

Tackling nonlinear price impact with linear strategies

Xavier Brokmann¹ | David Itkin²  | Johannes Muhle-Karbe³  | Peter Schmidt¹

¹Qube Research and Technologies, London, UK

²Department of Statistics, London School of Economics and Political Science, London, UK

³Department of Mathematics, Imperial College London, London, UK

Correspondence

Johannes Muhle-Karbe, Department of Mathematics, Imperial College London, London, UK.

Email: j.muhle-karbe@imperial.ac.uk

Abstract

Empirical studies in various contexts find that the price impact of large trades approximately follows a power law with exponent between 0.4 and 0.7. Yet, tractable formulas for the portfolios that trade off predictive trading signals, risk, and trading costs in an optimal manner are only available for quadratic costs corresponding to linear price impact. In this paper, we show that the resulting linear strategies allow to achieve virtually optimal performance also for realistic nonlinear price impact, if the “effective” quadratic cost parameter is chosen appropriately. To wit, for a wide range of risk levels, this leads to performance losses below 2% compared to a numerical algorithm proposed by Kolm and Ritter, run at very high accuracy. The effective quadratic cost depends on the portfolio risk and concavity of the impact function, but can be computed without any sophisticated numerics by simply maximizing an explicit scalar function.

1 | INTRODUCTION

Starting from the seminal work of Markowitz (1952), the central focus of portfolio optimization has been to balance expected returns and risk in an optimal manner. Yet, for strategies that trade actively based on signals that predict future price changes, trading costs can easily erode most or

This is an open access article under the terms of the [Creative Commons Attribution](https://creativecommons.org/licenses/by/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2024 The Author(s). *Mathematical Finance* published by Wiley Periodicals LLC.

all of the gross performance. Accordingly, it is crucial to account for trading costs in an effective manner when designing and implementing such portfolios.

For large funds, the key trading cost is the adverse price impact their trades generate. Large purchases (or sales) push prices up (or down), creating “slippage” relative to the hypothetical “unaffected price” without these trades. The relationship between trade sizes and price impact has been the subject of intensive study. In a number of different contexts, the price impact function is estimated to be concave and well described by a power law with an exponent between 0.4 and 0.7,¹ a relationship often dubbed the “square-root law.” For example, on a proprietary dataset of metaorders, Almgren et al. (2005) estimate a power law exponent of 0.6, which leads to impact costs that scale with an exponent of 1.6 of the trade sizes.

This range of trading costs falls squarely between the two trading-cost specifications for which general analytical results are available. Indeed, proportional costs (induced by bid–offer spreads, for example), formally correspond to the limiting case of a price impact function with exponent zero. In this regime, the effects of the costs are invariant with respect to the size of the portfolio and therefore can be studied using small-cost asymptotics even in very general settings (Kallsen & Muhle-Karbe, 2017; Martin, 2014; Soner & Touzi, 2013). The main finding of this literature is that the optimal policy corresponds to the minimal amount of trading required to remain in a “no trade region” around the frictionless optimizer.

The other tractable limiting regime are quadratic trading costs, which correspond to linear price impact. These lead to linear-quadratic control problems and in turn explicit solutions in rather general contexts (Abeille et al., 2016; Collin-Dufresne et al., 2020; Gârleanu & Pedersen, 2013). In particular, Gârleanu and Pedersen (2013) show that the optimal portfolio is a convex combination of two components. The first is the portfolio inherited from the previous time step, which acts as a slowdown to reduce turnovers and hence transaction costs. The second component is the expected Markowitz portfolio over all future time steps. When trading is costly, one does not directly aim for the frictionless optimal portfolio, but “in front of this moving target” along its expected trajectory, in order to further reduce future transaction costs.

For trading costs of power form with exponents between these two limiting cases, small-cost asymptotics of Guasoni and Weber (2020) and Cayé et al. (2020) suggest that the optimal policy interpolates between the optimal trading rules for linear and quadratic costs: one still trades towards a target portfolio at each time step, but the trading speed is slower the closer one gets to the target position. However, the effects of superlinear price impact costs depend on portfolio size. In particular, the target portfolio in the small-cost limit always equals the frictionless optimizer. Thus, the “aim in front of the target feature” is lost in the approximation, even though this is an important qualitative *and* quantitative feature of the solution for quadratic costs. This illustrates the limits of asymptotic methods for superlinear trading costs.²

Finding the optimal portfolio with nonlinear price impact, therefore, requires numerical methods. For example, Kolm and Ritter (2014) have recently proposed to deploy an algorithm inspired by the classical work of Viterbi (1967). This allows for the numerical calculation of the optimal,

¹ For single orders, compare, for example, Loeb (1983), Hasbrouck (1991), Hasbrouck and Seppi (2001), Lillo et al. (2002), Lillo and Mantegna (2003), and Potters and Bouchaud (2003). For metaorders, see, for example, Almgren et al. (2005), Bershova and Rakhlin (2013), Mastromatteo et al. (2014), Brokmann et al. (2015), Zarinelli et al. (2015), and Frazzini et al. (2018). If quotes are also used in the regression, then the relationship is more linear (Cont et al., 2013).

² Conversely, some results are also available for optimal trading with nonlinear price impact in the limit for zero risk aversion (Hey et al., 2023; Muhle-Karbe et al., 2024). However, the methods employed in these studies are crucially tied to the risk-neutral goal functionals used there.




FIGURE 1 P&Ls for Viterbi algorithm of Kolm and Ritter (2014) (solid blue), linear strategy with quadratic costs matched at the average trade size (dotted green) and quadratic costs chosen to maximize the P&L in the model with nonlinear impact costs (dashed orange). [Color figure can be viewed at wileyonlinelibrary.com]

multiperiod portfolio positions with a nonlinear impact term. However, running this algorithm at high accuracy is computationally expensive. In particular, comparative statics with respect to trading signals or other model parameters cannot be incorporated in an efficient manner, but have to be studied by rerunning the algorithm repeatedly from scratch.

This raises the question whether one can adapt the elegant closed-form expressions of Gârleanu and Pedersen (2013) to the model with nonlinear price impact in some form. In a more general context (with trading constraints and multiple predictors of future returns, for example), an approach in this spirit is proposed by Moallemi and Sağlam (2017). To wit, they show that if one only optimizes over a class of strategies that are linear in current and past realizations of the trading signals, then methods from convex optimization allow to speed up the numerical optimization considerably. In the present paper, we show that one can go one step further and in fact restrict attention to just the simple explicit trading rules of Gârleanu and Pedersen (2013), while barely sacrificing *any* performance relative to the numerical algorithm of Kolm and Ritter (2014).

The key idea is to consider the one parameter family of strategies of Gârleanu and Pedersen (2013), parametrized as a function of the quadratic cost parameter. The corresponding policies lead to substantial performance losses if this “effective” cost parameter is not chosen appropriately, for example, if one matches the trading costs at the expected trade size or uses a statistical estimate obtained by fitting a linear model to trade data (Hey et al., 2024).

However, *if* one acknowledges the nonlinearity of the actual price impact function and chooses the linear policy to take this into account in an optimal manner,³ then the performance losses become negligibly small relative to the optimal policy approximated by the algorithm of Kolm and Ritter (2014) at very high accuracy. Indeed, across a range of (realistic) levels of portfolio risks, we consistently find performance losses well below 2%. This is illustrated in Figure 1, which compares the expected profits net of transaction costs that can be achieved using the algorithm of Kolm and

³ In particular, the effective cost depends on the concavity parameter and the model’s gross Sharpe ratio (compare Figure 4).




FIGURE 2 Distribution of trade sizes for Viterbi algorithm of Kolm and Ritter (2014) (blue) and for linear strategy with quadratic costs chosen to maximize the P&L in the model with nonlinear impact costs (orange). [Color figure can be viewed at wileyonlinelibrary.com]

Ritter (2014) with the optimal linear policy and with a linear policy obtained by matching costs at the average trade size. Across a wide range of different risk levels, the optimal linear policy achieves virtually optimal performance, but does not require any sophisticated numerical methods for its deployment. Instead, the expected P&L and the associated risk of any linear policy can be computed in closed form, which allows to determine the optimal policy by simply maximizing an explicit scalar function.

It is important to note that the best linear strategy achieves this impressive performance despite not precisely matching the optimal trades. Instead, as illustrated in Figure 2, the linear approximation features more medium sized but less small and large trades, in line with the small-cost asymptotics of Guasoni and Weber (2020). Nevertheless, the Sharpe ratios generated by the two policies are virtually the same for realistic parameters, so that the linear policies can serve as simple yet powerful tools to account for nonlinear price impact.⁴

The present paper carries out this analysis for the most tractable setting where a single risky asset is traded based on a single predictive signal with linear AR(1) dynamics, and price impact only has an instantaneous but no transient effect on trade prizes. We focus on this concrete tractable setting here to highlight the simplicity of our approach. The latter clearly has the potential to be applied successfully in more general contexts, but we defer this to future research.

The remainder of the paper is organized as follows. Section 2 introduces our model with predictive signals, risk constraints, and trading costs. Section 3 in turn recapitulates the known results for portfolios without and with quadratic transaction costs. For trading costs of general power form, we also provide closed-form expressions for the performance of the linear strategies that are optimal for quadratic costs. In Section 4, the performance of the (optimal) linear policies is compared to the optimum computed with the algorithm of Kolm and Ritter (2014). For better readability, all derivations are collected in the appendices.

⁴In contrast, for smaller portfolios for which both impact costs and bid-ask spreads have a similar effect, more sophisticated approximations are required, compare Kolm and Ritter (2014).

2 | THE MODEL

2.1 | Predictor and returns

We consider a discrete-time model where two assets are traded at times $t = 0, 1, 2, \dots$. The first asset is safe and its price is normalized to one. The second asset is risky, with returns

$$r_{t+1} = \alpha f_t + \epsilon_{t+1}. \quad (1)$$

These risky returns are partially predictable through a signal (“predictor”) f_t independent from the iid return shocks ϵ_t (which have zero mean $\mathbb{E}[\epsilon_t] = 0$, and stationary variance $\mathbb{E}[\epsilon_t^2] = \sigma^2 > 0$).⁵ For tractability, we assume that the predictor f has linear AR(1) dynamics:

$$f_{t+1} = \rho f_t + \sqrt{1 - \rho^2} \eta_{t+1}. \quad (2)$$

Here, $\rho \in (0, 1)$ describes the persistence of the signal; the noise terms η_t driving it are iid Gaussian with zero mean and unit variance.

Remark 2.1. Models where asset returns are partially predictable through a persistent signal have a rich history, both in the academic literature (Barberis, 2000; Gârleanu & Pedersen, 2013; Kim & Omberg, 1996) and in studies rooted in investment practice (De Lataillade et al., 2012; Grinold, 2006; Martin & Schöneborn, 2011). A wide range of predictive variables have been put forward, ranging from dividend yields (Barberis, 2000), moving averages of past returns (Gârleanu & Pedersen, 2013), to order-flow imbalances (Lehalle & Neuman, 2019).

2.2 | Transaction costs

We now introduce trading costs to the model. To this end, we write q_t for the trader’s holdings after trading at time t and denote the corresponding trade at time t by $\Delta q_t = q_t - q_{t-1}$. As in Almgren and Chriss (2000) and many other studies, transaction costs are modeled by a symmetric convex function $c(\Delta q_t)$ of trades. One commonly used specification is $c(\Delta q) = \lambda_1 |\Delta q|$, which corresponds to the proportional transaction costs (“bid-ask spreads”) studied by Martin and Schöneborn (2011), De Lataillade et al. (2012), and Kallsen and Muhle-Karbe (2017), for example.

For $c(\Delta q) = \lambda_2 |\Delta q|^2$, one obtains the linear instantaneous price impact costs considered by, for example, Grinold (2006), Gârleanu and Pedersen (2013), Abeille et al. (2016), and Collin-Dufresne et al. (2020). This specification is particularly tractable since it leads to linear-quadratic control problems and in turn linear optimal feedback controls, see Section 3.2 below. However, the empirical literature consistently finds trading costs that fall between the linear and quadratic specifications. The typical parametric form is a power function

$$c(\Delta q) = \lambda_p |\Delta q|^p, \quad \text{for } p \in (1, 2),$$

⁵In particular, f_t is an optimal predictor in that ϵ and f are independent. For this reason, we also do not allow ϵ to be auto-correlated in time as one could then construct a better predictor f .

where the empirically relevant case is $p \in [1.4, 1.7]$, compare the references in the introduction. In our concrete examples, our baseline is the point estimate $p = 1.6$ from Almgren et al. (2005).

3 | PORTFOLIO OPTIMIZATION UNDER TRANSACTION COSTS

As in Gârleanu and Pedersen (2013), we consider a trader who chooses their trades Δq_t to maximize expected returns penalized for risk and transaction costs:

$$J_T(\Delta q) = \sum_{t=0}^T \mathbb{E} \left[q_t r_{t+1} - \frac{\gamma \sigma^2}{2} q_t^2 - c(\Delta q_t) \right], \quad (3)$$

for a risk aversion parameter $\gamma > 0$. To obtain stationary optimal policies, we assume the planning horizon lies in the distant future and focus on the ergodic limit

$$J(\Delta q) = \liminf_{T \rightarrow \infty} \frac{1}{T} J_T(\Delta q). \quad (4)$$

3.1 | Solution without transaction costs

Without transaction costs ($c = 0$), the goal functionals (3) and (4) both simplify to a sequence of single-period mean–variance optimization problems. The corresponding optimal holdings in turn are the *Markowitz portfolios*

$$q_t = \frac{\mathbb{E}_t[r_{t+1}]}{\gamma \sigma^2} = \frac{\alpha}{\gamma \sigma^2} f_t = \beta f_t. \quad (5)$$

As the return shocks ϵ_{t+1} are independent from the predictor f_t and the latter has unit variance, it follows that the portfolio returns and variance are

$$\begin{aligned} \mathbb{E}[q_t r_{t+1}] &= \beta \mathbb{E}[f_t(\alpha f_t + \epsilon_{t+1})] = \beta \alpha, \\ \mathbb{V}[q_t r_{t+1}] &= \beta^2 \left(\mathbb{E}[(\alpha f_t^2 + f_t \epsilon_{t+1})^2] - \alpha^2 \right) = \beta^2(\sigma^2 + \alpha^2(\kappa - 1)), \end{aligned}$$

where $\kappa = \mathbb{E}[f_t^4] = 3$ is the kurtosis of the Gaussian predictor f_t . The frictionless Sharpe ratio of the optimal policy in turn is time homogenous:

$$\text{SR} := \frac{\mathbb{E}[q_t r_{t+1}]}{\sqrt{\mathbb{V}[q_t r_{t+1}]}} = \frac{\alpha}{\sqrt{\sigma^2 + \alpha^2(\kappa - 1)}}. \quad (6)$$

Equivalently, $\alpha = \sigma \text{SR} / \sqrt{1 - \text{SR}^2(\kappa - 1)} = \sigma \text{SR} / \sqrt{1 - 2\text{SR}^2}$. A significant but reasonable annual Sharpe ratio is of the order of 2, leading to a daily (one period) Sharpe ratio of $\text{SR} \approx 2/16$. As a consequence, the parameter α in Equation (1) essentially controls the optimal Sharpe ratio:

$$\alpha \approx \sigma \text{SR}.$$

We use this relationship for the calibration of the model parameters in Section 4.1 below.

3.2 | Solution with quadratic costs

With transaction costs, the goal functionals (3) and (4) no longer separate into sequences of single-period problems. Instead, when rebalancing the portfolio is costly, each trade needs to take into account both present and future investment opportunities.

For quadratic transaction costs ($p = 2$), the optimal trades can be determined in closed form (Gârleanu & Pedersen, 2013). The optimal holdings q_t then are a weighted average between the trader's incoming holdings q_{t-1} and an average of the discounted future realizations of the Markowitz portfolio (5). As a consequence, the trader generally does not aim for the current Markowitz portfolio, but "in front of this moving target" along its expected trajectory.

When the predictor has the linear state dynamics (2), this aim portfolio can be computed in closed form; for the ergodic problem (4), the new optimal holdings then become a simple linear combination of the current predictor and the current holdings⁶:

$$q_t^* = K_f f_t + K_q q_{t-1}, \quad t \in \mathbb{N}, \quad (7)$$

where

$$K_f = \frac{\alpha}{2\lambda_2} \frac{1}{1 - \rho + \xi}, \quad K_q = \frac{1}{1 + \xi}, \quad \xi = \frac{\gamma\sigma^2}{4\lambda_2} \left(1 + \sqrt{1 + \frac{8\lambda_2}{\gamma\sigma^2}} \right). \quad (8)$$

The corresponding optimal trades are given by

$$\Delta q_t^* = K_f f_t + (K_q - 1)q_{t-1}, \quad t \in \mathbb{N}. \quad (9)$$

Remark 3.1. A natural generalization of the setting considered in the present study is a model with multiple risky assets. If these are all uncorrelated and trades in each asset do not influence the execution prices of the others then—as in Liu (2004)—the multivariate optimization problem separates into independent univariate problems of the form considered here. This observation is used in Section 4 to motivate parameter values that lead to a realistic tradeoff between alpha signals, risk, and Sharpe ratios net of costs.

When the risky assets are correlated and the impact of trades potentially propagates across assets, the model with linear price impact remains tractable (Gârleanu & Pedersen, 2013), up to dealing with more cumbersome matrix notation. In particular, the approximation from Section 3.3 could in principle be extended. However, there does not seem to be a generic nonlinear price impact model for multiple assets and it is also computationally challenging to deploy the numerical solver of Kolm and Ritter (2014) in this context to compute a benchmark. We, therefore, defer an analysis of this important extension for future research.

⁶ Formally, this corresponds to the limit of a zero time discount rate in the results of Gârleanu and Pedersen (2013); for the convenience of the reader, we provide a compact but self-contained derivation of this result in Appendix A.

3.3 | Approximation with power costs

When the trading costs are of power form $c(\Delta q) = \lambda_p |\Delta q|^p$ but not quadratic in line with empirical estimates, then closed-form solutions for the optimization problem (4) are only available in the limiting cases where either trading costs or risk aversion are small; compare Guasoni and Weber (2020), Cayé et al. (2020), and Hey et al. (2023). The optimal policy, therefore, generally needs to be computed numerically. In particular, changes in the signal parameters or risk aversion levels cannot be handled efficiently and the numerical method must be rerun from scratch for each new model configuration.

To ease this computational burden, one can try to limit the search for the optimal policy to a tractable subclass of strategies. For example, Moallemi and Sağlam (2017) consider policies that are linear in the current and all past values of the predictor. This allows to deploy efficient numerical methods from convex optimization to speed up the numerical optimization in quite general settings.

In the present paper, we go one step further and directly focus on the simple feedback policy (9) from the model with quadratic costs, which only depends linearly on current holdings and the current signal. To apply this policy in the model with subquadratic trading costs, the “effective” quadratic trading cost λ needs to be chosen in an appropriate manner. To this end, we consider the family of feedback policies

$$\Delta q_t(\lambda; \gamma) = K_f(\lambda; \gamma) f_t + (K_q(\lambda; \gamma) - 1) q_{t-1}(\lambda; \gamma). \quad (10)$$

Here K_f and K_q are given as in Equation (8) and we have highlighted the dependence on the effective quadratic cost λ , which we view as a tuning parameter, and the risk aversion parameter γ , which will be chosen to scale the strategy to a given level of risk.

The performance of such linear strategies in the model with general power costs $c(\Delta q) = \lambda_p |\Delta q|^p$, $p \in (1, 2)$ can be quantified explicitly; for better readability, the somewhat tedious derivation of this result is delegated to Appendix B.

Theorem 3.2. *For any $\gamma > 0$, we have*

$$J(\Delta q(\lambda; \gamma)) = \frac{\alpha K_f}{1 - K_q \rho} - \frac{\gamma}{2} \frac{\sigma^2 K_f^2 (1 + K_q \rho)}{(1 - K_q^2 \rho)(1 - K_q \rho)} - \frac{2^{p/2} \lambda_p \nu^p}{\sqrt{\pi}} \Gamma\left(\frac{p+1}{2}\right), \quad (11)$$

where $\Gamma(\cdot)$ is the gamma function and

$$\nu = K_f \left(1 - \frac{2\rho(1 - K_q)}{1 - K_q \rho} + \frac{(1 - K_q)^2 (1 + K_q \rho)}{(1 - K_q^2)(1 - K_q \rho)} \right)^{1/2}.$$

Moreover, the corresponding average one period risk is

$$R(\Delta q(\lambda; \gamma)) = \lim_{T \rightarrow \infty} \sqrt{\frac{1}{T} \sum_{t=0}^T \sigma^2 \mathbb{E} [q_t^2(\lambda; \gamma)]} = \sqrt{\frac{\sigma^2 K_f^2 (1 + K_q \rho)}{(1 - K_q^2)(1 - K_q \rho)}}. \quad (12)$$

(Here, we have dropped the arguments of $K_f = K_f(\lambda; \gamma)$ and $K_q = K_q(\lambda; \gamma)$ throughout to ease notation.)

In view of Theorem 3.2, the performance of the linear policy (10) for any value of the effective quadratic cost λ has a fully explicit expression in terms of the model parameters. This allows to conveniently assess the performance of particular policies, such as empirical point estimates where the cost function is assumed to be quadratic as in Gârleanu and Pedersen (2013) or quadratic costs chosen to match the power costs at the average trade size. More generally, the formulas from Theorem 3.2 can also be used to determine the *optimal* strategy in the linear class (10) by maximizing the explicit right-hand side of Equation (11).

Using expression (12), one can also fix the average portfolio risk to a given level. In the spirit of classical mean-variance optimization, one can in turn maximize the expected returns (now net of transaction costs) for this given level of risk. To this end, for each value λ of the quadratic cost, one first computes the risk aversion $\gamma(\lambda)$ that ensures Equation (12) matches the given risk constraint. Then, plugging $\gamma(\lambda)$ into the first and third terms on the right-hand side of Equation (11) and normalizing by the risk level yields the long-term average Sharpe ratio net of transaction costs for the given level of portfolio risk.

In summary, for transaction costs of general power form, the optimal policy in the linear class (9) can be derived in a straightforward manner without any need for sophisticated numerical optimization or analytical approximations. Moreover, it is equally simple to compare the performance of this strategy to any other linear policy. However, Theorem 3.2 says nothing about how close the optimal linear policy gets to the actual nonlinear optimizer. This is the problem we turn to next.

4 | NUMERICAL RESULTS

We now present the key results of the present paper, which compare the performance of the actual optimal policy and its linear approximations. Here, the actual optimizer is computed using the numerical algorithm of Kolm and Ritter (2014), run at very high accuracy.

4.1 | Model calibration

The model parameters are chosen to match those of a mid- to high-capacity trading system, trading daily a combination of predictors with a half-life of 5 days ($\rho = 0.87$) forecasting returns of a large pool of $N = 500$ assets with an overall annual Sharpe ratio $SR = 3$ (gross of trading costs). Each asset has a daily volatility $\sigma = 0.02$ in line with typical daily volatilities of global equity markets. Asset returns are assumed *uncorrelated*, so that Remark 3.1 applies and the full multivariate optimization problem separates into univariate subproblems. Assuming all assets are homogeneous (i.e., share identical alpha, ADV and volatility parameters), the overall yearly Sharpe ratio of $SR = 3$ is then achieved by setting the parameter α modulating the strength of the predictive signal for each asset to $\alpha = \sigma \times SR / \sqrt{256N} = 1.67 \times 10^{-4}$ (in daily units).

Following Almgren et al. (2005), trades of size Δq move market prices by a relative amount given by $\text{sign}(\Delta q)\eta\sigma|\Delta q/\text{ADV}|^{p-1}$, where η is a “universal coefficient of market impact” (which the authors determine to be $\eta = 0.142$), and the exponent $p = 1.6$ corresponds to price changes

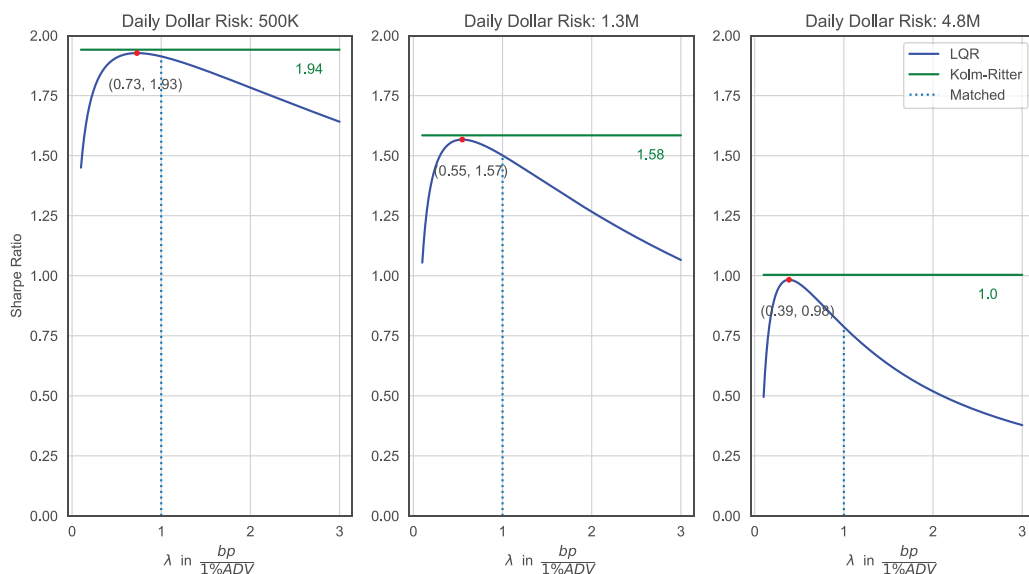


FIGURE 3 Performance of the LQR strategies (blue) for various daily dollar risk values relative to the optimal strategy computed by the Viterbi algorithm (green) of Kolm and Ritter (2014). The dashed line indicates the performance of the linear policy obtained by matching linear and nonlinear costs at the mean absolute trade size. [Color figure can be viewed at wileyonlinelibrary.com]

on the order of 1.8bp for trade sizes on the order of 1% of an asset's Average Daily Volume (ADV). This is in line with standard estimates published in academic papers (Bershova & Rakhlin, 2013; Brokmann et al., 2015; Frazzini et al., 2018; Mastromatteo et al., 2014; Zarinelli et al., 2015) as well as equity broker market reports; the effect of larger trading costs is explored in one of our robustness checks below.

The ADV, which is a proxy for the liquidity of an asset, is on the order of a few 10s of million USD for US equity markets. We assume an ADV of 15.8 million USD for each asset, a conservative yet realistic value (as a lower liquidity leads to higher costs and magnifies any effects related to market impact), resulting in impact costs given by $c(\Delta q) = \lambda_p |\Delta q|^p$ with $p = 1.6$ and $\lambda_p = 1.36 \times 10^{-7}$.

4.2 | Optimal linear policy


Figure 3 displays the Sharpe ratios net of costs for various daily dollar risk levels and the model parameters described in Section 4.1. We observe that the optimal LQR strategy performs remarkably well. Indeed, the loss in the net Sharpe ratio is only between 0.5 and 2% when compared to the KR strategy.

The performance gap depends on the specified daily dollar risk, and increases with this risk target. Moreover, the drop in performance around the optimal effective cost becomes much steeper for larger daily dollar risk values. This emphasizes that, for larger risk thresholds, one needs to be particularly careful both in choosing the model parameters and in performing the corresponding optimizations.

To verify that these qualitative findings hold across a variety of regimes we repeated the experiment with perturbed parameters. Table 1 compares the Sharpe ratio of the KR strategy against

TABLE 1 Comparison of the Kolm–Ritter and best LQR strategy with perturbed parameters.

Perturbed parameter value	Daily dollar risk	Kolm–Ritter Sharpe	Best LQR Sharpe	Percentage loss
$p = 1.52$	2.1M	2.08	2.07	0.52%
$p = 1.68$	0.7M	1.14	1.13	0.96%
$\rho = 0.79$ (Half-life = 3 days)	0.8M	1.27	1.25	1.69%
$\rho = 0.93$ (Half-life = 10 days)	2M	2.01	2.0	0.49%
SR = 2	0.8M	1.0	0.99	0.81%
SR = 4	1.8M	2.18	2.16	1.03%
$\lambda_p = 3.78 \times 10^{-7}$ (5bp cost for 1% ADV trade size)	0.67M	1.12	1.10	1.96%
$\lambda_p = 7.55 \times 10^{-7}$ (10bp cost for 1% ADV trade size)	0.4M	0.85	0.83	2.69%

**FIGURE 4** Effective cost parameter λ^* maximizing the Sharpe ratio net of power costs, plotted against the daily dollar risk for different choices of maximal frictionless Sharpe ratios. [Color figure can be viewed at wileyonlinelibrary.com]

the best LQR strategy when we change (i) the elasticity of the cost function p , (ii) the half-life of the signal determined by ρ , (iii) the Sharpe ratio of the frictionless Markowitz portfolio, and (iv) the magnitude of the nonlinear cost parameter λ_p . All of the experiments were done with the parameters described in Section 4.1 with only one parameter perturbed. We see that across the board, the Sharpe ratios obtained by the KR strategy outperform the best LQR approximation by at most a couple of percent. With larger trading costs λ_p , the gap increases somewhat (as larger impact costs have a similar effect to larger daily dollar risks), but remains moderate. This sensitivity check illustrates the effectiveness of the best LQR strategy even when the strategy is near its capacity.

Next, we explore how the optimal linear cost parameter depends on the risk level. Figure 4 plots the optimal LQR parameter across a range of daily dollar risks and for different frictionless




FIGURE 5 Contour plots for the Kolm–Ritter strategy (left panel) and the best LQR strategy (right panel); contours correspond to trade sizes. [Color figure can be viewed at wileyonlinelibrary.com]

Sharpe ratios. A first observation is that this effective linear cost is of the same order of magnitude across a wide range of risk targets—it only changes by a factor of about 3 when the risk target is increased by a factor of 50. Second, we see that the effective quadratic cost is smaller if either the frictionless Sharpe ratio and the daily dollar risk are large. The intuition is that either case requires more large trades, which are cheaper with subquadratic costs, unless the quadratic cost parameter is reduced accordingly.

4.3 | Comparison to Kolm–Ritter policy

A remarkable feature is that the LQR strategy is able to achieve a similar performance to the KR strategy without actually approximating the distribution of trades. Indeed, Figure 2 in the introduction plots the frequency of trade sizes Δq that the two strategies produce. The LQR distribution is known to be a centered Gaussian, whose variance can be worked out from Equation (10), while the trade distribution of the KR strategy was obtained by simulation. We see that the KR algorithm more frequently prescribes both very small and very large trades relative to the normally distributed LQR strategy. That is, it both has fatter tails and a more concentrated peak, in line with the small-cost asymptotics of Guasoni and Weber (2020).

The difference between the two strategies is further evidenced by the dependence on the current signal and holdings. Figure 5 provides contour plots for the optimal holdings q_t of both the KR and LQR strategies for different values of the current predictor f_t and current holdings q_{t-1} . The trades of the LQR strategy are linear in the current holdings and predictor, which leads to linear and equally spaced contour lines in the right panel of Figure 5. On the other hand, we observe that the contour lines of the KR strategy are not equally spaced (cf. the left panel of Figure 5). The spacing of the lines is wider for smaller values of the predictor or holdings and decreases for larger values. This illustrates that the KR algorithm makes—compared to the LQR strategy—smaller




FIGURE 6 Slices of the optimal trades $\Delta q_t(f_t, \cdot)$ for the Kolm–Ritter strategy (blue) and the optimal LQR strategy (orange) for 1.3M daily dollar risk. Here, the position q_{t-1} is expressed in units of its standard deviation. The main panels display the slice for all position values sampled. The inset plots zoom in to the regions where most realizations occur. [Color figure can be viewed at wileyonlinelibrary.com]

trades for smaller values of q_{t-1} and f_t but larger trades for larger values. This nonlinearity leads to the fat-tailed trade distribution of the KR strategy that we have already observed in Figure 2.

To illustrate this nonlinearity even more clearly, Figure 6 presents the slices of $\Delta q_t(f_t, \cdot)$ for the two strategies when the current predictor is $f_t = -1$, $f_t = 0$ or $f_t = 1$, respectively, that is, one standard deviation below, at, or one standard deviation above its average value. In line with the asymptotic results of Guasoni and Weber (2020) and Cayé et al. (2020), the main panel of Figure 6 shows that the KR strategy prescribes larger trades than its linear approximation when the current position is very large or small. However, unlike in the small-cost asymptotics, this effect becomes asymmetric when the current value of the predictor is nonzero. The intuition is that when both positive holdings and a negative predictor prescribe sales, then the resulting strong trading motive leads to a particularly large difference between the linear and nonlinear policies. In contrast, negative holdings and a negative predictor partially offset, so that the difference between the models matter to a much smaller extent.

In contrast to this behavior for large holdings, the inset panel of Figure 6 shows that when the current holdings are sufficiently close to a “target position” then, the trades of the optimal nonlinear policy are smaller than its linear counterpart. For a zero predictor, this happens when the holdings are close to zero, for positive or negative predictors, the target is shifted up or down. This slowdown of the optimal trading speed near the optimal allocation is reminiscent of the “no trade regions” that are optimal for linear costs, and in turn leads to the higher frequency of small trades depicted in Figure 2.

In summary, the policy computed using the algorithm of Kolm and Ritter (2014) takes into account the nonlinearity of the price impact function in various sensible ways that cannot be

reproduced using the linear approximation from Section 3.3. However, once the performance of the strategies is averaged across states, the difference nevertheless is surprisingly small.

5 | CONCLUSION

This paper studies how well the linear feedback policies that are optimal for quadratic transaction costs approximate the optimal performance in models with transaction costs of general power form in line with empirical studies. When the effective quadratic cost determining the linear policy is chosen naively, this approximation leads to substantial performance losses. But if this choice is optimized for the actual transaction costs, then the performance losses become surprisingly small.

This allows one to avoid using a computationally expensive numerical optimizer. Instead, nearly optimal performance can be achieved by simply optimizing over a one-parameter family of strategies. Thereby the powerful analytical toolbox developed for linear-quadratic models can also be brought to bear on models with nonlinear price impact costs.

In the present paper, we illustrate the power of this method in the simplest context with only instantaneous impact. The addition of impact decay as in Obizhaeva and Wang (2013) is a key direction for future research. Extending the analysis of this paper to a setting with many assets and trading signals is another important but challenging open problem. Finally, another crucial question is to provide a better theoretical understanding of why the approximation by simple linear policies works so well in the present context, and what are the scope and limitations of this approximation for more general settings.

ACKNOWLEDGMENTS

Pertinent comments of an anonymous associate editor and three anonymous referees are gratefully acknowledged.

DATA AVAILABILITY STATEMENT

Data sharing is not applicable to this article as no new data were created or analyzed in this study.

ORCID

David Itkin  <https://orcid.org/0000-0001-8643-574X>

Johannes Muhle-Karbe  <https://orcid.org/0000-0002-9165-2244>

REFERENCES

- Abeille, M., Sérié, E., Lazaric, A., & Brokmann, X. (2016). LQG for portfolio optimization. Preprint.
- Almgren, R., & Chriss, N. (2000). Optimal execution of portfolio transactions. *Journal of Risk*, 3, 5–39.
- Almgren, R., Thum, C., Hauptmann, E., & Li, H. (2005). Direct estimation of equity market impact. *Risk*, 18(7), 57–62.
- Barberis, N. (2000). Investing for the long run when returns are predictable. *Journal of Finance*, 55(1), 225–264.
- Bershova, N., & Rakhlin, D. (2013). The non-linear market impact of large trades: Evidence from buy-side order flow. *Quantitative Finance*, 13(11), 1759–1778.
- Brokmann, X., Sérié, E., Kockelkoren, J., & Bouchaud, J.-P. (2015). Slow decay of impact in equity markets. *Market Microstructure and Liquidity*, 1(2), 1550007.
- Cayé, T., Herdegen, M., & Muhle-Karbe, J. (2020). Trading with small nonlinear price impact. *Annals of Applied Probability*, 30(2), 706–746.

- Collin-Dufresne, P., Daniel, K., & Sağlam, M. (2020). Liquidity regimes and optimal dynamic asset allocation. *Journal of Financial Economics*, *136*(2), 379–406.
- Cont, R., Kukanov, A., & Stoikov, S. (2013). The price impact of order book events. *Journal of Financial Econometrics*, *12*(1), 47–88.
- De Lataillade, J., Deremble, C., Potters, M., & Bouchaud, J.-P. (2012). Optimal trading with linear costs. *Journal of Investment Strategies*, *1*(3), 91–115.
- Frazzini, A., Israel, R., & Moskowitz, T. J. (2018). Trading costs. Preprint.
- Gârleanu, N., & Pedersen, L. H. (2013). Dynamic trading with predictable returns and transaction costs. *Journal of Finance*, *68*(6), 2309–2340.
- Grinold, R. (2006). A dynamic model of portfolio management. *Journal of Investment Management*, *4*(2), 5.
- Guasoni, P., & Weber, M. H. (2020). Nonlinear price impact and portfolio choice. *Mathematical Finance*, *30*(2), 341–376.
- Hasbrouck, J. (1991). Measuring the information content of stock trades. *Journal of Finance*, *46*(1), 179–207.
- Hasbrouck, J., & Seppi, D. J. (2001). Common factors in prices, order flows, and liquidity. *Journal of Financial Economics*, *59*(3), 383–411.
- Hey, N., Bouchaud, J.-P., Mastromatteo, I., Muhle-Karbe, J., & Webster, K. (2024). The Cost of Misspecifying Price Impact. *Risk*, to appear.
- Hey, N., Mastromatteo, I., Muhle-Karbe, J., & Webster, K. (2023). Trading with concave price impact – theory and evidence. Preprint.
- Kallsen, J., & Muhle-Karbe, J. (2017). The general structure of optimal investment and consumption with small transaction costs. *Mathematical Finance*, *27*(3), 659–703.
- Kim, T. S., & Omberg, E. (1996). Dynamic nonmyopic portfolio behavior. *Review of Financial Studies*, *9*(1), 141–161.
- Kolm, P., & Ritter, G. (2014). Multiperiod portfolio selection and Bayesian dynamic models. *Risk*, *18*(7), 50–54.
- Lehalle, C.-A., & Neuman, E. (2019). Incorporating signals into optimal trading. *Finance and Stochastics*, *23*, 275–311.
- Lillo, F., Farmer, J. D., & Mantegna, R. N. (2002). Single curve collapse of the price impact function for the New York stock exchange. Preprint.
- Lillo, F., & Mantegna, R. N. (2003). Master curve for price-impact function. *Nature*, *421*, 129–130.
- Liu, H. (2004). Optimal consumption and investment with transaction costs and multiple risky assets. *Journal of Finance*, *59*(1), 289–338.
- Loeb, T. F. (1983). Trading cost: The critical link between investment information and results. *Financial Analysts Journal*, *39*(3), 39–44.
- Markowitz, H. (1952). Portfolio selection. *Journal of Finance*, *7*(1), 77–91.
- Martin, R. (2014). Optimal trading under proportional transaction costs. *Risk*, *18*(8), 54–59.
- Martin, R., & Schöneborn, T. (2011). Mean reversion pays, but costs. *Risk*, *15*(2), 96–101.
- Mastromatteo, I., Tóth, B., & Bouchaud, J.-P. (2014). Agent-based models for latent liquidity and concave price impact. *Physical Review E*, *89*(4), 23–25.
- Moallemi, C., & Sağlam, M. (2017). Dynamic portfolio choice with linear rebalancing rules. *Journal of Financial and Quantitative Analysis*, *52*(3), 1247–1278.
- Muhle-Karbe, J., Wang, Z., & Webster, K. (2024). Stochastic liquidity as a proxy for nonlinear price impact. *Operations Research*, *72*(2), 444–458.
- Obizhaeva, A., & Wang, J. (2013). Optimal trading strategy and supply/demand dynamics. *Journal of Financial Markets*, *16*(1), 1–32.
- Potters, M., & Bouchaud, J.-P. (2003). More statistical properties of order books and price impact. *Physica A*, *324*(1-2), 133–140.
- Soner, H. M., & Touzi, N. (2013). Homogenization and asymptotics for small transaction costs. *Siam Journal on Control and Optimization*, *51*(4), 2893–2921.
- Viterbi, A. J. (1967). Error bounds for convolutional codes and an asymptotically optimum decoding algorithm. *IEEE Transactions on Information Theory*, *13*(2), 260–269.
- Zarinelli, E., Treccani, M., Farmer, J. D., & Lillo, F. (2015). Beyond the square root: Evidence for logarithmic dependence of market impact on size and participation rate. *Market Microstructure and Liquidity*, *1*(2), 1550004.

How to cite this article: Brokmann, X., Itkin, D., Muhle-Karbe, J., & Schmidt, P. (2024). Tackling nonlinear price impact with linear strategies. *Mathematical Finance*, 1–19. <https://doi.org/10.1111/mafi.12449>

APPENDIX A: DERIVATION OF OPTIMAL QUADRATIC COST STRATEGY

For the convenience of the reader, this appendix provides a compact self-contained derivation of the optimal policy (9) for the ergodic goal function (4) in the case of quadratic costs. To this end, we first consider the finite horizon problem (3). Write $V_t(f_t, q_{t-1})$ for the value function of Equation (3) starting at time t from the initial signal f_t and holdings q_{t-1} before trading at time t . The value function satisfies the Bellman equation

$$V_t(f_t, q_{t-1}) = \sup_{\Delta q_t} \left\{ \alpha f_t (q_{t-1} + \Delta q_t) - \frac{\gamma \sigma^2}{2} (q_{t-1} + \Delta q_t)^2 - \lambda_2 (\Delta q_t)^2 + \mathbb{E}_t[V_{t+1}(f_{t+1}, q_t)] \right\}. \quad (\text{A.1})$$

After the last trading time T , the continuation value is evidently zero, $V_{T+1}(f_{T+1}, q_T) = 0$. With the quadratic ansatz $V_t(f, q) = \frac{1}{2}a(t)f^2 + b(t)fq - \frac{1}{2}c(t)q^2 + w(t)$ (for some deterministic functions $a(t)$, $b(t)$, $c(t)$ and $w(t)$ to be determined), we have

$$\begin{aligned} \mathbb{E}_t[V_{t+1}(f_{t+1}, q_t)] &= \frac{1}{2}a(t+1)\mathbb{E}_t[f_{t+1}^2] + b(t+1)\mathbb{E}_t[f_{t+1}](q_{t-1} + \Delta q_t) - \frac{1}{2}c(t+1)(q_{t-1} + \Delta q_t)^2 + w(t+1) \\ &= \frac{1}{2}a(t+1)(\rho^2 f_t^2 + 1 - \rho^2) + b(t+1)\rho f_t(q_{t-1} + \Delta q_t) - \frac{1}{2}c(t+1)(q_{t-1} + \Delta q_t)^2 + w(t+1). \end{aligned}$$

The pointwise maximizer of Equation (A.1) in turn becomes

$$\Delta q_t = \frac{\alpha + b(t+1)\rho}{\gamma \sigma^2 + 2\lambda + c(t+1)} f_t - \frac{c(t+1) + \gamma \sigma^2}{\gamma \sigma^2 + 2\lambda + c(t+1)} q_{t-1}. \quad (\text{A.2})$$

After plugging this back into the Bellman equation (A.1) and comparing coefficients for the terms proportional to f_t^2 , $f_t q_{t-1}$, q_{t-1}^2 and independent of these state variables, we obtain the following recursive system of Riccati equations:

$$\begin{aligned} a(t) &= a(t+1)\rho^2 + \frac{(\alpha + b(t+1)\rho)^2}{\gamma \sigma^2 + c(t+1) + 2\lambda}, & b(t) &= \frac{2\lambda(\alpha + b(t+1)\rho)}{\gamma \sigma^2 + c(t+1) + 2\lambda}, \\ c(t) &= \frac{2\lambda(\gamma \sigma^2 + c(t+1))}{\gamma \sigma^2 + c(t+1) + 2\lambda}, & w(t) &= \frac{1 - \rho^2}{2} a(t+1) + w(t+1), \end{aligned}$$

complemented by the terminal conditions $a(T+1) = b(T+1) = c(T+1) = w(T+1) = 0$. As the planning horizon T is postponed to infinity, the recursions for b and c (which already pin down the long-run optimal trades) converge to their stationary points, which satisfy

$$b = \frac{2\lambda(\alpha + b\rho)}{\gamma \sigma^2 + c + 2\lambda}, \quad c = \frac{2\lambda(\gamma \sigma^2 + c)}{\gamma \sigma^2 + c + 2\lambda}.$$

First solving for the positive solution of the equation for c and then solving for b yields

$$c = \frac{\gamma\sigma^2}{2} \left(\sqrt{1 + \frac{8\lambda}{\gamma\sigma^2}} - 1 \right), \quad b = \frac{2\alpha\lambda}{\gamma\sigma^2 + c + 2\lambda(1 - \rho)}.$$

Plugging these values into Equation (A.2) and rearranging then leads to the stationary policy (9) from the body of the paper that is optimal for the ergodic problem (4).

APPENDIX B: PROOF OF THEOREM 3.2

In this section, we prove Theorem 3.2. To this end, first note that as $\rho \in (0, 1)$, the predictor f is covariance stationary with constant mean $\mathbb{E}[f_t] = 0$ and autocovariance function $\text{Cov}[f_t, f_{t+k}] = \rho^k$ for $k \in \mathbb{N}$. In particular, the stationary variance of the predictor is $\text{Cov}[f_t, f_t] = 1$.

Next, we solve the difference equation (10) to obtain that

$$q_t(\lambda; \gamma) = K_f(\lambda; \gamma) \left(f_t + \sum_{s=1}^{t-1} K_q(\lambda; \gamma)^{t-s} f_s \right), \tag{B.1}$$

where we recall that $q_0 = 0$. Having the position size in terms of the model parameters and the signal f will allow us to obtain explicit expressions for the ergodic quantities of interest. To simplify the calculations, we assume without loss of generality that f is started at stationarity.

We start by computing the expected frictionless return $\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^T \mathbb{E}[q_t(\lambda; \gamma)r_{t+1}]$. Using Equation (B.1), we see that

$$\begin{aligned} \mathbb{E}[q_t(\lambda; \gamma)r_{t+1}] &= \alpha K_f(\lambda; \gamma) \left(\mathbb{E}[f_t^2] + \sum_{s=1}^{t-1} K_q(\lambda; \gamma)^{t-s} \mathbb{E}[f_t f_s] \right) \\ &= \alpha K_f(\lambda; \gamma) \left(1 + \sum_{s=1}^{t-1} (K_q(\lambda; \gamma)\rho)^{t-s} \right) \\ &= \alpha K_f(\lambda; \gamma) \left(1 + \frac{(K_q(\lambda; \gamma)\rho)^t - K_q(\lambda; \gamma)\rho}{K_q(\lambda; \gamma)\rho - 1} \right). \end{aligned} \tag{B.2}$$

Hence, it follows that

$$\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^T \mathbb{E}[q_t(\lambda; \gamma)r_{t+1}] = \frac{\alpha K_f(\lambda; \gamma)}{1 - K_q(\lambda; \gamma)\rho}. \tag{B.3}$$

Next we compute the risk term and establish Equation (12) in the process. Using Equation (B.1), we have that

$$\mathbb{E}[q_t^2(\lambda; \gamma)] = K_f^2(\lambda; \gamma) \left(\mathbb{E}[f_t^2] + 2 \sum_{s=1}^{t-1} K_q(\lambda; \gamma)^{t-s} \mathbb{E}[f_t f_s] + \mathbb{E} \left[\left(\sum_{s=1}^{t-1} K_q(\lambda; \gamma)^{t-s} f_s \right)^2 \right] \right).$$

The first two terms are handled as in Equation (B.2) and we obtain

$$\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^T \left(\mathbb{E} [f_t^2] + 2 \sum_{s=1}^{t-1} K_q(\lambda; \gamma)^{t-s} \mathbb{E} [f_t f_s] \right) = 1 + \frac{2K_q(\lambda; \gamma)\rho}{1 - K_q(\lambda; \gamma)\rho}. \tag{B.4}$$

It just remains to study the third term. We directly expand the square to see that

$$\begin{aligned} \mathbb{E} \left[\left(\sum_{s=1}^{t-1} K_q(\lambda; \gamma)^{t-s} f_s \right)^2 \right] &= \sum_{s=1}^{t-1} \sum_{r=1}^{t-1} K_q(\lambda; \gamma)^{t-s} K_q(\lambda; \gamma)^{t-r} \mathbb{E} [f_s f_r] \\ &= \sum_{s=1}^{t-1} \sum_{r=1}^{t-1} K_q(\lambda; \gamma)^{t-s} K_q(\lambda; \gamma)^{t-r} \rho^{|s-r|} \\ &= \sum_{s=1}^{t-1} \sum_{r=1}^s K_q(\lambda; \gamma)^{2t-s-r} \rho^{s-r} + \sum_{s=1}^{t-1} \sum_{r=s+1}^{t-1} K_q(\lambda; \gamma)^{2t-s-r} \rho^{r-s}. \end{aligned}$$

Direct computations repeatedly using the geometric sum formula yield

$$\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^T \mathbb{E} \left[\left(\sum_{s=1}^{t-1} K_q(\lambda; \gamma)^{t-s} f_s \right)^2 \right] = \frac{K_q^2(\lambda; \gamma)(1 + K_q(\lambda; \gamma)\rho)}{(1 - K_q^2(\lambda; \gamma))(1 - K_q(\lambda; \gamma)\rho)}. \tag{B.5}$$

Putting Equations (B.4) and (B.5) together we get

$$\begin{aligned} \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^T \sigma^2 \mathbb{E} [q_t^2(\lambda; \gamma)] &= \sigma^2 K_f^2(\lambda; \gamma) \left(1 + \frac{2K_q(\lambda; \gamma)\rho}{1 - K_q(\lambda; \gamma)\rho} + \frac{K_q^2(\lambda; \gamma)(1 + K_q(\lambda; \gamma)\rho)}{(1 - K_q^2(\lambda; \gamma))(1 - K_q(\lambda; \gamma)\rho)} \right) \\ &= \frac{\sigma^2 K_f^2(\lambda; \gamma)(1 + K_q(\lambda; \gamma)\rho)}{(1 - K_q^2(\lambda; \gamma))(1 - K_q(\lambda; \gamma)\rho)}, \end{aligned} \tag{B.6}$$

which proves (12). As a sanity check, we observe that as the parameter $\lambda \rightarrow 0$, this expression converges to $\frac{\alpha^2}{\gamma^2 \sigma^2}$, which is indeed the frictionless position variance.

Finally, we compute the contribution of the cost term $\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^T \lambda_p \mathbb{E} [|\Delta q_t(\lambda; \gamma)|^p]$. We recall from Equation (10) that

$$\Delta q_t(\lambda; \gamma) = K_f(\lambda; \gamma) f_t + (K_q(\lambda; \gamma) - 1) q_{t-1},$$

which is clearly a centered Gaussian random variable because f_t is.

Hence to pin down its distribution, we just need to compute its variance,

$$\mathbb{E} [\Delta q_t^2(\lambda; \gamma)] = K_f^2(\lambda; \gamma) \mathbb{E} [f_t^2] + 2K_f(\lambda; \gamma)(K_q(\lambda; \gamma) - 1) \mathbb{E} [f_t q_{t-1}] + (K_q(\lambda; \gamma) - 1)^2 \mathbb{E} [q_{t-1}^2].$$

We know that $\mathbb{E}[f_t^2] = 1$ and we previously computed $\mathbb{E}[q_t^2(\lambda; \gamma)]$ in the ergodic limit. Hence, it remains to compute the middle term. We have

$$\begin{aligned} \mathbb{E}[f_t q_{t-1}(\lambda; \gamma)] &= K_f(\lambda; \gamma) \left(\mathbb{E}[f_t f_{t-1}] + \sum_{s=1}^{t-2} K_q(\lambda; \gamma)^{t-1-s} \mathbb{E}[f_t f_s] \right) \\ &= K_f(\lambda; \gamma) \left(\rho + \sum_{s=1}^{t-2} K_q(\lambda; \gamma)^{t-1-s} \rho^{t-s} \right) \\ &= K_f(\lambda; \gamma) \left(\rho + \frac{K_q^2(\lambda; \gamma) \rho^2 - K_q^t(\lambda; \gamma) \rho^t}{K_q(\lambda; \gamma) - K_q^2(\lambda; \gamma) \rho} \right). \end{aligned}$$

In the ergodic limit, we get

$$\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^T \mathbb{E}[f_t q_{t-1}(\lambda; \gamma)] = \frac{K_f(\lambda; \gamma) \rho}{1 - K_q(\lambda; \gamma) \rho}.$$

Consequently, putting everything together we have that

$$\begin{aligned} \nu^2(\lambda; \gamma) &= \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^T \mathbb{E}[|\Delta q_t(\lambda; \gamma)|^2] = K_f^2(\lambda; \gamma) \left(1 - \frac{2\rho(1 - K_q(\lambda; \gamma))}{1 - K_q(\lambda; \gamma) \rho} \right. \\ &\quad \left. + \frac{(1 - K_q(\lambda; \gamma))^2 (1 + K_q(\lambda; \gamma) \rho)}{(1 - K_q^2(\lambda; \gamma))(1 - K_q(\lambda; \gamma) \rho)} \right). \end{aligned} \tag{B.7}$$

Hence from the Gaussian absolute moment formula, we have that

$$\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^T \lambda_p \mathbb{E}[|\Delta q_t(\lambda; \gamma)|^p] = \lambda_p \nu^p(\lambda; \gamma) \frac{2^{p/2} \Gamma\left(\frac{p+1}{2}\right)}{\sqrt{\pi}}. \tag{B.8}$$

Putting Equations (B.3), (B.6), and (B.8) together gives Equations (11) and (12) completing the proof.