Navigating the Carbon Markets: Optimal Procurement of Emission Allowances for Industrial Buyers

Nils Loehndorf[a] and Saswat Patra[*,b]

[a]*Centre for Transportation and Logistics, Massachusetts Institute of Technology,*
Email:nils@loehndorf.com
[b]*Department of Finance, Queen's Business School, Queen's University Belfast, GIFT City, India,*
Email:s.patra@qub.ac.uk

Abstract

This article considers the decision problem of a buyer who purchases emission allowances to compensate for carbon emissions of an industrial firm. The buyer decides about the timing and quantity of buying allowances on spot and futures markets and must trade-off the overall cost of compliance with the risk of non-compliance. As with any commodity, the buyer faces several risks, such as market, quantity, and regulatory risks. Carbon traders can use the model to calibrate a purchasing strategy that minimizes cost and matches organizational risk preferences. We formulate the problem as a two-stage optimization model to get stylized results and follow it up with a multi-stage stochastic programming problem considering forward curve dynamics, market impact cost due to limited liquidity, and demand uncertainty. We develop a solution approach based on stochastic dual dynamic programming and show how arbitrage-free discrete Markov processes can be trained on historical returns. Our results inform carbon buyers about finding the right mix of futures and spot that matches expected future demand of emission allowances as well as the firm's risk preferences. The article thereby contributes to the digitization of commodity procurement by enabling automation of decision processes. Model-driven decision-making thereby replaces manual, human-based analytics and risk management that is often error-prone, time-consuming and biased.

Keywords: Stochastic programming, Finance, Carbon emission trading; Commodity trading and operations management; Sustainable Operations

*Corresponding Author

Note: A recent version of the paper was presented at the 2024 Annual Conference of Commodities and Energy Market Association in Boston.

Navigating the Carbon Markets: Optimal strategy for industrial emission allowance procurement

Abstract: This article considers the decision problem of a buyer who purchases emission allowances to compensate for carbon emissions of an industrial firm. The buyer decides about the timing and quantity of buying allowances on spot and futures markets and must trade-off the overall cost of compliance with the risk of non-compliance. As with any commodity, the buyer faces several risks, such as market, quantity, and regulatory risks. The work is motivated by collaborating with a large chemical firm. Carbon traders can use the model to calibrate a purchasing strategy that minimizes cost and matches organizational risk preferences. We formulate the problem as a two-stage optimization model to get stylized results and follow it up with a multi-stage stochastic programming problem considering forward curve dynamics, market impact cost due to limited liquidity, and demand uncertainty. We develop a solution approach based on stochastic dual dynamic programming and show how arbitrage-free discrete Markov processes can be trained on historical returns. A backtest on historical price data demonstrates that the proposed model generates reasonable results compared to standard benchmark practitioners use. Our results inform carbon buyers about finding the right mix of futures and spot that matches expected future demand of emission allowances as well as the firm's risk preferences. The article thereby contributes to the digitization of commodity procurement by enabling automation of decision processes. Model-driven decision-making thereby replaces manual, human-based analytics and risk management that is often error-prone, time consuming and biased.

Key words: carbon emission trading; commodity trading and operations management; emission price models; multi-stage stochastic programming

1 Introduction

The European Union Emission Trading System (EU ETS) is a key policy instrument for reducing greenhouse gas emissions and mitigating climate change. The EU ETS operates by setting a cap on the total amount of CO2 emissions that can be emitted by participating industries, and allocates emissions allowances (EUAs) to those industries. Companies that exceed their allocated emissions must purchase additional EUAs, while those that emit less can sell their excess EUAs. The EU ETS has a well-established spot market, as well as a futures market that allows companies to hedge against price uncertainty by buying EUAs for future delivery at a fixed price.

Industrial buyers of EUAs who participate in the EU ETS face a number of challenges in terms of how to optimally purchase and manage their emission allowances. On the one hand, they want to reduce the cost of compliance by purchasing the minimum number of EUAs needed while, on the other hand, they also want to mitigate the exposure to risks associated with buying EUAs.

One challenge is that trading EUAs carries significant risks that directly impact a firm's financial performance:

1. The price of EUAs is highly volatile and sensitive to changes in economic activity, end-product prices, and renewable production (Ampudia et al. 2022).

2. Emission quantities are often unknown ahead of time and buying large quantities at short notice can be prohibitive, as the spot market can be a relatively thin market with large bid/ask spreads. Also, having bought excess allowances when demand drops ties up capital that is needed elsewhere.

3. Industrial buyers may be exposed to credit risk. This can be particularly relevant for purchases made on futures exchanges.(Eckert and Buli 2022).

4. The EU ETS is subject to ongoing regulatory changes, which can affect the availability and price of EUAs. Examples are periodic updating of medium and long-term caps as well as ad-hoc interventions (backloading). Koch et al. (2016) show that backloading in particular has resulted in sudden price jumps of EUA front-year futures.

Another challenge for industrial buyers is managing the trade-off between buying EUAs on the spot market vis-a-vis the futures market. The spot market allows buyers to purchase EUAs for immediate delivery at the current market price, while the futures market allows buyers to lock in a fixed price for EUAs that will be paid and delivered in the future. Buying EUAs on the spot market provides more flexibility, as it allows buyers to respond to changes in demand, but it may also be more risky due to limited liquidity and price uncertainty. Buying EUAs on the futures market may provide more price stability, but it is also more expensive and may tie up cash in the event of a decrease in emissions.

Model-driven approaches that guide the buying process and address these challenges are limited, so these decisions are often managed based on risk frameworks that exist for other types of commodities, which may not be suitable due to the unique nature of the EU ETS.

Our work aims to address this gap by providing a detailed model for optimal portfolio allocation and risk management that accounts for various sources of risk faced by buyers of EUAs.

We first present a stylized solution to the problem in a two-stage setting and then propose a multi-stage model along with a numerical solution strategy. Our model takes an operational view, where the buyer decides which products to buy on a weekly basis over a time horizon of several years under uncertainty about allowance demand and market prices. The model considers the trade-off between the firm's cost of capital and the markup for buying futures, as well as the effect of market impact cost from trading different products in different periods. In this way, the model captures the most important details of the operational decision process faced by buyers of EUAs.

Central to the solution of the multi-stage problem is a risk-averse stochastic programming model, where the level of risk-aversion, which often drives forward buying decisions, is controlled by a dynamic risk measure (Shapiro 2009). Given the high-dimensional nature of the decision problem, we propose an approach

based on Pereira and Pinto (1991)'s stochastic dual dynamic programming to solve the problem. We model the random data process of spot and futures prices as well as demand as a discrete Markov process that is trained on historical returns. This idea is new for this problem class, and it allows us to get an optimal solution for the problem (Girardeau et al. 2014).

We investigate the decision problem analytically (in the two-stage setting) and numerically (in the multi-stage setting) and identify the trade-offs involved in different purchasing strategies.

We find that the optimal strategy for large industrial buyers of carbon allowances is to use a mix of spot and futures with futures dominating spot purchases, whereas, for relatively small buyers, the optimal strategy is to use a mix of spot and futures with spot dominating futures in the portfolio. Further, the optimal risk-neutral buying strategy is to wait and use the futures market mostly to reduce transaction cost. In contrast, risk aversion prompts early buying of large quantities of allowances in the spot market, which contrast conventional wisdom that views the futures market as means to reduce risk.

As the multi-stage model captures the operational decision process of a carbon trader at a high level of detail, the model can be used to provide concise decision support for a carbon trader tasked with choosing the right time, volume, and product to cover future demand, as well as offer risk management a tool to calibrate a hedging strategy that matches the risk preferences of the organization. Our approach thereby contributes to the growing body of knowledge on digitization of decision process in commodity procurement.

The remainder of the paper is organized as follows. In Section 2, we provide a literature review on related work. We develop a two-stage analytical model for the risk-averse and risk-neutral buyer in section 3. In Section 4, we present a numerical formulation of the problem in a multi-stage setting. In Section 5, we describe how the model can be reformulated to become computationally tractable and introduce our solution approach. In Section 6, we discuss the results of the study, including the backtest results on historical data. We conclude the paper in Section 7 with a summary and an outlook on future work.

2 Related Literature

There is a rich body of literature on the economics of emission allowances that considers the environment as a public good (Ellerman et al. 2010, Tietenberg 2010). As we view the allowance market and its design as exogenous, this work is related to three streams of research: (i) emissions allowance trading, (ii) emission allowance pricing, and (iii) commodity trading and operations management.

2.1 Carbon Emission Trading

Previous research on optimal purchasing of emission allowances for polluting industries is scant, as this buying category is relatively new.

The work that is closest to ours is Zapletal et al. (2020, 2022) who model the emissions trading problem of a steel firm as a risk-averse multi-stage stochastic programming problem. While Zapletal et al. (2020) consider spot and futures prices and discretize random parameters with a scenario tree, Zapletal et al. (2022) additionally consider options and propose a solution based on stochastic dual dynamic programming. Both models, however, operate on a coarse time scale of one year and do not account for transaction cost, which limits their applicability as an operational model in carbon trading.

Other work on carbon trading is in Rong and Lahdelma (2007) who model the decision problem of coordinating power production and emission trading. The authors develop a scenario-based optimization approach. However, their work only considers the spot market, whereas recent adaptations to emission allowance trading, such as banking, are not taken into account.

Gong and Zhou (2013) investigate the joint problem of production planning and emissions trading by modeling the decision problem of a manufacturer who can choose between two production technologies with different emission intensities as a stochastic-dynamic program. Their work views the problem from a tactical planning perspective, where production decisions are endogenous, whereas our focus lies on modeling the operational decision process of a carbon trader who views emissions as exogenous.

Yuan et al. (2018) investigate the joint problem of production control and emission permit trading. The authors consider the case of a small price-taking player and consider fixed transaction cost due to administrative overhead, search, and bargaining cost. We do not make this assumption, as we consider that the firm operates a trading desk and must consider transaction cost in the form of market impact cost.

2.2 Emission Price Models

Our research employs historical emission allowance data to fit a stochastic process, aligning it with literature on emission allowance pricing.

Although emission allowances are not commodities, they are treated like commodities by regulators (Huettemann et al. 2010). Modeling and pricing for commodity trading and risk management is extensively covered in Geman (2009).

Seminal work on dynamic behaviour of emission allowance prices is in Seifert et al. (2008) who analyze spot price dynamics using a stochastic equilibrium model. They conclude that spot prices should be free from seasonal patterns and discounted prices should be martingales. Benz and Trück (2009) argue that log-returns exhibit jumps and propose GARCH or regime-switching models to capture this behaviour. Koch et al. (2016) explain such jumps by the reaction of allowance prices to regulatory news. In contrast to these authors, we fit a non-homogeneous Markov chain model to historical returns, which captures jumps through an increased transition probability between extreme price states.

Bredin and Parsons (2016) and Palao and Pardo (2021) study the relationship between allowance futures prices and spot prices and highlight the fact that there is a persistent negative convenience yield. Both model

futures prices as a log-linear function of spot price and time to maturity, where the slope is governed by the combination of the risk-free rate and the (unobservable) convenience yield. We use a similar approach by estimating the negative convenience yield and the joint dynamics of this slope and the spot price.

2.3 Commodity Trading and Operations Management

Our work adds to the literature on commodity trading in operations and supply chain management (e.g, Kleindorfer 2009, 2010). We will focus our review on multi-period models with inventory, as this is closest to our setting.

Much of the operations management literature takes an integrated view. Devalkar et al. (2011) consider a joint model of commodity buying, processing, and selling, and derive threshold policies for the optimal inventory levels. Devalkar et al. (2018) extend this work by introducing a dynamic risk measure. Although we do not consider an integrated model, our model uses the same dynamic risk measure.

Kouvelis et al. (2013) consider the joint problem of production planning and commodity trading as a problem of dynamic portfolio optimization. They introduce the notion of inter-period utility functions and use the mean-variance framework to model risk-aversion. Kouvelis et al. (2018) extend this work by proposing a time-consistent risk measure within the mean-variance framework.

Mandl and Minner (2020) and Mandl et al. (2022) propose to train threshold policies directly to historical commodity data to provide buyers with price and quantity targets. However, the authors neither model risk preferences nor transaction costs. Also, in contrast to their model, our model does not attempt to anticipate price trends but intends to reduce risk and transaction costs in an optimal way.

None of the above approaches consider market impact cost and the spread between the risk premium paid for futures and the cost of capital of the firm, which plays a crucial role in the optimal trading strategy.

2.4 Contribution

In contrast to the current body of literature, we do not take an integrated view but consider the demand for emission allowances as exogenous and beyond the buyer's control. This mirrors the practical constraints industrial buyers face with limited influence over the production planning process. We also attempt to model the operational decision process at a high level of detail and a decision frequency that reflects the realities these buyers face.

Our approach handles several complications that are relevant to large industrial buyers of emission allowances: (1) excess allowances can be stored at no cost and carried over to future allowance periods; (2) large buyers must account for lack of liquidity in the spot market and futures markets, which means that transaction costs must be taken into account; (3) spot purchases, as well as futures, can be used to satisfy demand within a compliance period so that cash flow risk must be weighed against cost-of-capital considerations.

In this way, our work advances the automation of decision processes, where model-driven decision-making replaces manual, human-based analytics that is often error-prone, biased, and time-consuming.

3 Two-stage Analytical Model

In this section, we set up a stylized model for a buyer of carbon allowances in a two-period framework. For our analysis, we consider a risk-neutral (*RN*) and risk-averse (*RA*) firm that must buy emission allowances to offset carbon emissions. Both firms are identical in all respects except for their risk preferences.

The firm can buy carbon allowances in the spot and futures markets in period one at price p_{1s} and p_{1f}, respectively, whereas in period two, trades are restricted to the spot market. The price in period two, p_{2s}, is a random variable with mean μ_p and variance σ_p^2).

Spot and the futures markets exhibit different levels of liquidity and trades in either market incur market impact cost, β_s and β_f, which impacts the overall transaction cost of participating in these markets.

The demand for carbon allowances d is a random variable unknown in period one, with mean μ_d and variance σ_d^2. The firm must meet the realized demand in period two based on allowance purchases, where x_{1s} and x_{2s} are spot purchases in period one and two, respectively, and x_{1f} denotes the purchase of futures in period one. We assume that short selling is not allowed, which means that spots and futures purchased in period one must be non-negative, whereas allowances purchased in excess of allowance demand can be sold at the spot price in period two.

We also account for the time value of money by discounting transactions in period two at factor $\delta \in (0,1]$ (based on the cost of capital of the firm). Unless stated otherwise, we assume that the market is in contango, which means that futures prices are higher than spot prices. This is typically the case in European carbon markets, since banking emission allowances is effectively costless.

The objective of the firm is to choose an optimal mix of spot and futures in period one to minimize the (risk-adjusted) cost of compliance, C.

3.1 Risk-neutral Buyer

Let us first consider the risk-neutral firm. In a two-period model, the objective of the risk-neutral buyer is to minimize C^{RN} given by:

$$\min_{x_{1f}^{RN}, x_{1s}^{RN}} (p_{1s} + \beta_s x_{1s}^{RN}) x_{1s}^{RN} + (p_{1f} + \beta_f x_{1f}^{RN}) x_{1f}^{RN} + \delta E[(d - x_{1s}^{RN} - x_{1f}^{RN})(p_{2s} + \beta_s (d - x_{1s}^{RN} - x_{1f}^{RN}))]$$

Differentiating partially w.r.t. x_{1f}^{RN} and x_{1s}^{RN}, we have $\partial C^{RN}/\partial x_{1s}^{RN} = 0$ and $\partial C^{RN}/\partial x_{1f}^{RN} = 0$ and:

$$2\beta_s(1+\delta)x_{1s}^{RN} + 2\beta_s\delta x_{1f}^{RN} = \delta\mu_p + 2\beta_s\delta\mu_d - p_{1s}$$

$$2\beta_s\delta x_{1s}^{RN} + (2\beta_s\delta + 2\beta_f)x_{1f}^{RN} = \delta\mu_p + 2\beta_s\delta\mu_d - p_{1f}$$

Solving for x_{1s}^{RN} and x_{1f}^{RN} and assuming $\delta\mu_p = p_{1s}$, we get:

$$x_{1f}^{RN} = \frac{2\delta\beta_s\mu_d}{2(\delta\beta_s+\beta_f+\delta\beta_f)} - \frac{(p_{1f}-p_{1s})(1+\delta)}{(\delta\beta_s+\beta_f+\delta\beta_f)} \;\; and \;\; x_{1s}^{RN} = \frac{2\delta\beta_f\mu_d}{2(\delta\beta_s+\beta_f+\delta\beta_f)} + \frac{\delta(p_{1f}-p_{1s})}{2(\delta\beta_s+\beta_f+\delta\beta_f)}$$

Lemma 1: *The risk-neutral buyer purchases futures in period one only if expected demand μ_d is greater than $\mu_d^{'} = \frac{(p_{1f}-p_{1s})(1+1/\delta)}{2\beta_s}$*

When the expected demand for the carbon allowances in period two is low (i.e., below threshold $\mu_d^{'}$), spot purchases have little effect on the spot price of the carbon allowances (in terms of transaction/market impact cost), and the cost of spot purchases is always lower than purchasing in the futures market. Thus, the risk-neutral buyer sticks to spot purchases when the expected demand is low. However, this demand threshold varies with the liquidity of the spot markets. If liquidity in the spot market decreases, market impact cost increases, and buying futures becomes relatively more attractive as it reduces the overall transaction costs.

Lemma 2: *The total quantity of spot and futures purchases by the risk-neutral buyer in period one is less than the expected demand μ_d*

Proof: See Appendix A for the proof.

This result tells us that the risk-neutral buyer prefers to wait for the demand to be realized rather than buying more allowances in period one than there is expected demand. Since the risk-neutral buyer is agnostic to the variance of price and demand, it is optimal to split purchases between period one and two to reduce overall market impact cost, thereby lowering the overall cost of compliance.

3.2 Risk-averse Buyer

Let us now consider the risk-averse buyer. Denote $\lambda > 0$ as the measure of the risk-aversion. In a two-period model, the objective of the firm is to minimize C^{RA} given by

$$\min_{x_{1f}^{RA}, x_{1s}^{RA}} (p_{1s}+\beta_s x_{1s}^{RA})x_{1s}^{RA} + (p_{1f}+\beta_f x_{1f}^{RA})x_{1f}^{RA} + \delta E[(d-x_{1s}^{RA}-x_{1f}^{RA})(p_{2s}+\beta_s(d-x_{1s}^{RA}-x_{1f}^{RA}))] + \delta\lambda Var[(d-x_{1s}^{RA}-x_{1f}^{RA})(p_{2s}+\beta_s(d-x_{1s}^{RA}-x_{1f}^{RA}))].$$

Differentiating partially w.r.t. x_{1f}^{RA} and x_{1s}^{RA}, we have $\partial C^{RA}/\partial x_{1s}^{RA} = 0$ and $\partial C^{RA}/\partial x_{1f}^{RA} = 0$, and

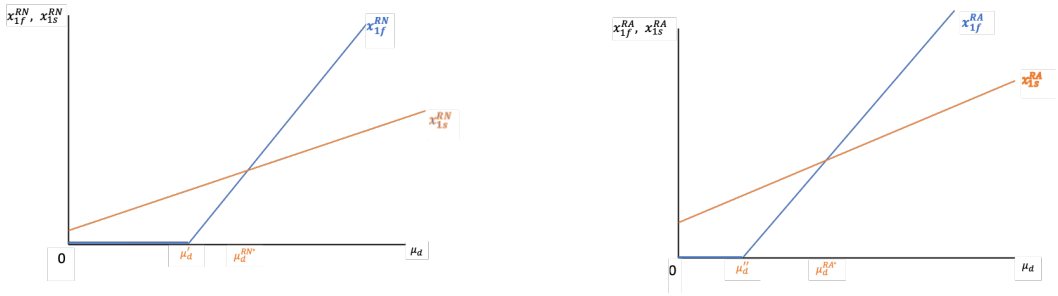
$$(2\beta_s(1+\delta) + 2\lambda\delta(\sigma_p^2 + 4\beta_s^2\sigma_d^2))x_{1s}^{RA} + 2\beta_s\delta x_{1f}^{RA} = \delta\mu_p + 2\beta_s\delta\mu_d - p_{1s} + \delta\lambda((\sigma_p^2 + 4\beta_s^2\sigma_d^2)\mu_d + 2\beta_s\mu_p\sigma_d^2),$$

and

$$(2\beta_s\delta + 2\beta_f + 2\lambda\delta(\sigma_p^2 + 4\beta_s^2\sigma_d^2))x_{1f}^{RA} + 2\beta_s\delta x_{1s}^{RA} = \delta\mu_p + 2\beta_s\delta\mu_d - p_{1f} + \delta\lambda((\sigma_p^2 + 4\beta_s^2\sigma_d^2)\mu_d + 2\beta_s\mu_p\sigma_d^2).$$

Substituting $a = \lambda(4\beta_s^2\sigma_d^2 + \sigma_p^2)$, $g = \delta\mu_p(1+2\lambda\beta_s\sigma_d^2)$ and solving for x_{1s}^{RA} and x_{1f}^{RA}, we get

$$x_{1f}^{RA} = \frac{(2\delta\beta_s + a\delta)(\beta_s + a\delta)\mu_d + \delta p_{1s}\beta_s - p_{1f}\beta_s - \delta p_{1f}(\beta_s + a) + g(\beta_s + a\delta)}{E}$$



(a) Risk-neutral (b) Risk-averse buyer

Figure 1 Mix of spot and futures buying as a function of expected demand

and

$$x_{1s}^{RA} = \frac{(2\delta\beta_s + a\delta)(\beta_f + a\delta)\mu_d + \delta p_{1f}\beta_s - p_{1s}\beta_f - \delta p_{1s}(\beta_s + a) + g(\beta_f + a\delta)}{E}.$$

Lemma 3: When the market is in contango, the risk-averse firm purchases futures in period one only if expected demand μ_d is greater than $\mu_d'' = \frac{(p_{1f}-p_{1s})}{(2\delta\beta_s+\delta a)}(1 + \frac{\delta\beta_s}{(\delta a+\beta_s)}) - \frac{2\beta_s\sigma_d^2}{(2\delta\beta_s+\delta a)}p_{1s}$

Similar to the risk-neutral buyer, the risk-averse buyer purchases exclusively spot if expected demand in period two is below threshold μ_d''. In this case, buying futures is suboptimal as the cost of buying spot is lower than the cost of buying futures (including transaction cost).

Lemma 4: The quantity of spot and futures purchased in period one by a risk-averse firm is higher than that of a risk-neutral firm

Proof: See Appendix A for the proof.

This lemma is a result of including the variance of the total cost in the objective. Unlike the risk-neutral buyer, the risk-averse buyer trades a decrease in uncertainty of price and demand for an increase in cost from buying more spot and futures in period one. By contrast, the risk-neutral buyer is indifferent towards risk and seeks to minimize the expected total cost of compliance.

PROPOSITION 1. *When futures are more liquid than spot markets, both kinds of buyers increase the share of futures in the portfolio with an increase in expected demand, μ_d, and futures dominate spot as μ_d crosses* $\mu_d^{RN^*} = \frac{(p_{1f}-p_{1s})(1+2\delta)}{2\delta(\beta_s-\beta_f)}$ *(*$\mu_d^{RA^*} = \frac{p_{1f}-p_{1s}}{\beta_s-\beta_f} + \frac{\beta_s p_{1f}-\beta_f p_{1s}}{(\beta_s-\beta_f)(2\delta\beta_s+\delta a)} - \frac{g}{2\delta\beta_s+\delta a}$*).*

Proof. See appendix A for the proof.

Both, risk-neutral and risk-averse buyer, have a spot-heavy portfolio when expected demand in period two is low. As expected demand, μ_d, crosses the given thresholds, $\mu_d^{RN^*}$ and $\mu_d^{RA^*}$, futures dominate spot purchases, because the higher liquidity of the futures markets leads to lower transactions cost. Further, the lower transaction costs for the futures outweigh the relatively higher prices and overall cost of compliance is reduced. Thus, both the buyers now prefer a futures heavy portfolio. Figure 1 depicts the relationship between spot and futures purchases and expected demand for risk-neutral as well as risk-averse buyers.

PROPOSITION 2. *When the spot market is more liquid than the futures market, buyers increase spot volumes with an increase in expected demand, μ_d, and spot purchases dominate futures purchases as μ_d crosses $\mu_d^{RN^*}$ ($\mu_d^{RA^*}$)*

Proof: See Appendix A for the proof.

When the market is in backwardation, spot dominates the futures when expected demand exceeds $\mu_d^{RN^*}$ ($\mu_d^{RA^*}$), because although market impact cost (transaction costs) in spot markets are lower than the futures, the spot price is higher than the futures. Hence, the overall cost of buying futures is lower when the expected demand is low. However, when the expected demand crosses the given thresholds $\mu_d^{RN^*}$ ($\mu_d^{RA^*}$), the lower market impact cost in spot markets outweigh the higher prices, thereby reducing the overall cost of compliance and spot dominate the futures in the portfolio of both the buyers.

In contrast to this, when the market is in contango, the portfolio for both buyers always have a higher proportion of spot than futures. In addition to this, the risk-neutral (risk-averse) buyer begins purchasing futures if expected demand exceeds $\mu_d^{'}$ ($\mu_d^{''}$), but since spot is more liquid than the futures, the proportion of futures in the portfolio for both buyers never exceeds spot purchases.

PROPOSITION 3. *When the futures market is more liquid than the spot markets, we find that the risk-averse buyer starts purchasing futures earlier when compared to a risk-neutral buyer, i.e. $\mu_d^{RA^*} < \mu_d^{RN^*}$. In addition to this, we also find that the futures dominate the spot in the portfolio relatively earlier for a risk-averse buyer vis-a-vis the risk neutral buyer i.e., $\mu_d^{''} < \mu_d^{'}$*

Proof: See Appendix A for the proof.

When the expected demand in period two is low i.e., $\mu_d < \mu_d^{''}$, both types of buyers purchase allowances only in the spot market, as futures are too expensive and high prices are not offset by lower transaction cost in the futures markets. As expected demand increases beyond $\mu_d^{''}$, the risk-averse buyer begins buying futures to reduce its market impact cost in the spot market. By contrast, the risk-neutral buyer begins buying futures only once expected demand crosses the higher threshold $\mu_d^{'} > \mu_d^{''}$. Since the risk-averse buyer is more sensitive to the increase in cost variability she shifts more volume to period one and thus begins trading in futures at lower expected demand than the risk-neutral buyer. With a further increase in expected demand i.e., $\mu_d > \mu_d^{RA^*}$, the portfolio of the risk-averse firm becomes dominated by futures while the portfolio of the risk-neutral firm is still dominated by spot purchases. In this case, lower market impact cost in the futures market offset the lower price of purchasing spot, so that the risk averse buyer purchases more allowances in the futures market than in the spot market. The same does not hold for the risk-neutral buyer who only begins to buy more futures than spot as expected demand crosses the threshold of $\mu_d^{RN^*}$. The reason for this is that the risk-neutral buyer shifts more volume to period two and in this way has lower transaction cost in period one than the risk-averse buyer.

4 Multi-stage Numerical Model

We now shift our attention to a multi-stage formulation of the decision problem that lends itself to a numerical solution. As before, we consider an industrial firm that buys emission allowances for each emitted ton of carbon dioxide that exceeds its allocated allowances during a compliance period (typically a calendar year). In the multi-stage model, the buyer tasked with carbon trading buys emission allowances from spot and futures market on a weekly basis for the current and the two subsequent years under uncertainty of allowance demand that accrues at the end of each year.

We assume that the firm only buys emission allowances and does not sell excess allowances, which is in line with organizational constraints of large industrial companies Zapletal et al. (2022). Instead, excess allowances are "banked" by being carried over from one compliance period to the next. The firm can buy emission allowances in the spot and futures markets. We assume that the firm only trades specific futures contracts, for example, the last future contract of any compliance period. This assumption is based on the observation that the December futures contracts tend to have significantly larger volumes than other EUA futures contracts.

We also assume that the market is always in contango, meaning that futures prices are always higher than spot and increasing with their time to maturity. This assumption is supported by Bredin and Parsons (2016) who argue that, in contrast to other commodities, theory of storage arguments cannot explain the difference, since holding EUAs is effectively free.

A common issue with modeling trading operations as stochastic-dynamic optimization problems is that even minor differences in expected prices invite speculative trading, effectively rendering the problem to become unbounded. To circumvent this issue, Gong and Zhou (2013) model carbon prices as a Markov chain and show that the process forms a martingale.In this work, we propose a data-driven approach by constructing a Markov process directly from empirical log-returns in such a way that the resulting discrete-state process becomes a martingale.

4.1 Model of Carbon Price and Emission Uncertainty

We model uncertainty in spot and futures prices as well as carbon emissions as a random data process.

4.1.1 Allowance demand.

The firm's need for emission allowances relies on the amount of CO2 which have been emitted within the current compliance period as well as emissions predicted for future compliance periods. We assume that emissions are closely tied to the production rate of the firm, e.g., how much steel, chemicals, cement, etc., is produced over a certain period of time. Changes in production rate are largely driven by changes in customer demand in the short-term as well as technological changes in the long-term, so that the need for emission allowances is uncertain to some extent.

We assume that the firm has access to a forecast of needed emission allowances. A model that is well-suited to model the evolution of forecast errors is the martingale model of forecast evolution, which was independently developed by Heath and Jackson (1994) and Graves et al. (1998). The model roots back the work by Hausman (1969) who model changes to forecasts as a sequence of independent log-normally distributed random variables. We adopt the latter approach in assuming that the relative change in forecasts is the same across all compliance periods, meaning that if we increase our forecast for the current compliance period, we also increase our forecast for any future compliance period, and vice versa.

A natural feature of the MMFE model is that changes in forecasts accumulate as time moves on, so that forecasts for the near future are more accurate than those for the more distant future, which also implies that future forward buying decisions are exposed to less quantity risk.

4.1.2 Spot and futures prices.

Denote $P_{t,j}$ as the price of a future in period t with maturity in period j, with $j > t$ and p_t as the spot price. The theory of storage explains the difference between futures and spot prices of a commodity by its cost-of-carry, that is the cost and benefit of holding the commodity now rather than later. The cost-of-carry is defined in terms of the risk-free interest rate, the cost of storing the commodity, and the convenience yield. As storage cost are effectively zero for emission allowances, futures prices on emission allowances are given by the following no-arbitrage condition,

$$\ln P_{tj} = \ln p_t + (r^f_{t,j} + r^y_{t,j})(j-t),\ j > t, \tag{1}$$

where $r^f_{t,j}$ is the risk-free interest rate and $r^y_{t,j}$ is the convenience yield (see Palao and Pardo 2021). As the convenience yield is unknown in reality, we will work with a simplified model that combines the two rates into a single factor r_t, which naturally leads to a linear model of futures price dependent on time to maturity,

$$\ln P_{tj} = \ln p_t + r_t(j-t),\ j > t. \tag{2}$$

An advantage of this formulation is that, we can easily estimate the slope from historical data using linear regression and thereby model the dynamics of futures prices as a function of spot price and slope.

4.2 Stochastic Process

The joint stochastic process of spot price, slope factor, and forecast error characterizes the process of exogenous model parameters. We model this process as a multivariate geometric random walk,

$$\ln S_t = \begin{pmatrix} \ln p_t \\ \ln r_t \\ \ln e_t \end{pmatrix} = \begin{pmatrix} \ln p_{t-1} \\ \ln r_{t-1} \\ \ln e_{t-1} \end{pmatrix} + \varepsilon_t, \quad \varepsilon_t \sim i.i.d, \quad t = 1, ..., T. \tag{3}$$

with known initial state, $S_1 = (p_1, r_1, e_1)'$, and random increments, ε_t, that follow some known joint distribution.

As we are not aiming for closed-form solutions but need the stochastic process for generating scenarios, we will use the distribution of empirical log-returns to model prices in Section 6. This approach has two main advantages:

1. Non-normal returns, which often occur in real-world financial data, can be naturally modeled.
2. The model can capture jumps in returns due to regulatory price changes, e.g., following backloading.

By modeling the joint stochastic process in this way, the model is able to capture important aspects of market dynamics and regulatory changes in emission allowances prices.

4.3 Market Impact Cost Model

To model market impact cost, we adopt the square root model of Almgren and Chriss (2001) which relates the absolute price impact of a trade to the market price in a square root fashion, i.e.,

$$p' - p = \eta\sigma\left(\frac{y}{V}\right)^{\beta}, \tag{4}$$

with $\beta \in (0,1],\ \eta \in (0,1]$ as a tunable parameters, σ as the standard deviation of the price, y as the execution volume and V as the total trading volume within a given trading period (typically a day).

Trading volumes can be estimated from historical volume figures which are often provided alongside daily price statistics. Denote $\hat{V}_{tj}$ as the estimated trading volume in period t of the product with maturity in j, where $\hat{V}_{tt}$ is the estimated spot volume.

As the above function is non-linear, we approximate (4) by a monotone increasing piecewise-constant function with K segments, for some large enough K. Denote y_{tj} as the total quantity of contract j in period t and z_{tjk} as the corresponding quantity in segment k. The market impact cost, Δ_{tj}, from trading quantity y_{tj} of product j in period t can be obtained by solving the following optimization problem,

$$\Delta_{tj} = \min_{z_{tjk}\geq 0} \sum_{k=1}^{K} \eta\sigma_t \left(\frac{k z_{tjk}}{\hat{V}_{tj}}\right)^{\beta}, \tag{5}$$

$$\sum_{k=1}^{K} z_{tjk} = y_{tj}, \tag{6}$$

for $t = 1,\ldots,T,\ j \geq t$. By setting $\beta = 1$, this problem can be simplified to a linear programming problem, which can then be readily incorporated into the multi-stage stochastic programming model defined below. While this simplification reduces the cost function to a linear form, it retains η as a tunable parameter.

4.4 Risk Preferences

In the following, we assume that risk preferences of the firm can be expressed by the nested conditional value-at-risk (CVaR), e.g. see Ruszczyński and Shapiro (2006).

For random profits $X_1,\ldots,X_T$, let $\mathrm{CVaR}_\alpha[X_t|X_{t-1}] = \alpha^{-1}\int_0^\alpha G^{-1}_{X_t|X_{t-1}}(t)dt$ where $G_{X_t|X_{t-1}}$ is the distribution function of X_t conditional on X_{t-1} and define the random variables

$$_{\alpha,\lambda}[X_t|X_{t-1}] = \lambda\mathrm{CVaR}_\alpha[X_t|X_{t-1}] + (1-\lambda)[X_t|X_{t-1}]$$

for $0 \leq \lambda \leq 1$ and $0 < \alpha \leq 1$. For a sequence of random variables, $X_1, \ldots, X_T$, the nested CVaR is defined as

$$X_1 +_{\alpha,\lambda} [X_2 +_{\alpha,\lambda} [X_3 +_{\alpha,\lambda} [\cdots] | X_2] | X_1], \tag{7}$$

i.e., by a convex combination of expectation and CVaR that recursively includes other convex combinations of expectation and CVaR. For $\lambda < 1$, the nested CVaR is a time-consistent extension of conventional CVaR to a dynamic setting (Shapiro 2009).

By varying parameters α and λ the nested CVaR can be fit a wide range of risk preferences. Unlike the variance used in the two-stage setting, nested CVaR has the advantage that it keeps the optimization problems computationally tractable and integrates well with the solution approach discussed in Section 5.

4.5 Model Formulation

We model the decision problem of the industrial buyer as a multi-stage stochastic programming problem,

$$\min_{x_1 \in {}_1(x_0,S_1)} C_1(x_1,S_1) +_{\alpha,\lambda} \left[\min_{x_2 \in {}_2(x_1,S_2)} C_2(x_2,S_2) + \cdots \right.$$
$$\left. \cdots +_{\alpha,\lambda} \left[\min_{x_T \in {}_T(x_{T-1},S_T)} C_T(x_T,S_T) \middle| S_{T-1} \right] \ldots \middle| S_1 \right], \tag{8}$$

where $C_t(\cdot,\cdot)$ is the immediate cost, ${}_t(\cdot)$ is the feasible set in $t = 1, \ldots, T$, and (s_1, x_0) is the deterministic initial state. The probability functionals ${}_{\alpha,\lambda}$ in (8) are conditioned on the collection of random vectors $(S_1, S_2, \ldots, S_{t-1})$, where $S_t = (p_t, r_t, e_t)$.

Furthermore, we assume relatively complete recourse, i.e., that the nested minimization problems are feasible for every x_{t-1}. Note that we formulate the problem in such a way that ${}_t$ depends on previous decisions $x_1, \ldots, x_{t-1}$ only by dependence on x_{t-1} and is independent of S_t. Also, the recursive structure of the above problem guarantees that decisions are non-anticipative, i.e., that x_t only depends on information available up until time t.

In the following, we define the decision variables that are contained in x_t as well as the feasible sets ${}_t$. Unless stated otherwise, we assume that $x_t \geq 0$.

We begin by establishing how the model treats trading decisions as well as the associated cash flows.

Denote by f_{tj} the futures contract position that matures in period j, where $j > t$, and y_{tj} as the buying quantity in period t. Accordingly, y_{tt} denotes spot trading decisions. Also, we assume that only buying is allowed, so that $y_{tj} \geq 0$.

As there are no weekly contracts, but only monthly contracts, and most buyers will likely only consider the December contract, we define the set $\mathcal{J}(t)$ as the set of contracts that are tradable in period t. To model a December contract, for example, we can include in this set a contract for the week in December in which that contract matures.

The portfolio of futures is tracked by the following balance equation,

$$f_{tj} = f_{t-1,j} + y_{tj}, \tag{9}$$

for $t = 1, \ldots, T$ and $j \in \mathcal{J}(t)$. As no money changes hands when futures are traded, we additionally track the cost of buying future j in period t,

$$c_{tj} = c_{t-1,j} + y_{tj} P_{tj} + \Delta_{jt}, \tag{10}$$

for $t = 1, \ldots, T$ and $j \in \mathcal{J}(t)$. These cost are added to the objective during the week when the future matures. In this way, we account for the trade-off between paying the convenience yield and the cost of capital. Moreover, this formulation allows us to measure liquidity risk by keeping track of margin payments.

The immediate cost is given by the discounted cost of buying spot as well as any cash flows from futures settlement,

$$C(x_t, S_t) = \gamma_t (y_{tt} p_t + \Delta_{tt} + c_{t,t-1}), \tag{11}$$

for $t = 1, \ldots, T$, with $\gamma_t = (1+\delta)^{\frac{-t}{52}}$ as discount factor and δ as the cost of capital adjusted by the risk-free rate[1].

The firm receives emission allowances through spot purchases or physical settlement of futures. Denote by b_t all allowances that have been banked and can be used within the current or any future compliance period. Denote d_t as the demand forecast for period t at the beginning of the planning horizon. The portfolio balance of banked emission allowances is given by

$$b_t = b_{t-1} - (d_t + e_t) + y_{tt} + f_{tt}, \tag{12}$$

for $t = 1, \ldots, T$. Note that demand is uncertain, since e_t is random.

5 Solution Method

In this section, we provide some theoretical background of the methods used to solve the multi-stage stochastic programming problem described in the previous section. In particular, we explain how to fit a discrete Markov process to historical returns and how to use stochastic dual dynamic programming (SDDP) to find an optimal policy.

SDDP goes back to the seminal work of Pereira and Pinto (1991). First risk-averse variants were proposed in Shapiro (2009) and Guigues and Römisch (2012). Convergence proofs of the algorithm can be found in Philpott and Guan (2008) for linear problems, Girardeau et al. (2014) for convex problems, and Guigues (2016) for risk-averse convex problems.

We will apply the Markovian version of SDDP described in Löhndorf and Shapiro (2019) and Löhndorf and Wozabal (2021) for the risk-averse case, which relies on discretization of the continuous stochastic process. As we will see, a discrete Markov process allows solving the multi-stage stochastic programming problem to optimality.

[1] We subtract the risk-free rate from the cost-of-capital, as we are using the risk-neutral measure to model the spot price process.

5.1 Fitting a Discrete Markov Process

Let us begin our exposition by describing how we can fit a discrete Markov process from a sample of historical returns.

To fit a discrete Markov process, we follow the two-step approach described in Löhndorf and Wozabal (2023) that first approximates an optimal partition of the unconditional distribution using a second-order learning algorithm and then estimates the transition probabilities between partition means using maximum likelihood estimation.

During the first step, the method chooses partition means of the empirical random process defined in (3), μ_{ti}, $i = 1, ..., N_t$, in a way that the average distance to the state space spanned by the empirical random process is minimized, which can be expressed as the following optimization problem,

$$\min_{\mu_{t1},\dots,\mu_{tN_t}} \int \min_{i=1,\dots,N_t} \|\xi_t - \mu_{ti}\|_p^p dP(\xi_t), \tag{13}$$

where $\|\cdot\|_p$ denotes the p-norm with $p \geq 2$. Given a solution $\mu_{t1}, ..., \mu_{tN_t}$ of (13), the space is partitioned into subsets Γ_{ti} consisting of points which are closest to μ_{ti}. The corresponding sets

$$\Gamma_{ti} := \{\xi_t : \|\xi_t - \mu_{ti}\|_p \leq \|\xi_t - \mu_{tj}\|_p,\ j = 1, ..., N_t,\ j \neq i\}, \tag{14}$$

represent the Voronoi diagram associated with respective means μ_{ti}, $i = 1, ..., N_t$. These means then serve as the states of the discrete Markov process, $\bar{S}_{ti} = \mu_{ti} = \mathbb{E}[S_t | S_t \in \Gamma_{ti}]$, $i = 1, \dots, N_t$, $t = 1, \dots, T$.

As finding the optimum to (13) is an $\mathcal{NP}$-hard problem, Löhndorf and Wozabal (2023) propose a quasi-Newton method which we will adopt here. For this, we draw M random sequences $(\hat{S}_t^m)_{t=1}^T$, $m = 1, ..., M$, from the empirical random process. With $(\theta_m)_{m=1}^M$ being a sequence of step sizes with $0 \leq \theta_m \leq 1$, $m = 1, ..., M$, partition means are updated recursively using

$$\mu_{ti}^m := \begin{cases} \mu_{ti}^{m-1} + \theta_m \left(\hat{S}_t^m - \mu_{ti}^{m-1}\right) & \text{if } i = \operatorname{argmin}\left\{||\hat{S}_t^m - \mu_{tj}^{m-1}||^2,\ j = 1, ..., N_t\right\}, \\ \mu_{ti}^{m-1} & \text{otherwise,} \end{cases} \tag{15}$$

with $\mu_{ti}^0 \equiv 0$, for $i = 1, ..., N_t$, $t = 2, \dots, T$, $m = 1, ..., M$.

Löhndorf and Wozabal (2023) show how to choose $(\theta_m)_{m=1}^M$ in a way that partition means converge for arbitrary choices of p and demonstrate its effectiveness for solving high-dimensional stochastic programming problems.

Let $\bar{S}_{ti} \equiv \bar{\mu}_{ti}^M$, $i = 1, \dots, N_t$, $t = 1, \dots, T$, denote the set of discrete states that result from (15) after M iterations. Denote Q_t as the $|N_t| \times |N_{t+1}|$ transition matrix between states in t and states in $t+1$ with transition probabilities q_{tij}.

To estimate transition probabilities, we again use the sample of M random sequences $(\hat{S}_t^m)_{t=1}^T$, $m = 1, ..., M$, from the empirical random process. Then, we count the frequency of transitions to cells in the Voronoi diagram Γ_{t+1}.

To ensure that the corresponding discrete process remains a martingale, we estimate transition probabilities going backwards in time, which allows us to correct the location of discrete states to match to the expected value of the successor states.

The maximum likelihood estimate of the probability of a state transition from $\bar{S}_{ti}$ to $\bar{S}_{t+1,j}$ is given by

$$q_{tij} = \frac{1}{M} \sum_{m=1}^{M} 1_{\Gamma_{tj}}(\tilde{S}^m_{t+1} | \bar{S}_{ti}),\ i = 1, \ldots, N_t,\ j = 1, \ldots, N_{t+1},\ t = T-1, \ldots, 1. \tag{16}$$

Since we need to correct the nodes in $t+1$ before we can compute the conditional expectation of the nodes in t, we begin in $T-1$ and work backwards. However, before correcting node i in t, we estimate transition probabilities, q_{tij}, using (16), and then set

$$\bar{S}_{ti} \equiv \sum_{j=1}^{N_{t+1}} q_{tij} \bar{S}_{t+1,j},\ i = 1, \ldots, N_t,\ t = T-1, \ldots, 1. \tag{17}$$

In this way, we obtain a discrete Markov process that is both a martingale and exhibits (approximately) minimal distance to the empirical random process.

5.2 Markovian Stochastic Dual Dynamic Programming

Let us briefly introduce SDDP, which we will use to solve the multi-stage stochastic programming problem defined in (8).

Given that $\bar{S}_t$ is Markovian, $\mathcal{P}_t(\bar{S}_t | \bar{S}_2, \ldots, \bar{S}_{t-1}) = \mathcal{P}_t(\bar{S}_t | \bar{S}_{t-1})$, we can formulate dynamic programming equations for the multi-stage stochastic programming problem,

$$V_t(x_{t-1}, \bar{S}_t) = \min_{x_t \in \mathcal{X}_t(x_{t-1}, S_t)} \left\{ C_t(x_t, S_t) + \mathcal{V}_{t+1}(x_t, \bar{S}_t) \right\},\ t = 2, ..., T, \tag{18}$$

where $C_t(\cdot, \cdot)$ is the (convex) immediate cost function and

$$\mathcal{V}_{t+1}(x_t, \bar{S}_t) = \mathbb{U}_{\alpha, \lambda | \bar{S}_t} \left[V_{t+1}(x_t, \bar{S}_{t+1}) \right],\ t = 2, ..., T, \tag{19}$$

with $\mathcal{V}_{T+1} \equiv 0$. We refer to functions $V_t(\cdot, \cdot)$ as value functions and $\mathcal{V}_t(\cdot, \cdot)$ as expected value functions.

The value functions, $V_t(x_{t-1}, \bar{S}_t)$, are convex in x_{t-1} if $\min\{C_t(x_t, \bar{S}_t) + \mathcal{V}_{t+1}(x_t, \bar{S}_t) | x_t \in \mathcal{X}_t(x_t, \bar{S}_{t-1})\}$ are convex in x_{t-1}. Convexity is preserved by the probability functionals defined in (8). The rest follows from backwards induction.

SDDP builds an outer approximation of the expected value function using cutting planes, which approximate the value from below. We refer to Löhndorf and Wozabal (2021) for a more detailed exposition of how cutting planes are used for value function approximation.

The algorithm alternates between a forward pass and a backward pass for a finite number of iterations, $i = 1, \ldots, I$, during which it adds cutting planes in states that are visited by the optimal policy.

During the forward pass, the algorithm generates a sample of states of the discrete stochastic process, $(\tilde{\tilde{S}}_t^i)_{t=1}^T$, along with a sample of incumbent decisions,

$$\tilde{x}_t^i = \min_{x_t \in \mathcal{X}_t(\tilde{x}_{t-1}^n, \tilde{\tilde{S}}_t^i)} \left\{ C(x_t, \tilde{\tilde{S}}_t^i) + \hat{\mathcal{V}}_t^{i-1}(x_t, \tilde{\tilde{S}}_t^i) \right\}, \tag{20}$$

with $\hat{\mathcal{V}}^0 \equiv 0$. During the backward pass, SDDP revisits these incumbent decisions to recursively improve the approximation,

$$\hat{V}_t(\tilde{x}_{t-1}^i, \bar{S}_{nt}) = \min_{x_t \in \mathcal{X}_t(\tilde{x}_{t-1}^i, \bar{S}_{nt})} \left\{ C(x_t, \bar{S}_{nt}) + \hat{\mathcal{V}}_t(x_t, \bar{S}_{nt}) \right\}, \tag{21}$$

with $\hat{\mathcal{V}}_T \equiv 0$ for $n = 1, ..., N_t$, $t = T, ..., 2$, $i = 1, ..., I$.

Since new cutting planes tighten the value function in regions that get explored under the incumbent policy, it can be shown that the algorithm converges to the optimal policy in a finite number of steps if the problem is linear (Philpott and Guan 2008, Guigues 2016).

To approximate the value function under the nested CVaR, we adopt the method described in Philpott et al. (2013) who propose to change the probability weights of the cutting planes during the backward pass, which gives those cuts more weight whose values at $\hat{x}_t^n$ are below the VaR. We refer to Shapiro (2011), Philpott et al. (2013), or Löhndorf and Wozabal (2021) for a detailed description.

6 Results

In this section, we show how an industrial buyer tasked with carbon trading can use the model to buy emission allowances as well as analyze the firm's exposure to risk. We also compare the model to a simple spot-buying strategy.

6.1 Data

The dataset used in this work was compiled from two major sources: the Intercontinental Exchange (ICE) and the European Energy Exchange (EEX). The ICE provided the historical daily closing prices and volumes of the EUA December contracts.. Simultaneously, the EEX data offered the corresponding daily spot prices for the same period.

6.2 Stochastic process

We use daily close prices and volumes of EUA December contracts between January 2010 and October 2022 along with the corresponding daily spot prices.

As the decision model operates in weekly time increments, we aggregate the data by week and store the first price and the total volume in each week. We calculate the implied interest rate, r_t, to model the price difference between spot and futures prices using simple linear regression (without intercept).

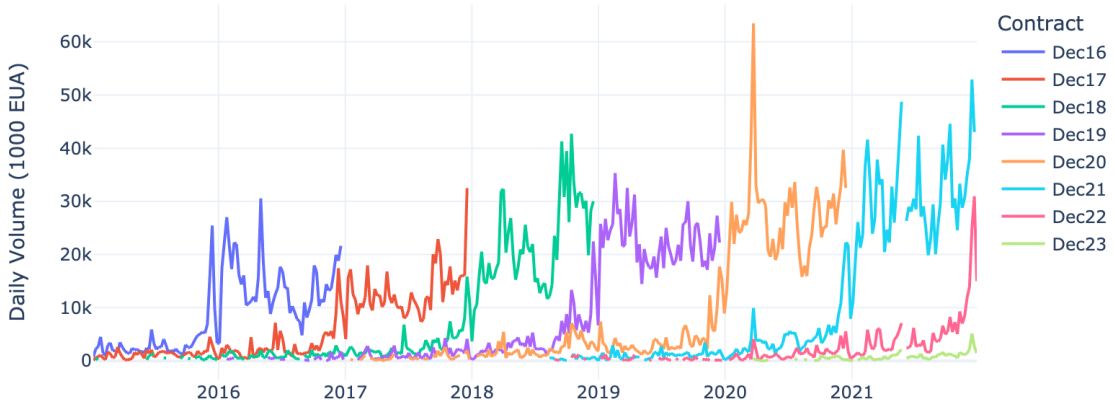


Figure 2 Daily volume (in 1000 EUAs) of December contracts traded between 2015 and 2021

As we do not have access to historical forecast errors, we use the STOXX Europe 600 Chemicals index as proxy for changes in allowance demand. Increases in output of the respective companies likely translates to an increase in allowance demand and vice versa.

The transformed data set provides us with approximately 670 records of parameters (p_t, r_t, d_t), which we can use to create a joint empirical distribution of log-returns.

6.2.1 Market volume forecast.

The average spot volume during the considered time period is relatively steady at 2820 lots (1000 EUA) per day. By contrast, trading volume in futures changes dramatically over time.

Figure 2 shows the daily volume of EUA futures traded between 2012 and 2022 in weekly ticks. We see that the volume of the front year contract is lowest two years ahead of maturity and highest when it is front year contract. There also appears to be some degree of seasonality during the front year period.

To predict the trading volume of each contract on any given trading day, we use linear regression with dummy variables for the number of months to maturity as explanatory variables. The model has an $R^2 = 0.75$ and will be used to inform the market impact model of the total volume required to calculate the price impact of a trade.

6.2.2 Market impact cost.

To model market impact cost, we set $\sigma = 0.0312$ and $\beta = 1$ and $\eta = 0.1$, the latter being of comparable magnitude as the η reported in Almgren et al. (2005).

Start date	Demand risk	Market impact	Avg discounted cost		
4-6			Opt Policy	Front Year	Front Month
2018-01-08	None	0.1	21.1	21.3	25.7
		1.0	20.6	23.9	28.3
2-6	SX4P	0.1	21.0	22.5	28.3
		1.0	20.5	25.0	31.0
2021-01-04	None	0.1	79.3	83.3	82.7
		1.0	79.7	92.5	92.1
2-6	SX4P	0.1	83.2	88.3	91.8
		1.0	84.1	97.5	101.4

Table 1 Cost comparison of the optimal policy with two 1/n spot buying strategies.

6.3 Demand data

For our case study, we assume a large industrial producer that purchases emission allowances to offset 1,000,000 tons of excess carbon emissions, and utilizes 5.0% cost of capital for its financial modeling.

6.4 Experimental Setup

We model the decision process starting on January 4, 2021, which marks the last observed data in our data set and account for futures buying up to three years until December 31, 2023, i.e. 156 weeks in total.

We use the QUASAR stochastic optimization software for algebraic modeling and for solving the multi-stage stochastic programming problem.

The computation time for the problem is negligible. For the decision problem described above, with $T = 156$ stages, $K = 500$ segments for the price impact function, and a Markov process with $N_t = 100,\ t = 2,\ldots,T$ discrete states, it takes roughly 10 seconds to fit the Markov process based on $S = 1,000,000$ samples, and 180 seconds for the stochastic dual dynamic programming algorithm to converge after $I = 500$ iterations on an Apple M1 Mac.

6.5 Comparison with Heuristics

We compare the optimal buying policy with two heuristic policies. The first policy (Front Year) is a 1/n strategy that purchases the expected demand starting at the beginning of the compliance period, whereas the second policy (Front Month) starts only in the last month of the compliance period. The Front Year and Front Month policies distribute the buying volumes evenly over the respective time periods, while adjusting buying volumes to changes in expected demand.

Table 1 summarizes the total cost in million EUR for the three policies for eight different cases, for which we varied the planning horizon (Start Date), whether or not to apply returns from the SX4P index to allowance demand (Demand Risk), and different η to capture price impact (Market Impact).

The two selected start dates mark very different price regimes. While EUA allowances prices were still relatively low at the beginning of 2018 at around €8.00, they quadrupled until the beginning of 2021 after a period of high price volatility between 2018 and 2020.

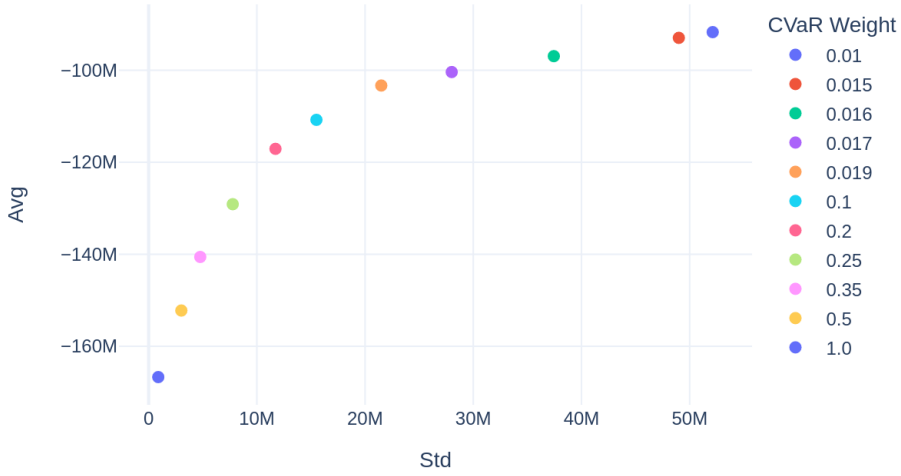


Figure 3 Efficient frontier for varying weights of risk aversion parameter λ

The results show that relying solely on spot trading, as seen in the two heuristic policies, is generally less effective, especially when the transaction costs are high. The optimal policy consistently outperforms these heuristic strategies in all eight cases studied, even in different price regimes and when facing varying levels of demand risk.

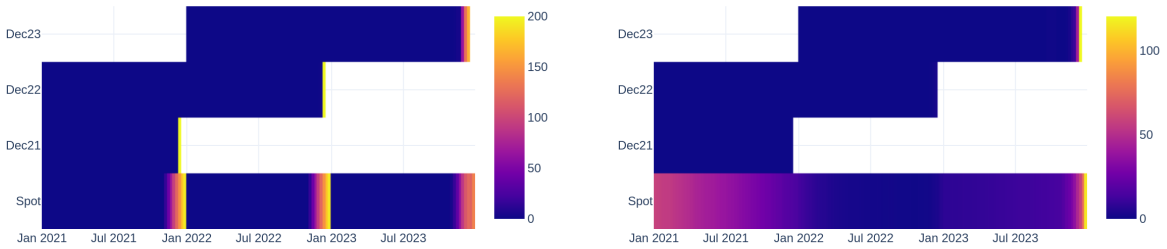
These findings underscore the importance of optimizing the purchase of emission allowances. While heuristic approaches may offer simplicity, their performance tends to be outstripped by a more elaborate optimized policy, particularly in more volatile market conditions.

6.6 Risk Analysis

Let us study the effect that the risk preferences of a firm have on the distribution of the average discounted cost.

Figure 3 illustrates the trade-off between risk and reward for different levels of risk aversion. The figure reveals an efficient frontier that depicts the average discounted cost and its standard deviation (as a proxy for risk) for varying weights of risk aversion parameter λ. Although higher values of λ effectively reduce risk, it comes at a high price prompting the firm to buy allowances earlier, leading to higher average discounted costs. At extreme levels of risk aversion, this strategy might not be in the organization's best interest due to excessively high costs.

Hence, the risk aversion parameter λ needs to be carefully calibrated to strike the right balance between risk and expected cost, and it should be set in line with the organization's risk preferences and financial capacity.



(a) Risk neutral case ($\alpha = 0.1,\ \lambda = 0$)

(b) Risk-averse case ($\alpha = 0.1,\ \lambda = 0.016$)

Figure 4 Outlook on future buying decisions from January 2021 to December 2023

6.7 Buying Decisions

The model gives buyers of emission allowances concise decision support, providing them with buying targets and an outlook on the timing and quantity of spot and futures trading.

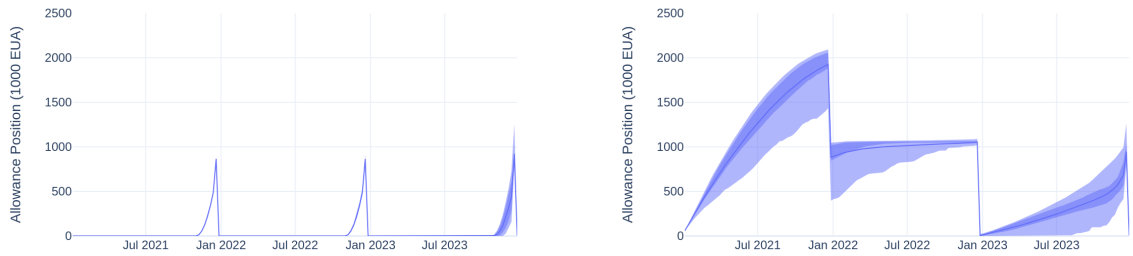
Figure 4 shows the trading volumes for spot and futures contracts throughout the planning horizon for the risk-neutral case (left) and a moderately risk-averse case (right). While we have the date of purchase on the X-axis, the Y-axis indicates the amount of spots and futures (maturing in Dec-21, Dec-22, and Dec-23) purchased. We see that the risk-neutral buyer waits until the last few weeks and then ramps up buying while splitting up trading between the spot and the December contract. The risk-averse buyer, by contrast, secures prices early in the planning horizon in the spot market and then complements this purchase by buying futures shortly before maturity. These results agree with the analytical results presented in Section 3.

This information indicates that the model successfully assists buyers of emission allowances to make informed decisions about when and how much to buy based on their risk preferences. The risk-neutral buyer postpones the purchase until the last few weeks before maturity. This buyer's strategy relies on the premise that they are aware of the time value of money but indifferent to risk and are willing to wait for allowance demand to materialize.

On the other hand, the risk-averse buyer begins securing prices earlier in the planning horizon, predominantly on the spot market. This indicates a strategy that seeks to mitigate potential future price volatility risks. Once these early purchases are secured, the risk-averse buyer supplements their purchases with futures contracts as the maturity date approaches. This mix of early spot purchases and later futures purchases helps them manage the potential risk and uncertainty of the emission allowances market.

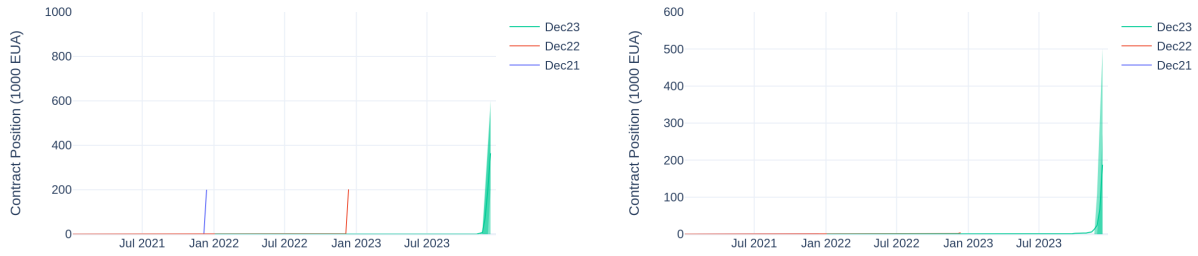
6.8 Portfolio of Allowances and Futures Contracts

In addition to the buyer's view, the model also provides insights into how the optimal portfolio dynamics of allowances and financial instruments would look like.



(a) Risk neutral case ($\alpha = 0.1,\ \lambda = 0$) (b) Risk-averse case ($\alpha = 0.1,\ \lambda = 0.016$)

Figure 5 Banked allowances between January 2021 to December 2023

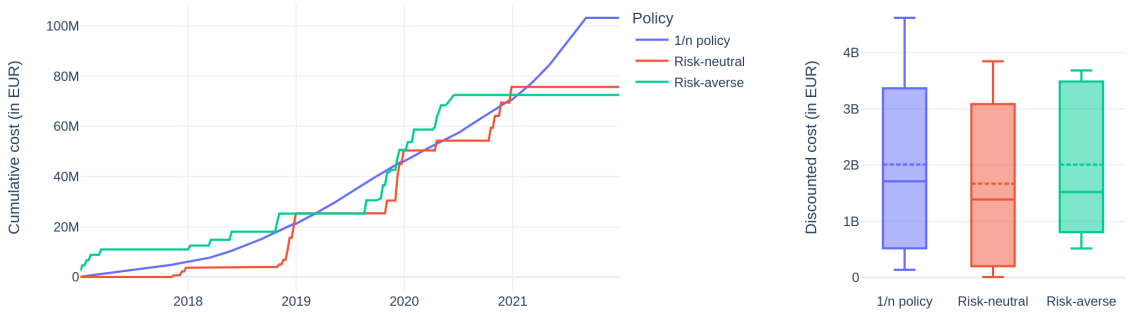


(a) Risk neutral case($\lambda = 0$) (b) Risk-averse case ($\lambda = 0.016$)

Figure 6 Portfolio of futures between January 2021 to December 2024

Figures 5 and 6 show how the model manages the portfolio of banked allowances and futures contracts over time for the risk-neutral case (left) and the risk-averse case (right). While the solid line represents the average, various shades indicate different quantiles (i.e., 0.05,0.25,0.75 and 0.95). The risk-neutral buyer appears to use the futures market more frequently than the risk-averse buyer, a strategy that might initially seem counter-intuitive for commodities. However, it's essential to consider the unique characteristics of emission allowances. Unlike many other commodities, emission allowances can be stored indefinitely without incurring any storage costs, which shifts the focus from physical storage considerations to financial aspects such as the time value of money and liquidity perception.

In this context, the risk-neutral buyer seems to favor futures contracts, likely driven by higher perceived liquidity in the futures market and relatively lower cost of capital when futures are close to maturity. The risk-averse buyer, on the other hand, shows a preference for the spot market, which can be attributed to the buyer's desire to mitigate price volatility risks. Given that the cost of storing emission allowances is



(a) Cumulative cost (b) Discounted cost

Figure 7 Cost from emission allowance trading for different trading policies

effectively zero, the risk-averse buyer considers the spot market more attractive, mainly if the cost of capital is relatively low.

6.9 Backtest with Historical Data

The model's performance was evaluated through a backtest conducted on historical data from January 2017 to December 2021. Over this rolling horizon, the model was applied with a planning period of three years, while recording only the first-stage decision. Unlike before, we assume deterministic allowance demand, making market prices the only random factor. The main objective of the model, therefore, becomes minimizing expected risk-adjusted discounted purchase cost while mitigating market impact costs.

The performance of the risk-neutral (with $\alpha = 0.1$ and $\lambda = 0$) and the risk-averse policy (with $\alpha = 0.1$ and $\lambda = 0.016$) was benchmarked against the 1/n Front Year heuristic, which initiates the purchase of emission allowances on the spot market at the beginning of each compliance period.

The result of this backtest for all three policies is illustrated in Figure 7. The left panel depicts cumulative cost, calculated as the inner product of realized prices and buying volumes, without taking market impact costs into account. The right panel represents the distribution of costs, discounted at the cost-of-capital rate to the onset of the compliance period.

The results reveal that the risk-averse policy performs best on average, followed closely by the risk-neutral policy, while the heuristic policy performs the worst. However, when we account for the cost of capital, the risk-neutral policy has the lowest average/median discounted costs gaining an advantage over the risk-averse and the heuristic policy by delaying cash outflows to the end of the compliance period. We also observe that while the median/average discounted costs of the optimal risk-averse policy is higher than the optimal risk-neutral policy, it has a lower variance incorporating the risk-preference of the buyer. This outcome is quite noteworthy, particularly in light of the market's overall upward trend and the exclusion of market impact costs from the analysis.

However, one should approach these results with a degree of caution. The model is not designed to predict or respond to market trends but rather considers the balance between market liquidity and cost of capital under spot and futures price uncertainty. Thus, the purpose of the backtest is not to contemplate the superiority of the optimal policy, but should serve as a reality check for the model and a validation of its soundness.

In conclusion, the backtest shows that the model yields sensible outcomes by confirming the result from the simulation analysis. The optimal policy for a risk-neutral buyer is to wait and then use a mix of spot and futures contracts to satisfy the firm's allowance demand, whereas for the risk-averse buyer, the optimal policy requires early investment to shield the firm against rising prices.

7 Conclusion

In this work, we present a new model that provides industrial buyers with a comprehensive approach to optimal procurement of emission allowances. The model takes the perspective of a buyer who decides about the timing and quantity of buying emission allowances on spot and futures markets to compensate for the firm's carbon emissions.

We formulate the problem as a two-stage analytical model followed by a multi-stage stochastic programming problem that considers price and demand uncertainty, forward curve dynamics, transaction costs, and cost of capital. To solve the problem numerically, we develop an approach based on Markovian stochastic dual dynamic programming and show how a discrete random process of price and demand can be estimated from historical returns.

We demonstrate how the model can be used by an industrial buyer to determine the right quantity and timing of buying emission allowances in spot and futures markets as well as to measure the firm's exposure to various sources of risk.

The results indicate that the best strategy for a large industrial buyer of carbon allowances is to have a mix of both spots and futures, with futures dominating the spots in the portfolio. In contrast, the optimal strategy for a relatively smaller buyer of carbon allowances is to have a portfolio with a higher proportion of spots. Further, we also find that a strategy wherein the buyer bides their time and then leverages a combination of spot and futures contracts to satisfy the firm's allowance demand yields the best outcome for a risk-neutral buyer, whereas, for a risk-averse buyer, it's always better to buy upfront.

Our findings underline the value of our model as a tool to inform decision-making around emission allowances procurement, efficiently navigating through elements of market liquidity, cost of capital, and price uncertainty. As such, the approach allows for effective risk management, cost optimization, and, ultimately, a greater alignment with the firm's broader sustainability and financial objectives.

Future work might expand the model to include additional strategies, such as the procurement of call options. Drawing from the work of Carmona and Hinz (2011), who have developed risk-neutral models for pricing emission allowance futures, could be a promising avenue for such development.

References

Almgren R, Chriss N (2001) Optimal execution of portfolio transactions. *Journal of Risk* 3:5–40.

Almgren R, Thum C, Hauptmann E, Li H (2005) Direct estimation of equity market impact. *Risk* 18(7):58–62.

Ampudia M, Bua G, Kapp D, Salakhova D (2022) The role of speculation during the recent increase in eu emissions allowance prices.

Benz E, Trück S (2009) Modeling the price dynamics of CO2 emission allowances. *Energy Economics* 31(1):4–15.

Bredin D, Parsons J (2016) Why is spot carbon so cheap and future carbon so dear? The term structure of carbon prices. *The Energy Journal* 37(3).

Carmona R, Hinz J (2011) Risk-neutral models for emission allowance prices and option valuation. *Management Science* 57(8):1453–1468.

Devalkar SK, Anupindi R, Sinha A (2011) Integrated optimization of procurement, processing, and trade of commodities. *Operations Research* 59(6):1369–1381.

Devalkar SK, Anupindi R, Sinha A (2018) Dynamic risk management of commodity operations: Model and analysis. *Manufacturing & Service Operations Management* 20(2):317–332.

Eckert V, Buli N (2022) Explainer: How margin calls came to threaten europe's energy firms. *Reuters* URL `https://www.reuters.com/business/energy/how-margin-calls-came-threaten-europes-energy-firms-2022-09-16/`.

Ellerman AD, Convery FJ, De Perthuis C (2010) *Pricing carbon: the European Union emissions trading scheme* (Cambridge University Press).

Geman H (2009) *Commodities and commodity derivatives: modeling and pricing for agriculturals, metals and energy* (John Wiley & Sons).

Girardeau P, Leclere V, Philpott A (2014) On the convergence of decomposition methods for multistage stochastic convex programs. *Mathematics of Operations Research* 40(1):130–145.

Gong X, Zhou SX (2013) Optimal production planning with emissions trading. *Operations Research* 61(4):908–924.

Graves SC, Kletter DB, Hetzel WB (1998) A dynamic model for requirements planning with application to supply chain optimization. *Operations Research* 46(3supplement-3):35–49.

Guigues V (2016) Convergence analysis of sampling-based decomposition methods for risk-averse multistage stochastic convex programs. *SIAM Journal on Optimization* 26(4):2468–2494.

Guigues V, Römisch W (2012) Sampling-based decomposition methods for multistage stochastic programs based on extended polyhedral risk measures. *SIAM Journal on Optimization* 22(2):286–312.

Hausman WH (1969) Sequential decision problems: A model to exploit existing forecasters. *Management Science* 16(2):93–111.

Heath DC, Jackson PL (1994) Modeling the evolution of demand forecasts ith application to safety stock analysis in production/distribution systems. *IIE transactions* 26(3):17–30.

Huettemann T, Fort J, Weinstein J (2010) Comment s71610-70 URL https://www.sec.gov/comments/s7-16-10/s71610-70.pdf.

Kleindorfer PR (2009) Integrating physical and financial risk management in supply management. *Risk Management in Commodity Markets: From Shipping to Agriculturals and Energy* 33–50.

Kleindorfer PR (2010) Commodity risk. *Encyclopedia of Quantitative Finance* .

Koch N, Grosjean G, Fuss S, Edenhofer O (2016) Politics matters: Regulatory events as catalysts for price formation under cap-and-trade. *Journal of Environmental Economics and Management* 78:121–139.

Kouvelis P, Li R, Ding Q (2013) Managing storable commodity risks: The role of inventory and financial hedge. *Manufacturing & Service Operations Management* 15(3):507–521.

Kouvelis P, Pang Z, Ding Q (2018) Integrated commodity inventory management and financial hedging: A dynamic mean-variance analysis. *Production and Operations Management* 27(6):1052–1073.

Löhndorf N, Shapiro A (2019) Modeling time-dependent randomness in stochastic dual dynamic programming. *European Journal of Operational Research* 2(1):650–661.

Löhndorf N, Wozabal D (2021) Gas storage valuation in incomplete markets. *European Journal of Operational Research* 288(1):318–330.

Löhndorf N, Wozabal D (2023) A quasi-newton algorithm for optimal discretization of markov processes. *Working paper published on Optimization Online* URL https://optimization-online.org/?p=22505.

Mandl C, Minner S (2020) Data-driven optimization for commodity procurement under price uncertainty. *Manufacturing & Service Operations Management* .

Mandl C, Nadarajah S, Minner S, Gavirneni S (2022) Data-driven storage operations: Cross-commodity backtest and structured policies. *Production and Operations Management* 31(6):2438–2456.

Palao F, Pardo Á (2021) The inconvenience yield of carbon futures. *Energy Economics* 101:105461.

Pereira M, Pinto L (1991) Multi-stage stochastic optimization applied to energy planning. *Mathematical Programming* 52(2):359–375.

Philpott A, de Matos V, Finardi E (2013) On solving multistage stochastic programs with coherent risk measures. *Operations Research* 61:957 – 970.

Philpott A, Guan Z (2008) On the convergence of stochastic dual dynamic programming and related methods. *Operations Research Letters* 36(4):450–455.

Rong A, Lahdelma R (2007) Co2 emissions trading planning in combined heat and power production via multi-period stochastic optimization. *European Journal of Operational Research* 176(3):1874–1895.

Ruszczyński A, Shapiro A (2006) Conditional risk mappings. *Mathematics of Operations Research* 31(3):544–561.

Seifert J, Uhrig-Homburg M, Wagner M (2008) Dynamic behavior of CO2 spot prices. *Journal of Environmental Economics and Management* 56(2):180–194.

Shapiro A (2009) On a time consistency concept in risk averse multistage stochastic programming. *Operations Research Letters* 37(3):143 – 147.

Shapiro A (2011) Analysis of stochastic dual dynamic programming method. *European Journal of Operational Research* 209(1):63 – 72.

Tietenberg T (2010) *Emissions trading: principles and practice* (Routledge).

Yuan Q, Chen Y, Yang J, Zhou Y (2018) Joint control of emissions permit trading and production involving fixed and variable transaction costs. *Production and Operations Management* 27(8):1420–1454.

Zapletal F, Šmíd M, Kopa M (2020) Multi-stage emissions management of a steel company. *Annals of Operations Research* 292(2):735–751.

Zapletal F, Šmíd M, Kozmík V (2022) Multi-stage stochastic optimization of carbon risk management. *Expert Systems with Applications* 201:117021.