

Optimal Trading with Linear Costs

Joachim de Lataillade, Cyril Deremble,
Marc Potters & Jean-Philippe Bouchaud

Capital Fund Management, 6 Boulevard Haussmann, 75009 Paris, France

Abstract

We consider the problem of the optimal trading strategy in the presence of linear costs, and with a strict cap on the allowed position in the market. Using Bellman's backward recursion method, we show that the optimal strategy is to switch between the maximum allowed long position and the maximum allowed short position, whenever the predictor exceeds a threshold value, for which we establish an exact equation. This equation can be solved explicitly in the case of a discrete Ornstein-Uhlenbeck predictor. We discuss in detail the dependence of this threshold value on the transaction costs. Finally, we establish a strong connection between our problem and the case of a quadratic risk penalty, where our threshold becomes the size of the optimal non-trading band.

1 Introduction

Contrarily to the efficient market dogma, prices have some degree of predictability, at least on short time scales. Statistical arbitrage strategies aim at exploiting this small predictability. However, costs make it difficult, although not impossible, to eke out a profit from these inefficiencies. Indeed, every trade is costly; the structure of these costs is actually quite complex. Some of them are related to various fees (market fees, brokerage fees, etc.) and are usually a small fraction of the traded quantity (typically 10^{-5} to 10^{-4} on liquid markets). We call these costs "linear", because they are simply proportionnal to the traded amount. Another source of linear costs is the bid-ask spread, which represents a few "basis points" (bp, i.e. 10^{-4} of the value of the contract). Much more subtle are impact induced costs, which come from the fact that a large order must be split in a sequence of small trades that are executed incrementally. But since each executed trade, on average, impacts the price in the direction of the trade, the average execution price is higher (in the case of a buy) than the decision price, leading to what is called "execution shortfall". This cost is clearly non-linear, since the price impact itself increases with the size Q of the trade. Empirical data suggests that the impact induced cost is on the order of $\sigma Q^{3/2}/\sqrt{V}$, where σ is the daily volatility and V the daily turnover (see e.g. [TLD⁺11]). This shows that for an order corresponding to 1% of the daily turnover, and for an asset with 2% daily volatility, the impact cost is on the order of 20 bp. But for much smaller orders, or for assets where the bid-ask spread is large, linear costs can be dominant.

The problem we want to address and solve in this work is to determine the optimal trading strategy when one has discovered a predictive statistical signal, in the presence of linear trading costs and a constraint on the maximum size of the position (both long and short). While the case of a quadratic risk control has been considered in the literature for linear costs [DN90, SS94, Con86],

quadratic costs (i.e. costs growing like Q^2) [DL07, AC00] or various impact-dependent cost [Alm03], we are not aware of any published solution in the case where the risk constraint is a cap on the position. This problem can be of significant interest in practice, because the risk of a trading system is sometimes handled completely outside of the system, through such a cap on the position, in order to reduce operational risk.

Our final solution for the optimal position π as a function of the predictor p is as follows: the size of the position $|\pi|$ should always be at the maximum allowed position M , with a sign that switches between -1 and $+1$ whenever the predictor exceeds a *threshold value* q^* (and vice-versa when the predictor becomes smaller than $-q^*$). We find an exact equation for the threshold q^* , that we solve explicitly in the case of a (discrete-time) Ornstein-Uhlenbeck predictor. In some limits we find that the threshold q^* scales as the one-third power of the cost parameter, a result already discussed in the literature but in the context of a quadratic risk constraint. We explain why our result is strongly connected with this alternative problem. We also check the validity of our results using numerical simulations.

2 Description of the problem

We consider an agent who trades a single asset, of current price Price_t . The position (signed number of shares/contracts) of the trader at time t is π_t . We assume that the agent has some signal p_t that predicts the next price change $r_t = \text{Price}_{t+1} - \text{Price}_t$, and is faced with the following constraints:

- His/her risk control system is simply a cap on the absolute size of his position : $|\pi| \leq M$, with no other risk control. M will be called the “MaxPos” of the agent.
- He/she has to pay linear costs ΓQ whenever he/she trades a quantity $Q = |\delta\pi| = |\pi_{t+1} - \pi_t|$

The agent wants to maximise his/her expected gains, by trading over a long period $[0, T]$ (we will later consider the limit $T \rightarrow \infty$).

We also assume the predictor to :

- have a linear predictability: $E[r_t|p_t] = A \cdot p_t$ with some constant A
- be positively auto-correlated: $\forall q, P(p_{t+1} > q|p_t)$ increases continuously with p_t
- be Markovian: $\forall \omega_{t+1}, P(\omega_{t+1}|p_t, p_{t-1}, \dots) = P(\omega_{t+1}|p_t)$ where ω_{t+1} is any event at $t + 1$
- be symmetric: $\forall q, P(p_{t+1} > q|p_t = p) = P(p_{t+1} < -q|p_t = -p)$
- be unbounded: $\forall q, \exists \epsilon_q > 0$ s.t. $P(p_{t+1} > q|p_t = 0) > \epsilon_q$.

Remarks:

- Without loss of generality, we can always set $A = 1$, so that p_t is in price units.
- Because of the positive auto-correlation of p_t , we can define an **integrated predictability** at $t = \infty$, depending on p_t :

$$p_\infty(p_t) = E[\text{Price}_\infty - \text{Price}_t \mid p_t] = \sum_{n=0}^{\infty} E[p_{t+n} \mid p_t]$$

This quantity indicates how much one will gain in the future if one keeps a fixed position $\pi_{t' \geq t} = \pi$: the expected gain is then $p_\infty(p_t) \pi$.

- Finally, one can consider the case where Condition (i) is not met, and note

$$\mathfrak{L}_t(p_t) = E[r_t|p_t]$$

the immediate predictability. Then, as we will see in Section 3.6, we can still express the solution of the system if we suppose that \mathfrak{L}_t is continuous, uneven, strictly increasing, and $\lim_{p_t \rightarrow \infty} \mathfrak{L}_t(p_t) > \Gamma$. But the parameters of this solution will of course depend on these functions \mathfrak{L}_t .

3 The Optimal Strategy

3.1 A naïve solution

At first sight, the solution to this problem seems straightforward: if the expected future gain (given by the integrated predictability) exceeds the trading cost per contract Γ , then one trades in the direction of the signal (if not already at the MaxPos), otherwise one does not. This solution obviously generates a positive average gain, but it has no reason to be the optimal solution. Indeed, because the predictor is auto-correlated in time, it might be worthy (and in general it will be) to wait for a larger value of the predictor, in order to grab the opportunities that have the most chances to get realised, and discard the others. As we shall see, the mistake in this naïve reasoning is not to compare the future gain with the cost, but rather comes from a wrong definition of the future gain, which does not include future trading decisions.

3.2 The Bellman method

To attack this problem, we will use Bellman's **optimal control theory**, or dynamic programming [Bel03], which consists in solving the problem backwards: by assuming one follows the optimal strategy for all future times $t' > t$, we can find the optimal solution at time t . As is usual in dynamic programming, we have a **control variable**¹ π_t , which needs to be optimised, and a **state variable** p_t , which parametrises the solution. The optimisation will be done through a **value function** $V_t(\pi, p)$, which gives the maximal expected gains between time t and $+\infty$, considering that the position at $t - 1$ is π and the predictor's value at t is p . The optimal solution of the system will be denoted $(\pi_t^*)_{t \in [0, T]}$.

Let us start with the simple case $t = T$, i.e. the optimal strategy at the last time step. In this case the expected future return is really $p_\infty(p)\pi_T$ where $p = p_T$, since no trading is allowed beyond that time. Any trade $\delta\pi$ induces a cost $\Gamma|\delta\pi|$, so:

- if $p_\infty(p) \geq \Gamma$ then $\pi_T^* = M$, $V_T(\pi, p) = p_\infty(p) \cdot M - \Gamma(M - \pi)$.
- if $p_\infty(p) \leq -\Gamma$ then $\pi_T^* = -M$, $V_T(\pi, p) = -p_\infty(p) \cdot M - \Gamma(M + \pi)$.
- if $|p_\infty(p)| < \Gamma$ then $\pi_T^* = \pi$, $V_T(\pi, p) = p_\infty(p) \cdot \pi$.

¹In the usual Bellman terminology the control variable is actually $\pi_t - \pi_{t-1}$, but here we want to insist on what we can control (the position) and what we cannot (the value of the predictor).

Hence, we recover exactly the naïve solution in this case, but this is only because there is no trading beyond $t = T$.

Now if we consider $t < T$, we have to maximise a quantity including immediate gains, costs and future gains. This leads to the following recurrence relation:

$$V_t(\pi, p) = \max_{|\pi'| \leq M} (p \cdot \pi' - \Gamma |\pi' - \pi| + \int P(p_{t+1} = p' \mid p_t = p) V_{t+1}(\pi', p') dp') \quad (1)$$

and π_t^* is the value of π' which realises this maximum when $\pi = \pi_{t-1}^*$ and $p = p_t$.

3.3 General solution

In what follows, we will need the following notations:

$$\begin{aligned} P(p'|p) &= P(p_{t+1} = p' \mid p_t = p) \\ \mathcal{P}_{>q}(p) &= P(p_{t+1} > q \mid p_t = p) \\ \mathcal{P}_{<q}(p) &= P(p_{t+1} < q \mid p_t = p) \end{aligned}$$

where the dependency on t is kept implicit.

Proposition 1. *There exist two functions $g(t, p)$ and $h(t, p)$ and a sequence $(q_t)_{t \in [0, T]}$ such that, for every $t \in [0, T]$, we have the following:*

$$\bullet \pi_t^* = \begin{cases} \pi_{t-1}^* & \text{if } |p_t| < q_t \\ M & \text{if } p_t \geq q_t \\ -M & \text{if } p_t \leq -q_t \end{cases} \quad (\text{with } \pi_{-1} = 0)$$

$$\bullet V_t(\pi, p) = \begin{cases} g(t, p)\pi + h(t, p) & \text{if } |p_t| < q_t \\ (g(t, p) - \Gamma) \cdot M + \Gamma\pi + h(t, p) & \text{if } p_t \geq q_t \\ (-g(t, p) - \Gamma) \cdot M - \Gamma\pi + h(t, p) & \text{if } p_t \leq -q_t \end{cases}$$

• $g(t, p)$ is a continuous, strictly increasing function of p which satisfies, for $t < T$:

$$g(t, p) = p + \Gamma \cdot [\mathcal{P}_{>q_{t+1}}(p) - \mathcal{P}_{<-q_{t+1}}(p)] + \int_{-q_{t+1}}^{q_{t+1}} P(p'|p) g(t+1, p') dp' \quad (2)$$

• q_t is such that $q_t \geq 0$ and

$$g(t, q_t) = \Gamma. \quad (3)$$

Proof. The proof is done by backwards recursion (*ie.* we assume that it is true for $t+1$ to prove that it is true for t). One can easily check from section (3.2) that the statement is true for $t = T$, with in particular $g(T, p) = p_\infty(p)$.

Let us then suppose that it is true for $t + 1$. We have the following:

$$\begin{aligned}
V_t(\pi, p) &= \max_{|\pi'| \leq M} (p \cdot \pi' - \Gamma |\pi - \pi'| + \int P(p'|p) V_{t+1}(\pi', p') dp') \\
&= \max_{|\pi'| \leq M} (p \cdot \pi' - \Gamma |\pi - \pi'| + \int_{-q_{t+1}}^{q_{t+1}} P(p'|p) \cdot [g(t+1, p')\pi' + h(t+1, p')] dp' \\
&\quad + \int_{q_{t+1}}^{+\infty} P(p'|p) \cdot [(g(t+1, p') - \Gamma) \cdot M + \Gamma\pi' + h(t+1, p')] dp' \\
&\quad + \int_{-\infty}^{-q_{t+1}} P(p'|p) \cdot [(-g(t+1, p') - \Gamma) \cdot M - \Gamma\pi' + h(t+1, p')] dp')
\end{aligned}$$

So if we set

$$g(t, p) = p + \Gamma \cdot [\mathcal{P}_{>q_{t+1}}(p) - \mathcal{P}_{<-q_{t+1}}(p)] + \int_{-q_{t+1}}^{q_{t+1}} P(p'|p) g(t+1, p') dp'$$

and

$$h(t, p) = \mathcal{H} + M \int_{q_{t+1}}^{+\infty} P(p'|p) \cdot [g(t+1, p') - \Gamma] dp' + M \int_{-\infty}^{-q_{t+1}} P(p'|p) \cdot [-g(t+1, p') - \Gamma] dp'$$

with

$$\mathcal{H} = \int_{-\infty}^{+\infty} P(p'|p) h(t+1, p') dp',$$

this gives us:

$$\begin{aligned}
V_t(\pi, p) &= \max_{|\pi'| \leq M} (g(t, p) \cdot \pi' - \Gamma |\pi - \pi'| + h(t, p)) \\
&= \max [\max_{\pi \leq \pi' \leq M} (g(t, p) - \Gamma)\pi' + \Gamma\pi, \max_{-M \leq \pi' \leq \pi} (g(t, p) + \Gamma)\pi' - \Gamma\pi] + h(t, p)
\end{aligned}$$

Using the definition of $g(t, p)$ above, we can prove the following :

- $g(t, 0) = 0$ by symmetry
- $g(t, p)$ is continuous and strictly increasing with p : indeed, if we rewrite Equation (2) as

$$g(t, p) = p + \int_{-\infty}^{\infty} P(p'|p) \bar{g}(t+1, p') dp'$$

where

$$\bar{g}(t+1, p') = \begin{cases} g(t+1, p') & \text{if } |p'| \leq q_{t+1} \\ \Gamma & \text{if } p' > q_{t+1} \\ -\Gamma & \text{if } p' < -q_{t+1} \end{cases}$$

then it suffices to see that $p \mapsto \bar{g}(t+1, p)$ is continuous and increasing, as well as the cumulative distribution functions $p \mapsto P_{>q}(p)$ for $q \in \mathbb{R}$

- $\lim_{p \rightarrow +\infty} g(t, p) = +\infty$ because for $p > 0$, we can show that $\mathcal{P}_{>q_{t+1}}(p) \geq \mathcal{P}_{<-q_{t+1}}(p)$ and $\int_{-q_{t+1}}^{q_{t+1}} P(p'|p) g(t+1, p') \geq 0$.

Thus, there exists a unique $q_t \geq 0$ which satisfies

$$g(t, q_t) = \Gamma$$

and we finally have:

- if $p \geq q_t$, then $g(t, p) \geq \Gamma$ so:

$$V_t(\pi, p) = \max [\Delta M + \Gamma \pi , (\Delta + \Gamma)\pi] + h(t, p)$$

with $\Delta = g(t, p) - \Gamma$. But in order to have $\Delta M + \Gamma \pi < (\Delta + \Gamma)\pi$ we need $\pi > M$ (or $\Delta = 0$), so the maximum is realised for $\pi' = M$, and

$$V_t(\pi, p) = (g(t, p) - \Gamma) \cdot M + \Gamma \pi + h(t, p)$$

Note that, technically, in the case $p = q_t$, then $\Delta = 0$ so that any $\pi' \geq \pi$ maximises $V_t(\pi, p)$: the optimum is not unique anymore. But for the sake of simplicity, we impose the solution $\pi' = M$ in that particular case.

- if $p \leq -q_t$, then $g(t, p) \leq -\Gamma$, and similarly we obtain a maximum for $\pi' = -M$, and

$$V_t(\pi, p) = (-g(t, p) - \Gamma) \cdot M - \Gamma \pi + h(t, p)$$

- if $|p| < q_t$, then $|g(t, p)| \leq \Gamma$, the maximum is realised for $\pi' = \pi$ and

$$V_t(\pi, p) = g(t, p)\pi + h(t, p)$$

□

3.4 The self-consistency equation

The solution given by Proposition 1 exhibits a dependency in t .

Let us now consider the case where $T \rightarrow \infty$, and suppose that the predictor is stationary, i.e.: $P(p_{t+1} = p' | p_t = p)$ is independent of t . Then we obtain a *telescopic* solution, where the dependency in t completely disappears, so that we only have a one-variable function g and a threshold q^* , satisfying the following equations:

$$g(p) = p + \Gamma \cdot [\mathcal{P}_{>q^*}(p) - \mathcal{P}_{<-q^*}(p)] + \int_{-q^*}^{q^*} P(p'|p) g(p') dp' \quad (4)$$

$$g(q^*) = \Gamma \quad (5)$$

Equation (4) is a **self-consistent** functional equation. The optimal solution to the system is then:

- if $p_t \geq q^*$ then $\pi_t^* = M$

- if $p_t \leq -q^*$ then $\pi_t^* = -M$
- if $|p_t| < q^*$ then $\pi_t^* = \pi_{t-1}^*$.

Thus, we obtain a very simple trading system, always saturated at $\pm \text{MaxPos}$, with a threshold to decide at each step whether we should revert the position or not. This of course looks a lot like the naïve solution from Section 3.1. The only difference lies in the value of the threshold q^* , defined by Equations (4) and (5), instead of $q_{\text{naïve}} = p_{\infty}^{-1}(\Gamma)$ for the naïve solution. Intuitively, those equations take our future trading into account, whereas the naïve solution does not.

If we look closely at Equation (4), its interpretation becomes transparent: $g(p)$ is equal to $1/2M$ times the expected difference in total future profit between the situation where $\pi = +M$ and the situation where $\pi = -M$. This difference is made up of:

- the term $2pM$ which represents the difference in immediate gain
- $2\Gamma M \cdot \mathcal{P}_{>q}(p)$ which represents the loss if the current position is $-M$ and in the next time step the predictor goes over the positive threshold q^* (hence π will go to $+M$)
- $2\Gamma M \cdot \mathcal{P}_{<-q}(p)$ which represents the loss if the current position is $+M$ and in the next time step the predictor goes below the negative threshold $-q^*$ (hence π will go to $-M$)
- $\int_{-q}^q P(p'|p) [2Mg(p')] dp'$ which is the expected difference in total future profit if, in the next step, the predictor remains between the two thresholds (leaving π unchanged).

Since the change of position between $-M$ and $+M$ costs $2\Gamma M$, it makes sense to compare $2Mg(p)$ with it and only trade when $g(p)$ is greater than Γ . Hence, $g(p)$ can be seen as the “gain per traded lot”.

Equation (4) already allows us to say that $g(p) \geq p$ for any $p \geq 0$ (and $g(p) \leq p$ for $p \leq 0$). As $g(q^*) = \Gamma$, this implies in particular that $q^* \leq \Gamma$. This property is actually rather intuitive: indeed, if the immediate gain is higher than the trading cost, then there is no reason not to trade the maximal possible amount.

3.5 Reformulation as a path integral

Although Equation (4) is easy to interpret, it proves very difficult to solve in concrete cases like the Ornstein-Uhlenbeck case that we will consider in Section 4. In the present section, we will extract an alternative equation for the optimal threshold, which, although more sophisticated than the self-consistency equation, will be easier to solve in practice.

Equation (4) can be rewritten by expanding the function g :

$$\begin{aligned}
g(p) = & p + \int_{-q^*}^{q^*} p' P(p'|p) dp' + \int_{-q^*}^{q^*} \int_{-q^*}^{q^*} p'' P(p''|p') P(p'|p) dp' dp'' + \dots \\
& + \Gamma \cdot \left[\int_{q^*}^{+\infty} P(p'|p) dp' + \int_{q^*}^{+\infty} \int_{-q^*}^{q^*} P(p''|p') P(p'|p) dp' dp'' + \dots \right] \\
& - \Gamma \cdot \left[\int_{-\infty}^{-q^*} P(p'|p) dp' + \int_{-\infty}^{-q^*} \int_{-q^*}^{q^*} P(p''|p') P(p'|p) dp' dp'' + \dots \right]
\end{aligned}$$

Let us now consider that the predictor starts at $p_0 = p$ at time t , and follows the infinite path (p_1, p_2, p_3, \dots) afterwards. The above expansion tells us that this path will contribute to $g(p)$ as long as $-q^* < p_i < q^*$, and will stop contributing as soon as $|p_i| \geq q^*$ for some $i > 0$. Moreover, its contribution is given by the sum $\sum_{i=0}^{n-1} p_i$, where $n > 0$ is the first index such that $|p_n| \geq q^*$, and by $\pm\Gamma$, depending on whether $p_n \geq q^*$ or $p_n \leq -q^*$. We will only consider predictors for which such an n exists, which is true with probability 1 thanks to Condition (v) in Section 2: indeed, if P_N is the probability for a path of length N , starting at $p_0 = p$, to satisfy $|p_i| < q$ for any $1 \leq i \leq N$, then $P_N \leq (1 - \epsilon_q)^N$ with $0 < \epsilon_q \leq 1$, so that $\lim_{N \rightarrow +\infty} P_N = 0$.

If we now set $P(p_0, \dots, p_n | p) = P(p_{t+i} = p_i, i \in [0, n] \mid p_t = p)$, this leads to the following equation:

$$g(p) = \sum_{n=0}^{\infty} \left[\int_{q^*}^{+\infty} \int_{-q^*}^{q^*} \dots \int_{-q^*}^{q^*} \left(\sum_{i=0}^{n-1} p_i + \Gamma \right) P(p_0, \dots, p_n | p) \prod_{i=0}^n dp_i \right. \\ \left. + \int_{-\infty}^{-q^*} \int_{-q^*}^{q^*} \dots \int_{-q^*}^{q^*} \left(\sum_{i=0}^{n-1} p_i - \Gamma \right) P(p_0, \dots, p_n | p) \prod_{i=0}^n dp_i \right]$$

Using now the fact that $g(q^*) = \Gamma$, we get:

$$\Gamma = \sum_{n=0}^{\infty} \left[\int_{q^*}^{+\infty} \int_{-q^*}^{q^*} \dots \int_{-q^*}^{q^*} \left(\sum_{i=0}^{n-1} p_i + \Gamma \right) P(p_0, \dots, p_n | q^*) \prod_{i=0}^n dp_i \right. \\ \left. + \int_{-\infty}^{-q^*} \int_{-q^*}^{q^*} \dots \int_{-q^*}^{q^*} \left(\sum_{i=0}^{n-1} p_i - \Gamma \right) P(p_0, \dots, p_n | q^*) \prod_{i=0}^n dp_i \right]$$

As we said above, there always exists, with probability 1, an integer n such that $|p_n| \geq q$, so:

$$1 = \sum_{n=0}^{\infty} \left[\int_{q^*}^{+\infty} + \int_{-\infty}^{-q^*} \right] \int_{-q^*}^{q^*} \dots \int_{-q^*}^{q^*} P(p_0, \dots, p_n | q^*) \prod_{i=0}^n dp_i.$$

which leads to:

$$\sum_{n=0}^{\infty} \left[\int_{q^*}^{+\infty} + \int_{-\infty}^{-q^*} \right] \int_{-q^*}^{q^*} \dots \int_{-q^*}^{q^*} \left(\sum_{i=0}^{n-1} p_i - 2\Gamma \cdot \mathbf{1}_{\{p_n < -q^*\}} \right) P(p_0, \dots, p_n | q^*) \prod_{i=0}^n dp_i = 0$$

where $\mathbf{1}$ is the indicator function.

This can be perhaps more gracefully expressed as a path integral: for a finite path $\phi : [0, n] \rightarrow \mathbb{R}$, we note $T_\phi = n$, $\phi_b = \phi(0)$, $\phi_e = \phi(n)$, $\mathcal{P}(\phi | p) = P(p_{t+z} = \phi(z), z \in [0, n] \mid p_t = p)$ and $\int_z \phi(z) dz = \sum_{i=0}^{n-1} \phi(i)$. The equation above can then be symbolically expressed as:

$$\boxed{\int_{\substack{|\phi_e| \geq q^* \\ \phi_b = q^* \\ -q^* < \phi(z) < q^*, z \in]0, T_\phi[}} \left[\int_z \phi(z) dz - 2\Gamma \cdot \mathbf{1}_{\{\phi_e \leq -q^*\}}(\phi) \right] P(\phi | q^*) \mathcal{D}\phi = 0} \quad (6)$$

Figure 1 sums up this reformulation of the problem: the value of q^* is such that the ‘‘penalty’’ 2Γ over all paths exiting through $-q^*$ is equal to the average gain (given by the sum of the values of the predictor) over all paths exiting either through q^* or $-q^*$.




Figure 1: Path integral representation

There is actually a direct interpretation for Equation (6), based on Figure 1, which is worth understanding. Imagine that we start with $p_t = p$ at time t , and the position before our trading decision is $\pi_{t-1} = -M$. Knowing how we will trade in the future (which depends on the optimal threshold q^*), we wonder whether it is worth reverting the position right now by buying $\Delta\pi = 2M$. Note that the reasoning below does not depend on the actual value of $\Delta\pi$, so one could as well consider buying just one share/contract.

Let us suppose that we do trade this quantity $\Delta\pi$, with a cost of $\Gamma \cdot \Delta\pi$. The next time we can possibly trade in the future is the first time $T > t$ such that $|p_T| \geq q^*$. If $p_T \geq q^*$ (path ϕ_1 on Figure 1), then we would have reverted our position anyway, so the cost of doing it early can be considered as null, and our gain is given $\Delta\pi$ times the values taken by the predictor, hence $\Delta\pi \int_z \phi(z) dz$. If on the contrary we have $p_T \leq -q^*$ (path ϕ_2 on Figure 1), then we will revert again our position to $-M$ by paying $\Gamma \cdot \Delta\pi$, hence a total cost of $2\Gamma \cdot \Delta\pi$, whereas the gain will also be given by $\Delta\pi \int_z \phi(z) dz$.

In the end, it is worth reverting the position to $\pi_t = M$ if, and only if:

$$\Delta\pi \cdot \int_{\substack{|\phi_e| \geq q^* \\ \phi_b = p \\ -q^* < \phi(z) < q^*, z \in]0, T_\phi[}} \left[\int_z \phi(z) dz - 2\Gamma \cdot \mathbf{1}_{\{\phi_e \leq -q^*\}}(\phi) \right] P(\phi|q^*) \mathcal{D}\phi \geq 0$$

As the optimal threshold determines exactly the limit between the trading zone and the no-trading zone, we recover Equation (6).

Hence, the rather simple problem we introduced in the present article has a very non-trivial solution, which is best described through the above path-integral formulation. Note that Equation (6) is completely general provided the assumptions of Section 2 are satisfied, it does not rely on any specific statistics of the predictor. In the next section, we will explicitly solve this equation when the predictor is Gaussian and follows a discrete Ornstein-Uhlenbeck evolution.

As a matter of notation, we will set:

$$\begin{aligned}\mathcal{L}(p) &= \int_{\substack{|\phi_e| \geq q^* \\ \phi_b = p \\ -q^* < \phi(z) < q^*, z \in]0, T_\phi[}} \left[\int_z \phi(z) dz \right] P(\phi|p) \mathcal{D}\phi \\ \mathcal{P}(p) &= \int_{\substack{\phi_e \leq -q^* \\ \phi_b = p \\ -q^* < \phi(z) < q^*, z \in]0, T_\phi[}} P(\phi|p) \mathcal{D}\phi,\end{aligned}$$

which can be interpreted, respectively, as the average contribution of all paths before exiting the channel $[-q^*, q^*]$, and as the probability for hitting the lower boundary $-q^*$ before the upper one q^* . In terms of these quantities, Equation (6) now writes:

$$\mathcal{L}(q^*) = 2\Gamma \cdot \mathcal{P}(q^*) \quad (7)$$

In some cases, both sides of this Equation will tend to be infinitesimal, so it is rather the ratio $\lim_{p \rightarrow q^*} \mathcal{L}(p) / \lim_{p \rightarrow q^*} \mathcal{P}(p)$ that we will ask to take the value 2Γ .

3.6 A note on non-linear predictability

By looking at the proof of Proposition 1, we note that the linearity of the predictability is not a crucial hypothesis. What we actually need is that $\mathfrak{L}_t(p_t) = E[r_t|p_t]$ satisfies the following properties, for any t :

- \mathfrak{L}_t is continuous and strictly increasing
- $\lim_{p_t \rightarrow \infty} \mathfrak{L}_t(p_t) > \Gamma$.

With these hypotheses it is possible to prove once again Proposition 1, except that Equation (2) becomes:

$$g(t, p) = \mathfrak{L}_t(p) + \Gamma \cdot [\mathcal{P}_{>q_{t+1}}(p) - \mathcal{P}_{<-q_{t+1}}(p)] + \int_{-q_{t+1}}^{q_{t+1}} P(p'|p) g(t+1, p') dp'$$

and the expression for $V_t(\pi, p)$ is similarly impacted.

If we want to consider the telescopic solution, then we need to have a predictability independent of t , that is: $\mathfrak{L}_t(p) = \mathfrak{L}(p)$. This gives, for the self-consistency equation:

$$\begin{aligned}g(p) &= \mathfrak{L}(p) + \Gamma \cdot [\mathcal{P}_{>q}(p) - \mathcal{P}_{<-q}(p)] + \int_{-q}^q P(p'|p) g(p') dp' \\ g(q) &= \Gamma\end{aligned}$$

This can again be solved using a path-integral formulation:

$$\int_{\substack{|\phi_e| \geq q^* \\ \phi_b = q^* \\ -q^* < \phi(z) < q^*, z \in]0, T_\phi[}} \left[\int_z \mathfrak{L}(\phi(z)) dz - 2\Gamma \cdot \mathbf{1}_{\{\phi_e \leq -q^*\}}(\phi) \right] P(\phi|q^*) \mathcal{D}\phi = 0$$

4 Application to an Ornstein-Uhlenbeck predictor

4.1 Definition

We will now focus on the case of a predictor following a discrete Ornstein-Uhlenbeck dynamics:

$$p_{t+1} - p_t = -\epsilon \cdot p_t + \beta \cdot \xi_t \quad (8)$$

where $(\xi_t)_{t \in \mathbb{R}}$ is a set of independent $\mathcal{N}(0, 1)$ Gaussian random variables.

One classical example of such a predictor is an exponential moving average of price returns:

$$p_t^{EMA} = K \sum_{t' < t} \rho^{t'-t-1} r_{t'}$$

If we suppose, as is usual, that the returns r_t are $\mathcal{N}(0, \sigma_r)$ random variables, then this gives an Ornstein-Uhlenbeck predictor with $\epsilon = 1 - \rho$ and $\beta = K\sigma_r$. Note however that the r_t must have some small correlations in order to be predictable! Therefore, in this case, the discussion in terms of an Ornstein-Uhlenbeck process is only consistent in the limit of small predictability, i.e. $K \ll 1$.

4.2 Properties & Orders of magnitude

Let us consider Equation (8) with the hypothesis that $\epsilon \ll 1$. Then we have $p_{t+1} \approx e^{-\epsilon} p_t + \beta \cdot \xi_t$, so that

$$E[p_{t+n}|p_t] \approx e^{-\epsilon n} p_t.$$

So, $\tau = \epsilon^{-1}$ is the **auto-correlation time** of the predictor p_t . The standard deviation of the predictor, i.e. its **average predictability**, $\sigma_p = \sqrt{E[p_t^2]}$, is given by $\beta/\sqrt{2\epsilon}$ (in the limit $\epsilon \ll 1$).

Hence:

- the smaller ϵ is, the longer the predictor takes to express itself
- the higher β is, the better the signal is (on average).

The integrated predictability is given by

$$\begin{aligned} p_\infty(p) &= \sum_{n=0}^{\infty} E[p_{t+n}|p_t] \approx \sum_{n=0}^{\infty} e^{-\epsilon n} p_t \\ &\approx p/\epsilon. \end{aligned}$$

This implies that the naïve threshold value is given by $q_{\text{naïve}} = \Gamma\epsilon$, while the integrated average predictability is:

$$\sigma_\infty = \frac{\beta}{\sqrt{2\epsilon^3}}$$

In practice, if a real price predictor is to be both useful and realistic, it should beat the trading costs when the predictor value is a few times its standard deviations. This allows to obtain a system which trades regularly, but not too often, compared to its auto-correlation time. For our Ornstein-Uhlenbeck predictor, this implies that when $p_t \propto \beta/\sqrt{2\epsilon}$, one should also have $p_t \epsilon^{-1} \propto \Gamma$. Therefore, the interesting regime for practical applications is:

$$\beta \propto \Gamma \epsilon^{3/2}.$$

In what follows, we will study the problem by distinguishing between two cases:

- If $\beta \gg \Gamma$, the predictor can easily beat its transaction costs at every step. This situation (which is not very realistic) requires us to keep a **discrete time** approach of the problem.
- If $\beta \ll \Gamma$, the predictor needs in general a large number of steps to beat the costs. This will lead us to a **continuous** formulation (and resolution) of the problem.

4.3 Discrete case: $\beta \gg \Gamma$

We already explained in Section 3.4 that we always have $q^* \leq \Gamma$. Consequently, whenever $\beta \gg \Gamma$, we also have $\beta \gg q^*$. This means that, starting at $p = q^*$, one will typically jump beyond q^* or $-q^*$ in just one step. Thus:

$$\begin{aligned} \mathcal{L}(q^*) &= q^* \\ \mathcal{P}(q^*) &= \int_{x^*}^{+\infty} \frac{e^{-x^2/2}}{\sqrt{2\pi}} dx \quad \text{with} \quad x^* = \frac{(2-\epsilon)q^*}{\beta} \end{aligned}$$

Since $\beta \gg q^*$, one has $x^* \ll 1$, and thus $\mathcal{P}(q^*) \approx 1/2$. Equation (7) finally gives:

$$q^* = \Gamma \tag{9}$$

Hence, if the volatility of each predictor change is very large compared to the trading costs, then one needs to be as selective as possible.

4.4 Continuous case: $\beta \ll \Gamma$

First, let us show why the condition $\beta \ll \Gamma$ leads us to express the problem in a continuous form. Under that condition, we cannot have an optimal threshold q^* of the same order as magnitude as β itself. Indeed, if this was the case, any time the predictor has the value q^* , it would have a significant probability to go below $-q^*$ at the next step since the predictor changes by an amount $\propto \beta$ at each time step. The optimal strategy would then require to resell everything at cost 2Γ , whereas the immediate gain would only be of the order of magnitude of β . Therefore, Equation (7) could not be satisfied.

Now, knowing that $q^* \gg \beta$, we need to evaluate $\mathcal{P}(q^*)$ and $\mathcal{L}(q^*)$. But for the predictor to go from $q^* \gg \beta$ to $-q^* \ll -\beta$ requires many steps. Therefore, one is effectively in the continuum limit, where the variation of the predictor at each time step is infinitesimal compared to q^* . We can then approximate the dynamics of the predictor by a drift-diffusion process:

$$dp = -\epsilon p dt + \beta dX_t \tag{10}$$

where $(X_t)_t$ is a Wiener process.

In such a continuous setting, the quantities $\mathcal{L}(q^*)$ and $\mathcal{P}(q^*)$ are actually ill-defined because the diffusion process starts on an absorbing boundary. This is a classical problem, which is handled by starting infinitesimally close to q^* . Therefore we consider $\mathcal{L}(p)$ and $\mathcal{P}(p)$ for $p < q^*$. It is easy to show that these two functions obey two **Kolmogorov backward equations**, that read:

$$\boxed{\frac{1}{2}\beta^2 \frac{\partial^2 \mathcal{L}}{\partial p^2} - \epsilon p \frac{\partial \mathcal{L}}{\partial p} = -p; \quad \frac{1}{2}\beta^2 \frac{\partial^2 \mathcal{P}}{\partial p^2} - \epsilon p \frac{\partial \mathcal{P}}{\partial p} = 0} \tag{11}$$

with boundary conditions: $\mathcal{L}(\pm q^*) = 0$ and $\mathcal{P}(q^*) = 0, \mathcal{P}(-q^*) = 1$.

4.4.1 Solution

Solving Equations (11) with their boundary conditions leads to:

$$\begin{aligned}\mathcal{L}(p) &= \frac{1}{\epsilon} \left(p - \frac{q}{I} \int_0^p e^{av^2} dv \right) \\ \mathcal{P}(p) &= \frac{1}{2} \left(1 - \frac{1}{I} \int_0^p e^{av^2} dv \right)\end{aligned}$$

with

$$I = \int_0^{q^*} e^{av^2} dv \quad \text{and} \quad a = \frac{\epsilon}{\beta^2}.$$

Setting now $p = q^* - u$ with $u \rightarrow 0$, one finds that Equation (7) becomes, to first order in u :

$$-\frac{u}{\epsilon} + \frac{uq^*}{\epsilon} \cdot \frac{e^{aq^{*2}}}{I} \approx \Gamma u \frac{e^{aq^{*2}}}{I}.$$

As expected, u disappears from the equation, to give the following solution for the threshold q^* :

$$\boxed{q^* = \frac{\beta}{\sqrt{\epsilon}} F^{-1} \left(\frac{\Gamma \epsilon^{3/2}}{\beta} \right) \quad \text{where} \quad F(x) = x - e^{-x^2} \int_0^x e^{v^2} dv} \quad (12)$$

Note that when $\epsilon \ll 1$, this equation can be expressed entirely in terms of the integrated predictability:

$$p_\infty(q^*) = \Gamma \cdot H \left(\frac{\sigma_\infty \sqrt{2}}{\Gamma} \right) \quad \text{where} \quad H(x) = x F^{-1} \left(\frac{1}{x} \right).$$

This means that we can find the optimal threshold for a predictor by studying only its total predictive power (if we suppose of course that it satisfies all the required properties).

4.4.2 Limits

One can now study the limits of Equation (12) for large and small values of the only remaining adimensional parameter $\eta = \Gamma \epsilon^{3/2} / \beta$. Interestingly, $\eta \sim 1$ is the regime mentioned above where predictability beats costs whenever the predictor's value is of the order of its rms.

The limiting behaviours of the function $F(x)$ are as follows:

- if $x \gg 1$, then $\int_0^x e^{v^2} dv \ll e^{x^2}$, so $F(x) \approx x$
- if $x \ll 1$, then $F(x) \simeq x - (1 - x^2) \int_0^x (1 - v^2) dv \approx \frac{2x^3}{3}$.

Therefore when $\eta \gg 1$, the threshold is simply given by $q^* = \Gamma \epsilon$. This result is rather intuitive: if β is very small then the predictability of the predictor is weak, compared to the trading cost. Hence, it makes sense to try to catch any profitable opportunity, without taking future trading into account. That is why we recover the naïve solution of Section (3.1).

If on the other hand $\beta \gg \Gamma \epsilon^{3/2}$ then $\eta \ll 1$, and $F^{-1}(\eta) \approx \sqrt[3]{\frac{3}{2} \cdot \eta}$, which finally gives

$$q^* = \sqrt[3]{\frac{3}{2} \cdot \Gamma \beta^2} \quad (13)$$

This is the result that would obtain with a predictor following a Brownian motion. Indeed, if β is large enough, the mean-reverting effect ϵ is not relevant, and the optimal threshold must consequently be independent of ϵ .

4.5 Shape of the global solution




Figure 2: Optimal threshold as a function of β

We summarize the various regimes in Figure 2, where we plot the optimal threshold q^* as a function of β , for ϵ and Γ fixed. One can see the three main characteristics of this solution:

- a constant threshold $q^* = \Gamma\epsilon$ for small values of $\beta \ll \Gamma\epsilon^{3/2}$
- a sublinear behaviour $q^* \propto \Gamma^{1/3}\beta^{2/3}$ in the intermediate regime between $\beta \sim \Gamma\epsilon^{3/2}$ and $\beta \sim \Gamma$
- a constant threshold $q^* = \Gamma$ for large values of $\beta \gg \Gamma$.

Note that these different results match seamlessly at the boundaries between the regimes. Indeed, $q^* = \sqrt[3]{\frac{3}{2} \cdot \Gamma\beta^2}$ becomes of the order of $\Gamma\epsilon$ when $\beta \sim \Gamma\epsilon^{3/2}$, and becomes of the order of Γ when $\beta \sim \Gamma$ is of the order of Γ . One can also compute systematic corrections to $q^* = \Gamma$ as an expansion in Γ/β :

$$q^* \approx \Gamma - (1 - \epsilon)\sqrt{\frac{2}{\pi}} \cdot \frac{\Gamma^2}{\beta} + \dots \quad \beta \gg \Gamma.$$

4.6 The case of a white noise predictor

To conclude, let us consider the special case of Equation (8) where $\epsilon = 1$. In that case, the predictor is a white noise in time: $E[p_t p_{t+1}] = 0$. Since we assume a perfect $p \mapsto -p$ symmetry, the self-consistency equation becomes simply $g(p) = p$, which trivially implies that $q^* = \Gamma$ in that case. This is also consistent with our explicit solutions above: when $\epsilon = 1$, the intermediate regime disappears and one indeed finds $q^* = \Gamma\epsilon = \Gamma$. This threshold is what we expect from such a system: without any auto-correlation, the best strategy is to trade as soon as the instantaneous predictability is above the trading cost. Note that in this case $p_\infty = p_t$, so this threshold also coincides with the naïve solution.

5 Numerical results

To check the robustness of our analytical results, we ran some simulations to determine the optimal threshold numerically, to be compared with the theoretical value we obtained in Section 4.

In what follows we set $\epsilon = 0.001$. We can set $\Gamma = 1$ without loss of generality, and the only remaining variable is β .

The algorithm to find the optimal threshold for a given value of β runs as follows:

- 1) We choose a set of threshold values q_1, \dots, q_n uniformly distributed over a reasonably large range (which contains the theoretical optimal threshold).
- 2) We generate a long random path $(p_t)_{t \in [0, T]}$ for the predictor, following the law given by Equation (8).
- 3) For each value of the threshold, we simulate the behaviour of the corresponding strategy, with a MaxPos of 1.
- 4) To obtain the P&L of each system, we calculate at each time step t the gain given by² $p_t \cdot \pi_t$ (where π_t is the position of the system) and the cost given by $\Gamma |\delta \pi_t|$ if there is a trade $\delta \pi_t$.
- 5) We select the threshold q_j with the maximal total P&L.
- 6) We choose new values q'_1, \dots, q'_n for the threshold, distributed around q_j , and we restart the algorithm with these values.

This loop is repeated several times, in order to get a sufficiently precise approximation of the optimal threshold. The result of this process is shown in Figures 3 and 4. The analytical solution in the continuous case, given by Equation (12), is easy to compute (the function $D(x) = e^{-x^2} \int_0^x e^{v^2} dv$ is a classic, called the **Dawson function**).

By comparing the two curves, one can check that the analytical solution is indeed a good fit of the simulation results. However, there are two interesting details to note:

- when the value of β becomes very small (Figure 3), the solution of the simulation becomes very noisy;
- there is a discrepancy between the analytical and simulated solutions when β increases (Figure 4) and, rather surprisingly, this happens in a regime where the inequality $\beta \ll \Gamma$ still seems to hold.

The reason for the first effect is obvious: when β becomes very small, then the predictor hardly ever beats its optimal threshold $\Gamma \epsilon$ during the course of the simulation, which implies that the real optimal threshold becomes very difficult to find: we need to use a huge number of steps in the simulation to find the proper solution.

The answer to the second remark is more subtle: if we go back to Section 4.4, we see that it is not really the hypothesis $\beta \ll \Gamma$ which implies that we are in a continuous setting, but rather the inequality $\beta \ll q^*$. Of course this second inequality was derived from the first one, but we did not

²Indeed, as we only consider expected values, there is no need to generate a random variable for the price return $r_t = \text{Price}_{t+1} - \text{Price}_t$ as a function of the predictor p_t : one can directly consider the mean of this variable, which is exactly p_t .




Figure 3: Optimal threshold for $0 < \beta < 1.2 \cdot 10^{-4}$




Figure 4: Optimal threshold for $0 < \beta < 0.003$

specify at that time the orders of magnitude involved in each inequality. In fact, if we consider for the threshold q^* the limit value $\sqrt[3]{\frac{3}{2} \cdot \Gamma \beta^2}$, one sees that $\frac{q^*}{\beta} = \kappa \sqrt[3]{\frac{\beta}{\Gamma}}$ with $\kappa \approx 1.15$. This means that our so-called continuum hypothesis breaks down much sooner than a naïve comparison between β and Γ would tell. In Figure 4, if we look at the point for which $\beta/\Gamma \approx 0.002$, one actually has $q^*/\Gamma \approx 0.144$; so the continuum hypothesis is already unwarranted for this value, and the solutions do not strictly coincide, even though the ratio between β and Γ is still extremely small.

6 Extension to other risk constraints

6.1 The “band” system

The problem we have presented and solved in the previous sections was rather specific, as it required a risk control based on the only constraint $|\pi| \leq M$. An alternative, more classical way to handle risk, is to consider a quadratic penalty $R(\pi_t) = \pi_t^2$ that represents risk aversion. The “utility” to be maximized is at each step t given by $g_t = r_t \pi_t - \lambda \pi_t^2$. Note that the penalty term could read, more

generally, $R(\pi_t) \propto |\pi_t/M|^z$. The case $z = 2$ is the above quadratic penalty, whereas the constraint on the maximum position formally corresponds to the limit $z \rightarrow \infty$.

The quadratic problem was considered in [MS11, Mar12], where it was showed that the optimal strategy in this setting is a **band**, also called a DT-NT-DT (Direct Trading - No Trading - Direct Trading) system. This system is defined as follows:

- there is a one-to-one correspondence between the values of the predictor and the position (in the quadratic case, this is simply given by: $\pi = p/2\lambda$)
- at each time step, one defines a band of size $2q^*$ around the value of the predictor
- if the image of the current position is inside the band then there is no trade; but if it is outside of the band then the trade should bring the image of the position to closest border of the band.

Note that we did not explicit here the position of the predictor p_t inside the band. It is actually proven in [MS11] that it does not have to be at the middle of the band, but it does when the trading cost becomes small. The band will be called **symmetric** in this case. This symmetric band is a property that will be necessary for our argument below to work.

This band system looks at first sight rather remote from the system we studied in the previous sections. Still it was shown in [MS11] that if the asset price follows a mean-reverting dynamics (which corresponds exactly to the case of a continuous Ornstein-Uhlenbeck predictor) then the optimal half-size of the band is given, with the notations of the present paper, by:

$$q^* = \sqrt[3]{\frac{3}{2} \cdot \Gamma \beta^2}$$

when the trading cost Γ is small. This is exactly the value of our threshold in the continuous setting, when costs are small!

Our goal in the next section is to explain that this is no coincidence. This will allow us not only to recover the results of [MS11], but also to extend them by giving the optimal solution under the condition that the band is symmetric.

6.2 Optimal size of the band

In what follows, we suppose that we already know that the optimal trading policy is a band system as described above, with the predictor being at the center of the band, and we will find the value of the half-band for this optimal system, under some condition on the current value of the predictor.

Consider that we start with $p_t = \mathbf{p}$, and the position before the next trading decision is at the lower border of the band: $\pi_{t-1} = (\mathbf{p} - q^*)/2\lambda$. Knowing that the future trading style is a band of optimal size $2q^*$, we wonder if it is worth buying an infinitesimal quantity $\delta\pi$ at t .

Let us suppose that we buy this quantity. The new risk penalty term will then be

$$\begin{aligned} R(\pi_t) &= \lambda(\pi_{t-1} + \delta\pi)^2 \\ &= R(\pi_{t-1}) + (\mathbf{p} - q^*) \cdot \delta\pi + O(\delta\pi^2) \end{aligned}$$

Now let us follow the scenarii given by Figure 5: by definition of the band system, we will not trade until the predictor's value becomes larger than \mathbf{p} or smaller than $\mathbf{p} - 2q^*$. In the former case,




Figure 5: Paths in the band system. The band is the region between $p - q$ and $p + q$, the paths ϕ_1 and ϕ_2 are the predictors trajectories defining the two possible ways to get out of the band if we start at the lower border.

then the trading cost is zero because we would have bought $\delta\pi$ anyway, whereas in the latter case we will have to sell it back, which gives a total cost of $2\Gamma \cdot \delta\pi$.

In both cases, if we note ϕ the path between t and the first T such that $p_T \geq \mathbf{p}$ or $p_T \leq \mathbf{p} - 2q^*$, then the gain of trading $\delta\pi$ is given by

$$\mathcal{G}_{\delta\pi} = \int_z (\phi(z) \cdot \delta\pi - \delta R) dz$$

where $\delta R = R(\pi_t) - R(\pi_{t-1}) \simeq (\mathbf{p} - q^*) \cdot \delta\pi$.

In the end, we have a positive gain by trading $\delta\pi$ if, and only if:

$$\delta\pi \cdot \int_{\substack{|\phi_e - \mathbf{p} + q^*| \geq q^* \\ \phi_b = \mathbf{p} \\ \mathbf{p} - 2q^* < \phi(z) < \mathbf{p}, z \in]0, T_\phi[}} \left[\int_z (\phi(z) - \mathbf{p} + q^*) dz - 2\Gamma \cdot \mathbf{1}_{\{\phi_e \leq \mathbf{p} - 2q^*\}}(\phi) \right] P(\phi|\mathbf{p}) \mathcal{D}\phi \geq 0$$

We can then do the change of variable $\psi(z) = \phi(z) - \mathbf{p} + q^*$, but in order to have $P(\psi|q^*) = P(\phi|\mathbf{p})$ we need to be in a case when the drift of the predictor can be neglected, i.e. when ϵ is small enough: $q^* \ll \beta/\sqrt{\epsilon}$. This gives (as $\delta\pi > 0$):

$$\int_{\substack{|\psi_e| \geq q^* \\ \psi_b = q^* \\ -q^* < \psi(z) < q^*, z \in]0, T_\psi[}} \left[\int_z \psi(z) dz - 2\Gamma \cdot \mathbf{1}_{\{\psi_e \leq -q^*\}}(\psi) \right] P(\psi|q^*) \mathcal{D}\psi \geq 0$$

The limit of the no-trading zone is the point where the cost of trading equilibrates precisely the gain, so the above inequality becomes an equality, and we recover exactly Equation (6)! This means that the same equation defines the optimal threshold in the quadratic risk-control case with small trading costs and in the ‘‘MaxPos’’ setting.

6.3 Consequences

As a consequence of our above result, we recover the optimal half-band size $q^* = \sqrt[3]{\frac{3}{2}} \cdot \Gamma \beta^2$ for small values of Γ , and in particular the two-third dependency on the costs explained in [Rog04]. This also gives us the optimal value for any Γ , provided that the band is symmetric and $q^* \ll \beta/\sqrt{\epsilon}$: this value will be given by Equation 12. This can be seen as a “variational solution”, where we get the optimal solution in a close-to-optimal subspace of the space of possible trading systems. But in reality, as shown in [MS11], there is a small shift of the center of the band, which is however of higher order in $\Gamma^{1/3}$.

Moreover, our result is more general than the case of an Ornstein-Uhlenbeck predictor: Equation (6) works for any price predictor, whatever its dynamics, provided it satisfies the hypotheses of Section 2. And the technique based on Kolmogorov backward equations, that we presented in Section 4.4, can easily be extended to these various dynamics.

Finally, it can be proven that for any reasonable risk-function $R(\pi)$ – for example $R(\pi) = |\pi|^z$ for any $z > 0$, if the trading system is a symmetric band then its half-size is again given by Equation (6). The limit $z \rightarrow \infty$ is singular in the sense that one loses the one-to-one correspondance between the predictor and the size of the trade, but formally our result holds for arbitrary z , and therefore applies to the “MaxPos” system considered in the previous sections.

7 Conclusion

We have considered and solved exactly the problem of the optimal trading strategy when one wants to follow a completely general Markovian predictor of the future returns of a single asset, in the presence of linear costs, and with a strict cap on the allowed position in the market. Using Bellman’s backward recursion method, we have shown that the optimal strategy is to switch between the maximum allowed long position and the maximum allowed short position, whenever the predictor exceeds a *threshold value*, for which we establish an exact, non-trivial, equation. This equation can be solved explicitly in the case of a discrete Ornstein-Uhlenbeck predictor. We discussed in detail the dependence of this threshold value on the transaction costs.

We also showed an unexpected relation between our problem and the problem where risk is handled dynamically, with an arbitrary risk penalty. The connection relies on the presence of a no-trading zone for the two problems, which allows a use of powerful Bellman techniques to calculate the optimal parameter.

There are various interesting extensions of our results that one can think of. One could consider the case of a predictor with jumps, and see how this affects the threshold value. One could consider an asymmetric risk constraint, where the maximum long and short positions are different (this could be relevant for option trading). But the most relevant extension would be to consider the case, important in practice, where the costs have both a linear component (coming from fees, bid-ask spread, etc.) and a quadratic component that would model impact. One may hope that an exact solution is still available in some regime, at least in the case of a single asset, or when the risk constraints do not couple different assets.

Acknowledgements

We thank Nicolas Bercot and Julianus Kockelkoren for fruitful discussions, Emeric Henry for reading the manuscript and Richard Martin for his precious comments and references.

References

- [AC00] R. F. Almgren and N. Chriss. Optimal execution of portfolio transactions. *Journal of Risk*, 3(2):5–39, 2000.
- [Alm03] R. F. Almgren. Optimal execution with nonlinear impact functions and trading-enhanced risk. *Applied Mathematical Finance*, 10(1):1–18, 2003.
- [Bel03] R. Bellman. *Dynamic Programming*. Dover Publications, March 2003.
- [Con86] G. M. Constantinides. Capital market equilibrium with transaction costs. *The Journal of Political Economy*, 94(4):842–862, 1986.
- [DL07] V. Durrleman and S. Lhermitte. Optimal arbitrage strategies in presence of market impacts. Technical report, 2007.
- [DN90] M. Davis and A. Norman. Portfolio selection with transaction costs. *Mathematics of Operations Research*, 15(4):676–713, October 1990.
- [Mar12] R. Martin. Optimal multifactor trading under proportional transaction costs. Technical report, 2012.
- [MS11] R. Martin and T. Schöneborn. Mean reversion pays, but costs. *Risk Magazine*, 2011.
- [Rog04] L.C.G. Rogers. *Why is the effect of proportional transaction costs $\mathcal{O}(\delta^{2/3})$?*, pages 303–308. AMS Contemporary Mathematics. 2004.
- [SS94] S E Shreve and H. M. Soner. Optimal investment and consumption with transaction costs. *The Annals of Applied Probability*, 4(3):609–692, 1994.
- [TLD⁺11] B. Tòth, Y. Lempérière, C. Deremble, J. de Lataillade, J. Kockelkoren, and J.-P. Bouchaud. Anomalous price impact and the critical nature of liquidity in financial mark. *Physical Review X*, 1:021006, 2011.