C_F^{MLN2}(t, K, \tau, \sigma_1, \sigma_2) = \lambda C_F^{LN}(t, K, \tau, \tau', \sigma_1) + (1 - \lambda) C_F^{LN}(t, K, \tau, \tau', \sigma_2) \tag{21}$$

The same results apply to put options.

The simplicity of the two-component mixture specification greatly facilitates the study of the component prices directly, which is clumsy to perform as the number of components increases. First, under $\mathcal{P}$, consider the expected terminal price of the allowance given $\mathcal{F}_t$ and a time frame of interest $[t, \tau]$:

$$E[S(\tau) \mid \mathcal{F}_t] = E\left[E\left[S(\tau) \mid \mathcal{F}_t, C\right]\right] = \lambda E[S_1(\tau) \mid \mathcal{F}_t] + (1 - \lambda) E[S_2(\tau) \mid \mathcal{F}_t] \tag{22}$$

It is easy to see that the expected terminal price is a linear interpolations of the expected component prices. Without loss of generality, assume that $E[S_1(\tau) \mid \mathcal{F}_t] < E[S_2(\tau) \mid \mathcal{F}_t]$, which implies:

$$E[S_1(\tau) \mid \mathcal{F}_t] < E[S(\tau) \mid \mathcal{F}_t] < E[S_2(\tau) \mid \mathcal{F}_t].$$

Then λ is the probability that the pessimistic market state dominates (i.e. $C = 1$), under which the allowance price is expected to decrease. The associated component price variable $S_1(t)$ is referred to as a downside component that captures the crash-o-phobia effect in the allowance market. While the level of crash-o-phobia depends on λ and $E[S_1(\tau) \mid \mathcal{F}_t]$, the possibility of large price drops is reflected. Opposite arguments apply to the upside component $S_2(t)$ associated with the optimistic market state (i.e. $C = 2$), which dominates with a probability of $1 - \lambda$.

From an economics perspective, the MLN-2 specification is clearly a simplified discretization of the continuum of possible market scenarios into a dichotomy, i.e., a pessimistic state which is established by bearish factors such as allowance surplus and production downturns, and an optimistic state which is established by bullish factors including allowance deficiency and tightening regulatory standards. The corresponding downside and upside component dynamics aptly describe the market's reactions to the random arrival of information under each state. This is consistent with the previous interpretation we have accorded to the MLN model and is also in alignment with the implications from modern finance theory where asset price movements are explained by a set of fundamental factors and a random component.

However, the above interpretations of the downside and upside components fail to apply to the model under $\mathcal{Q}$, where both expected terminal component prices are equal to the current price accumulated at the risk-free rate as required by the martingale restriction. Upon the calibration of the model to market option prices, the contrasting implications of the two components are reflected in their volatility parameter estimates. Therefore, to properly differentiate between the two components under the risk-neutral measure, we identify the downside component to be the one associated with a higher volatility parameter. This follows from the fact that the component price process with a higher volatility has a smaller drift term in the corresponding return process based on the property of the GBM. Therefore, we define the Share of Downside Return Risk (SDRR) over $[t, \tau]$, denoted by $\Xi(t, \tau)$, as:

$$\Xi(t, \tau) = \begin{cases} \dfrac{Var\left(\ln\left(\dfrac{S(\tau)}{S_t}\right)\right) - \lambda Var\left(\ln\left(\dfrac{S(\tau)}{S_t}\right) \mid C = 1\right)}{Var\left(\ln\left(\dfrac{S(\tau)}{S_t}\right)\right)}, & \sigma_1 < \sigma_2 \\ \dfrac{Var\left(\ln\left(\dfrac{S(\tau)}{S_t}\right)\right) - (1 - \lambda) Var\left(\ln\left(\dfrac{S(\tau)}{S_t}\right) \mid C = 2\right)}{Var\left(\ln\left(\dfrac{S(\tau)}{S_t}\right)\right)}, & \sigma_1 > \sigma_2 \end{cases} \tag{23}$$

[3] Fang (2012) demonstrates the appeal of the MLN-2 model in generating the volatility smiles in stock index options under a semi-hypothetical setting. Such capability is derived from the structural flexibility of the mixture specification and, hence, is market independent.

By working out the expression for each term and simplifying the result, we arrive at:

$$\Xi(t,\tau)=\begin{cases}\dfrac{(1-\lambda)\sigma_2^2+\frac{1}{4}\lambda(1-\lambda)\left(\sigma_1^2-\sigma_2^2\right)^2}{\lambda\sigma_1^2+(1-\lambda)\sigma_2^2+\frac{1}{4}\lambda(1-\lambda)\left(\sigma_1^2-\sigma_2^2\right)^2}, & \sigma_1<\sigma_2\\ \dfrac{\lambda\sigma_1^2+\frac{1}{4}\lambda(1-\lambda)\left(\sigma_1^2-\sigma_2^2\right)^2}{\lambda\sigma_1^2+(1-\lambda)\sigma_2^2+\frac{1}{4}\lambda(1-\lambda)\left(\sigma_1^2-\sigma_2^2\right)^2}, & \sigma_1>\sigma_2\end{cases} \tag{24}$$

Specifically, the SDRR measures the contribution to the risk in allowance return by the downside component implied by option prices. Conversely, it shows the impact of the downside component on option values, which is determined by both the downside variance and the difference between the two component variances.

In our analysis, we will implement the MLN-2 model as a stylized representation of the MLN class due to its numerical simplicity and ease of interpretation. The statistical fitness of the model for option valuation will be comparatively analyzed. In theory, the MLN-2 model is expected to outperform the benchmark model, since the latter is an instance of the former recoverable by pushing the parameter λ to either of its boundaries. The degree of such outperformance in practice is of interest to us to investigate in this paper.

7. Applications and numerical results

This section presents applications of the three open-phase allowance-price models proposed in this paper, followed by a set of numerical implementation of allowance option valuations of each model by using real market data.

7.1. Model applications

In the contemporary global emission economy, the majority of the established ETS are in open-trading phases where the markets are relatively stable and liquid. As a result, it is reasonable to view emission allowance as an alternative asset class contributing to the diversification of investment portfolios. This view, if it is accepted, substantially broadens the scope of application of the models, which can be implemented for a set of tasks, such as:

(1) Pricing, valuation, and risk management of allowance options and related derivatives.
(2) Generating allowance price scenarios by using calibrated parameters.
(3) Pricing, valuation, and hedging of guarantee riders in variable annuities backed by allowance-based funds.

In addition to option valuation runs that the models are intended to serve, the calibrated model of choice can also be used to price options that have limited active quotes. This often happens in the OTC market where the strike-maturity combination of interest serves particular risk-management purposes. Numerical efficiency is warranted by having closed-form valuation expressions, from which the option Greeks are easily derived for the design and implementation of proper hedging strategies in order to manage the market risk exposures in option positions.

Moreover, the allowance price processes in the models can be easily discretized to simulate price scenarios that serve a wide range of actuarial calculations. Risk-neutral allowance price scenarios are generated using parameters calibrated to market option prices, which can be subsequently used to price exotic contracts that lack a closed-form valuation expression. On the other hand, real-world allowance price scenarios are generated using parameters estimated from the spot market data, which is heuristically achieved using an estimation approach such as the conditional maximum likelihood estimation. These scenarios can then be used by insurers exposed to the allowance market for statutory capital calculations and risk-management purposes. The exposure may reside in both the asset side (e.g. having allowances in the asset portfolio) and the liability side (e.g. crediting a contractual rate indexed to allowance prices).

Finally, the models can be used for the pricing and valuation of benefit guarantee riders in variable annuity (VA) products backed by allowance-based funds. Joint specifications of the allowance price model and other actuarial assumptions such as mortality, lapse, utilization, and withdrawal rates are required. Under simplifications to the decrement and policyholder behavior assumptions, the guarantee values can be expressed in closed form as functions of put option values from the selected allowance price model. By extension, this leads to numerical efficiency in both liability hedging and valuation runs, where costly inner loop scenarios are otherwise required through a stochastic-on-stochastic framework. Note that the actual degree of simplifications required vary by guarantee riders.

In the following subsections, we provide a numerical illustration of allowance option valuation for the three reduced-form econometrics models proposed in this paper, which, to the best of our knowledge, has so far not been reported in prior studies in the literature. Additional discussion on the other applications of the models is provided in the subsequent section.

7.2. Data and methodologies

Due to data availability constraints, intra-day allowance option quotes are not available for our study. As an alternative approach, we manually collected and organized weekly closing option prices at Friday market closings from online databases reflecting ICE price quotes. To ensure robustness of the conclusions reached in this analysis, the same implementation is performed for two samples covering distinct periods. Sample A consists of the EU ETS Phase 3 call option prices for two allowance futures contracts settling in Decembers of 2018 and 2019, tickered CKZ18 and CKZ19 respectively. The information is gathered and aggregated over five months from November 17, 2017 to April 6, 2018, with the exclusion of November 25, 2017 and December 22, 2017 due to missing data points. Sample B consists of EU ETS Phase 4 call option prices for 3 allowance futures contracts settling in the Decembers of 2023 to 2025, tickered CKZ23, CKZ24, and CKZ25 respectively. The information is gathered over five weeks from October 13 to November 11, 2023. Compared to Sample A, the much wider range of traded strikes in Sample B ensures data sufficiency as well as an exhibition of the effect of volatility smile on the results.

To construct our sample, for each option, we identify a reasonable strike price range associated with the fewest number of noises, which defines the option data selected for the implementation. The entries are further cross-validated to ensure accuracy and reliability. We do not exclude any low-liquidity options from the sample set as their prices still carry a fair amount of information from the market. The option prices are then matched with the corresponding weekly closing futures price to create our final sample, where each row of data (a record) is a combination of an allowance future option price, the underlying allowance futures price, the option strike, and the option maturity. The futures settlement date is omitted as it does not enter into the option valuation expressions for any of the three models presented in our study. The strike prices range from €3 to €16 and €40 to €105 in increments of €0.5 for Sample A and B respectively. Finally, the risk-free rate r is specified to be 0%, which is the rate implied by the put–call parity relationship.

There are two notable advantages to using allowance futures option data as our sample of choice. First, both options as well as the underlying futures are standardized contracts actively traded in global marketplaces including the European Climate Exchange (ECX) and the Intercontinental Exchange (ICE), where the high liquidity guarantees both the availability and quality of the price data for the purpose of model calibration. Second, it minimizes the undue influence of

Table 2
Summary statistics of the sample data for model implementation.

Sample set	A training	A testing	B training	B testing
Number of records	594	378	844	246
Option prices				
Average	2.78	2.49	15.81	19.44
Min/Max	0.18/10.19	0.16/8.34	0.42/36.00	1.31/38.71
Moneyness				
Average	0.073	0.046	0.109	0.191
Min/Max	−0.797/1.472	−0.803/1.321	−0.213/0.542	−0.154/0.677

Table 3
Parameter estimates from futures option prices, Sample A.

Model	Parameter estimates		
LN	σ		
	0.5094		
SKM	σ	μ_3	μ_4
	0.5057	−0.0804	3.4564
MLN-2	λ	σ_1	σ_2
	0.7278	0.3848	0.8261

specifying the risk-free rate, which as one can verify only affects the futures option values from a discounting perspective for all three models. On the other hand, the disadvantage of this choice is that it results in a limited number of strike-maturity combinations for a comprehensive investigation. As an exchange-traded instrument, each futures contract is associated with only one available option maturity (i.e. 5 days before the futures settlement), precluding a static analysis of model performance over the maturity dimension.

Each sample is divided into a training set and a testing set. For Sample A, the training set consists of sample data two weeks apart starting on November 17, 2017, while the testing set consists of the remaining sample records. For Sample B, the training set consists of consecutive weekly sample data from October 11 to November 3, 2023, while the testing set is made up of the rest. This partition excels Sample A in allowing out-of-sample testing under market conditions beyond the training period. A summary statistic of the sample data is given in Table 2, where the moneyness of an option is defined as the logarithm of the ratio between the underlying allowance futures price and the strike price:

$$moneyness = \ln\left(\frac{F(t,\tau')}{K}\right) \tag{25}$$

In our implementation, the models are calibrated to the training set in the usual manner by minimizing sum of squared differences between the options' model prices and the corresponding market observations, where the search of the optimal parameter set within the prescribed domains are performed using a grid pattern search available in a software such as Matlab. The parameter estimates obtained are then used to calculate the model prices for the options in the testing set. Both in-sample and out-of-sample statistical fitness are comparatively analyzed across the models.

In addition, the estimates of the models-specific parameters and metrics are also examined. For the SKM model, these include the skewness and kurtosis coefficients (i.e. μ_3 and μ_4), the magnitudes of which are indicative of the extent of deviations in option-implied return distributions from the LN model specifications. Similarly, for the MLN-2 model, we focus on the estimates of the mixture weight λ, the differences between the component volatilities, and the Share of Downside Return Risk. Unfortunately, the lack of intra-day data precludes the analysis from achieving desirable statistical significance for these results.

Note that, in theory, both the SKM and MLN-2 models should perform at least as well as the benchmark model, which is essentially a special instance of the two. This numerical exercise aims at developing a better understanding of the level of outperformance in model fitness (if any) for an allowance option valuation.

7.3. Implementation and analysis

The models are calibrated to the training set by solving for the parameter values that minimize the Sum Square Errors (SSE) between the model and the actual market prices in the set, utilizing the option value expressions from Proposition 1, (10), and (21). The LN model is clearly the simplest one to calibrate, since it is associated with only one unknown parameter, making the search uni-dimensional. Both the SKM and MLN-2 models have three parameters that must be calibrated to the training data.

We first present the analysis for Sample A, where the parameter estimates are summarized in Table 3. The volatility estimate from the LN model is 0.5094, indicating a relatively high level of risk in allowance returns implied by the option price. For the SKM model, the volatility parameter estimate is very close to that under the LN model, which is expected since the SKM model adjusts for excess skewness and kurtosis only. The skewness parameter is close to 0, while the kurtosis parameter is above 3, pointing to respectively a lack of return asymmetry and moderately thick tail implied by option prices. Finally, for the MLN-2 model, the two component volatilities differ drastically, characterizing a distinct pair of return regimes that the allowance price may follow. λ is estimated to be 0.7282, which is far from either of its boundary value to support the validity of the mixture model specification. The corresponding Share of Downside Return Risk (SDRR) is 0.6498, indicating that over the calibration period, the downside component contributes to about 65% of the return variance on average as implied by the option prices. Accordingly, the crash-o-phobia effect is found to be material.[4]

The in-sample performances of the models are summarized in the first three columns of Table 4, where three discrepancy metrics are calculated, namely the Mean Square Error (MSE), the Mean Absolute Error (MAE), and the Mean Absolute Percentage Error (MAPE). The results show an overall satisfactory in-sample fit for all three models. As expected, both the SKM and the MLN-2 model outperform the LN model. At the same time, the MLN-2 model delivers recognizably more accurate prices than the SKM model across the three performance metrics, which can be attributed to the ability of the MLN-2 model to explicitly capture the crash-o-phobia effect in the option market in the absence of option-implied skewness. Fig. 4 shows the relative errors of the model prices against the option moneyness for the calibration set. Observe that the LN and SKM models show some tendency to underprice the out-of-money options, while a mild tendency of overpricing toward the out-of-money region can be observed for the MLN-2 model At-the-money and in-the-money options are quite accurately priced by all of the three models.

The out-of-sample performance of the models is assessed using the testing set, where the results are summarized in the last three columns of Table 4. All three models show good out-of-sample performance. The rankings in model fitness is overall consistent with the in-sample results except for the MAPE metrics based on which the MLN-2 model underperforms its rivals. This is apparently caused by the model's overpricing of a series of deeply out-of-money options on January 8, 2018, where the relative errors are inflated by extremely low market option prices. Fig. 5 shows the relative errors of the model prices

[4] Note that ideally, the model is best calibrated to the intra-day high-frequency price data to obtain a time series of parameter estimates and the SDRR metric, which will deliver a more comprehensive insight on how the option market's view varied over the study period.

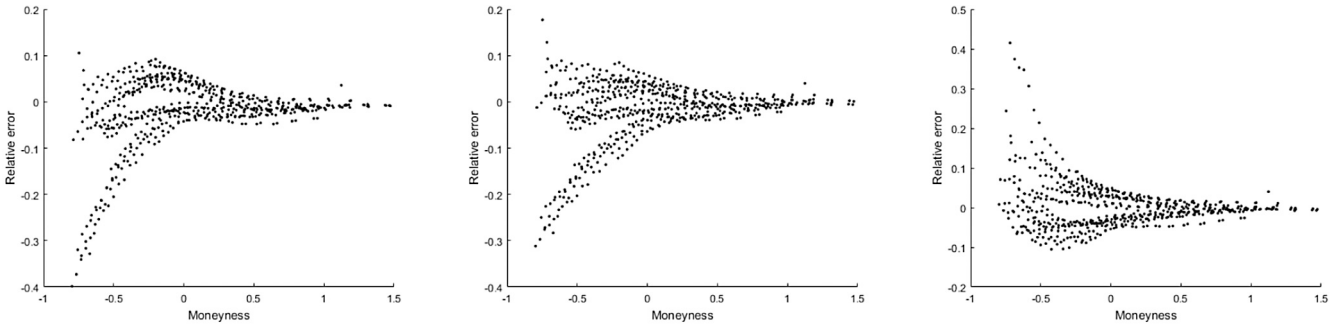

Fig. 4. In-sample relative error vs. moneyness, Sample A.

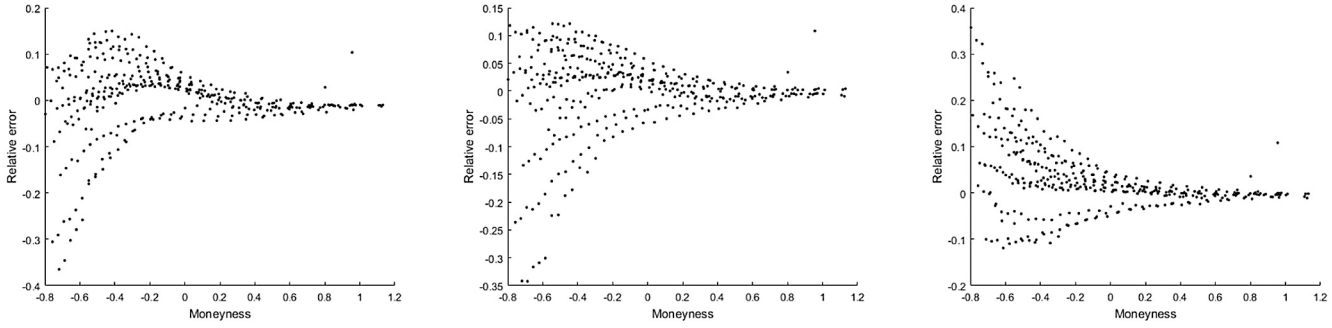

Fig. 5. Out-of-sample relative error vs. moneyness, Sample A.

Table 4
Performance metrics, Sample A.

Model	In-sample			Out-of-sample		
	LN	SKM	MLN-2	LN	SKM	MLN2
MSE	0.0062	0.0053	0.0046	0.0052	0.0044	0.0040
MAE	0.0676	0.0612	0.0577	0.0577	0.0511	0.0479
MAPE	4.56%	4.24%	3.79%	3.90%	3.68%	4.28%

against the option moneyness for the testing set. There is a more saturating dispersion in the relative errors for the out-of-money region, while the overall patterns observed are consistent with the in-sample results, reinforcing our previous argument made based on the in-sample information.

The analysis for Sample B reveals some further insights in the model fitness, primarily due to the different market conditions the sample period represents. Parameters estimates obtained by calibration to Sample B is shown in Table 5. The volatility estimates under the LN and SKM models are both around 0.33, much lower than those from sample A. However, the option-implied skewness under the SKM model is much more significant, which is consistent with the volatility skew illustrated in Section 5.1 for the calibration period. This may lead to enhanced SKM model fitness where return skewness is explicitly captured. The parameter estimates for the MLN-2 model still depict a distinct pair of volatility regimes that the price process may follow. λ is estimated to be 0.3303, which is not only far from its boundary values to support the mixture model specification, but also points to a high probability for the downside regime. The corresponding Share of Downside Return Risk (SDRR) is 0.8533, indicating that the downside component contributes to over 85% of the return variance on average as implied by the option prices over the calibration period.

Model performance metrics for Sample B are summarized in Table 6. While all 3 models showed reasonably satisfactory fitness, interesting insights can be drawn from the observations for each model. For the in-sample metrics, both the SKM and MLN-2 models present material fitness improvement over LN as expected. SKM clearly delivers the best in-sample fitness potentially from capturing the volatility skew mentioned above. However, for the out-of-sample metrics, the MLN-2 model appears to deliver the best overall results by slightly outperforming the SKM model. This is due to SKM's underpricing of a series of deeply out-of-money near-maturity options, consistent with its behavior identified from Sample A. In fact, the rank correlation between the out-of-sample relative errors from the SKM model and option moneyness

Table 5
Parameter estimates from futures option prices, Sample B.

Model	Parameter estimates		
LN	σ		
	0.3304		
SKM	σ	μ_3	μ_4
	0.3394	−0.5043	3.4617
MLN-2	λ	σ_1	σ_2
	0.3303	0.1836	0.6148

Table 6
Performance metrics, Sample B.

Model	In-sample			Out-of-sample		
	LN	SKM	MLN-2	LN	SKM	MLN2
MSE	0.1515	0.0281	0.0561	0.5995	0.3021	0.3298
MAE	0.3190	0.1136	0.1735	0.6395	0.4852	0.4696
MAPE	2.63%	1.26%	3.13%	4.31%	4.29%	3.51%

is approximately 0.94, indicating a high degree of association. On the other hand, the reduced fitness for the SKM model out-of-sample could also be caused by shifts in the volatility skew toward the testing period, which is deliberately set to avoid temporal overlap with the calibration period.

We conclude this section with a few remarks on the models in connection with the insights obtained from this study. First, using the three selected performance metrics over the two different sample periods, the MLN-2 model showed the best overall fitness. However, for periods associated with consistently significant return and volatility skewness, the SKM model may deliver the most accurate results especially for applications such as intraday price extrapolations. In addition, the selection of the model should be based on a particular application it is intended to serve. Though the MLN-2 model has the best overall fitness, it exhibits some tendency to overprice the out-of-money options. This is undesirable if the valuation results are used to set transaction prices in establishing long positions in the options. However, it may be an acceptable solution for an insurer in calculating the required capital of an allowance-indexed liability with option-like payoff to the counterparty.[5] Opposite arguments can be made for the SKM model.

As expected, there is clearly not a perfect-for-all choice. However the conclusions drawn from this study should offer useful insights to possible stakeholders, such as insurers and pension funds, on the tradeoffs involved. Finally, for insurance regulators, the theoretical and numerical findings on the models should hopefully shed light on the development of guidelines for allowance-linked insurance and annuity products, which is not explicitly covered in scope by existing policies.

8. Conclusion

In this paper, we have discussed the valuation of emission allowance options under an open trading phase, characterized by the possibility of inter-phase banking of allowances. We have performed a set of empirical exercises on historical allowance prices under EU ETS Phase 3 and EU ETS Phase 4. They represent the best example of open-phase markets among the in-force emission trading schemes, where the allowance returns exhibit similar stylized facts as the stock market including excess return kurtosis and volatility smiles. Three reduced-form econometrics models were introduced and analyzed systematically. These were the LN model, the SKM model, and the MLN-2 model respectively. Closed-form option price expressions were derived under each of these models, which share a common feature of being analytically-tractable. Numerical illustration was provided by calibrating the models to two samples of allowance futures call option prices collected for EU ETS Phase 3 and Phase 4. The MLN-2 model appears to deliver the most accurate fitness overall across the two sample periods. While all of the three models tend to misprice the out-of-money options, the SKM and MLN-2 models outperform the LN model in the in-sample and out-of-sample model fitness, i.e., based on performance/discrepancy metrics of the models calculated on the training set and the testing set. For periods of significant volatility skew, the SKM model appears to deliver the most accurate in-sample prices and could be best used for intra-day allowance option price extrapolations. Most interestingly, the MLN-2 model showed some tendency to overprice the out-of-money options, which are prone to be underpriced by the SKM model. These findings are important factors to consider in a model selection process for different applications, which may include the design, pricing, and valuation of allowance-linked annuity products. These additional applications of the models are likely to lead to a fruitful avenue for future research.

[5] This may include, for example, rate credited to policyholders of allowance-indexed fixed annuities.

CRediT authorship contribution statement

Mingyu Fang: Conceptualization, Data curation, Formal analysis, Investigation, Methodology, Resources, Software, Validation, Visualization, Writing – original draft. **Ken Seng Tan:** Conceptualization, Formal analysis, Funding acquisition, Investigation, Methodology, Resources, Supervision, Validation. **Tony S. Wirjanto:** Conceptualization, Formal analysis, Funding acquisition, Investigation, Methodology, Project administration, Resources, Supervision, Validation, Writing – original draft, Writing – review & editing.

Data references

The data that support the findings reported in this paper are available from the corresponding author upon request.

Acknowledgments

We thank the editor-in-Chief and a co-editor of the journal and, in particular, an anonymous reviewer, whose constructive comments and suggestions has led to a substantial improvement of both the content and presentation of the paper. In addition, we thank Ben Feng, David Saunders, Ilias Tsiakas, and Olaf Weber for comments and suggestions on the earlier versions of this paper. Finally, we also thank the chair and the participants at an invited session (VTA47 on Green Finance 2) of INFORMS Annual Meeting, Anaheim, California, USA, October 24–27, 2021. The usual disclaimer applies.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Funding

This study is supported by the Society of Actuaries Centers of Actuarial Excellence Research Grant [C013953].

Data references

The data that support the findings reported in this paper are available from the corresponding author upon request

Appendix. Proof of Proposition 2

Without loss of generality, it suffices to show the proof for a call option. Under a risk-neutral valuation:

$$C^{MLN}(t,K,\tau) = E\left[e^{-r(\tau-t)}(S(\tau)-K)^{+} \mid \mathcal{F}_t\right]. \tag{A.1}$$

Based on the model specification in Eqs. (11) and (13):

$$S(t) = \sum_{i=1}^{N} 1_{C=i} S_i(t), \quad i \in 1,2...N, \tag{A.2}$$

where

$$dS_i(t) = S_i(t)rdt + \sigma_i S_i(t) dZ(t). \tag{A.3}$$

Therefore, we have:

$$\begin{aligned} C^{MLN}(t,K,\tau) &= E\left[E\left[e^{-r(\tau-t)}(S(\tau)-K)^{+} \mid \mathcal{F}_t, C\right]\right] \qquad \text{(A.4)}\\ &= \sum_{i=1}^{N} \lambda_i E\left[e^{-r(\tau-t)}\left(S_i(\tau)-K\right)^{+} \mid \mathcal{F}_t\right] \\ &= \sum_{i=1}^{N} \lambda_i C^{LN}(t,K,\tau,\sigma_i). \qquad \text{(A.5)} \end{aligned}$$

where the last line follows from that each component price process being a GBM as specified in Eq. (A.3), coinciding with the specification for the LN model in Section 5.

References

Bjork, Tomas, 2004. Arbitrage Theory in Continuous Time, second ed. Oxford University Press.

Black, Fisher, Scholes, Myron, 1973. The pricing of options and corporate liabilities. J. Polit. Econ. 81 (3), 637–654.

Çetin, Umut, Verschuere, Michel, 2009. Pricing and hedging in carbon emission markets. Int. J. Theor. Appl. Finance 12 (8), 949–967.

Corrado, Charles, Su, Tie, 1996. Skewness and kurtosis in S & P 500 index returns implied by option prices. J. Financ. Res. 19 (2), 175–192.

Dupire, Bruno, 1994. Pricing with a smile. Risk 7, 18–20.

Engle, Robert, Mustafa, Chowdhury, 1992. Implied ARCH models from options prices. J. Econometrics 52 (1–2), 289–311.

Fang, Mingyu, 2012. Lognormal Mixture Model for Option Pricing with Applications to Exotic Options. UWSpace, A thesis presented to the University of Waterloo in fulfillment of the requirement for the degree of Master in Mathematics in Actuarial Science under the supervision of Dr. Tony S. Wirjanto.

Fang, Mingyu, Tan, Ken Seng, Wirjanto, Tony, 2022. Valuation of Carbon Emission Allowance Options under a Closed Trading Phase. Working Paper, University of Waterloo.

Hardy, Mary, 2001. A regime switching model of long term stock returns. N. Am. Actuar. J. 5 (1), 41–53.

He, Xinjiang, Zhu, Songping, 2015. Pricing European options with stochastic volatility under the minimal entropy martingale measure. European J. Appl. Math. 27 (2), 233–247.

Heston, Steven, 1993. A closed-form solution for options with stochastic volatility with applications to bond and currency options. Rev. Financ. Stud. 6 (2), 327–343.

Hull, John, White, Alan, 1987. The pricing of options on assets with stochastic volatilities. J. Finance 42 (2), 281–300.

Leisen, Dietmar, 2004. Mixed lognormal distributions for derivatives pricing and risk-management. Soc. Comput. Econ. 48.

Mnif, Walid, 2012. Incomplete Market Models of Carbon Emissions Markets (Electronic Thesis and Dissertation Repository). p. 975.

Neumann, Marco, 2002. Option Pricing under the Mixture of Distributions Hypothesis. Preliminary Discussion Paper 208.

Pellizzari, Paolo, Gamba, Andrea, 2002. Utility based pricing of contingent claims in incomplete markets. Appl. Math. Finance 9 (4), 241–260.

Zhang, Qiang, Han, Jiguang, 2013. Option pricing in incomplete markets. Appl. Math. Lett. 26 (10), 975–978.