Journal of Computational Finance **26**(1), 37–80
DOI: 10.21314/JCF.2022.018

Copyright Infopro Digital Limited 2022. All rights reserved. You may share using our article tools. This article may be printed for the sole use of the Authorised User (named subscriber), as outlined in our terms and conditions. https://www.infopro-insight.com/termsconditions/insight-subscriptions

Research Paper

Optimal trade execution with uncertain volume target

Julien Vaes and Raphael Hauser

Mathematical Institute, University of Oxford, Andrew Wiles Building,
Radcliffe Observatory Quarter, Woodstock Road, Oxford OX2 6GG, UK;
emails: julien.vaes@maths.ox.ac.uk, hauser@maths.ox.ac.uk

(Received May 27, 2020; revised September 13, 2021; accepted January 10, 2022)

ABSTRACT

In a seminal 2001 paper on the optimal execution of portfolio transactions, Almgren and Chriss defined the optimal trading strategy to liquidate a fixed volume of a single security under price uncertainty. Yet there exist situations, such as in the power market, in which the volume to be traded can only be estimated and becomes more accurate when approaching a specified delivery time. During the course of execution a trader should then constantly adapt their trading strategy to meet their fluctuating volume target. In this paper we develop a model that accounts for volume uncertainty and we show that a risk-averse trader benefits from delaying their trades. More precisely, we argue that the optimal strategy is a trade-off between early and late trades in order to balance risk that is associated with both price and volume. By incorporating a risk term related to the trading volume, the static optimal strategies that are suggested by our model avoid the blowup in the algorithmic complexity usually associated with dynamic programming solutions while yielding competitive performance.

Keywords: trading strategy; transactions costs; volume uncertainty; risk-aversion; conditional value-at-risk (CVaR).

Corresponding author: J. Vaes

Print ISSN 1460-1559 | Online ISSN 1755-2850
© 2022 Infopro Digital Risk (IP) Limited

1 INTRODUCTION

The optimal execution problem refers to the problem of finding the best trading strategy to ensure the transition from one portfolio to another within an allocated period of time. A trading strategy consists of buy and sell orders that are exchanged on the dedicated market. This problem of optimal execution is germane to markets where liquidity is insufficient. Indeed, in insufficiently liquid markets, price dynamics are sensitive to large trades. Hence, a rational trader should take into consideration the impact of their own trades, which is assumed to be adverse to their trading position in the sense that it increases their trading cost. In their seminal paper on optimal execution, Almgren and Chriss (2001) proposed a model that captures both the inherent random evolution of the prices and the effect of trades on the price dynamics. By adopting the risk–return trade-off defined by Markowitz (1952) for portfolio hedging, Almgren and Chriss (2001) added the nuance of risk that was missing from the optimal order execution problem stated in Bertsimas and Lo (1998). Almgren and Chriss showed that the more risk-averse a trader is, the quicker the acquisition of the desired position should be carried out in order to avoid the risk associated with the price dynamics. These strategies are static and time consistent: they are determined in advance of trading in the sense that they depend only on the information available prior to the start of the execution period. Also, if a trader recomputes the optimal strategy mid-course, the future trading plan remains unchanged, since future price fluctuations are assumed to be independent of the past realizations. Nevertheless, if a trader defines as part of their initial strategy a mid-course updating rule, aggressive in-the-money strategies in the sense of Kissell and Malamut (2005) lead to better risk–return trade-offs (Almgren and Lorenz 2007; Lorenz and Almgren 2011). As the variance penalizes both advantageous and adverse trading cost outcomes without any distinction, other risk measures have also been used in the literature, such as the expected utility (Schied and Schöneborn 2009; Schöneborn 2016) or the α-conditional value-at-risk (CVaR_α) (Butenko *et al* 2005; Feng *et al* 2012; Krokhmal and Uryasev 2007).

In the abovementioned papers, the total volume of securities to be traded is deterministic and is given as a parameter. However, in many situations a trader only disposes of an estimate of this volume, which varies throughout the course of execution; there is thus a need for a model that combines the risk associated with both the price dynamics and the uncertainty of the volume target. For example, the manager of an open-ended investment fund may issue or redeem shares to avoid the asset–liability mismatch that could ensue if shares were only traded in the secondary market. Orders are usually collected and aggregated during the trading day, and the issue or redemption of the net share order creates the need to acquire or dispose of a certain volume of the fund's investment assets, typically over the next two trading days. Our model allows the fund manager to front-load some of these trades and execute orders closer to the time they are first placed, before all share orders are netted out at the end of the day, thus greatly increasing the market liquidity for the fund's investors. Another example is the clearing of power futures markets, which must result in an outstanding volume close to the realized demand in every given delivery period, as the spot market has limited liquidity due to the physical constraints of the power plants used for generation.

To the best of our knowledge no previous model considers both these sources of uncertainty. Related papers are Cheng *et al* (2017, 2019) and Bulthuis *et al* (2017), which investigate the optimal strategy with uncertain order fills, ie, the risk for an order to be filled either incompletely or in excess (the latter being considered mainly for mathematical rather than practical reasons). In significant contrast to our situation, in their setting the magnitude of the uncertainty of the order fills is assumed to be either constant or proportional to the order size. In our case the uncertainty is independent of the trader's decisions: the updates in the volume forecast depend exclusively on extraneous variables. Our model consequently assumes that both sources of uncertainty (ie, the price dynamics and the forecast updates of the volume target) are inherent to the market and independent of the trader's trading strategy. Moreover, the model presented in this paper allows for more flexibility than the approaches adopted by Cheng *et al* (2017, 2019) and Bulthuis *et al* (2017), as it does not rely on the assumption that the trading periods are homogeneous; for instance, it allows for the asset's liquidity to vary over the course of the execution period. Finally, unlike the alternative approaches, the strategies obtained with our model are static, and their computation avoids relying on a computationally intensive dynamic programming approach.

The rest of this paper is structured as follows. In Section 2 we propose a model that incorporates the volume uncertainty in the CVaR_α equivalent formulation of the risk–return trade-off model by Almgren and Chriss (2001). We establish the relationship between the viability of the market and the uniqueness of a convex set of optimal trading strategies. We also adapt models from the literature to our problem and compare their relative advantages. Section 3 provides numerical evidence that our model achieves significantly better mean–CVaR_α trade-offs than when volume uncertainty is neglected, and it compares our approach against the current alternatives. We then illustrate how our model can be applied to power trading in Great Britain,[1] where power wholesalers exchange power future contracts to hedge against the risk associated with the volatility of the power demand. Finally, Sections 4 and 5

[1] The National Grid does not cover Northern Ireland, which is part of a single electricity market with the Republic of Ireland.

discuss the results, propose avenues for further research on this subject and state the conclusions of the paper.

2 MODEL

In order to fit into the framework described in the introductory examples, we formulate the optimal execution problem from the point of view of a trader who desires to acquire their position within a fixed time horizon; this is equivalent to the liquidation problem presented in Almgren and Chriss (2001) up to a change of sign in the temporary and permanent impacts that are induced by the trades, as presented hereinafter.

2.1 Optimal execution under price uncertainty

We use the following notation: $d_T > 0$ is the total volume to be traded by time T over m execution periods, S_0 is the initial security price and τ_i is the length of the trading period between the two consecutive discrete decision times t_{i-1} and t_i; with slight abuse of notation we also use τ_i in the sense of the ith trading period. Like Almgren and Chriss (2001) we consider that the price dynamics follows an arithmetic random walk, where the price step between decision times t_{i-1} and t_i varies according to the distribution of a random variable ξ_i. Finally, if, additionally to these assumptions, the temporary and permanent market impacts induced by the trades are inserted into the price dynamics, we find that the prices evolve as follows for $i \in \{1, \dots, m\}$:

$$\left.\begin{aligned} S_i &= S_{i-1} + \xi_i + \tau_i g\left(\frac{n_i}{\tau_i}\right), \\ \tilde{S}_i &= S_{i-1} + h\left(\frac{n_i}{\tau_i}\right), \end{aligned}\right\} \tag{2.1}$$

where S_i is the security price at decision time t_i, n_i is the volume of securities traded during trading period τ_i (bought if $n_i > 0$ and sold if $n_i < 0$), $\tilde{S}_i$ is the effective security price for the trades executed during trading period τ_i, and g and h respectively model the permanent and temporary price impacts as a function of the average trading rate over the trading interval. The liquidation cost is then a random variable given by (2.2), here formulated from a buyer's perspective:

$$\begin{aligned} \mathcal{C}(\boldsymbol{n}) &:= \sum_{i=1}^{m} n_i \tilde{S}_i - d_T S_0 \\ &= \sum_{i=1}^{m} \left(\xi_i + \tau_i g\left(\frac{n_i}{\tau_i}\right)\right)\left(d_T - \sum_{k=1}^{i} n_k\right) + \sum_{i=1}^{m} n_i h\left(\frac{n_i}{\tau_i}\right), \end{aligned} \tag{2.2}$$

where $\boldsymbol{n} = [n_1, \dots, n_m]^{\mathrm{T}}$ is the vector of the trade volumes n_i of period τ_i. Given the mean–variance framework of Markowitz (1952) and given a risk aversion that is parameterized by $\lambda_{\mathrm{Var}} \geqslant 0$, the optimal trading strategy deriving from the model by Almgren and Chriss (2001) is obtained by solving the following optimization problem:

$$\left.\begin{array}{ll} \underset{\boldsymbol{n}}{\text{minimize}} & \mathbb{E}[\mathcal{C}(\boldsymbol{n})] + \lambda_{\mathrm{Var}} \operatorname{Var}[\mathcal{C}(\boldsymbol{n})] \\ \text{subject to} & \mathbf{1}^{\mathrm{T}}\boldsymbol{n} = d_T, \end{array}\right\} \tag{2.3}$$

where $\mathbf{1}$ is a vector of "1"s, ie, $[1, \dots, 1]^{\mathrm{T}}$.

Gatheral and Schied (2013, Corollary 1) proved that the continuous formulation of this model is free of price manipulations, which is tantamount to the market viability, if the following conditions are satisfied:

(i) g is a linear nondecreasing function, ie, $g(v) = \gamma v$ with $\gamma \geqslant 0$; and

(ii) the function $f : x \mapsto xh(x)$ is convex.

These conditions are observed empirically, as shown in Almgren *et al* (2005). For the discrete formulation of Almgren and Chriss (2001), Proposition 2 in Huberman and Stanzl (2004) provides the following necessary conditions for a market with time-independent price-impact functions to be viable:

(i) the function g is linear; and

(ii) we have

$$h\left(\frac{q}{\tau}\right) - h\left(\frac{-q}{\tau}\right) \gtreqless g\left(\frac{q}{\tau}\right) \quad \text{for } q \gtreqless 0, \tag{2.4}$$

where q is the volume traded and τ is the length of each trading period, ie, for all $i \in \{1, \dots, m\}$: $\tau_i = \tau$; an interpretation of this condition is provided later on.

Almgren and Chriss (2001) assume linear permanent and temporary impact functions. We make similar assumptions for our model. However, we allow for the liquidity parameters to vary over time. Hence, we have time-dependent price-impact functions. The permanent and temporary impact functions g_i and h_i that are related to the trading period τ_i are noted as follows:

$$\left.\begin{aligned} g_i(v) &= \gamma_i v, \\ h_i(v) &= \epsilon_i \operatorname{sgn}(v) + \eta_i v, \end{aligned}\right\} \tag{2.5}$$

where v is the average trading rate over trading period τ_i. This choice of functions is relevant in that parameters ϵ_i and η_i might be considered as the fixed and variable costs of trading. A reasonable estimate for ϵ_i is the sum of half the bid–ask spread and the trading fees (Almgren and Chriss 2001). The parameter η_i can be interpreted as the gradient of a linear model for the volumes of orders in the limit order book as a function of the deviation from the best limit order; we assume that this parameter is the same when taking either short or long positions. As a consequence of the temporary impact, any market participant incurs an additional cost of

$$n h_i \left(\frac{n}{\tau_i} \right) = \epsilon_i |n| + \frac{\eta_i n^2}{\tau_i} \tag{2.6}$$

for trading n units of the security during the trading period τ_i.

2.2 Incorporation of the uncertainty of the volume target

In the case where d_T, the total volume to be traded by time T, only becomes fully deterministically known during the execution period, (2.3) must be adapted in order to ensure that the trader acquires the right number of units of the security. For this to happen we make the assumption that the volume target is perfectly known at the start of the last trading period τ_m. Henceforth we will denote the volume target by a capital letter, ie, D_T, to emphasize the fact that it is a random variable, unlike in Section 2.1.

In this paper we consider that the uncertainty associated with the total volume to be traded is defined as follows: for a given delivery time T, let D_i be the forecast at time t_i of the total volume to be traded by time $t_m = T$:

$$D_i := \mathbb{E}[D_T \mid \mathcal{F}_i], \tag{2.7}$$

where $(\mathcal{F}_i)_{i=1}^m$ is the filtration of sigma algebras that represent the information available at time t_i. Let us assume that two successive forecasts differ by a continuous random variable δ_i for $i < m$. At decision time t_{m-1} the volume target is assumed to be perfectly known. Hence, the forecasts of D_T evolve as follows over course of the entire execution period:

$$\left. \begin{aligned} D_i &= D_{i-1} + \delta_i \quad \text{if } i \in \{1, \dots, m-1\}, \\ D_m &= D_{m-1}, \end{aligned} \right\} \tag{2.8}$$

where $D_m = D_T$, the total volume to be traded by time T.

Obviously, in the presence of volume uncertainty, the mean–variance constraint in optimization problem (2.3) cannot be enforced, since only a rough estimate of the total volume to be traded is known at time t_0, ie, D_0. Nonetheless, the model by Almgren and Chriss (2001) may be useful as a planning tool when deployed in conjunction with a systematic recourse whenever an update on the volume target becomes available. A trader would then, at each trading period, recompute and update their strategy based on the newest available volume forecast. The main flaw in this approach is that a trader is unable to define a static strategy at time t_0 that ensures the satisfaction of constraint in (2.3). Indeed, the volumes to be traded in each trading period must constantly be updated. To get around this issue, we subsequently propose a new way of defining a trader's strategy. Further, to avoid a blowup in the algorithmic complexity of numerical evaluations of the model, we want to avoid approaches based on dynamic programming and to focus on models that account for the impact of recourse actions via the incorporation of a risk term.

First, unlike Almgren and Chriss (2001), who define a trader's strategy as the volumes to be traded in each trading period, we define a trader's strategy as the proportion y_i of the total volume D_T to be acquired over the course of each trading period τ_i. Naturally, the following constraint must hold to enforce the distribution of the entire volume target over all the trading periods:

$$\sum_{i=1}^{m} y_i = 1. \tag{2.9}$$

As a consequence, in the ideal situation where D_T is perfectly predictable, a trader would acquire $y_i D_T$ units of the security during trading period τ_i. Obviously, this definition must be adapted in the situation where D_T is uncertain, as a trader should take recourse in order to satisfy the constraint in (2.3). As a starting point we assume that a trader respects their initial strategy by trading during each trading period τ_i the planned proportion y_i of D_{i-1}, ie, the best estimate of the volume target D_T available. This only partially solves the problem engendered by the forecast updates, since past decisions cannot be modified a posteriori, which means that $\sum_{i=1}^{m} y_i D_{i-1} \neq D_T$. Indeed, if we denote by δ_k, $k < m$, the forecast update reported at the decision time t_k, then the forecast error related to δ_k due to the past decisions amounts to $\delta_k^{\varepsilon} := \delta_k \sum_{r=1}^{k} y_r$. This corresponds to the additional volume (which may be positive or negative) that needs to be traded over the remaining trading periods in order to match with D_T. We predicate that, when an update on the volume target occurs, the trader adjusts their strategy by redistributing the forecast error over the remaining trading periods according to a fixed distribution independent of the decision variables. Considering these corrections upfront allows us to partially account for future recourse actions without having to compute a costly dynamic programming simulation. In our approach we redistribute this additional volume over the future trading periods according to fixed proportions. We treat these fixed proportions as model parameters, and we leave the learning of optimal parameter values to future work. Let $\beta_{k,i}$ denote the proportion of the forecast error δ_k^{ε} to correct for at trading period τ_i. Evidently, $\beta_{k,i} = 0$ if $i \leqslant k$ and $\sum_{i=k+1}^{m} \beta_{k,i} = 1$ for all $k < m$. We will refer to $\boldsymbol{\beta} := (\beta_{k,i}) \in \mathbb{R}^{m-1\times m}$ as the redistribution matrix.

In the case where the volume target is uncertain, we define the cost of a trading strategy as the difference between the trading cost incurred at the end of the execution period by following the initial trading strategy under the consideration of the rough model on their recourse determined by the redistribution matrix $\boldsymbol{\beta}$ and the trading cost ideally obtained in an infinitely liquid market where the entire position D_T is traded at the start of the execution period. With these considerations, a trader's trading cost relating to a strategy $\boldsymbol{y} = [y_1, \dots, y_m]^{\mathrm{T}}$ is the following random variable:

$$\begin{aligned} \mathcal{C}(\boldsymbol{y}) &:= \sum_{i=1}^{m} n_i(\boldsymbol{y})\tilde{S}_i - D_T S_0 \\ &= \sum_{i=1}^{m} \biggl[\biggl(\xi_i + \tau_i g_i\biggl(\frac{n_i(\boldsymbol{y})}{\tau_i}\biggr)\biggr)\biggl(D_0 + \sum_{k=1}^{m} \delta_k - \sum_{k=1}^{i} n_k(\boldsymbol{y})\biggr)\biggr] \\ &\quad + \sum_{i=1}^{m} n_i(\boldsymbol{y}) h_i\biggl(\frac{n_i(\boldsymbol{y})}{\tau_i}\biggr), \end{aligned} \tag{2.10}$$

where $n_i(\boldsymbol{y})$ is a random variable representing the volume to be traded during trading period τ_i:

$$n_i(\boldsymbol{y}) = y_i D_{i-1} + \sum_{k=1}^{i-1} \beta_{k,i}\delta_k^{\varepsilon} = y_i D_0 + \sum_{k=1}^{i-1} \delta_k \biggl(y_i + \beta_{k,i} \sum_{r=1}^{k} y_r\biggr). \tag{2.11}$$

The number of positions $n_i(\boldsymbol{y})$ to be traded at trading period τ_i is composed of both the proportion y_i of the best volume estimate D_{i-1} and the position adjustments related to the previous volume forecast updates. We now have that if, for all $i \in \{1, \dots, m\}$, $n_i(\boldsymbol{y})$ is defined as in (2.11), then $\mathbf{1}^{\mathrm{T}}\boldsymbol{n}(\boldsymbol{y}) = D_T$, where $\boldsymbol{n}(\boldsymbol{y}) := [n_1(\boldsymbol{y}), \dots, n_m(\boldsymbol{y})]^{\mathrm{T}}$. From (2.11) we observe that, given any outcome ω of the sample space Ω, and thus any realization of the forecast updates, ie, $\boldsymbol{\delta}(\omega) := [\delta_1(\omega), \dots, \delta_{m-1}(\omega)]^{\mathrm{T}}$, the realization of the volumes to be traded, $\boldsymbol{n}(\boldsymbol{y};\omega) := [(n_1(\boldsymbol{y}))(\omega), \dots, (n_m(\boldsymbol{y}))(\omega)]^{\mathrm{T}}$, can be expressed as a linear combination of the decision variables $\boldsymbol{y}$. Formally, this mean that, for all $\omega \in \Omega$, we can find a (lower triangular) matrix $\boldsymbol{L}(\omega)$ such that

$$\boldsymbol{n}(\boldsymbol{y};\omega) = \boldsymbol{L}(\omega)\boldsymbol{y}. \tag{2.12}$$

Moreover, if $D_0 \neq 0$, $\boldsymbol{L}(\omega)$ is nonsingular, as all elements on the diagonal are nonzero (ie, $\boldsymbol{L}_{ii}(\omega) = D_0$ for all i). In the rest of the paper $\mathcal{C}(\boldsymbol{y};\omega)$ is shorthand for $(\mathcal{C}(\boldsymbol{y}))(\omega)$, the trading cost of executing strategy $\boldsymbol{y}$ given the outcome ω.

Finally, we consider that if a trader is guaranteed to not pay excessive prices in adverse times, the variance of their trading cost is not of the foremost importance. We thus consider that a risk-averse trader is more interested in minimizing their expected trading cost conditional on a quantile of worst-case scenarios than in minimizing the variance of their trading cost over all scenarios. We therefore quantify the risk of a trader's strategy with the CVaR_α risk measure:

$$\mathrm{CVaR}_\alpha[\mathcal{C}(\boldsymbol{y})] = \frac{1}{\beta_{\boldsymbol{y}}} \int_{\Omega_{\boldsymbol{y}}} \mathcal{C}(\boldsymbol{y};\omega)\,\mathrm{d}\mathbb{P}(\omega), \tag{2.13}$$

where $\beta_{\boldsymbol{y}} = \mathbb{P}[\mathcal{C}(\boldsymbol{y}) \geqslant \mathrm{Var}_\alpha[\mathcal{C}(\boldsymbol{y})]]$, $\Omega_{\boldsymbol{y}} = \{\omega \in \Omega : \mathcal{C}(\boldsymbol{y};\omega) \geqslant \mathrm{Var}_\alpha[\mathcal{C}(\boldsymbol{y})]\}$ and

$$\mathrm{Var}_\alpha[\mathcal{C}(\boldsymbol{y})] = \min\{\gamma \in \mathbb{R} \mid F_{\mathcal{C}(\boldsymbol{y})}(\gamma) \geqslant 1-\alpha\}, \tag{2.14}$$

where

$$F_{\mathcal{C}(\boldsymbol{y})}(\gamma) = \mathbb{P}[\mathcal{C}(\boldsymbol{y}) \leqslant \gamma] \tag{2.15}$$

is the cumulative distribution function (CDF) of $\mathcal{C}(\boldsymbol{y})$. The CVaR_α is a coherent risk measure (Artzner *et al* 1999; Rockafellar and Uryasev 2002) that focuses on the proportion α of extreme costs and that can be interpreted as the expectation of the costs conditional on them exceeding the threshold $\mathrm{Var}_\alpha[\mathcal{C}(\boldsymbol{y})]$. Given a risk-aversion parameter $\lambda_{\mathrm{CVaR}} \in [0, 1]$, a trader thus tries to minimize the mean–CVaR_α trade-off of the total trading cost:

$$\left.\begin{aligned} &\underset{\boldsymbol{y}}{\text{minimize}} && \varphi_\alpha^{\lambda_{\mathrm{CVaR}}}(\boldsymbol{y}) := (1-\lambda_{\mathrm{CVaR}})\mathbb{E}[\mathcal{C}(\boldsymbol{y})] + \lambda_{\mathrm{CVaR}}\mathrm{CVaR}_\alpha[\mathcal{C}(\boldsymbol{y})] \\ &\text{subject to} && \mathbf{1}^{\mathrm{T}}\boldsymbol{y} = 1. \end{aligned}\right\} \tag{2.16}$$

Note that in optimization problem (2.16) a trader's recourse is estimated upfront rather than being simulated via dynamic programming, to avoid the curse of dimensionality that plagues stochastic optimization. Nevertheless, in contrast to the model by Almgren and Chriss (2001), model (2.16) is well defined since it adapts to the total volume variability and is thus guaranteed to satisfy the constraint that the total traded volume equals D_T.

2.2.1 A necessary condition for market viability

As previously mentioned, Huberman and Stanzl (2004) provide the necessary condition (2.4) for a market to be viable if the trading periods are homogeneous (ie, a market with $g_i = g$, $h_i = h$ and $\tau_i = \tau$ for every trading period τ_i, $i \in \{1, \dots, m\}$). In the case of linear impact functions (ie, (2.5)), condition (2.4) then becomes

$$\left(\epsilon \operatorname{sgn}\left(\frac{q}{\tau}\right) + \eta\left(\frac{q}{\tau}\right)\right) - \left(\epsilon \operatorname{sgn}\left(\frac{-q}{\tau}\right) + \eta\left(\frac{-q}{\tau}\right)\right) \gtreqless \gamma q \quad \text{for } q \gtreqless 0. \tag{2.17}$$

This condition intuitively means that the temporary cost incurred by buying and reselling any same amount q of shares on different trading periods should always be greater than the cost-saving opportunity engendered by the price shift caused by the permanent impact. For any vector $\boldsymbol{\epsilon} := [\epsilon_1, \dots, \epsilon_m]^{\mathrm{T}}$ such that $\epsilon_k \geqslant 0$ for all k, condition (2.17) is equivalent to $\eta > \frac{1}{2}\gamma\tau$, which is the sufficient condition for having a unique optimal trading strategy in the model by Almgren and Chriss (2001).

With the rationale that the market should not allow any price manipulation, which we define as a round-trip strategy with negative expected cost, condition (2.17) can be generalized in the more general framework where the impact functions g_i and h_i depend on the trading period τ_i.

DEFINITION 2.1 (Round-trip strategy) A round-trip strategy $\boldsymbol{n}$ is a strategy where the sum of all the trades executed during the execution period equates to zero, ie, $\mathbf{1}^{\mathrm{T}}\boldsymbol{n} = 0$.

Note that a round-trip strategy is defined in terms of the volumes to be traded during each trading period rather than the proportions $\boldsymbol{y}$. Indeed, in the case where the total volume to be traded D_T is certain and equals zero, the only round-trip strategy based on the proportions of D_T would be equivalent to a no trade; hence, defining a round-trip strategy in terms of $\boldsymbol{n}$ offers more flexibility. Based on Definition 2.1 and on the price manipulation definition from Huberman and Stanzl (2004, Definition 1), we define a price manipulation as follows.

DEFINITION 2.2 (Price manipulation) A (risk-neutral) price manipulation is a round-trip strategy $\boldsymbol{n}$ with a strictly negative expected cost, ie,

$$\mathbb{E}[\mathcal{C}(\boldsymbol{n})] < 0. \tag{2.18}$$

LEMMA 2.3 *If a market is free of price manipulation, then the matrix*

$$\boldsymbol{M} := \boldsymbol{E} - \boldsymbol{\Gamma}$$

is positive semi-definite, ie, $\boldsymbol{M} \succeq 0$*, where*

$$\boldsymbol{E} = \begin{bmatrix} \dfrac{2\eta_1}{\tau_1} + \dfrac{2\eta_m}{\tau_m} & \dfrac{2\eta_m}{\tau_m} & \cdots & \dfrac{2\eta_m}{\tau_m} \\ \dfrac{2\eta_m}{\tau_m} & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & \dfrac{2\eta_m}{\tau_m} \\ \dfrac{2\eta_m}{\tau_m} & \cdots & \dfrac{2\eta_m}{\tau_m} & \dfrac{2\eta_{m-1}}{\tau_{m-1}} + \dfrac{2\eta_m}{\tau_m} \end{bmatrix}$$

and

$$\boldsymbol{\Gamma} = \begin{bmatrix} 2\gamma_1 & \gamma_2 & & \gamma_i & & \gamma_{m-1} \\ \gamma_2 & 2\gamma_2 & & \vdots & & \vdots \\ & & \ddots & \vdots & & \vdots \\ \gamma_i & \dots & \dots & 2\gamma_i & & \vdots \\ & & & & \ddots & \vdots \\ \gamma_{m-1} & \dots & \dots & \dots & \dots & 2\gamma_{m-1} \end{bmatrix}.$$

PROOF To prove Lemma 2.3, we show that if matrix $\boldsymbol{M}$ is not positive semi-definite, then we can find a price manipulation. Given an outcome $\omega \in \Omega$, the trading cost of a round-trip strategy $\boldsymbol{n}$ is given by

$$\mathcal{C}(\boldsymbol{n};\omega) = \sum_{i=1}^{m} \left[(\xi_i(\omega) + \gamma_i n_i)\left(-\sum_{k=1}^{i} n_k \right)\right] + \sum_{i=1}^{m} \left(\epsilon_i |n_i| + \frac{\eta_i n_i^2}{\tau_i} \right). \quad (2.19)$$

As $\boldsymbol{n}$ defines a round-trip strategy, the trading cost can be expressed in terms of the first $m-1$ components of $\boldsymbol{n}$ by using the fact that $\sum_{i=1}^{m} n_i = 0 \implies n_m = -\sum_{i=1}^{m-1} n_i$:

$$\begin{aligned} \mathcal{C}(\boldsymbol{n};\omega) = {} & \sum_{i=1}^{m-1} \left[(\xi_i(\omega) + \gamma_i n_i)\left(-\sum_{k=1}^{i} n_k \right)\right] \\ & + \sum_{i=1}^{m-1} \left(\epsilon_i |n_i| + \frac{\eta_i n_i^2}{\tau_i} \right) + \left(\epsilon_m \left| \sum_{i=1}^{m-1} n_i \right| + \frac{\eta_m (\sum_{i=1}^{m-1} n_i)^2}{\tau_m} \right). \end{aligned} \quad (2.20)$$

This trading cost can be split into two parts: $\mathrm{Lin}(\boldsymbol{n};\omega)$, which grows linearly with $\boldsymbol{n}$ and depends on the uncertainty outcome ω; and a quadratic term in $\boldsymbol{n}$ that is independent of ω,

$$\begin{aligned} \mathcal{C}(\boldsymbol{n};\omega) &= \mathrm{Lin}(\boldsymbol{n};\omega) + \sum_{i=1}^{m-1} \left(\frac{\eta_i n_i^2}{\tau_i} \right) + \frac{\eta_m (\sum_{i=1}^{m-1} n_i)^2}{\tau_m} - \sum_{i=1}^{m-1} \left(\gamma_i n_i \left(\sum_{k=1}^{i} n_k \right)\right) \\ &= \mathrm{Lin}(\boldsymbol{n};\omega) + \tfrac{1}{2} \boldsymbol{n}_{[1:m-1]}^{\mathrm{T}} \boldsymbol{M} \boldsymbol{n}_{[1:m-1]}, \end{aligned} \quad (2.21)$$

where $\boldsymbol{n}_{[1:m-1]}$ is the vector of the first $m-1$ components of $\boldsymbol{n}$. If matrix $\boldsymbol{M}$ is not positive semi-definite, there exists a direction $\boldsymbol{d} \in \mathbb{R}^{m-1}$ along which the quadratic form is concave for every $\omega \in \Omega$, as $\boldsymbol{M}$ is independent of ω; any normalized eigen-vector associated with one of the negative eigenvalues of $\boldsymbol{M}$ is such a direction. In the following let ζ^- be a negative eigenvalue of $\boldsymbol{M}$ and let $\boldsymbol{v}$ be an associated normalized eigenvector. As a consequence, there exists a $\kappa \in \mathbb{R}$, such that the round-trip strategy

$$\tilde{\boldsymbol{n}} := [\kappa v_1, \dots, \kappa v_{m-1}, -\kappa(\mathbf{1}^{\mathrm{T}}\boldsymbol{v})]^{\mathrm{T}}$$

is a price manipulation. Indeed, we can find a $\kappa \in \mathbb{R}$ such that $\mathcal{C}(\tilde{\boldsymbol{n}};\omega) < 0$ for all $\omega \in \Omega$, which implies that $\tilde{\boldsymbol{n}}$ is a price manipulation (ie, $\mathbb{E}[\mathcal{C}(\tilde{\boldsymbol{n}})] < 0$). The existence of such a κ is justified by the fact that when $\kappa \to \infty$, and thus when $\|\tilde{\boldsymbol{n}}\|_2 \to \infty$ the trading cost is dominated by its quadratic part (ie, $\mathcal{C}(\tilde{\boldsymbol{n}};\omega) \simeq \frac{1}{2}\tilde{\boldsymbol{n}}_{[1:m-1]}^{\mathrm{T}}\boldsymbol{M}\tilde{\boldsymbol{n}}_{[1:m-1]} = \frac{1}{2}\kappa^2\zeta^- < 0$). □

2.2.2 Set of optimal strategies

Analogously to model (2.3), we can prove that the objective function $\varphi_\alpha^{\lambda_{\mathrm{CVaR}}}$ of optimization problem (2.16) is convex on its feasible set $\mathcal{Y} := \{\boldsymbol{y} \in \mathbb{R}^m \mid \mathbf{1}^{\mathrm{T}}\boldsymbol{y} = 1\}$ if the market excludes any price manipulation. The convexity of the objective function on $\mathcal{Y}$ for any value of $\lambda_{\mathrm{CVaR}} \in [0, 1]$ is equivalent to the simultaneous convexity of both functions $f_1 \colon \mathbb{R}^m \to \mathbb{R}; \boldsymbol{y} \mapsto \mathbb{E}[\mathcal{C}(\boldsymbol{y})]$ and $f_2 \colon \mathbb{R}^m \to \mathbb{R}; \boldsymbol{y} \mapsto \mathrm{CVaR}_\alpha[\mathcal{C}(\boldsymbol{y})]$ on $\mathcal{Y}$.

LEMMA 2.4 *If the permanent and temporary impact functions are given by (2.5) and if $\boldsymbol{M} \succeq 0$ (respectively, $\boldsymbol{M} \succ 0$), then function $f_1 \colon \mathbb{R}^m \to \mathbb{R}, \boldsymbol{y} \mapsto \mathbb{E}[\mathcal{C}(\boldsymbol{y})]$ is (respectively, strictly) convex on $\mathcal{Y}$.*

PROOF By definition we have $\mathbb{E}[\mathcal{C}(\boldsymbol{y})] = \int_\Omega \mathcal{C}(\boldsymbol{y};\omega)\,\mathrm{d}\mathbb{P}(\omega)$. Hence, if for any outcome $\omega \in \Omega$ function $\mathcal{C}(\cdot\,;\omega) \colon \mathbb{R}^m \to \mathbb{R}; \boldsymbol{y} \mapsto \mathcal{C}(\boldsymbol{y};\omega)$ is (strictly) convex in $\boldsymbol{y}$, then Lemma 2.4 holds since the expectation of $\mathcal{C}(\boldsymbol{y})$ can be seen as a nonnegative weighted sum of the $\mathcal{C}(\boldsymbol{y};\omega)$. Given (2.10), $\mathcal{C}(\boldsymbol{y};\omega)$ is written as follows:

$$\begin{aligned}\mathcal{C}(\boldsymbol{y};\omega) = \sum_{i=1}^{m}\Bigg[\bigg(\xi_i(\omega) + \tau_i g_i\bigg(\frac{n_i(\boldsymbol{y};\omega)}{\tau_i}\bigg)\bigg)\bigg(D_0 + \sum_{k=1}^{m}\delta_k(\omega) - \sum_{k=1}^{i} n_k(\boldsymbol{y};\omega)\bigg)\Bigg] \\ + \sum_{i=1}^{m} n_i(\boldsymbol{y};\omega)h_i\bigg(\frac{n_i(\boldsymbol{y};\omega)}{\tau_i}\bigg). \qquad (2.22)\end{aligned}$$

Since for any $\omega \in \Omega$ there exists a linear transformation between $\boldsymbol{y}$ and $\boldsymbol{n}(\boldsymbol{y};\omega)$ as shown in (2.12), proving the convexity of $\mathcal{C}(\cdot\,;\omega)$ in $\boldsymbol{y}$ is equivalent to proving the convexity of $\tilde{\mathcal{C}}(\cdot\,;\omega)$ in $\boldsymbol{n}(\boldsymbol{y};\omega)$ (Boyd and Vandenberghe 2004), where, using (2.5),

$\tilde{\mathcal{C}}(\cdot\,;\omega)$ is given by

$$\begin{aligned}\tilde{\mathcal{C}}(\boldsymbol{n}(\boldsymbol{y};\omega);\omega) &= \sum_{i=1}^{m}\Bigg[(\xi_i(\omega)+\gamma_i n_i(\boldsymbol{y};\omega))\bigg(D_0+\sum_{k=1}^{m}\delta_k(\omega)-\sum_{k=1}^{i}n_k(\boldsymbol{y};\omega)\bigg)\Bigg]\\ &\quad+\sum_{i=1}^{m}\bigg[\epsilon_i|n_i(\boldsymbol{y};\omega)|\frac{\eta_i n_i^2(\boldsymbol{y};\omega)}{\tau_i}\bigg]\\ &= \sum_{i=1}^{m}\xi_i(\omega)\bigg(D_T(\omega)-\sum_{k=1}^{i}n_k(\boldsymbol{y};\omega)\bigg)+\sum_{i=1}^{m}\gamma_i n_i(\boldsymbol{y};\omega)D_T(\omega)\\ &\quad+\sum_{i=1}^{m}\epsilon_i|n_i(\boldsymbol{y};\omega)|-\sum_{i=1}^{m}\gamma_i n_i(\boldsymbol{y};\omega)\bigg(\sum_{k=1}^{i}n_k(\boldsymbol{y};\omega)\bigg)\\ &\quad+\sum_{i=1}^{m}\frac{\eta_i n_i^2(\boldsymbol{y};\omega)}{\tau_i},\end{aligned}\tag{2.23}$$

where $D_T(\omega) := D_0 + \sum_{k=1}^{m}\delta_k(\omega)$. It is straightforward that the first three terms of (2.23) are convex in $\boldsymbol{n}(\boldsymbol{y};\omega)$; the last two can be combined in a similar manner as in the proof of Lemma 2.3, ie,

$$\begin{aligned}&-\sum_{i=1}^{m}\gamma_i n_i(\boldsymbol{y};\omega)\bigg(\sum_{k=1}^{i}n_k(\boldsymbol{y};\omega)\bigg)+\sum_{i=1}^{m}\frac{\eta_i n_i^2(\boldsymbol{y};\omega)}{\tau_i}\\ &\quad=\tfrac{1}{2}\boldsymbol{n}_{[1:m-1]}^{\mathrm{T}}\boldsymbol{M}\boldsymbol{n}_{[1:m-1]}-\gamma_m\bigg(D_T(\omega)-\sum_{i=1}^{m-1}n_i(\boldsymbol{y};\omega)\bigg)D_T(\omega)\\ &\qquad+\frac{\eta_m}{\tau_m}D_T(\omega)\bigg(D_T(\omega)-2\sum_{i=1}^{m-1}n_i(\boldsymbol{y};\omega)\bigg),\end{aligned}\tag{2.24}$$

where $\boldsymbol{n}_{[1:m-1]} := [n_1(\boldsymbol{y};\omega),\dots,n_{m-1}(\boldsymbol{y};\omega)]^{\mathrm{T}}$. Hence, the quadratic part of the trading costs is (respectively, strictly) convex if $\boldsymbol{M}\succeq 0$ (respectively, $\boldsymbol{M}\succ 0$); this completes the proof. □

LEMMA 2.5 *If the permanent and temporary impact functions are given by (2.5) and if $\boldsymbol{M}\succeq 0$ (respectively, $\boldsymbol{M}\succ 0$), then function $f_2\colon\mathbb{R}^m\to\mathbb{R};\boldsymbol{y}\mapsto$ $\mathit{CVaR}_\alpha[\mathcal{C}(\boldsymbol{y})]$ is (respectively, strictly) convex on $\mathcal{Y}$.*

PROOF Any real-valued coherent risk measure $\varrho\colon\mathbb{R}^m\to\mathbb{R};\boldsymbol{y}\mapsto\varrho[\mathcal{C}(\boldsymbol{y})]$ as defined in Artzner *et al* (1999) can be expressed in its dual representation, ie,

$$\varrho[\mathcal{C}(\boldsymbol{y})] = \sup_{\mu\in\mathcal{D}}\mathbb{E}^{\mu}[\mathcal{C}(\boldsymbol{y})],\tag{2.25}$$

where $\mathcal{D}$ is a convex subset of probability measures on Ω (Shapiro *et al* 2009, Theorems 6.4 and 6.6). As a consequence, as CVaR_α is a coherent risk measure (Rockafellar and Uryasev 2002, Corollary 12), the desired result is shown by combining the (respectively, strict) convexity of $\mathcal{C}(\,\cdot\,;\omega)$ for any ω with the fact that a function that is defined as the pointwise supremum of (respectively, strictly) convex functions is (respectively, strictly) convex (Boyd and Vandenberghe 2004). □

THEOREM 2.6 *If the permanent and temporary impact functions are given by (2.5), if* $\lambda_{\mathrm{CVaR}} \in [0,1]$ *and if* $\boldsymbol{M} \succeq 0$ *(respectively,* $\boldsymbol{M} \succ 0$*), then the objective function of optimization problem (2.16) is (respectively, strictly) convex on its feasible set* $\mathcal{Y}$.

PROOF The proof is straightforward, based on Lemmas 2.4 and 2.5. □

Similarly to model (2.3), the convexity and the nonemptiness of the feasible set $\mathcal{Y}$ and Theorem 2.6 imply, under the conditions that $\lambda_{\mathrm{CVaR}} \in [0,1]$ and that $\boldsymbol{M} \succeq 0$, the existence and uniqueness of a convex set $\mathcal{Y}^\star \subseteq \mathcal{Y}$ of optimal execution strategies for the model defined in (2.16). If $\boldsymbol{M} \succ 0$, the set of optimal strategies is a singleton (ie, $\mathcal{Y}^\star := \{\boldsymbol{y}^\star\}$).

In other words, the necessary condition $\boldsymbol{M} \succeq 0$ for the absence of price manipulation, and thus for having a viable market where no trader has interest in impacting the price dynamics with artificial trades, is tantamount to the convexity of optimization problem (2.16), and thus to a convex set of optimal strategies for a trader, considering both price and volume uncertainty (PVU).

2.3 Comparison with other approaches

2.3.1 Mean–variance with recourse

In the case where the volume target is uncertain, the model by Almgren and Chriss (2001) can be deployed in conjunction with a systematic recourse whenever a forecast update on the volume target becomes available. We will show that such a strategy can be reproduced in our model with an appropriate choice of $\boldsymbol{y}$ and $\boldsymbol{\beta}$. Similarly to model (2.16), a trading strategy $\boldsymbol{n}$ in model (2.3) can be equivalently expressed as the proportions y_i of the fixed volume target d_T to be traded during each trading period τ_i, $i \in \{1,\dots,m\}$.

LEMMA 2.7 *If* ξ_i, $i \in \{1,\dots,m\}$*, are draws from independent random variables each with zero-mean (ie,* $\mathbb{E}[\xi_i] = 0$*), if the permanent and temporary impact functions are given by (2.5), and if the temporary impact parameter* ϵ_i *is constant over (ie, independent of) all trading periods (ie, for all* $i \in \{1,\dots,m\}$ *such that* $\epsilon_i = \epsilon$*), then the optimal strategy* $\boldsymbol{y}^\star$ *of optimization problem (2.3) expressed in terms of the proportions of* d_T *is independent of* d_T.

PROOF Based on (2.2) the expectation and variance of the trading cost are given by

$$\begin{aligned}\mathbb{E}[\mathcal{C}(\boldsymbol{n})] &= \sum_{i=1}^{m} \gamma_i n_i \left(d_T - \sum_{k=1}^{i} n_k \right) + \sum_{i=1}^{m} \left(\epsilon |n_i| + \frac{\eta_i n_i^2}{\tau_i} \right) \\ &= \sum_{i=1}^{m} \gamma_i y_i \left(1 - \sum_{k=1}^{i} y_k \right) d_T^2 + \sum_{i=1}^{m} \left(\epsilon |y_i d_T| + \frac{\eta_i}{\tau_i} (y_i d_T)^2 \right), \\ \mathrm{Var}[\mathcal{C}(\boldsymbol{n})] &= \sum_{i=1}^{m} \mathrm{Var}[\xi_i] \left(d_T - \sum_{k=1}^{i} n_k \right)^2 = \sum_{i=1}^{m} \mathrm{Var}[\xi_i] \left(1 - \sum_{k=1}^{i} y_k \right)^2 d_T^2.\end{aligned}$$

Moreover, if we assume without loss of generality that $d_T > 0$, then the optimal volumes $\boldsymbol{n}^\star$ of optimization problem (2.3) are positive (Almgren and Chriss 2001). Therefore, adding the nonnegativeness constraint on the decision variables $\boldsymbol{n}$ to optimization problem (2.3) does not alter its optimal solution. Hence, we can replace $\sum_{i=1}^{m} \epsilon |n_i|$ with $\sum_{i=1}^{m} \epsilon n_i = \epsilon d_T$. As a consequence, all the terms in the objective function of optimization problem (2.3) that depend on the decision variables $\boldsymbol{y}$ are multiplied by d_T^2, which means that d_T has no impact on the optimal solution $\boldsymbol{y}^\star$. □

Lemma 2.7 implies, that under PVU, the strategy of a trader that consists in deploying the model by Almgren and Chriss (2001) in conjunction with a systematic recourse whenever an update on the volume target becomes available is time consistent. More precisely, given the optimal strategy $\boldsymbol{y}^{\star,t_0}$ of (2.3) computed at time t_0, if we recompute the optimal strategy $\boldsymbol{y}^{\star,t_{i-1}}$ of (2.3) at a subsequent decision time t_{i-1}, $i \in \{2, \dots, m\}$, it would simply be the continuation of $\boldsymbol{y}^{\star,t_0}$ from time t_{i-1} to time t_m, ie,

$$\boldsymbol{y}^{\star,t_{i-1}} := [y_i^{\star,t_{i-1}}, \dots, y_m^{\star,t_{i-1}}]^{\mathrm{T}} = \left(\frac{1}{\sum_{r=i}^{m} y_r^{\star,t_0}} \right) \boldsymbol{y}_{[i:m]}^{\star,t_0}, \tag{2.26}$$

where $\boldsymbol{y}_{[i:m]}^{\star,t_0} := [y_i^{\star,t_0}, \dots, y_m^{\star,t_0}]^{\mathrm{T}}$. Note that the rescaled proportions $\boldsymbol{y}^{\star,t_{i-1}}$, $i \in \{2, \dots, m\}$, are independent of past realizations (price moves and forecast updates) and can consequently be computed beforehand at the start of the execution period, ie, at time t_0. If $\boldsymbol{n}^{\mathrm{rec}} := [n_1^{\mathrm{rec}}, \dots, n_m^{\mathrm{rec}}]^{\mathrm{T}}$ denotes the traded volumes of a trader who updates their strategy by solving (2.3) whenever an update on the volume target occurs, then the volume n_i^{rec} traded on trading period τ_i, $i \in \{1, \dots, m\}$, is equal to the proportion $y_i^{\star,t_{i-1}}$ of the best volume target estimate D_{i-1} minus the volumes already traded, ie, $n_i^{\mathrm{rec}} = y_i^{\star,t_{i-1}} (D_{i-1} - \sum_{k=1}^{i-1} n_k^{\mathrm{rec}})$.

LEMMA 2.8 *If ξ_i, $i \in \{1, \dots, m\}$, are draws from independent random variables, each with zero-mean, ie, $\mathbb{E}[\xi_i] = 0$, if the permanent and temporary impact functions* *are given by (2.5), and if the temporary impact parameter ϵ_i is constant over all trading periods, ie, for all $i \in \{1, \dots, m\}$ such that $\epsilon_i = \epsilon$, then there exist a strategy $\bar{\boldsymbol{y}}$ and a redistribution matrix $\bar{\boldsymbol{\beta}} := (\beta_{k,i}) \in \mathbb{R}^{m-1 \times m}$ such that the trading volumes given by (2.11) are equal to $\boldsymbol{n}^{\text{rec}}$.*

PROOF Let us define $\bar{\boldsymbol{y}} := \boldsymbol{y}^{\star,t_0} > 0$, the optimal solution of optimization problem (2.3) at time t_0, and define $\bar{\boldsymbol{\beta}} := (\bar{\beta}_{k,i}) \in \mathbb{R}^{m-1 \times m}$ as the redistribution matrix implied by $\boldsymbol{y}^{\star,t_0}$, ie, the matrix whose components are defined as follows:

$$\bar{\beta}_{k,i}(\boldsymbol{y}^{\star,t_0}) := \begin{cases} 0 & \text{if } i \leqslant k, \\ \dfrac{y_i^{\star,t_0}}{\sum_{r=k+1}^{m} y_r^{\star,t_0}} & \text{otherwise.} \end{cases} \tag{2.27}$$

Then, $\bar{\boldsymbol{y}}$ and $\bar{\boldsymbol{\beta}}$ are such that the trading volumes given by (2.11) are equal to $\boldsymbol{n}^{\text{rec}}$. Indeed, using the fact that $\sum_{i=1}^{m} y_i^{\star,t_0} = 1$, we have that the volumes $\boldsymbol{n}^{\text{rec}}$ are given by

$$\begin{aligned} n_1^{\text{rec}} &= y_1^{\star,t_0} D_0, \\ n_2^{\text{rec}} &= \frac{y_2^{\star,t_0}}{\sum_{r=2}^{m} y_r^{\star,t_0}} (D_0 + \delta_1 - y_1^{\star,t_0} D_0) \\ &= \frac{y_2^{\star,t_0}}{1 - y_1^{\star,t_0}} ((1 - y_1^{\star,t_0}) D_0 + \delta_1) = y_2^{\star,t_0} D_0 + \frac{y_2^{\star,t_0}}{1 - y_1^{\star,t_0}} \delta_1, \\ n_3^{\text{rec}} &= \frac{y_3^{\star,t_0}}{\sum_{r=3}^{m} y_r^{\star,t_0}} \left(D_0 + \delta_1 + \delta_2 - y_1^{\star,t_0} D_0 - y_2^{\star,t_0} D_0 - \frac{y_2^{\star,t_0}}{1 - y_1^{\star,t_0}} \delta_1 \right) \\ &= y_3^{\star,t_0} D_0 + \frac{y_3^{\star,t_0}}{1 - y_1^{\star,t_0}} \delta_1 + \frac{y_3^{\star,t_0}}{1 - y_1^{\star,t_0} - y_2^{\star,t_0}} \delta_2, \\ &\;\vdots \\ n_i^{\text{rec}} &= y_i^{\star,t_0} D_0 + \sum_{k=1}^{i-1} \left(\frac{y_i^{\star,t_0}}{1 - \sum_{r=1}^{k} y_r^{\star,t_0}} \right) \delta_k. \end{aligned}$$

Using the definitions of $\bar{\boldsymbol{y}}$ and $\bar{\boldsymbol{\beta}}$, we obtain

$$\begin{aligned} n_i^{\text{rec}} &= \bar{y}_i D_0 + \sum_{k=1}^{i-1} \left(\frac{(1 - \sum_{r=1}^{k} \bar{y}_r) \bar{y}_i + \bar{y}_i (\sum_{r=1}^{k} \bar{y}_r)}{1 - \sum_{r=1}^{k} \bar{y}_r} \right) \delta_k \\ &= \bar{y}_i D_0 + \sum_{k=1}^{i-1} \left(\bar{y}_i + \bar{\beta}_{k,i} \sum_{r=1}^{k} \bar{y}_r \right) \delta_k, \end{aligned} \tag{2.28}$$

which is exactly the expression of the traded volumes in (2.16) for the strategy $\bar{\boldsymbol{y}}$ and the redistribution matrix $\bar{\boldsymbol{\beta}}$ (see (2.11)). □

Lemma 2.8 implies that if we were to optimize (2.16) on the set of strategies $\mathcal{Y}$ and the set of redistribution matrixes $\mathcal{B} := \{\boldsymbol{\beta} \in \mathbb{R}^{m-1\times m} \mid \beta_{k,i} = 0$ if $i \leqslant k$ and $\forall k \in \{1,\dots,m-1\}$: $\sum_{i=k+1}^{m} \beta_{k,i} = 1\}$, then there exists a strategy $(\boldsymbol{y}, \boldsymbol{\beta}) \in \mathcal{Y} \times \mathcal{B}$ such that the volumes traded by following this strategy are the same as those of a trader recomputing the optimal strategy in the mean–variance framework whenever an update on the volume target is released. Hence, a wisely chosen redistribution matrix $\boldsymbol{\beta}$ can mimic the action of taking recourse. A sensible choice for the redistribution matrix $\boldsymbol{\beta}$ would be, for instance, the matrix implied by the optimal risk-neutral strategy $\boldsymbol{y}^{\star,t_0,\lambda_{\mathrm{Var}}=0}$ in the mean–variance framework (optimization problem (2.3)) at time t_0 (ie, the matrix defined as in (2.27) but with $\boldsymbol{y}^{\star,t_0,\lambda_{\mathrm{Var}}=0}$).

2.3.2 Exponential utility and mean–quadratic variation

The uncertainty of the volume target has, to the best of our knowledge, not yet been considered in the literature. Some papers deal with a conceptually similar problem to ours but rely on assumptions that are discarded in our model, allowing for greater generality as explained hereinafter. Cheng *et al* (2017, 2019) incorporate in the optimal execution problem the uncertainty related to order fills, which is the risk for an order to be under- or overexecuted. Bulthuis *et al* (2017) extend this model by including, additionally to the market orders, limit orders with uncertain fill rates. They also add penalties to tailor the orders' sign and magnitude. For comparison, we consider the model by Cheng *et al* (2017), as we make similar assumptions on the market; the model is formulated here as a liquidation problem (ie, from a seller's perspective rather than the buyer's perspective adopted earlier).

Their model is as follows. Let x_t denote the number of shares the trader holds at a time $t \in [0, T]$ during the execution process, and let x_0 denote the initial amount of shares. Similarly to Almgren and Chriss (2001), the fair price S_t is driven by the stochastic differential equation

$$\mathrm{d}S_t = \gamma\,\mathrm{d}x_t + \mu\,\mathrm{d}t + \sigma\,\mathrm{d}W_t \iff S_t = S_0 - \gamma x_0 + \gamma x_t + \mu t + \sigma W_t, \quad (2.29)$$

and the transacted price $\tilde{S}_t$ is given by

$$\tilde{S}_t = S_t - \eta v_t, \quad (2.30)$$

where W_t is a Brownian motion, v_t denotes the rate at which the trader is scheduled to trade at the instant t, μ parameterizes the price drift, and γx_t, $\gamma \geqslant 0$, and $-\eta v_t$, $\eta > 0$, refer to the permanent and temporary impacts of the trades, respectively. To take into account the uncertainty regarding the order fills, Cheng *et al* (2017) add to the dynamics of the position x_t a noise driven by another Brownian motion Z_t (potentially correlated with W_t), ie,

$$\mathrm{d}x_t = -v_t\,\mathrm{d}t + m(v_t)\,\mathrm{d}Z_t, \tag{2.31}$$

where the diffusion $m(v_t)$ characterizes the magnitude of the uncertainty of order fills. At each time $t \in [0, T]$, the profit and loss (P&L) $\Pi_t(x)$ generated by the strategy x until time t is

$$\Pi_t(x) = x_t(S_t - S_0) - \int_0^t (\tilde{S}_u - S_0)\,\mathrm{d}x_u. \tag{2.32}$$

The first term on the right-hand side of (2.32) represents the change in fair value of the positions that still have to be transacted (it is beneficial for the trader if $S_t > S_0$), whereas the second term is the cost incurred due to the price risk and the price impact of the shares traded up to time t. Using (2.29) and (2.30), the P&L can be reformulated as

$$\Pi_t(x) = x_t(\gamma(x_t - x_0) + \mu t + \sigma W_t) - \int_0^t (\gamma(x_u - x_0) + \mu u + \sigma W_u - \eta v_u)\,\mathrm{d}x_u. \tag{2.33}$$

To discourage any residual positions at the terminal time T, a penalty term $f : x_T \to -\beta x_T^2$, $\beta \geqslant \eta > \gamma/2$, is added to $\Pi_T(x)$, which gives the final P&L $\tilde{\Pi}_T(x) := \Pi_T(x) + f(x_T)$. In the case where the magnitude of the uncertainty of order fills is constant ($m(v_t) = m_0$) and where there is no drift ($\mu = 0$), Cheng *et al* (2017) provide closed-form expressions for the optimal trading rate $v_t^\star$ and the liquidation trajectory $x_t^\star$ of a risk-averse trader that assesses risk in two distinct ways, on the one hand by maximizing the expected utility of the final P&L (ie, $\sup_{v\in\mathcal{A}} \mathbb{E}[U(\tilde{\Pi}_T(x))]$ with exponential utility (EU) function $U(x) = 1 - \exp(-\theta_{\mathrm{EU}} x)/\theta_{\mathrm{EU}}$), and on the other hand by maximizing the expected P&L penalized by its quadratic variation (QV) (Forsyth *et al* 2012), $\sup_{v\in\mathcal{A}} \mathbb{E}[\tilde{\Pi}_T(x) - \lambda_{\mathrm{QV}}\mathrm{QV}[\Pi_T(x)]]$. In both cases, $\mathcal{A}$ denotes the set of admissible controls. The risk-aversion parameters are given by $\theta_{\mathrm{EU}} > 0$ (respectively, $\lambda_{\mathrm{QV}} > 0$); the larger the θ_{EU} (respectively, the λ_{QV}), the more risk-averse the trader.

Based on the optimal trading rate $v_t^\star$ given by Cheng *et al* (2017), it is straightforward to obtain the optimal trading rate of the execution problem formulated from the perspective of a buyer perspective desiring to acquire an uncertain amount of shares D_T. Indeed, the fluctuations in the number of positions of the trader due to the uncertain order fills can be seen as fluctuations in the uncertain volume target. Hence, if the dynamics of the volume target forecast $D_t := \mathbb{E}[D_T \mid \mathcal{F}_t]$ is given by $\mathrm{d}D_t = m_0\,\mathrm{d}X_t$, where X_t is a Brownian motion and $(\mathcal{F}_t)_{t\in[0,T]}$ is the filtration of sigma algebras, representing the information available at time t, if w_t denotes the trading rate at the instant t, and if z_t denotes the number of shares left to be traded in order to match the current forecast D_t on the volume target, then the dynamics of z_t is given by

$$\mathrm{d}z_t = -w_t\,\mathrm{d}t + m_0\,\mathrm{d}X_t, \tag{2.34}$$

where the initial amount of shares to be traded, z_0, is equal to $D_0 := \mathbb{E}[D_T \mid \mathcal{F}_0]$. Note that, in contrast to v_t, a positive value for w_t corresponds to the acquisition and not the liquidation of shares. Consequently, the impact of the trades on the price dynamics and the transacted price coincides with that of Cheng *et al* (2017) up to a change of sign, and similarly for the price drift, the rationale being that a positive price drift in a liquidation program has a similar beneficial effect in terms of costs as a negative price drift in an acquisition program. The fair S_t dynamics, the transacted price $\tilde{S}_t$ and the P&L $\Pi_t(z)$ are then given by

$$\begin{aligned} \mathrm{d}S_t &= -\gamma\,\mathrm{d}z_t - \mu\,\mathrm{d}t - \sigma\,\mathrm{d}W_t \\ &\iff S_t = S_0 + \gamma(z_0 - z_t) - \mu t - \sigma W_t, \end{aligned} \tag{2.35}$$

$$\tilde{S}_t = S_t + \eta w_t, \tag{2.36}$$

$$\Pi_t(z) = z_t(S_0 - S_t) - \int_0^t (S_0 - \tilde{S}_u)\,\mathrm{d}z_u. \tag{2.37}$$

Similarly to (2.32), the first term on the right-hand side of (2.37) represents the change in fair value of the positions that still have to be acquired (it is beneficial for the trader if $S_t < S_0$), whereas the second term is the cost incurred due to the price risk and the price impact of the shares traded up to time t. Using (2.35) and (2.36), the P&L can be reformulated as

$$\begin{aligned} \Pi_t(z) = {} & z_t(\gamma(z_t - z_0) + \mu t + \sigma W_t) \\ & - \int_0^t (\gamma(z_u - z_0) + \mu u + \sigma W_u - \eta w_u)\,\mathrm{d}z_u. \end{aligned} \tag{2.38}$$

We observe that (2.33) and (2.38) are similar. Therefore, if a risk-averse trader maximizes either the expected utility of the final P&L (ie, $\tilde{\Pi}_T(z) = \Pi_T(z) + f(z_T)$) or the expected P&L that is penalized by its quadratic variation, then the optimal trading rate $w_t^\star$ in an acquisition program subjected to an uncertain volume target will be the same as the optimal trading rate $v_t^\star$ in a liquidation program subjected to uncertain order fills. If, in addition, we assume the absence of price drift and that any remaining positions are prohibited (ie, $\mu = 0$ and $\beta \to \infty$), then the continuous execution problem formulated above essentially tackles the same problem as model (2.16) in the absence of a fixed cost (ie, for all $i \in \{1, \dots, m\}$ such that $\epsilon_i = 0$), when the price shifts ξ_i and the forecast updates δ_i, $i \in \{1, \dots, m\}$, follow a zero-mean normal distribution of variances $\tau_i\sigma^2$ and $\tau_i m_0^2$ (ie, $\xi_i \sim \mathcal{N}(0, \tau_i\sigma^2)$ and $\delta_i \sim \mathcal{N}(0, \tau_i m_0^2)$), respectively.

To conclude, there are many advantages of our approach. First, we showed in Lemma 2.8 that the traded volumes resulting from the strategy of recomputing the optimal strategy in the mean–variance framework whenever a volume target update becomes available can be reproduced in our model with an appropriate updating rule $\boldsymbol{\beta}$. Second, our model is significantly more versatile than that of Cheng *et al* (2017). Indeed, to provide closed-form expressions for the optimal trading rate, Cheng *et al* (2017) rely on the assumptions that both the price dynamics and the uncertain order fills are modeled with Brownian motions and that the price volatility (σ), the volatility of the updates on the volume target (m_0) and the price impact parameters (γ and η) are constant over the entire execution period. All these considerations are not needed in our model, which allows us to consider a wider range of situations, such as the case where the liquidity profile is not constant over the course of the execution period or where the forecast updates on the volume target are correlated. In these situations finding closed-form expressions of the optimal trading rate for the model proposed by Cheng *et al* (2017) would be significantly more challenging. Alternatively, the optimal rate could be obtained by using a dynamic programming approach. However, this would result in a significantly higher computational cost than our approach. Finally, as we will show in Section 3, despite the fact that the CVaR_α risk measure is not time consistent (which may lead to suboptimal strategies (Rudloff *et al* 2014)), the strategies obtained with our model have competitive performance compared with the optimal solution provided by the dynamical programming approach suggested in Cheng *et al* (2017). This suggests that integrating the uncertainty related to the volume target in the estimate of the trading cost of a strategy (see (2.10)) and taking partial recourse via a predetermined rule $\boldsymbol{\beta}$ lead to intuitive and competitive trading strategies while avoiding the computational cost of more convoluted approaches.

3 NUMERICAL RESULTS

In this section we analyze the influence of incorporating the uncertainty related to the volume target into the trade execution problem. We provide some numerical evidence that, when the uncertainty related to the volume target is considered, a trader should adapt their execution program by delaying their trades. To illustrate the results we apply our model in two distinct test cases. In both cases it is assumed that the market has a single stock with an initial price $S_0 = 50$ and a median daily trading volume V of 5 million shares.

CASE 1 (Constant bid–ask spread) We assume the market liquidity remains constant over all the trading periods; we also a bid–ask spread b_i equal to 0.25% of the initial asset price S_0 for each trading period τ_i (ie, for all $i \in \{1, \dots, m\}$ such that $b_i = 0.125$).

CASE 2 (U-shaped intraday bid–ask spread profile) Previous papers (see, for example, Chan *et al* 1995) suggest that the traded volumes and bid–ask spreads of New York Stock Exchange stocks often follow a U-shaped pattern over the trading day, the trading activity being concentrated at the market opening and closure. Consequently, the price-impact slopes (ie, η_i) are larger in the first and last periods than in the middle periods (Huberman and Stanzl 2005). To capture such behavior let us assume the following profile for the bid–ask spread over the five trading periods: $\boldsymbol{b} := [b_1, \dots, b_5]^{\mathrm{T}} = [0.15, 0.125, 0.125, 0.2, 0.25]^{\mathrm{T}}$, corresponding to 0.3%, 0.25%, 0.25%, 0.4% and 0.5% of the initial asset price S_0.

For both cases we assume that the volatility σ of the asset is constant during the entire execution period and is equal to 0.95, and that the price shift ξ_i, $i \in \{1, \dots, m\}$, follows a zero-mean Gaussian distribution of variance $\tau_i \sigma^2$ (ie, $\xi_i \sim \mathcal{N}(0, \tau_i \sigma^2)$). For the temporary impact, the fixed cost factor ϵ_i is assumed to amount to one-half of the bid–ask spread (ie, $\epsilon_i = 0.5 b_i$ for all $i \in \{1, \dots, m\}$). Additionally, Almgren and Chriss (2001) suggest that a trader incurs a price impact equal to one bid–ask spread for every percent of the daily volume (ie, $\eta_i = b_i/(0.01V)$ for all $i \in \{1, \dots, m\}$), which we also adopt in our paper.

As for the permanent price impact, Almgren and Chriss (2001) assume that the price effects become significant when 10% of the daily volume is traded. In this application we suggest that the effect is significant if it corresponds to five times the bid–ask spread of the trading period, whereas Almgren and Chriss (2001) consider a price shift equal to one bid–ask spread to be significant. This increase in scale allows us to illustrate more clearly the contribution of volume uncertainty to the risk adopted by the trader, but it is also not completely unrealistic for illiquid assets. We therefore set $\gamma_i = 5b_i/0.1V$ for all $i \in \{1, \dots, m\}$. For smaller values of γ_i, the impact of volume uncertainty is slightly smaller, but due to the cumulative effect of sequential trades, even small savings in the trading cost make a large difference over time.

Finally, we assume for both cases that the initial volume target D_0 amounts to 0.5 million shares and that for each trading period τ_i except the last one the forecast update follows a zero-mean Gaussian distribution of variance $\tau_i \nu^2$, with ν being equal to 10% of the initial volume target ($\nu := 0.1 D_0$) (ie, for all $i \in \{1, \dots, m-1\}$ such that $\delta_i \sim \mathcal{N}(0, \tau_i \nu^2)$). As suggested in Section 2.3, we assume the residual volumes to be traded due to the forecast updates are redistributed over the future trading periods based on the matrix $\boldsymbol{\beta}$ implied by the risk-neutral optimal trading strategy $\boldsymbol{y}^{\star, t_0, \lambda_{\mathrm{Var}}=0}$ in the mean–variance framework (optimization problem (2.3)) computed at time t_0, ie, $\boldsymbol{\beta} = (\beta_{k,i}(\boldsymbol{y}^{\star,t_0,\lambda_{\text{Var}}=0})) \in \mathbb{R}^{m-1\times m}$ with

$$\beta_{k,i}(\boldsymbol{y}^{\star,t_0,\lambda_{\text{Var}}=0}) = \begin{cases} 0 & \text{if } i \leqslant k, \\ \dfrac{y_i^{\star,t_0,\lambda_{\text{Var}}=0}}{\sum_{r=k+1}^{m} y_r^{\star,t_0,\lambda_{\text{Var}}=0}} & \text{otherwise.} \end{cases} \tag{3.1}$$

In practice the exact distributions of the price moves and the forecast updates are unknown and are estimated from historical data (see, for example, Section 3.4). In the following we denote by $\hat{\cdot}$ any trader's estimation of the ground truth parameter. For example, $\tau_i \hat{\nu}_i{}^2$ is the trader-estimated value of $\tau_i \nu_i^2$, the variance of the forecast update released after trading period τ_i. Table 1 summarizes the market details of both cases.

3.1 Optimal strategies and impact of the liquidity profile

Given the market parameters of Table 1, the condition $\boldsymbol{M} \succ 0$ of Theorem 2.6 is verified, which means that for any risk-aversion parameter $\lambda_{\text{CVaR}} \in [0, 1]$ the optimal trading strategy in the mean–CVaR framework is uniquely defined. Parts (a) and (b) of Figure 1 depict, for both liquidity profiles, the optimal strategies of model (2.16) when only price uncertainty (PU) is considered, whereas parts (c) and (d) depict the optimal trading strategies when both sources of uncertainty are taken into account (ie, PVU). The difference between the optimal strategies in these two situations is represented for both cases in parts (e) and (f), respectively,. This difference illustrates how the optimal strategies under PU should only be adjusted when a trader additionally takes into account the uncertainty related to the volume target.

The impact of considering the uncertainty related to the volume to be traded on the optimal strategies is consistent with intuition: if the volume to be traded is uncertain, waiting to receive a more precise estimate of this volume is advantageous in terms of risk reduction. Indeed, we observe that a trader should mitigate their traded volume of the first trading period and spread it over the next periods. Moreover, the more risk-averse the trader, the bigger the impact on their optimal strategy.

These plots also illustrate that the impact of considering volume uncertainty on the optimal execution strategy is nonmonotonic due to two different subsets of extreme events: those in which volume forecasts are increased in late trading periods and those where volume forecasts are downsized. The first situation requires carrying out enough early trades to avoid having to upsize the traded volume in the latest stages to a level that is strongly affected by liquidity constraints. The second situation makes it necessary to avoid overtrading in the early trading periods and having to trade out of these positions later. The two effects tend to counterbalance each other, with the result of reallocating some of the early trades to the middle section of the trading

TABLE 1 Market parameters.

Acquisition time	$T = 5$ (days)
Number of trading periods	$m = 5$
Trading period length	for all $i \in \{1, \dots, m\}$, $\tau_i = T/m = 1$ (day)
Initial asset price	$S_0 = 50$ (US$/share)
Asset's volatility	$\sigma = 0.95$ ((US$/share)/day$^{1/2}$)
Price shifts distribution	for all $i \in \{1, \dots, m\}$, $\xi_i \sim \mathcal{N}(0, \tau_i \sigma^2)$
Daily volume	$V = 5 \times 10^6$ (share)
Initial volume target forecast	$D_0 = 0.5 \times 10^6$ (share)
Volume forecast updates distribution	for all $i \in \{1, \dots, m-1\}$, $\delta_i \sim \mathcal{N}(0, \tau_i \nu^2)$ with $\nu = 0.1 D_0$
Bid–ask spread	Case 1: for all $i \in \{1, \dots, m\}$, $b_i = 0.125$ (US$/share)
	Case 2: $\boldsymbol{b} = [0.15, 0.125, 0.125, 0.2, 0.25]$ (US$/share)
Fixed cost	for all $i \in \{1, \dots, m\}$, $\epsilon_i = 0.5 b_i$ (US$/share)
Impact at 1% of market	for all $i \in \{1, \dots, m\}$, $\eta_i = b_i / 0.01V$ ((US$/share)/(share/day))
Permanent impact parameter	for all $i \in \{1, \dots, m\}$, $\gamma_i = 5 b_i / 0.1V$ ((US$/share)^2)
CVaR$_\alpha$ parameter	$\alpha = 0.1$

FIGURE 1 Optimal trading strategies under different assumptions for both test cases.

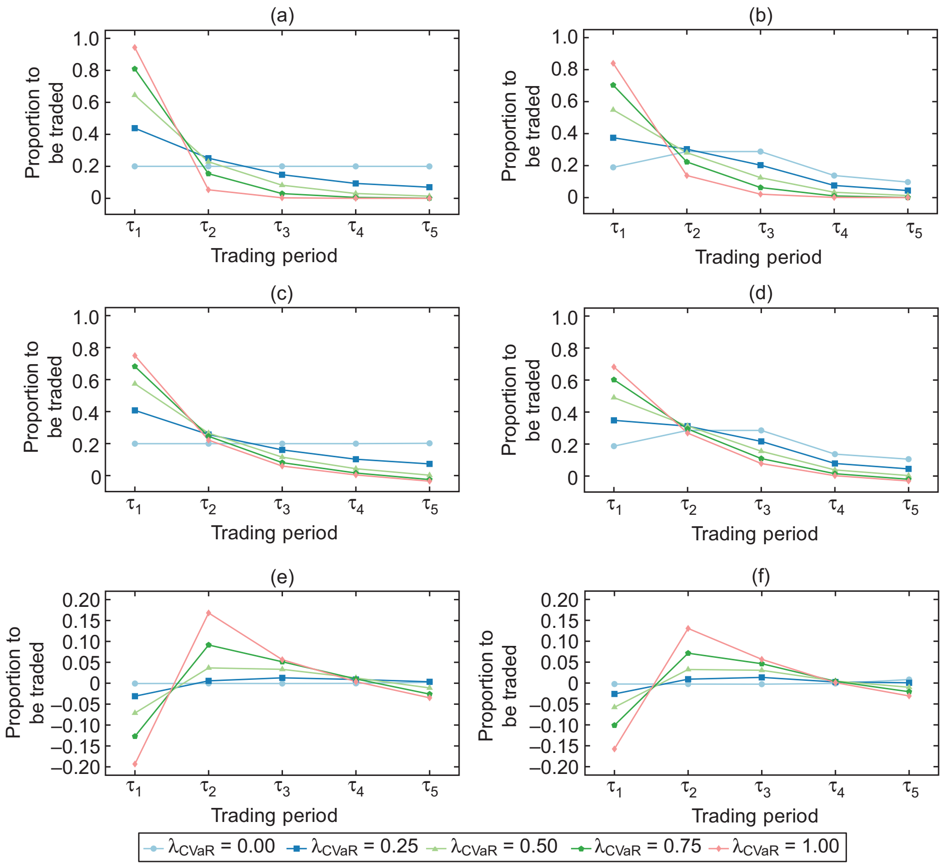

(a) Case 1: optimal trading strategies under PU only (ie, for all $i \in \{1,\dots,m-1\}$ such that $\hat{\nu}_i = 0$). (b) Case 2: optimal trading strategies under PU only (ie, for all $i \in \{1,\dots,m-1\}$ such that $\hat{\nu}_i = 0$). (c) Case 1: optimal trading strategies under PVU (ie, for all $i \in \{1,\dots,m-1\}$ such that $\hat{\nu}_i = \nu$). (d) Case 2: optimal trading strategies under PVU (ie, for all $i \in \{1,\dots,m-1\}$ such that $\hat{\nu}_i = \nu_i$). (e) Case 1: the difference between the strategies represented in part (c) and those in part (a). (f) Case 2: the difference between the strategies represented in part (d) and those in part (b).

period. This effect is picked up due to the incorporation of recourse estimates and the use of the CVaR_α measure that focuses on extreme events.

The liquidity profile has an impact on both the optimal strategies and the adjustments to the strategies when considering the uncertainty related to the volume target in addition to the PU. First, compared with the situation where liquidity remains constant, a risk-neutral trader should shift their trades from less liquid to more liquid periods. This is not as simple for risk-averse traders, since the more risk-averse a trader is, the greater the impact of adversarial scenarios on their objective function, and consequently the more the trader factors in the impact of market liquidity in conjunction with uncertainty. The decreasing influence of market liquidity is observed in parts (b) and (d) of Figure 1, as the shape of the risk-neutral strategy gradually vanishes with increasing risk-aversion parameter in favor of front-loading to cope with PU. Second, the liquidity profile also impacts the adjustments to the strategies when introducing volume uncertainty into the model, as illustrated in parts (e) and (f) of Figure 1 and in Figure 2; the latter depicts the adjustments to several other liquidity profiles as well as the optimal strategies under PVU.

To cope with volume uncertainty, a trader has to consider the two different subsets of the extreme events described above. On the one hand, the less liquid the last trading periods, the more sensitive the trader is to the risk of having to upsize the traded volume of the last trading periods, and therefore the smaller the planned proportions of the volume to be traded over these periods. On the other hand, a risk-averse trader considers the risk of overtrading by mitigating the trading quantity of the first trading period and redistributing it over the following periods. The less liquid the market is during these next trading periods, the greater the risk of overtrading and therefore the more beneficial the decrease in the trading quantity of the first trading period in terms of risk reduction. Figure 2 also shows that the adjustment in the trading proportion of the first trading period is redistributed into the next trading periods according to the liquidity profile: the more liquid one trading period is relative to another, the more favorable it is to increase its planned trade volume allocation rather than that of the other periods.

3.2 Model performance

In this section we analyze the impact of considering volume uncertainty as part of the model on the strategy performance in case 1; similar observations can be made for case 2. We compare the performance of the strategies obtained when minimizing the mean–$\mathrm{CVaR}_{0.1}$ trade-off in two distinct cases. The first case corresponds to a trader that wrongly assumes that the total volume to be traded is fixed and will not change this assumption in the course of the execution of their strategy (ie, for all $i \in \{1, \dots, m-1\}$ such that $\hat{\nu}_i = 0$). As a consequence, the corresponding strategies are those obtained if the trader considers that uncertainty only arises from the price dynamics (see Figure 1(a) and Table 2(a)). The second case corresponds to the situation where the trader correctly estimates the variability of the volume target forecasts, ie, for all $i \in \{1, \dots, m-1\}$ such that $\hat{\nu}_i = \nu_i$, and takes it into account when defining their trading strategy (see Figure 1(c) and Table 2(b)). Table 2 reports for both cases the expectation, the CVaR_α for $\alpha = 0.1$, 0.05 and 0.01 and the

FIGURE 2 Analysis of the impact of the bid–ask spread profile on the strategy adjustments of an exclusively risk-focused trader ($\lambda_{\text{CVaR}} = 1$) when considering the uncertainty related to the volume target.

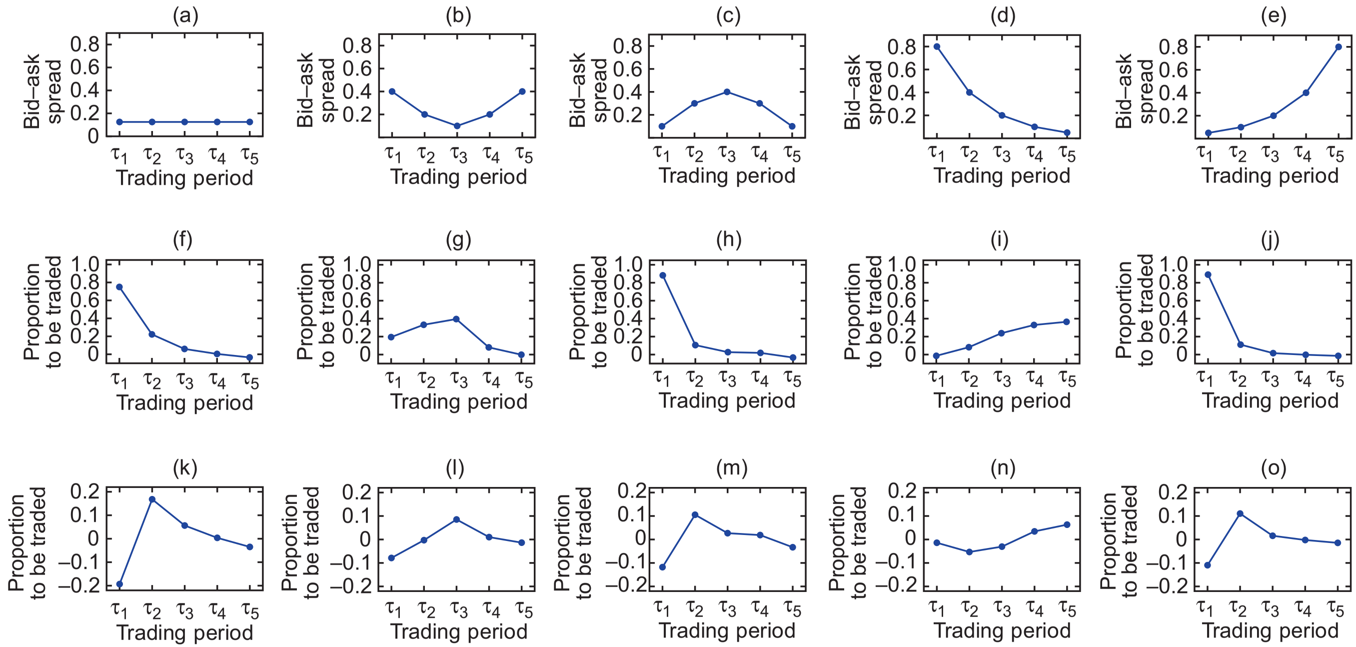

(a)–(e) Bid–ask spread profile. (f)–(j) Optimal strategy in the mean–$\text{CVaR}_{0.1}$ (PVU) framework. (k)–(o) Strategy adjustments between the strategies under PU and PVU.

TABLE 2 Performance of the strategies obtained by solving (2.16) for different risk-aversion parameter values λ_{CVaR}.

(a) Price dynamics

λ_{CVaR}	Expectation (US$)	$\text{CVaR}_{0.1}$ (US$)	$\text{CVaR}_{0.05}$ (US$)	$\text{CVaR}_{0.01}$ (US$)	Variance (US$^{2})	$\varphi_{0.1}^{\lambda_{\text{CVaR}}}$ (US$)
0.00	2.973×10^5	1.351×10^6	1.586×10^6	2.093×10^6	3.085×10^{11}	2.973×10^5
0.25	3.412×10^5	1.062×10^6	1.242×10^6	1.641×10^6	1.334×10^{11}	5.216×10^5
0.50	4.303×10^5	9.445×10^5	1.088×10^6	1.419×10^6	6.393×10^{10}	6.874×10^5
0.75	5.268×10^5	9.320×10^5	1.055×10^6	1.349×10^6	3.930×10^{10}	8.307×10^5
1.00	6.273×10^5	9.833×10^5	1.096×10^6	1.368×10^6	3.143×10^{10}	9.833×10^5

(b) Price dynamics and total volume to be traded

λ_{CVaR}	Expectation (US$)	$\text{CVaR}_{0.1}$ (US$)	$\text{CVaR}_{0.05}$ (US$)	$\text{CVaR}_{0.01}$ (US$)	Variance (US$^{2})	$\varphi_{0.1}^{\lambda_{\text{CVaR}}}$ (US$)
0.00	2.973×10^5	1.353×10^6	1.589×10^6	2.096×10^6	3.096×10^{11}	2.973×10^5
0.25	3.331×10^5	1.084×10^6	1.268×10^6	1.676×10^6	1.464×10^{11}	5.210×10^5
0.50	4.015×10^5	9.638×10^5	1.113×10^6	1.456×10^6	7.877×10^{10}	6.826×10^5
0.75	4.586×10^5	9.274×10^5	1.058×10^6	1.366×10^6	5.474×10^{10}	8.102×10^5
1.00	4.982×10^5	9.211×10^5	1.041×10^6	1.330×10^6	4.502×10^{10}	9.211×10^5

In part (a) a trader considers the only source of uncertainty to be the price dynamics (strategies under PU only depicted in Figure 1(a)). In part (b) a trader considers that the uncertainty arises from both the price dynamics and the total volume to be traded (strategies under PVU depicted in Figure 1(c)).

variance of the trading cost incurred when following the strategies that are obtained by solving model (2.16).

Figure 3 compares the probability density functions (pdfs) of the trading cost of both situations. The more risk-averse a trader, the more pronounced the difference between the densities is. Table 2 and Figure 3 clearly illustrate the benefit of including the uncertainty related to the volume target in the equivalent CVaR_α formulation of the risk–return trade-off model by Almgren and Chriss (2001): a trader that correctly assesses the uncertainty of the volume target achieves better mean–$\text{CVaR}_{0.1}$ trade-offs.

3.3 Comparison of the performance of the different frameworks

In the previous section we noted that a trader who correctly assesses the volume uncertainty achieves better mean–$\text{CVaR}_{0.1}$ trade-offs at the expense of a greater variance in their trading costs. In this section we compare our model with the ones in the

FIGURE 3 Comparison of the pdfs of the trading cost when the uncertainty related to the total volume to be traded is or is not taken into account in model (2.16) (the pdfs are estimated based on 10^8 realizations of the underlying processes).

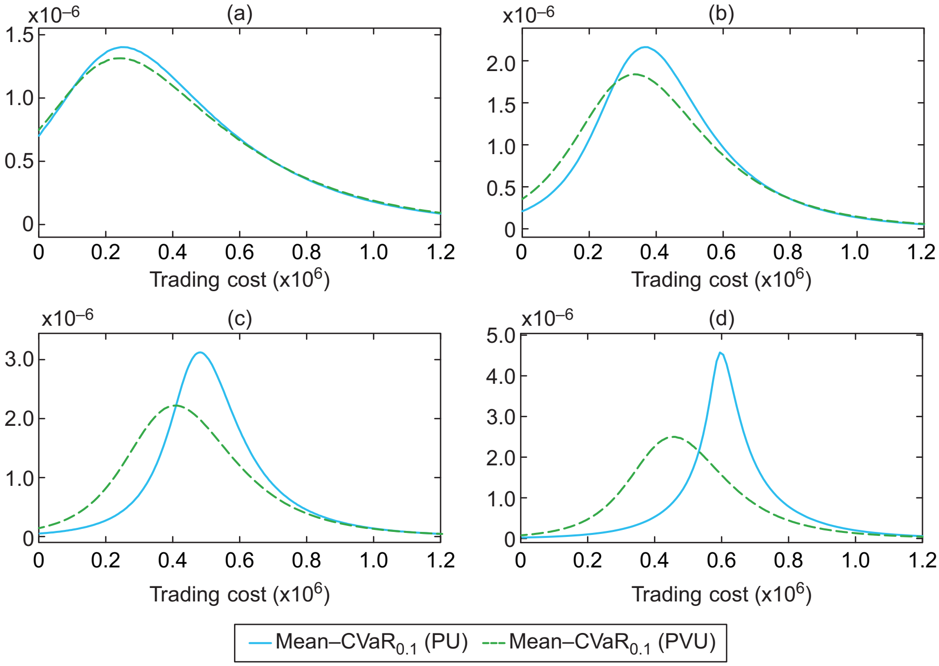

(a) $\lambda_{\text{CVaR}} = 0.25$. (b) $\lambda_{\text{CVaR}} = 0.5$. (c) $\lambda_{\text{CVaR}} = 0.75$. (d) $\lambda_{\text{CVaR}} = 1.0$.

literature. To evaluate the relative performance of each framework, we decided to use as comparison criteria the expectation and the CVaR_α of the trading costs, the variance not being of utmost importance, as justified hereafter. Additionally, we also compare the pdf and the CDF of the trading costs of the different frameworks.

3.3.1 Mean–variance with recourse

First, as mentioned previously the mean–variance framework from Almgren and Chriss (2001) is in practice not directly applicable in the presence of an uncertain volume target due to the constraint of optimization problem (2.3). However, a natural solution to circumvent this issue is to use it in conjunction with a systematic recourse at every trading period. Since the objective function in the mean–variance and mean–CVaR_α frameworks differs, it is not straightforward to compare the performance of the two models as we do not have a mapping between λ_{Var} and λ_{CVaR}. However, the

TABLE 3 Performance of the strategies derived from the mean–variance framework with recourse for different risk-aversion parameter values λ_{Var}.

λ_{Var} (1/US$)	Expectation (US$)	$\text{CVaR}_{0.1}$ (US$)	$\text{CVaR}_{0.05}$ (US$)	$\text{CVaR}_{0.01}$ (US$)	Variance (US$^{2})
0.0	2.974×10^5	1.351×10^6	1.587×10^6	2.094×10^6	3.087×10^{11}
10^{-6}	3.626×10^5	1.016×10^6	1.182×10^6	1.554×10^6	1.076×10^{11}
10^{-5}	5.616×10^5	9.336×10^5	1.042×10^6	1.302×10^6	3.335×10^{10}
10^{-4}	6.716×10^5	9.935×10^5	1.088×10^6	1.319×10^6	2.722×10^{10}
$\to +\infty$	6.912×10^5	1.008×10^6	1.102×10^6	1.328×10^6	2.698×10^{10}

optimal strategy in the mean–variance framework converges when the risk aversion increases. Hence, the parameter values $\lambda_{\text{CVaR}} = 1$ and $\lambda_{\text{Var}} \to +\infty$ both correspond to a trader that is exclusively risk focused; these values enable us to compare the performance of both models. The strategy of such a risk-focused trader in the mean–variance framework with recourse consists of trading the entire initial forecast of the total volume to be traded during the first trading period and then, during the subsequent trading periods τ_i, $i \in \{2, \dots, m\}$, trading an amount of shares δ_{i-1}, which corresponds to the forecast update unveiled at time t_{i-1}. This strategy can be reproduced in our model by choosing $\boldsymbol{y} = [1, 0, \dots, 0]^{\text{T}}$ and $\boldsymbol{\beta} = (\beta_{k,i})$ with $\beta_{k,i} = 1$ if $i = k + 1$ and 0 otherwise. The last rows of Table 2(a) and Table 3 report the performance of an exclusively risk-focused trader in both models; note that in both situations the trader only integrates PU into their estimate of the trading cost, but the actual trading costs are evaluated based on a market that is subjected to both sources of uncertainty. The last line of Table 2(b), which uses the same value of λ_{CVaR}, illustrates the benefit of factoring in the trade volume uncertainty: both the expectation and the $\text{CVaR}_{0.1}$ of the trading cost are significantly reduced. The standard deviation of the trading cost slightly increases from 1.773×10^5 to 2.122×10^5, but this is irrelevant: indeed, Figure 4, which depicts the probability densities of the trading cost for both models, shows empirically that the random variable that represents the trading cost in the mean–variance framework with recourse is first-order stochastically dominant over the trading cost of the strategy that results from model (2.16). As a consequence, a trader who is driven by profit should adopt the strategy proposed by model (2.16), as the increased variance is due to increased downward risk, which is beneficial.

3.3.2 EU and mean–QV

We have previously seen that the model proposed by Cheng *et al* (2017) could be slightly transformed in order to tackle a similar problem to ours in continuous time.

FIGURE 4 Comparison of the pdfs and CDFs of the trading cost of a risk-averse trader computing their strategy with the mean–variance framework with recourse for $\lambda_{\text{Var}} \to \infty$ and with the mean–CVaR$_{0.1}$ (PU and PVU) framework (model (2.16)) for $\lambda_{\text{CVaR}} = 1$.

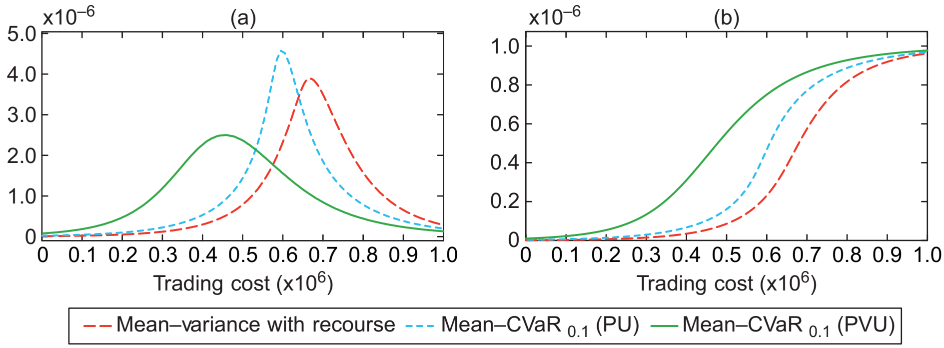

The pdfs are estimated based on 10^8 realizations of the underlying processes.

In their model Cheng *et al* assume that a risk-averse trader maximizes either the EU of their final P&L $\tilde{\Pi}_T(z)$ or the expectation of the final P&L that is penalized by its quadratic variation. In both situations they provide closed-form expressions for the optimal trading trajectory $z_t^\star$ when no residual volume is tolerated (ie, when $\beta \to \infty$). When allowing enough trading periods (eg, $m = 10^2$), the discrete approximation of the strategy derived from the model by Cheng *et al* (2017) has nearly exactly the same final P&L as that of the continuous strategy. Hence, to compare their model with ours we consider case 1 described in Table 1, with the difference being that we assume one trading day ($T = 1$) divided into 100 trading periods ($m = 10^2$). Note that the final P&L $\tilde{\Pi}_T(z)$ of a strategy z is equivalent to minus its trading cost.

Cheng *et al* (2017) first argue that the strategy resulting from the EU maximization can be regarded as a version of an adaptive Almgren–Chriss strategy. When the uncertainty of the volume target disappears (ie, $m_0 \to 0$), the optimal strategy, given the risk-aversion parameter θ_{EU} obtained via the EU framework, recovers the classical Almgren–Chriss strategy with parameter value $\lambda_{\text{Var}} = \theta_{\text{EU}}/2$. Hence, we can consider that a risk-averse trader with risk-aversion parameter λ_{Var} in the mean–variance framework would have a risk-aversion parameter θ_{EU} equal to $2\lambda_{\text{Var}}$ in the EU framework; this provides a mapping between λ_{Var} and θ_{EU} that allows the comparison of both frameworks across the range of risk-aversion levels. By comparing parts (a) and (b) of Table 4 we observe that the EU framework produces results of the same order as the mean–variance framework with recourse.

TABLE 4 Trading cost comparison of the trading strategies for the different frameworks studied with an execution period subdivided into $m = 10^2$ trading periods.

(a) Mean–variance with recourse

λ_{Var} (1/US$)	Expectation (US$)	$\text{CVaR}_{0.1}$ (US$)	$\text{CVaR}_{0.05}$ (US$)	$\text{CVaR}_{0.01}$ (US$)	Variance (US$^{2})
0	8.133×10^5	1.411×10^6	1.532×10^6	1.784×10^6	1.006×10^{11}
10^{-6}	8.154×10^5	1.399×10^6	1.519×10^6	1.765×10^6	9.557×10^{10}
10^{-5}	9.155×10^5	1.409×10^6	1.510×10^6	1.722×10^6	6.792×10^{10}
10^{-4}	2.071×10^6	2.424×10^6	2.496×10^6	2.647×10^6	3.547×10^{10}
$\to +\infty$	6.311×10^7	6.330×10^7	6.334×10^7	6.341×10^7	9.803×10^9

(b) EU

θ_{EU}	Expectation (US$)	$\text{CVaR}_{0.1}$ (US$)	$\text{CVaR}_{0.05}$ (US$)	$\text{CVaR}_{0.01}$ (US$)	Variance (US$^{2})
$\to 0$	7.861×10^5	1.264×10^6	1.348×10^6	1.512×10^6	7.423×10^{10}
2×10^{-6}	7.883×10^5	1.252×10^6	1.334×10^6	1.493×10^6	7.002×10^{10}
2×10^{-5}	9.354×10^5	1.313×10^6	1.379×10^6	1.510×10^6	4.631×10^{10}
2×10^{-4}	3.466×10^6	3.640×10^6	3.670×10^6	3.730×10^6	9.814×10^9
$\to +\infty$	2.780×10^7	2.794×10^7	2.797×10^7	2.802×10^7	7.081×10^9

(c) Mean–QV

λ_{QV} (1/US$)	Expectation (US$)	$\text{CVaR}_{0.1}$ (US$)	$\text{CVaR}_{0.05}$ (US$)	$\text{CVaR}_{0.01}$ (US$)	Variance (US$^{2})
$\to 0$	7.860×10^5	1.263×10^6	1.346×10^6	1.508×10^6	7.389×10^{10}
10^{-6}	7.885×10^5	1.255×10^6	1.336×10^6	1.495×10^6	7.075×10^{10}
10^{-5}	8.791×10^5	1.276×10^6	1.346×10^6	1.482×10^6	5.133×10^{10}
10^{-4}	1.637×10^6	1.902×10^6	1.948×10^6	2.038×10^6	2.282×10^{10}
$\to +\infty$	2.534×10^6	2.741×10^6	2.777×10^6	2.848×10^6	1.393×10^{10}

(d) Mean–$\text{CVaR}_{0.1}$ (PVU)

λ_{CVaR}	Expectation (US$)	$\text{CVaR}_{0.1}$ (US$)	$\text{CVaR}_{0.05}$ (US$)	$\text{CVaR}_{0.01}$ (US$)	Variance (US$^{2})
0.00	8.132×10^5	1.409×10^6	1.530×10^6	1.780×10^6	1.001×10^{11}
0.25	8.153×10^5	1.398×10^6	1.517×10^6	1.764×10^6	9.529×10^{10}
0.50	8.206×10^5	1.387×10^6	1.503×10^6	1.742×10^6	9.035×10^{10}
0.75	8.286×10^5	1.382×10^6	1.495×10^6	1.729×10^6	8.597×10^{10}
1.00	8.396×10^5	1.380×10^6	1.491×10^6	1.722×10^6	8.186×10^{10}

While Cheng *et al* (2017) employ the EU for the sake of tractability, they also provide, in the special case where the uncertainty of order fills is modeled with a diffusion process, a closed-form solution of the optimal trading rate in the mean–QV framework by solving the dynamic program analytically. They motivate the use of the quadratic variation, since, as pointed out by Almgren (2012), it captures the bulk of uncertainty. The optimal trading rate in this model can be regarded as an adaptive Almgren–Chriss strategy, with the difference that the characteristic timescales of liquidation are slightly different due to the uncertain order fills (Cheng *et al* 2017). When $m_0 \to 0$, the optimal strategy given the risk-aversion parameter λ_{QV} recovers the classical Almgren–Chriss strategy with parameter value $\lambda_{\mathrm{Var}} = \lambda_{\mathrm{QV}}$. This provides a mapping between λ_{Var} and λ_{QV}. The performance of the strategies derived with this model, given different risk-aversion parameter values λ_{QV}, is presented in Table 4(c): the more risk-averse the trader, the smaller the variance of their trading cost. Compared with the two previous models (mean–variance with recourse and EU maximization), for large risk-aversion parameters we observe a gain of one order of magnitude for both the expectation and the $\mathrm{CVaR}_{0.1}$ of the trading costs, whereas the variance remains (approximately) the same order.

Finally, we consider the model presented in this paper, where a trader considers the uncertainty arising from both the price dynamics and the updates on the volume target (ie, the mean–$\mathrm{CVaR}_{0.1}$ (PVU) framework). The performance of the strategies obtained with this model is shown in Table 4(d).

In a first step, as a sanity check, let us consider a risk-neutral trader (ie, $\lambda_{\mathrm{Var}} = 0$, $\theta_{\mathrm{EU}} \to 0$, $\lambda_{\mathrm{QV}} = 0$ and $\lambda_{\mathrm{CVaR}} = 0$). In this case all the frameworks minimize the same objective function, ie, the expectation of the trading costs. The first row of parts (a) and (d) of Table 4 shows that the mean–variance with recourse and the mean–$\mathrm{CVaR}_{0.1}$ (PVU) frameworks produce nearly identical results; this is coherent with Lemma 2.8. Indeed, since the redistribution matrix $\boldsymbol{\beta}$ was chosen as that implied by the risk-neutral strategy of the mean–variance framework computed at time t_0 (see (3.1)), the optimal strategy in the mean–variance framework with recourse belongs to the feasible set of strategies of the mean–$\mathrm{CVaR}_{0.1}$ (PVU) framework. Besides, we observe that the EU and mean–QV frameworks perform better than the two previous frameworks. Since both frameworks solve the same dynamical program, the discrepancy between the values reported in parts (b) and (c) of Table 4 for a risk-neutral trader can be attributed to numerical errors. This better performance arises from the fact that these frameworks solve the dynamic program and consequently reach the optimal trading strategy. We can therefore estimate the suboptimality gap of our model to $(8.129 \times 10^5 - 7.861 \times 10^5)/7.855 \times 10^5 \simeq 3.5\%$. This can be seen as the price to pay for taking only partial recourse via a predetermined rule.

In a second step we compare our model with the mean–QV framework for an exclusively risk-focused trader, which can be modeled by taking $\lambda_{\mathrm{QV}} \to +\infty$ and

FIGURE 5 Comparison of the pdfs and CDFs of the trading cost of a risk-averse trader computing their strategy with the mean–QV framework proposed by Cheng *et al* (2017) for $\lambda_{\text{QV}} \to \infty$ and with the mean–$\text{CVaR}_{0.1}$ (PVU) framework (model (2.16)) for $\lambda_{\text{CVaR}} = 1$.

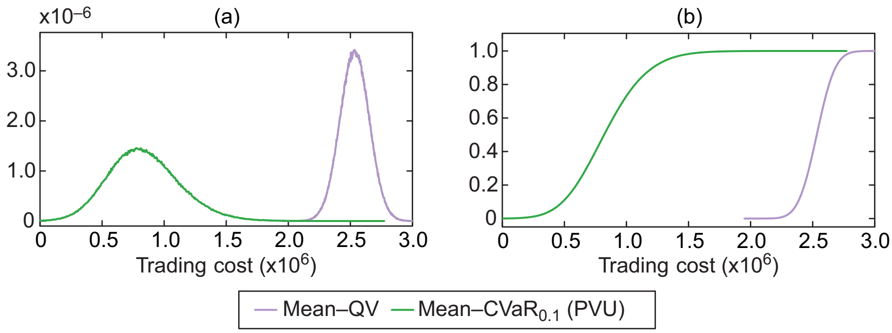

The pdfs are estimated based on 10^6 realizations of the underlying processes.

$\lambda_{\text{CVaR}} = 1$. Parts (c) and (d) of Table 4 illustrate that if a trader allows for a large variance of their trading cost, then a significant reduction in the expectation and the CVaR_α of the trading costs can be achieved by using our new model compared with the mean–QV framework. Similarly to the comparison with the mean–variance framework with recourse, we observe empirically in Figure 5 that the random variable representing the trading cost in the mean–QV framework is first-order stochastically dominant over the trading cost of the strategy resulting from model (2.16). Again, letting the variance increase enables our model to shift the distribution of the trading cost to smaller values, which is beneficial. In fact, the increase in variance is due to a larger spread in lower trading costs.

Finally, Table 4 reports that a risk-neutral trader in the mean–QV framework achieves a smaller value for the $\text{CVaR}_{0.1}$ of the trading costs compared with a trader that exclusively aims to minimize $\text{CVaR}_{0.1}$ in the mean–$\text{CVaR}_{0.1}$ (PVU) framework (ie, $\lambda_{\text{CVaR}} = 1$). This may call into question the advantage of our model. However, two important points should be brought to the reader's attention. First, despite the apparent better performance of the mean–QV framework, we should keep in mind that the analytical solution to the dynamic program in order to obtain the optimal mean–QV trading rate is available only for special well-behaved cases. In the general case, obtaining an analytical solution to the dynamic program would be more challenging, if not impossible, and obtaining a numerical solution would be significantly more computationally expensive than our approach. Second, the quadratic variation fails to correctly take into account the risk aversion of a trader. Indeed, as

FIGURE 6 Comparison of the pdfs and CDFs of the trading cost of a risk-averse trader computing their strategy with the mean–QV framework proposed by Cheng *et al* (2017) for different risk-aversion parameter values.

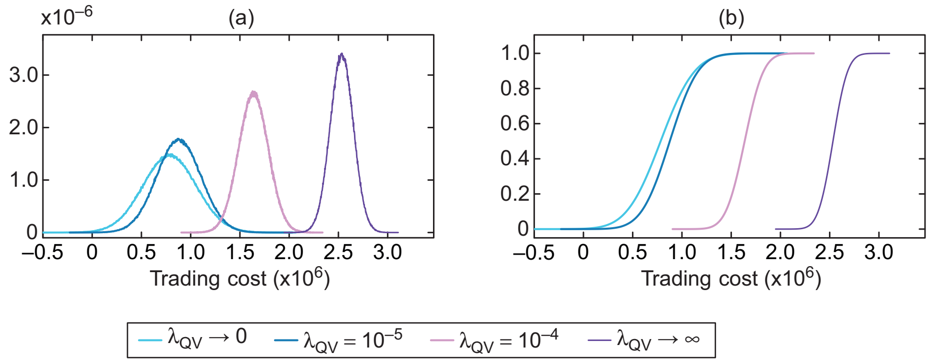

The pdfs are estimated based on 10^6 realizations of the underlying processes.

represented in Figure 6 and reported in Table 4(c), we generally observe that the greater the risk-aversion level of the trader (ie, λ_{QV}), the greater the expectation and $\mathrm{CVaR}_{0.1}$ of their trading costs. Hence, the decrease in the variance of the trading costs is coupled with a shift in the distribution of the trading costs to greater values, which is a rationale that would not be encountered in practice.

In conclusion, our model is significantly more versatile, as it does not rely on any assumptions on the underlying distributions that model the uncertainty related to the price dynamics and the forecast updates of the volume target. It is also computationally cheaper than the alternative approaches, and it coherently models the risk-aversion level of a trader: the expectation of the worst-case trading costs decreases as the risk-aversion level of the trader increases.

3.4 Real-world application: power trading

One of the main advantages of our model is that it does not rely on any particular assumptions on the distributions of the price moves and forecast updates. It can consequently be applied to a wider range of applications than the models of Almgren and Chriss (2001) and Cheng *et al* (2017). Here, we show how our model can be applied to power trading, where for each half-hourly settlement period power wholesalers (respectively, producers) trade power future contracts with the objective of achieving a contractual net position that matches the expected power (respectively, power generation) demand of their end-consumers.

FIGURE 7 Power trading time line and modeling assumptions.

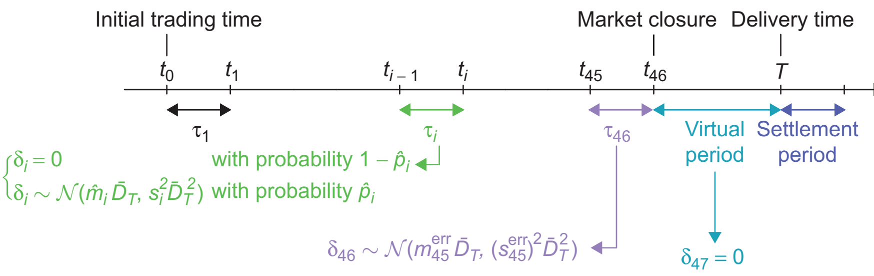

In Great Britain for every given settlement period the market participants can trade down future contracts from several years ahead until market closure, which occurs one hour before the start of delivery (ie, the start of the settlement period). For any given settlement period starting at time T, a wholesaler aims to achieve a final net position close to the realized power demand D_T of their end-consumers. If they supply power to a fraction ρ of the population, the power demand of their customers is here assumed to be given by $D_T = \rho V_T$, where V_T denotes the national power demand during the settlement period. Any residual volume resulting from the disparity between their contractual position against the realized demand D_T is charged at the imbalance price, which results from the balancing mechanism operated by the system operator after the market closure. For any half-hour settlement period, there is consequently no obligation for a market participant to balance their contractual positions against the expected demand (or generation). However, the exposure to this imbalance price drives the incentive to avoid any residual volume.

Here we investigate how a risk-averse power wholesaler in Great Britain should acquire power positions ahead of a given settlement period in order to minimize their costs while factoring in risk as suggested by our model (2.16). The market is modeled as follows. The settlement period is one half-hour long and starts at the delivery time T. The trading window spans from 2 to 48 half-hours ahead of delivery time; the initial trading time is denoted by t_0. The trading window is subdivided into 46 half-hour-long trading periods. For consistency with the notation introduced in Section 2.2, the ith trading period is denoted by τ_i and corresponds to the half-hour period that starts $48 - i + 1$ half-hours ahead of the settlement period (see Figure 7).

To make their trading decisions, the market participants use the operational data relating to the electricity balancing and settlement arrangements in Great Britain published by the Balancing Mechanism Reporting Service. Using the data from the past five years, Figure 8 shows the trend in the forecast updates of the national power

FIGURE 8 Analysis of the forecast updates of the national power demand in Great Britain.

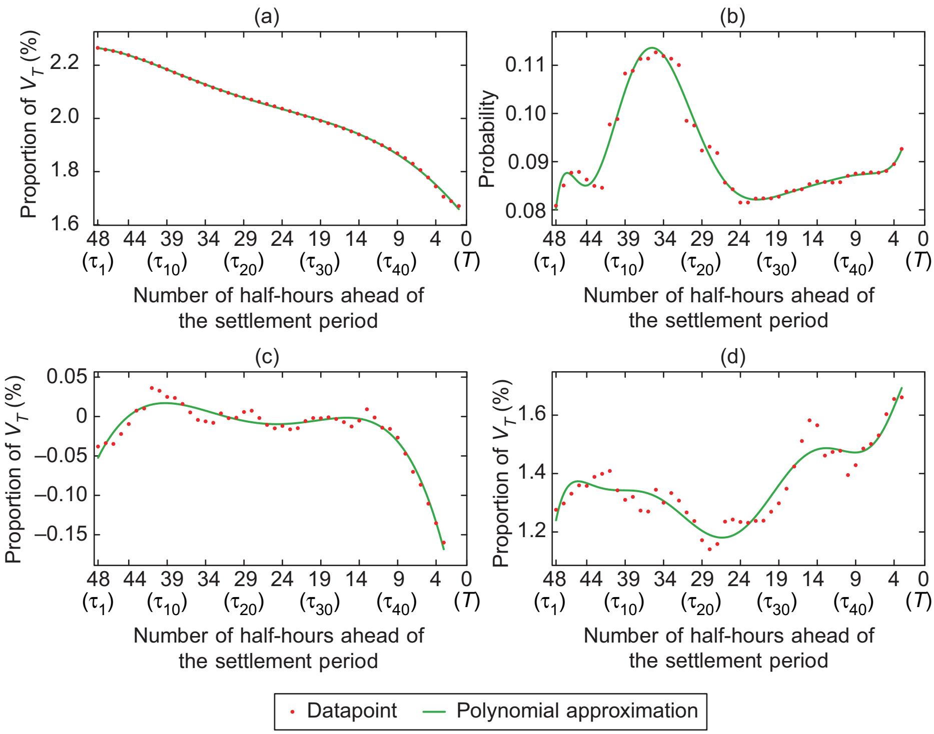

(a) Mean relative forecast error. (b) Probability forecast update. (c) Mean of the relative forecast update. (d) Standard deviation of the relative forecast update.

demand in Great Britain: part (a) illustrates the mean of the relative forecast error, which is computed as $|V_i - V_T|/V_T$, where V_T denotes the realized national demand and V_i denotes the demand forecast unveiled after the ith trading period τ_i (ie, at decision time t_i). Part (b) represents the probability p_i of having a forecast update of the power demand at time t_i. Conditional on there being a forecast update, parts (c) and (d) of Figure 8 respectively represent the mean m_i and the standard deviation s_i of the relative forecast update given by $(V_i - V_{i-1})/V_T$. For all graphs we have fitted a polynomial approximation of the data points; let $\hat{p}_i$, $\hat{m}_i$ and $\hat{s}_i$ denote the value of the polynomial approximating the probability, the mean and the standard deviation of the forecast update that comes to the trader's knowledge at time t_i, respectively. Based on these values we assume that, for each trading period τ_i with $i \leqslant 45$, a wholesaler estimates the distribution of the related forecast update on the power demand of their customers as follows. First, there is a probability $\hat{p}_i$ of having a forecast update. Second, conditional on a forecast update occurring, the forecast update δ_i is assumed to follow the normal distribution $\mathcal{N}(\hat{m}_i \bar{D}_T, \hat{s}_i^2 \bar{D}_T^2)$, where $\bar{D}_T = \rho \bar{V}_T$ and $\bar{V}_T$ denote the average realized national power demand. Finally, because the market closure happens one hour before the start of delivery, a trader cannot take into consideration in their trading strategy any forecast update that is released after the decision time t_{45}. Therefore, the residual volume that should have been traded and that will be exchanged at the imbalance price is given by the random variable $\rho(V_T - V_{45})$. This is taken into consideration in our model by assuming that the forecast update δ_{46} of the last trading period follows the normal distribution $\mathcal{N}(\hat{m}_{45}^{\text{err}} \bar{D}_T, (\hat{s}_{45}^{\text{err}})^2 \bar{D}_T^2)$, where $\hat{m}_{45}^{\text{err}}$ and $\hat{s}_{45}^{\text{err}}$ are, respectively, the estimates of the average and standard deviation of the relative forecast error at decision time t_{45}, ie, $(V_T - V_{45})/V_T$. In addition to these 46 trading periods, we add a virtual 47th trading period, which corresponds to the balancing period after the market closure. Note that during this balancing period we artificially assume that there is no forecast update on the demand, as it has already been taken into consideration in the forecast update distribution of τ_{46}. To incentivize the trader to trade the entirety of their position before the balancing period τ_{47}, we penalize any outstanding volume by fixing a larger temporary impact parameter η_{47} for the balancing period τ_{47} compared with the other trading periods; a value three times larger is used here (see Table 5). In a similar vein as the penalty term that is parameterized by β in the model by Cheng *et al* (2017), the larger the η_{47}, the more any outstanding volume is penalized. As suggested in Section 2.3, we assume that the residual volumes to be traded due to the forecast updates are redistributed over the future trading periods based on the redistribution matrix $\boldsymbol{\beta}$ that is implied by the risk-neutral optimal trading strategy $\boldsymbol{y}^{\star,t_0,\lambda_{\text{Var}}=0}$ in the mean–variance framework.

From the perspective of the price distributions we consider an initial price S_0 of 90. For the sake of simplicity (the data not being publicly available) we assume that the volatility σ of the asset is constant during the entire execution period and is equal to 0.9 and that for each trading period τ_i the price shift ξ_i follows the normal distribution $\mathcal{N}(0, \tau_i \sigma^2)$. The rest of the parameters are fixed as follows: the average realized national power demand V_T equals 3×10^4 megawatts (MW). The wholesaler provides power to $\rho = \frac{1}{3}$ of the population, has an initial position of zero and has an initial estimate of the power demand of the end-consumers equal to $D_0 = \bar{D}_T = \rho \bar{V}_T = 10^4$. Finally, the bid–ask spread of each trading period is assumed to be equal to 0.25% of the initial asset price S_0 (ie, for all $i \in \{1, \dots, m\}$ such that $b_i = 0.225$); the values used for temporary and permanent impact parameters can be found in Table 5.

Similarly to Figure 1, Figure 9(a) depicts the trading strategies returned by our model for a risk-averse wholesaler in Great Britain when they consider the

TABLE 5 Power market parameters.

Number of trading periods	$m = 47$
Trading period length	for all $i \in \{1, \dots, m\}$, $\tau_i = \frac{1}{48}$ (day)
Initial future contract price	$S_0 = 90$ (£/MW)
Asset's volatility	$\sigma = 0.9$ ((£/MW)/day$^{1/2}$)
Price shifts distribution	for all $i \in \{1, \dots, m\}$, $\xi_i \sim \mathcal{N}(0, \tau_i \sigma^2)$
Average realized national power demand	$\bar{V}_T = 3 \times 10^4$ (MW)
Proportion of the population supplying	$\rho = \frac{1}{3}$
Initial forecast of the customers' power demand	$D_0 = \bar{D}_T = \rho \bar{V}_T = 10^5$ (MW)
Volume forecast updates distribution	for all $i \in \{1, \dots, 45\}$, $\delta_i \sim \mathcal{N}(\hat{m}_i \bar{D}_T, \hat{s}_i^2 \bar{D}_T^2)$
	$\delta_{46} \sim \mathcal{N}(\hat{m}_{45}^{\text{err}} \bar{D}_T, (\hat{s}_{45}^{\text{err}})^2 \bar{D}_T^2)$
	$\delta_{47} = 0$
Bid–ask spread	for all $i \in \{1, \dots, m\}$, $b_i = 0.25\% S_0 = 0.225$ (£/MW)
Fixed cost	for all $i \in \{1, \dots, m\}$, $\epsilon_i = 0.5 b_i = 0.1125$ (£/MW)
Impact at 2% of market	for all $i \in \{1, \dots, 46\}$, $\eta_i = b_i / (0.02 \bar{V}_T) = 3.75 \times 10^{-5}$ ((£/MW)/(MW/day))
	$\eta_{47} = 3 b_{47} / (0.02 \bar{V}_T) = 11.25 \times 10^{-5}$ ((£/MW)/(MW/day))
Permanent impact parameter	for all $i \in \{1, \dots, m\}$, $\gamma_i = (5 b_i)/(0.1 \bar{V}_T) = 3.75 \times 10^{-5}$(£/MW^2)
CVaR$_\alpha$ parameter	$\alpha = 0.1$

FIGURE 9 Trading power in Great Britain.

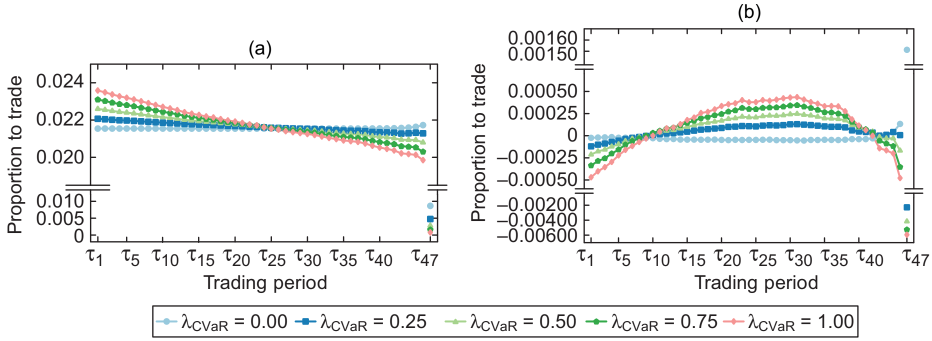

(a) Optimal trading strategies under PVU. (b) Difference between the strategies under PVU and PU.

uncertainty related to both the price dynamics and the power demand of the end-consumers. For different risk-aversion levels Figure 9(b) depicts how the trading strategy is modified compared with the situation where the wholesalers wrongly assume that the power demand of the end-consumers is perfectly known at the start of the execution period. The difference in performance of the trading strategies obtained in both situations is listed in Table 6. When we consider the uncertainty related to the power demand, better mean–$\text{CVaR}_{0.1}$ trade-offs are achieved for any risk level. We do not analyze the results further, as similar comments can be made compared with the previous test cases (see Sections 3.1 and 3.2).

4 DISCUSSION

Our choice to use the CVaR_α as a risk measure in model (2.16) was guided by its widespread use and ease of interpretation, despite the controversy that this risk measure may lead to time-inconsistent optimal policies: the policy obtained by optimizing the model at an intermediate stage of the execution period is not necessarily equivalent to the continuation of the optimal policy computed at the initial time, t_0 (Rudloff *et al* 2014; Shapiro 2009). In Section 3, however, we saw that, due to the forward planning of recourse with a predetermined trade volume reallocation rule, our model competes favorably with established alternatives, yielding lower trading costs in expectation and CVaR_α. Our merit function accounts for the trade-off between the partially conflicting goals of lowering the expectation and the CVaR_α of trading costs, and it is computationally cheaper than either taking recourse under the mean–variance framework whenever forecast updates on the volume target become

TABLE 6 Trading power in Great Britain: performance of the strategies obtained with (2.16) for different risk-aversion parameter values λ_{CVaR}.

(a) Price dynamics

λ_{CVaR} (1/US$)	Expectation (US$)	$\text{CVaR}_{0.1}$ (US$)	$\text{CVaR}_{0.05}$ (US$)	$\text{CVaR}_{0.01}$ (US$)	Variance (US$^{2})	$\varphi_{0.1}^{\lambda_{\text{CVaR}}}$ (US$)
0.00	6.102×10^4	7.572×10^4	7.945×10^4	8.841×10^4	5.558×10^7	6.102×10^4
0.25	6.103×10^4	7.566×10^4	7.938×10^4	8.832×10^4	5.503×10^7	6.469×10^4
0.50	6.105×10^4	7.562×10^4	7.933×10^4	8.825×10^4	5.449×10^7	6.834×10^4
0.75	6.109×10^4	7.559×10^4	7.929×10^4	8.819×10^4	5.396×10^7	7.196×10^4
1.00	6.114×10^4	7.559×10^4	7.928×10^4	8.818×10^4	5.345×10^7	7.559×10^4

(b) Price dynamics and power demand

λ_{CVaR} (1/US$)	Expectation (US$)	$\text{CVaR}_{0.1}$ (US$)	$\text{CVaR}_{0.05}$ (US$)	$\text{CVaR}_{0.01}$ (US$)	Variance (US$^{2})	$\varphi_{0.1}^{\lambda_{\text{CVaR}}}$ (US$)
0.00	6.101×10^4	7.606×10^4	7.993×10^4	8.922×10^4	5.751×10^7	6.101×10^4
0.25	6.110×10^4	7.531×10^4	7.885×10^4	8.737×10^4	5.282×10^7	6.465×10^4
0.50	6.122×10^4	7.510×10^4	7.850×10^4	8.668×10^4	5.094×10^7	6.816×10^4
0.75	6.134×10^4	7.503×10^4	7.837×10^4	8.638×10^4	4.988×10^7	7.161×10^4
1.00	6.144×10^4	7.502×10^4	7.832×10^4	8.623×10^4	4.914×10^7	7.502×10^4

(a) A trader considers that the only source of uncertainty is the price dynamics. (b) A trader considers that the uncertainty arises from both the price dynamics and the power demand.

available or solving a dynamic program in order to optimize the mean–QV trade-off. These advantages of our model outweigh the disadvantage that the time inconsistency of the mean–CVaR_α risk measure may lead to suboptimal strategies (Rudloff *et al* 2014).

The model proposed in this paper should thus be considered as an alternative to the traditional methods for its competitive performance, its ease of interpretation and its relatively cheap computational cost. If we desire to take full recourse in order to further improve the performance of our model, we could replace the CVaR_α term with the dynamic coherent risk measure (Riedel 2004) induced by the CVaR_α risk measure. Lin *et al* (2015) explored such an approach in the context of the traditional trade execution problem under PU only. An extension to the case with volume uncertainty would imply the applicability of the dynamic programming principle, which guarantees the time consistency of the optimal strategies but at the cost of a greater computational time. This would be a useful extension of our work, and it is left for future research.

Among the parameters for our model we need to specify a redistribution matrix to decide how any residual trading volumes that become apparent after forecast updates are redistributed over future trading periods. In the numerical section of this paper we chose this redistribution matrix according to the trade volume proportions implied by the risk-neutral optimal trading strategy $\boldsymbol{y}^{\star,t_0,\lambda_{\mathrm{Var}}=0}$ under the mean–variance framework of Almgren and Chriss (2001). This choice is not guaranteed to be optimal, and we could optimize over the redistribution matrix as well. Although we do not address this issue in the present work, we believe it is of interest for future research. The numerical simulations we carried out suggest that mild additional cost savings can be achieved.

Finally, in this paper we assumed that the random variables that model forecast updates are identically distributed. In practice this assumption is often not justified, such as in power markets, where empirical data show that the standard deviation of forecast updates is time dependent and follows a "∩" pattern over the trading window: forecast updates at long time horizons are relatively small, since they are mainly based on the seasonal trend. At shorter timescales, that is, for time horizons ranging from a couple of weeks to several hours ahead of the delivery time, the forecast updates are more significant, since weather forecasts are available and are steadily becoming more precise. Finally, the last forecast updates represent slight adjustments to ensure that the market clearing constraint is satisfied. Another example where inhomogeneous forecast updates occur is in open-ended funds, where the fund manager – if notified as and when share remissions and purchases are ordered during the trading day – could proactively trade in the underlying assets on the trading day itself rather than on the next day, so as to increase the liquidity of the fund. In this case the standard deviation of the updates on the volume target often follows a "∪" profile instead. Further investigation is thus needed to analyze the impact of the forecast uncertainty profile on the performance of our approach relative to the model by Almgren and Chriss (2001). Since the optimal redistribution coefficients depend on the profile of the standard deviation of the forecast updates, these two points should be jointly addressed in further investigations.

5 CONCLUSION

In this paper we considered the optimal trade execution problem in a setting where price uncertainty grows in time, but unlike what is assumed in the existing literature, the required trade volume is also uncertain and only becomes known at the end of the execution period. We assume that forecasts of increasing accuracy become available, so that volume uncertainty decreases over time while price uncertainty increases. The model presented in this paper is designed to manage both uncertainties via risk terms, so as to make them precomputable and to avoid the combinatorial blowup of dynamic programming approaches. We have demonstrated that the model has desirable convexity properties, guaranteeing the existence and uniqueness of a convex set of optimal strategies. We show that it produces significantly lower transaction costs than classical trade execution approaches, even if the latter allow for recourse whenever a new forecast becomes available.

DECLARATION OF INTEREST

The authors report no conflicts of interest. The authors alone are responsible for the content and writing of the paper.

ACKNOWLEDGEMENTS

We thank the anonymous reviewer for their careful reading and constructive feedback; their input allowed us to improve our paper substantially. The work of Julien Vaes was supported by the Alan Turing Institute under a Turing Doctoral Studentship (grant TU/C/000022). The work of Raphael Hauser was supported by the Alan Turing Institute under the Engineering and Physical Sciences Research Council grant EP/N510129/1.

REFERENCES

Almgren, R. (2012). Optimal trading with stochastic liquidity and volatility. *SIAM Journal on Financial Mathematics* **3**, 163–181 (https://doi.org/10.1137/090763470).

Almgren, R., and Chriss, N. (2001). Optimal execution of portfolio transactions. *The Journal of Risk* **3**(2), 5–39 (https://doi.org/10.21314/JOR.2001.041).

Almgren, R., and Lorenz, J. (2007). Adaptive arrival price. *Trading* **2007**(1), 59–66. URL: https://guides.pm-research.com/content/iijtrading/2007/1/59.full.pdf.

Almgren, R., Thum, C., Hauptmann, E., and Li, H. (2005). Direct estimation of equity market impact. *Risk* **18**(7), 58–62. URL: https://api.semanticscholar.org/CorpusID:18312213.

Artzner, P., Delbaen, F., Eber, J. M., and Heath, D. (1999). Coherent measures of risk. *Mathematical Finance* **90**(3), 203–228 (https://doi.org/10.1111/1467-9965.00068).

Bertsimas, D., and Lo, A. W. (1998). Optimal control of execution costs. *Journal of Financial Markets* **10**(1), 1–50 (https://doi.org/10.1016/S1386-4181(97)00012-8).

Boyd, S., and Vandenberghe, L. (2004). *Convex Optimization*. Cambridge University Press (https://doi.org/10.1017/CBO9780511804441).

Bulthuis, B., Concha, J., Leung, T., and Ward, B. (2017). Optimal execution of limit and market orders with trade director, speed limiter, and fill uncertainty. *International Journal of Financial Engineering* **40** (2–3), Paper 1750020 (https://doi.org/10.1142/S2424786317500207).

Butenko, S., Golodnikov, A., and Uryasev, S. (2005). Optimal security liquidation algorithms. *Computational Optimization and Applications* **32**(1), 9–27 (https://doi.org/10.1007/s10589-005-2052-9).

Chan, K., Chung, Y. P., and Johnson, H. (1995). The intraday behavior of bid–ask spreads for NYSE stocks and CBOE options. *Journal of Financial and Quantitative Analysis* **30**(3), 329–346 (https://doi.org/10.2307/2331344).

Cheng, X., Di Giacinto, M., and Wang, T. H. (2017). Optimal execution with uncertain order fills in Almgren–Chriss framework. *Quantitative Finance* **17**(1), 55–69 (https://doi.org/10.1080/14697688.2016.1185531).

Cheng, X., Di Giacinto, M., and Wang, T. H. (2019). Optimal execution with dynamic risk adjustment. *Journal of the Operational Research Society* **70**(10), 1662–1677 (https://doi.org/10.1080/01605682.2019.1644143).

Feng, Y., Rubio, F., and Palomar, D. P. (2012). Optimal order execution for algorithmic trading: a CVaR approach. In *Proceedings of the IEEE 13th International Workshop on Signal Processing Advances in Wireless Communications*, pp. 480–484. IEEE Press, Piscataway, NJ (https://doi.org/10.1109/SPAWC.2012.6292954).

Forsyth, P. A., Kennedy, J. S., Tse, S. T., and Windcliff, H. (2012). Optimal trade execution: a mean quadratic variation approach. *Journal of Economic Dynamics and Control* **36**(12), 1971–1991 (https://doi.org/10.1016/j.jedc.2012.05.007).

Gatheral, J., and Schied, A. (2013). Dynamical models of market impact and algorithms for order execution. In *Handbook on Systemic Risk*, Fouque, J.-P., and Langsam, J. A. (eds), pp. 579–599. Cambridge University Press (https://doi.org/10.1017/CBO9781139151184.030).

Huberman, G., and Stanzl, W. (2004). Price manipulation and quasi-arbitrage. *Econometrica* **72**(4), 1247–1275 (https://doi.org/10.1111/j.1468-0262.2004.00531.x).

Huberman, G., and Stanzl, W. (2005). Optimal liquidity trading. *Review of Finance* **9**(2), 165–200 (https://doi.org/10.1007/s10679-005-7591-5).

Kissell, R., and Malamut, R. (2005). Algorithmic decision-making framework. *Journal of Trading* **1**(1), 12–21 (https://doi.org/10.3905/jot.2006.609171).

Krokhmal, P., and Uryasev, S. (2007). A sample-path approach to optimal position liquidation. *Annals of Operations Research* **152**, 193–225 (https://doi.org/10.1007/s10479-006-0143-3).

Lin, Q., Chen, X., and Peña, J. (2015). A trade execution model under a composite dynamic coherent risk measure. *Operations Research Letters* **43**(1), 52–58 (https://doi.org/10.1016/j.orl.2014.11.005).

Lorenz, J., and Almgren, R. (2011). Mean–variance optimal adaptive execution. *Applied Mathematical Finance* **18**(5), 395–422 (https://doi.org/10.1080/1350486X.2011.560707).

Markowitz, H. (1952). Portfolio selection. *Journal of Finance* **70**(1) 77–91 (https://doi.org/10.1111/j.1540-6261.1952.tb01525.x).

Riedel, F. (2004). Dynamic coherent risk measures. *Stochastic Processes and Their Applications* **112**(2), 185–200 (https://doi.org/10.1016/j.spa.2004.03.004).

Rockafellar, R. T., and Uryasev, S. (2002). Conditional value-at-risk for general loss distributions. *Journal of Banking and Finance* **26**(7), 1443–1471 (https://doi.org/10.1016/S0378-4266(02)00271-6).

Rudloff, B., Street, A., and Valladão, D. M. (2014). Time consistency and risk averse dynamic decision models: definition, interpretation and practical consequences. *European Journal of Operational Research* **234**(3), 743–750 (https://doi.org/10.1016/j.ejor.2013.11.037).

Schied, A., and Schöneborn, T. (2009). Risk aversion and the dynamics of optimal liquidation strategies in illiquid markets. *Finance and Stochastics* **13**(2), 181–204 (https://doi.org/10.1007/s00780-008-0082-8).

Schöneborn, T. (2016). Adaptive basket liquidation. *Finance and Stochastics* **20**(2), 455–493 (https://doi.org/10.1007/s00780-016-0290-6).

Shapiro, A. (2009). On a time consistency concept in risk averse multistage stochastic programming. *Operations Research Letters* **37**(3), 143–147 (https://doi.org/10.1016/j.orl.2009.02.005).

Shapiro, A., Dentcheva, D., and Ruszczyński, A. (2009). *Lectures on Stochastic Programming: Modeling and Theory*. MOS–SIAM Series on Optimization. SIAM, Philadelphia, PA (https://doi.org/10.1137/1.9780898718751).