

NBER WORKING PAPER SERIES

DYNAMIC TRADING WITH PREDICTABLE RETURNS AND TRANSACTION
COSTS

Nicolae B. Garleanu
Lasse H. Pedersen

Working Paper 15205
<http://www.nber.org/papers/w15205>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2009

We are grateful for helpful comments from Darrell Duffie, Pierre Collin-Dufresne, Andrea Frazzini, Esben Hedegaard, Anthony Lynch, Ananth Madhavan (discussant), Andrei Shleifer, and Humbert Suarez, as well as from seminar participants at Stanford Graduate School of Business, University of California at Berkeley, Columbia University, NASDAQ OMX Economic Advisory Board Seminar, University of Tokyo, and the Journal of Investment Management Conference. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2009 by Nicolae B. Garleanu and Lasse H. Pedersen. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Dynamic Trading with Predictable Returns and Transaction Costs
Nicolae B. Garleanu and Lasse H. Pedersen
NBER Working Paper No. 15205
July 2009
JEL No. G11,G12

ABSTRACT

This paper derives in closed form the optimal dynamic portfolio policy when trading is costly and security returns are predictable by signals with different mean-reversion speeds. The optimal updated portfolio is a linear combination of the existing portfolio, the optimal portfolio absent trading costs, and the optimal portfolio based on future expected returns and transaction costs. Predictors with slower mean reversion (alpha decay) get more weight since they lead to a favorable positioning both now and in the future. We implement the optimal policy for commodity futures and show that the resulting portfolio has superior returns net of trading costs relative to more naive benchmarks. Finally, we derive natural equilibrium implications, including that demand shocks with faster mean reversion command a higher return premium.,

Nicolae B. Garleanu
Haas School of Business
F628
University of California, Berkeley
Berkeley, CA 94720
and NBER
garleanu@haas.berkeley.edu

Lasse H. Pedersen
NYU Stern Finance
44 West Fourth Street
Suite 9-190
New York, NY 10012
and NBER
lpederse@stern.nyu.edu

Active investors and asset managers — such as hedge funds, mutual funds, and proprietary traders — try to predict security returns and trade to profit from their predictions. Such dynamic trading often entails significant turnover and trading costs. Hence, any active investor must constantly weigh the expected excess return to trading against the risk and costs of trading. An investor often uses different return predictors, e.g., value and momentum predictors, and these have different prediction strengths and mean-reversion speeds, or, said differently, different “alphas” and “alpha decays.” The alpha decay is important because it determines how long the investor can enjoy high expected returns and, therefore, affects the trade-off between returns and transactions costs. For instance, while a momentum signal may predict that the IBM stock return will be high over the next month, a value signal might predict that Cisco will perform well over the next year. The optimal trading strategy must consider these dynamics.

This paper addresses how the optimal trading strategy depends on securities’ current expected returns, the evolution of expected returns in the future, their risks and correlations, and their trading costs. We present a closed-form solution for the optimal portfolio rebalancing rule taking these considerations into account.

The optimal trading strategy is intuitive: The best new portfolio is a combination of 1) the current portfolio (to reduce turnover), 2) the optimal portfolio in the absence of trading costs (to get part of the best current risk-return trade-off), and 3) the expected optimal portfolio in the future (a dynamic effect). Said differently, the best portfolio is a weighted average of the current portfolio and a “target portfolio” that combines portfolios 2) and 3).

Consistent with this decomposition, an investor facing transaction costs trades more aggressively on persistent signals than on fast mean-reverting signals: the benefits from the former accrue over longer periods, and are therefore larger. As is natural, transaction costs inhibit trading, both currently and in the future. Thus, target portfolios are conservative given the signals, and trading towards the target portfolio is slower when transaction costs are large.

The key role played by each return predictor’s mean reversion is an important implication

of our model. It arises because transaction costs imply that the investor cannot easily change his portfolio and, therefore, must consider his optimal portfolio both now and in the future. In contrast, absent transaction costs, the investor can re-optimize at no cost and needs to consider only the current investment opportunities (and possible hedging demands) without regard to alpha decay.

We first solve the model in discrete time. One may wonder, however, whether the length of the discrete-time interval between trading opportunities is important for the model, how different models with different interval lengths fit together, and what happens as this length approaches zero, that is, with continuous trading. To answer these questions, we present a continuous-time version of the model and show how the discrete-time solutions approach the continuous-time solution. An additional benefit of the continuous-time model is that the solution is even simpler, making applications of the model even easier.

As one such application, we embed the continuous-time model in an equilibrium setting. Rational investors facing transaction costs trade with several groups of noise traders who provide a time-varying excess supply or demand of assets. We show that, in order for the market to clear, the investors must be offered return premia depending on the properties of the noise-traders' positions. In particular, the noise trader positions that mean revert more quickly generate larger alphas in equilibrium, as the rational investors must be compensated for incurring more transaction costs per time unit. Long-lived supply fluctuations, on the other hand, give rise to smaller and more persistent alphas.

Finally, we illustrate our results empirically in the context of commodity futures markets. We use returns over the past 5 days, 12 months, and 5 years to predict returns. The 5-day signal is quickly mean reverting (fast alpha decay), the 12-month signal mean reverts more slowly, whereas the 5-year signal is the most persistent. We calculate the optimal dynamic trading strategy taking transaction costs into account and compare its performance to the optimal portfolio ignoring transaction costs and to a class of strategies that perform static (one-period) transaction-cost optimization. Our optimal portfolio performs the best net of transaction costs among all the strategies that we consider. Its net Sharpe ratio is about

20% better than that of the best strategy among all the static strategies. Our strategy's superior performance is achieved by trading at an optimal speed and by trading towards a target portfolio that is optimally tilted towards the more persistent return predictors.

We also study the impulse-response of the security positions following a shock to return predictors. While the no-transaction-cost position immediately jumps up and mean reverts with the speed of the alpha decay, the optimal position increases more slowly to minimize trading costs and, depending on the alpha decay speed, may eventually become larger than the no-transaction-cost position as the optimal position is sold more slowly.

Our paper is related to several large strands of literature. First, a large literature studies portfolio selection with return predictability in the absence of trading costs (see, e.g., Campbell and Viceira (2002) and references therein). A second strand of literature derives the optimal trade execution, treating what to trade as given exogenously (see, e.g., Perold (1988), Bertsimas and Lo (1998), Almgren and Chriss (2000), Obizhaeva and Wang (2006), and Engle and Ferstenberg (2007)). A third strand of literature, starting with Constantinides (1986), considers the optimal portfolio selection with trading costs, but without return predictability.¹ Constantinides (1986) considers a single risky asset in a partial equilibrium and studies trading-cost implications for the equity premium. Equilibrium models with trading costs include Amihud and Mendelson (1986), Vayanos (1998), Vayanos and Vila (1999), Lo, Mamaysky, and Wang (2004), Gârleanu (2009), and Acharya and Pedersen (2005), who also consider time-varying trading costs. Liu (2004) determines the optimal trading strategy for an investor with constant absolute risk aversion (CARA) and many independent securities with both fixed and proportional costs (without predictability). The assumptions of CARA and independence across securities imply that the optimal position for each security is independent of the positions in the other securities. In a fourth (and most related) strand of literature, using calibrated numerical solutions, trading costs are combined with incomplete

¹Davis and Norman (1990) provide a more formal analysis of Constantinides' model. Also, Gârleanu (2009) and Lagos and Rocheteau (2006) show how search frictions and payoff mean-reversion impact how close one trades to the static portfolio. Our continuous-time model with bounded-variation trading shares features with Longstaff (2001) and, in the context of predatory trading, by Brunnermeier and Pedersen (2005) and Carlin, Lobo, and Viswanathan (2008). See also Oehmke (2009).

markets by Heaton and Lucas (1996), and with predictability and time-varying investment opportunity by Balduzzi and Lynch (1999), Lynch and Balduzzi (2000), Jang, Koo, Liu, and Loewenstein (2007), and Lynch and Tan (2008). Grinold (2006) derives the optimal steady-state position with quadratic trading costs and a single predictor of returns per security. Like Heaton and Lucas (1996) and Grinold (2006), we also rely on quadratic trading costs.

We contribute to the literature in several ways. We provide a closed-form solution for a model with multiple correlated securities and multiple return predictors with different mean-reversion speeds, uncovering the role of alpha decay; derive new equilibrium implications; and demonstrate the model’s empirical importance using real data.

We end our discussion of the related literature by noting that quadratic programming techniques are also used in macroeconomics and other fields, and, usually, the solution comes down to algebraic matrix Riccati equations (see, e.g., Ljungqvist and Sargent (2004) and references therein). We solve our model explicitly, including the Riccati equations, in both discrete and continuous time.

The paper is organized as follows. Section 1 lays out a general discrete-time model, provides a closed-form solution, and presents various related results and examples. Section 2 solves the analogous continuous-time model and shows how it is approached by the discrete-time model as the time interval between trades becomes small. Section 3 studies the model’s equilibrium implications. Section 4 applies our framework to a trading strategy for commodity futures, and Section 5 concludes.

1 Discrete-Time Model

We first present the model, then solve it and provide additional results and examples.

1.1 General Discrete-Time Framework

We consider an economy with S securities traded at each time $t = 1, 2, 3, \dots$. The securities' price changes between times t and $t + 1$, $p_{t+1} - p_t$, are collected in a vector r_{t+1} given by

$$r_{t+1} = \mu_t + \alpha_t + u_{t+1}, \quad (1)$$

where μ_t is the “fair return,” e.g., from the CAPM, u_{t+1} is an unpredictable zero-mean noise term with variance $\text{var}_t(u_{t+1}) = \Sigma$, and α_t (alpha) is the predictable excess return, i.e., known to the investor already at time t , and given by

$$\alpha_t = Bf_t \quad (2)$$

$$\Delta f_{t+1} = -\Phi f_t + \varepsilon_{t+1}. \quad (3)$$

Here, f is a $K \times 1$ vector of factors that predict returns, B is a $S \times K$ matrix of factor loadings, Φ is a $K \times K$ positive-definite matrix of mean-reversion coefficients for the factors, and ε_{t+1} is the shock affecting the predictors with variance $\text{var}_t(\varepsilon_{t+1}) = \Omega$. Naturally, $\Delta f_{t+1} = f_{t+1} - f_t$ is the change in the factors.

The interpretation of these assumptions is straightforward: the investor analyzes the securities and his analysis results in forecasts of excess returns. The most direct interpretation is that the investor regresses the return on security s on factors f which could be past returns over various horizons, valuation ratios, and other return-predicting variables:

$$r_{t+1}^s = \mu_t^s + \sum_k \beta^{sk} f_t^k + u_{t+1}^s, \quad (4)$$

and thus estimates each variable's ability to predict returns as given by β^{sk} (collected in the matrix B). Alternatively, one can think of the factors as an analyst's overall assessment of a security (possibly based on a range of qualitative information) and B as the strength of these assessments in predicting returns.

We note that each factor f^k can in principle predict each security. However, the model

can easily be simplified to the special case in which there are different factors for different securities, as we discuss in Example 2 below. We further note that Equation (3) means that the factors and alphas mean-revert to zero. This is a natural assumption since an excess return that is always present should be viewed as compensation for risk, not reward for security analysis. Hence, such average returns are part of the fair return μ_t . That said, an intercept term can be accommodated, e.g., as a constant factor $f_t^1 = 1$.

Trading is costly in this economy and the transaction cost (TC) associated with trading $\Delta x_t = x_t - x_{t-1}$ shares is given by

$$TC(\Delta x_t) = \frac{1}{2} \Delta x_t^\top \Lambda \Delta x_t, \quad (5)$$

where Λ is a symmetric positive-definite matrix measuring the level of trading costs.² Trading costs of this form can be thought of as follows. Trading Δx_t shares moves the (average) price by $\frac{1}{2} \Lambda \Delta x_t$, and this results in a total trading cost of Δx_t times the price move, which gives TC . Hence, Λ (actually $1/2 \Lambda$ for convenience) is a multi-dimensional version of Kyle's lambda.

The investor's objective is to choose the dynamic trading strategy (x_0, x_1, \dots) to maximize the present value of all future expected alphas, penalized for risks and trading costs:

$$\max_{x_0, x_1, \dots} E_0 \left[\sum_t (1 - \rho)^t \left(x_t^\top \alpha_t - \frac{\gamma}{2} x_t^\top \Sigma x_t - \frac{1}{2} \Delta x_t^\top \Lambda \Delta x_t \right) \right], \quad (6)$$

where $\rho \in (0, 1)$ is a discount factor, and γ is the risk aversion coefficient.

There are several natural interpretations of this objective. First, we can envision an investor who is compensated based on his performance relative to a benchmark b . Under this interpretation, x_t is the deviation of the total portfolio, which we can denote by x^* , from the benchmark portfolio, that is, $x_t = x_t^* - b$. Hence, $x_t^\top \alpha$ measures the excess return over the benchmark, and $x_t^\top \Sigma x_t$ measures the variance of the tracking error relative to the

²The assumption that Λ is symmetric is without loss of generality. To see this, suppose that $TC(\Delta x_t) = \frac{1}{2} \Delta x_t^\top \bar{\Lambda} \Delta x_t$, where $\bar{\Lambda}$ is not symmetric. Then, letting Λ be the symmetric part of $\bar{\Lambda}$, i.e., $\Lambda = (\bar{\Lambda} + \bar{\Lambda}^\top)/2$, Λ generates the same trading costs as $\bar{\Lambda}$.

benchmark.

Another interpretation centers on a hedge fund manager who cares about his total net return (so x_t is the total portfolio), but is committed to achieving “alpha” (i.e., as opposed to making time-invariant bets based on constant risk premia).

A final interpretation concerns a “standard” investor who considers the return of his total portfolio (i.e., not just the alpha above the fair return μ or over a benchmark). Under this interpretation, x_t is the total portfolio, and we eliminate μ from (4), letting instead the first factor be constant, $f_t^1 = 1$, and $\beta^{s1} = \mu^s$. Thus, $\beta^{s1} f_t^1 = \mu^s$ captures the average security returns. It follows that α incorporates the entire return in the objective function and, in fact, this objective can be justified in a standard set-up with exponential utility for consumption and normally-distributed price changes, under certain conditions.

These interpretations are linked naturally. To see this, we can take the final interpretation and find an optimal “total portfolio” x_t^* (i.e., the solution with $f_t^1 = 1$, $\beta^{s1} = \mu^s$, and no μ term in (4), and compare it to the alpha-maximizing portfolio x_t from the first interpretation (i.e., the solution of (6) as is). It can be shown that, if $x_0^* = (\gamma\Sigma)^{-1}\mu + x_0$, then $x_t^* = (\gamma\Sigma)^{-1}\mu + x_t$ for all t , that is, the total-return-maximizing portfolio x_t^* is the “benchmark” plus the optimal deviation x_t from this benchmark, where the benchmark is given as the constant Markowitz portfolio relative to the average returns μ .

1.2 Solution and Results

We solve the model using dynamic programming. We start by introducing a value function $V(x_{t-1}, f_t)$ measuring the value of entering period t with a portfolio of x_{t-1} securities and observing return-predicting factors f_t . The value function solves the Bellman equation:

$$V(x_{t-1}, f_t) = \max_{x_t} \left\{ x_t^\top \alpha_t - \frac{\gamma}{2} x_t^\top \Sigma x_t - \frac{1}{2} \Delta x_t^\top \Lambda \Delta x_t + (1 - \rho) E_t[V(x_t, f_{t+1})] \right\}. \quad (7)$$

We guess, and later verify, that the solution has a quadratic form:

$$V(x_t, f_{t+1}) = -\frac{1}{2} x_t^\top A_{xx} x_t + x_t^\top A_{xf} f_{t+1} + \frac{1}{2} f_{t+1}^\top A_{ff} f_{t+1} + a_0, \quad (8)$$

where we need to derive the scalar a_0 , the symmetric matrices A_{xx} and A_{ff} , and the matrix A_{xf} .

The model in its most general form can be solved explicitly as we state in the following proposition. The expressions for the coefficient matrices (A_{xx}, A_{xf}) are somewhat long, so we leave them in the Appendix, but they become simple in the special cases discussed below, and, in continuous time, they are relatively simple even in the most general case.

Proposition 1 *The optimal dynamic portfolio x_t is a “matrix-weighted average” of the current position and a target portfolio:*

$$x_t = (I - \Lambda^{-1}A_{xx})x_{t-1} + \Lambda^{-1}A_{xx} \text{target}_t, \quad (9)$$

with

$$\text{target}_t = A_{xx}^{-1}A_{xf}f_t. \quad (10)$$

The matrix A_{xx} is positive definite; A_{xx} and A_{xf} are stated explicitly in (A.13) and (A.19).

An alternative characterization of the optimal portfolio is a weighted average of the current position, x_{t-1} , the optimal position in the absence of transaction costs, $(\gamma\Sigma)^{-1}Bf_t$, and the expected target next period, $E_t(\text{target}_{t+1}) = A_{xx}^{-1}A_{xf}(I - \Phi)f_t$:

$$\begin{aligned} x_t &= [\Lambda + \gamma\Sigma + (1 - \rho)A_{xx}]^{-1} \\ &\times [\Lambda x_{t-1} + \gamma\Sigma((\gamma\Sigma)^{-1}Bf_t) + (1 - \rho)A_{xx}(A_{xx}^{-1}A_{xf}(I - \Phi)f_t)]. \end{aligned} \quad (11)$$

The proposition provides expressions for the optimal portfolio that are natural and relatively simple. The optimal trade Δx_t follows directly from the proposition as

$$\Delta x_t = \Lambda^{-1}A_{xx}(\text{target}_t - x_t). \quad (12)$$

The optimal trade is proportional to the difference between the current portfolio and the target portfolio, and the trading speed decreases in the trading cost Λ .

We discuss the intuition behind the result further under the additional assumption that

$\Lambda = \lambda \Sigma$ for some number $0 < \lambda \in \mathbb{R}$, which simplifies the solution further. This means that the trading-cost matrix is proportional to the return variance-covariance matrix. This trading cost is natural and, in fact, implied by the model of dealers in Gârleanu, Pedersen, and Poteshman (2008). To understand this, suppose that a dealer takes the other side of the trade Δx_t for a single period and can “lay it off” thereafter, and that alpha is zero conditional on the dealer’s information. Then the dealer’s risk is $\Delta x_t^\top \Sigma \Delta x_t$ and the trading costs is the dealer’s compensation for risk, depending on the dealer’s risk aversion λ . Under this assumption, we derive the following simple and intuitive optimal trading strategy.

Proposition 2 *When the trading cost is proportional to the amount of risk, $\Lambda = \lambda \Sigma$, then the optimal new portfolio x_t is a weighted average of the current position x_{t-1} and a moving “target portfolio”*

$$x_t = \left(1 - \frac{a}{\lambda}\right) x_{t-1} + \frac{a}{\lambda} \text{target}_t \quad (13)$$

where $\frac{a}{\lambda} < 1$ and

$$\text{target}_t = (\gamma \Sigma)^{-1} B \left(I + \frac{a(1-\rho)}{\gamma} \Phi \right)^{-1} f_t \quad (14)$$

$$a = \frac{-(\gamma + \lambda \rho) + \sqrt{(\gamma + \lambda \rho)^2 + 4\gamma \lambda (1 - \rho)}}{2(1 - \rho)} \quad (15)$$

The target is the optimal position in the absence of trading costs if the return-predictability coefficients were $B \left(I + \frac{a(1-\rho)}{\gamma} \Phi \right)^{-1}$ instead of B .

Alternatively, x_t is a weighted average of the current position, x_{t-1} , the optimal position in the absence of trading costs, $\text{static}_t = (\gamma \Sigma)^{-1} B f_t$, and the expected target in the future, $E_t(\text{target}_{t+1}) = (\gamma \Sigma)^{-1} B \left(I + \frac{a(1-\rho)}{\gamma} \Phi \right)^{-1} (I - \Phi) f_t$:

$$x_t = \frac{\lambda}{\lambda + \gamma + (1 - \rho)a} x_{t-1} + \frac{\gamma}{\lambda + \gamma + (1 - \rho)a} \text{static}_t + \frac{(1 - \rho)a}{\lambda + \gamma + (1 - \rho)a} E_t(\text{target}_{t+1}). \quad (16)$$

This result provides a simple and appealing trading rule. Equation (13) states that the optimal portfolio is between the existing one and an optimal target, where the weight on the

target a/λ decreases in trading costs λ because higher trading costs imply that one must trade more slowly. The weight on the target increases in γ because a higher risk aversion means that it is more important not to let one's position stray too far from its optimal level.

The alternative characterization (16) provides a similar intuition and comparative statics, and separates the target into the current Markowitz static optimal position without transaction costs and the expected future target. The weight on the future target is small if the trading cost λ is small (because this makes a small) or if the agent is very impatient such that ρ is close to 1. We note that while the weights on the current position x_{t-1} appear different in (13) and (16), they are, naturally, the same. The optimal trading policy is illustrated in Figure 1.

The optimal trading is simpler yet under the additional (and rather standard) assumption that the mean reversion of each factor f^k only depends on its own level (not the level of the other factors), that is, $\Phi = \text{diag}(\phi^1, \dots, \phi^K)$ is diagonal, so that Equation (3) simplifies to scalars:

$$\Delta f_{t+1}^k = -\phi^k f_t^k + \varepsilon_{t+1}^k. \quad (17)$$

Under these assumptions we have:

Proposition 3 *If $\Lambda = \lambda\Sigma$ and $\Phi = \text{diag}(\phi^1, \dots, \phi^K)$, then the optimal portfolio is the weighted average (13) of the current portfolio x_{t-1} and a target portfolio, which is the optimal portfolio without trading costs with each factor f_t^k scaled depending on its alpha decay ϕ^k :*

$$\text{target}_t = (\gamma\Sigma)^{-1}B \left(\frac{f_t^1}{1 + \phi^1(1 - \rho)a/\gamma}, \dots, \frac{f_t^K}{1 + \phi^K(1 - \rho)a/\gamma} \right)^\top \quad (18)$$

We see that the target portfolio is very similar to the optimal portfolio without transaction costs $(\gamma\Sigma)^{-1}Bf_t$. The transaction costs imply first that one optimally only trades part of the way towards the target, and, second, that the target down-weights each return-predicting factor more the higher is its alpha decay ϕ^k . Down-weighting factors reduces the size of the position, and, more importantly, changes the relative importance of the different factors as

illustrated in Figure 2. Naturally, giving more weight to the more persistent factors means that the investor trades towards a portfolio that not only has a high alpha now, but also is expected to have a high alpha for a longer time in the future.

We next provide a few examples.

Example 1: Timing a single security

An interesting and simple case is when there is only one security. This occurs when an investor is timing his long or short view of a particular security or market. In this case, the assumption that $\Lambda = \lambda\Sigma$ from Propositions 2–3 is without loss of generality since all parameters are just scalars. In the scalar case, we use the notation $\Sigma = \sigma^2$ and $B = (\beta^1, \dots, \beta^K)$. Assuming that Φ is diagonal, we can apply Proposition 3 directly to get the optimal timing trade:

$$x_t = \left(1 - \frac{a}{\lambda}\right) x_{t-1} + \frac{a}{\lambda} \frac{1}{\gamma\sigma^2} \sum_{i=1}^K \frac{1}{1 + \phi^i(1 - \rho)a/\gamma} \beta^i f_t^i. \quad (19)$$

Example 2: Relative-value trades based on security characteristics

It is natural to assume that the agent uses certain characteristics of each security to predict its returns. Hence, each security has its own return-predicting factors (whereas, in the general model above, all the factors could influence all the securities). For instance, one can imagine that each security is associated with a value characteristic (e.g., its own book-to-market) and a momentum characteristic (its own past return). In this case, it is natural to let the alpha for security s be given by

$$\alpha_t^s = \sum_i \beta^i f_t^{i,s}, \quad (20)$$

where $f_t^{i,s}$ is characteristic i for security s (e.g., IBM's book-to-market) and β^i be the predictive ability of characteristic i (i.e., how book-to-market translates into future expected return, for any security), which is the same for all securities s . Further, we assume that

characteristic i has the same mean-reversion speed for each security, that is, for all s ,

$$\Delta f_{t+1}^{i,s} = -\phi^i f_t^{i,s} + \varepsilon_{t+1}^{i,s}. \quad (21)$$

We collect the current values of characteristic i for all securities in a vector $f_t^i = (f_t^{i,1}, \dots, f_t^{i,S})^\top$, e.g., the book-to-market of security 1, book-to-market of security 2, etc.

This setup based on security characteristics is a special case of our general model. To map it into the general model, we stack all the various characteristic vectors on top of each other into f :

$$f_t = \begin{pmatrix} f_t^1 \\ \vdots \\ f_t^I \end{pmatrix}. \quad (22)$$

Further, we let $I_{S \times S}$ be the S -by- S identity matrix and can express B using the kronecker product:

$$B = \beta^\top \otimes I_{S \times S} = \begin{pmatrix} \beta^1 & 0 & 0 & & \beta^I & 0 & 0 \\ 0 & \ddots & 0 & \dots & 0 & \ddots & 0 \\ 0 & 0 & \beta^1 & & 0 & 0 & \beta^I \end{pmatrix}. \quad (23)$$

Thus, $\alpha_t = B f_t$. Also, let $\Phi = \text{diag}(\phi \otimes 1_{S \times 1}) = \text{diag}(\phi^1, \dots, \phi^1, \dots, \phi^I, \dots, \phi^I)$. With these definitions, we apply Proposition 3 to get the optimal characteristic-based relative-value trade as

$$x_t = \left(1 - \frac{a}{\lambda}\right) x_{t-1} + \frac{a}{\lambda} (\gamma \Sigma)^{-1} \sum_{i=1}^I \frac{1}{1 + \phi^i (1 - \rho) a / \gamma} \beta^i f_t^i. \quad (24)$$

Example 3: Static model

When the investor completely discounts the future, i.e., $\rho = 1$, he only cares about the

current period and the problem is static. The investor simply solves

$$\max_{x_t} x_t^\top \alpha_t - \frac{\gamma}{2} x_t^\top \Sigma x_t - \frac{\lambda}{2} \Delta x_t^\top \Sigma \Delta x_t \quad (25)$$

with a solution that specializes Proposition 2:

$$x_t = \frac{\lambda}{\gamma + \lambda} x_{t-1} + \frac{\gamma}{\gamma + \lambda} (\gamma \Sigma)^{-1} \alpha_t. \quad (26)$$

To recover the optimal dynamic weight on the current position x_{t-1} from (16), one must lower the trading cost λ to $\frac{1}{1+(1-\rho)a/\gamma} \lambda$ to account for the future benefits of the position. Alternatively, one can increase risk aversion, or do some combination.

Interestingly, however, with multiple return-predicting factors, no choice of risk aversion γ and trading cost λ recovers the dynamic solution. This is because the static solution treats all factors the same, while the dynamic solution gives more weight to factors with slower alpha decay. We show empirically in Section 4 that even the best choice of γ and λ in a static model may perform significantly worse than our dynamic solution.

To recover the dynamic solution in a static setting, one must change not just γ and λ , but additionally the alphas $\alpha_t = B f_t$ by changing B as described in Propositions 2–3.

Example 4: Today's first signal is tomorrow's second signal

Suppose that the investor is timing a single market using each of the several past daily returns to predict the next return. In other words, the first signal f_t^1 is the daily return for yesterday, the second signal f_t^2 is the return the day before yesterday, and so on, so that the last signal used yesterday is ignored today. In this case, the trader already knows today what some of her signals will look like in the future. Today's yesterday is tomorrow's day-before-yesterday:

$$\begin{aligned} f_{t+1}^1 &= \varepsilon_{t+1}^1 \\ f_{t+1}^k &= f_t^{k-1} \quad \text{for } k > 1 \end{aligned}$$

The matrix Φ is therefore not diagonal, but has the form

$$I - \Phi = \begin{pmatrix} 0 & & & 0 \\ 1 & 0 & & \\ & \ddots & \ddots & \\ 0 & & 1 & 0 \end{pmatrix}.$$

Suppose for simplicity that all signals are equally important for predicting returns $B = (\beta, \dots, \beta)$ and use the notation $\Sigma = \sigma^2$. Then we can use Proposition 2 to get the optimal trading strategy

$$\begin{aligned} x_t &= \left(1 - \frac{a}{\lambda}\right) x_{t-1} + \frac{1}{\sigma^2} B((\gamma + \lambda + (1 - \rho)a)I - \lambda(1 - \rho)(I - \Phi))^{-1} f_t \\ &= \left(1 - \frac{a}{\lambda}\right) x_{t-1} + \frac{\beta}{\sigma^2 (\gamma + \lambda + (1 - \rho)a)} \sum_{k=1}^K (1 - z^{K+1-k}) f_t^k \\ &= \left(1 - \frac{a}{\lambda}\right) x_{t-1} + \frac{\beta(\lambda - a)}{\lambda^2 \sigma^2} \sum_{k=1}^K (1 - z^{K+1-k}) f_t^k \end{aligned}$$

where $z = \frac{\lambda(1-\rho)}{\gamma+\lambda+(1-\rho)a} < 1$. Hence, the optimal portfolio gives the largest weight to the first signal (yesterday's return), the second largest to the second signal, and so on. This is intuitive, since the first signal will continue to be important the longest, the second signal the second longest, and so on.

2 Continuous-Time Model

We next present the continuous-time version of our model. The continuous-time model is convenient since it has an even simpler solution and, therefore, it constitutes a useful workhorse model for applications — e.g., our equilibrium analysis. We show below that the continuous-time model obtains naturally as the limit of discrete-time models.

The securities have prices p with dynamics

$$dp_t = (\mu_t + \alpha_t) dt + du_t \quad (27)$$

where, as before, μ_t is the “fair return,” the random “noise” u is a martingale (e.g., a Brownian motion) with drift zero and instantaneous variance covariance matrix $\text{var}_t(du_t) = \Sigma dt$, and the predictable return α is given by

$$\alpha_t = Bf_t \quad (28)$$

$$df_t = -\Phi f_t dt + d\varepsilon_t. \quad (29)$$

The vector f contains the factors that predict returns, B contains the factor loadings, Φ is the matrix of mean-reversion coefficients, and the noise term ε is a martingale (e.g., a Brownian motion) with drift zero and instantaneous variance-covariance matrix $\text{var}_t(d\varepsilon_t) = \Omega dt$.

The agent chooses his trading intensity $\tau_t \in \mathbb{R}^S$, which determines the rate of change³ of his position x_t :

$$dx_t = \tau_t dt. \quad (30)$$

The cost per time unit of trading τ_t shares per time unit is

$$TC(\tau_t) = \frac{1}{2} \tau_t^\top \Lambda \tau_t \quad (31)$$

and the investor chooses his optimal trading strategy to maximize the present value of the future stream of alphas, penalized for risk and trading costs:

$$\max_{(\tau_s)_{s \geq t}} E_t \int_t^\infty e^{-\rho(s-t)} \left(x_s^\top \alpha_s - \frac{\gamma}{2} x_s^\top \Sigma x_s - \frac{1}{2} \tau_s^\top \Lambda \tau_s \right) ds. \quad (32)$$

³We only consider smooth portfolio policies because discreet jumps in positions or quadratic variation would be associated with infinite trading costs in our setting. E.g., if the agent trades n shares over a time period of Δ_t , then the cost is $\int_0^{\Delta_t} TC(\frac{n}{\Delta_t}) dt = \frac{1}{2} \Lambda \frac{n^2}{\Delta_t}$ which approaches infinity as Δ_t approaches 0.

The value function $V(x, f)$ of the investor solves the Hamilton-Jacoby-Bellman (HJB) equation

$$\rho V = \sup_{\tau} \left\{ x^{\top} B f - \frac{\gamma}{2} x^{\top} \Sigma x - \frac{1}{2} \tau^{\top} \Lambda \tau + \frac{\partial V}{\partial x} \tau + \frac{\partial V}{\partial f} (-\Phi f) + \frac{1}{2} \text{tr} \left(\Omega \frac{\partial^2 V}{\partial f \partial f^{\top}} \right) \right\}. \quad (33)$$

Maximizing this expression with respect to the trading intensity results in

$$\tau = \Lambda^{-1} \frac{\partial V}{\partial x}^{\top}.$$

It is natural to conjecture a quadratic form for the value function:

$$V(x, f) = -\frac{1}{2} x^{\top} A_{xx} x + x^{\top} A_{xf} f + \frac{1}{2} f^{\top} A_{ff} f + A_0.$$

We verify the conjecture as part of the proof to the following proposition.

Proposition 4 *The optimal portfolio x_t tracks a moving “target portfolio” $A_{xx}^{-1} A_{xf} f_t$ with a tracking speed of $\Lambda^{-1} A_{xx}$. That is, the optimal trading intensity $\tau_t = \frac{dx_t}{dt}$ is*

$$\tau_t = \Lambda^{-1} A_{xx} (A_{xx}^{-1} A_{xf} f_t - x_t), \quad (34)$$

where the positive definite matrix A_{xx} and the matrix A_{xf} are given by

$$A_{xx} = -\frac{\rho}{2} \Lambda + \Lambda^{\frac{1}{2}} \left(\gamma \Lambda^{-\frac{1}{2}} \Sigma \Lambda^{-\frac{1}{2}} + \frac{\rho}{4} I \right)^{\frac{1}{2}} \Lambda^{\frac{1}{2}} \quad (35)$$

$$\text{vec}(A_{xf}) = (\rho I + \Phi^{\top} \otimes I_K + I_S \otimes (A_{xx} \Lambda^{-1}))^{-1} \text{vec}(B). \quad (36)$$

As in discrete time, the optimal trading strategy has a particularly simple form when trading costs are proportional to the variance of the fundamentals:

Proposition 5 *If trading costs are proportional to the amount of risk, $\Lambda = \lambda \Sigma$, then the optimal trading intensity $\tau_t = \frac{dx_t}{dt}$ is*

$$\tau_t = \frac{a}{\lambda} (\text{target}_t - x_t) \quad (37)$$

with

$$\text{target} = (\gamma\Sigma)^{-1} B \left(I + \frac{a}{\gamma} \Phi \right)^{-1} f_t \quad (38)$$

$$a = \frac{-\rho\lambda + \sqrt{\rho^2\lambda^2 + 4\gamma\lambda}}{2}. \quad (39)$$

In words, the optimal portfolio x_t tracks target_t with speed $\frac{a}{\lambda}$. The tracking speed decreases with the trading cost λ and increases with the risk-aversion coefficient γ .

If each factor's alpha decay only depends on itself, $\Phi = \text{diag}(\phi^1, \dots, \phi^K)$, then the target is the optimal portfolio without transaction costs with each return-predicting factor f_t down-weighted more the higher is the trading cost λ and the higher is its alpha decay speed ϕ^k :

$$\text{target}_t = (\gamma\Sigma)^{-1} B \left(\frac{f_t^1}{1 + a\phi^1/\gamma}, \dots, \frac{f_t^K}{1 + a\phi^K/\gamma} \right)^\top. \quad (40)$$

When the agent is very patient, that is, $\rho = 0$, the expressions are even simpler. The coefficient a is simply $a = \sqrt{\gamma\lambda}$, and the tracking speed is $\frac{a}{\lambda} = \sqrt{\frac{\gamma}{\lambda}}$ which clearly decreases with trading costs λ and increases with risk aversion γ .

2.1 Connection between Discrete and Continuous Time

The continuous-time model, and therefore solution, are readily seen to be the limit of their discrete-time analogues when parameters are chosen consistently, adjusted for the length of the time interval between successive trading opportunities.

Proposition 6 *Consider the discrete-time model of Section 1 with parameters defined to*

depend on the time interval Δ_t in the following way:

$$\hat{\Sigma}(\Delta_t) = \Sigma \Delta_t \quad (41)$$

$$\hat{\Omega}(\Delta_t) = \Omega \Delta_t \quad (42)$$

$$\hat{\Lambda}(\Delta_t) = \Delta_t^{-1} \Lambda \quad \text{or} \quad \hat{\lambda}(\Delta_t) = \Delta_t^{-2} \lambda \quad (43)$$

$$\hat{B}(\Delta_t) = B \Delta_t \quad (44)$$

$$\hat{\Phi}(\Delta_t) = 1 - e^{-\Phi \Delta_t} \quad (45)$$

$$\hat{\rho}(\Delta_t) = 1 - e^{-\rho \Delta_t} \quad (46)$$

$$\hat{\gamma}(\Delta_t) = \gamma. \quad (47)$$

Then, given the initial position x_0 , the discrete-time solution converges to the continuous-time solution as Δ_t approaches zero: The optimal discrete-time position converges to the continuous-time one, i.e., $\hat{x}_t \rightarrow x_t$ a.s., as does the optimal trade per time unit, i.e., $\Delta \hat{x}_t / \Delta_t \rightarrow \tau_t$ a.s.

We note that Equations (41)–(42) simply state that the variance is proportional to time. The adjustment to the trading cost in Equation (43) is different for the following reason. Suppose that one can trade twice as frequently and consider trading over two time periods. The same total amount as previously can be traded now by splitting the order in half. With a quadratic trading cost, this leads to a total trading cost over the two periods of $2 \cdot TC(\Delta x / 2) = 2 \cdot TC(\Delta x) / 4 = TC(\Delta x) / 2$. Hence, in order for the total trading costs to be independent of the trading frequency, Λ must double when the frequency doubles, explaining the equation for Λ . Another way to say this is that the trading cost over 1 time period should depend on the intensity of trade $\Delta \hat{x}_t / \Delta_t$ and the length of the time period Δ_t so that $TC = \Delta_t \frac{\Delta \hat{x}_t}{\Delta_t} \Lambda \frac{\Delta \hat{x}_t}{\Delta_t} = \hat{x}_t' \frac{\Lambda}{\Delta_t} \hat{x}_t = \hat{x}_t' \hat{\Lambda} \hat{x}_t$, and this means that $\hat{\Lambda} = \frac{\Lambda}{\Delta_t}$. When trading costs are proportional to Σ , the equation for λ simply follows from the previous analysis and $\Lambda = \lambda \Sigma$.

2.2 Persistent Price Impact

In some cases trading may have a significant persistent price impact in addition to the transitory trading cost that we have studied so far. For this, we consider an investor that can transact at a price $\bar{p}_t = p_t + D_t$ by paying a transitory trading cost TC . Here, p is the price without the effect of the investor's own trading (as before), TC is as before, and the new term D_t captures the accumulated price distortion due to the investor's trades. Trading with intensity τ pushes prices by $C^\top \tau$, and the price distortion mean reverts at a speed (or “resiliency”) R :

$$dD_t = -RD_t dt + C^\top \tau_t dt \quad (48)$$

The investor's objective is as before (i.e., (32)), but now the securities' alpha (i.e., expected return $E_t(d\bar{p}_t)$) incorporates both the effect of predictability of p by the factors f_t and of the predictability due to price distortions (dD_t):

$$\alpha_t = Bf_t - RD_t + C^\top \tau_t. \quad (49)$$

The value function now becomes quadratic in the extended state variable (x, D, f) :

$$V(x, f, D) = -\frac{1}{2}x^\top A_{xx}x + x^\top A_{xf}f + \frac{1}{2}f^\top A_{ff}f + x^\top A_{xD}D + f^\top A_{fD}D + \frac{1}{2}D^\top A_{DD}D + A_0.$$

We solve the HJB equation as before.

Proposition 7 *The optimal portfolio x_t tracks a moving “target portfolio” with a tracking speed of $\Lambda^{-1}(A_{xx} - CA_{Dx} - C)$. That is, the optimal trading intensity $\tau_t = \frac{dx_t}{dt}$ is*

$$\tau_t = \Lambda^{-1}(A_{xx} - CA_{Dx} - C)[target_t - x_t], \quad (50)$$

with

$$target_t = (A_{xx} - CA_{Dx} - C)^{-1}((A_{xD} + CA_{DD})D_t + (A_{xf} + CA_{Df})f_t). \quad (51)$$

where the coefficient matrices A solve (A.41) in the appendix.

We see that the optimal trading policy has a similar structure to before, but the persistent price impact changes both the speed of trading and the target portfolio. One can naturally also solve the discrete time model with persistent price impact in the same way, and the result is analogous.⁴

3 Equilibrium Implications

In this section we study the restrictions placed on a security's return properties by the market equilibrium. More specifically, we consider a situation in which an investor facing transaction costs absorbs a residual supply specified exogenously and analyze the relationship implied between the characteristics of the supply dynamics and the return alpha.

For simplicity, we consider a model with one security in which $L \geq 1$ groups of (exogenously given) noise traders hold positions z_t^l (net of the aggregate supply) given by

$$dz_t^l = \kappa (f_t^l - z_t^l) dt \quad (52)$$

$$df_t^l = -\psi_l f_t^l dt + dW_t^l. \quad (53)$$

In addition, the Brownian motions W^l satisfy $\text{var}_t(dW_t^l)/dt = \Omega_{ll}$. It follows that the aggregate noise-trader holding, $z_t = \sum_l z_t^l$, satisfies

$$dz_t = \kappa \left(\sum_{l=1}^L f_t^l - z_t \right) dt. \quad (54)$$

We conjecture that the investor's inference problem is as studied in Section 2, where f given by $f \equiv (f^1, \dots, f^L, z)$ is a linear return predictor and B is to be determined. We verify the conjecture and find B as part of Proposition 8 below.

⁴One can also solve the model without transitory trading costs, $\Lambda = 0$. In this case, the optimal trading involves infinite turnover (non-zero quadratic variation) because buying and immediately selling is not costly when the only friction is persistent price impact in continuous time. It therefore appears more realistic to have non-zero transitory trading shocks.

Given the definition of f , the mean-reversion matrix Φ is given by

$$\Phi = \begin{pmatrix} \psi_1 & 0 & \cdots & 0 \\ 0 & \psi_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ -\kappa & -\kappa & \cdots & \kappa \end{pmatrix}. \quad (55)$$

Suppose that the only other investors in the economy are the investors considered in Section 2, facing transaction costs given by $\Lambda = \lambda\sigma^2$. In this simple context, an equilibrium is defined as a price process and market-clearing asset holdings that are optimal for all agents given the price process. Since the noise traders' positions are optimal by assumption as specified by (52)–(53), the restriction imposed by equilibrium is that the dynamics of the price are such that, for all t ,

$$x_t = -z_t \quad (56)$$

$$dx_t = -dz_t. \quad (57)$$

Using (37), these equilibrium conditions lead to

$$\frac{a}{\lambda}\sigma^{-2}B(a\Phi + \gamma I)^{-1} + \frac{a}{\lambda}e_{L+1} = -\kappa(\mathbf{1} - 2e_{L+1}), \quad (58)$$

where $e_{L+1} = (0, \dots, 0, 1) \in \mathbb{R}^{L+1}$ and $\mathbf{1} = (1, \dots, 1) \in \mathbb{R}^{L+1}$. It consequently follows that, if the investor is to hold $-f_t^L$ at time t for all t , then the factor loadings must be given by

$$\begin{aligned} B &= \sigma^2 \left[-\frac{\lambda}{a}\kappa(\mathbf{1} - 2e_{L+1}) - e_{L+1} \right] (a\Phi + \gamma I) \\ &= \sigma^2 [-\lambda\kappa(\mathbf{1} - 2e_{L+1}) - ae_{L+1}] \left(\Phi + \frac{\gamma}{a}I \right). \end{aligned} \quad (59)$$

For $l \leq L$, we calculate B_l further as

$$\begin{aligned} B_l &= -\sigma^2\kappa(\lambda\psi_l + \lambda\gamma a^{-1} + \lambda\kappa - a) \\ &= -\lambda\sigma^2\kappa(\psi_l + \rho + \kappa), \end{aligned} \quad (60)$$

while

$$B_{L+1} = \sigma^2(\rho\lambda\kappa + \lambda\kappa^2 - \gamma). \quad (61)$$

Using this, it is straightforward to see the following key equilibrium implications:

Proposition 8 *The market is in equilibrium if and only if $x_0 = -z_0$ and the security's alpha is given by*

$$\alpha_t = \sum_{l=1}^L \lambda \sigma^2 \kappa (\psi_l + \rho + \kappa) (-f_t^l) + \sigma^2 (\rho\lambda\kappa + \lambda\kappa^2 - \gamma) z_t. \quad (62)$$

The coefficients $\lambda \sigma^2 \kappa (\psi_k + \rho + \kappa)$ are positive and increase in the mean-reversion parameters ψ_k and κ and in the trading costs $\lambda \sigma^2$. In other words, noise trader selling ($f_t^k < 0$) increases the alpha, and especially so if its mean reversion is faster and if the trading cost is larger.

Naturally, noise-trader selling increases the expected excess return (alpha), while noise-trader buying lowers the alpha, since the arbitrageurs need to be compensated to take the other side of the trade. Interestingly, the effect is larger when trading costs are larger and for noise-trader shocks with faster mean reversion because such shocks are associated with larger trading costs for the arbitrageurs.

4 Application: Dynamic Trading of Commodity Futures

In this section we apply our framework to trading commodity futures using real data.

4.1 Data

We consider 15 different liquid commodity futures, which do not have tight restrictions on the size of daily price moves (limit up/down). In particular, we collect data on Aluminum, Copper, Nickel, Zinc, Lead, and Tin from the London Metal Exchange (LME), on Gas Oil

from the Intercontinental Exchange (ICE), on WTI Crude, RBOB Unleaded Gasoline, and Natural Gas from the New York Mercantile Exchange (NYMEX), on Gold and Silver from the New York Commodities Exchange (COMEX), and on Coffee, Cocoa, and Sugar from the New York Board of Trade (NYBOT). (This excludes futures on various agriculture and livestock that have tight price limits.)

We consider the sample period 01/01/1996 – 01/23/2009, for which we have data on all the commodities.⁵ Every day, we compute for each commodity the price change of the most liquid futures contract (among the available contract maturities), and normalize the series such that each commodity’s price changes have annualized volatility of 10%. We abstract from the cost of rolling from one futures contract to the next. (In the real world, there is a separate roll market with small transaction costs, far smaller than the cost of independently selling the “old” contract and buying the “new” one.)

4.2 Predicting Returns and Other Parameter Estimates

We use the characteristic-based model described in Example 2 in Section 1, where each commodity characteristic is its own past returns at various horizons. Hence, to predict returns, we run a pooled panel regression:

$$r_{t+1}^s = \begin{matrix} 0.000 \\ (-0.02) \end{matrix} + \begin{matrix} 0.011 \\ (1.4) \end{matrix} f_t^{5D,s} + \begin{matrix} 0.037 \\ (4.6) \end{matrix} f_t^{1Y,s} - \begin{matrix} 0.015 \\ (-1.85) \end{matrix} f_t^{5Y,s} + u_{t+1}^s, \quad (63)$$

where the left hand side is the commodity price changes and the right hand side contains the return predictors: f^{5D} is the average past five days’ price changes, divided by the past month’s standard deviation of price changes, f^{1Y} is the past year’s average price change divided by the past year’s standard deviation, and f^{5Y} is the analogous quantity for a five-year window. We report the OLS t-statistics in brackets.

We see that price changes show continuation at short and medium frequencies and re-

⁵Our return predictors use moving averages of price data lagged up to five years, which are available for most commodities except some of the LME base metals. In the early sample when some futures do not have a complete lagged price series, we use the average of the available data.

versal over long horizons.⁶ The goal is to see how an investor could optimally trade on this information, taking transaction costs into account. Of course, these (in-sample) regression results are only available now and a more realistic analysis would consider rolling out-of-sample regressions. However, using the in-sample regression allows us to focus on portfolio optimization. (Indeed, using out-of-sample return forecasts would add noise to the evaluation of the optimization gains of our method.)

The return predictors are chosen so that they have very different mean reversion:

$$\begin{aligned}\Delta f_{t+1}^{5D,s} &= -0.1977 f_t^{5D,s} + \varepsilon_{t+1}^{5D,s} \\ \Delta f_{t+1}^{1Y,s} &= -0.0034 f_t^{1Y,s} + \varepsilon_{t+1}^{1Y,s} \\ \Delta f_{t+1}^{5Y,s} &= -0.0010 f_t^{5Y,s} + \varepsilon_{t+1}^{5Y,s}.\end{aligned}\tag{64}$$

These mean reversion rates correspond to a 3-day half life for the 5-day signal, a 205-day half life for the 1-year signal, and a 701-day half life for the 5-year signal.⁷

We estimate the variance-covariance matrix Σ using daily price changes over the full sample. We set the absolute risk aversion to $\gamma = 10^{-9}$, which we can think of as corresponding to a relative risk aversion of 1 for an agent with 1 billion dollars under management. We set the time discount rate to $\rho = 1 - \exp(-0.02/260)$ corresponding to a 2 percent annualized rate. Finally, we set the transaction cost matrix to $\Lambda = \lambda \Sigma$, where we consider $\lambda = 5 \times 10^{-7}$ as well as a higher λ that is twice as large.

4.3 Dynamic Portfolio Selection with Trading Costs

We consider three different trading strategies: the optimal trading strategy given by Equation (24) (“**optimal**”), the optimal trading strategy in the absence of transaction costs

⁶Asness, Moskowitz, and Pedersen (2008) document 12-month momentum and 5-year reversals of commodities and other securities. These results are robust and hold both for price changes and returns. The 5-day momentum is less robust. For instance, for certain specifications using percent returns, the 5-day coefficient switches sign to reversal. This robustness is not important for our study due to our focus on optimal trading rather than out-of-sample return predictability.

⁷The half life is the time it is expected to take for half the signal to disappear. It is computed as $\log(0.5)/\log(1 - .1977)$ for the 5-day signal.

(“**no-TC**”), and the optimal portfolio in a static (i.e., one-period) model with transaction costs given by Equation (26) (“**static**”). For the static portfolio we use a modified λ such that the coefficient on x_{t-1} is the same as for the optimal portfolio, which is numerically almost the same as choosing λ to maximize the static portfolio’s net Sharpe Ratio.

The performance of each strategy as measured by the Sharpe Ratio (SR) is reported in Table 1. The cumulative excess return of each strategy scaled to 10% annualized volatility is depicted in Figure 3, and Figure 4 shows the cumulative net returns. We see that, naturally, the highest SR before transaction cost is achieved by the no-TC strategy, and the optimal and static portfolios have similar drops in gross SR due to their slower trading. After transaction costs, however, the optimal portfolio is the best, significantly better than the best possible static strategy, and the no-TC strategy incurs enormous trading costs.

It is interesting to consider the driver of the superior performance of the optimal dynamic trading strategy relative to the best possible static strategy. The key to the out-performance is that the dynamic strategy gives less weight to the 5-day signal because of its fast alpha decay. The static strategy simply tries to control the overall trading speed, but this is not sufficient: it either incurs large trading costs due to its “fleeting” target (because of the significant reliance on the 5-day signal), or it trades so slowly it is difficult to capture the alpha. The dynamic strategy overcomes this problem by trading somewhat fast, but trading mainly according to the more persistent signals.

To illustrate the difference in the positions of the different strategies, Figure 5 shows the positions over time of two of the commodity futures, namely Crude and Gold. We see that the optimal portfolio is a much more smooth version of the no-TC strategy, thus reducing trading costs.

4.4 Response to New Information

It is instructive to trace the response to a shock to the return predictors, namely to $\varepsilon_t^{i,s}$ in Equation (64). Figure 6 shows the responses to shocks to each return-predictor factor, namely the 5-day factor, the 1-year factor, and the 5-year factor.

The first panel shows that the no-TC strategy immediately jumps up after a shock to the 5-day factor and slowly mean reverts as the alpha decays. The optimal strategy trades much more slowly and never accumulates nearly as large a position. Interestingly, since the optimal position also trades more slowly out of the position as the alpha decays, the lines cross as the optimal strategy eventually has a larger position than the no-TC strategy.

The second panel shows the response to the 1-year factor. The no-TC jumps up and decays, whereas the optimal position increases more smoothly and catches up as the no-TC starts to decay. The third panel shows the same for the 5Y signal, except that the effects are slower and with opposite sign since 5-year returns predict future reversals.

5 Conclusion

This paper provides a highly tractable framework for studying optimal trading strategies in light of various return predictors, risk and correlation considerations, as well as transaction costs. We derive an explicit closed-form solution for the optimal trading policy and highlight several useful and intuitive results. The optimal portfolio tracks a “target portfolio,” which is analogous to the optimal portfolio in the absence of trading costs in its tradeoff between risk and return, but different since more persistent return predictors are weighted more heavily relative to return predictors with faster alpha decay. The optimal strategy is not to trade all the way to the target portfolio, since this entails too high transaction costs. Instead, it is optimal to take a smoother and more conservative portfolio that moves in the direction of the target portfolio while limiting turnover.

Our framework constitutes a powerful tool to optimally combine various return predictors taking into account their evolution over time, decay rate, and correlation, and trading off their benefits against risks and transaction costs. Such trade-offs are at the heart of the decisions of “arbitrageurs” that help make markets efficient as per the efficient market hypothesis. Arbitrageurs’ ability to do so is limited, however, by transaction costs, and our model provides a tractable and flexible framework for the study of the dynamic implications of this limitation.

We illustrate this feature by embedding our setting in an equilibrium model with several “noise traders” who trade in and out of their positions with varying mean-reversion speeds. In equilibrium, a rational arbitrageur – with trading costs and using the methodology that we derive – needs to take the other side of these noise-trader positions to clear the market. We solve the equilibrium explicitly and show how noise trading leads to return predictability and return reversals. Further, we show that noise-trader demand that mean-reverts more quickly leads to larger return predictability. This is because a fast mean reversion is associated with high transaction costs for the arbitrageurs and, consequently, they must be compensated in the form of larger return predictability.

We implement our optimal trading strategy for commodity futures. Naturally, the optimal trading strategy in the absence of transaction costs has a larger Sharpe ratio gross of fees than our trading policy. However, net of trading costs our strategy performs significantly better since it incurs far lower trading costs while still capturing much of the return predictability and diversification benefits. Further, the optimal dynamic strategy is significantly better than the best static strategy – taking dynamics into account significantly improves performance.

In conclusion, we provide a tractable solution to the dynamic trading strategy in a relevant and general setting that we believe to have many interesting applications.

A Further Analysis and Proofs

Given the linear dynamics of x , the position x_t can be expressed easily as a function of the initial condition and the exogenous path. These results can be used to provide a simple proof of Proposition 6.

Proposition 9 *In discrete time, the optimal dynamic portfolio x_t can be written as a function of the initial position x_0 and the return-predicting factors f_s between time 0 and the current time t :*

$$x_t = M_1^t x_0 + \sum_{s=1}^t M_1^{t-s} M_2 f_s, \quad (\text{A.1})$$

where

$$M_1 = ((1 - \rho)A_{xx} + \gamma\Sigma + \Lambda)^{-1}\Lambda = I - \Lambda^{-1}A_{xx} \quad (\text{A.2})$$

$$M_2 = ((1 - \rho)A_{xx} + \gamma\Sigma + \Lambda)^{-1}(B + (1 - \rho)A_{xf}(I - \Phi)) = \Lambda^{-1}A_{xf}. \quad (\text{A.3})$$

Proposition 10 *In continuous time, the optimal dynamic portfolio x_t can be written in terms of the initial position x_0 and the path of realized factors f_s between 0 and the current time t :*

$$x_t = e^{-\Lambda^{-1}A_{xx}t}x_0 + \int_{s=0}^t e^{-\Lambda^{-1}A_{xx}(t-s)}\Lambda^{-1}A_{xf}f_s ds. \quad (\text{A.4})$$

Proof of Propositions 1, 2, and 3. We calculate the expected future value function as

$$\begin{aligned} E_t[V(x_t, f_{t+1})] &= -\frac{1}{2}x_t^\top A_{xx}x_t + x_t^\top A_{xf}(I - \Phi)f_t + \frac{1}{2}f_t^\top (I - \Phi)^\top A_{ff}(I - \Phi)f_t \\ &\quad + \frac{1}{2}E_t(\varepsilon_{t+1}^\top A_{ff}\varepsilon_{t+1}) + a_0. \end{aligned} \quad (\text{A.5})$$

The agent maximizes the quadratic objective $-\frac{1}{2}x^\top J_t x_t + x_t^\top j_t + d_t$ with

$$\begin{aligned} J_t &= \gamma\Sigma + \Lambda + (1 - \rho)A_{xx} \\ j_t &= (B + (1 - \rho)A_{xf}(I - \Phi))f_t + \Lambda x_{t-1} \\ d_t &= -\frac{1}{2}x_{t-1}^\top \Lambda x_{t-1} + (1 - \rho) \left(\frac{1}{2}f_t^\top (I - \Phi)^\top A_{ff}(I - \Phi)f_t + \frac{1}{2}E_t(\varepsilon_{t+1}^\top A_{ff}\varepsilon_{t+1}) + a_0 \right). \end{aligned} \quad (\text{A.6})$$

The maximum value is attained by

$$x_t = J_t^{-1}j_t, \quad (\text{A.7})$$

which proves (11).

The maximum value is equal to $V(x_{t-1}, f_t) = \frac{1}{2}j_t^\top J_t^{-1}j_t + d_t$ and combining this with (8) we obtain an equation that must hold for all x_{t-1} and f_t . This implies the following restrictions on the coefficient matrices:

$$-A_{xx} = \Lambda(\gamma\Sigma + \Lambda + (1 - \rho)A_{xx})^{-1}\Lambda - \Lambda \quad (\text{A.8})$$

$$A_{xf} = \Lambda(\gamma\Sigma + \Lambda + (1 - \rho)A_{xx})^{-1}(B + (1 - \rho)A_{xf}(I - \Phi)) \quad (\text{A.9})$$

$$\begin{aligned} A_{ff} &= (B + (1 - \rho)A_{xf}(I - \Phi))^\top (\gamma\Sigma + \Lambda + (1 - \rho)A_{xx})^{-1} (B + (1 - \rho)A_{xf}(I - \Phi)) \\ &\quad + (1 - \rho)(I - \Phi)^\top A_{ff}(I - \Phi). \end{aligned} \quad (\text{A.10})$$

We next derive the coefficient matrices A_{xx} , A_{xf} , and A_{ff} by solving these equations. For this, we first rewrite Equation (A.8) by letting $Z = \Lambda^{-\frac{1}{2}}A_{xx}\Lambda^{-\frac{1}{2}}$ and $M = \Lambda^{-\frac{1}{2}}\Sigma\Lambda^{-\frac{1}{2}}$, which yields

$$Z = I - (\gamma M + I + (1 - \rho)Z)^{-1},$$

which is a quadratic with an explicit solution. Since all solutions Z can be written as a limit of polynomials of M , Z and M commute and the quadratic can be sequentially rewritten as

$$\begin{aligned} (1 - \rho)Z^2 + Z(I + \gamma M - (1 - \rho)I) &= \gamma M \\ \left(Z + \frac{1}{2(1 - \rho)}(\rho I + \gamma M) \right)^2 &= \frac{\gamma}{1 - \rho}M + \frac{1}{4(1 - \rho)^2}(\rho I + \gamma M)^2, \end{aligned}$$

resulting in

$$Z = \left(\frac{\gamma}{1 - \rho}M + \frac{1}{4(1 - \rho)^2}(\rho I + \gamma M)^2 \right)^{\frac{1}{2}} - \frac{1}{2(1 - \rho)}(\rho I + \gamma M) \quad (\text{A.11})$$

$$A_{xx} = \Lambda^{\frac{1}{2}} \left[\left(\frac{\gamma}{1 - \rho}M + \frac{1}{4(1 - \rho)^2}(\rho I + \gamma M)^2 \right)^{\frac{1}{2}} - \frac{1}{2(1 - \rho)}(\rho I + \gamma M) \right] \Lambda^{\frac{1}{2}}, \quad (\text{A.12})$$

that is,

$$\begin{aligned} A_{xx} &= \left(\frac{\gamma}{1 - \rho}\Lambda^{\frac{1}{2}}\Sigma\Lambda^{\frac{1}{2}} + \frac{1}{4(1 - \rho)^2}(\rho^2\Lambda^2 + 2\rho\gamma\Lambda^{\frac{1}{2}}\Sigma\Lambda^{\frac{1}{2}} + \gamma^2\Lambda^{\frac{1}{2}}\Sigma\Lambda^{-1}\Sigma\Lambda^{\frac{1}{2}}) \right)^{\frac{1}{2}} \\ &\quad - \frac{1}{2(1 - \rho)}(\rho\Lambda + \gamma\Sigma). \end{aligned} \quad (\text{A.13})$$

Note that the positive definite choice of solution Z is only one that results in a positive definite matrix A_{xx} .

In the case $\Lambda = \lambda\Sigma$ for some scalar $\lambda > 0$, the solution is $A_{xx} = a\Sigma$, where a solves

$$-a = \frac{\lambda^2}{\gamma + \lambda + (1 - \rho)a} - \lambda, \quad (\text{A.14})$$

or

$$(1 - \rho)a^2 + (\gamma + \lambda\rho)a - \lambda\gamma = 0, \quad (\text{A.15})$$

with solution

$$a = \frac{\sqrt{(\gamma + \lambda\rho)^2 + 4\gamma\lambda(1 - \rho)} - (\gamma + \lambda\rho)}{2(1 - \rho)}. \quad (\text{A.16})$$

The other value-function coefficient determining optimal trading is A_{xf} , which solves the linear equation (A.9). To write the solution explicitly, we note first that, from (A.8),

$$\Lambda(\gamma\Sigma + \Lambda + (1 - \rho)A_{xx})^{-1} = I - A_{xx}\Lambda^{-1}. \quad (\text{A.17})$$

Using the general rule that $\text{vec}(XYZ) = (Z^\top \otimes X) \text{vec}(Y)$, we re-write (A.9) in vectorized form:

$$\text{vec}(A_{xf}) = \text{vec}((I - A_{xx}\Lambda^{-1})B) + ((1 - \rho)(I - \Phi)^\top \otimes (I - A_{xx}\Lambda^{-1})) \text{vec}(A_{xf}), \quad (\text{A.18})$$

so that

$$\text{vec}(A_{xf}) = (I - (1 - \rho)(I - \Phi)^\top \otimes (I - A_{xx}\Lambda^{-1}))^{-1} \text{vec}((I - A_{xx}\Lambda^{-1})B). \quad (\text{A.19})$$

In the case $\Lambda = \lambda\Sigma$, the solution is

$$\begin{aligned} A_{xf} &= \lambda B((\gamma + \lambda + (1 - \rho)a)I - \lambda(1 - \rho)(I - \Phi))^{-1} \\ &= \lambda B((\gamma + \lambda\rho + (1 - \rho)a)I + \lambda(1 - \rho)\Phi)^{-1} \end{aligned} \quad (\text{A.20})$$

$$= B \left(\frac{\gamma}{a} + (1 - \rho)\Phi \right)^{-1}. \quad (\text{A.21})$$

Finally, A_{ff} is calculated from the linear equation (A.10), which is of the form

$$A_{ff} = Q + (1 - \rho)(I - \Phi)^\top A_{ff}(I - \Phi) \quad (\text{A.22})$$

with

$$Q = (B + (1 - \rho)A_{xf}(I - \Phi))^\top (\gamma\Sigma + \Lambda + (1 - \rho)A_{xx})^{-1} (B + (1 - \rho)A_{xf}(I - \Phi))$$

a positive-definite matrix.

The solution is easiest to write explicitly for diagonal Φ , in which case

$$A_{ff,ij} = \frac{Q_{ij}}{1 - (1 - \rho)(1 - \Phi_{ii})(1 - \Phi_{jj})}. \quad (\text{A.23})$$

In general,

$$\text{vec}(A_{ff}) = (I - (1 - \rho)(I - \Phi)^\top \otimes (I - \Phi)^\top)^{-1} \text{vec}(Q). \quad (\text{A.24})$$

One way to see that A_{ff} is positive definite is to iterate (A.22) starting with $A_{ff}^0 = 0$, given that $I \geq I - \Phi$.

Having computed the coefficient matrices, finishing the proof is straightforward. Equation (16) follows directly from (11). Equation (13) follows from (16) by using the equations for A_{xf} and a , namely (A.9) and (A.14). ■

Proof of Propositions 4 and 5. Given the conjectured value function, the optimal choice τ equals

$$\tau_t = -\Lambda^{-1}A_{xx}x_t + \Lambda^{-1}A_{xf}f_t,$$

Once this is inserted in the HJB equation, it results in the following equations defining the value-function coefficients (using the symmetry of A_{xx}):

$$-\rho A_{xx} = A_{xx}\Lambda^{-1}A_{xx} - \gamma\Sigma \quad (\text{A.25})$$

$$\rho A_{xf} = -A_{xx}\Lambda^{-1}A_{xf} - A_{xf}\Phi + B \quad (\text{A.26})$$

$$\rho A_{ff} = A_{xf}^\top \Lambda^{-1} A_{xf} - 2A_{ff}\Phi. \quad (\text{A.27})$$

Pre- and post-multiplying (A.25) by $\Lambda^{-\frac{1}{2}}$, we obtain

$$-\rho Z = Z^2 + \frac{\rho^2}{4}I - C, \quad (\text{A.28})$$

that is,

$$\left(Z + \frac{\rho}{2}I\right)^2 = C, \quad (\text{A.29})$$

where

$$Z = \Lambda^{-\frac{1}{2}} A_{xx} \Lambda^{-\frac{1}{2}} \quad (\text{A.30})$$

$$C = \gamma \Lambda^{-\frac{1}{2}} \Sigma \Lambda^{-\frac{1}{2}} + \frac{\rho^2}{4} I. \quad (\text{A.31})$$

This leads to

$$Z = -\frac{\rho}{2} I + C^{\frac{1}{2}} \geq 0, \quad (\text{A.32})$$

implying that

$$A_{xx} = -\frac{\rho}{2} \Lambda + \Lambda^{\frac{1}{2}} \left(\gamma \Lambda^{-\frac{1}{2}} \Sigma \Lambda^{-\frac{1}{2}} + \frac{\rho^2}{4} \right)^{\frac{1}{2}} \Lambda^{\frac{1}{2}}. \quad (\text{A.33})$$

The solution for A_{xf} follows from Equation (A.26), using the general rule that $\text{vec}(XYZ) = (Z^\top \otimes X) \text{vec}(Y)$:

$$\text{vec}(A_{xf}) = (\rho I + \Phi^\top \otimes I_K + I_S \otimes (A_{xx} \Lambda^{-1}))^{-1} \text{vec}(B)$$

If $\Lambda = \lambda \Sigma$, then $A_{xx} = a \Sigma$ with

$$-\rho a = a^2 \frac{1}{\lambda} - \gamma \quad (\text{A.34})$$

with solution

$$a = -\frac{\rho}{2} \lambda + \sqrt{\gamma \lambda + \frac{\rho^2}{4} \lambda^2}. \quad (\text{A.35})$$

In this case, (A.26) yields

$$\begin{aligned} A_{xf} &= B \left(\rho I + \frac{a}{\lambda} I + \Phi \right)^{-1} \\ &= B \left(\frac{\gamma}{a} I + \Phi \right)^{-1}, \end{aligned}$$

where the last equality uses (A.34).

Then we have

$$\tau_t = \frac{a}{\lambda} [\Sigma^{-1} B (a\Phi + \gamma I)^{-1} f_t - x_t] \quad (\text{A.36})$$

It is clear from (A.35) that $\frac{a}{\lambda}$ decreases in λ and increases in γ . ■

Proof of Proposition 6. We prove this proposition in two main steps. We use the notation from Proposition 9.

(i) It holds that

$$\begin{aligned} M_1(\Delta_t) &= I - (\Lambda^{-1} A_{xx} + O(\Delta_t)) \Delta_t \\ M_2(\Delta_t) &= (\Lambda^{-1} A_{xf} + O(\Delta_t)) \Delta_t \end{aligned}$$

as $\Delta_t \rightarrow 0$.

(ii) $M_1(\Delta_t)^{\frac{t}{\Delta_t}} \rightarrow e^{-\Lambda^{-1} A_{xx} t}$ uniformly on $[0, T]$ for any $T > 0$. For any continuous path u , $\hat{x}_t \rightarrow x_t$ uniformly on $[0, T]$ for any $T > 0$. It then follows immediately from (9) and (34) that $\frac{\Delta x_t}{\Delta t} \rightarrow \tau_t$. ■

Proof of Proposition 7. The HJB equation is

$$\begin{aligned} \rho V &= \max_{\tau} \left\{ x^\top (Bf - RD + C^\top \tau) - \frac{\gamma}{2} x^\top \Sigma x + \frac{\partial V}{\partial f} (-\Phi f) + \frac{\partial V}{\partial x} \tau + \frac{\partial V}{\partial D} (-RD + C^\top \tau) \right\} \\ &= \max_{\tau} \left\{ x^\top Bf - \frac{\gamma}{2} x^\top \Sigma x - \frac{1}{2} \tau^\top \Lambda \tau + \tau^\top (-Q_{1x}x + Q_{1D}D + Q_{1f}f) + \frac{\partial V}{\partial f} (-\Phi f) - \right. \\ &\quad \left. (D^\top A_{DD} + f^\top A_{fD} + x^\top A_{xD}) RD \right\}, \end{aligned} \quad (\text{A.37})$$

with

$$\begin{aligned} -Q_{1x} &= -A_{xx} + CA_{Dx} + C \\ Q_{1D} &= A_{xD} + CA_{DD} \\ Q_{1f} &= A_{xf} + CA_{Df}. \end{aligned} \quad (\text{A.38})$$

It follows immediately that

$$\begin{aligned}\tau &= \Lambda^{-1} Q_{1x} [target - x] \\ &= \Lambda^{-1} (A_{xx} - CA_{Dx} - C) [target - x],\end{aligned}\tag{A.39}$$

with

$$\begin{aligned}target &= [Q_{1x}^{-1} Q_{1D}] D + [Q_{1x}^{-1} Q_{1f}] f \\ &= (A_{xx} - CA_{Dx} - C)^{-1} ((A_{xD} + CA_{DD}) D + (A_{xf} + CA_{Df}) f).\end{aligned}\tag{A.40}$$

The coefficient matrices solve the system:

$$\begin{aligned}-\rho A_{xx} &= -\gamma \Sigma + Q_{1x}^\top \Lambda^{-1} Q_{1x} \\ &= -\gamma \Sigma + (A_{xx} - CA_{Dx} - C)^\top \Lambda^{-1} (A_{xx} - CA_{Dx} - C) \\ \rho A_{xD} &= -Q_{1x}^\top \Lambda^{-1} Q_{1D} - A_{xD} R - R \\ &= -(A_{xx} - CA_{Dx} - C)^\top \Lambda^{-1} (A_{xD} + CA_{DD}) - A_{xD} R - R \\ \rho A_{DD} &= Q_{1D}^\top \Lambda^{-1} Q_{1D} - 2A_{DD} R \\ &= (A_{xD} + CA_{DD})^\top \Lambda^{-1} (A_{xD} + CA_{DD}) - 2A_{DD} R \\ \rho A_{xf} &= B - Q_{1x}^\top \Lambda^{-1} Q_{1f} - A_{xf} \Phi \\ &= B - (A_{xx} - CA_{Dx} - C)^\top \Lambda^{-1} (A_{xf} + CA_{Df}) - A_{xf} \Phi \\ \rho A_{Df} &= -A_{Df} \Phi - R^\top A_{Df} + (A_{xD} + CA_{DD})^\top \Lambda^{-1} (A_{xf} + CA_{Df}) \\ \rho A_{ff} &= Q_{1f}^\top \Lambda^{-1} Q_{1f} - 2A_{ff} \Phi \\ &= (A_{xf} + CA_{Df})^\top \Lambda^{-1} (A_{xf} + CA_{Df}) - 2A_{ff} \Phi.\end{aligned}\tag{A.41}$$

We note that the first three equations above have to be solved simultaneously for A_{xx} , A_{xD} , and A_{DD} ; there is no closed-form solution. The complication is due to the fact that current trading affects the persistent price component D (that is, $C \neq 0$). ■

Proof of Proposition 8. Suppose that $\alpha_t = Bf_t$ with B given by (59) and apply Proposition 5 to conclude that, if $x_t = -f_t^{K+1}$, then $dx_t = -df_t^{K+1}$. The comparative-static results are immediate. ■

References

- Acharya, V., and L. H. Pedersen, 2005, “Asset Pricing with Liquidity Risk,” *Journal of Financial Economics*, 77, 375–410.
- Almgren, R., and N. Chriss, 2000, “Optimal Execution of Portfolio Transactions,” *Journal of Risk*, 3, 5–39.
- Amihud, Y., and H. Mendelson, 1986, “Asset Pricing and the Bid-Ask Spread,” *Journal of Financial Economics*, 17, 223–249.
- Asness, C., T. Moskowitz, and L. H. Pedersen, 2008, “Value and Momentum Everywhere,” *Working paper AQR, NYU, and Chicago*.
- Balduzzi, P., and A. Lynch, 1999, “Transaction costs and predictability: some utility cost calculations,” *Journal of Financial Economics*, 52(1), 47–78.
- Bertsimas, D., and A. W. Lo, 1998, “Optimal Control of Execution Costs,” *Journal of Financial Markets*, 1, 1–50.
- Brunnermeier, M. K., and L. H. Pedersen, 2005, “Predatory Trading,” *Journal of Finance*, 60(4), 1825–1863.
- Campbell, J. Y., and L. M. Viceira, 2002, *Strategic Asset Allocation Portfolio Choice for Long-Term Investors*. Oxford University Press.
- Carlin, B. I., M. Lobo, and S. Viswanathan, 2008, “Episodic Liquidity Crises: Cooperative and Predatory Trading,” *Journal of Finance*, 62, 2235–2274.
- Constantinides, G. M., 1986, “Capital Market Equilibrium with Transaction Costs,” *Journal of Political Economy*, 94, 842–862.
- Davis, M., and A. Norman, 1990, “Portfolio selection with transaction costs,” *Mathematics of Operations Research*, 15, 676–713.

- Engle, R., and R. Ferstenberg, 2007, “Execution Risk,” *Journal of Portfolio Management*, 33, 34–45.
- Gârleanu, N., 2009, “Portfolio Choice and Pricing in Imperfect Markets,” *Journal of Economic Theory*, 144(2), 532–564.
- Gârleanu, N., L. H. Pedersen, and A. Potesman, 2008, “Demand-Based Option Pricing,” *Review of Financial Studies* (forthcoming).
- Grinold, R., 2006, “A Dynamic Model of Portfolio Management,” *Journal of Investment Management*, 4, 5–22.
- Heaton, J., and D. Lucas, 1996, “Evaluating the effects of incomplete markets on risk sharing and asset pricing,” *Journal of Political Economy*, 104, 443–487.
- Jang, B., H. Koo, H. Liu, and M. Loewenstein, 2007, “Liquidity Premia and Transaction Costs,” *Journal of Finance*, 62, 2329–2366.
- Lagos, R., and G. Rocheteau, 2006, “Search in Asset Markets,” Working paper, NYU.
- Liu, H., 2004, “Optimal consumption and investment with transaction costs and multiple assets,” *Journal of Finance*, 59, 289–338.
- Ljungqvist, L., and T. Sargent, 2004, *Recursive Macroeconomic Theory*, 2nd edition. MIT press.
- Lo, A., H. Mamaysky, and J. Wang, 2004, “Asset prices and trading volume under fixed transaction costs,” *Journal of Political Economy*, 112, 1054–1090.
- Longstaff, F. A., 2001, “Optimal Portfolio Choice and the Valuation of Illiquid Securities,” *The Review of Financial Studies*, 14, 407–431.
- Lynch, A., and P. Balduzzi, 2000, “Predictability and transaction costs: the impact on rebalancing rules and behavior,” *Journal of Finance*, pp. 2285–2309.

- Lynch, A., and S. Tan, 2008, “Explaining the Magnitude of Liquidity Premia: The Roles of Return Predictability, Wealth Shocks and State-dependent Transaction Costs.,” *Working paper, NYU*.
- Obizhaeva, A., and J. Wang, 2006, “Optimal trading strategy and supply/demand dynamics,” *Working paper, MIT*.
- Oehmke, M., 2009, “Dynamic Arbitrage in Illiquid Markets,” *Working paper, Columbia*.
- Perold, A., 1988, “The Implementation Shortfall: Paper Versus Reality,” *Journal of Portfolio Management*, 14, 4–9.
- Vayanos, D., 1998, “Transaction Costs and Asset Prices: A Dynamic Equilibrium Model,” *Review of Financial Studies*, 11, 1–58.
- Vayanos, D., and J.-L. Vila, 1999, “Equilibrium Interest Rate and Liquidity Premium with Transaction Costs,” *Economic Theory*, 13, 509–539.

Table 1: **Performance of Trading Strategies Before and After Transaction Costs.** This table shows the annualized Sharpe ratio gross and net of trading costs for the optimal trading strategy in the absence of trading costs (“no TC”), our optimal dynamic strategy (“optimal”), and a strategy that optimizes a static one-period problem with trading costs (“static”). Panel A illustrates this for a low transaction cost parameter, while Panel B has a high one.

Panel A: Low Transaction Costs

	no TC	optimal	static
gross SR	0.79	0.63	0.64
net SR	−18	0.54	0.44

Panel B: High Transaction Costs

	no TC	optimal	static
gross SR	0.79	0.57	0.58
net SR	−22	0.45	0.33

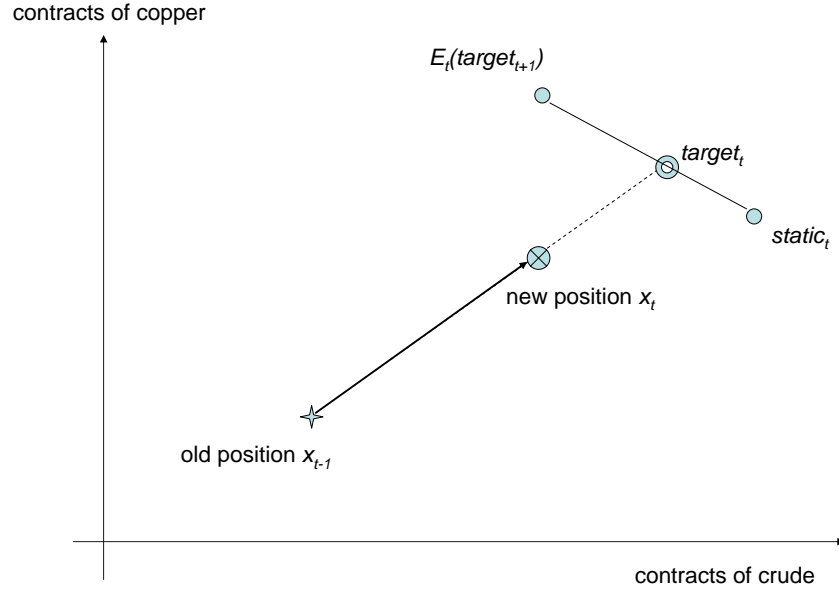


Figure 1: **Optimal Trading Strategy: Triangulating between Current, Static, and Future Portfolios.** This figure shows how the optimal trade moves from the existing position x_{t-1} towards the target, trading only part of the way to the target to limit transactions costs. The target is an average of the static Markowitz portfolio and the expected future target, which depends on the optimal portfolio in the future including the return predictors' expected alpha decay.

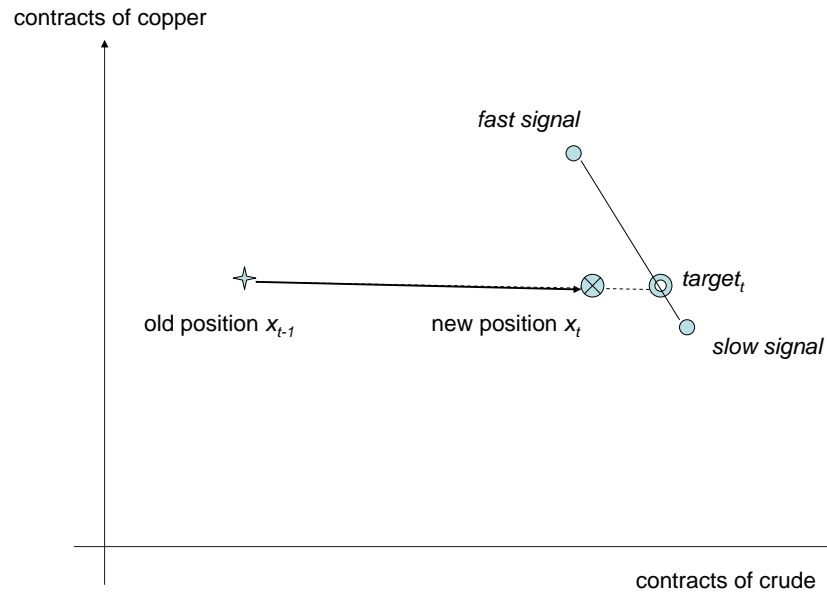


Figure 2: **Optimal trading strategy: Down-Weight Fast-Decay Factors.** This figure shows how the optimal trade moves from the existing position x_{t-1} towards the target, which put more weight on slow persistent factors relative to factors with fast alpha decay.




Figure 3: **Cumulative Excess Returns Gross of Transactions Costs.** This figure shows the cumulative excess returns before transactions costs for the the optimal trading strategy in the absence of trading costs (“no TC”), our optimal dynamic strategy (“optimal”), and a strategy that optimizes a static one-period problem with trading costs (“static”).




Figure 4: **Cumulative Excess Returns Net of Transactions Costs.** This figure shows the cumulative excess returns after transactions costs for the the optimal trading strategy in the absence of trading costs (“no TC”), our optimal dynamic strategy (“optimal”), and a strategy that optimizes a static one-period problem with trading costs (“static”).




Figure 5: **Positions in Crude and Gold Futures.** This figure shows the positions in crude and gold for the the optimal trading strategy in the absence of trading costs (“no TC”) and our optimal dynamic strategy (“optimal”) using high and low transactions costs.

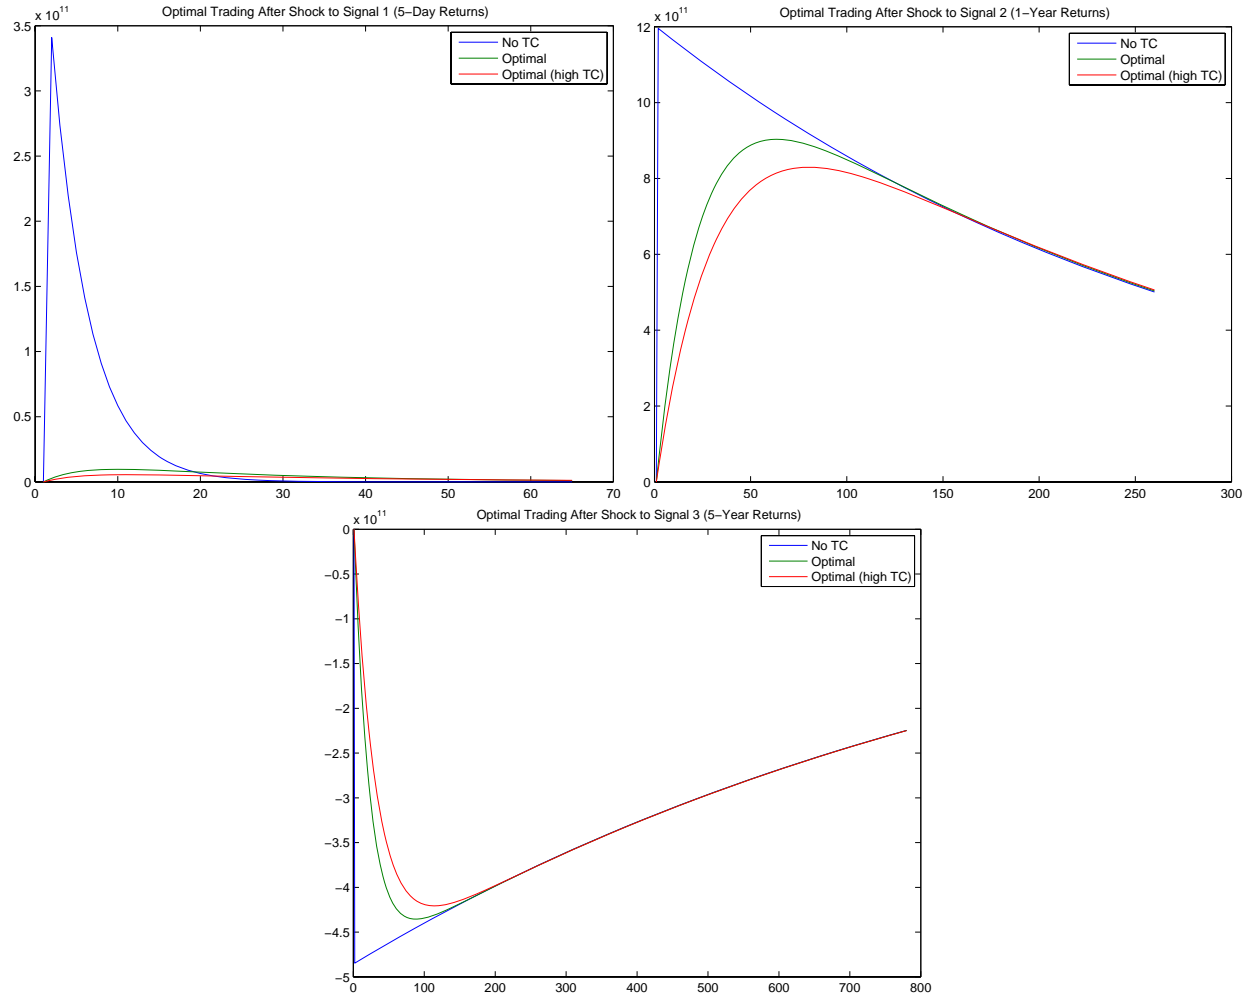


Figure 6: **Optimal Trading in Response to Shock to Return Predicting Signals.**

This figure shows the response in the optimal position following a shock to a return predictor as a function of the number of days since the shock. The top left panel does this for a shock to the fast 5-day return predictor, the top right panel considers a shock to the 12-month return predictor, and the bottom panel to the 5-year predictor. In each case, we consider the response of the optimal trading strategy in the absence of trading costs (“no TC”) and our optimal dynamic strategy (“optimal”) using high and low transactions costs.