

Dynamic Trading with Predictable Returns and Transaction Costs*

Nicolae Gârleanu and Lasse Heje Pedersen[†]

August, 2012

Abstract

We derive a closed-form optimal dynamic portfolio policy when trading is costly and security returns are predictable by signals with different mean-reversion speeds. The optimal strategy is characterized by two principles: 1) aim in front of the target and 2) trade partially towards the current aim. Specifically, the optimal updated portfolio is a linear combination of the existing portfolio and an “aim portfolio,” which is a weighted average of the current Markowitz portfolio (the moving target) and the expected Markowitz portfolios on all future dates (where the target is moving). Intuitively, predictors with slower mean reversion (alpha decay) get more weight in the aim portfolio. We implement the optimal strategy for commodity futures and find superior net returns relative to more naive benchmarks.

*We are grateful for helpful comments from Kerry Back, Darrell Duffie, Pierre Collin-Dufresne, Andrea Frazzini, Esben Hedegaard, Hong Liu (discussant), 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, New York University, the University of Copenhagen, Rice University, University of Michigan Ross School, Yale University SOM, the Bank of Canada, and the Journal of Investment Management Conference.

[†]Gârleanu is at Haas School of Business, University of California, Berkeley, NBER, and CEPR; e-mail: garleanu@haas.berkeley.edu. Pedersen (corresponding author) is at New York University, Copenhagen Business School, FRIC, AQR Capital Management, NBER, and CEPR, 44 West Fourth Street, NY 10012-1126; e-mail: lpederse@stern.nyu.edu, <http://www.stern.nyu.edu/~lpederse/>.

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 transaction costs. Hence, any active investor must constantly weigh the expected benefit of trading against its costs and risks. 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 time 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.

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 transaction costs. We present a closed-form solution for the optimal dynamic portfolio strategy, giving rise to two principles: 1) aim in front of the target and 2) trade partially towards the current aim.

To see the intuition for these portfolio principles, note that the investor would like to keep his portfolio close to the optimal portfolio in the absence of transaction costs, which we call the “Markowitz portfolio” (or the tangency portfolio). The Markowitz portfolio is a *moving target* since the return-predicting factors change over time. Due to transaction costs, it is obviously not optimal to trade all the way to the target all the time. Hence, transaction costs make it optimal to slow down the speed of trading and, interestingly, to modify the aim, not trading directly towards the *current* Markowitz portfolio. Indeed, the optimal strategy is to trade towards an “aim portfolio,” which is a weighted average of the current Markowitz portfolio (the moving target) and the expected Markowitz portfolios on all future dates (where the target is moving).

While new to finance, these portfolio principles have close analogues in other fields such as the guidance of missiles towards moving targets, shooting, and sports. For example, related

dynamic programming principles are used to guide missiles to an enemy airplane in so-called “lead homing” systems. Similarly, hunters are reminded to “lead the duck” when aiming their weapon.¹ The most famous example from the sports world is perhaps the following quote from the “great one”:

“A great hockey player skates to where the puck is going to be, not where it is.”

— Wayne Gretzky

Another way to state our portfolio principle is that the best new portfolio is a combination of 1) the current portfolio (to reduce turnover), 2) the Markowitz portfolio (to partly get the best current risk-return trade-off), and 3) the expected optimal portfolio in the future (a dynamic effect).

Figure 1 illustrates this natural trading rule. The solid line illustrates the expected path of the Markowitz portfolio, starting with large positions in both security 1 and security 2, and gradually converging towards its long-term mean (e.g., the market portfolio). The aim portfolio is a weighted-average of the current and future Markowitz portfolios so it lies in the “convex hull” of the solid line or, equivalently, between the current Markowitz portfolio and the expected aim portfolio next period. The optimal new position is achieved by trading partially towards this aim portfolio.

In this example, curve in the solid lines means that the Markowitz position in security 1 decays more slowly as the predictor that currently “likes” security 1 is more persistent. Therefore, the aim portfolio has a larger position in security 1, and, consequently, the optimal trade buys more shares in security 1 than it would otherwise. We show that it is in fact a more general principle that predictors with slower mean reversion (alpha decay) get more weight in the aim portfolio. An investor facing transaction costs should trade more aggressively on persistent signals than on fast mean-reverting signals: the benefits from the former accrue over longer periods, and are therefore larger.

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

¹We thank Kerry Back for this analogy.

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.

Our specification of transaction costs is sufficiently rich to allow for both purely transitory and persistent costs. Persistent transaction costs means that trading leads to a market impact and this effect on prices persists for a while. Indeed, since we focus on market impact costs, it may be more realistic to consider such persistent effects, especially over short time periods. We show that, with persistent transaction costs, the optimal strategy remains to trade partially towards an aim portfolio and to aim in front of the target, though the precise trading strategy is different and more involved. Furthermore, we offer micro-foundations for each type of transaction costs, rooted in the inventory costs of liquidity providers who take the other side of the trades, and illustrate the continuous-time limits (in the appendix).

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 aim 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 reduced more slowly.

The paper is organized as follows. Section 1 describes how our paper contributes to the portfolio-selection literature that started with Markowitz (1952). We provide a closed-form solution for a model with multiple correlated securities and multiple return predictors with different mean-reversion speeds. The closed-form solution illustrates several intuitive portfolio principles that are difficult to see in the models following Constantinides (1986), where the solution requires complex numerical techniques even with a *single* security and *no* return predictors (i.i.d. returns). Indeed, we uncover the role of alpha decay and the intuitive aim-in-front-of-the-target and trade-towards-the-aim principles, and our empirical analysis suggests that these principles are useful.

Section 2 lays out the model with temporary transaction costs and lays out the solution method. Section 3 shows the optimality of aiming in front of the target, and trading partially towards the aim. Section 4 solves the model with persistent transaction costs. Section 5 provides a number of theoretical applications while Section 6 applies our framework empirically to trading commodity futures. Section 7 concludes. Appendix A contains continuous-time versions of the temporary and persistent transaction cost models, Appendix B develops micro foundations for these costs, Appendix C shows the connection between discrete and continuous time, Appendix D derives equilibrium implications of the model, showing that transaction costs can help explain high short-term alphas and return reversals. All proofs are in Appendix E.

1 Related Literature

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). Alpha decay plays no role in this literature, and nor does it play a role in the literature on optimal portfolio selection with trading costs but without return predictability following Constantinides (1986).

This latter literature models transaction costs as proportional bid-ask spreads and relies

on numerical solutions. Constantinides (1986) considers a single risky asset in a partial equilibrium and studies trading-cost implications for the equity premium.² 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.

Our trade-partially-towards-the-aim strategy is qualitatively different from the optimal strategy with proportional or fixed transaction costs, which exhibits periods of no trading. Our strategy mimics a trader who is continuously “floating” limit orders close to the mid-quote — a strategy that is used in practice. The trading speed (the limit orders’ “fill rate” in our analogy) depends on how large transaction costs the trader is willing to accept (i.e., on where the limit orders are placed).

In a third (and most related) strand of literature, using calibrated numerical solutions, trading costs are combined with incomplete 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.

A fourth strand of literature derives the optimal trade execution, treating the asset and quantity 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)).

²Davis 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).

Finally, 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.

2 Model and Solution

We consider an economy with S securities traded at each time $t \in \{0, 1, 2, \dots\}$. The securities' price changes between times t and $t + 1$ in excess of the risk-free return, $p_{t+1} - (1 + r^f)p_t$, are collected in an $S \times 1$ vector r_{t+1} given by

$$r_{t+1} = Bf_t + u_{t+1}. \quad (1)$$

Here, f_t is a $K \times 1$ vector of factors that predict returns,³ B is an $S \times K$ matrix of factor loadings, and u_{t+1} is the unpredictable zero-mean noise term with variance $\text{var}_t(u_{t+1}) = \Sigma$.

The return-predicting factor f_t is known to the investor already at time t and it evolves according to

$$\Delta f_{t+1} = -\Phi f_t + \varepsilon_{t+1}, \quad (2)$$

where $\Delta f_{t+1} = f_{t+1} - f_t$ is the change in the factors, Φ is a $K \times K$ matrix of mean-reversion coefficients for the factors, and ε_{t+1} is the shock affecting the predictors with variance $\text{var}_t(\varepsilon_{t+1}) = \Omega$. We impose on Φ standard conditions sufficient to ensure that f is stationary.

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 the factors f that could be past returns over various horizons, valuation ratios, and other return-predicting variables, and

³The unconditional mean excess returns are also captured in the factors f . E.g., one can let the first factor be a constant, $f_t^1 = 1$ for all t , such that the first column of B contains the vector of mean returns. (In this case, the shocks to the first factor are zero, $\varepsilon_t^1 = 0$.)

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

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

$$TC(\Delta x_t) = \frac{1}{2} \Delta x_t^\top \Lambda \Delta x_t, \quad (3)$$

where Λ is a symmetric positive-definite matrix measuring the level of trading costs.⁴ 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, which can also be justified by inventory considerations (e.g., Grossman and Miller (1988) or Greenwood (2005) for the multi-asset case). While this transaction-cost specification is chosen partly for tractability, the empirical literature generally finds transaction costs to be convex (e.g., Engle, Ferstenberg, and Russell (2008), Lillo, Farmer, and Mantegna (2003)), with some researchers actually estimating quadratic trading costs (e.g., Breen, Hodrick, and Korajczyk (2002)).

Most of our results hold with this general transaction cost function, but some of the resulting expressions are simpler in the following special case.

Assumption A. Transaction costs are proportional to the amount of risk, $\Lambda = \lambda \Sigma$.

This assumption means that the transaction cost matrix Λ is some scalar $\lambda > 0$ times the variance-covariance matrix of returns, Σ , as is natural and, in fact, implied by the model of Gârleanu, Pedersen, and Poteshman (2008) as well as the micro-foundation that we provide

⁴The 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., $\Lambda = (\bar{\Lambda} + \bar{\Lambda}^\top)/2$, generates the same trading costs as $\bar{\Lambda}$.

in Appendix B. To understand this, suppose that a dealer takes the other side of the trade Δx_t , holds this position for a period of time and “lays it off” at the end of the period. Then the dealer’s risk is $\Delta x_t^\top \Sigma \Delta x_t$ and the trading cost is the dealer’s compensation for risk, depending on the dealer’s risk aversion reflected by λ .

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 excess returns, penalized for risks and trading costs:

$$\max_{x_0, x_1, \dots} E_0 \left[\sum_t (1 - \rho)^{t+1} \left(x_t^\top r_{t+1} - \frac{\gamma}{2} x_t^\top \Sigma x_t \right) - \frac{(1 - \rho)^t}{2} \Delta x_t^\top \Lambda \Delta x_t \right], \quad (4)$$

where $\rho \in (0, 1)$ is a discount factor, and γ is the risk-aversion coefficient.⁵

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:

$$V(x_{t-1}, f_t) = \max_{x_t} \left\{ -\frac{1}{2} \Delta x_t^\top \Lambda \Delta x_t + (1 - \rho) \left(x_t^\top E_t[r_{t+1}] - \frac{\gamma}{2} x_t^\top \Sigma x_t + E_t[V(x_t, f_{t+1})] \right) \right\}. \quad (5)$$

The model in its general form can be solved explicitly:

Proposition 1 *The model has a unique solution and the value function is given by*

$$V(x_t, f_{t+1}) = -\frac{1}{2} x_t^\top A_{xx} x_t + x_t^\top A_{xf} f_{t+1} + \frac{1}{2} f_{t+1}^\top A_{ff} f_{t+1} + A_0. \quad (6)$$

*The coefficient matrices A_{xx} , A_{xf} , A_{ff} are stated explicitly in (E.9), (E.12), and (E.15), and A_{xx} is positive definite.*⁶

⁵Said differently, the investor has mean-variance preferences over the change in his wealth W_t each time period, net of the risk-free return: $\Delta W_{t+1} - r^f W_t = x_t^\top r_{t+1} - TC_{t+1}$.

⁶Note that A_{xx} and A_{ff} can always be chosen to be symmetric.

3 Results: Aim in Front of the Target

We next explore the properties of the optimal portfolio policy, which turns out to be intuitive and relatively simple. The core idea is that the investor aims to achieve a certain position, but trades only partially towards this “aim portfolio” due to transaction costs. The aim portfolio itself combines the current optimal portfolio in the absence of transaction costs and the expected future such portfolios. The formal results are stated in the following propositions.

Proposition 2 (Trade Partially Towards the Aim) (i) *The optimal portfolio is*

$$x_t = x_{t-1} + \Lambda^{-1} A_{xx} (aim_t - x_{t-1}), \quad (7)$$

which implies trading at a proportional rate given by the the matrix $\Lambda^{-1} A_{xx}$ towards the “aim portfolio,”

$$aim_t = A_{xx}^{-1} A_{xft}. \quad (8)$$

(ii) *Under Assumption A, the optimal trading rate is the scalar $a/\lambda < 1$, where*

$$a = \frac{-(\gamma(1-\rho) + \lambda\rho) + \sqrt{(\gamma(1-\rho) + \lambda\rho)^2 + 4\gamma\lambda(1-\rho)^2}}{2(1-\rho)}. \quad (9)$$

The trading rate is decreasing in transaction costs λ and increasing in risk aversion γ .

This proposition provides a simple and appealing trading rule. The optimal portfolio is a weighted average of the existing portfolio x_{t-1} and the aim portfolio:

$$x_t = \left(1 - \frac{a}{\lambda}\right) x_{t-1} + \frac{a}{\lambda} aim_t. \quad (10)$$

The weight of the aim portfolio — which we also call the “trading rate” — determines how far the investor should rebalance towards the aim. Interestingly, the optimal portfolio always rebalances by fixed fraction towards the aim (i.e., the trading rate is independent of the current portfolio x_{t-1} or past portfolios). The optimal trading rate is naturally greater

if transaction costs are smaller. Said differently, high transaction costs imply that one must trade more slowly. Also, the trading rate is greater if risk aversion is larger, since a larger risk aversion makes the risk of deviating from the aim more painful (and a larger absolute risk aversion can also be viewed as an investor with less capital, for whom transaction costs play a smaller role).

Next, we want to understand the aim portfolio. The aim portfolio in our dynamic setting turns out to be closely related to the optimal portfolio in a static model without transaction costs ($\Lambda = 0$), which we call the *Markowitz* portfolio. In agreement with the classical findings of Markowitz (1952),

$$\text{Markowitz}_t = (\gamma \Sigma)^{-1} B f_t. \quad (11)$$

Proposition 3 (Aim in Front of the Target) *(i) The aim portfolio is the weighted average of the current Markowitz portfolio and the expected future aim portfolio. Under Assumption A, this can be written as follows, letting $z = \gamma/(\gamma + a)$:*

$$\text{aim}_t = z \text{Markowitz}_t + (1 - z) E_t(\text{aim}_{t+1}). \quad (12)$$

(ii) The aim portfolio can also be expressed as the weighted average of the current Markowitz portfolio and the expected Markowitz portfolios at all future times. Under Assumption A,

$$\text{aim}_t = \sum_{\tau=t}^{\infty} z(1 - z)^{\tau-t} E_t(\text{Markowitz}_\tau). \quad (13)$$

The weight z of the current Markowitz portfolio decreases with transaction costs λ and increases in risk aversion γ .

We see that the aim portfolio is a weighted average of current and future expected Markowitz portfolios. While without transaction costs, the investor would like to hold the Markowitz portfolio to earn the highest possible risk-adjusted return, with transaction costs the investor needs to economize on trading and thus trade partially towards the aim and, as a result, he needs to adjust his aim in front of the target. Proposition 3 shows that the

optimal aim portfolio is an exponential average of current and future Markowitz portfolio, where the weight on the current (and near term) Markowitz portfolio is larger if transaction costs are smaller.

A graphical illustration of optimal trading rule. The optimal trading policy is illustrated in detail in Figure 2. Panel A of Figure 2 shows how the optimal *first* trade is derived, Panel B how the expected *second* trade, and Panel C the *entire path* of expected future trades. Let's first understand Panel A. The solid curve is the expected path of future Markowitz portfolios. Since expected returns mean revert, the expected Markowitz portfolio converges to its long-term mean, illustrated at the origin of the figure. In this example, asset 2 loads on a factor that decays the fastest, so the future Markowitz positions are expected to have relatively larger positions in asset 1. As a result of the general alpha decay and transaction costs, the current aim portfolio has smaller positions than the Markowitz portfolio and, as a result of the differential alpha decay, the aim portfolio loads more on in asset 1. The optimal new position is found by moving partially towards the aim portfolio. Panel B shows that the expected next trade is towards the new aim, using the same logic as before. Panel C traces out the entire paths of expected future positions. The optimal strategy is to chase a moving target, adjusting the aim for alpha decay and trading patiently by always edging partially towards the aim.⁷

To further understand the aim portfolio, we can characterize the effect of the future expected Markowitz portfolios in terms of the different trading signals (or factors), f_t , and their mean reversion speeds. Indeed, a more persistent factor has a larger effect on future Markowitz portfolios than a factors that quickly mean reverts. Indeed, the central relevance of signal persistence in the presence of transaction costs is one of the distinguishing features of our analysis.

Proposition 4 (Weight Signals Based on Alpha Decay) *(i) Under Assumption A, the aim portfolio is the Markowitz portfolio built as if the signals f were scaled down based on*

⁷The parameters underlying these figures are $f_0 = (1, 1)^\top$, $B = I_{2 \times 2}$, $\phi_1 = 0.1$, $\phi_2 = 0.4$, $\Sigma = I_{2 \times 2}$, $\gamma = 0.5$, $\rho = 0.05$, $\Lambda = 2\Sigma$.

their mean reversion Φ :

$$aim_t = (\gamma\Sigma)^{-1} B \left(I + \frac{a}{\gamma} \Phi \right)^{-1} f_t. \quad (14)$$

(ii) If the matrix Φ is diagonal, $\Phi = \text{diag}(\phi^1, \dots, \phi^K)$, then the aim portfolio simplifies as the Markowitz portfolio with each factor f_t^k scaled down based on its own alpha decay ϕ^k :

$$aim_t = (\gamma\Sigma)^{-1} B \left(\frac{f_t^1}{1 + \phi^1 a/\gamma}, \dots, \frac{f_t^K}{1 + \phi^K a/\gamma} \right)^\top. \quad (15)$$

This proposition shows explicitly the close link between the optimal dynamic aim portfolio in light of transaction costs and the classic Markowitz portfolio. The aim portfolio resembles the Markowitz portfolio, but the factors are scaled down based on transaction costs (captured by a), risk aversion (γ), and, importantly, the mean-reversion speed of the factors (Φ).

The aim portfolio is particularly simple under the rather standard assumption that the dynamics of each factor f^k depend only 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 (2) simplifies to scalars:

$$\Delta f_{t+1}^k = -\phi^k f_t^k + \varepsilon_{t+1}^k. \quad (16)$$

The resulting aim portfolio is very similar to the Markowitz portfolio, $(\gamma\Sigma)^{-1} B f_t$. Hence, transaction costs imply first that one optimally only trades part of the way towards the aim, and, second, that the aim 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. This feature is also seen in Figure 2. The convex J-shape of the path of expected future Markowitz portfolios indicates that the factors that predict a high return for asset 2 decay faster than those that predict asset 1. To make this point in a different way, if the expected returns of the two assets decayed equally fast, then the Markowitz portfolio would be expected to move linearly towards its long-term mean. Since the aim portfolio downweights the faster decaying factors, the investor trades less towards asset 2. To see this graphically, note that the aim lies below

the line joining the Markowitz portfolio with the origin, thus downweighting asset 2 relative to asset 1. Naturally, giving more weight to the more persistent factors means that the investor trades towards a portfolio that not only has a high expected return now, but also is expected to have a high expected return for a longer time in the future.

We end this section by considering what portfolio an investor ends up owning if he always follows our optimal strategy:

Proposition 5 (Position Homing In) *Suppose that the agent has followed the optimal trading strategy from time $-\infty$ until time t . Then the current portfolio is an exponentially weighted average of past aim portfolios. Under Assumption A,*

$$x_t = \sum_{\tau=-\infty}^t \frac{a}{\lambda} \left(1 - \frac{a}{\lambda}\right)^{t-\tau} aim_{\tau} \quad (17)$$

We see that the optimal portfolio is an exponentially weighted average of current and past aim portfolios. Clearly, the history of the past expected returns affects the current position since the investor trades patiently to economize on transaction costs. The proposition further implies that the investor can compute the exponentially weighted average of past aim portfolios and always trade *all the way* to this portfolio (assuming that his initial portfolio starts at the right place, otherwise the first trade will be suboptimal).

4 Persistent Transaction Costs

In some cases the impact of trading on prices may have a non-negligible *persistent* component. If an investor trades weekly and the current prices are unaffected by his trades during the previous week, then the temporary transaction cost model above is appropriate. However, if the frequency of trading is larger than the resiliency of prices, then the investor will be affected by persistent price impact costs.

To study this situation, we extend the model by letting the price be given by $\bar{p}_t = p_t + D_t$ and the investor incur the cost associated with the persistent price distortion D_t in addition to the temporary trading cost TC from before. Hence, the price \bar{p}_t is the sum of the price

p_t without the persistent effect of the investor's own trading (as before) and the new term D_t , which captures the accumulated price distortion due to the investor's (previous) trades. Trading an amount Δx_t pushes prices by $C\Delta x_t$ such that the price distortion becomes $D_t + C\Delta x_t$, where C is Kyle's lambda for persistent price moves. Further, the price distortion mean reverts at a speed (or "resiliency") R . Hence, the price distortion next period ($t + 1$) is:

$$D_{t+1} = (I - R)(D_t + C\Delta x_t). \quad (18)$$

The investor's objective is as before, with a natural modification due to the price distortion:

$$E_0 \left[\sum_t (1 - \rho)^{t+1} \left(x_t^\top [Bf_t - (R + r^f)(D_t + C\Delta x_t)] - \frac{\gamma}{2} x_t^\top \Sigma x_t \right) + (1 - \rho)^t \left(-\frac{1}{2} \Delta x_t^\top \Lambda \Delta x_t + x_{t-1}^\top C \Delta x_t + \frac{1}{2} \Delta x_t^\top C \Delta x_t \right) \right]. \quad (19)$$

Let us explain the various new terms in this objective function. The first term is the position x_t times the expected excess return of the price $\bar{p}_t = p_t + D_t$ given inside the inner square brackets. As before, the expected excess return of p_t is Bf_t . The expected excess return due to the post-trade price distortion is

$$D_{t+1} - (1 + r^f)(D_t + C\Delta x_t) = -(R + r^f)(D_t + C\Delta x_t). \quad (20)$$

The second term is the penalty for taking risk as before. The three terms on the second line of (19) are discounted at $(1 - \rho)^t$ because these cash flows are incurred at time t , not time $t + 1$. The first of these is the temporary transaction cost as before. The second reflects the mark-to-market gain from the old position x_{t-1} from the price impact of the new trade, $C\Delta x_t$. The last term reflects that the traded shares Δx_t are assumed to be executed at the average price distortion, $D_t + \frac{1}{2}C\Delta x_t$. Hence, the traded shares Δx_t earn a mark-to-market gain of $\frac{1}{2}\Delta x_t^\top C\Delta x_t$ as the price moves up an additional $\frac{1}{2}C\Delta x_t$.

The value function is now quadratic in the extended state variable $(x_{t-1}, y_t) \equiv (x_{t-1}, f_t, D_t)$:

$$V(x, y) = -\frac{1}{2}x^\top A_{xx}x + x^\top A_{xy}y + \frac{1}{2}y^\top A_{yy}y + A_0.$$

As before, there exists a unique solution to the Bellman equation and the following proposition characterizes the optimal portfolio strategy.

Proposition 6 *The optimal portfolio x_t is*

$$x_t = x_{t-1} + M^{rate} (aim_t - x_{t-1}), \quad (21)$$

which tracks an aim portfolio, $aim_t = M^{aim}y_t$, that depends on the return-predicting factors and the price distortion, $y_t = (f_t, D_t)$. The coefficient matrices M^{rate} and M^{aim} are stated in the appendix.

The optimal trading policy has a similar structure to before, but the persistent price impact changes both the trading rate and the aim portfolio. The aim is now a weighted average of current and expected future Markowitz portfolios, as well as the current price distortion.

Figure 3 graphically illustrates the optimal trading strategy with temporary and persistent price impact. Panel A uses the parameters from Figures 1–2 with only temporary transaction costs, Panel B has both temporary and persistent transaction costs, while Panel C has purely persistent price impact. Specifically, suppose that the Kyle’s lambda for temporary price impact is $\Lambda = w\tilde{\Lambda}$ and Kyle’s lambda for persistent price impact is $C = (1 - w)\tilde{\Lambda}$, where we vary w to determine how much of the price impact is temporary vs. persistent and where $\tilde{\Lambda}$ is a fixed matrix. Panel A has $w = 1$ (pure temporary costs), Panel B has $w = 0.5$ (both temporary and persistent costs), and Panel C has $w = 0$ (pure persistent costs).⁸

We see that the optimal portfolio policy with persistent transaction costs also tracks the Markowitz portfolio while aiming in front of the target. In Proposition 9 in the Appendix

⁸The parameters from Figures 1–2 are given in Footnote 7 and the additional parameters are $D_0 = 0$, the resiliency $R = 0.1$, and the risk-free rate given by $(1 + r^f)(1 - \rho) = 1$. As further interpretation of Figure 3, note that temporary price impact corresponds to a persistent impact with complete resiliency, $R = 1$. (This holds literally under the natural restriction that the risk-free is the inverse of the discount rate, $(1 + r^f)(1 - \rho) = 1$.) Hence, Panel A has a price impact with complete resiliency, Panel C has a price impact with slow resiliency, and Panel B has two kinds of price impact with, respectively, fast and slow resiliency.

we show more generally that the optimal portfolio under persistent price impact depends on the expected future Markowitz portfolios (i.e., aims in front of the target). This is similar to the case of temporary price impact, but what is different with purely persistent price impact is that the initial trade is larger and, even in continuous time, there can be jumps in the portfolio. This is because, when the price impact is persistent, the trader incurs a transaction cost based on the entire cumulative trade, and therefore is more willing to incur it early in order to start collecting the benefits of a better portfolio. (The resilience still makes it cheaper to postpone part of trade, however). Furthermore, the cost of buying a position and selling it shortly thereafter is much smaller with persistent price impact.

5 Theoretical Applications

We next provide a few simple and useful examples of our model.

Example 1: Timing a single security

An 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, Assumption A ($\Lambda = \lambda\Sigma$) is without loss of generality since all parameters are scalars, and we use the notation $\sigma^2 = \Sigma$ and $B = (\beta^1, \dots, \beta^K)$. Assuming that Φ is diagonal, we can apply Proposition 4 directly to get the optimal timing portfolio:

$$x_t = \left(1 - \frac{a}{\lambda}\right) x_{t-1} + \frac{a}{\lambda} \frac{1}{\gamma\sigma^2} \sum_{i=1}^K \frac{\beta^i}{1 + \phi^i a/\gamma} f_t^i. \quad (22)$$

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 expected return for security s be given by

$$E_t(r_{t+1}^s) = \sum_i \beta^i f_t^{i,s}, \quad (23)$$

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 ,

$$\Delta f_{t+1}^{i,s} = -\phi^i f_t^{i,s} + \varepsilon_{t+1}^{i,s}. \quad (24)$$

We collect the current values of characteristic i for all securities in a vector $f_t^i = \left(f_t^{i,1}, \dots, f_t^{i,S} \right)^\top$, 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 :

$$f_t = \begin{pmatrix} f_t^1 \\ \vdots \\ f_t^I \end{pmatrix}. \quad (25)$$

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

$$B = \beta^\top \otimes I_{S \times S} = \begin{pmatrix} \beta^1 & 0 & 0 & & \beta^I & 0 & 0 \\ 0 & \ddots & 0 & \dots & 0 & \ddots & 0 \\ 0 & 0 & \beta^1 & & 0 & 0 & \beta^I \end{pmatrix}. \quad (26)$$

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 4 to get the optimal characteristic-based relative-value trade as

$$x_t = \left(1 - \frac{a}{\lambda}\right) x_{t-1} + \frac{a}{\lambda} (\gamma \Sigma)^{-1} \sum_{i=1}^I \frac{1}{1 + \phi^i a / \gamma} \beta^i f_t^i. \quad (27)$$

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

$$\max_{x_t} x_t^\top E_t(r_t) - \frac{\gamma}{2} x_t^\top \Sigma x_t - \frac{\lambda}{2} \Delta x_t^\top \Sigma \Delta x_t \quad (28)$$

with a solution:

$$x_t = \frac{\lambda}{\gamma + \lambda} x_{t-1} + \frac{\gamma}{\gamma + \lambda} (\gamma \Sigma)^{-1} E_t(r_{t+1}) = x_{t-1} + \frac{\gamma}{\gamma + \lambda} (\text{Markowitz}_t - x_{t-1}). \quad (29)$$

This optimal *static* portfolio in light of transaction costs differs from our optimal dynamic portfolio in two ways: (i) The weight on the current portfolio x_{t-1} is different, and (ii) the aim portfolio is different since in the static case the aim portfolio is the Markowitz portfolio. Problem (i) with the static portfolio, namely that it prescribes a suboptimal trading rate because it does not account for the future benefits of the position can be fixed by changing the transaction-cost parameter λ (or risk aversion γ or both).

However, problem (ii) cannot be fixed in this way. Interestingly, 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 6 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 expected returns $E_t(r_{t+1}) = B f_t$ by changing B as described in Proposition 4.

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 for K past time periods. 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:

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

Said differently, the matrix Φ 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 $\sigma^2 = \Sigma$. Then we can use Proposition 4 to get the optimal trading strategy

$$\begin{aligned} x_t &= \left(1 - \frac{a}{\lambda}\right) x_{t-1} + \frac{a}{\lambda} \frac{1}{\sigma^2} B (\gamma + a\Phi)^{-1} f_t \\ &= \left(1 - \frac{a}{\lambda}\right) x_{t-1} + \frac{a}{\lambda} \frac{\beta}{\gamma\sigma^2} \sum_{k=1}^K \left(1 - \left(\frac{a}{\gamma+a}\right)^{K+1-k}\right) f_t^k. \end{aligned} \tag{30}$$

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.

6 Empirical Application: Dynamic Trading of Commodity Futures

In this section we illustrate our approach using data on commodity futures. We show how dynamic optimizing can improve performance in an intuitive way, and how it changes the way new information is used.

6.1 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, as seen in Table 1, 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.⁹

For each commodity and each day, we collect the futures price measured in U.S. dollars per contract. For instance, if the gold price is \$1,000 per ounce, the price per contract is \$100,000, since each contract is for 100 ounces. Table 1 provides summary statistics on each contract’s average price, the standard deviation of price changes, the contract multiplier (e.g., 100 ounces per contract in the case of gold), and daily trading volume.

We use the most liquid futures contract of all maturities available. By always using data on the most liquid futures, we are implicitly assuming that the trader’s position is always held in these contracts. Hence, we are assuming that when the most liquid futures nears maturity and the next contract becomes more liquid, the trader “rolls” into the next contract,

⁹Our 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.

i.e., replaces the position in the near contract with the same position in the far contract. Given that rolling does not change a trader’s net exposure, it is reasonable to abstract from the transaction costs associated with rolling. (Traders in the real world do in fact behave like this. There is a separate roll market, which entails far smaller costs than independently selling the “old” contract and buying the “new” one.) When we compute price changes, we always compute the change in price of a given contract (not the difference between the new contract and the old one), since this corresponds to an implementable return. Finally, we collect data on the average daily trading volume per contract as seen in the last column of Table 1. Specifically, we receive an estimate of the average daily volume of the most liquid contract traded electronically and outright (i.e., not including calendar-spread trades) in December 2010 from an asset manager based on underlying data from Reuters.

6.2 Predicting Returns and Other Parameter Estimates

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

$$r_{t+1}^s = 0.001 + 10.32 f_t^{5D,s} + 122.34 f_t^{1Y,s} - 205.59 f_t^{5Y,s} + u_{t+1}^s, \quad (31)$$

(0.17) (2.22) (2.82) (-1.79)

where the left-hand side is the daily 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 five days’ standard deviation of daily price changes, f^{1Y} is the past year’s average daily price change divided by the past year’s standard deviation, and f^{5Y} is the analogous quantity for a five-year window. Hence, the predictors are rolling Sharpe ratios over three different horizons, and, to avoid dividing by a number close to zero, the standard deviations are winsorized below the average tenth percentile of standard deviations. We estimate the regression using feasible generalized least squares and report the t-statistics in brackets.

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

versal over long horizons.¹⁰ 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 the economic insights underlying our novel portfolio optimization. Indeed, the in-sample analysis allows us to focus on the benefits of giving more weight to signals with slower alpha decay, without the added noise in the predictive power of the signals arising when using out-of-sample return forecasts.)

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

$$\begin{aligned}
\Delta f_{t+1}^{5D,s} &= -0.2519 f_t^{5D,s} + \varepsilon_{t+1}^{5D,s} \\
\Delta f_{t+1}^{1Y,s} &= -0.0034 f_t^{1Y,s} + \varepsilon_{t+1}^{1Y,s} \\
\Delta f_{t+1}^{5Y,s} &= -0.0010 f_t^{5Y,s} + \varepsilon_{t+1}^{5Y,s}.
\end{aligned} \tag{32}$$

These mean reversion rates correspond to a 2.4-day half life for the 5-day signal, a 206-day half life for the 1-year signal, and a 700-day half life for the 5-year signal.¹¹

We estimate the variance-covariance matrix Σ using daily price changes over the full sample, shrinking the correlations 50% towards zero. 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, to choose the transaction-cost matrix Λ , we make use of price-impact estimates from the literature. In particular, we use the estimate from Engle, Ferstenberg, and Russell (2008) that trades amounting to 1.59% of the daily volume in a stock have a price impact of about 0.10%. (Breen, Hodrick, and Korajczyk (2002) provides a similar estimate.) Further,

¹⁰Erb and Harvey (2006) document 12-month momentum in commodity futures prices. Asness, Moskowitz, and Pedersen (2008) confirm this finding and also document 5-year reversals. 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.

¹¹The half life is the time it is expected to take for half the signal to disappear. It is computed as $\log(0.5)/\log(1 - 0.2519)$ for the 5-day signal.

Greenwood (2005) finds evidence that market impact in one security spills over to other securities using the specification $\Lambda = \lambda\Sigma$, where we recall that Σ is the variance-covariance matrix. We calibrate Σ as the empirical variance-covariance matrix of price changes, where the covariance are shrunk 50% towards zero for robustness.

We choose the scalar λ based on the Engle, Ferstenberg, and Russell (2008) estimate by calibrating it for each commodity and then computing the mean and median across commodities. Specifically, we collect data on the trading volume of each commodity contract as seen in last column of Table 1 and then calibrate λ for each commodity as follows. Consider, for instance, unleaded gasoline. Since gasoline has a turnover of 11,320 contracts per day and a daily price-change volatility of \$1,340, the transaction cost per contract when one trades 1.59% of daily volume is $1.59\% \times 11,320 \times \lambda^{Gasoline} / 2 \times 1,340^2$, which is 0.10% of the average price per contract of \$48,000 if $\lambda^{Gasoline} = 3 \times 10^{-7}$.

We calibrate the trading costs for the other commodities similarly, and obtain a median value of 5.0×10^{-7} and a mean of 8.4×10^{-7} . There are significant differences across commodities (e.g., the standard deviation is 1.0×10^{-6}), reflecting that these estimates are based on turnover while the specification $\Lambda = \lambda\Sigma$ assumes that transaction costs depend on variances. While our model is general enough to handle transaction costs that depend on turnover (e.g., by using these calibrated λ 's in the diagonal of the Λ matrix), we also need to estimate the spill-over effects (i.e., the off-diagonal elements). Since Greenwood (2005) provides the only estimate of these transaction-cost spill-overs in the literature using the assumption $\Lambda = \lambda\Sigma$ and since real-world transaction costs likely depend on variance as well as turnover, we stick to this specification and calibrate λ as the median across the estimates for each commodity. Naturally, other specifications of the transaction-cost matrix would give slightly different results, but our main purpose is simply to illustrate the economic insights that we have proved in general theoretically.

We also consider a more conservative transaction cost estimate of $\lambda = 10 \times 10^{-7}$. Alternatively, this more conservative analysis can be interpreted as the trading strategy of a larger investor (i.e., we could have equivalently reduced the absolute risk aversion γ).

6.3 Dynamic Portfolio Selection with Trading Costs

We consider three different trading strategies: the optimal trading strategy given by Equation (27) (“optimal”), the optimal trading strategy in the absence of transaction costs (“Markowitz”), and a number of trading strategies based on a static (i.e., one-period) transaction-cost optimization as in Equation (29) (“static optimization”). The static portfolio optimization results in trading partially towards the Markowitz portfolio (as opposed to a aim that depends on signals’ alpha decays) and we consider ten different trading speeds in seen in Table 2. Hence, under the static optimization, the updated portfolio is a weighted average of the Markowitz portfolio (with weight denoted “weight on Markowitz”) and the current portfolio.

Table 2 reports the performance of each strategy as measured by, respectively, its Gross Sharpe Ratio and its Net Sharpe Ratio (i.e., its Sharpe ratio after accounting for transaction costs). Panel A reports these numbers using our base-case transaction-cost estimate (discussed above), while Panel B uses our high transaction-cost estimate. We see that, naturally, the highest SR before transaction costs is achieved by the Markowitz strategy. 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 Markowitz 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 return. 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 4 shows the positions over time of two of the commodity futures, namely Crude and Gold. We

see that the optimal portfolio is a much smoother version of the Markowitz strategy, thus reducing trading costs while at the same time capturing most of the excess return. Indeed, the optimal position tends to be long when the Markowitz portfolio is long and short when the Markowitz portfolio is short, and to be larger when the expected return is large, but moderates the speed and magnitude of trades.

6.4 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 (32). Figure 5 shows the responses to shocks to each return-predicting factor, namely the 5-day factor, the 1-year factor, and the 5-year factor.

The first panel shows that the Markowitz 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 Markowitz strategy.

The second panel shows the response to the 1-year factor. The Markowitz jumps up and decays, whereas the optimal position increases more smoothly and catches up as the Markowitz 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.

7 Conclusion

This paper provides a highly tractable framework for studying optimal trading strategies in the presence of several return predictors, risk and correlation considerations, as well as transaction costs. We derive an explicit closed-form solution for the optimal trading policy, which gives rise to several intuitive results. The optimal portfolio tracks an “aim 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 aim 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 aim 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 dynamic 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.

The model’s tractability makes it potentially useful for future research on return predictability and transaction costs. As one such application of the model, we illustrate in Appendix D how transaction costs can lead to large alphas for short time periods in equilibrium. Indeed, we model an equilibrium with several “noise traders” who trade in and out of their positions with varying mean-reversion speeds and a rational arbitrageur — with trading costs and using the methodology that we derive — who takes 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 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. This can help explain the short-term return reversals documented by Lehman (1990) and Lo and MacKinlay (1990), and their relation to transaction costs documented by Nagel (2011).

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 — i.e., 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. The main insights for portfolio selection can be summarized by the rules that one should aim in front of the target and trade partially towards the current aim.

A Continuous-Time Models

In this section we present continuous-time counterparts to the models with temporary and persistent price impact costs, as well as the case where both types of price impact is present. The continuous-time solutions are in fact even simpler than the discrete-time solutions. Both with temporary and persistent costs, we derive the optimal dynamic portfolio policy explicitly and show how it captures the notion of aiming in front of the target.

A.1 Purely Temporary Transaction Costs

The securities have fundamental prices p with dynamics

$$dp_t = (r^f p_t + B f_t) dt + du_t, \tag{A.1}$$

where, as before, the random “noise” u is a martingale (e.g., a Brownian motion) with instantaneous variance-covariance matrix $\text{var}_t(du_t) = \Sigma dt$, and the predictable component of the excess return is $B f_t$ with

$$df_t = -\Phi f_t dt + d\varepsilon_t. \tag{A.2}$$

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 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 change¹² of his position x_t :

$$dx_t = \tau_t dt. \tag{A.3}$$

¹²We only consider smooth portfolio policies here because discrete jumps in positions or quadratic variation would be associated with infinite trading costs in this 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 transitory cost per time unit of trading τ_t shares per time unit is

$$TC(\tau_t) = \frac{1}{2}\tau_t^\top \Lambda \tau_t. \quad (\text{A.4})$$

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

$$\max_{(\tau_s)_{s \geq t}} E_t \int_t^\infty e^{-\rho(s-t)} \left(x_s^\top B f_s - \frac{\gamma}{2} x_s^\top \Sigma x_s - \frac{1}{2} \tau_s^\top \Lambda \tau_s \right) ds. \quad (\text{A.5})$$

We conjecture a quadratic value function, just as in discrete time: $V(x, f) = -\frac{1}{2}x^\top A_{xx}x + x^\top A_{xf}f + \frac{1}{2}f^\top A_{ff}f + A_0$. The following proposition records the solution to the model; we leave the proof for the more general statement (Proposition 8) in Section A.2.

Proposition 7 *The optimal portfolio x_t tracks a moving “aim 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*

$$\tau_t = \Lambda^{-1}A_{xx} \left(A_{xx}^{-1}A_{xf}f_t - x_t \right), \quad (\text{A.6})$$

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

$$A_{xx} = -\frac{\rho}{2}\Lambda + \Lambda^{\frac{1}{2}} \left(\gamma\Lambda^{-\frac{1}{2}}\Sigma\Lambda^{-\frac{1}{2}} + \frac{\rho^2}{4}I \right)^{\frac{1}{2}} \Lambda^{\frac{1}{2}} \quad (\text{A.7})$$

$$\text{vec}(A_{xf}) = \left(\rho I + \Phi^\top \otimes I_K + I_S \otimes (A_{xx}\Lambda^{-1}) \right)^{-1} \text{vec}(B). \quad (\text{A.8})$$

Under Assumption A, the optimal trading intensity simplifies to

$$\tau_t = \frac{a}{\lambda} (aim_t - x_t), \quad (\text{A.9})$$

with

$$aim_t = \Sigma^{-1}B(\gamma I + a\Phi)^{-1}f_t \quad (\text{A.10})$$

$$a = \frac{1}{2} \left(-\rho\lambda + \sqrt{\rho^2\lambda^2 + 4\gamma\lambda} \right). \quad (\text{A.11})$$

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

A.2 Temporary and Persistent Transaction Costs

We modify the set-up above by adding persistent transaction costs. Specifically, the agent transacts at price $\bar{p}_t = p_t + D_t$, where the distortion D_t evolves according to

$$dD_t = -RD_t dt + Cdx_t = -RD_t dt + C\tau_t dt. \quad (\text{A.12})$$

The agent's objective now becomes

$$\max_{(\tau_s)_{s \geq t}} E_t \int_t^\infty e^{-\rho(s-t)} \left(x_s^\top (\alpha_s - (r + R)D_s + C\tau_s) - \frac{\gamma}{2} x_s^\top \Sigma x_s - \frac{1}{2} \tau_s^\top \Lambda \tau_s \right) ds. \quad (\text{A.13})$$

We conjecture a quadratic value function, as before, in the state variable $(x_t, f_t, D_t) \equiv (x_t, y_t)$ — specifically, we write $V(x, y) = -\frac{1}{2}x^\top A_{xx}x + x^\top A_{xy}y + \frac{1}{2}y^\top A_{yy}y + A_0$.

Proposition 8 *The optimal trading intensity has the form*

$$\tau_t = \bar{M}^{rate} (\bar{M}^{aim} y_t - x_t) \quad (\text{A.14})$$

for appropriate matrices \bar{M}^{rate} and \bar{M}^{aim} defined in the proof.

A.3 Purely Persistent Costs

The set-up is as above, but now we take $\Lambda = 0$. Under this assumption, it no longer follows that x_t has to be of the form $dx_t = \tau_t dt$ for some τ . Indeed, with purely persistent price impact costs, the optimal portfolio policy can have jumps and infinite quadratic variation (i.e., “wiggle” like a Brownian motion). To derive an economically meaningful objective for the agent, we proceed from the discrete-time model and consider the limit as the time interval goes to zero.

The discrete-time objective is given by

$$E_0 \left[\sum_t e^{-\rho(t+\Delta t)} \left(x_t^\top (Bf_t \Delta t - (R + r^f) (D_t + C\Delta x_t) \Delta t) - \frac{\gamma}{2} x_t^\top \Sigma x_t \Delta t \right) + \frac{1}{2} e^{-\rho(t+\Delta t)} (2x_{t-\Delta t}^\top C\Delta x_t + \Delta x_t^\top C\Delta x_t) e^{\rho\Delta t} \right], \quad (\text{A.15})$$

which in the limit, given the definition of D as a lcr1 process, becomes

$$E_t \int_t^\infty e^{-\rho(s-t)} \left(x_s^\top (\alpha_s - (r + R) (D_{s-} + C\Delta x_s)) - \frac{\gamma}{2} x_s^\top \Sigma x_s \right) ds + E_t \int_t^\infty e^{-\rho(s-t)} x_{s-}^\top C dx_s + \frac{1}{2} E_t \int_t^\infty e^{-\rho(s-t)} d[x_s, Cx_s], \quad (\text{A.16})$$

where the price distortion evolves as

$$dD_t = -RD_t dt + Cdx_t. \quad (\text{A.17})$$

A helpful observation in this case is that making a large trade Δx over an infinitesimal time interval has an easily described impact on the value function. In fact, the ability to liquidate one's position instantaneously, and then take a new position, at no cost relative to trading directly to the new position implies

$$V(x, D, f) = V(0, D - Cx, f) - \frac{1}{2} x^\top Cx. \quad (\text{A.18})$$

This intuitive conjecture is verified by constructing a solution to the associated HJB equation that satisfies this condition.

Before stating the result for this version of the model, we introduce the following natural counterpart to Assumption A:

Assumption B. The persistent price impact is proportional to the level of risk, $C = c\Sigma$, where c is a scalar. The resiliency R is a scalar.

Proposition 9 *A value function exists of the quadratic form (E.67), where A_{DD} and A_{Df} are given as solutions to (E.73) and (E.74), respectively. The optimal portfolio is given by (E.71) with $D^0 = D_- - Cx_-$:*

$$x = J^{-1} [(B - C^\top R^\top A_{Df}) f - ((r^f + R) + C^\top R^\top A_{DD}) (D_- - Cx_-)]. \quad (\text{A.19})$$

Under Assumption B, the optimal portfolio is a linear combination of the current price distortion as represented by D^0 , the current Markowitz portfolio, and the exponentially weighted average of future expected Markowitz portfolios, $\int^\infty e^{-(\rho - bRc + R)t} E[\text{Markowitz}_t | f_0] dt$.

B Micro-Foundation for Transaction Costs

In this section we provide formal arguments for our modeling choices by presenting micro-foundations for the types of costs we consider and for their dependence on the period length. As a consequence, we also show that the discrete-time model, when specified in accordance to the micro-foundations, tends to the continuous-time model as the period length goes to zero.

B.1 Temporary Price Impact

To obtain a temporary price impact of trades endogenously, we consider an economy populated by three types of investors: (i) the trader whose optimization problem we study in the paper, referred throughout this section as “the trader,” (ii) “market makers,” who act as intermediaries, and (iii) “end users,” on whom market makers eventually unload their positions as described below.

The temporary price impact is due to the market makers’ inventories. We assume that there are a mass-one continuum of market makers indexed by the set $[0, h]$ and they arrive for the first time at the market at a time equal to their index. The market operates only at discrete times Δ_t apart,¹³ and the market makers trade at the first trading opportunity.

¹³We make the simplifying assumption that $\frac{h}{\Delta_t}$ is an integer.

Once they trade — say, at time t — market makers must spend h units of time gaining access to end users. At time $t + h$, therefore, they unload their inventories at a price p_{t+h} described below, and rejoin the market immediately thereafter. It follows that at each trading date in the market there is always a mass $\frac{\Delta t}{h}$ of competing market makers that clear the market.

The price p , the competitive price of end users, follows an exogenous process and corresponds to the fundamental price in the body of the paper. Market makers take this price as given and trade a quantity q to maximize a quadratic utility:

$$\max_q \left\{ \hat{E}_t [q(p_{t+h} - e^{rh}\hat{p}_t)] - \frac{\gamma^M}{2} \text{Var}_t [q(p_{t+h} - e^{rh}\hat{p}_t)] \right\}, \quad (\text{B.1})$$

where \hat{p}_t is the market price at time t and r is the (continuously-compounded) risk-free rate over the horizon. \hat{E} denotes expectations under the probability measure obtained from the market makers' beliefs using their (normalized) marginal utilities corresponding to $q = 0$ as Radon-Nikodym derivative. Consequently,

$$\hat{E}_t [p_{t+h}] = e^{rh} p_t,$$

so that the maximization problem becomes

$$\max_q \left\{ q(p_t - \hat{p}_t) - e^{-rh} \frac{\gamma^M}{2} \text{Var}_t [qp_{t+h}] \right\}. \quad (\text{B.2})$$

The price \hat{p} is set so as to satisfy the market-clearing condition

$$0 = \Delta x_t + q \frac{\Delta t}{h}. \quad (\text{B.3})$$

Since p is exogenous and Gaussian with variance V_h h periods ahead that can be calculated easily,¹⁴ the maximization problem yields

$$\hat{p}_t = p_t + e^{-rh} \gamma^M V_h \frac{\Delta x_t}{\Delta t} h. \quad (\text{B.4})$$

¹⁴The resulting value is $V_h = \Sigma h + B N_h \Omega N_h^\top B^\top$, where $N_h = \int_0^h \int_u^h e^{-\Phi(t-u)} dt du = \Phi^{-1} h - \Phi^{-2} (I - e^{-\Phi h})$ if Φ is invertible. (Note that the first term, Σh , is of order h , while the second of order h^2 .)

Consequently, if the trader trades an amount Δx_t , he trades at the unit price of p_t and pays an additional transaction cost of

$$e^{-rh}\gamma^M \Delta x_t^\top V_h \frac{\Delta x_t}{\Delta_t} h,$$

which has the quadratic form posited in the body of the paper.

Two cases suggest themselves naturally when considering the choice for the holding period h as a function of Δ_t . In the first case, a decreasing Δ_t is thought of as an improvement in the trading technology, attention, etc., of all market participants, and therefore h decreases as Δ_t does — in its simplest form, $h = \Delta_t$, which yields a transaction cost of the order Δ_t^2 . Generally, as long as $h \rightarrow 0$ as $\Delta_t \rightarrow 0$, the transaction costs also vanishes.

The second case is that of a constant h : the dealers need a fixed amount of time to lay off a position regardless of the frequency with which our original traders access the market. It follows, in this case, that the price impact does not vanish as Δ_t becomes small: in the continuous-time limit ($\Delta_t \rightarrow 0$), the per-unit-of-time transaction cost is proportional to

$$\lim_{\Delta_t \rightarrow 0} \frac{\Delta x_t^\top}{\Delta_t} V_h h \frac{\Delta x_t}{\Delta_t} = \tau^\top V_h h \tau, \tag{B.5}$$

as assumed in Section A.1. One can therefore interpret Δ_t in this case as the frequency with which the researcher observes the world, which does not impact (to the first order) equilibrium quantities — in particular, flow trades and costs.

B.2 Persistent Price Impact

A similar model, but with a different specification of the market makers, can be used to justify a persistent price impact. Consider therefore the same model as in the previous section, but suppose now that market makers do not hold their inventories for a deterministic number h of time units, but rather manage to deplete them, through trade with end users at price p , at a constant rate ψ . Thus, between two trading dates with the trader, a market maker's

inventory evolves according to

$$\Delta I_{t+1} = -\psi I_t \Delta_t + q_t, \quad (\text{B.6})$$

where, in equilibrium,

$$q_t = \Delta x_t.$$

The market makers continue to maximize a quadratic objective:

$$\max_q \left\{ \hat{E}_t \sum_{s=t+n\Delta_t} e^{-r(s-t)} \left(\psi I_s^\top p_s \Delta_t - q_s^\top \hat{p}_s - \frac{\gamma^M}{2} I_s^\top V_{\Delta_t} I_s \right) \right\}, \quad (\text{B.7})$$

subject to (B.6) and expectations about q described below. Note that the market maker's objective depends (positively) on the expected cash flows $\psi I_s^\top p_s \Delta_t - q_s^\top \hat{p}_s$ due to future trades with the end user and the trader and negatively on the risk of his inventory.

We assume that market makers cannot predict the trader's order flow Δx . More specifically, according to their probability distribution,

$$\hat{E} [\Delta x_t | \mathcal{F}_s, s < t] = 0 \quad (\text{B.8})$$

$$\hat{E} [(\Delta x_t)^2 | \mathcal{F}_s, s < t] = v. \quad (\text{B.9})$$

Moments of q_s and I_s follow immediately.

The first-order condition with respect to q_t is

$$0 = \hat{E}_t \sum_{s=t+n\Delta_t} e^{-r(s-t)} \left(\psi p_s^\top - \gamma^M I_s^\top \frac{V_{\Delta_t}}{\Delta_t} \right) \frac{\partial I_s}{\partial q_t} \Delta_t - \hat{p}_t^\top. \quad (\text{B.10})$$

Using the fact that $\frac{\partial I_s}{\partial q_t} = (1 - \psi \Delta_t)^{s-t}$, the first-order condition yields

$$\hat{p}_t = \hat{E}_t \sum_{s=t+n\Delta_t} e^{-r(s-t)} (1 - \psi \Delta_t)^{s-t} \left(\psi p_s - \gamma^M \frac{V_{\Delta_t}}{\Delta_t} I_s \right) \Delta_t. \quad (\text{B.11})$$

Using the facts that $\hat{E}_t[e^{-r(s-t)}p_s] = p_t$ and $\hat{E}_t[I_s] = (1 - \psi\Delta_t)^{t-s}I_t$, we obtain

$$\hat{p}_t = p_t - \kappa_I I_t \tag{B.12}$$

for a constant matrix

$$\kappa_I = \sum_{n=0}^{\infty} e^{-rn\Delta_t} (1 - \psi\Delta_t)^{2n\Delta_t} \gamma^M \frac{V_{\Delta_t}}{\Delta_t} \Delta_t. \tag{B.13}$$

The price \hat{p}_t is only the price at the end of trading date t — the price at which the last unit of the q_t shares is traded. We assume that, during the trading date, orders of infinitesimal size come to market sequentially and the market makers' expectation is that the remainder of date- t order flow aggregates to zero — thus, the order flow is a martingale. It follows that the price paid for the k^{th} percentile of the order flow q_t is $p_t - \kappa_I(I_{t-1} + kq_t)$. This mechanism ensures that round-trip trades over very short intervals do not have transaction-cost implications, just as in Section 4.

We note that, as $\Delta_t \rightarrow 0$, (B.12) continues to hold with

$$\kappa_I = \frac{\gamma^M}{r} \lim_{\Delta_t \rightarrow 0} \frac{V_{\Delta_t}}{\Delta_t} \tag{B.14}$$

$$= \frac{\gamma^M}{r} \Sigma. \tag{B.15}$$

This price specification is the same as in Section A.3, with

$$D_t = -\kappa_I I_t \tag{B.16}$$

$$\begin{aligned} dD_t &= -\kappa_I dI_t \\ &= -\kappa_I (-\psi I_t + \tau_t) dt \\ &= -\kappa_I \psi \kappa_I^{-1} D_t dt - \kappa_I \tau_t dt \\ &= -RD_t dt + C^\top \tau_t dt. \end{aligned} \tag{B.17}$$

B.3 Temporary and Persistent Price Impact

The two types of price impact can obtain simultaneously in this model, provided that two kinds of market makers coexist and interact in particular ways. Specifically, suppose that one group of market makers transact with the trader. After a period of length h , these market makers clear their inventories with a second group of market makers, who specialize in locating end users and trading with them. As in Section B.2, these market makers deplete their inventories only gradually (at a constant rate), giving rise to a persistent impact. The trader must compensate both groups of market makers for the risk taken, resulting in the two price-impact components.

C 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. There is one choice to be made concerning the adjustment, namely that of Λ . The model in Section B.1, in particular, suggests the following two natural possibilities:

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

$$\hat{\Lambda}(\Delta_t) = \Delta_t\Lambda \quad \text{or} \quad \hat{\lambda}(\Delta_t) = \lambda. \quad (\text{C.1}')$$

The adjustments in Equations (C.1)–(C.1') are as implied by (B.4), the relation that justifies costs quadratic in Δx . They differ with respect to the assumption made about the holding period h : Equation (C.1) obtains with constant h , while Equation (C.1') with $h = \Delta_t$. These adjustments follow from Section B.1 directly in the case in which trading costs are proportional to the variance Σ , which gives the dependence of $\hat{\lambda}$ on Δ_t . The dependence of $\hat{\Lambda}$ follows from $\Lambda = \lambda\Sigma$ in this case, and in general by analogy.

The other adjustments are immediate. For instance, Equations (C.2)–(C.3) simply state

that the variance is proportional to time:

$$\hat{\Sigma}(\Delta_t) = \Sigma \Delta_t \tag{C.2}$$

$$\hat{\Omega}(\Delta_t) = \Omega \Delta_t \tag{C.3}$$

$$\hat{B}(\Delta_t) = B \Delta_t \tag{C.4}$$

$$\hat{\Phi}(\Delta_t) = \Phi \Delta_t \tag{C.5}$$

$$\hat{\rho}(\Delta_t) = \rho \Delta_t \tag{C.6}$$

$$\hat{\gamma}(\Delta_t) = \gamma \tag{C.7}$$

$$\hat{R}(\Delta_t) = R \Delta_t \tag{C.8}$$

$$\hat{C}(\Delta_t) = C. \tag{C.9}$$

$$\hat{r}^f(\Delta_t) = r^f \Delta_t \tag{C.10}$$

We are ready to state the result.

Proposition 10 (i) *Consider the discrete-time model of Section 4 with parameters defined to depend on the time interval Δ_t as stated by (C.1) and (C.2)–(C.9). Then*

$$\lim_{\Delta_t \rightarrow 0} \frac{M^{rate}(\Delta_t)}{\Delta_t} = \bar{M}^{rate} \tag{C.11}$$

$$\lim_{\Delta_t \rightarrow 0} M^{aim}(\Delta_t) = \bar{M}^{aim}, \tag{C.12}$$

where \bar{M}^{rate} and \bar{M}^{speed} are the continuous-time matrix coefficients of Proposition 8.

Given that f is continuous almost everywhere on every path almost surely, (C.11) and (C.12) imply that, fixing (x_t, f_t, D_t) , $\lim_{\Delta_t \rightarrow 0} \frac{x_{t+\Delta_t} - x_t}{\Delta_t} \rightarrow \tau_t$ a.s.

(ii) If (C.1') holds instead of (C.1), then the solution approaches that of Section A.3, summarized in Proposition 9.

Equation (C.12) states that, given the current conditions, the aim portfolio in discrete time approaches the one in continuous time as the time period between trades becomes small. Equation (C.11) concerns the tracking speed: the per-unit-of-time fraction of the distance to the aim covered by a trade is the same, for small Δ_t , as in continuous time. Consequently,

the optimal solutions have the same dynamics, in the limit.

D 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 excess return.

For simplicity, we consider a model set in continuous time, as detailed in Section A.1, featuring 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

$$dz_t^l = \kappa (f_t^l - z_t^l) dt \tag{D.1}$$

$$df_t^l = -\psi_l f_t^l dt + dW_t^l. \tag{D.2}$$

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

$$dz_t = \kappa \left(\sum_{l=1}^L f_t^l - z_t \right) dt. \tag{D.3}$$

We conjecture that the investor's inference problem is as studied in Section A.1, 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 11 below.

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

$$\Phi = \begin{pmatrix} \psi_1 & 0 & \cdots & 0 \\ 0 & \psi_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ -\kappa & -\kappa & \cdots & \kappa \end{pmatrix}. \tag{D.4}$$

Suppose that the only other investors in the economy are the investors considered in Section A.1, 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 (D.1)–(D.2), the restriction imposed by equilibrium is that the dynamics of the price are such that, for all t ,

$$x_t = -z_t \tag{D.5}$$

$$dx_t = -dz_t. \tag{D.6}$$

Using (A.9), these equilibrium conditions lead to

$$\frac{a}{\lambda}\sigma^{-2}B(a\Phi + \gamma I)^{-1} + \frac{a}{\lambda}e_{L+1} = -\kappa(\mathbf{1} - 2e_{L+1}), \tag{D.7}$$

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 $-z_t = -f_t^{L+1}$ at time t for all t , then the factor loadings must be given by

$$B = \sigma^2 \left[-\frac{\lambda}{a}\kappa(\mathbf{1} - 2e_{L+1}) - e_{L+1} \right] (a\Phi + \gamma I). \tag{D.8}$$

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

$$\begin{aligned} B_l &= -\sigma^2\kappa(\lambda\psi_l + \lambda\gamma a^{-1} + \lambda\kappa - a) \\ &= -\lambda\sigma^2\kappa(\psi_l + \rho + \kappa), \end{aligned} \tag{D.9}$$

while

$$B_{L+1} = \sigma^2(\rho\lambda\kappa + \lambda\kappa^2 - \gamma). \tag{D.10}$$

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

Proposition 11 *The market is in equilibrium if and only if $x_0 = -z_0$ and the security's*

expected excess return is given by

$$\frac{1}{dt} E_t[dp_t - r^f p_t dt] = \sum_{l=1}^L \lambda \sigma^2 \kappa (\psi_l + \rho + \kappa) (-f_t^l) + \sigma^2 (\rho \lambda \kappa + \lambda \kappa^2 - \gamma) z_t. \quad (\text{D.11})$$

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.

E Proofs

Proof of Proposition 1. Assuming that the value function is of the posited form, we calculate the expected future value function as

$$\begin{aligned} E_t[V(x_t, f_{t+1})] &= -\frac{1}{2} x_t^\top A_{xx} x_t + x_t^\top A_{xf} (I - \Phi) f_t + \frac{1}{2} f_t^\top (I - \Phi)^\top A_{ff} (I - \Phi) f_t \\ &\quad + \frac{1}{2} E_t(\varepsilon_{t+1}^\top A_{ff} \varepsilon_{t+1}) + A_0. \end{aligned} \quad (\text{E.1})$$

Let $\bar{\rho} = 1 - \rho$ and $\bar{\Lambda} = \bar{\rho}^{-1} \Lambda$. The agent maximizes the quadratic objective $-\frac{1}{2} x_t^\top J_t x_t + x_t^\top j_t + d_t$ with

$$\begin{aligned} J_t &= \gamma \Sigma + \bar{\Lambda} + A_{xx} \\ j_t &= (B + A_{xf} (I - \Phi)) f_t + \bar{\Lambda} x_{t-1} \\ d_t &= -\frac{1}{2} x_{t-1}^\top \bar{\Lambda} x_{t-1} + \frac{1}{2} f_t^\top (I - \Phi)^\top A_{ff} (I - \Phi) f_t + \frac{1}{2} E_t(\varepsilon_{t+1}^\top A_{ff} \varepsilon_{t+1}) + A_0. \end{aligned} \quad (\text{E.2})$$

The maximum value is attained by

$$x_t = J_t^{-1} j_t, \quad (\text{E.3})$$

and it is equal to $V(x_{t-1}, f_t) = \frac{1}{2} j_t^\top J_t^{-1} j_t + d_t$. Combining this fact with (6) we obtain an equation that must hold for all x_{t-1} and f_t , which implies the following restrictions on the coefficient matrices:¹⁵

$$-\bar{\rho}^{-1} A_{xx} = \bar{\Lambda}(\gamma\Sigma + \bar{\Lambda} + A_{xx})^{-1} \bar{\Lambda} - \bar{\Lambda} \quad (\text{E.4})$$

$$\bar{\rho}^{-1} A_{xf} = \bar{\Lambda}(\gamma\Sigma + \bar{\Lambda} + A_{xx})^{-1} (B + A_{xf}(I - \Phi)) \quad (\text{E.5})$$

$$\begin{aligned} \bar{\rho}^{-1} A_{ff} &= (B + A_{xf}(I - \Phi))^\top (\gamma\Sigma + \bar{\Lambda} + A_{xx})^{-1} (B + A_{xf}(I - \Phi)) \\ &\quad + (I - \Phi)^\top A_{ff} (I - \Phi). \end{aligned} \quad (\text{E.6})$$

The existence of a solution to this system of Riccati equations can be established using standard results, e.g., as in Ljungqvist and Sargent (2004). In this case, however, we can derive explicit expressions, as follows. We start by letting $Z = \bar{\Lambda}^{-\frac{1}{2}} A_{xx} \bar{\Lambda}^{-\frac{1}{2}}$ and $M = \bar{\Lambda}^{-\frac{1}{2}} \Sigma \bar{\Lambda}^{-\frac{1}{2}}$ and rewriting Equation (E.4) as

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

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

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

¹⁵Remember that A_{xx} and A_{ff} can always be chosen symmetric.

resulting in

$$Z = \left(\bar{\rho}\gamma M + \frac{1}{4}(\rho I + \gamma M)^2 \right)^{\frac{1}{2}} - \frac{1}{2}(\rho I + \gamma M) \quad (\text{E.7})$$

$$A_{xx} = \bar{\Lambda}^{\frac{1}{2}} \left[\left(\bar{\rho}\gamma M + \frac{1}{4}(\rho I + \gamma M)^2 \right)^{\frac{1}{2}} - \frac{1}{2}(\rho I + \gamma M) \right] \bar{\Lambda}^{\frac{1}{2}}, \quad (\text{E.8})$$

that is,

$$A_{xx} = \left(\bar{\rho}\gamma \bar{\Lambda}^{\frac{1}{2}} \Sigma \bar{\Lambda}^{\frac{1}{2}} + \frac{1}{4}(\rho^2 \bar{\Lambda}^2 + 2\rho\gamma \bar{\Lambda}^{\frac{1}{2}} \Sigma \bar{\Lambda}^{\frac{1}{2}} + \gamma^2 \bar{\Lambda}^{\frac{1}{2}} \Sigma \bar{\Lambda}^{-1} \Sigma \bar{\Lambda}^{\frac{1}{2}}) \right)^{\frac{1}{2}} - \frac{1}{2}(\rho \bar{\Lambda} + \gamma \Sigma). \quad (\text{E.9})$$

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

The other value-function coefficient determining optimal trading is A_{xf} , which solves the linear equation (E.5). To write the solution explicitly, we note first that, from (E.4),

$$\bar{\Lambda}(\gamma \Sigma + \bar{\Lambda} + A_{xx})^{-1} = I - A_{xx} \Lambda^{-1}. \quad (\text{E.10})$$

Using the general rule that $\text{vec}(XYZ) = (Z^\top \otimes X) \text{vec}(Y)$, we rewrite (E.5) in vectorized form:

$$\begin{aligned} \text{vec}(A_{xf}) &= \bar{\rho} \text{vec}((I - A_{xx} \Lambda^{-1})B) \\ &\quad + \bar{\rho}((I - \Phi)^\top \otimes (I - A_{xx} \Lambda^{-1})) \text{vec}(A_{xf}), \end{aligned} \quad (\text{E.11})$$

so that

$$\text{vec}(A_{xf}) = \bar{\rho} (I - \bar{\rho}(I - \Phi)^\top \otimes (I - A_{xx} \Lambda^{-1}))^{-1} \text{vec}((I - A_{xx} \Lambda^{-1})B). \quad (\text{E.12})$$

Finally, A_{ff} is calculated from the linear equation (E.6), which is of the form

$$\bar{\rho}^{-1} A_{ff} = Q + (I - \Phi)^\top A_{ff} (I - \Phi) \quad (\text{E.13})$$

with

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

a positive-definite matrix.

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

$$A_{ff,ij} = \frac{\bar{\rho} Q_{ij}}{1 - \bar{\rho}(1 - \Phi_{ii})(1 - \Phi_{jj})}. \quad (\text{E.14})$$

In general,

$$\text{vec}(A_{ff}) = \bar{\rho} (I - \bar{\rho}(I - \Phi)^\top \otimes (I - \Phi)^\top)^{-1} \text{vec}(Q). \quad (\text{E.15})$$

One way to see that A_{ff} is positive definite is to iterate (E.13) starting with $A_{ff}^0 = 0$.

We conclude that the posited value function satisfies the Bellman equation. ■

Proof of Proposition 2. Differentiating the Bellman equation (5) with respect to x_{t-1} gives

$$-A_{xx}x_{t-1} + A_{xf}f_t = \Lambda(x_t - x_{t-1}),$$

which clearly implies (7) and (8).

In the case $\Lambda = \lambda \Sigma$ for some scalar $\lambda > 0$ and letting $\bar{\lambda} = \bar{\rho}^{-1}\lambda$, the solution to the value-function coefficients is $A_{xx} = a\Sigma$, where a solves a simplified version of (E.4):

$$-\bar{\rho}^{-1}a = \frac{\bar{\lambda}^2}{\gamma + \bar{\lambda} + a} - \bar{\lambda}, \quad (\text{E.16})$$

or

$$a^2 + (\gamma + \bar{\lambda}\rho)a - \lambda\gamma = 0, \quad (\text{E.17})$$

with solution

$$a = \frac{\sqrt{(\gamma + \bar{\lambda}\rho)^2 + 4\gamma\lambda} - (\gamma + \bar{\lambda}\rho)}{2}. \quad (\text{E.18})$$

It follows immediately that $\Lambda^{-1}A_{xx} = a/\lambda$. ■

Proof of Proposition 3. We show that

$$aim_t = (\gamma\Sigma + A_{xx})^{-1} (\gamma\Sigma \times Markowitz_t + A_{xx} \times E_t(aim_{t+1})) \quad (\text{E.19})$$

by using (8), (E.5), and (E.4) successively to write

$$\begin{aligned} aim_t &= A_{xx}^{-1} A_{xf} f_t \\ &= A_{xx}^{-1} \Lambda (\gamma\Sigma + \bar{\Lambda} + A_{xx})^{-1} (\gamma\Sigma \times Markowitz_t + A_{xx} \times E_t(aim_{t+1})) \\ &= (\gamma\Sigma + A_{xx})^{-1} (\gamma\Sigma \times Markowitz_t + A_{xx} \times E_t(aim_{t+1})). \end{aligned} \quad (\text{E.20})$$

Equation (12) follows immediately as a special case.

For part (ii), we iterate (E.19) forward to obtain

$$\begin{aligned} aim_t &= (\gamma\Sigma + A_{xx})^{-1} \times \\ &\quad \sum_{\tau=t}^{\infty} (A_{xx} (\gamma\Sigma + A_{xx})^{-1})^{(\tau-t)} \gamma\Sigma \times E_t(Markowitz_{\tau}), \end{aligned}$$

which specializes to (13). ■

Proof of Proposition 4. In the case $\Lambda = \lambda\Sigma$, equation (E.5) is solved by

$$\begin{aligned} A_{xf} &= \lambda B((\gamma + \bar{\lambda} + a)I - \lambda(I - \Phi))^{-1} \\ &= \lambda B((\gamma + \bar{\lambda}\rho + a)I + \lambda\Phi)^{-1} \end{aligned} \quad (\text{E.21})$$

$$= B\left(\frac{\gamma}{a} + \Phi\right)^{-1}, \quad (\text{E.22})$$

and thus the aim portfolio is

$$aim_t = (a\Sigma)^{-1}B \left(\frac{\gamma}{a} + \Phi \right)^{-1} f_t, \quad (\text{E.23})$$

which is the same as (14). Equation (15) is immediate. ■

Proof of Proposition 5. Rewriting (7) as

$$x_t = (I - \Lambda^{-1}A_{xx})x_{t-1} + \Lambda^{-1}A_{xx} \times aim_t \quad (\text{E.24})$$

and iterating this relation backwards gives

$$x_t = \sum_{\tau=-\infty}^t (I - \Lambda^{-1}A_{xx})^{t-\tau} \Lambda^{-1}A_{xx} \times aim_\tau. \quad (\text{E.25})$$

■

Proof of Proposition 6. We start by defining

$$\begin{aligned} \Pi &= \begin{bmatrix} \Phi & 0 \\ 0 & R \end{bmatrix}, & \tilde{C} &= (1 - R) \begin{bmatrix} 0 \\ C \end{bmatrix}, \\ \tilde{B} &= \begin{bmatrix} B & -(R + r^f) \end{bmatrix}, & & \\ \tilde{\Omega} &= \begin{bmatrix} \Omega & 0 \\ 0 & 0 \end{bmatrix}, & \tilde{\varepsilon}_t &= \begin{bmatrix} \varepsilon_t \\ 0 \end{bmatrix}. \end{aligned} \quad (\text{E.26})$$

It is useful to keep in mind that $y_t = (f_t^\top, D_t^\top)^\top$ (a column vector). Given this definition, it follows that

$$E_t[y_{t+1}] = (I - \Pi)y_t + \tilde{C}(x_t - x_{t-1}). \quad (\text{E.27})$$

The conjectured value function is

$$V(x_{t-1}, y_t) = -\frac{1}{2}x_{t-1}^\top A_{xx}x_{t-1} + x_{t-1}^\top A_{xy}y_t + \frac{1}{2}y_t^\top A_{yy}y_t + A_0, \quad (\text{E.28})$$

so that

$$\begin{aligned}
E_t [V(x_t, y_{t+1})] &= -\frac{1}{2}x_t^\top A_{xx}x_t + x_t^\top A_{xy} \left((I - \Pi)y_t + \tilde{C}(x_t - x_{t-1}) \right) + \quad (\text{E.29}) \\
&\quad \frac{1}{2} \left((I - \Pi)y_t + \tilde{C}(x_t - x_{t-1}) \right)^\top A_{yy} \left((I - \Pi)y_t + \tilde{C}(x_t - x_{t-1}) \right) + \\
&\quad \frac{1}{2}E_t [\tilde{\varepsilon}_{t+1}^\top A_{yy}\tilde{\varepsilon}_{t+1}] + A_0.
\end{aligned}$$

The trader consequently chooses x_t to solve

$$\begin{aligned}
\max_x \left\{ x^\top \tilde{B}y_t - x^\top (R + r^f)C(x - x_{t-1}) - \frac{\gamma}{2}x^\top \Sigma x \right. \\
+ \frac{1}{2}\bar{\rho}^{-1} (x^\top Cx - x_{t-1}^\top Cx_{t-1} - (x - x_{t-1})^\top \Lambda(x - x_{t-1})) \\
- \frac{1}{2}x^\top A_{xx}x + x^\top A_{xy} \left((I - \Pi)y_t + \tilde{C}(x - x_{t-1}) \right) \quad (\text{E.30}) \\
\left. + \frac{1}{2} \left((I - \Pi)y_t + \tilde{C}(x - x_{t-1}) \right)^\top A_{yy} \left((I - \Pi)y_t + \tilde{C}(x - x_{t-1}) \right) \right\},
\end{aligned}$$

which is a quadratic of the form $-\frac{1}{2}x^\top Jx + x^\top j_t + d_t$, with

$$J = \gamma\Sigma + \bar{\Lambda} + (2(R + r^f) - \bar{\rho}^{-1})C + A_{xx} - 2A_{xy}\tilde{C} - \tilde{C}^\top A_{yy}\tilde{C} \quad (\text{E.31})$$

$$\begin{aligned}
j_t &= \tilde{B}y_t + (\bar{\Lambda} + (R + r^f)C)x_{t-1} + A_{xy} \left((I - \Pi)y_t - \tilde{C}x_{t-1} \right) + \quad (\text{E.32}) \\
&\quad \tilde{C}A_{yy} \left((I - \Pi)y_t - \tilde{C}x_{t-1} \right)
\end{aligned}$$

$$\equiv S_x x_{t-1} + S_y y_t \quad (\text{E.33})$$

$$\begin{aligned}
d_t &= -\frac{1}{2}x_{t-1}^\top \bar{\Lambda}x_{t-1} - \frac{1}{2}\bar{\rho}^{-1}x_{t-1}^\top Cx_{t-1} + \quad (\text{E.34}) \\
&\quad \frac{1}{2} \left((I - \Pi)y_t - \tilde{C}x_{t-1} \right)^\top A_{yy} \left((I - \Pi)y_t - \tilde{C}x_{t-1} \right).
\end{aligned}$$

Here,

$$S_x = \bar{\Lambda} + (R + r^f)C - A_{xy}\tilde{C} - \tilde{C}^\top A_{yy}\tilde{C} \quad (\text{E.35})$$

$$S_y = \tilde{B} + A_{xy}(I - \Pi) + \tilde{C}^\top A_{yy}(I - \Pi). \quad (\text{E.36})$$

The value of x attaining the maximum is given by

$$x_t = J^{-1}j_t, \quad (\text{E.37})$$

and the maximal value is

$$\frac{1}{2}j_t^\top J^{-1}j_t + d_t = V(x_{t-1}, y_t) - A_0 \quad (\text{E.38})$$

$$= -\frac{1}{2}x_{t-1}^\top A_{xx}x_{t-1} + x_{t-1}^\top A_{xy}y_t + \frac{1}{2}y_t^\top A_{yy}y_t. \quad (\text{E.39})$$

The unknown matrices have to satisfy a system of equations encoding the equality of all coefficients in (E.39). Thus,

$$-\bar{\rho}^{-1}A_{xx} = S_x^\top J^{-1}S_x - \bar{\Lambda} - \bar{\rho}^{-1}C + \tilde{C}^\top A_{yy}\tilde{C} \quad (\text{E.40})$$

$$\bar{\rho}^{-1}A_{xy} = S_x^\top J^{-1}S_y - \tilde{C}^\top A_{yy}(I - \Pi) \quad (\text{E.41})$$

$$\bar{\rho}^{-1}A_{yy} = S_y^\top J^{-1}S_y + (I - \Pi)^\top A_{yy}(I - \Pi). \quad (\text{E.42})$$

For our purposes, the more interesting observation is that the optimal position x_t is rewritten as

$$x_t = x_{t-1} + \underbrace{(I - J^{-1}S_x)}_{M^{\text{rate}}} \left(\underbrace{(I - J^{-1}S_x)^{-1} (J^{-1}S_y)}_{\text{aim}_t = M^{\text{aim}}y_t} y_t - x_{t-1} \right). \quad (\text{E.43})$$

■

Proof of Proposition 8. Using the notation (E.26), the HJB equation is

$$\begin{aligned} \rho V &= \max_{\tau} \left\{ x^\top (\tilde{B}y + C\tau) - \frac{\gamma}{2}x^\top \Sigma x - \frac{1}{2}\tau^\top \Lambda \tau + \frac{\partial V}{\partial x}\tau + \frac{\partial V}{\partial y}(-\Pi y + \tilde{C}\tau) + e \right\} \\ &= \max_{\tau} \left\{ x^\top \tilde{B}y - \frac{\gamma}{2}x^\top \Sigma x - \frac{1}{2}\tau^\top \Lambda \tau + \tau^\top (Q_x x + Q_y y) - \frac{\partial V}{\partial y}\Pi y + e \right\}, \end{aligned}$$

where

$$e = \frac{1}{2} \text{tr} \left(\tilde{\Omega} \frac{\partial^2 V}{\partial y \partial y^\top} \right) \quad (\text{E.44})$$

$$Q_x = -A_{xx} + \tilde{C}^\top A_{xy}^\top + C^\top$$

$$Q_y = A_{xy} + \tilde{C}^\top A_{yy}. \quad (\text{E.45})$$

It follows immediately that

$$\begin{aligned} \tau_t &= -\Lambda^{-1} Q_x [aim_t - x_t] \\ &= \Lambda^{-1} \left(A_{xx} - \tilde{C}^\top A_{xy}^\top - C^\top \right) [aim_t - x_t] \\ &\equiv \bar{M}^{rate} [aim_t - x_t], \end{aligned} \quad (\text{E.46})$$

with

$$\begin{aligned} aim_t &= -(Q_x^{-1} Q_y) y_t \\ &= \left(A_{xx} - \tilde{C}^\top A_{xy}^\top - C^\top \right)^{-1} \left(A_{xy} + \tilde{C}^\top A_{yy} \right) y_t \\ &\equiv \bar{M}^{aim} y_t. \end{aligned} \quad (\text{E.47})$$

The coefficient matrices solve the system

$$\begin{aligned} -\rho A_{xx} &= -\gamma \Sigma + Q_x^\top \Lambda^{-1} Q_x \\ &= -\gamma \Sigma + \left(A_{xx} - A_{xy} \tilde{C} - C \right) \Lambda^{-1} \left(A_{xx} - \tilde{C}^\top A_{xy}^\top - C^\top \right) \\ \rho A_{xy} &= Q_x^\top \Lambda^{-1} Q_y + \tilde{B} - A_{xy} \Pi \\ &= - \left(A_{xx} - A_{xy} \tilde{C} - C \right) \Lambda^{-1} \left(A_{xy} + \tilde{C}^\top A_{yy} \right) + \tilde{B} - A_{xy} \Pi \\ \rho A_{yy} &= Q_y^\top \Lambda^{-1} Q_y - 2A_{yy} \Pi \\ &= \left(A_{xy}^\top + A_{yy} \tilde{C} \right) \Lambda^{-1} \left(A_{xy} + \tilde{C}^\top A_{yy} \right) - 2A_{yy} \Pi. \end{aligned} \quad (\text{E.48})$$

We note that the equations above have to be solved simultaneously for A_{xx} , A_{xy} , and A_{yy} ; there is no closed-form solution in general. The complication is due to the fact that current

trading affects the persistent price component D (that is, $C \neq 0$). The special case $C = 0$ (and $D_0 = 0$) — i.e., purely transitory costs — does lead to closed-form solutions — which are even slightly simpler than in discrete time — as follows.

Specializing (E.46) and (E.47) to this case, we obtain

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

while (E.48) becomes

$$-\rho A_{xx} = A_{xx}\Lambda^{-1}A_{xx} - \gamma\Sigma \tag{E.49}$$

$$\rho A_{xf} = -A_{xx}\Lambda^{-1}A_{xf} - A_{xf}\Phi + B \tag{E.50}$$

$$\rho A_{ff} = A_{xf}^\top\Lambda^{-1}A_{xf} - 2A_{ff}\Phi. \tag{E.51}$$

Pre- and post-multiplying (E.49) by $\Lambda^{-\frac{1}{2}}$, we obtain

$$-\rho Z = Z^2 + \frac{\rho^2}{4}I - Y, \tag{E.52}$$

that is,

$$\left(Z + \frac{\rho}{2}I\right)^2 = y, \tag{E.53}$$

where

$$Z = \Lambda^{-\frac{1}{2}}A_{xx}\Lambda^{-\frac{1}{2}} \tag{E.54}$$

$$Y = \gamma\Lambda^{-\frac{1}{2}}\Sigma\Lambda^{-\frac{1}{2}} + \frac{\rho^2}{4}I. \tag{E.55}$$

This leads to

$$Z = -\frac{\rho}{2}I + Y^{\frac{1}{2}} \geq 0, \tag{E.56}$$

implying that

$$A_{xx} = -\frac{\rho}{2}\Lambda + \Lambda^{\frac{1}{2}} \left(\gamma\Lambda^{-\frac{1}{2}}\Sigma\Lambda^{-\frac{1}{2}} + \frac{\rho^2}{4} \right)^{\frac{1}{2}} \Lambda^{\frac{1}{2}}. \quad (\text{E.57})$$

The solution for A_{xf} follows from Equation (E.50), using the general rule that $\text{vec}(XYZ) = (Z^\top \otimes X) \text{vec}(Y)$:

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

Just as in discrete time, the form of the solution further simplifies under Assumption A. Given $\Lambda = \lambda\Sigma$, $A_{xx} = a\Sigma$ with

$$-\rho a = a^2 \frac{1}{\lambda} - \gamma, \quad (\text{E.58})$$

with positive solution

$$a = -\frac{\rho}{2}\lambda + \sqrt{\gamma\lambda + \frac{\rho^2}{4}\lambda^2}. \quad (\text{E.59})$$

In this case, (E.50) 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 (E.58).

Then, we have

$$\tau_t = \frac{a}{\lambda} [\Sigma^{-1} B (a\Phi + \gamma I)^{-1} f_t - x_t]. \quad (\text{E.60})$$

It is clear from (E.59) that $\frac{a}{\lambda}$ decreases in λ and increases in γ . ■

Proof of Proposition 9. Let's start with the complete problem:

$$\begin{aligned}
V(x_t, D_t, f_t) &= E_t \int_t^\infty e^{-\rho(s-t)} \left(x_s^\top (\alpha_s - (r + R)D_s) - \frac{\gamma}{2} x_s^\top \Sigma x_s \right) ds \\
&\quad + E_t \int_t^\infty e^{-\rho(s-t)} x_{s-}^\top C dx_s + \frac{1}{2} E_t \int_t^\infty e^{-\rho(s-t)} d[x_s, Cx_s].
\end{aligned} \tag{E.61}$$

The HJB equation is

$$\begin{aligned}
0 = \sup_{\Delta x, \mu, \sigma} &\left\{ (x + \Delta x)^\top (\alpha - (r^f + R)(D + C\Delta x)) dt - \frac{\gamma}{2} x^\top \Sigma x dt + x^\top C \mu dt + \frac{1}{2} \sigma^\top C \sigma dt \right. \\
&\quad - \rho V dt + V_x \mu dt + V_D (-R(D + C\Delta x) + C\mu) dt + V_f (-\Phi f) dt \\
&\quad + \frac{1}{2} \text{tr}(V_{xx} \sigma \sigma^\top) dt + \frac{1}{2} \text{tr}(V_{DD} C \sigma \sigma^\top C^\top) dt + \text{tr}(V_{xD} C \sigma \sigma^\top) dt \\
&\quad + \text{tr}(V_{xf} \Omega^{\frac{1}{2}} \sigma^\top) dt + \text{tr}(V_{Df} \Omega^{\frac{1}{2}} \sigma^\top C^\top) dt + \frac{1}{2} \text{tr}(V_{ff} \Omega) dt \\
&\quad \left. + x^\top C \Delta x + \frac{1}{2} \Delta x^\top C \Delta x + V(x + \Delta x, D + C\Delta x, f) - V(x, D, f) \right\}.
\end{aligned}$$

(For clarity, we mention how this equation is derived heuristically: the change between t and $t + \Delta t$ in the value function is due to a possible 'jump' trade Δx at t and the continuous trade given by a diffusion governed by μ and σ .)

Note that, under (A.18),

$$V_x = -V_D C - x^\top C \tag{E.62}$$

$$V_{xx} = C^\top V_{DD} C - C \tag{E.63}$$

$$V_{xD} = -V_{DD} C \tag{E.64}$$

$$V_{xf} = -V_{Df} C, \tag{E.65}$$

and it follows that (i) the term in Δx (last line) is identically zero; (ii) the term in μdt is identically zero; (iii) the term in $\Omega^{\frac{1}{2}} \sigma dt$ is identically zero; (iv) the term in $\sigma \sigma^\top dt$ is identically zero. The remaining terms are of order dt and, with $\hat{V}(D, f) = V(0, D, f)$, yield

the HJB equation

$$0 = \sup_x \left\{ -\rho \hat{V}(D^0, f) + \frac{\rho}{2} x^\top Cx + x^\top Bf - x^\top (r^f + R)(D + C(x - x_-)) - \frac{\gamma}{2} x^\top \Sigma x - \hat{V}_D R(D^0 + Cx) - \hat{V}_f \Phi f + \text{tr}(\hat{V}_{ff} \Omega) \right\}, \quad (\text{E.66})$$

where $x_- = x_{t-}$ and $D_- = D_{t-}$ are the state variables and $D^0 = D_- - Cx_- = D_- + C(x - x_-) - Cx$.

To provide slightly more clarity concerning (E.66), we note that the first two terms in (E.66) equal the value function decay $-\rho V(\hat{x}, \hat{D}, f)$, while the remaining terms represent the flow benefit from taking position x for the next infinitesimal time period: the expected excess return, the risk cost, the distortion decay summed with the opportunity cost of funds (the risk-free rate), from which the position x will suffer over dt , and the change over time in V induced by the decay of D and of f , as well as the convexity adjustment for f . Note that, in order for the problem to be well defined, it is necessary that $\rho < 2(r^f + R) + \gamma\Sigma$ — otherwise, the agent gains too much from pushing the prices up currently relative to the perceive cost of the risk and the decay in the distortion.

We conjecture a quadratic form for the value function \hat{V} :

$$\hat{V}(D, f) = \frac{1}{2} D^\top A_{DD} D + D^\top A_{Df} f + f^\top A_{ff} f + A_0. \quad (\text{E.67})$$

Let¹⁶

$$J = \gamma\Sigma + (2r^f - \rho + R)C + C^\top R^\top \quad (\text{E.68})$$

$$j = (B - C^\top R^\top A_{Df}) f - (C^\top R^\top A_{DD} + r^f + R) D^0. \quad (\text{E.69})$$

¹⁶Note that we choose the symmetric version of J , which makes the solution A_{DD} to (E.73) symmetric.

It follows that

$$x = J^{-1} (Bf - (r^f + R)D^0 - C^\top R^\top V_D^\top) \quad (\text{E.70})$$

$$= J^{-1}j \quad (\text{E.71})$$

and the HJB equation becomes

$$0 = \frac{1}{2}j^\top J^{-1}j - V_D R D^0 - V_f \Phi f - \rho V + \frac{1}{2}tr(V_{ff}\Omega). \quad (\text{E.72})$$

The constant matrices A_{DD} and A_{Df} are computed in the usual way:

$$\rho A_{DD} = (A_{DD}RC + r^f + R^\top) J^{-1} (C^\top R^\top A_{DD} + r^f + R) - A_{DD}R - R^\top A_{DD} \quad (\text{E.73})$$

$$\rho A_{Df} = (A_{DD}RC + r^f + R^\top) J^{-1} (-B + CR^\top A_{Df}) - R^\top A_{Df} - A_{Df}\Phi. \quad (\text{E.74})$$

Suppose now that Assumption B holds. Then we have $A_{DD} = a_{DD}\Sigma^{-1}$ and

$$\rho a_{DD} = (cRa_{DD} + r^f + R)^2 (\gamma + (2R + 2r^f - \rho)c)^{-1} - 2a_{DD}R, \quad (\text{E.75})$$

or

$$0 = a_{DD}^2 c^2 R^2 - a_{DD} ((2R + \rho)(\gamma + (2r^f + 2R - \rho)c) - 2cR(r^f + R)) + (r^f + R)^2,$$

with explicit solutions. The appropriate solution here is the lower one, which gives the economically meaningful outcome for large γ , for instance.

Maintaining Assumption B, A_{Df} can be written down easily:

$$A_{Df} = -\frac{a_{DD}RC + r^f + R}{\gamma + c(2R + 2r^f - \rho)}\Sigma^{-1}B \left(\rho - \frac{a_{DD}RC + r^f + R}{\gamma + c(2R + 2r^f - \rho)}Rc + R + \Phi \right)^{-1}. \quad (\text{E.76})$$

With

$$b = \frac{a_{DD}RC + r^f + R}{\gamma + c(2R + 2r^f - \rho)},$$

$A_{Df}f$ becomes

$$\begin{aligned} A_{Df}f &= -b\Sigma^{-1}B(\rho - bRc + R + \Phi)^{-1}f \\ &= -b\gamma \int_0^\infty e^{-(\rho - bRc + R)t} E[\text{Markowitz}_t | f_0 = f] dt. \end{aligned} \quad (\text{E.77})$$

Taken together, (E.69), (E.71), and (E.77) prove the last assertion of the proposition. ■

Proof of Proposition 10. (i) We work with the characterization of solutions provided in the proof of Proposition 6. We show that, as $\Delta_t \rightarrow 0$, Equations (E.40)–(E.42) tend to their counterparts in (E.48), which implies that the solutions also do.

To keep the proof short, we prove the claim only for (E.40); the other two equations work similarly. We first rewrite this equation as

$$-A_{xx} = \bar{\rho}S_x^\top J^{-1}S_x - \hat{\Lambda}(\Delta_t) - C + \bar{\rho}\tilde{C}^\top A_{yy}\tilde{C} \quad (\text{E.78})$$

$$= \bar{\rho}(S_x - J)^\top J^{-1}(S_x - J) - \hat{\Lambda}(\Delta_t) - \bar{\rho}J + 2\bar{\rho}S_x + \bar{\rho}\tilde{C}^\top A_{yy}\tilde{C}, \quad (\text{E.79})$$

and then rearrange it, using (E.31) and (E.35), as

$$-A_{xx}(1 - \bar{\rho}) = \bar{\rho}(S_x - J)^\top J^{-1}(S_x - J) - \bar{\rho}\gamma\Sigma\Delta_t. \quad (\text{E.80})$$

Dividing through by Δ_t and ignoring terms in Δ_t in $S_x - J$ and $J\Delta_t$, we obtain

$$-\rho A_{xx} = -\gamma\Sigma + \left(A_{xx} - A_{xy}\tilde{C} - C\right)\Lambda^{-1}\left(A_{xx} - \tilde{C}^\top A_{xy}^\top - C^\top\right), \quad (\text{E.81})$$

the same as in continuous time.

Having established that the value-function coefficients in discrete time have as limit their counterparts in continuous time, we now note that, when letting Δ_t go to 0, the rate term M^{rate} is given by

$$\Lambda^{-1}\left(A_{xx} - C - \tilde{C}A_{xy}^\top\right), \quad (\text{E.82})$$

while the aim term M^{aim} by

$$(A_{xx} - C^\top A_{xy}^\top - C^\top)^{-1} (A_{xy} + \tilde{C}^\top A_{yy}). \quad (\text{E.83})$$

These expressions are the same as obtained in continuous time.

(ii) Let's start with the case $\Lambda = 0$. Convergence follows by construction: the objective (A.16) is the limit of the discrete-time objective (19). (We omit the formal details.) The result follows for general Λ because the cost portion due to Λ ,

$$-\frac{1}{2}E_0 \left[\sum_n \Delta x_{n\Delta t}^\top \hat{\Lambda}(\Delta t) \Delta x_{n\Delta t} \right] = -\frac{1}{2}E_0 \left[\sum_n \Delta x_{n\Delta t}^\top \Lambda \Delta x_{n\Delta t} \Delta t \right], \quad (\text{E.84})$$

has limit

$$-\frac{1}{2}E_0 \left[\int \Delta x_t^\top \Lambda \Delta x_t dt \right], \quad (\text{E.85})$$

which equals zero for strategies x that jump on zero-measure sets, as is the case with the optimal policy (A.19). ■

Proof of Proposition 11. Suppose that $E_t[dp_t - r^f p_t dt] = B f_t dt$ with B given by (D.8) and apply the special case of Proposition 7 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.
- Breen, W. J., L. S. Hodrick, and R. A. Korajczyk, 2002, “Predicting Equity Liquidity,” *Management Science*, 48(4), 470–483.
- 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.
- Engle, R., R. Ferstenberg, and J. Russell, 2008, "Measuring and Modeling Execution Cost and Risk," *Working paper, University of Chicago*.
- Erb, C., and C. R. Harvey, 2006, "The Strategic and Tactical Value of Commodity Futures," *Financial Analysts Journal*, 62, 69–97.
- 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. Poteshman, 2008, "Demand-Based Option Pricing," *Review of Financial Studies (forthcoming)*.
- Greenwood, R., 2005, "Short and Long Term Demand Curves for Stocks: Theory and Evidence," *Journal of Financial Economics*, 75, 607–650.
- Grinold, R., 2006, "A Dynamic Model of Portfolio Management," *Journal of Investment Management*, 4, 5–22.
- Grossman, S., and M. Miller, 1988, "Liquidity and Market Structure," *Journal of Finance*, 43, 617–633.
- 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.

- Lehman, B. N., 1990, “Fads, Martingales, and Market Efficiency,” *Quarterly Journal of Economics*, 105, 128.
- Lillo, F., J. D. Farmer, and R. N. Mantegna, 2003, “Master Curve for Price-Impact Function,” *Nature*, 421, 129–130.
- 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., and A. MacKinlay, 1990, “When are contrarian profits due to stock market overreaction?,” *Review of financial studies*, 3(2), 175–205.
- 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*.
- Markowitz, H. M., 1952, “Portfolio Selection,” *The Journal of Finance*, 7, 77–91.
- Nagel, S., 2011, “Evaporating liquidity,” *Review of financial studies*, p. forthcoming.
- 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.

Commodity	Average Price Per Contract	Standard Deviation of Price Changes	Contract Multiplier	Daily Trading Volume (Contracts)
Aluminum	44,561	637	25	9,160
Cocoa	15,212	313	10	5,320
Coffee	38,600	1,119	37,500	5,640
Copper	80,131	2,023	25	12,300
Crude	40,490	1,103	1,000	151,160
Gasoil	34,963	852	100	37,260
Gold	43,146	621	100	98,700
Lead	23,381	748	25	2,520
Natgas	50,662	1,932	10,000	46,120
Nickel	76,530	2,525	6	1,940
Silver	36,291	893	5,000	43,780
Sugar	10,494	208	112,000	25,700
Tin	38,259	903	5	NaN
Unleaded	47,967	1,340	42,000	11,320
Zinc	36,513	964	25	6,200

Table 1: **Summary Statistics.** For each commodity used in our empirical study, the first column reports the average price per contract in U.S. dollars over our sample period 01/01/1996–01/23/2009. For instance, since the average gold price is \$431.46 per ounce, the average price per contract is \$43,146 since each contract is for 100 ounces. Each contract’s multiplier (100 in the case of gold) is reported in the third column. The second column reports the standard deviation of price changes. The fourth column reports the average daily trading volume per contract, estimated as the average daily volume of the most liquid contract traded electronically and outright (i.e., not including calendar-spread trades) in December 2010.

	Panel A: Benchmark Transaction Costs		Panel B: High Transaction Costs	
	Gross SR	Net SR	Gross SR	Net SR
Markowitz	0.83	-9.84	0.83	-10.11
Dynamic optimization	0.62	0.58	0.58	0.53
Static optimization				
Weight on Markowitz = 10%	0.63	-0.41	0.63	-1.45
Weight on Markowitz = 9%	0.62	-0.24	0.62	-1.10
Weight on Markowitz = 8%	0.62	-0.08	0.62	-0.78
Weight on Markowitz = 7%	0.62	0.07	0.62	-0.49
Weight on Markowitz = 6%	0.62	0.20	0.62	-0.22
Weight on Markowitz = 5%	0.61	0.31	0.61	0.00
Weight on Markowitz = 4%	0.60	0.40	0.60	0.19
Weight on Markowitz = 3%	0.58	0.46	0.58	0.33
Weight on Markowitz = 2%	0.52	0.46	0.52	0.39
Weight on Markowitz = 1%	0.36	0.33	0.36	0.31

Table 2: **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: Construction of Current Optimal Trade




Figure 1: **Optimal Trading Strategy: Aim in Front of the Target.** This figure shows how the optimal trade moves the portfolio from the existing position x_{t-1} towards the aim portfolio, trading only part of the way to the aim portfolio to limit transactions costs. The aim portfolio is an average of the current Markowitz portfolio (the optimal portfolio in the absence of transaction costs) and the expected future aim portfolio, which reflects how the Markowitz portfolio is expected to mean-revert to its long-term level (the lower, left corner of the figure).




Figure 2: **Aim in Front of the Target by Underweighting Fast-Decay Factors.** This figure shows how the optimal trade moves from the existing position x_{t-1} towards the aim, which puts relatively more weight on assets loading on persistent factors. Relative to the Markowitz portfolio, the weight in the aim of asset 2 is lower because asset 2 has a faster-decaying alpha, as is apparent in the lower expected weight it receives in future Markowitz portfolios.

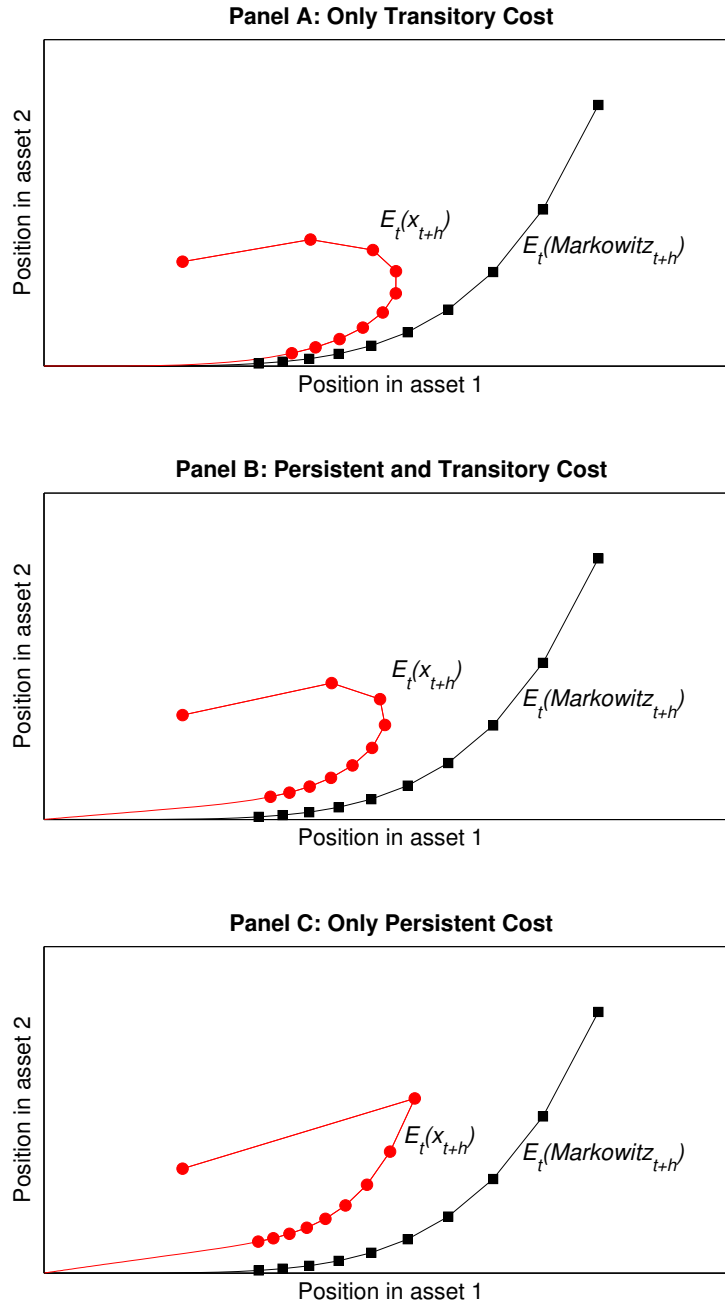


Figure 3: **Aim in Front of the Target with Persistent Costs.** This figure shows the optimal trade when part of the transaction cost is persistent. In panel A, the entire cost is transitory, as in Figures 1 and 2. In panel B, half of the cost is transitory, while the other half is persistent, with a half life of 6.9 periods. In panel C, the entire cost is persistent.




Figure 4: **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 (“Markowitz”) and our optimal dynamic strategy (“optimal”).

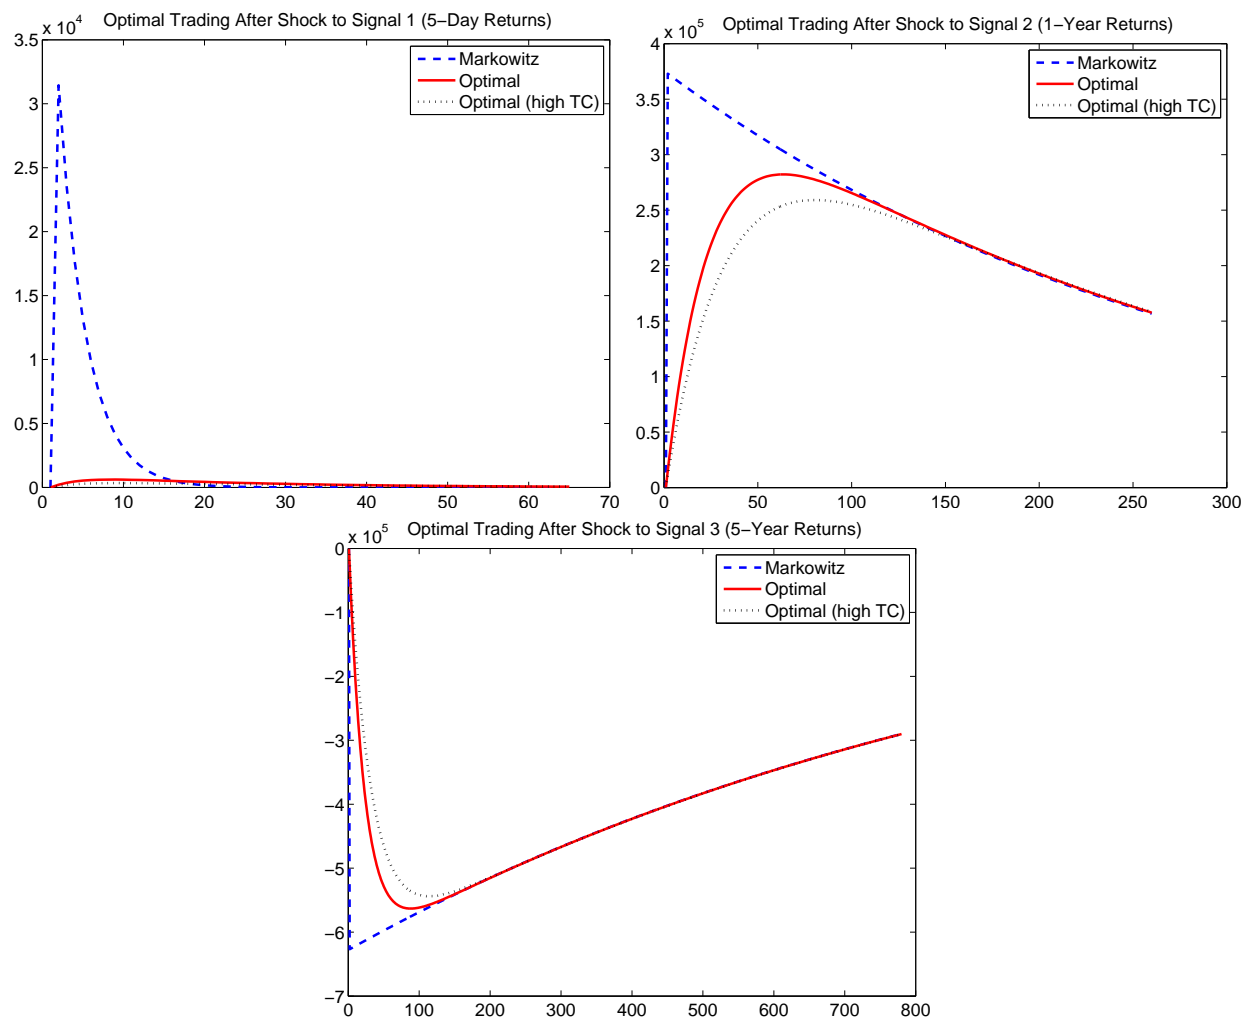


Figure 5: **Optimal Trading in Response to Shock to Return Predicting Signals.** This figure shows the response in the optimal position following a shock to a return predictor as a function of the number of days since the shock. The top left panel does this for a shock to the fast 5-day return predictor, 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 (“Markowitz”) and our optimal dynamic strategy (“optimal”) using high and low transactions costs.