Optimal trade execution with unknown drift

Martin Forde

12th Sep 2024

Abstract

We show how existing results for optimal trading strategies with linear temporary price impact/exponential resilience or proportional transaction costs can be easily adapted for the more realistic situation when the drift of the asset is unknown, so we have to project to the observable filtration generated by the asset price process, using results from non-linear filtering theory. In particular, we observe that an arithmetic Brownian motion P with unknown (constant) drift μ is the continuation of a generalized bridge process under $\mathcal{F}^P$ with the true drift is replaced with its unbiased estimate over a fixed time window[1].

1 Introduction

Many price impact articles (and textbooks on the continuous-time Kalman filter) consider a semimartingale price process with a drift process which is an OU process, but since the drift process is not directly observable, we cannot easily estimate its paramaters, and even if the drift process were observable, we can still e.g. only compute MLE or GMM estimates for its parameters which will typically have non-small sample variance unless the time window under consideration is large (i.e. years in practice) and the model is well specified over this large time window (which will seldom be the case in practice)[2]. One can use the Kalman filter combined with the E-M algorithm to do this[3] (for which there built-in functions in Python for example), but from practical experience, we do not recommend since the sample variance of the estimate for the mean reversion speed of the OU process will be too large.

The alternate approach to this kind of problem (which we do not pursue here) is to use limit order book imbalance to predict mid-price movements (see e.g. [CDJ18], [CDO23], [PRS23] and references therein), see also [AD13] for an interesting approach when a moving-average term appears in the drift.

1.1 Outline of article

In the next subsection we introduce our Bachelier model with unknown drift, and make the canonical choice of Gaussian prior for μ at $t = 0$ based on a price history of P going back to some time $t_0 < 0$. We then re-write the price process P in the form in (2) where the drift is the obvious conditional unbiased estimate of μ; from a standard result in filtering theory, the $\bar{W}$ term is a Brownian motion with respect to $\mathcal{F}^P$, and we note that P is the continuation of a Brownian bridge process under $\mathcal{F}^P$, for which we have the usual explicit solution given in (3). In Section 2, we show numerically how the unknown drift affects the optimal trading strategy for an agent subject to linear temporary and transient price impact (with exponential kernel) using the main result in [NV22] and a large liquidation penalty. In section 3, we turn our attention to the Bachelier version of the classical Merton problem but with unknown drift and we give a necessary and sufficient condition for the agent to have positive expected P&L when trading with uncertain drift, and we compute the leading order asymptotic approximation for the width of the no-trade region for the case of non-zero transaction costs as in [KM15], but for the case of unknown drift.

1.2 Model setup

In this note, we consider a financial market living on a stochastic basis $(\Omega, \mathcal{F}, \mathcal{F}_t, \mathbb{P})$, where the filtration $\mathbb{F} = \{\mathcal{F}_t\}_{t_0 \leq t \leq T}$ satisfies the usual conditions. $\mathbb{P}$ is the objective probability measure, and we assume the basis carries a one-dimensional $\mathbb{P}$-Brownian motion W. We consider a financial market with a single asset with $\mathbb{P}$-dynamics

$$P_t \;=\; P_0 + \mu t + \sigma W_t \tag{1}$$

and $\mathcal{F}^P$ will denote the filtration generated by P (augmented by $\mathbb{P}$-null sets). We assume that σ is known and μ is unknown to a financial agent and that P has been observed continuously since $t_0 < 0$. The assumption that σ is known is natural since it can be computed from the observed quadratic variation of P over any subset of $[t_0, 0]$ which

[1]We thank Faycal Drissi, Lane Hughston, Rohan Hobbs and Leandro Sanchez Betancourt for many interesting discussions

[2]Note that if $dS_t = Y_t dt + \sigma dW_t$, where Y is an OU process: $dY_t = \kappa(\theta - Y_t)dt + \sigma_Y dB_t$ with $dW_t dB_t = \rho dt$ and σ is constant, then S is also a Gaussian process

[3]We thank Leandro Sanchez Betancourt for pointing this out

can be estimated from the realized variance of P (see e.g. Ait-Sahalia&Jacod[AJ14] for details and convergence results in this vein).

We specify an initial distribution $f(\mu)$ for μ at $t=0$, and assume that W is independent of μ. The natural/canonical choice here (which we will henceforth assume unless stated otherwise) is that $\mu \sim N(\frac{P_0-P_{t_0}}{0-t_0}, \frac{\sigma^2}{|t_0|})$, since this is what we obtain when applying the usual confidence interval approach to estimate μ using that $P_0 - P_{t_0} \sim N(\mu t, \sigma^2|t_0|)$, or from a Bayesian standpoint, $f(\mu)$ is the posterior $f(\mu|P_0)$ for μ using Bayes' theorem if the initial (prior) distribution for μ at t_0 is $U([-n,n])$, and we then let $n\to\infty$ (i.e. we have a flat prior for μ at t_0, see Appendix A for details). See also [BGP19] and [Driss22] for a Bayesian approach to this type of problem using a Gaussian prior for μ. We then use the classical filtering result in where the true drift is replaced by its conditional expectation.

From the final part of the Appendix, we know that $\mathbb{E}(\mu|\mathcal{F}^P_t) = \frac{P_t-P_{t_0}}{t-t_0}$. Hence the process $\bar{W}_t$ defined by

$$P_t \;=\; P_0 + \int_0^t \frac{P_u - P_{t_0}}{u-t_0} du + \sigma \bar{W}_t \tag{2}$$

is an $\mathcal{F}^P_t$-Brownian motion, see e.g. Eq 4.5 in Björk et al.[BDL10], Theorem 6.1 in [Chi] or Proposition 2 in [BGP19]. As mentioned in [BGP19], $\bar{W}$ is known as the *innovation process* in filtering theory.

Remark 1.1 This result appears counter intuitive since (2) does not contain μ, but the result is saying that P in (2) has the same law as P in (1) with the understanding that μ in (1) is $N(\frac{P_0-P_{t_0}}{0-t_0}, \frac{\sigma^2}{|t_0|})$, and we can easily verify this equivalence by comparing their sample covariances with Monte Carlo.

1.3 Basic properties of P under $\mathcal{F}^P$

Proposition 1.1 *Under $\mathcal{F}^P$, P satisfies the same (linear) SDE as the continuation of a Brownian bridge process constrained to be at P_{t_0} at $t=t_0$, but here $t\geq 0$ and $t_0<0$, and has the explicit solution:*

$$P_t \;=\; P_0 \frac{t_0 - t}{t_0} + P_{t_0}\frac{t}{t_0} + (t_0-t)\sigma \int_0^t \frac{d\bar{W}_s}{t_0 - s}\,.$$

Proof. *See Appendix B (or Eq 5.6.23 in [KS91]).* ■

Remark 1.2 From (1) we know that $\mathbb{E}(P_t^2)<\infty$ hence P_t is finite a.s., so P cannot explode in finite time, despite the apparent mean-fleeing behaviour of P around P_{t_0} under $\mathcal{F}^P$ in (2). In particular, P is a Markov process with respect to $\mathcal{F}^P$, and we see that $\mathbb{E}(P_t|P_s)$ satisfies the ODE

$$\frac{d}{dt}\mathbb{E}(P_t|P_s) \;=\; \frac{\mathbb{E}(P_t|P_s) - P_{t_0}}{t-t_0}$$

with solution

$$\mathbb{E}(P_t|P_s) \;=\; P_{t_0} + \frac{P_s - P_{t_0}}{s-t_0}(t-t_0) \;=\; P_{t_0} + \hat{\mu}_s(t-t_0) \tag{3}$$

for $t_0 \leq 0 \leq s \leq t$ (this expression will be needed in (4) below).

2 Application to price impact problems

2.1 Unconstrained problem

We can now apply many well known price impact methods/results to P but working under $\mathcal{F}^P$ - e.g., for an agent subject to linear temporary price impact with no liquidation penalty where the price paid per share at time t is $S_t = P_t + kv_t$ and v_t is the trading speed, using the same pointwise optimization argument with optional projection as in section 4.1 of [FSS22], we know that the optimal buying speed with no liquidation penalty is $v^*_t = \frac{\xi_t}{2k}$ where

$$\xi_t \;=\; \mathbb{E}(P_T - P_t|\mathcal{F}^P_t) \;=\; \hat{\mu}_t(T-t) \;=\; \frac{P_t - P_{t_0}}{t-t_0}(T-t)$$

and with a non-zero transaction cost of size ε

$$v^*_t \;=\; \frac{1}{2k}(\xi_t - \varepsilon\,\mathrm{sgn}(\xi_t))\,1_{|\xi_t|\geq\varepsilon}\,.$$

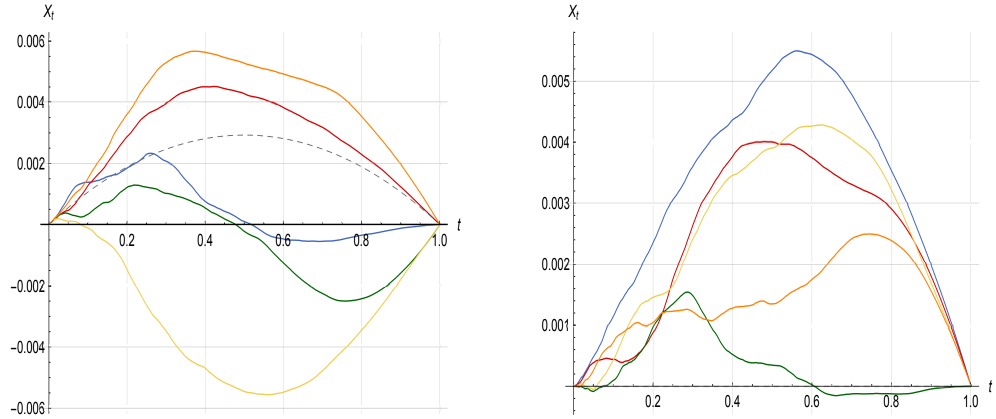

Figure 1: Here we have simulated the optimal stock holding X_t fives times for the problem in [NV22] but with μ unknown for parameters $\mu = 0.05$ (left) and $\mu = 0$ (right), with fixed parameters $\sigma = .2$, $\kappa = 1$, $\rho = 1$, $\lambda = 1$, $\phi = 0$, $\varrho = 1000$, $Y_0 = 0$, $P_0 = 1$ and $t_0 = -1$ with $P_{t_0} = P_0 - \mu|t_0|$ in both plots. The grey dashed line is the solution when μ is known (which is deterministic, and is identically zero in the second plot) so we see that not knowing the true value of μ significantly affects the optimal strategy, and in the second plot we see it leads to trading when the agent would not trade with perfect information about μ. Note that the terminal liquidation penalty coefficient $\varrho \gg 1$ here, so we are close to the case of perfect liquidation (which is also a round trip since $X_0 = 0$). (Mathematica code available on request)

2.2 Temporary price impact and exponential resilience with liquidation and running inventory penalties

If P belongs to the class of special semimartingales as defined in section 2 in [NV22], and the agent is subject to temporary price impact plus transient price impact under the propagator model with exponential resilience as in [NV22] with a running inventory penalty and finite liquidation penalty, then the standard variational and optional projection argument used to derive the main Theorem 3.2 in [NV22] still works under the filtration $\mathcal{F}^P$, so we just need to compute

$$\frac{1}{ds}\mathbb{E}(dA_s|\mathcal{F}^P_t) \;=\; \mathbb{E}(\hat{\mu}_s|\mathcal{F}^P_t) \;=\; \mathbb{E}(\frac{P_s - P_{t_0}}{s-t_0}|P_t) \;=\; \frac{P_{t_0} + \hat{\mu}_t(s-t_0) - P_{t_0}}{s-t_0} \;=\; \hat{\mu}_t \;=\; \frac{P_t - P_{t_0}}{t-t_0} \quad (4)$$

for $s \geq t$ (where the third equality follows from (3) with s and t swapped round); note this expression is needed for Eq 3.6 in [NV22] (see numerical results in Figure 1). A similar (but simpler) formula (also just requiring $\mathbb{E}(dA_s|\mathcal{F}^P_t)$) appears in Theorem 3.1 in [BMO20] for the case when there is no resilience, and the aforementioned formulae in [BMO20] and [NV22] both require computing a matrix exponential.

We can similarly extend to the case of two or N-agents using section 2 in [NV23] to compute the (open-loop) Nash equilibrium, although such results require the (possibly unrealistic) assumption that multiple agents can see each others trading speeds. In the second plot of Figure 1, we see there is non-zero trading activity for a round trip even though the agent would not trade with perfect information about μ (since the true $\mu = 0$ in this case), so in general for round trip problems there will be a critical range (μ_-, μ_+) of μ-values straddling zero, inside which the expected P&L is negative if μ is unknown to the agent). More generally, when $X_0 \neq 0$ this critical range need not include zero due to liquidation penalties.

Remark 2.1 For all the price impact problems considered, all that matters ultimately is $\mathbb{E}(\mu_s|\mathcal{F}^P_t)$ so we can replace the Brownian motion W above with any sufficiently well behaved martingale M, and the choice of martingale does not affect the optimal trading strategy (unless we start using non-linear utility functions). In this case, since

$$P_0 - P_{t_0} \;=\; \mu(0-t_0) + M_0 - M_{t_0}$$

so the natural choice of initial distribution for μ now is the law of $(P_0 - P_{t_0} - (M_0 - M_{t_0}))/(0-t_0)$ with P_0 and P_{t_0} taking their observed values. Note also that Theorem 3.6 in [NV22] does not require P to be continuous, so one could assume P is a Lévy process, e.g. a CGMY process.

3 The Merton problem with unknown drift

In this section, we remove the friction (i.e. the price impact) but we now allow the agent to be risk-averse by using a non-linear utility function. Assuming

$$dS_t \;=\; \mu dt + \sigma dW_t \quad (5)$$

with unknown μ as in section 2 for P, we consider the Merton problem with $r = 0$, and let ϕ_t denote the agent's stock holding at time t, which we assume has to be $\mathcal{F}_t^S$-adapted (note we are working in a Bachelier setting here because otherwise the solution for $\phi^*(S,t)$ below becomes rather cumbersome). Then the total wealth of the agent X_t evolves as

$$dX_t = \phi_t dS_t$$

so the HJB equation for the value function $V(S,x,t) = \sup_{\phi\in\mathcal{A}} \mathbb{E}_{S,X,t}(U(X_T))$ is

$$V_t + \mu(S)V_S + \frac{1}{2}\sigma^2 V_{SS} + \sup_\phi[\phi\mu(S)V_x + \frac{1}{2}\sigma^2\phi^2 V_{xx} + \sigma^2\phi V_{Sx}] = 0$$

with $\mu(S) := \frac{S-S_{t_0}}{t-t_0}$, and we can then solve for ϕ^* in feedback form, and then re-write as a non-linear PDE. For the case when $U(x) = -e^{-\alpha x}$, using the ansatz $V(S,x,t) = -e^{-\alpha(x+w(S,t))}$, we find that

$$w_t + \frac{1}{2}\sigma^2 w_{SS} + \frac{1}{2}\frac{(S-S_{t_0})^2}{(t-t_0)^2\alpha\sigma^2} = 0$$

for which the terminal condition is $w(S,T) = 0$ if there is no liquidation penalty. This can be solved in closed-form (using Feynman-Kac) to give

$$w(S,t) = \frac{1}{2\alpha}(\frac{(T-t)(S^2 - 2SS_{t_0} + S_{t_0}^2 + (t_0-t)\sigma^2)}{(t-t_0)(T-t_0)\sigma^2} + \log\frac{T-t_0}{t-t_0})$$

and

$$\phi^*(S,t) = \frac{\hat{\mu}_t}{\alpha\sigma^2} = \frac{S_t - S_{t_0}}{\alpha\sigma^2(t-t_0)}. \qquad (6)$$

Note that

$$\mathbb{E}(\int_0^T \phi^*(S,t)dS_t) = \mathbb{E}(\int_0^T \frac{S_t - S_{t_0}}{\alpha\sigma^2(t-t_0)}\mu dt) = (S_0 - S_{t_0})\frac{\mu}{\alpha\sigma^2}\frac{\log(T-t_0)}{\log|t_0|}$$

so we have negative expected P&L if and only if $\mu(S_0 - S_{t_0}) < 0$.

Note this is of the same form as the solution $\bar{\phi} = \frac{\mu}{\alpha\sigma^2}$ for the problem when μ is known (note this is for the Bachelier model in (5) here not the case when S is Geometric Brownian motion for which the solution is $\frac{\mu}{\alpha\sigma^2 S_t}$), but now μ has been replaced with $\hat{\mu}_t$. Thus even if the true $\mu = 0$, we see that it is optimal for the agent to trade with partial information about μ (see simulation in Figure 2), although $\phi^*(S,t) \to \frac{\mu}{\alpha\sigma^2}$ as $t \to \infty$. Note ϕ_t^* is not generally equal to the "myopic" trading strategy $\mu_t/(\alpha\sigma_t^2)$ when S follows a more general Itô process of the form $dS_t = \mu_t dt + \sigma_t dW_t$, see e.g. [KO96].

Remark 3.1 For the case of log utility $U(x) = \log x$ when S is a general semimartingale $dS_t = S_t(\mu_t dt + \sigma_t dW_t)$, it is well known that $\phi_t^* = \frac{\mu_t}{\sigma_t^2}$, so in this case $\phi^* = \frac{\hat{\mu}_t^*}{\sigma_t^2}$ (this is known as the growth optimal portfolio).

Remark 3.2 In practice, one could argue that one should not start trading unless we have already rejected the null hypothesis that $\mu = 0$.

3.1 Adding small proportionate transaction costs

From (6) we see that

$$d\phi^*(S_t,t) = \frac{dS_t}{\alpha\sigma^2(t-t_0)}$$

so $\frac{d\langle\phi^*\rangle_t}{d\langle S\rangle_t} = \frac{1}{\alpha^2\sigma^2(t-t_0)^2}$, and (from the formal computations in section 2.1 in [KM15], or section 4.1 in [KL13]) the leading order term for the optimal trading strategy with proportional transaction costs of size $\varepsilon \ll 1$ and fixed time horizon $T > 0$ with exponential utility function as above is to engage in the minimal amount of trading to keep ϕ_t within $\bar{\phi} \pm \Delta\phi_t$, where

$$\Delta\phi_t^* = \pm(\frac{3}{2\alpha}\frac{d\langle\phi^*\rangle_t}{d\langle S\rangle_t}S_t)^{\frac{1}{3}}\varepsilon^{\frac{1}{3}}$$

and we see that the no-trade region (NTR) shrinks when t goes large or when S_t goes small (see numerics below). Note that the leading order width of the NTR is here is zero when the drift is known, since ϕ^* is constant in this case.

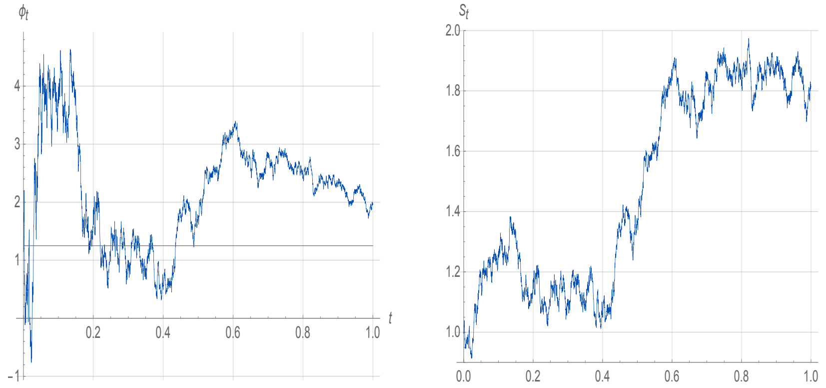

Figure 2: On the left we see a Monte Carlo simulation of the optimal stock holding ϕ_t^* with μ unknown (blue) and μ known (grey), and the corresponding stock price process S_t (right plot) for the Merton problem with exp utility and unknown drift with $t_0 = -1$, $S_{t_0} = .95$, $S_0 = 1$, $T = 20$, $\alpha = 1$, $\sigma = .2$ and true $\mu = 0.05$. Since $S_0 > S_{t_0}$ by assumption, the agent initially held a long position but went short as the stock price went down.

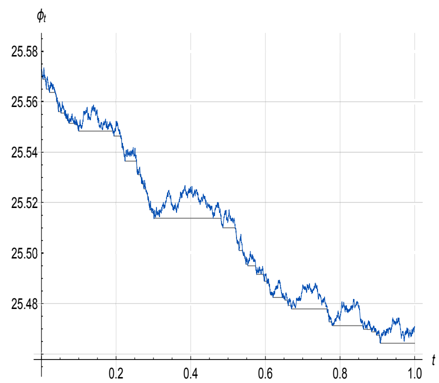

Figure 3: Here we have plotted the upper boundary for the No-Trade region (blue) with proportional transaction costs of size ε, and the optimal stock holding ϕ_t^* (in grey) when the true drift is unknown, for $t_0 = -1$; $\mu = 0.05$, $\sigma = .2$, $\varepsilon = .005$; $S_{t_0} = S_0 - \mu|t_0|$, $S_0 = 1$; $T = 1$, $\alpha = 1$, and ϕ starting on the upper boundary. Note that a smaller investor needs to choose a larger α value to ensure a smaller $\bar{\phi}$ since we are working with exp utility.

3.2 Concluding remarks

We have shown how known results for optimal trading strategies with linear temporary price impact/exponential resilience or proportional transaction costs are easily adapted for the more realistic situation when the drift of the asset is unknown (by switching to the filtration generated by the price process P) and warned against the common trend of using an OU process for the drift as it's almost impossible to estimate its parameters with low sample variance in practice.

In Section 2, we saw numerically how the unknown drift can make the optimal liquidation strategy highly stochastic for an agent subject to linear temporary and transient price impact (with exponential kernel) using the main result in [NV22], even though the solution is deterministic when the drift is known. In particular, not knowing the the drift leads to non-zero trading activity for round trips when the true drift is zero, which is clearly sub-optimal since the agent has no "edge" on the market in this case. Going forward, one could look to numerically solve the (double obstacle) free boundary problem (FBP) for the transaction costs problem in Section 3.1 with unknown drift (see e.g. [DY09] for the known drift case), or the (related) problem of computing the optimal trading stategy for a liquidity taker in a liquidity pool with an automated market maker (e.g. the Uniswap protocol), where the transaction cost is now the pool fee (which is a stochastic permanent price impact problem). The domain for the PDE for these problems is a cube since the relevant state variables are (t, S_t, Y_t) where Y is the agents risky wealth, or one could also look to learn the upper and lower free boundaries using a DNN.

References

[AD13] Appleby, J. and J.Daniels, "Long Memory and Financial Market Bubble dynamics in Affine Stochastic Differential Equations with Average Functionals", arxiv preprint, 2013

[AJ14] Aït-Sahalia, Y. and J.Jacod, "High-Frequency Financial Econometrics", Princeton University Press

[AN22] Abi-Jaber, E. and Neuman E., "Optimal Liquidation with Signals: the General Propagator case", preprint (2022)

[BGP19] Bismuth, A., Gueant, O., Pu, J., "Portfolio choice, portfolio liquidation, and portfolio transition under drift uncertainty", *Mathematics and Financial Economics*, 13, 661–719, 2019

[CDJ18] Cartea, A., R.Donnelly and S.Jaimungal, "Enhancing trading strategies with order book signals", *Applied Mathematical Finance*, 25:1, 2018

[CDO23] Cartea, A., F.Drissi and P.Osselin, "Bandits for Algorithmic Trading with Signals", SSRN preprint, 2023

[Chi] Chigansky, P., "Introduction to Nonlinear Filtering", lecture notes

[BMO20] Belak, C., J.Muhle-Karbe and K.Ou, "Liquidation in Target Zone Models, *Market Microstructure and Liquidity*, Vol. 4 (2020), No. 03, 1950010

[BDL10] Björk, T., M.H.A.Davis and C.Landén, "Optimal investment under partial information", *Mathematical Methods of Operations Research*, volume 71, pages 371–399, 2010

[DY09] Dai, M. and F.Yi, "Finite-horizon optimal investment with transaction costs: A parabolic double obstacle problem", *J. Differential Equations*, 246, 1445–1469, 2009

[Driss22] Drissi, F., "Solvability of differential riccati equations and applications to algorithmic trading with signals", *Applied Mathematical Finance*, 29, 457–493, 2022.

[FSS22] Forde,M., L.Sánchez-Betancourt and B.Smith, "Optimal trade execution for Gaussian signals with power-law resilience", *Quantitative Finance*, 22(3), 585-596, 2022

[GM15] Guasoni, P. and J.Muhle-Karbe, "Long Horizons, High Risk Aversion, and Endogeneous Spreads", *Mathematical Finance*, 25, no. 4 p. 724-753, 2015.

[GGMS14] Gerhold, S., P.Guasoni, J.Muhle-Karbe, "Transaction Costs, Trading Volume, and the Liquidity Premium", *Finance and Stochastics*, Vol. 18 (2014), No. 1, pp. 1-37

[KS91] I.Karatzas and S.Shreve (1991), Brownian motion and Stochastic Calculus, Springer-Verlag.

[KL13] Kallsen, J. and S.Li, "Portfolio Optimization under Small Transaction Costs: a Convex Duality Approach", preprint, 2013.

[KM15] Kallsen, J. and J.Muhle-Karbe, "Option Pricing and Hedging with Small Transaction Costs", *Mathematical Finance*, Vol. 25 (2015), No. 4, pp. 702-723, 2015

[KO96] Tong Suk Kim and Edward Omberg, "Dynamic Nonmyopic Portfolio Behavior" *The Review of Financial Studies*, Vol. 9, No. 1 (Spring, 1996)

[Kut04] Kutoyants Yu. A., "Statistical Inference for Ergodic Diffusion Processes", Springer Series in Statistics, London, 2004, 496p.

[NV22] Neuman, E. and M.Voß, "Optimal Signal-Adaptive Trading with Temporary and Transient Price Impact", *SIAM J. Financial Mathematics*, Vol. 13, No. 2, 551-575, 2022.

[NV23] Neuman, E. and M.Voß, "Trading with the crowd", *Mathematical Finance*, Vol. 33 (2023), No. 3, pp. 548-617

[PRS23] Pulido, S., M.Rosenbaum and E.Sfendourakis, "Understanding the least well-kept secret of high-frequency trading", arxiv preprint, 2023

[WW97] Whalley, A.E. and P.Wilmott, "An Asymptotic Analysis of an Optimal Hedging model for Option Pricing with transaction costs", *Mathematical Finance*, Vol. 7, No. 3 (July 1997), 307–324

A Computing $\mathbb{E}(\mu|(P_s)_{s\in[t_0,t]})$

If $M_t = \mu t + \sigma W_t$, then for a flat prior for μ on $\mathbb{R}$ at t_0 (which is clearly an improper prior), using Bayes' formula and Girsanov's theorem, the posterior $p(\mu|(P_s)_{s\in[t_0,t]})$ of μ at $t>0$ (given $(P_s)_{s\in[t_0,t]}$) is

$$\begin{aligned} p(\mu|(P_s)_{s\in[t_0,t]}) &\propto \text{Likelihood function of } (P_s)_{s\in[t_0,t]} \\ &= \frac{1}{\sigma}e^{\int_{t_0}^t \gamma dP_s - \frac{1}{2}\int_{t_0}^t \gamma^2 ds}\mathbb{Q}_0(\frac{d(P-P_0)}{\sigma}) \\ &= \frac{1}{\sigma}e^{\frac{\mu}{\sigma}(\frac{P_t-P_{t_0}}{\sigma})-\frac{1}{2}(\frac{\mu}{\sigma})^2(t-t_0)}\mathbb{Q}_0(\frac{d(P-P_0)}{\sigma}) \\ &= const.\times e^{-\frac{(\mu-\bar{\mu})^2}{2\sigma^2/(t-t_0)}}\mathbb{Q}_0(\frac{d(P-P_0)}{\sigma}) \end{aligned}$$

where $\mathbb{Q}_0$ denotes the Wiener measure on $(C([t_0,t]),\mathcal{B}(C([t_0,t])),\mathbb{Q}_0)$, $\gamma=\mu/\sigma$ and $\bar{\mu}=\frac{P_t-P_{t_0}}{t-t_0}$, so the posterior for μ is $N(\frac{P_t-P_{t_0}}{t-t_0},\frac{\sigma^2}{t-t_0})$ as one would expect, so (formally at least) $\mathbb{E}(\mu|(P_s)_{s\in[t_0,t]})=\frac{P_t-P_0}{t-t_0}$.

If we modify this analysis to instead use a prior at $t=0$ which is $N(\frac{P_0-P_{t_0}}{0-t_0},\sigma^2/|t_0|)$, then we also find that $\mathbb{E}(\mu|(P_s)_{s\in[0,t]})=\frac{P_t-P_{t_0}}{t-t_0}$.

B Proof of SDE for P with respect to $\mathcal{F}^P$

Appying Ito's lemma to (3) we see that

$$\begin{aligned}
dP_t &= -\frac{P_0}{t_0}dt + P_{t_0}\frac{dt}{t_0} + (t_0-t)\sigma\frac{d\bar{W}_t}{t_0-t} - \sigma\int_0^t \frac{d\bar{W}_s}{t_0-s}dt \\
&= (P_{t_0}-P_0)\frac{dt}{t_0} + \sigma d\bar{W}_t - \sigma\int_0^t \frac{d\bar{W}_s}{t_0-s}dt \\
&= (P_{t_0}-P_0)\frac{dt}{t_0} + \sigma d\bar{W}_t - (\frac{P_t}{t_0-t} - \frac{P_0}{t_0} - P_{t_0}\frac{t}{t_0(t_0-t)})dt \\
&= P_{t_0}\frac{dt}{t_0} + \sigma d\bar{W}_t - (\frac{P_t}{t_0-t} - P_{t_0}\frac{t}{t_0(t_0-t)})dt \\
&= \frac{P-P_{t_0}}{t-t_0}dt + \sigma d\bar{W}_t
\end{aligned}$$