Optimal Turnover, Liquidity and Autocorrelation

Gordon Ritter

joint work with Bastien Baldacci and Jerome Benveniste

October 24, 2022

New York University	Courant Institute of Mathematics
Columbia University	Dept of Mathematics / MAFN
Baruch College	Department of Mathematics and MFE program
University of Chicago	Financial Mathematics program
Ritter Alpha, LP	President, Chief Investment Officer

https://cims.nyu.edu/~ritter/

https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3281235

This presentation is confidential and intended solely for attendees of The IAQF/Thalesians, and such persons' professional advisors. This presentation may not be reproduced or distributed. This presentation is for informational purposes only and does not constitute an offer to sell, or a solicitation of an offer to buy, any security or instrument in or to participate in any trading strategy with the presenter or in any fund or account advised by the presenter (each, a "fund"). If such offer is made, it will only be made by means of a confidential offering memorandum, which would contain material information (including certain risks of investment and conflicts of interest) not discussed in this presentation and which would supersede and which would qualify in its entirety the information discussed in this presentation. Any decision to invest in a fund should be made after reviewing the relevant confidential offering memorandum, conducting such investigations as the investor deems necessary and consulting the investor's own legal, accounting and tax advisors in order to make an independent determination of the suitability and consequences of an investment. No representation or warranty, express or implied, as to the accuracy or completeness of the information discussed in this presentation is made. It is believed to be reasonably accurate and current for the purposes of the illustrations provided but has not been independently verified. There can be no assurance that a fund's objectives will be achieved or that a fund will avoid substantial or complete losses. An investment in a fund is speculative and involves substantial risk of loss. An investment should only be considered by persons who can afford a loss of their entire investment. Past performance is not necessarily indicative of future results. Any statements herein regarding the viability or likelihood of success of trading techniques or approaches, future events or other similar statements constitute only subjective views, are based upon expectations or beliefs, should not be relied on, are subject to change due to a variety of factors and involve inherent risks and uncertainties, both general and specific, many of which cannot be predicted or quantified. Actual results could differ materially from those discussed, contemplated by or underlying the statements made during this presentation. In light of these risk and uncertainties, there can be no assurance that these statements are or will prove to be accurate or complete in any way.

Introduction

- The steady-state turnover of a trading strategy is of clear interest to practitioners and portfolio managers, as is the steady-state Sharpe ratio.
- In this article, we show that in a convenient Gaussian process model, the steady-state turnover can be computed explicitly, and obeys a clear relation to the liquidity of the asset and to the autocorrelation of the alpha forecast signals.
- The steady-state optimal turnover is given by

$$\gamma\sqrt{n+1}$$

where γ is a liquidity-adjusted notion of risk-aversion, and n is the ratio of mean-reversion speed to γ.

- The work that we shall discuss was published in the June 2022 issue of Risk magazine.

  https://www.risk.net/risk-magazine/jun-2022

  https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4018447

  An earlier preprint may be found on SSRN.

- One of the central problems faced by institutional investment managers is the proper management of trading costs.
- For very active strategies, the most significant source of trading costs is usually market impact; see among others Almgren et al. (2005) and Bouchaud (2009).
- The manager may reduce overall market impact cost by reducing turnover, but this potentially also reduces the manager's ability to monetize time-sensitive trading opportunities.

- The proper level of turnover of a real strategy can depend on many factors.
- Nevertheless, in a class of Gaussian process models with linear price impact[1], we derive a simple explicit relation (2.15) between the optimal turnover, the autocorrelation of the trading signal, the investor's risk-aversion, and the liquidity and volatility of the underlying asset.

[1] This class of models has been the subject of an intensive literature, see among others Almgren and Chriss (2001) and Gârleanu and Pedersen (2016) and the reference books Cartea, Jaimungal, and Ricci (2014) and Guéant (2016)

- Our explicit results, while mathematically elegant, only apply to the case of linear price impact, leading to quadratic total cost.
- Empirical studies support the conclusion that the price impact of large orders is in fact proportional to the square root of the order's participation; see Tóth et al. (2011) and references therein.
- Nevertheless, we join Gârleanu and Pedersen (2013) in the belief that it is worthwhile to develop explicit formulas which apply to the linear case, because they can be used as heuristics, approximations to more complex models, or to develop bounds.

Historically, one of the first explicit formulas relating approximate investment performance to forecast accuracy was the Grinold (1989) fundamental law of active management, which asserts that

$$\mathrm{IR} \approx \mathrm{IC}\sqrt{N}$$

where

1. IR is the information ratio,
2. IC is the information coefficient defined as the correlation of a single signal to the dependent return, and
3. N is the effective number of independent bets.

- Bets may be independent either because they pertain to statistically independent investments, or because they pertain to different periods of time, or both.
- Unfortunately the Grinold (1989) formula cannot relate the turnover of the strategy to any other meaningful quantity.
- The formula always prefers increasing turnover, because if asset returns are serially independent, then trading more often increases N with no perceived cost.
- Liquidity of the asset is not an input to the formula.

- In order to better understand the fundamental relation between optimal turnover, liquidity, and autocorrelation of alpha signals, we work in a continuous-time stochastic process model.
- Letting x_t denote an investor's holdings of a risky asset at time t, denominated in dollars (or any convenient numeraire), define the *steady-state turnover* as the following limit

$$\lim_{t\to\infty} \frac{\mathbb{E}|\dot{x}_t|}{\mathbb{E}|x_t|} \tag{2.1}$$

where $\dot{x}_t$ denotes the time-derivative.

- The steady-state turnover of the *optimal* strategy is of clear interest to practitioners and portfolio managers, as is the steady-state Sharpe ratio.
- In what follows, we show that in a convenient Gaussian process model, the steady-state turnover can be computed explicitly, and obeys a clear relation to the liquidity of the asset and to the autocorrelation of the alpha forecast signals.
- Thus, we contribute to the literature on optimal execution, in the spirit of Cartea and Jaimungal (2016) and Lehalle and Neuman (2019), by providing closed-form expressions of important trading metrics in a general linear-quadratic framework.

- Let λ denote the linear price impact coefficient in the tradition of Kyle (1985), let κ denote the investor's absolute risk-aversion, and let σ denote the instantaneous volatility.
- In the special case of an Ornstein-Uhlenbeck model with mean-reversion speed ϕ, we find that steady-state optimal turnover is given by

$$\sqrt{\frac{\sigma\left(\phi\sqrt{\kappa\lambda}+\kappa\sigma\right)}{\lambda}}.$$

- Before deriving these explicit formulas, we present Theorem 2.1, which is arguably the most general result on optimal trading strategies for quadratic trading-cost models in the setting of a mean-quadratic-variation objective.

A very general result on quadratic costs

- Presently we generalize the well-known result of Gârleanu and Pedersen (2016), which assumed the return-predictor process is a Markovian jump diffusion, to any square-integrable process.
- Our theorem also generalizes the main result of Almgren and Chriss (2001), and thus has applications to optimal execution of algorithmic orders.
- Such a general result is of intrinsic interest in its own right, but we need it specifically later in this paper, to derive results on steady-state optimal turnover and strategy performance.
- The proof is also of interest, as it shows that the problem can be converted into a convex optimization problem in an infinite-dimensional space, and is thus amenable to standard convex optimization techniques as per Baldacci and Benveniste (2020).

- Fix a probability measure space $(\Xi, \mathbb{P})$ where $\mathbb{P}$ is a probability measure on a σ-algebra of events.
- The outcomes in this space are various possible trajectories of the market and of our trading within it.
- The filtration

$$\mathcal{F} = \{\mathcal{F}_t : t \geq 0\}$$

  denotes, as usual, the information that is available at time t.
- An *adapted process* is a stochastic process x such that each x_t is $\mathcal{F}_t$-measurable.
- Since it is used often, we define a notation for the conditional expectation:

$$\mathbb{E}_t[Y] := \mathbb{E}[Y \mid \mathcal{F}_t],$$

  where Y is any random variable.

- Let $\mathcal{H}$ denote the usual real Hilbert space of all adapted $\mathbb{R}^N$-valued processes on $[0, \infty)$ which are integrable in the mean-square sense:

$$\mathbb{E} \int_0^\infty ||\mu_t||^2 dt < \infty.$$

- Also let $\mathcal{A}$ be the Sobolev space of $\mathbb{R}^N$-valued adapted processes that are almost surely differentiable and whose derivative lies in $\mathcal{H}$.
- We use the symbol $\cdot$ for the Euclidean inner product in $\mathbb{R}^N$, and $\langle \, , \, \rangle$ for the standard inner product on $\mathcal{H}$:

$$\langle x, y \rangle = \mathbb{E} \int_0^\infty x_t \cdot y_t \, dt \, .$$

- In the following, the positive-definite matrix Ω denotes the quadratic co-variation of the asset return process; with this interpretation, the risk term in our objective function is integrated instantaneous variance.
- This corresponds with the treatment of integrated variance in Almgren and Chriss (2001), and also corresponds directly to the objective function of Gârleanu and Pedersen (2016) in the zero-discounting case.

- Let $\mu \in \mathcal{H}$. let Ω be a positive-definite $N \times N$ covariance matrix.
- Let Λ be a positive-definite $N \times N$ market-impact coefficient matrix.
- Let $\kappa > 0$ be a risk-aversion coefficient.

Theorem 2.1

For any $x_0 \in \mathbb{R}^N$*, there is a unique solution to the optimization problem*

$$\max_{x \in \mathcal{A};\, x(0)=x_0} \mathbb{E} \int_0^\infty \left[\mu_t \cdot x_t - \frac{1}{2} \dot{x} \cdot \Lambda \dot{x} - \frac{\kappa}{2} x_t \cdot \Omega x_t \right] dt, \qquad (2.2)$$

given by the solution, with boundary condition $x(0) = x_0$*, to the stochastically-forced ODE system*

$$\dot{x}_t = -\Gamma x_t + b_t \qquad (2.3)$$

where

$$\Gamma = (\kappa \Lambda^{-1} \Omega)^{1/2} \text{ and} \qquad (2.4)$$

$$b_t = \int_t^\infty e^{\Gamma(t-s)} \Lambda^{-1} \mathbb{E}_t \mu_s \, ds. \qquad (2.5)$$

- The objective in (2.2) may be unbounded if

$$\mathbb{E}\int_0^\infty ||\mu_s||^2\, ds = \infty$$

  as will be the case in many examples of interest, including stationary processes.
- Nevertheless, under mild growth conditions on μ which are satisfied by most realistic forecasts, an investor with an infinite horizon and no discounting will arrive at a well-defined investment strategy by using (2.3), or equivalently by considering finite-horizon problems and letting the horizon tend to infinity.

- Note b_t plays a role analogous to the classical concept of a forcing function (a function which depends on time, but not on x).
- The forcing function is stochastic because b_t depends on

$$\mathbb{E}_t \mu_s,$$

which is a conditional expectation in the filtration, hence a stochastic process.
- One may wonder whether $\Lambda^{-1}\Omega$ actually possesses a real matrix square root, which is required for Γ to be real.
- Fortunately

$$\begin{aligned}\Gamma &:= (\kappa\Lambda^{-1}\Omega)^{1/2} \\ &= \kappa\Lambda^{-1/2}\left(\Lambda^{-1/2}\Omega\Lambda^{-1/2}\right)^{1/2}\Lambda^{1/2}.\end{aligned}$$

so that the only matrix that we take the square root of, is symmetric and positive semi-definite.

Relation to Garleanu and Pedersen

- Theorem 2.1 represents a strict generalization of Gârleanu and Pedersen (2016) Proposition 1, to predictor processes which are not necessarily Markov, and to the case of zero discounting.
- For Markov predictors with zero discounting, our solution agrees with Gârleanu and Pedersen (2016) exactly.

- The Gârleanu and Pedersen (2016) solution is given in terms of matrices known as $\bar{M}^{\text{rate}}$ and b which have relatively complicated expressions, due to discounting.
- In Gârleanu and Pedersen (2016) equations (9)-(10), they express

$$\bar{M}^{\text{rate}} = \Lambda^{-1} A_{xx}, \quad \text{where}$$

$$A_{xx} = -\frac{\rho}{2}\Lambda + \Lambda^{1/2}\left(\kappa\Lambda^{-1/2}\Omega\Lambda^{-1/2} + \frac{\rho^2}{4} I\right)^{1/2}\Lambda^{1/2}$$

$$b := \kappa A_{xx}^{-1}\Omega$$

where ρ is the discount factor that GP inserted under the integral in (2.2).

- This expression becomes much simpler in the limit $\rho \to 0$.

- In our framework, which explicitly works in the limit $\rho \to 0$, $\bar{M}^{\text{rate}}$ and b simplify to

$$\bar{M}^{\text{rate}} = b = \Gamma.$$

- These expressions are quite a bit simpler than G-P's version, because of the zero discounting.
- It turns out that the discount factor, which most portfolio managers probably view as arbitrary, is responsible for a large complication in the resulting expressions.

- We then have from (2.5),

$$\begin{aligned}\bar{M}^{\mathrm{aim}} &= \Gamma^{-1}\Lambda^{-1}\int_t^{\infty} e^{\Gamma(t-s)}\mathbb{E}_t\mu_s\,ds \\ &= \Gamma\int_t^{\infty} e^{\Gamma(t-s)}\mathbb{E}_t\left[(\kappa\Omega)^{-1}\mu_s\right]\,ds\end{aligned}$$

- The latter expression agrees with equation (12) of Gârleanu and Pedersen (2016), who denote the inner matrix by Markowitz_t.

- Theorem 2.1 enables researchers to derive explicit optimal trading strategies for a wide range of models by relating them to equivalent ODE systems.
- Moreover, the maximization in (2.2) is not simply over static, predetermined trading plans, but over the substantially larger class of all admissible stochastic processes.

- This contributes to the literature on optimal execution in a linear-quadratic framework in the spirit of Lehalle and Neuman (2019).
- In our case, Theorem 2.1 facilitates the derivation of simple explicit formulas for the steady-state turnover and information ratio of the optimal strategy, as we shall present in the coming sections.

Steady-state calculations for Gaussian processes

- We will be interested in detailed calculations for single-asset trading paths.
- On account of (2.3), we are interested in a Gaussian process x defined by the stochastically-forced linear ODE:

$$\dot{x}_t = -\gamma x_t + a_t,$$

for some constant $\gamma > 0$, where a_t is a process satisfying

$$\mathbb{E}_t a_s = e^{-\theta(t-s)} a_t \text{ for all } s > t, \tag{2.6}$$

for some $\theta > 0$.
- It follows from (2.6) that

$$\lim_{s \to \infty} \mathbb{E}_t[a_s] = 0,$$

so the steady-state mean of the process a vanishes, and hence the same must hold for x and $\dot{x}$.

- For simplicity let us assume that the unconditional mean of a_t is zero for all t.
- Under these assumptions, the Gaussian process x_t has certain steady-state properties which can be derived in closed form.
- Since x_t is prototypical solution of an optimal trading problem that we discuss later on, the steady-state properties of x_t are clearly of crucial importance to understanding the optimal steady-state turnover.

- For notational convenience, let us define the following three functions:

$$h(t) = \mathbb{E}a_t x_t, \; g(t) = \mathbb{E}x_t^2, \text{ and } v(t) = \mathbb{E}\dot{x}_t^2. \tag{2.7}$$

- The steady-state values of these functions, denoted by $\bar{h}, \bar{g}, \bar{v}$, can be derived by standard calculations for Gaussian processes.

$$\bar{h} = \frac{c_0}{\theta + \gamma}.$$

$$\bar{g} = \frac{\bar{h}}{\gamma}.$$

$$\bar{v} = \bar{h}\theta$$

where c_0 is the steady-state variance of a_t.

- Since differentiation is a linear operator, the derivative of a Gaussian process is another Gaussian process (see Williams and Rasmussen (2006) section 9.4 and references therein).
- Since $\{x_t\}$ is a Gaussian process, each x_t is a Gaussian random variable, as is each $\dot{x}_t$.
- In our example both x_t and $\dot{x}_t$ have mean zero.

- From classical properties of the mean-zero normal distribution, one then has

$$\mathbb{E}[|x_t|]^2 = (\pi/2)\mathbb{E}[x_t^2],$$

and the same holds for $\dot{x}_t$ with the same constant, $\pi/2$.
- Hence the steady-state turnover, as defined by (2.1), exists and is given by

$$\lim_{t\to\infty} \frac{\mathbb{E}|\dot{x}_t|}{\mathbb{E}|x_t|} = \frac{\sqrt{\bar{v}}}{\sqrt{\bar{g}}} = \sqrt{\theta\gamma}. \tag{2.8}$$

Single-asset trading strategies

- Suppose an investor with risk-aversion $\kappa > 0$ is trading a single financial asset and solves

$$\sup_x \mathbb{E}\left[\int_0^\infty \mu_t x_t - \frac{1}{2}\kappa\sigma^2 x_t^2 - \frac{1}{2}\lambda \dot{x}_t^2 \, dt\right] \tag{2.9}$$

  where x, μ are assumed to satisfy the technical hypotheses in Theorem 2.1, and the supremum is over all such x.
- The use of λ to denote a linear price impact coefficient is inspired by, and consistent with, the notation of Kyle (1985); indeed such a coefficient is often referred to as *Kyle's lambda*.
- We are interested in the steady-state properties of solutions to (2.9), for which the key result is the following corollary of Theorem 2.1.

Corollary 2.2

The optimal trading strategy solving (2.9) is the solution of the ODE with stochastic coefficients:

$$\dot{x}_t = -\gamma x_t + m_t, \tag{2.10}$$

where:

$$\gamma = \sqrt{\frac{\kappa\sigma^2}{\lambda}}; \tag{2.11}$$

$$m_t = \int_t^\infty \lambda^{-1} e^{-\gamma(s-t)} \mathbb{E}_t \mu_s \, ds. \tag{2.12}$$

- Corollary 2.2 is a straightforward consequence of applying Theorem 2.1 to the single-asset case, where $N = 1$ and the $N \times N$ matrices Λ, Ω are just scalars λ, σ^2.
- In particular (2.11) may be seen as the reduction of (2.4).

- Assume now that the alpha-forecast process μ_t is an Ornstein-Uhlenbeck (O-U) process:

$$d\mu_t = -\phi\mu_t dt + \nu \, dW. \qquad (2.13)$$

- In this context ϕ is called the *speed of mean-reversion*, and

$$\ln(2)/\phi$$

  is the half-life.
- Then m_t in (2.12) satisfies the exponential-decay condition (2.6), with $\theta = \gamma + \phi$:

$$\mathbb{E}_t m_s = e^{-(\gamma+\phi)(s-t)} m_t. \qquad (2.14)$$

- Hence the calculations leading to (2.8) apply.
- In particular, the optimal steady-state turnover exists and is given by (2.8) which becomes:

$$\text{optimal turnover} = \gamma\sqrt{\phi/\gamma + 1}. \tag{2.15}$$

- Plugging in (2.11) it is easily verified that (2.15) is equal to

$$\sqrt{\sigma(\phi\sqrt{\kappa\lambda} + \kappa\sigma)/\lambda},$$

and that γ is the scalar version of the Garleanu-Pedersen rate matrix.

- Eq. (2.15), while only strictly valid in the case of quadratic costs, can be used by practitioners as an easily-remembered rule of thumb.
- One can see, e.g. from (2.14), that γ and ϕ each have dimensions of inverse time, and so optimal turnover is also in the same units as γ, inverse time.
- As such it represents the fraction of steady-state book size which is traded in the given time unit.

Practical example

- For example, if

$$\kappa = 10^{-6},$$

then for an asset with $\sigma = 0.01/\text{day}$,

$$\lambda = 10 \text{ bps per } 1\% \text{ of ADV},$$

and average daily volume (ADV) is \$ 10 million, then one has

$$\gamma = 0.1/\text{day}.$$

- If $\phi = 0.2/\text{day}$, corresponding to a half-life about 3.5 days, then $\phi/\gamma = 2$ and optimal turnover (2.15) is about

$$\gamma\sqrt{3} \approx 17.3\%/\text{day}.$$

Steady-state information ratio

- From (2.12) and the fact that

$$\mathbb{E}_t \mu_s = e^{-\phi(s-t)} \mu_t,$$

we get the relation

$$m_t = \lambda^{-1}(\gamma + \phi)^{-1} \mu_t. \tag{2.16}$$

- Thus, using (2.7) with m_t in place of a_t, and also using (2.16) to solve for μ_t, we find that the expected *ex ante* rate of profitability is

$$\mathbb{E}[x_t \mu_t] = \lambda(\gamma + \phi)\mathbb{E}[x_t m_t] = \lambda(\gamma + \phi)h(t) \tag{2.17}$$

- We define the *steady-state information ratio (IR)* as

$$\text{IR} = \lim_{t\to\infty} \frac{\mathbb{E}[x_t\mu_t - (\lambda/2)\dot{x}_t^2]}{\sqrt{\mathbb{E}[\sigma^2 x_t^2]}}$$

- The same calculations that were used in deriving (2.8) now allow us to find, in the special case of the O-U process given in (2.13) for $d\mu_t$, that

$$\text{IR} = \frac{\nu}{2\sigma}\sqrt{\frac{\gamma}{2\phi(\phi+2\gamma)}} \tag{2.18}$$

- In the numerical example from above, with $\phi/\gamma = 2$ one has

$$\text{IR} = \nu/(8\gamma\sigma) \approx 0.5\nu/\sigma.$$

Statistically independent assets

- The extension to a portfolio of N *statistically independent* assets, all governed by mean-reverting dynamics with similar half-lives and liquidity parameters, may be approximated by multiplying (2.18) by $\sqrt{N}$ as per Grinold (1989).
- In general the assets in a portfolio will not be statistically independent, but there are special cases which may be approximated using the independence assumption.

- For example, in a market where the covariance matrix is well-described by a multi-factor model (Ross, 1976), then a strictly factor-neutral strategy may be modeled as trading synthetic "residual assets," which are defined as baskets that realize the factor model's residual returns.
- For example, if one starts with a position in stock i, and then adds an appropriate combination of pure factor portfolios to completely hedge stock i's exposures to all common factors, one has a representation of the i-th residual asset as this hedged basket.

- Residual assets will be statistically independent if the factor model's covariance matrix represents the true covariance of the data-generating process.
- For strategies modeled as dynamic portfolios of residual assets, the approximation of multiplying (2.18) by $\sqrt{N}$ is reasonable.
- Thus one obtains the closed-form generalization of the classic "fundamental law of active management" to a modern context, in which the trader is aware of market impact and acts in a dynamically optimal way with respect to a mean-quadratic-variation objective.

Extensions and future work

- In optimal turnover calculations, liquidity (and hence, price impact) plays a central role; hence, it is natural that in multi-asset extensions, we treat the case of cross-asset impact.
- In this context, a trading model of the form described above is said to optimize with *cross-impact* if the instantaneous utility function features a term of the form

$$\frac{1}{2}\dot{x} \cdot \Lambda \dot{x},$$

  where Λ is not a strictly diagonal matrix.
- If Λ is diagonal we say there is no cross-impact in the model.

- If Λ is not positive semi-definite, then one can identify a trade list with negative cost.
- More generally, various assumptions for the matrix Λ are possible, corresponding to the various no-arbitrage assumptions one might wish to make.
- For example Schneider and Lillo (2019) found that for bounded decay kernels and no dynamic arbitrage, Λ must be symmetric.
- Tomas, Mastromatteo, and Benzaquen (2022) present a cogent summary of the various no-arbitrage assumptions, and empirical evidence thereof.

- For strategies trading a single asset in a Gaussian process background, we derived a simple explicit formula relating the three fundamental variables of trading with alpha signals: autocorrelation, liquidity, and optimal turnover.
- Specifically, optimal turnover is given by

$$\gamma\sqrt{n+1}$$

where γ is a liquidity-adjusted risk-aversion parameter, and n is the ratio of mean-reversion speed to γ.
- This single-asset result is interesting and provides an easily-remembered "rule of thumb," but it is not fully satisfying, since most real trading strategies involve multiple assets which cannot be viewed as statistically independent.
- The existence of a simple, explicit formula governing the single-asset case leaves a gap, where a multi-asset generalization is needed.

- In a future article, which is nearing readiness to submit, we plan to fill that gap by deriving an explicit formula for the steady-state optimal turnover in the portfolio case.
- We assume that the alpha signals have the same mean-reversion rate, denoted ϕ, for each asset being traded.
- Even then, to derive steady-state optimal trading rates for each asset in closed form is non-trivial because the trading rate in asset i is a function of the trading rates in all the other assets due to covariance and cross-impact.

- The key insight that enables further progress is that the stochastic differential equations governing the system partially decouple in a basis of eigenvectors for the matrix

$$\Gamma = (\kappa \Lambda^{-1} \Omega)^{1/2}.$$

- This is a nice basis for almost any calculation one might want to do concerning the behavior at optimality of solutions of the stochastic variational problem, because by Theorem 2.1, those solutions solve the stochastic ODE system

$$\dot{x}_t = -\Gamma x_t + b_t,$$

and we are using a basis in which this system decouples.

- The same technique is sometimes used in other types of analysis of systems of coupled ODEs, where it goes under various names such as "the decoupling technique".
- Nevertheless, its application to optimal portfolio turnover in the context of Theorem 2.1 appears to be novel.
- Perhaps one can think of the eigenvectors of Γ as "liquidity-adjusted principal components" since if Λ were proportional to the identity matrix, they would simply be the principal component directions of covariance Ω.

Almgren, Robert and Neil Chriss (2001). "Optimal execution of portfolio transactions". In: *Journal of Risk* 3, pp. 5–40.

Almgren, Robert et al. (2005). "Direct estimation of equity market impact". In: *Risk* 18.7, pp. 58–62.

Baldacci, Bastien and Jerome Benveniste (2020). "A note on Almgren-Chriss optimal execution problem with geometric Brownian motion". In: *Market Microstructure and Liquidity (to appear)*.

Bouchaud, Jean-Philippe (2009). "Price impact". In: *arXiv preprint arXiv:0903.2428*.

Cartea, Alvaro and Sebastian Jaimungal (2016). "Incorporating order-flow into optimal execution". In: *Mathematics and Financial Economics* 10.3, pp. 339–364.

Cartea, Álvaro, Sebastian Jaimungal, and Jason Ricci (2014). "Buy low, sell high: A high frequency trading perspective". In: *SIAM Journal on Financial Mathematics* 5.1, pp. 415–444.

Gârleanu, Nicolae and Lasse Heje Pedersen (2013). "Dynamic trading with predictable returns and transaction costs". In: *The Journal of Finance* 68.6, pp. 2309–2340.

— (2016). "Dynamic portfolio choice with frictions". In: *Journal of Economic Theory* 165, pp. 487–516.

Grinold, Richard C (1989). "The fundamental law of active management". In: *The Journal of Portfolio Management* 15.3, pp. 30–37.

Guéant, Olivier (2016). *The Financial Mathematics of Market Liquidity: From optimal execution to market making*. Vol. 33. CRC Press.

Kyle, Albert S (1985). "Continuous auctions and insider trading". In: *Econometrica: Journal of the Econometric Society*, pp. 1315–1335.

Lehalle, Charles-Albert and Eyal Neuman (2019). "Incorporating signals into optimal trading". In: *Finance and Stochastics* 23.2, pp. 275–311.

Ross, Stephen A (1976). "The arbitrage theory of capital asset pricing". In: *Journal of economic theory* 13.3, pp. 341–360.

Schneider, Michael and Fabrizio Lillo (2019). "Cross-impact and no-dynamic-arbitrage". In: *Quantitative Finance* 19.1, pp. 137–154.

Tomas, Mehdi, Iacopo Mastromatteo, and Michael Benzaquen (2022). "How to build a cross-impact model from first principles: Theoretical requirements and empirical results". In: *Quantitative Finance* 22.6, pp. 1017–1036.

Tóth, Bence et al. (2011). "Anomalous price impact and the critical nature of liquidity in financial markets". In: *Physical Review X* 1.2, p. 021006.

Williams, Christopher K and Carl Edward Rasmussen (2006). *Gaussian processes for machine learning*. Vol. 2. 3. MIT press Cambridge, MA.