Mathematics and Financial Economics (2021) 15:315–352
https://doi.org/10.1007/s11579-020-00278-7

Equilibrium effects of intraday order-splitting benchmarks

Jin Hyuk Choi[1] · **Kasper Larsen**[2] · **Duane J. Seppi**[3]

Received: 25 November 2019 / Accepted: 14 August 2020 / Published online: 4 September 2020

Abstract
This paper presents a continuous-time model of intraday trading, pricing, and liquidity with dynamic TWAP and VWAP benchmarks. The model is solved in closed-form for the competitive equilibrium and also for non-price-taking equilibria. The intraday trajectories of TWAP trading targets cause predictable intraday patterns of price pressure, and randomness in VWAP target trajectories induces additional randomness in intraday price-pressure patterns. TWAP and VWAP trading both reduce market liquidity and increase price volatility relative to just terminal trading targets alone. The model is computationally tractable, which lets us provide a number of numerical illustrations.

Keywords Dynamic trading · TWAP · VWAP · Portfolio rebalancing · Liquidity · Market-maker inventory · Equilibria · Market microstructure

JEL Classification G12 · G14

Dynamic order-splitting strategies are a prominent feature of present-day financial markets. As described in O'Hara [35], large asset managers use sequences of small *child orders* to trade on large latent meta *parent demands*. These strategies include heuristically mimicking

The authors have benefited from helpful comments from Frank Riedel, an anonymous associate editor and referee, Yashar Barardehi, Umut Cetin, Chris Frei, Steve Karolyi, Ajitesh Mehta, Johannes Muhle-Karbe, Dan Ocone, Tom Ruchti, Chester Spatt, Harvey Stein, Kim Weston, Ariel Zetlin-Jones, and seminar participants at Carnegie Mellon University, Columbia University, IAQF/Thalesians, and Intech. Jin Hyuk Choi is supported by the National Research Foundation of Korea (NRF) and the Korea government (MSIT) (Nos. 2020R1C1C1A01014142 and 2017R1E1A1A03070732). Kasper Larsen is supported by the National Science Foundation (NSF) under Grant No. DMS 1812679 (2018–2021). Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation (NSF).

✉ Kasper Larsen
KL756@math.rutgers.edu

Jin Hyuk Choi
jchoi@unist.ac.kr

Duane J. Seppi
ds64@andrew.cmu.edu

1 Ulsan National Institute of Science and Technology (UNIST), Ulsan, South Korea

2 Rutgers University, New Brunswick, USA

3 Carnegie Mellon University, Pittsburgh, USA

the time-weighted average price (TWAP) or volume-weighted average price (VWAP) and also optimized strategies from formal execution-cost minimization problems (see, e.g., [1,2]). The scale of order splitting in present-day markets is substantial. The Financial Insights [18] trading survey reports that one popular strategy, VWAP execution orders, represents around 50% of all institutional trading.[1] Dynamic trading strategies can be price-sensitive in that they respond to the price of liquidity as investors trade off price improvement and trading profits vs. tracking error relative to intraday trading targets.[2] In addition, high frequency market makers use dynamic strategies to supply liquidity while attempting to keep their inventory close to an intraday target of zero.[3] The result is a complex ecosystem in which asset managers, other position-taking traders, and high frequency market makers all use dynamic trading strategies to supply and demand liquidity given different benchmarked intraday targets.

Our paper is the first to model the equilibrium impact of TWAP, VWAP, and other dynamic benchmarks on intraday trading and liquidity. Specifically, we focus on liquidity effects of order splitting rather than on information aggregation as in Kyle [30]. We take intraday trading target trajectories and penalties as inputs to our model, and then show how they affect price dynamics and liquidity over the trading day via their effect on aggregate order imbalances. In particular, we model a market with multiple price-sensitive investors with heterogeneous trading targets who follow optimal continuous-time dynamic strategies. We solve for equilibria in closed-form. Our main results are as follows:

- Order-splitting benchmarks like TWAP, Almgren-Chriss targets, and VWAP lead to predictable intraday price pressure due to persistant latent trading-demand imbalances. In particular, aggregate latent buy (sell) imbalances lead to positive (negative) price pressure at the market open that then decays, i.e., trends down (up) on average, over the trading day.
- Trading benchmarks reduce intraday liquidity and increase price volatility relative to when investors just have terminal end-of-day trading targets. This is because penalties for intraday deviations from trading target trajectories reduce the inventory-holding flexibility of investors over the day.
- Comparative static results lead to empirical predictions about price dynamics. For example, market illiquidity and price volatility in a competitive Radner equilibrium are increasing in investor order-execution costs and inventory penalties for trading deviations from trading target trajectories.
- In a competitive Radner equilibrium, price-sensitive investors deviate from their ideal trading targets by a fraction of the aggregate latent trading-demand imbalance. In non-price-taking equilibria, investor target deviations can also depend on individual investor targets.

Our research is related to several prior research literatures. First, there is a large optimal control literature on optimal order execution with exogenous price impact. This includes Bertsimas and Lo [6], Almgren and Chriss [1,2], Gatheral and Schied [22], Engle et al. [17], Predoiu et al. [37], Boulatov et al. [8] and other research surveyed in Gatheral and Shied

[1] Large asset managers conduct dynamic trading strategies using in-house trading desks and also via principal and agency trading with external brokers.

[2] Madhavan [32] discusses price improvement on order execution relative to VWAP. Domowitz and Yegerman [13] estimate empirical order-execution costs benchmarked relative to VWAP.

[3] Hagströmer and Nordén [26] and Menkveld [34] show that high-frequency (HFT) market makers are an important source of intraday liquidity. A common feature of HFT market makers is that they have "very short time-frames for establishing and liquidating positions" SEC [41], which is consistent with a zero target inventory level. Weller [45] shows further that liquidity over the trading day is provided by a network of liquidity providers with slower and faster trading latencies who shift inventory between themselves over holding periods of different lengths. This behavior is also consistent with zero intraday inventory targets.

[23]. In contrast, we model dynamic trading benchmarks and their effects in an equilibrium framework. Second, the practice of benchmarking order-execution quality with VWAP and other metrics is described in Berkowitz et al. [5] and Madhavan [32].[4] Baldauf et al. [3] show that VWAP benchmarking is optimal for certain principal-agent problems in delegated order execution. In contrast, we model the market effects of trading benchmarks rather than the reasons why investors have such benchmarks. Third, Korajczyk and Murphy [29], van Kervel and Menkveld [43], and van Kervel et al. [42] document empirical interactions between dynamic trading by different investors. Our work is most closely related to van Kervel et al. [42], which shows theoretically and empirically how dynamic trading strategies interact across multiple strategic investors. In contrast, we model the equilibrium price effects of dynamic trading as well as trading interaction effects.

Our analysis extends the costly inventory model of market making—see, e.g., Garman [21], Stoll [40], and Grossman and Miller [24]—by allowing for endogenous arriving orders from optimized order-splitting strategies. In particular, our model highlights the idea that, with endogenous investor trading demand, price pressure in equilibrium prices is not just due to inventory-holding costs of liquidity providers on executed transactions. Rather, equilibrium price pressure also depends on the underlying latent investor trading-demand imbalances. In order for markets to clear, equilibrium prices deter investors from trading on trading demands for which there is no counterparty. This is a new perspective relative to more transactional market microstructure models, but one that is consistent with empirical evidence in van Kervel et al. [42] that large investors trade less in the direction of aggregate imbalances in underlying parent trading demands. Our modeling approach is related to previous research by Brunnermeier and Pedersen [9] and Carlin et al. [10] on optimal rebalancing and predatory trading. There are two main differences: First, investors in our model are subject to penalties tied to intraday trading target trajectories rather than to just a single hard terminal constraint at the end of the day. Second, there are no ad hoc intraday liquidity providers in our model. Instead, all intraday liquidity is provided endogenously by rational optimizing investors. As a result, there is no predatory trading in our model even when investors are strategic.[5] Our paper also builds on the Vayanos [44] dynamic trading model. Like Vayanos [44]—and the continuous-time extension in Sannikov and Skrzypacz [39]—our model includes investors who smooth a series of random intraperiod shocks but, in addition, we also have investors with private heterogeneous trading target trajectories they would like to follow over the day. The discrete-time model in Du and Zhu [14] also has quadratic penalties but with zero targets and has private dividend information (whereas dividend information is public in our model).

Our model is also related to Gârleanu and Pedersen [20] and the multi-investor extension in Bouchard et al. [7]. In these competitive equilibrium models, traders incur penalties based on their stock holdings and their trading rates. Our model differs because we allow for possible price-impact of individual trades and for non-zero trading targets that are ex ante private information.

Lastly, there is no asymmetric information about future asset cash-flow fundamentals in our model. Thus, our analysis here on trading and non-informational price pressure is complementary to Choi et al. [11], which studies order-splitting and dynamic rebalancing in a Kyle [30] style market in which a strategic informed investor with long-lived private

[4] Implementation shortfall is another alternative trading benchmark; see Perold [36].

[5] Predatory trading is a manipulative strategy to buy (sell) to raise (lower) prices artificially and then unwind these positions at a profit given price support from predictable persistent buying (selling) by a large investor rebalancing its portfolio. In Brunnermeier and Pedersen [9], predatory trading is possible because of ad hoc liquidity providers who trade using exogenous linear schedules that do not rationally anticipate predictable future price changes later in the day given earlier trading.

information and a strategic rebalancer with a hard terminal trading target both follow dynamic trading strategies.

1 Model primitives

The goal of our analysis is to model the equilibrium effects of TWAP and other intraday trading benchmarks on prices and trading. Thus, the necessary model ingredients are investors who have heterogeneous trading targets and penalties for deviating from their intraday targets and who are price-sensitive so that market-clearing can be endogenized. Moreover, given the high empirical frequency of child orders in dynamic order-execution strategies in Korajczyk and Murphy [29] and van Kervel and Menkveld [43], modeling these effects in continuous time is appropriate.

We develop a continuous-time equilibrium model with a unit time horizon in which trade takes place at each time point $t \in [0, 1]$. This can be interpreted as one trading day. There are two securities: A money market account with a constant unit price (i.e., the account pays a zero interest rate) and a stock with an endogenously determined price process $S = (S_t)_{t\in[0,1]}$. Information about future dividends is generated by an exogenous standard Brownian motion $D = (D_t)_{t\in[0,1]}$ with a given known initial value $D_0 \in \mathbb{R}$, zero drift, and volatility normalized to one. Here D_t denotes the expected present value at time $t \leq 1$ of subsequent future dividends, where D_1 is the expectation at the end of the day. The market microstructure literature calls D_t the "fundamental asset value." In our model, the difference $S_t - D_t$ is price pressure required to clear the stock market.

Investors in our model all follow optimal trading strategies that trade off trading profits/costs and different types of inventory-holding penalties. For a generic investor i, let $\theta_{i,t}$ denote the investor's actual stock holdings at time $t \in [0, 1]$. By subtracting any initial stock position, we can, without loss of generality, normalize initial stock positions $\theta_{i,-}$ to zero. Let $\theta^c_{i,-}$ denote the trader's initial cash money-market balance. Formally, the optimal stock holdings for each investor i over the day solve an optimization problem

$$V_i(X_{i,0}) := \sup_{\theta_i\in\mathcal{A}_i} \mathbb{E}\Big[X_{i,1} - L_{i,1}\,\Big|\,\mathcal{F}_{i,0}\Big], \tag{1.1}$$

where $X_{i,1}$ is the investor's terminal wealth at $t = 1$ given her trading gains or losses over the day given by the wealth process

$$X_{i,t} := \theta^c_{i,-} + \int_0^t \theta_{i,u}\, dS_u, \quad t \in [0, 1], \tag{1.2}$$

where $X_{i,0} := \theta^c_{i,-}$ is investor i's initial wealth (i.e., initial cash balance in her money market account) and where $L_{i,1}$ in (1.1) is an investor-specific terminal cumulative inventory penalty at $t = 1$ that accrues over the day. The initial information set $\mathcal{F}_{i,0}$ in (1.1) is formally defined below in (1.12), and the admissible set of controls $\mathcal{A}_i$ is defined below in Definition 2.1.

Two types of investors trade in our model with different holding penalties:

- There are $M \in \mathbb{N}$ price-sensitive *targeted investors* whose terminal penalty $L_{i,1}$ accrues via a penalty process

$$L_{i,t} := \int_0^t \kappa(s)\big(\theta_{i,s} - \gamma(s)\tilde{a}_i\big)^2 ds, \quad t \in [0, 1], \quad i \in \{1, \dots, M\}, \tag{1.3}$$

  where $\tilde{a}_i$ is investor i's end-of-day target stock holdings and $\gamma(t)\tilde{a}_i$ is a trajectory of intraday target stock holdings. The function $\kappa(t)$ describes the severity of penalties for deviations of actual holdings $\theta_{i,t}$ over time from the target trajectory $\gamma(t)\tilde{a}_i$. The variables $(\tilde{a}_i, \theta^c_{i,-})$ are private knowledge of trader i but are partially revealed in equilibrium to other investors. At this point, we make no distributional assumptions about $(\tilde{a}_1, \dots, \tilde{a}_M)$ except that they are independent of the dividend-information Brownian motion D.[6]

We differentiate between two types of targeted investors based on their trading targets $\tilde{a}_i$. We refer to traders with targets $\tilde{a}_i \neq 0$ as *rebalancers*. These are large asset managers who use dynamic order-execution algorithms to trade on their targets. Traders with $\tilde{a}_i = 0$ do not need to trade per se but can provide liquidity, and so we call these *intraday liquidity providers*.

The target ratio $\gamma(t)$ is a deterministic function that gives how much of her daily target $\tilde{a}_i$ investor i would ideally hold at each time t during the day. The trajectory $\gamma(t)\tilde{a}_i$ can be thought of as a time-weighted average price (TWAP) strategy or as, more generally, a variation on an Almgren-Chriss trading strategy.[7] Realistically, the target ratio $\gamma(t)$ should be non-decreasing over the day with $\gamma(1) = 1$ at the end of the day. For example, for a standard TWAP strategy, the target ratio grows linearly with time as $\gamma(t) := t$. Alternatively, for a more generalized strategy, $\gamma(t)$ might follow the shape of the expected cumulative volume curve over the trading day. The assumption that $\gamma(t)$ is deterministic simplifies the analysis. Later, Section 5 extends the model to allow for stochastic target ratios γ_t that are then related to VWAP trading in Section 6.

- There are $\overline{M} \in \mathbb{N}$ price-sensitive *realtime hedgers* indexed by $i \in \{M+1, \dots, M+\overline{M}\}$. These hedgers have no private information. Their cumulative penalty $L_{i,1}$ in (1.1) accrues via the process

$$L_{i,t} := \int_0^t \overline{\kappa}(s)\big(\theta_{i,s} - \epsilon B_s\big)^2 ds, \quad t \in [0,1], \quad i \in \{M+1, \dots, M+\overline{M}\}, \tag{1.4}$$

which is similar to the common-value specification in Sannikov and Skrzypacz [39]. In (1.4), the process $B = (B_t)_{t\in[0,1]}$ with $B_0 = 0$ is a standard Brownian motion representing a publicly observable risk-factor that is independent of all other random variables, and $\epsilon \geq 0$ scales the hedger target holdings relative to the risk-factor B_t.[8] The hedgers' penalty-severity function $\overline{\kappa}(s)$ is potentially different from $\kappa(s)$ in (1.3). In the special case $\overline{\kappa}(s) := 0$, the hedgers reduce to risk-neutral *Merton investors* with potential price impact. Another special case is $\epsilon := 0$ in which the hedgers coincide with the intraday liquidity providers discussed above but with a possibly different penalty-severity function $\overline{\kappa}$.

Trading by the hedgers injects intraday randomness (due to B_t) in the aggregate trading-demand imbalances over the day. However, unlike noise traders, the hedgers are price-sensitive and optimize their trades.

Investors optimize their trading in actual markets given a variety of considerations. These include both the underlying utility gains from changing their holdings and also bid-ask and other order-execution costs (if taking liquidity) or profits (if providing liquidity), inventory and risk-management costs, and predictable wealth effects due to trading with/against price

[6] We do not need to impose a Gaussian structure on the private information variables because our investors' optimal strategies do not involve filtering. This is because—as we shall see in (1.11) below—the equilibria we construct initially reveal a sufficient statistic for how investor rebalancing targets affect price dynamics.

[7] It is a "variation" in the sense that Almgren and Chriss [1,2] solve for optimal orders given specific risk and cost objectives, whereas we represent investor i's risk and execution-cost objectives implicitly via the $L_{i,t}$ penalty.

[8] For simplicity, we assume ϵ is constant. Alternatively, it could be a time-dependent function $\epsilon(t)$.

pressure. Our analysis decomposes trading optimization for targeted investors into two parts: First, we interpret the intraday trajectory $\gamma(t)\tilde{a}_i$ as a partially optimized strategy given the end-of-day latent target $\tilde{a}_i$ and given order-execution, inventory, and risk-management costs but omitting wealth effects due to price pressure.[9] Second, the targeted investor's optimization problem in (1.1) then adjusts the partially optimized strategy $\gamma(t)a_i$ for price-pressure effects to obtain the fully optimized strategy $\theta_{i,t}$. In particular, (1.1) includes the expected gain/loss from how the holdings $\theta_{i,t}$ comove with/against price pressure and also an expected penalty, given $L_{i,t}$ in (1.3), which is a reduced-form representation of incremental increased order-execution, inventory, and risk-management costs given the deviation of $\theta_{i,t}$ from the partially optimized trajectory $\gamma(t)\tilde{a}_i$ over the day.[10] An analogous logic justifies the hedger trading optimization (1.1) with the penalty in (1.4) except that now the partially optimized hedger strategy—given order-execution, risk-management, and inventory costs and inefficiency costs from incomplete hedging but not wealth effects from price pressure—is the stochastic process ϵB_t.

Given this motivation, there are several specific points to note: First, the decomposition leading to (1.1) makes our model mathematically tractable, because the quantities $\gamma(t)$, ϵ, $\kappa(t)$, and $\overline{\kappa}(t)$ describing investor trading preferences are inputs in our analysis rather than something to be solved for formally.

Second, our analysis makes a distinction between price pressure $S_t - D_t$, which is explicitly modeled, and order-execution costs—like bid-ask spread add-ons, brokerage fees, and exchange fees—that are implicitly modeled as one of the determinants of the penalties $L_{i,t}$ in (1.3) and (1.4).[11] In particular, the target ratio function $\gamma(t)$, hedging scalar ϵ, and penalty-severity functions $\kappa(t)$ and $\overline{\kappa}(t)$ all implicitly depend on order-execution costs as well as on inventory and other costs. Intuitively, variation in the target ratio $\gamma(t)$ and penalty severity $\kappa(t)$ over the day reflects, in part, time patterns in order-execution costs (e.g., U-shaped average bid-ask spreads, see McInish and Wood [33], and price impacts, see Barardehi and Bernhardt [4]) and an intensifying inventory-holding preference to reach the target $\tilde{a}_i$ towards the end of the day. Our analysis allows, in particular, for penalty-severity functions $\kappa(t)$ that explode towards the end of the trading day as $t \uparrow 1$ as well as for bounded penalty severities.

Third, the targeted investors and hedgers are all price-sensitive, and so they adjust their stock holdings $\theta_{i,t}$ relative to their intraday trading targets in response to premia and discounts in prices in order to clear the market over time given the aggregate latent trading-demand imbalances $\gamma(t)\tilde{a}_\Sigma + \bar{M}\epsilon B_t$. In particular, their intraday targets are soft rather than hard constraints. Thus, there are two competing drivers in the objective (1.1). On one hand, neglecting the penalty term $L_{i,1}$ in (1.1), investor i would maximize her expected trading profit from liquidity provision. On the other hand, neglecting the wealth term $X_{i,1}$ in (1.1) means investor i would minimize the penalty $L_{i,1}$. In this case, targeted investors $i \in \{1, \dots, M\}$ would use the strategy $\gamma(t)\tilde{a}_i$, and the hedgers $i \in \{M+1, \dots, M+\overline{M}\}$ would hold ϵB_t. The equilibrium strategies $\hat{\theta}_{i,t}$ in (3.8) below strike an optimal balance between these two competing

[9] In other words, $\gamma(t)\tilde{a}_i$ is the ideal trading trajectory if investor i could trade at prices D_t with just order-execution, risk, and inventory costs but without the price impact of order-flow imbalances.

[10] To keep the model parsimonious, we assume rebalancers and liquidity providers have the same penalty-severity function $\kappa(t)$ (where $\gamma(t)$ does not matter for liquidity providers for whom $\tilde{a}_i = 0$). In a richer model, the liquidity-provider penalty severity might differ due to additional market-making considerations such as funding costs, risk-aversion, moral hazard costs due to risk limits arising from in-firm principal-agent conflicts, and the fact that liquidity providers are more likely than rebalancers to earn than to pay the bid-ask spread.

[11] The implicit order-execution costs are add-ons on top of S_t for individual transactions of individual investors, whereas $S_t - D_t$ is market-wide price pressure in all transactions for all investors.

forces where the penalty-severity functions $\kappa(t)$ and $\overline{\kappa}(t)$ determine the relative importance of these two forces over the trading day [0, 1].

Fourth, an important point about the targeted investors is that they care, not just about their terminal trading targets $\tilde{a}_i$, but about the entire path of their holdings over the day relative to their intraday target trajectories $\gamma(t)\tilde{a}_i$. In other words, the rebalancers' latent stock demand over the day given $\tilde{a}_i \neq 0$ has a time-varying trend controlled by $\gamma(t)$. This aspect of our model—which allows us to study TWAP and other forms of targeted trading—is new relative to the previous literature, which has modeled constant zero targets Du and Zhu [14], driftless stochastic targets Sannikov and Skrzypacz [39], Vayanos [44], and fixed terminal targets but no intraday targets Brunnermeir and Pedersen [9].

Market clearing in the model takes the following form: Recall that all initial stock positions $\theta_{i,-}$ have been normalized to zero. Given an aggregate stock supply normalized to zero, market clearing requires equilibrium investor-holdings $\hat{\theta}_{i,t}$ over the day to satisfy[12]

$$\sum_{i=1}^{M+\overline{M}} \hat{\theta}_{i,t} = 0 \quad \text{for all times } t \in [0, 1]. \tag{1.5}$$

The money market is also normalized to be in zero supply. By Walras law, clearing in the stock market implies clearing in the money market because investors use self-financing strategies. In particular, we assume the initial money-market endowments $\theta^c_{1,-}, \dots, \theta^c_{M+\overline{M},-}$ clear the money market

$$\sum_{i=1}^{M+\overline{M}} \theta^c_{i,-} = 0. \tag{1.6}$$

The analysis below distinguishes between conceptually different stock-price processes. Let $S = (S_t)_{t\in[0,1]}$ denote a generic stock-price process defined in terms of its dynamics

$$dS_t := \mu_t dt + dN_t, \quad S_0 := \hat{S}_0, \tag{1.7}$$

where $\mu = (\mu_t)_{t\in[0,1]}$ is a generic drift process, $N = (N_t)_{t\in[0,1]}$ is a fixed martingale, and $\hat{S}_0$ is a fixed initial stock price. The equilibrium price process we seek to determine is denoted by $\hat{S} = (\hat{S}_t)_{t\in[0,1]}$ and has two boundary conditions at times $t \in \{0, 1\}$:

First, the equilibrium stock price $\hat{S}_1$ at the end of the trading day at $t = 1$ is pinned down by a reduced-form end-of-day requirement[13]

$$\hat{S}_1 = D_1 + \varphi_0 \epsilon B_1 + \varphi_1 \tilde{a}_\Sigma \tag{1.8}$$

where $\varphi_0, \varphi_1 \in \mathbb{R}$ are exogenous constants and the total target imbalance is denoted by

$$\tilde{a}_\Sigma := \sum_{i=1}^{M} \tilde{a}_i. \tag{1.9}$$

[12] It is possible to extend our model to include noise-trader orders such that the floating stock supply becomes an exogenous stochastic process $a(t) + b(t)Z_t + c(t)B_t$ where a, b, and c are deterministic functions of time $t \in [0, 1]$, B_t is the risk-factor Brownian motion in (1.4), and Z_t is a Brownian motion independent of all other random variables.

[13] If the terminal restriction (1.8) is eliminated, our model becomes simpler because the stock volatility becomes a free parameter and can, for example, be set to be one. The fact that competitive Radner equilibrium models without dividends have free volatilities is well-known; see, e.g., Theorem 4.6.3 in Karatzas and Shreve [28].

When our model is applied to a short time horizon (e.g., a trading day), the end-of-day price $\hat{S}_1$ is assumed to be the overnight valuation that clears the stock market given the latent trading-demand imbalances due to the aggregate rebalancer and hedger targets. In particular, market-clearing at $t = 1$ at the end of the day is assumed to reflect overnight stock-holding by the rebalancers, market-makers, and hedgers and also possibly additional net stock demand from overnight liquidity providers who do not trade during the day. If $\varphi_1 > 0$, then a positive aggregate latent target $\tilde{a}_\Sigma$ pushes up the end-of-day price $\hat{S}_1$ relative to D_1 in order for markets to clear. Similarly, if $\varphi_0 > 0$, then a positive hedging target ϵB_1 also raises $\hat{S}_1$ relative to D_1. Both effects are qualitatively natural for how latent trading-demand imbalances might affect end-of-day prices. A special case of (1.8) is

$$\hat{S}_1 = D_1. \tag{1.10}$$

This case applies if D_1 is a liquidating dividend paid at time $t = 1$ (as in, e.g., Grossman and Stiglitz [25], Kyle [30]) or, alternatively, if there are no overnight liquidity effects in pricing at $t = 1$. For $(\hat{S}_t)_{t\in[0,1]}$ to satisfy (1.8) or (1.10), the price dynamics in (1.7) must be restricted as time approaches maturity (i.e., as $t \uparrow 1$). As we shall see in Theorem 3.2 below, the equilibrium stock-price process is linear in $(D_t, \tilde{a}_\Sigma, B_t)$ with time-varying deterministic coefficients. The terminal restriction (1.8) gives boundary conditions for these time-varying coefficient functions. This allows us to derive endogenous intraday price and investor-holding processes that are consistent with the assumed end-of-day price in (1.8).

Second, all investors act as price-takers at time $t = 0$. This means the initial price $S_0 = \hat{S}_0$ in (1.7) is unaffected by individual investor holdings $\theta_{i,0}$. This is for tractability. In addition, in equilibrium, the initial orders of investors cause the endogenous opening price $\hat{S}_0$ to adjust to clear the market given the underlying latent trading-demand imbalance due to the aggregate target $\tilde{a}_\Sigma$. From a modeling perspective, the initial stock price $\hat{S}_0$ at time $t = 0$ is required to reveal the aggregate target $\tilde{a}_\Sigma$ defined in (1.9) in the sense that

$$\sigma(\hat{S}_0) = \sigma(\tilde{a}_\Sigma). \tag{1.11}$$

The measurability requirement (1.11) allows investors to avoid filtering and thereby keeps the model tractable.

The information structure of our model is as follows: For tractability, all traders have homogeneous beliefs in the sense that they all believe the processes (D, B) are the same independent Brownian motions. At each time $t \in [0, 1]$ over the trading day, the dividend-information process D_t and the real-time hedging factor B_t are publicly observed. At each time $t \in [0, 1]$, each investor i chooses a stock-holding position $\theta_{i,t}$ adapted to the filtration[14]

$$\mathcal{F}_{i,t} := \begin{cases} \sigma(S_u, D_u, B_u, \tilde{a}_i, \theta^c_{i,-})_{u\in[0,t]}, & i = 1, \dots, M, \\ \sigma(S_u, D_u, B_u, \theta^c_{i,-})_{u\in[0,t]}, & i = M+1, \dots, M+\overline{M}. \end{cases} \tag{1.12}$$

While the risk factor B_t is observable, the requirement $S_0 = \hat{S}_0$ in (1.7) and the measurability property (1.11) ensure the total latent target $\tilde{a}_\Sigma$ can be inferred from S_0.

[14] As usual in continuous-time models, certain stochastic integrals need to be martingales; see Definition 2.1 below for details. Also, as usual, we have implicitly augmented (1.12) with null-sets to ensure that the "usual conditions" hold (see, e.g., Protter [38] for details).

2 Individual optimization problems

This section gives a precise description of the optimization problem in (1.1) for the targeted investors and hedgers. Our analysis proceeds in two steps: first, we describe technical properties of the set of admissible holding strategies $\mathcal{A}_i$ in (1.1). Second, we construct perceptions of off-equilibrium market-clearing prices for a generic investor i in (1.1) when using an arbitrary holding strategy $\theta_i \in \mathcal{A}_i$. Section 3 then derives equilibrium strategies $\hat{\theta}_{i,t}$ for investor i given these price beliefs.

Step 1 The set of admissible strategies $\mathcal{A}_i$ for investor i in (1.1) is defined as follows where a càglàd process is a left-continuous process with right limits:

Definition 2.1 For a given stock-price process S_t with drift μ_t and martingale N_t as in (1.7) with a predictable quadratic variation process $\langle N\rangle_t$, we define $\mathcal{F}_{i,t}$ by (1.12). A càglàd process $\theta_i = (\theta_{i,t})_{t\in[0,1]}$ adapted to $\mathcal{F}_{i,t}$ is *admissible*, and we write $\theta_i \in \mathcal{A}_i$, if the following integrability condition holds:

$$\mathbb{E}\left[\int_0^1 \left(|\theta_{i,t}\mu_t|dt + \theta_{i,t}^2 d\langle N\rangle_t\right)\Big|\,\mathcal{F}_{i,0}\right] < \infty, \quad i \in \{1,\ldots,M+\overline{M}\}. \tag{2.1}$$

□

It is well-known that an integrability condition like (2.1) rules out doubling strategies (see, e.g., Chapters 5 and 6 in [15]). The equilibrium stock-holding process in (3.8) below is not bounded (because it depends linearly on B_t) but does still satisfy (2.1). There are two reasons for the left-continuity requirement placed on $\theta_{i,t}$. First, left-continuity of the investor holding paths is sufficient to ensure that the state-process Y_t appearing in (2.2) below is left-continuous (ultimately, Y_t in (2.2) becomes the solution $Y_t^{\theta_i}$ to the stock-market clearing equation (2.5)). In turn, Y_t's left-continuity is sufficient for investors to infer Y_t from past and current stock-price observations (see Lemma A.1 in "Appendix A"). Second, in Example 5.3 in Section 5, which is an extension of our model to stochastic targets, the martingale N_t is only càdlàg (i.e., right-continuous with left limits). When N_t has such points of discontinuity, left-continuity of $\theta_{i,t}$ and the integrability condition (2.1) are sufficient to ensure martingality of the stochastic integral $\int_0^t \theta_{i,u}dN_u$, $t \in [0, 1]$ (see, e.g., p. 171 in Protter [38]).[15]

Step 2 This step constructs perceptions for a generic investor i about the market-clearing prices $S_t^{\theta_i}$ she faces given arbitrary holdings $\theta_{i,t}$. In equilibrium, price dynamics and price perceptions must agree. Off equilibrium, however, price perceptions must be reasonable. We describe two different cases of reasonable off-equilibrium price beliefs. One case is standard pricing-taking. This is the simplest version of our model. The second case has price-impact in the sense that investors are strategic and believe their individual holdings affect off-equilibrium market-clearing prices. In particular, although we assume investors are price-takers at the market open at $t = 0$, our second case allows for price impact of investor holdings $\theta_{i,t}$ over the rest of the day for $t \in (0, 1)$. In other words, our strategic traders actively take into account the price-impact their trades have when they solve for their individual optimal trading strategies.

Case of price-taking perceptions Optimal holdings for investors with price-taking beliefs are derived from (1.1) under the assumption that the off-equilibrium prices $S_t^{\theta_i}$ each investor i perceives herself as facing, the martingale N_t, and the perceived price drift $\mu_t^{\theta_i}$ given

15 Verification of optimality of conjectured optimizers in continuous-time stochastic control problems always involves proving certain stochastic integrals are martingales. This is illustrated in Duffie [15, Chapter 9C].

holdings $\theta_{i,t}$ for investor i are all unaffected by her arbitrary holdings $\theta_{i,t}$. The equilibrium construction for this case in Sections 3 and 4 then derives an equilibrium price process such that the market clears given investor optimal holdings given the equilibrium price process with price-taking beliefs.

Case of price-impact perceptions There are many ways to model price-impact. For example, stock prices are affine functions of investor orders in the discrete-time models in Vayanos [44, Eq. (4.2)], and Kyle [30, Eq. (3.12)]. A discrete-time version of our model also takes price changes ΔS_{t_n} as affine functions of the holdings θ_{i,t_n}. However, modeling off-equilibrium prices in continuous-time as functions of the levels of arbitrary holdings over time is technically difficult because, while $\theta_{i,t}$ can serve as an integrand, integration with respect to $d\theta_{i,t}$ might be ill-posed. Thus, in our continuous-time model of price impact, we assume that the perceived price drift $\mu_t^{\theta_i}$ is an affine function of the holdings $\theta_{i,t}$, but that the martingale N_t in (1.7) and the initial value $S_0 = \hat{S}_0$ are both independent of the holding process $\theta_{i,t}$. Thus, our model of price-impact is an affine drift specification in the general continuous-time price-impact setting in Cvitanić and Cuoco [12]. Given the perceived off-equilibrium price impacts, we derive optimal holding strategies for investors in our model.

Our eventual goal is to construct a Subgame Perfect Nash equilibrium, so the off-equilibrium prices associated with off-equilibrium holdings by a given investor i must be consistent with off-equilibrium beliefs for other investors such that, given the other investors' optimal off-equilibrium holdings given their beliefs, the market clears to produce the prices perceived by investor i. We denote the other investors who respond to off-equilibrium holdings by investor i using an index $j \neq i$. In this setting, we construct perceived market-clearing prices for investor i and the associated perceived responses $\theta_{j,t}$ of other investors $j \neq i$ to arbitrary holdings $\theta_{i,t}$ by investor i.

Our derivation of optimal investor holding strategies involves two different off-equilibrium stock-price processes in the filtrations in (1.12) and in the investor wealth dynamics in (1.2). First, let S_t^Y denote the off-equilibrium prices as perceived by a generic investor j with $j \neq i$ given investor j's own holdings $\theta_{j,t}$ and given a general state-process Y_t representing the perceived net effect of arbitrary holdings by other investors.[16] Second, let $S_t^{\theta_i}$ denote off-equilibrium prices perceived by investor i given investor i's holdings $\theta_{i,t}$ and given the state-process $Y_t^{\theta_i}$ induced by the market-clearing responses of investors $j \neq i$ to i's holdings $\theta_{i,t}$. The difference between these two perceived price processes is that S_t^Y is for an arbitrary exogenous state-process Y_t whereas $S_t^{\theta_i}$ is for an endogenous state-process $Y_t^{\theta_i}$ given the effect of $\theta_{i,t}$ on market-clearing. The prices S_t^Y and $S_t^{\theta_i}$ are linked with each other and with the equilibrium prices $\hat{S}_t$. First, the initial stock prices S_0^Y and $S_0^{\theta_i}$ are always pinned down to be $\hat{S}_0$ in (1.7). Thus, all investors act like price takers at time $t = 0$. However, the perceived off-equilibrium terminal prices $S_1^{\theta_i}$ and S_1^Y are not required to satisfy the equilibrium terminal condition (1.8). This is permissible because these are perceived off-equilibrium prices that can differ from equilibrium prices. Second, the equilibrium holdings $\hat{\theta}_{i,t}$ for all traders $i \in \{1, .., M+\overline{M}\}$ must produce the same equilibrium stock price $S_t^{\hat{\theta}_i} = \hat{S}_t$, which in turn must satisfy condition (1.8) at time $t = 1$.

The aspect of price beliefs that matters for modeling the impact of off-equilibrium holdings on market-clearing is how $\theta_{i,t}$ for an investor i at time t affects the market-clearing stock-price drift μ_t and, thus, the investment attractiveness of holding stock at time t. Thus, we start this step with an assumption that the perceived general price process $S^Y = (S_t^Y)_{t\in[0,1]}$

[16] The filtration $\mathcal{F}_{i,t}$ in (1.12) also depends on S and thus $\mathcal{F}_{i,t}$ varies depending on which stock-price process $S \in \{\hat{S}, S^{\theta_i}, S^Y\}$ we consider in the various steps of the equilibrium construction.

for other investors $j \neq i$ in trader j's optimization problem (1.1) has the form in (1.7) where the perceived stock-price drift $\mu = \mu^Y$ is defined as[17]

$$\mu_t^Y := \begin{cases} \nu_0(t)Y_t + \nu_1(t)\gamma(t)\tilde{a}_\Sigma + \nu_2(t)B_t + \nu_3(t)\theta_{j,t} + \nu_4(t)\gamma(t)\tilde{a}_j, & j \in \{1,\dots,M\}, \\ \bar{\nu}_0(t)Y_t + \bar{\nu}_1(t)\gamma(t)\tilde{a}_\Sigma + \bar{\nu}_2(t)B_t + \bar{\nu}_3(t)\theta_{j,t}, & j \in \{1+M,\dots,M+\overline{M}\}, \end{cases} \tag{2.2}$$

where $Y = (Y_t)_{t\in[0,1]}$ is an exogenous state-process with càglàd paths and $\nu_0(t),\dots,\nu_4(t)$ and $\bar{\nu}_0(t),\dots,\bar{\nu}_3(t)$ are continuous deterministic functions with $\nu_0(t) \neq 0$, $\bar{\nu}_0(t) \neq 0$ such that

$$\nu_3(t) < \kappa(t) \text{ when } M \geq 1 \text{ and } \bar{\nu}_3(t) < \overline{\kappa}(t) \text{ when } \overline{M} \geq 1, \quad t \in [0,1). \tag{2.3}$$

The martingale N in the price process (1.7) for S^Y has dynamics

$$dN_t := dD_t + \zeta(t)\epsilon dB_t, \quad N_0 := 0, \tag{2.4}$$

where $\zeta(t)$ is a continuous deterministic function of time for $t \in [0,1]$.

The coefficients $\nu_0(t),\dots,\bar{\nu}_3(t)$ and $\zeta(t)$ in (2.2) and (2.4) describe the subjective perceptions (i.e., off-equilibrium beliefs) of investor j about the price process she faces. We assume all investors perceive the same pricing coefficients. The coefficients $\nu_3(t)$ and $\bar{\nu}_3(t)$ represent, specifically, an investor's perception of how her own holdings $\theta_{j,t}$ directly affect prices off-equilibrium. In the discussion here, the perceived price coefficients are taken as given. However, they are endogenized in the equilibrium construction in Sect. 3.

Because the state-process Y_t is general (e.g., Y_t need not be Markovian), it is crucial for tractability that investors have linear utilities as in (1.1). With linear preferences, investors' optimal holding decisions at each time $t \in [0,1]$ only depend on the perceived price drift and the associated holding penalties accruing at that time t. Thus, investors optimize pointwise given the impact their holdings have on the perceived price drift.[18,19] These arguments lead to Lemma A.1 in "Appendix A", which gives the optimal response holdings $\theta_{j,t}^Y$ in (A.1) for targeted investors and hedgers given a perceived price process for investor j of the form described in (2.2) through (2.4).[20] The restrictions in (2.3) ensure the second-order condition for optimality of (A.1) is satisfied (without the second-order condition there is no optimizer).

The optimal-response holdings $\theta_{j,t}^Y$ in Lemma A.1 in "Appendix A" are for an arbitrary left-continuous state-process Y_t. However, to construct perceived prices for investor i that also clear the stock market, the price perceptions of investor i must take into account the fact

17 The equilibrium construction in Sect. 3 explicitly solves for perceived prices with this conjectured form that are consistent with equilibrium.

18 More specifically, investors maximize the expectation of the integrand of a Riemann integral of a quadratic function of $\theta_{i,t}$ at each time $t \in [0,1]$. See Eqs. (A.2) and (A.3) in "Appendix A".

19 An optimal trading strategy is typically computed as the solution to a Hamilton-Bellman-Jacobi equation that takes intertemporal trade-offs into account. However, in our model the intraday target trajectory $\gamma(t)\tilde{a}_i$ and hedge positions ϵB_t are interpreted as the solutions to the targeted investors' and hedgers' intertemporal partial optimization problems with order-execution, risk, and inventory costs but when price pressure from aggregate order-flow imbalances is ignored. As a result, the full optimization problem in (1.1) involves a series of separable trade-offs between incremental penalties for deviations of $\theta_{i,t}$ from the target trajectory and hedge positions at each time $t \in [0,1]$ and the contemporaneous price drift (expected capital gain) at time t when investors adjust their position to trade on price pressure.

20 In Lemma A.1 in "Appendix A" with price-impact, the state-process Y_t must be inferable, which is ensured by Y_t having left-continuous paths, which follows from investor holdings having left-continuous paths as required by Definition 2.1.

that her holdings $\theta_{i,t}$ also affect the state-process Y_t perceived by other investors $j \neq i$ via the market-clearing condition:

$$\theta_{i,t} + \sum_{j\neq i, j=1}^{M+\overline{M}} \theta_{j,t}^{Y} = 0, \quad t \in [0, 1], \tag{2.5}$$

where $\theta_{j,t}^{Y}$ denotes optimal responses of investors $j \neq i$ in (A.1) in Lemma A.1 in "Appendix A" to a state-process Y_t. Thus, given arbitrary holdings $\theta_{i,t}$, we can then solve (2.5) for the associated state-process, which we denote by $Y_t^{\theta_i}$. The state-process $Y_t^{\theta_i}$ ensures that the off-equilibrium optimal-response holdings for traders $j \in \{1, \dots, M + \overline{M}\}\backslash\{i\}$ clear the stock market as trader i varies her off-equilibrium holdings $\theta_{i,t}$. The perceived off-equilibrium market-clearing stock-price process $S_t^{\theta_i}$ associated with $\theta_{i,t}$ has a drift process $\mu_t^{\theta_i}$ given by (2.2) with $Y_t := Y_t^{\theta_i}$. However, the initial price $S_0^{\theta_i} := \hat{S}_0$ and martingale N_t in $dS_t^{\theta_i}$ in (1.7) do not depend on $\theta_{i,t}$.

Lemma 2.2 *Assume that $\nu_0(t) \neq 0$, $\bar{\nu}_0(t) \neq 0$, $t \in [0, 1]$, and that (2.3) holds. The off-equilibrium stock-price process perceived by investor $i \in \{1, \dots, M + \overline{M}\}$ for arbitrary holdings $\theta_i \in \mathcal{A}_i$ is defined by*

$$S_t^{\theta_i} := \hat{S}_0 + \int_0^t \mu_u^{\theta_i} du + N_t, \quad t \in [0, 1], \tag{2.6}$$

where the stock-price drift is defined by

$$\mu_t^{\theta_i} := \begin{cases} \frac{2\kappa\nu_0(\bar{\nu}_3-\overline{\kappa})+(\bar{\nu}_3-\overline{\kappa})\nu_0\nu_4+\overline{M}(\kappa-\nu_3)(\nu_1\bar{\nu}_0-\nu_0\bar{\nu}_1)}{(M-1)\nu_0(\overline{\kappa}-\bar{\nu}_3)+\overline{M}\bar{\nu}_0(\kappa-\nu_3)}\gamma\tilde{a}_\Sigma - \frac{\overline{M}(\kappa-\nu_3)(\nu_0\bar{\nu}_2-\nu_2\bar{\nu}_0+2\overline{\kappa}\nu_0\epsilon)}{(M-1)\nu_0(\overline{\kappa}-\bar{\nu}_3)+\overline{M}\bar{\nu}_0(\kappa-\nu_3)}B_t \\ \quad + \frac{\kappa(-2\overline{\kappa}\nu_0+2\nu_0\bar{\nu}_3+\nu_3\overline{M}\bar{\nu}_0)-\nu_3((M+1)\nu_0(\bar{\nu}_3-\overline{\kappa})+\nu_3\overline{M}\bar{\nu}_0)}{(M-1)\nu_0(\overline{\kappa}-\bar{\nu}_3)+\overline{M}\bar{\nu}_0(\kappa-\nu_3)}\theta_{i,t} \\ \quad + \frac{2\kappa\nu_0(\overline{\kappa}-\bar{\nu}_3)+M\nu_0\nu_4(\overline{\kappa}-\bar{\nu}_3)+\nu_4\overline{M}\bar{\nu}_0(\kappa-\nu_3)}{(M-1)\nu_0(\overline{\kappa}-\bar{\nu}_3)+\overline{M}\bar{\nu}_0(\kappa-\nu_3)}\gamma\tilde{a}_i, \quad i \in \{1, \dots, M\}, \\ -\frac{(\overline{\kappa}-\bar{\nu}_3)(\bar{\nu}_0(2\kappa+M\nu_1+\nu_4)-M\nu_0\bar{\nu}_1)}{M\nu_0(\overline{\kappa}-\bar{\nu}_3)+(\overline{M}-1)\bar{\nu}_0(\kappa-\nu_3)}\gamma\tilde{a}_\Sigma - \frac{M(\overline{\kappa}-\bar{\nu}_3)(\nu_2\bar{\nu}_0-\nu_0\bar{\nu}_2)+2\overline{\kappa}(\overline{M}-1)\bar{\nu}_0\epsilon(\kappa-\nu_3)}{M\nu_0(\overline{\kappa}-\bar{\nu}_3)+(\overline{M}-1)\bar{\nu}_0(\kappa-\nu_3)}B_t \\ \quad + \frac{2\overline{\kappa}\bar{\nu}_0(\nu_3-\kappa)+\overline{\kappa}M\nu_0\bar{\nu}_3-M\nu_0\bar{\nu}_3^2+(\overline{M}+1)\bar{\nu}_0\bar{\nu}_3(\kappa-\nu_3)}{M\nu_0(\overline{\kappa}-\bar{\nu}_3)+(\overline{M}-1)\bar{\nu}_0(\kappa-\nu_3)}\theta_{i,t}, \quad i \in \{M+1, \dots, M+\overline{M}\}, \end{cases} \tag{2.7}$$

and where the martingale N_t is as in (2.4), and given an initial stock price $\hat{S}_0$ that satisfies (1.11). The price process (2.6) clears the stock market in the sense that (2.5) holds.

We note three consequences of Lemma 2.2: First, because the state-process $Y_t^{\theta_i}$ in (A.5) in "Appendix A" is affine in $\theta_i \in \mathcal{A}_i$, we see that $Y_t^{\theta_i}$ has left-continuous paths because $\theta_{i,t}$ has left-continuous paths (see Definition 2.1). Furthermore, from (A.1) in Lemma A.1 in "Appendix A", investor j's optimal response $\theta_{j,t}^{Y^{\theta_i}}$ for $j \neq i$ is also affine in trader i's off-equilibrium holdings $\theta_{i,t}$. Second, because the price drift (2.7) is affine in $\theta_{i,t}$, the corresponding optimization problem (1.1) is a quadratic problem when the stock price is defined as in (2.6). These two properties are used to derive the equilibrium holdings $\hat{\theta}_i \in \mathcal{A}_i$ in Theorem 3.2 in Section 3. Third, the case of price-taking perceptions is a special case of the price-impact case in which the coefficients multiplying $\theta_{i,t}$ in (2.7) are zero. Our equilibrium construction in Section 3 is for the general price-impact case. Section 4 then derives conditions on the price-perception coefficients in (2.2) for equilibrium with price-taking perceptions.

3 Equilibria with deterministic targets

Having described the individual investor optimization problems in Sect. 2, this section formally defines and constructs an equilibrium.

Definition 3.1 (*Equilibrium*) A *Subgame Perfect Nash equilibrium* consists of perceived price coefficients given by deterministic functions $\nu_0(t), \dots, \nu_4(t)$, and $\bar{\nu}_0(t), \dots, \bar{\nu}_3(t)$ in (2.2) with $\nu_0(t) \neq 0$ and $\bar{\nu}_0(t) \neq 0$, an initial stock price $\hat{S}_0$ satisfying (1.11), and a martingale $N = (N_t)_{t\in[0,1]}$ such that, given the perceived stock-price process (2.6), the resulting optimal stock-holding processes $\hat{\theta}_{1,t}, \dots, \hat{\theta}_{M+\overline{M},t}$ from (1.1) satisfy the following conditions:

(i) The optimal holdings $\hat{\theta}_i \in \mathcal{A}_i, i \in \{1, \dots, M+\overline{M}\}$, satisfy the intraday market-clearing condition (1.5).

(ii) When $\theta_{i,t}$ is set to the optimizer $\hat{\theta}_{i,t}$ in (2.7), the resulting stock-price drifts $\mu_t^{\hat{\theta}_i}$ are the same $\hat{\mu}_t$ for all investors $i \in \{1, \dots, M + \overline{M}\}$. The corresponding equilibrium stock-price process from (2.6) with $\theta_{i,t} = \hat{\theta}_{i,t}$ is denoted by $\hat{S}_t$.

(iii) The stock-price process $\hat{S}_t$ satisfies the terminal price condition (1.8) at time $t = 1$ for given constants $\varphi_0, \varphi_1 \in \mathbb{R}$.

□

Our equilibrium concept is stronger than Nash because it involves beliefs for each investor i about the perceptions of other investors $j \neq i$ that determine investor j's optimal responses to hypothetical off-equilibrium holdings by investor i.

Our main result is Theorem 3.2 below, which gives restrictions ensuring equilibrium existence (proof is in "Appendix A"). As we shall see, there are two degrees of freedom in the perceived price coefficients $\nu_0(t), \dots, \nu_4(t)$ and $\bar{\nu}_0(t), \dots, \bar{\nu}_3(t)$, and so there are multiple (indeed, infinitely many) equilibria. This situation is similar to Vayanos [44, Sec. 5] and Sannikov and Skrzypacz [39]. Keeping the price impact coefficients $\nu_3(t)$ and $\bar{\nu}_3(t)$ as the two free parameters simplifies the exposition. In Sect. 4 we use the mathematical flexibility of these free functions to consider equilibria with different amounts of competition and strategic behavior.

The equilibrium stock-price process $\hat{S}_t$ will be shown to have the form

$$\hat{S}_t := D_t + g(t)\tilde{a}_\Sigma + \zeta(t)\epsilon B_t, \tag{3.1}$$

for two continuously differentiable deterministic functions $g, \zeta : [0, 1] \to \mathbb{R}$. Consequently, the equilibrium stock-price drift $\hat{\mu}_t$ and martingale N_t in (1.7) are given by

$$\begin{aligned} \hat{\mu}_t &= g'(t)\tilde{a}_\Sigma + \zeta'(t)\epsilon B_t, \\ dN_t &= dD_t + \zeta(t)\epsilon dB_t, \quad N_0 = 0. \end{aligned} \tag{3.2}$$

We note three features about (3.1). First, equilibrium prices are expressed in (3.1) as functions of the underlying trading-demand variables $\tilde{a}_\Sigma$ and B_t. In particular, prices are functions of market-clearing investor holdings, which are functions of the underlying latent trading-demand variables $\tilde{a}_i$ (which aggregate to $\tilde{a}_\Sigma$) and B_t. The intuition is that equilibrium prices depend on the underlying latent total trading demand, which includes both trades that occur in equilibrium and also trading-demand imbalances that prices deter so that markets clear. Making the role of latent trading demand—and especially demand imbalances due to intraday TWAP trading targets—explicit is one of the contributions of our analysis. Second, given $\tilde{a}_\Sigma$, the equilibrium stock-price process is Gaussian. More specifically, the price process (3.1) is a Bachelier model with time-dependent coefficients.[21] Third, the individual investors' perceived price processes in (2.6) exhibit path dependency in the sense that the path of $(B_u)_{u\in[0,t]}$ is needed to determine $S_t^{\theta_i}$ for $t \in [0,1]$ for an arbitrary holding process $\theta_{i,t}$. However, only the current value of B_t affects the equilibrium prices $\hat{S}_t$ in (3.1).

Theorem 3.2 *Let $\gamma : [0,1] \to [0,\infty)$ be a continuous function, let $\nu_3, \bar{\nu}_3 : [0,1) \to \mathbb{R}$ be continuous functions, let $\kappa, \overline{\kappa} : [0,1) \to (0,\infty)$ be continuous and integrable functions; i.e.,*

$$\int_0^1 \big(\kappa(t) + \overline{\kappa}(t)\big)dt < \infty, \tag{3.3}$$

that satisfy the second-order conditions (2.3), and let there be at least $M + \overline{M} \geq 2$ investors. In addition, assume that (1.6) holds in the money market. Define the functions in (3.1) by

$$\begin{aligned} g(t) &:= \varphi_1 - \int_t^1 \mu_1(u)\gamma(u)du, \\ \zeta(t) &:= \varphi_0 - \int_t^1 \mu_2(u)du, \end{aligned} \tag{3.4}$$

for $t \in [0,1]$ where[22]

$$\begin{aligned} \mu_1 &:= -\frac{2\kappa\big(2(M+\overline{M})\overline{\kappa} - (1+M+\overline{M})\bar{\nu}_3\big)}{M\big(2(M+\overline{M})\overline{\kappa} - (1+M+\overline{M})\bar{\nu}_3\big) + \overline{M}\big(2(M+\overline{M})\kappa - (1+M+\overline{M})\nu_3\big)}, \\ \mu_2 &:= -\frac{2\overline{M}\overline{\kappa}\big(2(M+\overline{M})\kappa - (1+M+\overline{M})\nu_3\big)}{M\big(2(M+\overline{M})\overline{\kappa} - (1+M+\overline{M})\bar{\nu}_3\big) + \overline{M}\big(2(M+\overline{M})\kappa - (1+M+\overline{M})\nu_3\big)}. \end{aligned} \tag{3.5}$$

Provided that $g(0) \neq 0$, price-perception functions $\nu_0 \neq 0$, ν_1, ν_2, ν_4, and $\bar{\nu}_0$, $\bar{\nu}_1$, $\bar{\nu}_2$ satisfying

$$\bar{\nu}_0 = \frac{\nu_0(\overline{\kappa} - \bar{\nu}_3)}{\kappa - \nu_3}. \tag{3.6}$$

and

$$\begin{aligned} \nu_4 &= -\frac{2\kappa\Big(2\kappa + (M+\overline{M}-2)(1+M+\overline{M})\nu_3\Big)}{(M+\overline{M})\Big(2(M+\overline{M})\kappa - (1+M+\overline{M})\nu_3\Big)}, \\ \bar{\nu}_1 &= \frac{(\overline{\kappa}-\bar{\nu}_3)\nu_1}{\kappa-\nu_3} + \frac{(\overline{\kappa}-\bar{\nu}_3)(\nu_3-\bar{\nu}_3-(M+\overline{M})(2\kappa-2\overline{\kappa}-\nu_3+\bar{\nu}_3))(2\kappa+\nu_4)}{(\kappa-\nu_3)(M\overline{M}(2(\kappa+\overline{\kappa})-\nu_3)+\overline{M}(2\overline{M}\kappa-(1+\overline{M})\nu_3)+M^2(2\overline{\kappa}-\bar{\nu}_3)-M(1+\overline{M})\bar{\nu}_3)}, \\ \bar{\nu}_2 &= \frac{(\overline{\kappa}-\bar{\nu}_3)\nu_2}{\kappa-\nu_3} + \frac{2\epsilon\overline{\kappa}(-2\overline{M}(M+\overline{M})(\kappa-\overline{\kappa})-2\overline{\kappa}+\overline{M}(1+M+\overline{M})\nu_3+2\bar{\nu}_3-((M-1)M+3M\overline{M}+2\overline{M}^2)\bar{\nu}_3)}{M\overline{M}(2(\kappa+\overline{\kappa})-\nu_3)+\overline{M}(2\overline{M}\kappa-(1+\overline{M})\nu_3)+M^2(2\overline{\kappa}-\bar{\nu}_3)-M(1+\overline{M})\bar{\nu}_3}, \end{aligned} \tag{3.7}$$

together with N_t in (3.2) form an equilibrium in which:

(i) Investor equilibrium holdings are given by

$$\hat{\theta}_{i,t} = \begin{cases} \alpha_1(t)\gamma(t)\tilde{a}_\Sigma + \alpha_2(t)\epsilon B_t + \alpha_3(t)\gamma(t)\tilde{a}_i, & i \in \{1,\dots,M\}, \\ \bar{\alpha}_1(t)\gamma(t)\tilde{a}_\Sigma + \bar{\alpha}_2(t)\epsilon B_t, & i \in \{M+1,\dots,M+\overline{M}\}, \end{cases} \tag{3.8}$$

[21] While the equilibrium stock price can be negative with positive probability, such Gaussian models have been widely used in the market microstructure literature by, e.g., Grossman and Stiglitz [25] and Kyle [30]. Gaussian models are also widely used in the optimal order-execution literature including Almgren and Chriss [1]; see, e.g., the discussion in the Gatheral and Schied [23, Section 3.1] survey.

[22] For notational brevity, the time arguments for $\mu_1(t)$, $\mu_2(t)$, $\kappa(t)$, $\overline{\kappa}(t)$, $\nu_3(t)$, and $\bar{\nu}_3(t)$ are suppressed.

where

$$
\begin{aligned}
\alpha_1 &:= \frac{(M+\overline{M}-1)\mu_1}{2(M+\overline{M})\kappa-(1+M+\overline{M})\nu_3},\\
\alpha_2 &:= \frac{(M+\overline{M}-1)\mu_2}{2(M+\overline{M})\kappa-(1+M+\overline{M})\nu_3},\\
\alpha_3 &:= \frac{2\kappa(M+\overline{M}-1)}{2(M+\overline{M})\kappa-(1+M+\overline{M})\nu_3},\\
\bar{\alpha}_1 &:= \frac{(M+\overline{M}-1)\mu_1}{2(M+\overline{M})\overline{\kappa}-(1+M+\overline{M})\bar{\nu}_3},\\
\bar{\alpha}_2 &:= -\frac{M(M+\overline{M}-1)\mu_2}{\overline{M}\big(2(M+\overline{M})\kappa-(1+M+\overline{M})\nu_3\big)}.
\end{aligned}
\tag{3.9}
$$

(ii) The equilibrium stock price $\hat{S}_t$ is given by (3.1) with $g(t)$ and $\zeta(t)$ from (3.4). The associated equilibrium price drift in (3.2) is

$$
\hat{\mu}_t := \mu_1(t)\gamma(t)\tilde{a}_\Sigma + \mu_2(t)\epsilon B_t, \tag{3.10}
$$

with $\mu_1(t)$ and $\mu_2(t)$ from (3.5).

Remark 3.1 We note several properties of this equilibrium here:

1. From (3.1), the initial equilibrium stock price at $t=0$ is

$$
\begin{aligned}
\hat{S}_0 &= D_0 + g(0)\tilde{a}_\Sigma + \zeta(0)\epsilon B_0\\
&= D_0 + g(0)\tilde{a}_\Sigma,
\end{aligned}
\tag{3.11}
$$

where the second equality follows because $B_0 = 0$. Therefore, whenever $g(t)$ from (3.4) satisfies $g(0) \neq 0$, the aggregate target $\tilde{a}_\Sigma$ can be inferred from $\hat{S}_0$ and vice versa given that D_0 is public information. Thus, when $g(0) \neq 0$, we have $\sigma(\hat{S}_0) = \sigma(\tilde{a}_\Sigma)$ as required in (1.11). By (3.4), the condition $g(0) \neq 0$ is equivalent to

$$
\varphi_1 \neq \int_0^1 \mu_1(u)\gamma(u)du, \tag{3.12}
$$

where μ_1 is from (3.5). Consequently, for given functions $(\gamma, \kappa, \overline{\kappa}, \nu_3, \bar{\nu}_3)$ satisfying the second-order conditions (2.3), there is just one value of the terminal-price coefficient φ_1 for which $g(0) = 0$. For all other φ_1, we have $g(0) \neq 0$ and an equilibrium exists.
2. The equilibrium stock-price drift $\hat{\mu}_t$ and price levels $\hat{S}_t$ have the following qualitative properties: The second-order conditions (2.3) ensure

$$
\begin{aligned}
2(M+\overline{M})\kappa - (1+M+\overline{M})\nu_3 > \kappa > 0,\\
2(M+\overline{M})\overline{\kappa} - (1+M+\overline{M})\bar{\nu}_3 > \overline{\kappa} > 0,
\end{aligned}
\tag{3.13}
$$

so that the equilibrium price-drift coefficients in (3.5) can be signed with $\mu_1(t) < 0$ and $\mu_2(t) < 0$. This is intuitive. Larger latent aggregate trading-demand targets $\tilde{a}_\Sigma$ and larger hedging needs ϵB_t mean that the equilibrium price drift $\hat{\mu}_t$ must be lower in order to incentivize price-sensitive targeted investors and hedgers to adjust their holdings to clear the market. Given negative price-drift coefficients $\mu_1(t)$ and $\mu_2(t)$, it then follows from (3.4) that, given end-of-day terminal pricing coefficients $\varphi_0, \varphi_1 \geq 0$, the intraday price-level coefficients $g(t)$ and $\zeta(t)$ are both positive.

3. The investor holding coefficients $\alpha_1(t), \ldots, \bar{\alpha}_2(t)$ in (3.9) can also be signed using (3.5) and (3.13). The targeted-investor coefficients $\alpha_1(t)$ and $\alpha_2(t)$ on the aggregate imbalance state variables $\tilde{a}_\Sigma$ and ϵB_t are both negative. This is intuitive because investors reduce their personal holdings in response to the lower price drifts induced in equilibrium by positive latent aggregate imbalances. A similar intuition applies for the negative hedger coefficient $\bar{\alpha}_1(t)$ on the latent aggregate rebalancer imbalance $\tilde{a}_\Sigma$. The coefficient $\alpha_3(t)$ on the targeted investor's personal target $\tilde{a}_i$ is, as expected, positive. Similarly, the sign of the hedger's coefficient $\bar{\alpha}_2(t)$ on the hedging imbalance state variable is positive. This is the net effect of ϵB_t both as a personal target in the hedger penalty $L_{i,t}$ in (1.4) and as a state variable for the impact of the aggregate latent hedger imbalance on the price drift in the investor trading profits $X_{i,t}$ in (1.2). Recall that all initial stock positions have been normalized to zero for all traders (with no loss of generality). Therefore, from (3.8), there are initial discrete orders (i.e., block trades $\theta_{i,0} - \theta_{i,-} \neq 0$) at $t = 0$. This is related to why the initial price $\hat{S}_0$ fully reveals $\tilde{a}_\Sigma$. However, afterwards trading evolves continuously for $t \in (0, 1]$.[23]
4. Inserting $\bar{\nu}_0(t)$ from (3.6) into (2.7) and rearranging using (3.8) and (3.9) for the equilibrium holdings $\hat{\theta}_{i,t}$ and (3.5) and (3.10) for the equilibrium price drift $\hat{\mu}_t$ lets us express the perceived price drift for an investor i in (2.7) as

$$\mu_t^{\theta_i} = \begin{cases} \hat{\mu}_t + \frac{\nu_3(t)(M+\overline{M}+1)-2\kappa(t)}{M+\overline{M}-1}\big(\theta_{i,t} - \hat{\theta}_{i,t}\big), & i \in \{1, \ldots, M\}, \\ \hat{\mu}_t + \frac{\bar{\nu}_3(t)(M+\overline{M}+1)-2\bar{\kappa}(t)}{M+\overline{M}-1}\big(\theta_{i,t} - \hat{\theta}_{i,t}\big), & i \in \{M+1, \ldots, M+\overline{M}\}. \end{cases} \tag{3.14}$$

This is not surprising given the equilibrium requirement that each investor perceives the same price drift in equilibrium, i.e., $\mu_t^{\hat{\theta}_i} = \hat{\mu}_t$ for all $i \in \{1, \ldots, M + \overline{M}\}$. Because the equilibrium holdings $\hat{\theta}_{i,t}$ include an $\tilde{a}_i$ term, the representation (3.14) explains why there is an investor-specific $\tilde{a}_i$ term in the perceived price drift $\mu_t^{\theta_i}$ in (2.7).
5. The specific price-perception functions ν_0, ν_1, and ν_2 for the targeted investors are irrelevant in Theorem 3.2 in the sense that, given ν_3 and $\bar{\nu}_3$ and provided $\nu_0(t) \neq 0$, all different $\nu_0 \neq 0,\ \nu_1$, and ν_2 produce the same equilibrium prices and investor holdings. To provide some intuition, consider the drift in (2.2) for $j \in \{1, \ldots, M + \bar{M}\}$. This drift is overparameterized because both Y and θ_j are expressible in terms of $(\tilde{a}_j, \tilde{a}_\Sigma, B)$ when Y is replaced by the solution Y^{θ_j} of (2.5). We also note here that, while ν_0, ν_1, and ν_2 do not affect equilibrium prices and investor holdings in (3.1) [given (3.4) and (3.5)] and (3.8) [given (3.9)], they do pin down the other perceived price coefficients $\nu_4, \overline{\nu}_0, \overline{\nu}_1$, and $\overline{\nu}_2$ via (3.6) and (3.7), which are related to market-clearing and the common perceptions of equilibrium price conditions from Definition 3.1.

There is an important difference between prices in equilibrium and perceived off-equilibrium prices. Both in-equilibrium and off-equilibrium, the stock-price drifts $\hat{\mu}_t$ and $\mu_t^{\theta_i}$ are determined such that the market clears at each time t. In the one case, this is part of the definition of equilibrium, and, in the other case, it is a reasonable off-equilibrium belief. However, in equilibrium, the price level $\hat{S}_t$ adjusts at time t to be consistent with the required market-clearing drift $\hat{\mu}_t$ and the terminal price condition in (1.8). For example, a large target $\tilde{a}_\Sigma$ leads from (3.1) and (3.4) to a high opening price $\hat{S}_0$ at $t = 0$ so that the intraday drift $\hat{\mu}_t$ from (3.5) and (3.10) can be predictably low with prices drifting down in expectation over the day to the terminal price $\hat{S}_1$ in (1.8). Similarly, a positive random shock to B_t at time t leads from (3.1) and (3.4) to a random increase in prices $\hat{S}_t$ such that the market-clearing

[23] The discontinuous model in Example 5.3 below has optimal discrete orders throughout the trading day.

drift $\hat{\mu}_t$ can have a random decrease as required in (3.5) and (3.10). In contrast, for perceived off-equilibrium prices for investors with price-impact, we only specify how the perceived off-equilibrium price drifts $\mu_t^{\theta_i}$ in (2.7) change at each time t so that the market clears given an investor's holdings $\theta_{i,t}$. However, the perceived off-equilibrium price level $S_t^{\theta_i}$ at time t is not affected by $\theta_{i,t}$ at time t. This simplifies the specification of off-equilibrium beliefs and is still reasonable for price beliefs since market-clearing depends on the price drift at time t, not on the contemporaneous price level (as per the discussion about the off-equilibrium price-impact model before (2.7)).

Price pressure $\hat{S}_t - D_t$ is the effect of intraday imbalances in latent trading demand. It can be positive or negative depending on the aggregate target imbalance $\tilde{a}_\Sigma$ and the hedging need ϵB_t. From (3.1), price pressure has a deterministic trend $g(t)\tilde{a}_\Sigma$ over the day given the total rebalancing target imbalance $\tilde{a}_\Sigma$ and a stochastic component $\zeta(t)\epsilon B_t$ due to the randomly evolving hedging target. For example, $g(0)\tilde{a}_\Sigma$ is the initial price impact of the aggregate latent trading-target imbalance $\tilde{a}_\Sigma$ revealed by the opening order-flow at time $t = 0$. Thereafter, the price impact $g(t)\tilde{a}_\Sigma$ of the given aggregate imbalance $\tilde{a}_\Sigma$ varies predictably over time with $g(t)$. A similar phenomenon applies to price pressure due to hedging demand. In particular, $\zeta(t)\epsilon dB_t$ is the immediate impact of an innovation ϵdB_t in hedging demand at time t, and then $\zeta(s)\epsilon dB_t$ is the predictable continuation impact of ϵdB_t at times $s > t$ later in the day. From (3.4) and (3.5), we see that $g(t)$ (the $\tilde{a}_\Sigma$ coefficient in $\hat{S}_t$) is affine in $\gamma(t)$, and that $\zeta(t)\epsilon$ (the B_t coefficient in $\hat{S}_t$) is linear in ϵ (because $\zeta(t)$ does not depend on (γ, ϵ)). Thus, the path of intraday price pressure $\hat{S}_t - D_t$ at different times t during the day has an intertemporal factor-type structure where $\tilde{a}_\Sigma$ is a common factor (which is different on different days but fixed over the course of a given day) that causes price pressure to change deterministically over the day given intraday variation in $g(t)$, and B_t is a martingale that cause price pressure $\hat{S}_t - D_t$ to change randomly over the day. Thus, at a given time $s \in [0, 1)$, the future random price pressure at subsequent times $t \in (s, 1]$ given $\tilde{a}_\Sigma$ have the following conditional means and variances:

$$
\begin{aligned}
\mathbb{E}[\hat{S}_t - D_t | \sigma(\tilde{a}_\Sigma, B_u)_{u\in[0,s]}] &= g(t)\tilde{a}_\Sigma + \zeta(t)\epsilon B_s, \\
\mathbb{V}[\hat{S}_t - D_t | \sigma(\tilde{a}_\Sigma, B_u)_{u\in[0,s]}] &= \zeta(t)^2\epsilon^2(t - s).
\end{aligned}
\tag{3.15}
$$

As a result, positive (negative) latent aggregate trading targets $\tilde{a}_\Sigma$ lead to predictable positive (negative) price-pressure trends $\mathbb{E}[\hat{S}_t|\sigma(\tilde{a}_\Sigma)] - D_t$ over the day. In addition, randomness in the intraday hedging factor B_t produces randomness in equilibrium price pressure, where higher hedging factors B_t are associate with higher prices and lower price drifts. Again, this is intuitive.

Empirical predictions Security prices are often decomposed econometrically into an informational component that follows a martingale and a residual liquidity effect (as in, e.g., Hasbrouck [27]). A standard interpretation is that liquidity effects in prices decay predictably over time as liquidity supply from initial liquidity providers is first depleted by arriving order-flow imbalances and then replenished as order-flow imbalances are subsequently dispersed and absorbed by the broader market. Our model has two new empirical implications: First, our price pressure $\hat{S}_t - D_t$ is driven by both arriving orders and also by what those orders reveal about the underlying latent trading-demand imbalances $\tilde{a}_\Sigma$ and ϵB_t. Second, the source of intraday predictability in our price pressure differs from the standard liquidity-supply interpretation. Price predictability here reflects the net effect of predictable time variation in latent trading demand (i.e., the targets $\gamma(t)\tilde{a}_\Sigma$ and ϵB_t) as well as predictability in liquidity supply (controlled by $\kappa(t)$ and $\bar{\kappa}(t)$).

Empirical predictions from our model are testable using different types of data. First, the formula for equilibrium prices in (3.1) predicts that intraday prices are driven by three components: (1) time-varying intraday price pressure (controlled by $g(t)$) due to a daily aggregate parent target $\tilde{a}_\Sigma$ that is constant over the day but varies across days (depending on each day's different realized $\tilde{a}_\Sigma$), (2) additional intraday price pressure due to a Brownian motion (B_t) with heteroscedastic price effects (controlled by $\zeta(t)\epsilon$), and (3) a homoscedastic random walk reflecting fundamental information (controlled by the dividend state-process D_t). This representation imposes restrictions that can be estimated and tested using standard intraday price and order-flow data (e.g., TAQ). In particular, it is not necessary for the econometrician to know the realized daily imbalances $\tilde{a}_\Sigma$. Rather, a multi-day sample of intraday data can be viewed as a collection of panels of intraday data (i.e., one intraday path for each day in the sample) and then to identify the functions $g(t)$ and $\zeta(t)\epsilon$ using cross-day and intraday variance decompositions.

Second, the first testing strategy can be further refined given individual-investor order data (e.g., from IIROC), in which case it would be possible to measure the daily realizations of the aggregate large-investor imbalance $\tilde{a}_\Sigma$ for investors following dynamic strategies on different days.

Third, the price coefficients $g(t)$ and $\zeta(t)$ in (3.4) along with (3.5) give predictions for how the intraday effects (g and ζ) vary across different days given daily variation in the numbers of targeted rebalancers and hedgers (M and $\bar{M}$) and given daily variation in order-execution costs due to variation in bid-ask spreads (since order-execution costs implicitly affect κ and $\bar{\kappa}$). This relation can be estimated using standard data on bid-ask quotes (to identify days with lower or higher liquidity on which benchmark deviations are more or less costly) and given individual-trader data (e.g., from IIROC) to identify investor types and then to use that identification to measure the changing number of identified hedgers and rebalancers who actively trade on any given day.

4 Competitive equilibrium and welfare

From Theorem 3.2, the deterministic functions $\nu_3(t)$ and $\bar{\nu}_3(t)$ for perceived own-order price-impacts are two degrees of freedom in our model. By imposing additional economically-motivated structure on $\nu_3(t)$ and $\bar{\nu}_3(t)$, we can identify unique equilibria corresponding to different forms of competition and market power. From Section 2, price taking is a special case of interest.

In the competitive Radner equilibrium, investors act like price-takers over the whole day—not just at the market open—in that the perceived prices $S_t^{\theta_i}$ for each investor i are unaffected by her holdings $\theta_{i,t}$. Hence, the competitive Radner equilibrium is obtained by requiring that the coefficient in front of $\theta_{i,t}$ in the perceived drift (2.7) for $dS_t^{\theta_i}$ is zero. From (3.14), this requirement on $\nu_3(t)$ and $\bar{\nu}_3(t)$ is seen to imply

$$
\begin{aligned}
\frac{\nu_3(t)(M+\overline{M}+1)-2\kappa(t)}{M+\overline{M}-1} &= 0,\\
\frac{\bar{\nu}_3(t)(M+\overline{M}+1)-2\overline{\kappa}(t)}{M+\overline{M}-1} &= 0,
\end{aligned}
\tag{4.1}
$$

which gives the competitive-equilibrium perceived pricing coefficients[24]

$$\nu_3^*(t) = \frac{2\kappa(t)}{1+M+\overline{M}}, \quad \bar{\nu}_3^*(t) = \frac{2\overline{\kappa}(t)}{1+M+\overline{M}}. \tag{4.2}$$

We show next that the competitive Radner equilibrium is the welfare-maximizing equilibrium. While there are many ways to measure social welfare (see, e.g., Section 6.1 in Vayanos [44]), we follow Du and Zhu [14, Eq. 42] and consider maximizing the expected aggregate certainty-equivalent for the $M+\overline{M}$ investors. The certainty equivalent $\text{CE}_i \in \mathbb{R}$ for investor i is defined by

$$\text{CE}_i := V_i(X_{i,0}), \quad i \in \{1,\dots,M+\overline{M}\}, \tag{4.3}$$

where the value functions V_i are defined in (1.1). The aggregate expected social welfare objective is given by

$$\sup_{\nu_3(t),\bar{\nu}_3(t)} \sum_{i=1}^{M+\overline{M}} \mathbb{E}[\text{CE}_i], \tag{4.4}$$

subject to $(\nu_3, \bar{\nu}_3)$ satisfying requirements (2.3). The expectation in the objective (4.4) is ex ante in the sense that it is taken over random daily investor variables $(\tilde{a}_1,\dots,\tilde{a}_M)$ and $(\theta^c_{1,-},\dots,\theta^c_{M+\overline{M},-})$. The following theorem shows that the competitive Radner equilibrium (4.2) attains (4.4).

Theorem 4.1 *In the setting of Theorem 3.2, let the random private targets* $(\tilde{a}_1,\dots,\tilde{a}_M)$ *be square integrable and not perfectly correlated, and let* $M \geq 2$. *The competitive Radner equilibrium with (4.2) is the welfare-maximizing equilibrium in that the unique maximizers of (4.4) are given by* ν_3^* *and* $\bar{\nu}_3^*$ *in (4.2). The corresponding optimal holding strategies (3.8) are given by*

$$\hat{\theta}_{i,t} = \begin{cases} \gamma(t)\tilde{a}_i - \frac{\overline{\kappa}(t)\left(\gamma(t)\tilde{a}_\Sigma + \overline{M}\epsilon B_t\right)}{\overline{M}\kappa(t)+M\overline{\kappa}(t)}, & i=1,\dots,M, \\ \frac{-\kappa(t)\gamma(t)\tilde{a}_\Sigma + M\overline{\kappa}(t)\epsilon B_t}{\overline{M}\kappa+M\overline{\kappa}(t)}, & i = M+1,\dots,M+\overline{M}, \end{cases} \tag{4.5}$$

and the corresponding equilibrium stock-price drift coefficients in (3.10) are given by

$$\mu_1(t) := -\frac{2\kappa(t)\overline{\kappa}(t)}{\overline{M}\kappa(t)+M\overline{\kappa}(t)}, \quad \mu_2(t) := -\frac{2\kappa(t)\overline{\kappa}(t)\overline{M}\epsilon}{\overline{M}\kappa(t)+M\overline{\kappa}(t)}. \tag{4.6}$$

The proof (see “Appendix A”) uses the non-trivial correlations between $(\tilde{a}_1,\dots,\tilde{a}_M)$ to produce a strict inequality in Cauchy-Schwarz's inequality. This ensures that the second-order condition for optimality corresponding to (4.4) holds.

24 At first glance, it might seem that price-taking would mean price perceptions for investor i with $\nu_3(t)=0$ (for targeted investors) and $\bar{\nu}_3(t)=0$ (for hedgers). However, investor i's holdings $\theta_{i,t}$ have both a direct effect (ν_3 and $\bar{\nu}_3$) on perceived market-clearing prices in (2.6) and also an indirect effect via the optimal responses $\theta^{Y^{\theta_i}}_{j,t}$ of other investors $j \neq i$ to $\theta_{i,t}$ via the perceived market-clearing condition (2.5) and the endogenous perceived state-process $Y^{\theta_i}_t$ solving (2.5). The conditions in (4.1) ensure that investor holdings $\theta_{i,t}$ have no perceived price impact net of both effects.

A natural question concerns the impact of benchmarks like TWAP and Almgren-Chriss targets on competitive financial markets. By inserting (4.2) into (3.5), the price coefficients (3.4) become

$$\begin{aligned} g(t) &= \varphi_1 + 2\int_t^1 \frac{\kappa(u)\overline{\kappa}(u)}{\overline{\kappa}(u)M + \kappa(u)\overline{M}}\gamma(u)du, \\ \zeta(t) &= \varphi_0 + 2\overline{M}\int_t^1 \frac{\kappa(u)\overline{\kappa}(u)}{\overline{\kappa}(u)M + \kappa(u)\overline{M}}du. \end{aligned} \tag{4.7}$$

From (4.7), when $\varphi_0 > 0$ and $\varphi_1 > 0$, then both functions $g(t)$ and $\zeta(t)$ are positive (as per Remark 3.1.2), and equilibrium prices $\hat{S}_t$ differ from D_t. In the special case of $\varphi_0 := \varphi_1 := 0$, the formulas in (4.7) show that if either κ or $\overline{\kappa}$ is set to zero (but not both) over $[0, 1]$, then g and ζ become 0, and, thus, $\hat{S}_t$ becomes D_t. This is because liquidity supply becomes infinite with $L_{i,1} = 0$ for some investors.

For $u \in [0, 1)$, the integrand term $\frac{\kappa(u)\overline{\kappa}(u)}{\overline{\kappa}(u)M+\kappa(u)\overline{M}}$ in (4.7) is increasing in $\kappa(u)$ because its derivative with respect to $\kappa(u)$ is $\frac{M\overline{\kappa}(u)^2}{(\overline{\kappa}(u)M+\kappa(u)\overline{M})^2} \geq 0$. Therefore, given $M, \overline{M} \geq 1$, the functions $g(t)$ and $\zeta(t)$ in (4.7) are increasing in κ in the sense that $\kappa(u) \leq \kappa^\circ(u)$ for $u \in [0, t]$ produces corresponding ordered solutions $g(t) \leq g^\circ(t)$ and $\zeta(t) \leq \zeta^\circ(t)$. Hence, when the incremental order-execution costs and inventory and risk-management penalties for trading deviations from the target trajectory $\gamma(t)\tilde{a}_i$ are large, then predictable trends in intraday price pressure $g(t)\tilde{a}_\Sigma$ are large, and market illiquidity $\zeta(t)$ for hedging trading imbalances ϵB_t is high, and price-pressure variance $\zeta(t)^2\epsilon^2 t$ is high. This leads to the following comparison result:

Corollary 4.2 *In the competitive Radner equilibrium with ν_3^* and $\bar{\nu}_3^*$ from (4.2), the function $\zeta(t)$ in (4.7) is increasing in $\kappa(t)$ and $\bar{\kappa}(t)$. Consequently, for a fixed function $\overline{\kappa}(t)$ and a time point $t^\circ \in (0, 1)$, for two penalty-severity functions $\kappa(t)$ and $\kappa^\circ(t)$ ordered such that $\kappa(t) < \kappa^\circ(t)$ for $t \in [0, t^\circ)$ and $\kappa(t) = \kappa^\circ(t)$ for $t \in [t^\circ, 1]$, the illiquidity $\zeta(t)$ and price volatility in (3.15) are less in the market with $\kappa(t)$ than in the market with $\kappa^\circ(t)$. The same is true for analogous $\bar{\kappa}(t)$ and $\bar{\kappa}^\circ(t)$ given a fixed $\kappa(t)$.*

Proof The first claim follows from ζ's representation in (4.7). The second claim follows from the representation of $\hat{S}_t - D_t$ in (3.1). The third claim follows from symmetry. □

In addition, in the competitive equilibrium, predictable trends in intraday price pressure $g(t)\tilde{a}_\Sigma$ are increasing in the target ratio $\gamma(t)$, and price volatility is increasing in the hedging scalar ϵ.

For markets to clear, the equilibrium holdings $\hat{\theta}_{i,t}$ of targeted investors and hedgers differ from their targets. In the competitive equilibrium, these differences, from (4.5), are

$$\begin{aligned} \hat{\theta}_{i,t} - \gamma(t)\tilde{a}_i &= -\frac{\gamma(t)\tilde{a}_\Sigma + \overline{M}\epsilon B_t}{\overline{M}\frac{\kappa(t)}{\overline{\kappa}(t)} + M}, \quad i = 1, \ldots, M, \\ \hat{\theta}_{i,t} - \epsilon B_t &= -\frac{\gamma(t)\tilde{a}_\Sigma + \overline{M}\epsilon B_t}{\overline{M} + M\frac{\overline{\kappa}(t)}{\kappa(t)}}, \quad i = M+1, \ldots, M+\overline{M}. \end{aligned} \tag{4.8}$$

Thus, competitive rebalancers and liquidity providers split the aggregate imbalances $\gamma(t)\tilde{a}_\Sigma + \overline{M}\epsilon B_t$ with the hedgers independently of their individual targets $\tilde{a}_i$. The sharing is pro rata adjusted for their differential penalty severities. In this context, intraday liquidity providers (i.e., targeted investors with targets $\tilde{a}_i = 0$) absorb a share of the aggregate demand imbalance, but not the full imbalance given their inventory-holding penalties from (1.3). Our results are consistent with evidence in van Kervel et al. [42] that dynamic trading by large investors involves reciprocal liquidity provision (if investor targets are in opposite directions and net out) or reduced trading (if their targets are in the same direction). The limits of the differences in (4.8) as time $t \uparrow 1$ depend on the limiting behavior of the relative penalty severities $\frac{\overline{\kappa}(t)}{\kappa(t)}$. Section 7 gives examples of $\kappa(t)$ and $\overline{\kappa}(t)$ and illustrates different possible limits of (4.8).

In non-competitive equilibria, investors behave strategically over $t \in (0, 1]$ in that they perceive their holdings have an impact on prices.[25] In particular, the $\theta_{i,t}$-coefficients in the perceived drift (2.7) for $dS_t^{\theta_i}$ are non-zero. A natural non-competitive specification is that these drift coefficients would be negative so that increased positive holdings $\theta_{i,t}$ would depress perceived price drifts. Once again, the individual non-competitive strategies aggregate such that the resulting equilibrium prices only depend on the aggregate state variables $\tilde{a}_\Sigma$ and ϵB_t. However, if investors are not price-takers and (4.2) does not hold, then the individual rebalancer targets $\tilde{a}_i$ can appear in the rebalancer deviation $\hat{\theta}_{i,t} - \gamma(t)\tilde{a}_i$ in addition to $\tilde{a}_\Sigma$ and ϵB_t.

Empirical predictions Equation (4.7) and Corollary 4.2 lead to a set of empirical predictions about changing conditions across different days in a competitive market. Suppose that on different days there are different numbers of investors M (rebalancers and liquidity providers) and $\overline{M}$ (hedgers) with different penalty severities κ and $\overline{\kappa}$. Our model predicts the intraday price-pressure variance $\zeta(t)^2\epsilon^2 t$ and market illiquidity $\zeta(t)$ should be higher on days on which there are fewer rebalancers and liquidity providers M and larger penalty severities κ and $\overline{\kappa}$. These results are not normative critiques of TWAP trading but rather positive predictions about the empirical effect of targeted trading on daily market dynamics.

5 Equilibria with stochastic targets

Because our traders have linear preferences, they behave myopically in the sense that the solution to their individual optimization problems (1.1) is found by maximizing pointwise at each time $t \in [0, 1]$. Therefore, Theorem 3.2(i) continues to hold for equilibrium investor holdings $\hat{\theta}_{i,t}$ when the deterministic target ratio function $\gamma(t)$ is replaced with an arbitrary stochastic process $\gamma = (\gamma_t)_{t\in[0,1]}$ that is independent of the Brownian motions (D, B) and the private targets $(\tilde{a}_1, \ldots, \tilde{a}_M)$. A natural interpretation is that random intraday fluctuations in implicit bid-ask order-execution costs lead to changes in the target ratio γ_t. Once again, the rebalancer penalty process $L_{i,t}$ is a reduced-form for incremental order-execution, inventory, and risk-management costs relative to the now stochastic target trajectory $\gamma_t\tilde{a}_i$. Provided γ_t is observable for all investors at time $t \in [0, 1]$, we can re-define the filtrations in (1.12) as

$$\mathcal{F}_{i,t} := \begin{cases} \sigma(S_u, D_u, B_u, \gamma_u, \tilde{a}_i, \theta^c_{i,-})_{u\in[0,t]}, & i = 1, \ldots, M, \\ \sigma(S_u, D_u, B_u, \gamma_u, \theta^c_{i,-})_{u\in[0,t]}, & i = M+1, \ldots, M+\overline{M}. \end{cases} \tag{5.1}$$

Lemma A.1 in "Appendix A" (trader j's optimal response) continues to hold word-for-word when $\gamma(t)$ is replaced with γ_t and the martingale N_t in (2.4) is replaced by either of the two martingales in (5.6) and (5.13) below. From the proof of Lemma A.1, the reason this

[25] Even with non-price-taking behavior over $t \in (0, 1]$, our model still requires price-taking at the market open $t = 0$ such that (1.11) holds for the opening price S_0 both on- and off-equilibrium. Allowing for non-price-taking behavior at $t = 0$ would complicate the measurability condition in (1.11), and, thus, might require investors to filter over time to estimate $\tilde{a}_\Sigma$ rather than being able to infer it from S_0. Modeling targeted trading in this more complicated learning environment would be an interesting future extension.

extension is possible is that traders with linear utilities solve for their optimal holdings via pointwise optimization at each $t \in [0, 1]$.

Theorem 3.2(ii) for equilibrium prices $\hat{S}_t$ needs to be adjusted depending on the specification of the stochastic dynamics of the target-ratio process γ_t. We consider two specific examples of stochastic target ratios. Both examples have a zero initial value $\gamma_0 = 0$, and both tie down the terminal target ratio by requiring $\gamma_1 = 1$ at the end of the day.

Example 5.1 (Brownian bridge) In this target specification, the deterministic target ratio $\gamma(t)$ is replaced with the stochastic target-ratio process γ_t defined as the Brownian bridge process solving the stochastic differential equation

$$d\gamma_t := \frac{1-\gamma_t}{1-t}dt + dZ_t, \quad \gamma_0 := 0, \tag{5.2}$$

where $Z = (Z_t)_{t\in[0,1]}$ is an independent standard Brownian motion. At time s, the conditional expected future target ratio at time $t > s$ is

$$\mathbb{E}_s[\gamma_t] = \gamma_s + \frac{t-s}{1-s}(1-\gamma_s), \quad 0 \le s < t < 1. \tag{5.3}$$

Moreover, the drift in (5.2) ensures that $\gamma_t \to 1$ almost surely as $t \uparrow 1$. To derive the appropriate version of Theorem 3.2(ii), we redefine the candidate stock price (3.1) as

$$\hat{S}_t := D_t + \big(h(t) + \sigma(t)\gamma_t\big)\tilde{a}_\Sigma + \zeta(t)\epsilon B_t, \quad t \in [0, 1], \tag{5.4}$$

for three deterministic functions $h(t)$, $\sigma(t)$, and $\zeta(t)$ solving the system of linear ODEs:

$$\begin{aligned} \sigma'(t) &= \frac{\sigma(t)}{1-t} + \mu_1(t), \quad \sigma(1) = 0,\\ h'(t) &= -\frac{\sigma(t)}{1-t}, \quad h(1) = \varphi_1,\\ \zeta'(t) &= \mu_2(t), \quad \zeta(1) = \varphi_0, \end{aligned} \tag{5.5}$$

where the deterministic functions μ_1 and μ_2 are unchanged from (3.5). The linear ODE system (5.5) is triangular in the following sense: First, the ODE for $\sigma(t)$ in (5.5) is explicitly solved in (5.9) below. Second, given $\sigma(t)$, the remaining two ODEs in (5.5) for $h(t)$ and $\zeta(t)$ are solved by integrating, where their solutions are given in (5.10) below and (3.4). Furthermore, the martingale N_t in (3.2) is redefined as follows:

$$dN_t := dD_t + \zeta(t)\epsilon dB_t + \sigma(t)\tilde{a}_\Sigma dZ_t, \quad N_0 := 0. \tag{5.6}$$

Thus, a qualitatively new feature with a stochastic target ratio is that intraday price pressure due to the target imbalance $\tilde{a}_\Sigma$ is now random due to the target-ratio shocks dZ_t in (5.2).

Theorem 5.2 (Brownian bridge) *Under the assumptions of Theorem 3.2, $h(0) \neq 0$, and when the target ratio γ_t is a Brownian bridge, there exists an equilibrium in which:*

(i) *The perceived price parameters $\nu_0, \ldots, \nu_4$ and $\bar{\nu}_0, \ldots, \bar{\nu}_3$ are as in Theorem 3.2.*
(ii) *Investor equilibrium holdings are given by*

$$\hat{\theta}_{i,t} = \begin{cases} \alpha_1(t)\gamma_t\tilde{a}_\Sigma + \alpha_2(t)\epsilon B_t + \alpha_3(t)\gamma_t\tilde{a}_i, & i \in \{1, \ldots, M\},\\ \bar{\alpha}_1(t)\gamma_t\tilde{a}_\Sigma + \bar{\alpha}_2(t)\epsilon B_t, & i \in \{M+1, \ldots, M+\overline{M}\}, \end{cases} \tag{5.7}$$

where the deterministic functions $\alpha_1, \alpha_2, \alpha_3, \bar{\alpha}_1$, and $\bar{\alpha}_2$ are given by (3.8).

(iii) The equilibrium stock price is defined by (5.4) *with the martingale* N_t *given by* (5.6) *for deterministic functions* h, σ*, and* ζ *that are the unique solutions of the linear ODE system in* (5.5)*, and the price drift is given by*

$$\hat{\mu}_t := \mu_1(t)\gamma_t\tilde{a}_\Sigma + \mu_2(t)\epsilon B_t, \tag{5.8}$$

where the deterministic functions μ_1 *and* μ_2 *are given by* (3.5)*. Furthermore, the linear ODE for* σ *in* (5.5) *is uniquely solved by*

$$\sigma(t) = -\frac{1}{1-t}\int_t^1 (1-u)\mu_1(u)du, \quad t \in [0,1), \tag{5.9}$$

which satisfies $\lim_{t\uparrow 1}\sigma(t) = 0$*. The solution* (5.9) *ensures that* $\frac{\sigma(t)}{1-t}$ *is integrable; hence,* $h(t)$ *in* (5.5) *is found by integration for* $t \in [0,1]$*:*

$$h(t) = \varphi_1 + \int_t^1 \frac{\sigma(u)}{(1-u)}du. \tag{5.10}$$

The solution for $\zeta(t)$ *is identical to* (3.4)*.*

□

The coefficient $h(t) + \sigma(t)\gamma_t$ giving the price impact of the target imbalance $\tilde{a}_\Sigma$ with a Brownian bridge target ratio γ_t in (5.4) is related to the corresponding coefficient $g(t)$ with a deterministic target ratio $\gamma(t)$ as follows: Consider a deterministic target ratio equal to the ex ante expected Brownian bridge target ratio $\gamma(t) := \mathbb{E}[\gamma_t] = t$ and let $g(t)$ denote the associated deterministic price impact in (3.4). Then, for $t \in [0,1]$:

$$\begin{aligned}\mathbb{E}[h(t) + \sigma(t)\gamma_t] &= h(t) + \sigma(t)\gamma(t)\\ &= g(t),\end{aligned} \tag{5.11}$$

where the last equality holds because at $t = 1$ we have $g(1) = \varphi_1 = h(1)$ and for $t < 1$ the time-derivatives of both sides of (5.11) agree.[26] In other words, the price impact $h(t)+\sigma(t)\gamma_t$ in (5.4) with a stochastic target ratio γ_t has the same ex ante expected price impact as in the corresponding deterministic model with $\gamma(t) = \mathbb{E}[\gamma_t]$ plus additional randomness.

An unrealistic feature of the Brownian bridge target ratio is that γ_t can be negative as well as bigger than one with positive probability at times $t \in (0,1)$. Our next target-ratio process does not have these problems. The following construction is based on gamma processes.[27] Such pure jump processes have a long history of applications in option pricing theory (see, e.g., Madan et al. [31]).

Example 5.3 (Gamma bridge) The following model is based on Frei and Westray [19]. In this model variation, the deterministic target ratio $\gamma(t)$ is replaced with a stochastic target-ratio process γ_t that is a càdlàg gamma bridge process starting at $\gamma_0 = 0$ and ending at $\gamma_1 = 1$. In between the gamma bridge increases via a series of positive jumps that are dense on (0, 1]. Corollary 1 in Émery and Yor [16] ensures that γ_t has predictable intensity $\frac{1-\gamma_{t-}}{1-t}$ and that the

[26] For $t \in (0,1)$, Eq. (3.4) produces $g'(t) = \mu_1(t)\gamma(t) = \mu_1(t)t$ whereas (5.9) and (5.10) produce $\sigma'(t) = \frac{\sigma(t)}{1-t} + \mu_1(t)$ and $h'(t) = -\frac{\sigma(t)}{1-t}$. Combining these with the product rule produces the claim.

[27] Recall that a Lévy process $l = (l_t)_{t\in[0,1]}$ with $l_0 := 0$ and gamma distributed increments $l_t - l_s$, $0 \le s < t \le 1$, is called a gamma process. In our case, the mean and variance are normalized to one. A gamma bridge process is then defined by $\gamma_t := \frac{l_t}{l_1}$ for $t \in [0,1]$. See Émery and Yor [16] as well as Frei and Westray [19] for more details.

quadratic variation process $[\gamma]_t$ has predictable intensity $\frac{(1-\gamma_{t-})^2}{(1-t)(2-t)}$ where $\gamma_{t-} := \lim_{s\uparrow t}\gamma_s$ for $t \in (0,1]$ is the left-continuous version of the càdlàg process γ_t. In other words, the compensated processes

$$\gamma_t - \int_0^t \frac{1-\gamma_{s-}}{1-s}ds, \quad \text{and} \quad [\gamma]_t - \int_0^t \frac{(1-\gamma_{s-})^2}{(1-s)(2-s)}ds, \quad t \in [0,1], \tag{5.12}$$

are martingales.[28]

The equilibrium holdings $\hat{\theta}_{i,t}$ in the gamma bridge model are unchanged from (3.8) in Theorem 3.2(i). To derive the appropriate version of Theorem 3.2(ii) for the equilibrium stock price $\hat{S}_t$, the candidate stock price (3.1) is again redefined by (5.4) for deterministic functions (h,σ,ζ), and the martingale N_t is redefined as

$$dN_t := dD_t + \zeta(t)\epsilon dB_t + \sigma(t)\tilde{a}_\Sigma\big(d\gamma_t - \tfrac{1-\gamma_{t-}}{1-t}dt\big), \quad N_0 := 0. \tag{5.13}$$

Consider now the requirements in Definition 2.1 in this setting. First, the quadratic variation process $[N]_t$ of N_t has dynamics given by

$$d[N]_t = \big(1+\zeta(t)^2\epsilon^2\big)dt + \sigma(t)^2\tilde{a}_\Sigma^2 d[\gamma]_t. \tag{5.14}$$

Consequently, the second martingale in (5.12) produces dynamics for the predictable quadratic variation process (i.e., $[N]_t$'s compensator; see, e.g., p. 122 in Protter [38]) as

$$d\langle N\rangle_t = \Big(1+\zeta(t)^2\epsilon^2 + \sigma(t)^2\tilde{a}_\Sigma^2\frac{(1-\gamma_{t-})^2}{(1-t)(2-t)}\Big)dt. \tag{5.15}$$

Second, while holdings $\theta_{i,t}$ are required to be adapted to the filtration $\mathcal{F}_{i,t}$ defined in (5.1), the left-continuity requirement in Definition 2.1 (part of the càglàd requirement on admissible $\theta_{i,t}$) prevents trader $i \in \{1,\ldots,M\}$ from using holdings that depend on γ_t such as, e.g., $\tilde{a}_i\gamma_t$. This is because gamma bridges are not left-continuous. However, $\theta_{i,t}$ can depend on the left-continuous version γ_{t-}.[29]

Theorem 5.4 (Gamma bridge) *Under the assumptions of Theorem 3.2, $h(0) \neq 0$, and when the target ratio γ_t is a gamma bridge, there exists an equilibrium in which:*

(i) *The perceived price parameters $v_0,\ldots,v_4$ and $\bar{v}_0,\ldots,\bar{v}_3$ are as in Theorem 3.2.*

(ii) *Investor equilibrium holdings are given by*

$$\hat{\theta}_{i,t} = \begin{cases} \alpha_1(t)\tilde{a}_\Sigma\gamma_{t-} + \alpha_2(t)\epsilon B_t + \alpha_3(t)\tilde{a}_i\gamma_{t-}, & i \in \{1,\ldots,M\}, \\ \bar{\alpha}_1(t)\tilde{a}_\Sigma\gamma_{t-} + \bar{\alpha}_2(t)\epsilon B_t, & i \in \{M+1,\ldots,M+\overline{M}\}, \end{cases} \tag{5.16}$$

where the deterministic functions $\alpha_1, \alpha_2, \alpha_3, \bar{\alpha}_1$, and $\bar{\alpha}_2$ are given by (3.8).

(iii) *The equilibrium stock price is defined by (5.4) with the martingale N_t defined in (5.13) for deterministic functions h, σ, and ζ given as the unique solutions of the linear ODE system in (5.5), and the price drift is given by*

$$\hat{\mu}_t := \mu_1(t)\tilde{a}_\Sigma\gamma_{t-} + \mu_2(t)\epsilon B_t, \tag{5.17}$$

where the deterministic functions μ_1 and μ_2 are again defined by (3.5).

28 For simplicity, the underlying gamma process is normalized to have unit mean and unit variance, which among other properties gives us $\mathbb{E}[\gamma_t] = t$ and $\mathbb{E}[\gamma_t^2] = \frac{1}{2}t(1+t)$. However, the following analysis can easily be modified to include a parameter $m \in (0,\infty)$ by redefining the predictable intensity of the quadratic variation process $[\gamma]_t$ to be $\frac{(1-\gamma_{t-})^2}{(1-t)(1+m(1-t))}$ and leaving γ's predictable intensity as $\frac{1-\gamma_{t-}}{1-t}$. This would give us $\mathbb{E}[\gamma_t^2] = \frac{t(1+mt)}{1+m}$ whereas $\mathbb{E}[\gamma_t] = t$ is as before.

29 This admissibility restriction is essentially an assumption about how quickly investors can act on γ_t.

□

The equilibria with the Brownian bridge and gamma bridge target-ratio processes γ_t in Theorem 5.2 and Theorem 5.4 both have the following comparative statics:

Corollary 5.5 *In the setting of Theorems 5.2 and 5.4, the dynamics of the predictable quadratic variation process of the equilibrium price pressure $\hat{S}_t - D_t$ and the predictable quadratic cross-variation processes between $\hat{S}_t - D_t$ and (γ, B) are given by*

$$\begin{aligned} d\langle \hat{S} - D\rangle_t &= \sigma(t)^2 \tilde{a}_\Sigma^2 d\langle\gamma\rangle_t + \zeta(t)^2\epsilon^2 dt,\\ d\langle \hat{S} - D, \gamma\rangle_t &= \sigma(t)\tilde{a}_\Sigma d\langle\gamma\rangle_t,\\ d\langle \hat{S} - D, B\rangle_t &= \zeta(t)\epsilon dt. \end{aligned} \tag{5.18}$$

In the Brownian bridge model (5.2) we have $d\langle\gamma\rangle_t = dt$ whereas in the gamma bridge model the second martingale in (5.12) produces $d\langle\gamma\rangle_t = \frac{(1-\gamma_{t-})^2}{(1-t)(2-t)}dt$.

Proof From Theorems 5.2 and 5.4 the equilibrium stock-price process is (5.4) for both the Brownian and gamma bridge processes. The variations (5.18) follow from the representation of $\hat{S}_t - D_t$ from (5.4). □

The formulas in (5.18) show that comovement between price pressure and the underlying sources of randomness are completely determined by the solutions of the ODEs in (5.5). Consequently, variances and covariances for the price pressure $\hat{S}_t - D_t$ can be expressed in terms of these functions. For example, in the gamma bridge model, the price-pressure variance is

$$\begin{aligned} \mathbb{V}[\hat{S}_t - D_t|\sigma(\tilde{a}_\Sigma)] &= \mathbb{E}\Big[\big(\sigma(t)(\gamma_t - t)\tilde{a}_\Sigma + \zeta(t)\epsilon B_t\big)^2 |\sigma(\tilde{a}_\Sigma)\Big]\\ &= \sigma(s)^2\mathbb{E}[(\gamma_t - t)^2]\tilde{a}_\Sigma^2 + \zeta(t)^2\epsilon^2 t\\ &= \sigma(s)^2\frac{t(1-t)}{2}\tilde{a}_\Sigma^2 + \zeta(t)^2\epsilon^2 t, \end{aligned} \tag{5.19}$$

where the last equality uses $\mathbb{E}[\gamma_t] = t$ and $\mathbb{E}[\gamma_t^2] = \frac{1}{2}t(1+t)$. Section 7 provides numerical examples of the solutions to the ODEs in (5.5).

Empirical predictions In the competitive Radner equilibrium where ν_3^* and $\bar{\nu}_3^*$ are given by (4.2), the representation of ζ in (4.7) continues to hold whereas σ in (5.9) becomes

$$\sigma(t) = \frac{1}{1-t}\int_t^1 \frac{2(1-u)\kappa(u)\overline{\kappa}(u)}{\overline{M}\kappa(u) + M\overline{\kappa}(u)}du, \quad t \in [0,1). \tag{5.20}$$

Consequently, in the competitive Radner equilibrium, the comparison result for price volatility given in Corollary 4.2 continue to hold (i.e., price volatility is increasing in κ).

6 Connection to VWAP benchmark trading

The stochastic target ratio γ_t in Section 5 is driven by randomly evolving market conditions that affect intraday order-execution costs. This section shows how the stochastic target ratio γ_t can be linked to VWAP trading. To do this, we augment our earlier model so that aggregate market volume is now the sum of two components: First, as before, there is trading by the rebalancers, liquidity providers, and hedgers. Second, there is now additional trading by a new group of investors that we introduce into the model. These are a large number of other buyers and sellers—who we call the *crowd*—with inelastic trading demands that are assumed to naturally net to zero. For example, these could be retail investors and small asset managers trading for naturally offsetting personal reasons. In our augmented VWAP model, these two groups trade alongside each other. Since trading by the crowd is naturally offsetting, it has no impact on aggregate trading-demand imbalances. However, it does contribute to the aggregate market volume. We assume here that crowd volume affects (or is correlated with) order-execution costs and, thus, affects the trading target trajectories $\gamma_t \tilde{a}_i$ of the targeted investors via γ_t. In contrast, trading by our targeted investors affects aggregate order imbalances and, thus, market-clearing prices. However, we assume there is no feedback loop whereby their trading affects order-execution costs and, thus, their (exogenous in our model) intraday trading target trajectories.[30]

To this end, let $\mathrm{vol} = (\mathrm{vol}_t)_{t\in[0,1]}$ be an exogenous stochastic process for the stock's cumulative crowd volume vol_t over the interval $[0, t]$ for times $t \in [0, 1]$. The relative cumulative volume process is then defined as the ratio

$$\frac{\mathrm{vol}_t}{\mathrm{vol}_1}, \quad t \in [0, 1], \tag{6.1}$$

which is zero initially, has non-decreasing paths, and has terminal value one. Because vol_1 at time $t = 1$ cannot be observed at times $t < 1$, the volume ratio (6.1) also cannot be observed at times $t < 1$. Consequently, the ratio (6.1) cannot be used as a state-process.

The VWAP objective replacing (1.1) for targeted investors is

$$\sup_{\theta_i \in \mathcal{A}_i} \mathbb{E}\left[X_{i,1} - \int_0^1 \kappa(t) \Big(\frac{\mathrm{vol}_t}{\mathrm{vol}_1} \tilde{a}_i - \theta_{i,t} \Big)^2 dt \,\Big|\, \mathcal{F}_{i,0} \right], \quad i \in \{1, \dots, M\}. \tag{6.2}$$

The idea behind (6.2) is that intraday fluctuations in the relative volume ratio (6.1) affect the ex post intraday trading target trajectory $\frac{\mathrm{vol}_t}{\mathrm{vol}_1} \tilde{a}_i$. However, because investor i's position $\theta_{i,t}$ is adapted to the filtration $\mathcal{F}_{i,t}$ and because investor i cannot use her holdings $\theta_{i,t}$ to manipulate the intraday volume weights $\frac{\mathrm{vol}_t}{\mathrm{vol}_1}$, the optimization problem (6.2) can be replaced, given linear utilities, with the equivalent problem:[31]

$$\sup_{\theta_i \in \mathcal{A}_i} \mathbb{E}\left[X_{i,1} - \int_0^1 \kappa(t) \Big(\mathbb{E}\left[\frac{\mathrm{vol}_t}{\mathrm{vol}_1} \Big| \mathcal{F}_{i,t} \right] \tilde{a}_i - \theta_{i,t} \Big)^2 dt \,\Big|\, \mathcal{F}_{i,0} \right]. \tag{6.3}$$

We model $\mathbb{E}\left[\frac{\mathrm{vol}_t}{\mathrm{vol}_1} \Big| \mathcal{F}_{i,t} \right]$ directly as a gamma bridge γ_t. In that case, (6.3) becomes (1.1) when $\mathcal{F}_{i,t}$ is defined by (5.1) and $\gamma(t)$ is replaced by γ_t. Frei and Westray [19] model the realized relative volume curve $\frac{\mathrm{vol}_t}{\mathrm{vol}_1}$ used for VWAP benchmarking by the gamma bridge γ_t from Example 5.3. However, as discussed on page 617 in Frei and Westray [19], this presents a potential problem because the realized relative volume curve in (6.1) cannot be observed prior to the end of the trading day. In contrast, the intraday expected relative volume curve $\mathbb{E}\left[\frac{\mathrm{vol}_t}{\mathrm{vol}_1} | \mathcal{F}_{i,t} \right]$ in (6.3) in our model is—by definition—observable at times $t \in [0, 1]$. Thus, we model it as a gamma bridge.

Empirical implications: Prices in the gamma bridge model in Section 5 include a stochastic response to random changes in the target ratio γ_t, which includes the effect of changing

[30] In practice, VWAP benchmarking for an investor often excludes the investor's own trading from the measure of volume used contractually in VWAP benchmarking (see, e.g., Madhavan [32], Exhibit 1). For simplicity, we keep ϵ constant for hedgers. Allowing for a stochastic ϵ_t is also possible.

[31] The optimization problems in (6.2) and (6.3) are different and yield different objective values, but the two problems are equivalent in that they share the same maximizer. This is because the objectives in (6.2) and (6.3) differ only because of terms that do not involve $\theta_{i,t}$.

order-execution and inventory costs on the optimal target trajectory for targeted investors. Our VWAP analysis links these fluctuations in the stochastic target ratio γ_t to intraday fluctuations in the expected intraday volume ratio $\mathbb{E}\left[\frac{\text{vol}_t}{\text{vol}_1}\middle|\mathcal{F}_{i,t}\right]$. Thus, our VWAP model predicts that a positive shock to the expected daily volume ratio at time t increases the magnitude $|\hat{S}_t - D_t|$ of price pressure due to increased trading-demand imbalances due to an increase in the magnitude of the aggregate intraday trading target $|\gamma_t \tilde{a}_\Sigma|$.

7 Numerics

This section presents numerics for the competitive Radner equilibrium where (4.1) holds. Our analysis uses the two bridge models for the target ratio γ_t in Examples 5.1 and 5.3. The objects of interest are: (i) properties of the equilibrium price process $\hat{S}_t$ in (5.4), and (ii) how the rebalancers, liquidity providers, and hedgers optimally trade. The numerical properties here illustrate the analytic derivations in Section 5.

Our analysis uses the terminal stock-price restriction (1.10) (i.e., $\varphi_0 := \varphi_1 := 0$ in (1.8)), $M := 10$ targeted investors, and $\overline{M} := 10$ hedgers. The private targets $(\tilde{a}_1, \ldots, \tilde{a}_M)$ here are independent and have ex ante moments

$$\mathbb{E}[\tilde{a}_i] = 0, \quad \mathbb{E}[\tilde{a}_i^2] = 1, \quad i \in \{1, \ldots, M\}. \tag{7.1}$$

We consider penalty-severity functions in (1.3) and (1.4) defined by

$$\begin{aligned} \kappa(t) &:= \frac{(1-p)}{(1-t)^p}, \quad t \in [0,1), \quad p \in [0,1), \\ \overline{\kappa}(t) &:= \frac{(1-\bar{p})}{(1-t)^{\bar{p}}}, \quad t \in [0,1), \quad \bar{p} \in [0,1), \end{aligned} \tag{7.2}$$

which are parameterized by p and $\bar{p}$. For comparison purposes, these functions both integrate to one over $t \in [0, 1]$. A natural baseline is $p := 0$, where the penalty severity is constant over the day. For $p \in (0, 1)$, the penalty functions explode as $t \uparrow 1$ at various rates but still satisfy the integrability condition (3.3). When p is close to one (e.g., $p := 0.99$), the intraday penalty severities are negligible (i.e., close to zero) until close to the end of the day. In this case, our model mimics the situation where trader $i \in \{1, \ldots, M\}$ faces no intraday penalties but just a quasi-hard terminal constraint $\theta_{i,1} = \tilde{a}_i$. Figure 1 illustrates some of the $\kappa(t)$ functions used in our numerical analysis.

Our first topic is equilibrium pricing. From Theorems 3.2, 5.2, and 5.4, the price function $\zeta(t)$ in (3.1) and (5.4) is identical for the deterministic $\gamma(t)$ model and for both the Brownian and gamma bridge γ_t models and is given in (4.7). The price-loading functions $\sigma(t)$ and $h(t)$ do not appear in the deterministic $\gamma(t)$ model but are the same for both the Brownian and gamma bridge models and are given in (5.20) and (5.10). Figure 2 shows the price-loading functions $\sigma(t)$, $h(t)$, and $\zeta(t)$ for different values of p and $\bar{p}$. The signs of $\sigma(t)$, $h(t)$, and $\zeta(t)$ are all positive (from (5.9), (5.10), and Remark 3.1.2). We note that the greater the penalty severity $\kappa(t)$ is, the more sensitive prices are to shocks in the amount ϵB_t driving the hedger trading demand, which the targeted investors must provide. The values $\zeta(t)$ and $h(t)$ converge to φ_0 and φ_1 as $t \uparrow 1$, which in these numerics are, for simplicity, taken to be zero.

Fig. 1 Examples of penalty severity functions $\kappa(t)$ in (7.2). The lines are: $p := 0$ (——), $p := 0.1$ (– –), $p := 0.5$ (– · –), $p := 0.99$ (– · · –)

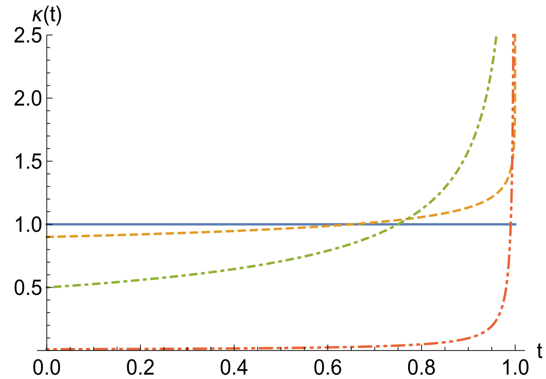

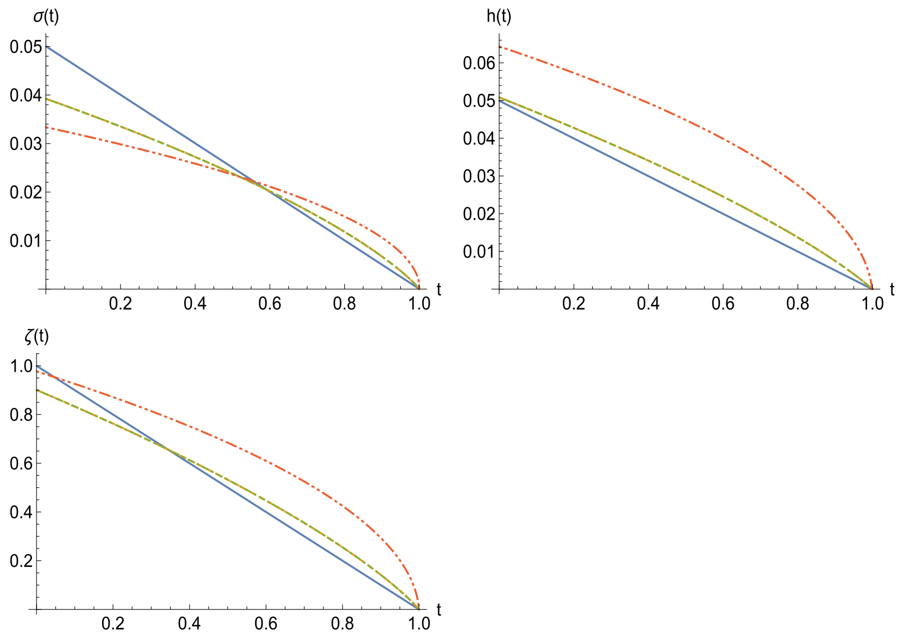

Fig. 2 Equilibrium price functions $\sigma(t)$, $h(t)$ and $\zeta(t)$ in (5.20), (5.10), and (4.7) for the competitive equilibrium (4.2). The parameters are given by (7.1)–(7.2), $\epsilon := 1$, and the discretization divides the day $t \in [0, 1]$ into 1,000 trading rounds. The lines are $p = \bar{p} = 0$ (——), $p = 0.5,\ \bar{p} = 0$ (– –), $p = 0,\ \bar{p} = 0.5$ (– · –), $p = \bar{p} = 0.5$ (– · · –)

Our second topic is the equilibrium stock holdings. For $i \in \{1, \ldots, M\}$, the optimal VWAP strategies for $i \in \{1, \ldots, M\}$ in (5.7) and (5.16) give expected holdings over the day

$$\begin{aligned}\mathbb{E}[\hat{\theta}_{i,t}|\sigma(\tilde{a}_i, \tilde{a}_\Sigma)] &= \alpha_1(t)\mathbb{E}[\gamma_{t-}]\tilde{a}_\Sigma + \alpha_2(t)\epsilon\mathbb{E}[B_t] + \alpha_3(t)\mathbb{E}[\gamma_{t-}]\tilde{a}_i \\ &= \alpha_1(t)t\tilde{a}_\Sigma + \alpha_3(t)t\tilde{a}_i.\end{aligned} \tag{7.3}$$

Consequently, from (7.3), trader $i \in \{1, \ldots, M\}$ expects ex ante to deviate from her target trajectory $\tilde{a}_i\gamma_{t-}$ by

$$\begin{aligned}\mathbb{E}[\hat{\theta}_{i,t} - \tilde{a}_i\gamma_{t-}|\sigma(\tilde{a}_i, \tilde{a}_\Sigma)] &= \alpha_1(t)t\tilde{a}_\Sigma + \big(\alpha_3(t) - 1\big)t\tilde{a}_i \\ &= \alpha_1(t)t\tilde{a}_\Sigma,\end{aligned} \tag{7.4}$$

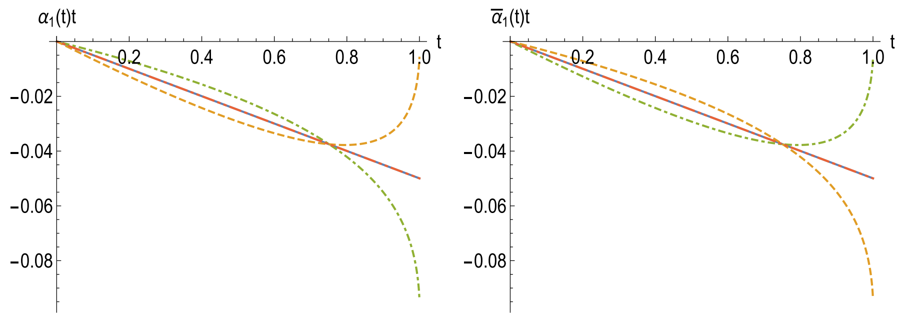

Fig. 3 $\tilde{a}_\Sigma$-coefficients in the conditional expected deviations (7.4) and (7.5) for the competitive equilibrium (4.2). The parameters are given by (7.1)–(7.2), $\epsilon := 1$, and the discretization divides the day $t \in [0, 1]$ into 1,000 trading rounds. The lines are: $p = \bar{p} = 0$ (——), $p = 0.5,\ \bar{p} = 0$ (−−), $p = 0,\ \bar{p} = 0.5$ (−·−), $p = \bar{p} = 0.5$ (−··−)

where the last equality follows from inserting the competitive values (4.2) into the expression for α_3 given in (3.9), which gives $\alpha_3(t) = 1$ here. Figure 3 shows the coefficient on $\tilde{a}_\Sigma$ in the expected deviations between targeted investor i's holdings up through time $t \in [0, 1]$ relative to her corresponding target. Similarly, for hedger $i \in \{M + 1, \ldots, M + \overline{M}\}$, the optimal strategies (5.7) and (5.16) give ex ante expected hedger target deviations

$$\mathbb{E}[\hat{\theta}_{i,t} - \epsilon B_t | \sigma(\tilde{a}_\Sigma)] = \bar{\alpha}_1(t) t \tilde{a}_\Sigma. \tag{7.5}$$

Finally, we illustrate that when $p \neq \bar{p}$, the targeted investors' and hedgers' soft target constraints can become hard constraints. In other words, if the targeted investor penalty severity $\kappa(t)$ explodes faster as $t \uparrow 1$ than the hedger penalty severity $\overline{\kappa}(t)$, the targeted investors hit their targets $\tilde{a}_i$ with probability one at the end of trading. To see this, we use (4.8) to compute the limit. Because $\gamma_1 = 1$, the terminal deviation for targeted investors $i \in \{1, \ldots, M\}$ is

$$\hat{\theta}_{i,1} - \tilde{a}_i = \begin{cases} 0, & \bar{p} < p, \\ -\frac{\tilde{a}_\Sigma + \overline{M}\epsilon B_1}{\overline{M} + M}, & \bar{p} = p, \\ -\frac{\tilde{a}_\Sigma + \overline{M}\epsilon B_1}{M}, & \bar{p} > p. \end{cases} \tag{7.6}$$

For hedgers $i \in \{M + 1, \ldots, M + \overline{M}\}$, the terminal deviation is

$$\hat{\theta}_{i,1} - \epsilon B_1 = \begin{cases} -\frac{\tilde{a}_\Sigma + \overline{M}\epsilon B_1}{\overline{M}}, & \bar{p} < p, \\ -\frac{\tilde{a}_\Sigma + \overline{M}\epsilon B_1}{\overline{M} + M}, & \bar{p} = p, \\ 0 & \bar{p} > p. \end{cases} \tag{7.7}$$

8 Extension to inhomogeneous investors

We can allow for inhomogeneity within the realtime hedgers. To illustrate this, we split the group of real-time hedgers into two homogenous subgroups with respectively $\bar{M}$ and M° investors. Therefore, we replace (1.4) with hedger penalties

$$L_{i,t} := \begin{cases} \int_0^t \overline{\kappa}(s)\big(\theta_{i,s} - \epsilon B_s\big)^2 ds, & i \in \{M + 1, \ldots, M + \bar{M}\}, \\ \int_0^t \kappa^\circ(s)\big(\theta_{i,s} - \epsilon^\circ B_s\big)^2 ds, & i \in \{M + \bar{M} + 1, \ldots, M + \bar{M} + M^\circ\}, \end{cases} \tag{8.1}$$

for $t \in [0, 1]$. In the penalty processes in (8.1), the two subgroups' penalty-severity functions $\kappa^\circ(t)$ and $\bar{\kappa}(t)$ can differ and their scalars ϵ° and ϵ can differ.

Up to minor modifications, the equilibrium structure in Theorem 3.2 continues to hold, but for this model extension there are now three free perceived price-impact parameters $(\nu_3, \bar{\nu}_3, \nu_3^\circ)$. The equilibrium stock-price drift in (3.2) becomes

$$\begin{aligned}\hat{\mu}_t =& -\frac{(M+\bar{M}+M^\circ)C\bar{C}C^\circ(2\kappa+\nu_4)}{2\big((M+\bar{M}+M^\circ)^2-1\big)(\kappa-\nu_3)(M^\circ C\bar{C}+\bar{M}CC^\circ+M\bar{C}C^\circ)}\gamma\tilde{a}_\Sigma \\ &-\frac{2C(\bar{M}\bar{\kappa}C^\circ\epsilon+M^\circ\kappa^\circ\bar{C}\epsilon^\circ)}{M^\circ C\bar{C}+\bar{M}CC^\circ+M\bar{C}C^\circ}B_t,\end{aligned} \tag{8.2}$$

where

$$\begin{aligned}C(t) &:= 2(M+\bar{M}+M^\circ)\kappa(t)-(1+M+\bar{M}+M^\circ)\nu_3(t),\\ \bar{C}(t) &:= 2(M+\bar{M}+M^\circ)\bar{\kappa}(t)-(1+M+\bar{M}+M^\circ)\bar{\nu}_3(t),\\ C^\circ(t) &:= 2(M+\bar{M}+M^\circ)\kappa^\circ(t)-(1+M+\bar{M}+M^\circ)\nu_3^\circ(t),\\ \nu_4 &:= -\frac{4\kappa^2+2(1+M+\bar{M}+M^\circ)(-2+M+\bar{M}+M^\circ)\kappa\nu_3}{(M+\bar{M}+M^\circ)C}.\end{aligned} \tag{8.3}$$

Likewise, the equilibrium investor holdings in (3.8) become

$$\begin{aligned}\hat{\theta}_t =& -\frac{(M+\bar{M}+M^\circ)\bar{C}C^\circ(2\kappa+\nu_4)}{2(1+M+\bar{M}+M^\circ)(\kappa-\nu_3)(M^\circ C\bar{C}+\bar{M}CC^\circ+M\bar{C}C^\circ)}\gamma\tilde{a}_\Sigma\\ &-\frac{2(-1+M+\bar{M}+M^\circ)(\bar{M}\bar{\kappa}C^\circ\epsilon+M^\circ\kappa^\circ\bar{C}\epsilon^\circ)}{M^\circ C\bar{C}+\bar{M}CC^\circ+M\bar{C}C^\circ}B_t\\ &+\frac{(M+\bar{M}+M^\circ)(2\kappa+\nu_4)}{2(1+M+\bar{M}+M^\circ)(\kappa-\nu_3)}\gamma\tilde{a}_i,\\ \hat{\bar{\theta}}_t =& -\frac{(M+\bar{M}+M^\circ)CC^\circ(2\kappa+\nu_4)}{2(1+M+\bar{M}+M^\circ)(\kappa-\nu_3)(M^\circ C\bar{C}+\bar{M}CC^\circ+M\bar{C}C^\circ)}\gamma\tilde{a}_\Sigma\\ &+\frac{2(-1+M+\bar{M}+M^\circ)(\bar{\kappa}(MC^\circ+M^\circ C)\epsilon-M^\circ\kappa^\circ C\epsilon^\circ)}{M^\circ C\bar{C}+\bar{M}CC^\circ+M\bar{C}C^\circ}B_t,\\ \hat{\theta}_t^\circ =& -\frac{(M+\bar{M}+M^\circ)C\bar{C}(2\kappa+\nu_4)}{2(1+M+\bar{M}+M^\circ)(\kappa-\nu_3)(M^\circ C\bar{C}+\bar{M}CC^\circ+M\bar{C}C^\circ)}\gamma\tilde{a}_\Sigma\\ &+\frac{2(-1+M+\bar{M}+M^\circ)(\kappa^\circ(M\bar{C}+\bar{M}C)\epsilon^\circ-\bar{M}\bar{\kappa}C\epsilon)}{M^\circ C\bar{C}+\bar{M}CC^\circ+M\bar{C}C^\circ}B_t.\end{aligned} \tag{8.4}$$

We conjecture that the model structure scales linearly with as many free perceived hedger price-impact coefficients $\bar{\nu}_3$ as there are heterogenous hedger subgroups.

9 Discussion

Two features of the mathematical structure of our model are particularly important for tractability. The first is that the aggregate target $\tilde{a}_\Sigma$ is inferable at time $t = 0$ and so, as a result, there is no need for investors to filter trading data over time to estimate the other investors' private targets. The second is the linear preference structure, which leads to the optimal controls $\hat{\theta}_{i,t}$ at time $t \in [0, 1]$ being solutions to pointwise optimization problems.

There are good reasons to think the qualitative properties of our analysis are robust even though the specific functional forms of prices and trading strategies depend on our modeling assumptions (e.g., Brownian motion dynamics). First, investor perceptions about how prices respond to off-equilibrium orders are likely a key factor determining the equilibrium form. Thus, price-sensitive investors should still trade off intraday target-tracking inventory penalties and perceived price-impact costs, even if price impact is non-linear. Second, rational forward-looking liquidity provision is still likely to lead to the absence of manipulative predatory trading. Third, intraday liquidity is likely to be reduced by intraday trading target penalties relative to just terminal end-of-day target penalties.

Lastly, we comment on numerical implementation. The model is characterized by low dimensional state-processes that makes numerics fast to perform. Furthermore, the model's linear structure makes the numerics stable (coupled linear ODEs). We have experimented extensively with the numerics and have not found any instability concerns.

10 Conclusion

This paper has solved for continuous-time Subgame Perfect Nash equilibria with endogenous liquidity provision and intraday trading targets. We show how TWAP, VWAP, and other trading benchmarks induce intraday patterns in investor positions and in price dynamics. There are also potential extensions of our model. First, it would be interesting to extend the model to allow for more heterogeneity in the investor optimization problems. For example, $\gamma(t)$ and ϵ could be replaced with different ratios $\gamma_i(t)$ for targeted investors and ϵ_i or a stochastic process ϵ_t for hedgers. Second, perhaps the most pertinent extension would be to allow for additional randomness such that the initial equilibrium stock price $\hat{S}_0$ cannot fully reveal the aggregate target $\tilde{a}_\Sigma$. Such an extension would require filtering to learn about trading-demand imbalances. Third, the Brownian motion driving hedger demands might be private information rather than publicly observable.

Proofs

Lemma A.1 (Trader j's optimal response) *Let $\nu_0(t) \neq 0$, $\bar{\nu}_0(t) \neq 0$, and assume that* (2.3) *holds. Fix an exogenous state-process with càglàd paths $Y = (Y_t)_{t\in[0,1]}$ and fix a trader index $j \in \{1, \dots, M + \overline{M}\}$. When the stock price S in the wealth* (1.2) *and filtration $\mathcal{F}_{i,t}$ in* (1.12) *is $S := S^Y$ with drift* (2.2) *and martingale* (2.4)*, the optimizer for* (1.1) *over $\mathcal{A}_j$ is*

$$\theta^Y_{j,t} := \begin{cases} \frac{1}{2\left(\kappa(t)-\nu_3(t)\right)}\Big(\nu_0(t)Y_t + \nu_1(t)\gamma(t)\tilde{a}_\Sigma + \nu_2(t)B_t + \big(\nu_4(t) + 2\kappa(t)\big)\gamma(t)\tilde{a}_j\Big), & j \in \{1, \dots, M\}, \\ \frac{1}{2\left(\bar{\kappa}(t)-\bar{\nu}_3(t)\right)}\Big(\bar{\nu}_0(t)Y_t + \bar{\nu}_1(t)\gamma(t)\tilde{a}_\Sigma + \big(\bar{\nu}_2(t) + 2\epsilon\bar{\kappa}(t)\big)B_t\Big), & j \in \{M+1, \dots, M+\overline{M}\}, \end{cases} \tag{A.1}$$

provided that $\theta^Y_{j,t}$ satisfies the integrability condition (2.1).

Proof of Lemma A.1 Because utilities are linear, we have the following representation for an arbitrary control $\theta_j \in \mathcal{A}_j$ for targeted investors $j \in \{1, \dots, M\}$

$$\begin{aligned} &\mathbb{E}[X_{j,1} - L_{j,1}|\mathcal{F}_{j,0}] \\ &\quad = X_{j,0} + \mathbb{E}\Big[\int_0^1 \Big(\theta_{j,s}\big(\nu_0(s)Y_s + \nu_1(s)\gamma(s)\tilde{a}_\Sigma + \nu_2(s)B_s + \nu_3(s)\theta_{j,s} + \nu_4(s)\gamma(s)\tilde{a}_j\big) \\ &\quad - \kappa(s)\big(\theta_{j,s} - \gamma(s)\tilde{a}_j\big)^2 \Big) ds \,\Big|\, \mathcal{F}_{j,0} \Big]. \end{aligned} \tag{A.2}$$

For hedgers $j \in \{M+1, \dots, M+\overline{M}\}$, we have a similar representation

$$\begin{aligned} &\mathbb{E}[X_{j,1} - L_{j,1}|\mathcal{F}_{j,0}] \\ &\quad = X_{j,0} + \mathbb{E}\Big[\int_0^1 \Big(\theta_{j,s}\big(\bar{\nu}_0(s)Y_s + \bar{\nu}_1(s)\gamma(s)\tilde{a}_\Sigma + \bar{\nu}_2(s)B_s + \bar{\nu}_3(s)\theta_{j,s}\big) \\ &\quad - \overline{\kappa}(s)\big(\theta_{j,s} - \epsilon B_s\big)^2 \Big) ds \,\Big|\, \mathcal{F}_{j,0} \Big]. \end{aligned} \tag{A.3}$$

The martingale property of $\int \theta_j dN$ is used to eliminate $\mathbb{E}[\int_0^1 \theta_{j,s} dN_s]$ in both (A.2) and (A.3). The integrands in the ds-integrals are quadratic in $\theta_{j,s}$. Consequently, the supremum in (1.1) is achieved by maximizing the integrands in (A.2) and (A.3) pointwise over $\theta_{j,s}$ at each state and at each time $s \in [0, 1]$. This gives the optimal holdings in (A.1).

The left-continuity of Y's paths allows investor j to infer Y from past and current observations of S^Y and (D, B). To see this, it suffices to show that observing

$$\int_0^t \nu_0(s)Y_s ds \text{ for } j \in \{1, \dots, M\} \text{ and } \int_0^t \bar{\nu}_0(s)Y_s ds \text{ for } j \in \{M+1, \dots, M+\overline{M}\}, \tag{A.4}$$

over time $t \in [0, 1)$ is sufficient for investor j to infer Y_t. The integrals in (A.4) are well-defined because $\nu_0(s)$ and $\bar{\nu}_0(s)$ are continuous on $[0, 1)$ and Y has càglàd paths (hence, Y's paths are locally bounded). The ability to infer Y_t from (A.4) follows directly from computing the time derivative from the left of (A.4). This left-derivative is $\nu_0(t)Y_t$ for $j \in \{1, \dots, M\}$ and $\bar{\nu}_0(t)Y_t$ for $j \in \{M+1, \dots, M+\overline{M}\}$ by the left-continuity requirement placed on Y's paths. Consequently, Y_t is $\mathcal{F}_{i,t}$ measurable when $S := S^Y$ in (1.12) where S^Y is defined using (2.2). Therefore, Y_t can be used as a state-variable for investor j in (A.1). □

Comment: Because our traders are penalized for deviations of their holdings from intraday targets, investor optimal stock holdings are given in terms of levels rather than trading rates. This property allows our investors to absorb trading noise with only finite quadratic variation such as the Brownian motion dynamics in, e.g., Kyle [30].

Proof of Lemma 2.2 We define the state-process

$$Y_t^{\theta_i} := \begin{cases} \frac{2\kappa(\bar{\nu}_3-\overline{\kappa})-(\overline{\kappa}-\bar{\nu}_3)\big((M-1)\nu_1+\nu_4\big)+\overline{M}\bar{\nu}_1(\nu_3-\kappa)}{(M-1)\nu_0(\overline{\kappa}-\bar{\nu}_3)+\overline{M}\bar{\nu}_0(\kappa-\nu_3)}\gamma\tilde{a}_\Sigma - \frac{(M-1)\nu_2(\overline{\kappa}-\bar{\nu}_3)+\overline{M}(\kappa-\nu_3)(\bar{\nu}_2+2\overline{\kappa}\epsilon)}{(M-1)\nu_0(\overline{\kappa}-\bar{\nu}_3)+\overline{M}\bar{\nu}_0(\kappa-\nu_3)}B_t \\ \quad -\frac{2(\kappa-\nu_3)(\overline{\kappa}-\bar{\nu}_3)}{(M-1)\nu_0(\overline{\kappa}-\bar{\nu}_3)+\overline{M}\bar{\nu}_0(\kappa-\nu_3)}\theta_{i,t} + \frac{(\overline{\kappa}-\bar{\nu}_3)(2\kappa+\nu_4)}{(M-1)\nu_0(\overline{\kappa}-\bar{\nu}_3)+\overline{M}\bar{\nu}_0(\kappa-\nu_3)}\gamma\tilde{a}_i, \quad i \in \{1, \dots, M\}, \\ -\frac{2\kappa(\overline{\kappa}-\bar{\nu}_3)+(\overline{\kappa}-\bar{\nu}_3)(M\nu_1+\nu_4)+(\overline{M}-1)\bar{\nu}_1(\kappa-\nu_3)}{M\nu_0(\overline{\kappa}-\bar{\nu}_3)+(\overline{M}-1)\bar{\nu}_0(\kappa-\nu_3)}\gamma\tilde{a}_\Sigma + \frac{M\nu_2(\bar{\nu}_3-\overline{\kappa})-(\overline{M}-1)\bar{\nu}_2(\kappa-\nu_3)-2\overline{\kappa}(\overline{M}-1)\epsilon(\kappa-\nu_3)}{M\nu_0(\overline{\kappa}-\bar{\nu}_3)+(\overline{M}-1)\bar{\nu}_0(\kappa-\nu_3)}B_t \\ \quad -\frac{2(\kappa-\nu_3)(\overline{\kappa}-\bar{\nu}_3)}{M\nu_0(\overline{\kappa}-\bar{\nu}_3)+(\overline{M}-1)\bar{\nu}_0(\kappa-\nu_3)}\theta_{i,t}, \quad i \in \{M+1, \dots, M+\overline{M}\}, \end{cases} \tag{A.5}$$

so that the corresponding drift process (2.2) perceived by investor $j \neq i$ becomes

$$\begin{cases} \nu_0(t)Y_t^{\theta_i} + \nu_1(t)\gamma(t)\tilde{a}_\Sigma + \nu_2(t)B_t + \nu_3(t)\theta_{j,t} + \nu_4(t)\gamma(t)\tilde{a}_j, & j \in \{1,\dots,M\}, \\ \bar{\nu}_0(t)Y_t^{\theta_i} + \bar{\nu}_1(t)\gamma(t)\tilde{a}_\Sigma + \bar{\nu}_2(t)B_t + \bar{\nu}_3(t)\theta_{j,t}, & j \in \{1+M,\dots,M+\overline{M}\}. \end{cases} \tag{A.6}$$

Thus, trader j's optimal response holdings $\theta_{j,t}^{Y^{\theta_i}}$ to $Y_t^{\theta_i}$ are given by (A.1) in Lemma A.1. By summing these holdings, we see that (2.5) holds. □

Proof of Theorem 3.2 We first insert S^{θ_i} from (2.6) into the objective (1.1) to get

$$\begin{aligned} &\mathbb{E}\Big[\int_0^1 \Big(\theta_{i,t}\mu_t^{\theta_i} - \kappa(t)\big(\theta_{i,t} - \gamma(t)\tilde{a}_i\big)^2\Big)dt\Big], \quad i \in \{1,\dots,M\}, \\ &\mathbb{E}\Big[\int_0^1 \Big(\theta_{i,t}\mu_t^{\theta_i} - \overline{\kappa}(t)\big(\theta_{i,t} - \epsilon B_t\big)^2\Big)dt\Big], \quad i \in \{M+1,\dots,M+\overline{M}\}, \end{aligned} \tag{A.7}$$

where $\mu_t^{\theta_i}$ is defined in (2.7). In (A.7), the expectation $\mathbb{E}[\int_0^1 \theta_{i,t}dN_t]$ drops out because $\theta_i \in \mathcal{A}_i$ ensures that the stochastic integrals $\int \theta_i dN$ are martingales. The integrands in (A.7) are quadratic functions of $\theta_{i,t}$ and the quadratic terms produce the second-order conditions

$$\begin{aligned} &\frac{(\kappa - \nu_3)((M+1)\nu_0(\overline{\kappa} - \bar{\nu}_3) + \overline{M}\bar{\nu}_0(\kappa - \nu_3))}{(M-1)\nu_0(\overline{\kappa} - \bar{\nu}_3) + \overline{M}\bar{\nu}_0(\kappa - \nu_3)} > 0, \\ &\frac{(\overline{\kappa} - \bar{\nu}_3)(M\nu_0(\overline{\kappa} - \bar{\nu}_3) + (\overline{M}+1)\bar{\nu}_0(\kappa - \nu_3))}{M\nu_0(\overline{\kappa} - \bar{\nu}_3) + (\overline{M}-1)\bar{\nu}_0(\kappa - \nu_3)} > 0. \end{aligned} \tag{A.8}$$

Provided (A.8) holds, the pointwise maximizers of the integrands in (A.7) are given by

$$\hat{\theta}_{i,t} := \begin{cases} \Big(\tilde{a}_i\gamma(2\kappa + \nu_4)(M\nu_0(\overline{\kappa} - \bar{\nu}_3) + \overline{M}\bar{\nu}_0(\kappa - \nu_3)) + \tilde{a}_\Sigma\gamma\big(2\kappa\nu_0(\bar{\nu}_3 - \overline{\kappa}) - \overline{\kappa}\nu_0\nu_4 \\ \quad +\nu_0\nu_4\bar{\nu}_3 + \overline{M}(\kappa - \nu_3)(\nu_1\bar{\nu}_0 - \nu_0\bar{\nu}_1)\big) + B_t\overline{M}(\kappa - \nu_3)(\nu_2\bar{\nu}_0 - \nu_0(\bar{\nu}_2 + 2\overline{\kappa}\epsilon))\Big) \\ \quad \Big/\Big(2(\kappa - \nu_3)((M+1)\nu_0(\overline{\kappa} - \bar{\nu}_3) + \overline{M}\bar{\nu}_0(\kappa - \nu_3))\Big), \quad i \in \{1,\dots,M\}, \\ \frac{\tilde{a}_\Sigma\gamma\big(M\nu_0\bar{\nu}_1 - \bar{\nu}_0(2\kappa + M\nu_1 + \nu_4)\big) + B_t M(\nu_0\bar{\nu}_2 - \nu_2\bar{\nu}_0 + 2\overline{\kappa}\nu_0\epsilon)}{2(M\nu_0(\overline{\kappa} - \bar{\nu}_3) + (\overline{M}+1)\bar{\nu}_0(\kappa - \nu_3))}, \quad i \in \{M+1,\dots,M+\overline{M}\}. \end{cases} \tag{A.9}$$

Summing (A.9) over $i \in \{1,\dots,M+\overline{M}\}$ shows that the stock market clears in the sense that (1.5) holds when $\bar{\nu}_0$ is defined as in (3.6).

Expression (3.6) allows us to re-write the second-order conditions in (A.8) as

$$\begin{aligned} &\frac{(\kappa - \nu_3)(M + \overline{M} + 1)}{M + \overline{M} - 1} > 0, \\ &\frac{(\overline{\kappa} - \bar{\nu}_3)(M + \overline{M} + 1)}{M + \overline{M} - 1} > 0. \end{aligned} \tag{A.10}$$

Because we have assumed (2.3), the second-order conditions (A.10) hold. Consequently, the pointwise optimizers are given in (A.9).

To ensure that the resulting stock-price drift processes $\mu_t^{\hat{\theta}_i}$ are the same for all investors $i \in \{1, \dots, M+\overline{M}\}$ and, in particular, ensuring that the individual private targets $(\tilde{a}_1, \dots, \tilde{a}_M)$ do not appear in $\mu_t^{\hat{\theta}_i}$, we define the deterministic functions in (3.7). By inserting (3.6) and (3.7) into (A.9) and simplifying, gives

$$\hat{\theta}_{i,t} = \begin{cases} \frac{2(M+\overline{M}-1)}{M}\Big(\frac{\overline{M}(\tilde{a}_\Sigma\gamma\kappa - B_t\overline{\kappa}M\epsilon)}{M^2(2\overline{\kappa}-\bar{\nu}_3)+M\overline{M}(2(\kappa+\overline{\kappa})-\nu_3)-M(\overline{M}+1)\bar{\nu}_3+\overline{M}(2\kappa\overline{M}-\nu_3(\overline{M}+1))} \\ \quad +\frac{\gamma\kappa(\tilde{a}_iM-\tilde{a}_\Sigma)}{2\kappa(M+\overline{M})-\nu_3(M+\overline{M}+1)}\Big), \quad i \in \{1,\dots,M\}, \\ \frac{2(M+\overline{M}-1)(B_t\overline{\kappa}M\epsilon-\tilde{a}_\Sigma\gamma\kappa)}{M^2(2\overline{\kappa}-\bar{\nu}_3)+M\overline{M}(2(\kappa+\overline{\kappa})-\nu_3)-M(\overline{M}+1)\bar{\nu}_3+\overline{M}(2\kappa\overline{M}-\nu_3(\overline{M}+1))}, \quad i \in \{M+1,\dots,M+\overline{M}\}. \end{cases} \tag{A.11}$$

Because the remaining functions ν_0, ν_1, and ν_2 do not appear in (A.11), these functions are irrelevant in the sense that different ν_0, ν_1, and ν_2 functions all produce the same equilibrium prices and investor holdings provided $\nu_0 \neq 0$. The functions $\alpha_1, \alpha_2, \alpha_3, \bar{\alpha}_1, \bar{\alpha}_2$ in (3.9) are found by matching $(\tilde{a}_\Sigma, B_t, \tilde{a}_i)$ coefficients in (A.11), which also produces the expression for $\hat{\theta}_{i,t}$ in (3.8).

At this point, the stock-price drift processes $\mu_t^{\hat{\theta}_i}$ are all identical, and we define $\hat{\mu}_t$ as their common value. The representation (3.10) of $\hat{\mu}_t$ follows from inserting (A.11) into $\mu_t^{\hat{\theta}_i}$ and matching $(\tilde{a}_\Sigma, B_t)$ coefficients. The resulting functions $\mu_1(t)$ and $\mu_2(t)$ are given in (3.5), which, we show next, are integrable. The bounds (3.13) produce

$$\begin{aligned} |\mu_1| &\leq \min\Big\{\frac{2\kappa}{M}, \frac{2(2(M+\overline{M})\overline{\kappa}-(1+M+\overline{M})\bar{\nu}_3)}{\overline{M}}\Big\}, \\ |\mu_2| &\leq \min\Big\{2\overline{\kappa}, \frac{2\overline{M}(2(M+\overline{M})\kappa-(1+M+\overline{M})\nu_3)}{M}\Big\}. \end{aligned} \tag{A.12}$$

The integrability (3.3) assumed of κ and $\overline{\kappa}$ ensure integrability of the upper bounds in (A.12). Therefore, the ODEs (3.4) can be solved by integrating. To see that the functions $\alpha_1, \alpha_2, \alpha_3, \bar{\alpha}_1, \bar{\alpha}_2$ in (3.9) are uniformly bounded we again use (A.12):

$$\begin{aligned} |\alpha_1| &\leq \frac{M+\overline{M}-1}{\kappa}|\mu_1| \leq \frac{2(M+\overline{M}-1)}{M}, \\ |\alpha_2| &\leq \frac{2\overline{M}(M+\overline{M}-1)}{M}, \\ |\alpha_3| &\leq 2(M+\overline{M}-1), \\ |\bar{\alpha}_1| &\leq \frac{2(M+\overline{M}-1)}{\overline{M}}, \\ |\bar{\alpha}_2| &\leq \frac{M(M+\overline{M}-1)}{\overline{M}\overline{\kappa}}|\mu_2| \leq \frac{2M(M+\overline{M}-1)}{\overline{M}}. \end{aligned}$$

Because $\alpha_1, \alpha_2, \alpha_3, \bar{\alpha}_1, \bar{\alpha}_2$ are bounded functions, the process $\hat{\theta}_{i,t}$ in (3.8) is admissible in the sense of Definition 2.1, and its optimality follows.

The last step of the proof establishes the terminal price condition (1.8). Itô's lemma produces the dynamics of $\hat{S}_t$ in (3.1) to be

$$\begin{aligned} d\hat{S}_t &= g'(t)\tilde{a}_\Sigma dt + dD_t + \zeta'(t)\epsilon B_t dt + \zeta(t)\epsilon dB_t \\ &= \hat{\mu}_t dt + dD_t + \zeta(t)\epsilon dB_t, \end{aligned} \tag{A.13}$$

where the last equality uses (3.4). The terminal conditions in (3.4) produce the terminal price restriction (1.8). □

Proof of Theorem 4.1 The Cauchy-Schwarz inequality ensures that when $M \geq 2$ and $(\tilde{a}_1, \ldots, \tilde{a}_M)$ are not perfectly correlated, we have

$$M \sum_{i=1}^{M} \mathbb{E}[\tilde{a}_i^2] > \mathbb{E}[\tilde{a}_\Sigma^2]. \tag{A.14}$$

We can rewrite (4.4) as

$$\sum_{i=1}^{M+\overline{M}} \mathbb{E}[\text{CE}_i] = \int_0^1 F\big(\nu_3(t), \bar{\nu}_3(t)\big) dt, \tag{A.15}$$

where the function $F : \mathbb{R}^2 \mapsto \mathbb{R}$ is defined by

$$\begin{aligned} F\big(\nu_3(t), \bar{\nu}_3(t)\big) &:= \sum_{i=1}^{M} \mathbb{E}\Big[\hat{\theta}_{i,t}\hat{\mu}_t - \kappa(t)\big(\gamma(t)\tilde{a}_i - \hat{\theta}_{i,t}\big)^2\Big] \\ &\quad + \sum_{i=M+1}^{M+\overline{M}} \mathbb{E}\Big[\hat{\theta}_{i,t}\hat{\mu}_t - \overline{\kappa}(t)(\hat{\theta}_{i,t} - \epsilon B_t)^2\Big] \\ &= -\sum_{i=1}^{M} \mathbb{E}\Big[\kappa(t)\big(\gamma(t)\tilde{a}_i - \hat{\theta}_{i,t}\big)^2\Big] - \sum_{i=M+1}^{M+\overline{M}} \mathbb{E}\Big[\overline{\kappa}(t)(\hat{\theta}_{i,t} - \epsilon B_t)^2\Big], \end{aligned} \tag{A.16}$$

and where $\hat{\theta}_{i,t}$ is as in Theorem 3.2. The equality in (A.16) follows from $\sum_{i=1}^{M+\bar{M}} \hat{\theta}_{i,t} = 0$. Therefore, our goal is to find the maximum point of $F(\nu_3, \bar{\nu}_3)$ under the restrictions $\nu_3 < \kappa$ and $\bar{\nu}_3 < \overline{\kappa}$. We observe that

$$\begin{aligned} &\frac{\partial}{\partial \bar{\nu}_3} F(\nu_3, \bar{\nu}_3) \\ &\quad = -\frac{4\overline{M}(M+\overline{M}-1)(M+\overline{M}+1)(\mathbb{E}[\tilde{a}_\Sigma^2]\gamma^2\kappa^2 + M^2 t\epsilon^2\overline{\kappa}^2)\Big(M(1+M+\overline{M})\bar{\nu}_3 + \overline{M}(1+M+\overline{M})\nu_3 - 2\overline{M}\kappa - 2M\overline{\kappa}\Big)}{\Big(\overline{M}\big(2(M+\overline{M})\kappa - (1+M+\overline{M})\nu_3\big) + M\big(2(M+\overline{M})\overline{\kappa} - (1+M+\overline{M})\bar{\nu}_3\big)\Big)^3}. \end{aligned} \tag{A.17}$$

Because the denominator in (A.17) is positive by (3.13), we have

$$\frac{\partial F}{\partial \bar{\nu}_3}(\nu_3, \bar{\nu}_3) > 0, \tag{A.18}$$

for $(\nu_3, \bar{\nu}_3)$ satisfying

$$\nu_3 < \frac{2\kappa}{1 + M + \overline{M}} - \frac{M(M + \overline{M} - 1)\overline{\kappa}}{\overline{M}(1 + M + \overline{M})} \quad \text{and} \quad \bar{\nu}_3 < \overline{\kappa}. \tag{A.19}$$

Therefore,

$$\sup_{\bar{\nu}_3 < \overline{\kappa}} F(\nu_3, \bar{\nu}_3) = F(\nu_3, \overline{\kappa}) \quad \text{when} \quad \nu_3 < \frac{2\kappa}{1 + M + \overline{M}} - \frac{M(M + \overline{M} - 1)\overline{\kappa}}{\overline{M}(1 + M + \overline{M})}. \tag{A.20}$$

Long but elementary computations produce

$$\frac{\partial}{\partial \nu_3} F(\nu_3, \overline{\kappa}) > 0, \quad \text{for} \quad \nu_3 < \frac{2\kappa}{1 + M + \overline{M}} - \frac{M(M + \overline{M} - 1)\overline{\kappa}}{\overline{M}(1 + M + \overline{M})}. \tag{A.21}$$

From (A.20) and (A.21), we conclude that the maximum point of F satisfies the inequality $\nu_3 \geq \frac{2\kappa}{1+M+\overline{M}} - \frac{M(M+\overline{M}-1)\overline{\kappa}}{\overline{M}(1+M+\overline{M})}$, which is equivalent to

$$\frac{2\overline{M}\kappa + 2M\overline{\kappa} - \overline{M}(1+M+\overline{M})\nu_3}{M(1+M+\overline{M})} \leq \overline{\kappa}. \tag{A.22}$$

Under the restriction in (A.22), the derivative (A.17) implies that we have

$$\sup_{\bar{\nu}_3<\overline{\kappa}} F(\nu_3, \bar{\nu}_3) = F(\nu_3, \frac{2\overline{M}\kappa + 2M\overline{\kappa} - \overline{M}(1+M+\overline{M})\nu_3}{M(1+M+\overline{M})}). \tag{A.23}$$

Taking the derivative with respect to ν_3 of F in (A.16) produces

$$\begin{aligned} &\frac{\partial}{\partial \nu_3} F(\nu_3, \tfrac{2\overline{M}\kappa+2M\overline{\kappa}-\overline{M}(1+M+\overline{M})\nu_3}{M(1+M+\overline{M})}) \\ &\quad = -\frac{4(M\sum_{i=1}^M \mathbb{E}[\tilde{a}_i^2]-\mathbb{E}[\tilde{a}_\Sigma^2])(M+\overline{M}-1)(1+M+\overline{M})\gamma^2\kappa^2\Big((1+M+\overline{M})\nu_3-2\kappa\Big)}{M\Big(2(M+\overline{M})\kappa-(1+M+\overline{M})\nu_3\Big)^3}. \end{aligned} \tag{A.24}$$

Because of (A.14), we can conclude that $F(\nu_3, \frac{2\overline{M}\kappa+2M\overline{\kappa}-\overline{M}(1+M+\overline{M})\nu_3}{M(1+M+\overline{M})})$ is maximized at $\nu_3^* = \frac{2\kappa}{1+M+\overline{M}}$, and the corresponding $\bar{\nu}_3^*$ is $\frac{2\overline{\kappa}}{1+M+\overline{M}}$. This maximizer is unique by the previous arguments.

Finally, inserting $(\nu_3^*, \bar{\nu}_3^*)$ into F in (A.16) produces the total welfare

$$\sup_{\nu_3(t),\bar{\nu}_3(t)} \sum_{i=1}^{M+\overline{M}} \mathbb{E}[\mathrm{CE}_i] = -\int_0^1 \frac{\kappa(t)\overline{\kappa}(t)\big(\gamma(t)^2\mathbb{E}[\tilde{a}_\Sigma^2]+\epsilon^2\overline{M}^2 t\big)}{\overline{M}\kappa(t)+M\overline{\kappa}(t)}dt, \tag{A.25}$$

the optimal holding strategies in (4.5), and the common drift in (4.6). □

Proof of Theorem 5.2 This proof is similar to the proof of Theorem 3.2, so here we only outline the key difference. Itô's product rule produces the dynamics of the right-hand-side of (5.4) to be

$$\begin{aligned} &h'(t)\tilde{a}_\Sigma dt + dD_t + \zeta'(t)\epsilon B_t dt + \zeta(t)\epsilon dB_t + \tilde{a}_\Sigma\big(\sigma(t)d\gamma_t + \gamma_t\sigma'(t)dt\big) \\ &\quad = \Big(h'(t)\tilde{a}_\Sigma + \zeta'(t)\epsilon B_t + \tilde{a}_\Sigma\big(\sigma(t)\frac{1-\gamma_t}{1-t} + \sigma'(t)\gamma_t\big)\Big)dt + dD_t + \zeta(t)\epsilon dB_t + \sigma(t)\tilde{a}_\Sigma dZ_t \\ &\quad = \hat{\mu}_t dt + dD_t + \zeta(t)\epsilon dB_t + \sigma(t)\tilde{a}_\Sigma dZ_t, \end{aligned} \tag{A.26}$$

where the last equality uses the ODEs (5.5). The terminal conditions in (5.5) produce the terminal price restriction (1.8).

By computing the derivative in formula (5.9), we see that the ODE for σ in (5.5) holds. The zero terminal condition for σ follows from

$$|\sigma(t)| \leq \frac{1}{1-t}\int_t^1 (1-u)|\mu_1(u)|du \leq \int_t^1 |\mu_1(u)|du, \quad t \in [0,1),$$

which converges to zero as $t \uparrow 1$ because $\mu_1(u)$ is integrable over $u \in [0, 1]$. To see that $\frac{\sigma(t)}{1-t}$ is integrable, the representation (5.9) produces

$$\begin{aligned}\int_0^1 \left|\frac{\sigma(t)}{1-t}\right| dt &\leq \int_0^1 \int_t^1 \frac{1-u}{(1-t)^2} |\mu_1(u)| du dt \\ &= \int_0^1 u|\mu_1(u)| du < \infty,\end{aligned} \tag{A.27}$$

where the equality uses Tonelli's theorem. □

Proof of Theorem 5.4 This proof is similar to the proofs of Theorems 3.2 and 5.2, so here we only outline the key difference. Itô's product rule produces the dynamics of the right-hand-side of (5.4) to be

$$\begin{aligned}&\tilde{a}_\Sigma h'(t)dt + dD_t + \zeta'(t)\epsilon B_t dt + \zeta(t)\epsilon dB_t + \tilde{a}_\Sigma\big(\sigma(t)d\gamma_t + \sigma'(t)\gamma_{t-}dt\big) \\ &= \Big(h'(t)\tilde{a}_\Sigma + \zeta'(t)\epsilon B_t + \tilde{a}_\Sigma\big(\sigma(t)\tfrac{1-\gamma_{t-}}{1-t} + \sigma'(t)\gamma_{t-}\big)\Big)dt \\ &+ dD_t + \zeta(t)\epsilon dB_t + \sigma(t)\tilde{a}_\Sigma\big(d\gamma_t - \tfrac{1-\gamma_{t-}}{1-t}dt\big) \\ &= \hat{\mu}_t dt + dD_t + \zeta(t)\epsilon dB_t + \sigma(t)\tilde{a}_\Sigma\big(d\gamma_t - \tfrac{1-\gamma_{t-}}{1-t}dt\big),\end{aligned} \tag{A.28}$$

where the last equality uses the ODEs (5.5). The terminal conditions in (5.5) produce the terminal price restriction (1.8). □

References

1. Almgren, R., Chriss, N.: Value under liquidation. Risk **12**, 61–63 (1999)
2. Almgren, R., Chriss, N.: Optimal execution of portfolio transactions. J. Risk **3**, 5–39 (2000)
3. Baldauf, M., Frei, C., Mollner, J.: Principal trading arrangements. Working paper, University of Alberta (2020)
4. Barardehi, Y., Bernhardt, D.: Re-appraising intraday trading patterns: what you didn't know you didn't know. Working paper, Chapman University (2018)
5. Berkowitz, S., Logue, D., Noser, E.: The total cost of transactions on the NYSE. J. Finance **43**, 97–112 (1988)
6. Bertsimas, D., Lo, A.: Optimal control of execution costs. J. Financ. Mark. **1**, 1–50 (1998)
7. Bouchard, B., Fukasawa, M., Herdegen, M., Muhle-Karbe, J.: Equilibrium returns with transaction costs. Finance Stoch. **22**, 569–601 (2018)
8. Boulatov, A., Bernhardt, D., Larionov, I.: Predatory and defensive trading in a dynamic model of optimal execution by multiple traders. Working paper, University of Illinois (2016)
9. Brunnermeier, M.K., Pedersen, L.H.: Predatory trading. J. Finance **60**, 1825–1863 (2005)
10. Carlin, B.I., Lobo, M.S., Viswanathan, S.: Episodic liquidity crises: cooperative and predatory trading. J. Finance **62**, 2235–2274 (2007)
11. Choi, J.H., Larsen, K., Seppi, D.: Information and trading targets in a dynamic market equilibrium. J. Financ. Econ. **132**, 22–49 (2018)
12. Cvitanić, J., Cuoco, D.: Optimal consumption choices for a "large" investor. J. Econ. Dyn. Control **22**, 401–436 (1998)
13. Domowitz, I., Yegerman, H.: The cost of algorithmic trading: a first look at comparative performance. In: Bruce, B.R. (ed.) Algorithmic Trading: Precision, Control, Execution, pp. 30–40. Institutional Investor, London (2005)
14. Du, S., Zhu, H.: What is the optimal trading frequency in financial markets? Rev. Econ. Stud. **84**, 1606–1651 (2017)
15. Duffie, D.: Dynamical Asset Pricing Theory, vol. 3. Princeton University Press, Princeton (2001)
16. Émery, M., Yor, M.: A parallel between Brownian bridges and gamma bridges. Publ. Res. Inst. Math. Sci. **40**, 669–688 (2004)
17. Engle, R., Ferstenberg, R., Russell, J.: Measuring and modeling execution cost and risk. J. Portf. Manag. **38**, 14–28 (2012)

18. Financial Insights: Marching up the learning curve: the second buy side algorithmic trading survey. Financial Insights special report (2006)
19. Frei, C., Westray, N.: Optimal execution of a VWAP order: a stochastic control approach. Math. Finance **25**, 612–639 (2015)
20. Gârleanu, N., Pedersen, L.H.: Dynamic portfolio choice with frictions. J. Econ. Theory **165**, 487–516 (2016)
21. Garman, M.B.: Market microstructure. J. Financ. Econ. **3**, 257–275 (1976)
22. Gatheral, J., Schied, A.: Optimal trade execution under geometric Brownian motion in the Almgren and Chriss framework. Int. J. Theor. Appl. Finance **14**, 353–368 (2011)
23. Gatheral, J., Schied, A.: Dynamical models for market impact and algorithms for optimal order execution. In: Fouque, J.-P., Langsam, J. (eds.) Handbook on Systemic Risk, pp. 579–602. Cambridge University Press, Cambridge (2013)
24. Grossman, S.J., Miller, M.H.: Liquidity and market structure. J. Finance **43**, 617–633 (1988)
25. Grossman, S.J., Stiglitz, J.E.: On the impossibility of informationally efficient markets. Am. Econ. Rev. **70**, 393–408 (1980)
26. Hagströmer, B., Nordén, L.: The diversity of high-frequency traders. J. Financ. Mark. **16**, 741–770 (2013)
27. Hasbrouck, J.: Measuring the information content of stock trades. J. Finance **46**, 179–207 (1991)
28. Karatzas, I., Shreve, S.E.: Methods of Mathematical Finance. Springer, New York (1998)
29. Korajczyk, R.A., Murphy, D.: High-frequency market making to large institutional trades. Rev. Financ. Stud. **32**, 1034–1067 (2019)
30. Kyle, A.: Continuous auctions and insider trading. Econometrica **53**, 1315–1336 (1985)
31. Madan, D., Carr, P., Chang, E.: The variance gamma process and option pricing. Eur. Finance Rev. **2**, 79–105 (1998)
32. Madhavan, A.: VWAP strategies. Trans. Perform. Spring **2002**, 32–38 (2002)
33. McInish, T., Wood, R.: An analysis of intraday patterns in Bid/Ask spreads for NYSE stocks. J. Finance **47**, 753–764 (1999)
34. Menkveld, A.: High frequency trading and the new market makers. J. Financ. Mark. **16**, 712–740 (2013)
35. O'Hara, M.: High frequency market microstructure. J. Financ. Econ. **116**, 257–270 (2015)
36. Perold, A.F.: The implementation shortfall: paper versus reality. J. Portf. Manag. **14**, 4–9 (1988)
37. Predoiu, S., Shaikhet, G., Shreve, S.E.: Optimal execution in a general one-sided limit-order book. SIAM J. Financ. Math. **2**, 183–212 (2011)
38. Protter, P.E.: Stochastic Integration and Differential Equations, vol. 2. Springer, New York (2004)
39. Sannikov, Y., Skrzypacz, A.: Dynamic trading: price inertia and front-running. Working paper, Stanford University Graduate School of Business (2016)
40. Stoll, H.R.: The pricing of security dealer services: an empirical study of NASDAQ stocks. J. Finance **33**, 1153–1172 (1978)
41. U.S. Securities and Exchange Commission: concept release on equity market structure #34-61358 (2010)
42. van Kervel, V., Kwan, A., Westerholm, J.: Order splitting and searching for a counterparty. Working paper, Pontificia Universidad Católica de Chile (2018)
43. van Kervel, V., Menkveld, A.: High-frequency trading around large institutional orders. J. Finance **74**, 1091–1137 (2018)
44. Vayanos, D.: Strategic trading and welfare in a dynamic market. Rev. Econ. Stud. **66**, 219–254 (1999)
45. Weller, B.: Intermediation chains and specialization by speed: evidence from commodity futures markets. Working paper, University of Chicago (2013)