

Dynamic Trading with Predictable Returns and Transaction Costs

Gârleanu, Nicolae; Heje Pedersen, Lasse

Document Version
Accepted author manuscript

Published in:
Journal of Finance

DOI:
[10.1111/jofi.12080](https://doi.org/10.1111/jofi.12080)

Publication date:
2013

License
Unspecified

Citation for published version (APA):
Gârleanu, N., & Heje Pedersen, L. (2013). Dynamic Trading with Predictable Returns and Transaction Costs. *Journal of Finance*, 68(6), 2309–2340. <https://doi.org/10.1111/jofi.12080>

[Link to publication in CBS Research Portal](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

If you believe that this document breaches copyright please contact us (research.lib@cbs.dk) providing details, and we will remove access to the work immediately and investigate your claim.

Download date: 14. Jul. 2025



Dynamic Trading with Predictable Returns and Transaction Costs

Nicolae Gârleanu and Lasse Heje Pedersen

Journal article (Accepted manuscript*)

Please cite this article as:

Gârleanu, N., & Heje Pedersen, L. (2013). Dynamic Trading with Predictable Returns and Transaction Costs. *Journal of Finance*, 68(6), 2309-2340. <https://doi.org/10.1111/jofi.12080>

This is the peer reviewed version of the article, which has been published in final form at DOI:
<https://doi.org/10.1111/jofi.12080>

This article may be used for non-commercial purposes in accordance with
[Wiley Terms and Conditions for Self-Archiving](#)

* This version of the article has been accepted for publication and undergone full peer review but has not been through the copyediting, typesetting, pagination and proofreading process, which may lead to differences between this version and the publisher's final version AKA Version of Record.

Uploaded to [CBS Research Portal](#): March 2020

Dynamic Trading with Predictable Returns and Transaction Costs

NICOLAE GÂRLEANU and LASSE HEJE PEDERSEN*

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.

Keywords: dynamic trading, predictability, transaction costs, portfolio choice

JEL Classification: G11, G12

*Gârleanu is at Haas School of Business, University of California, Berkeley, NBER, and CEPR, and Pedersen is at New York University, Copenhagen Business School (FRIC Center for Financial Frictions), AQR Capital Management, NBER, and CEPR. 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), Mikkel Heje Pedersen, 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, University of Copenhagen, Rice University, University of Michigan Ross School, Yale University School of Management, the Bank of Canada, and the Journal of Investment Management Conference. Pedersen gratefully acknowledges support from the European Research Council (ERC grant no. 312417).

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, for example, value and momentum predictors, and these have different prediction strengths and mean-reversion speeds — put differently, different “alphas” and “alpha decays.” The alpha decay is important because it determines how long the investor can enjoy high expected returns, and therefore affects the trade-off between returns and transaction 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, securities’ risks and return 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.” 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 trading and, interestingly, to modify the aim, and thus not to trade 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).

Panel A of Figure 1 illustrates the construction of the optimal portfolio of two securities.¹ 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 (for example, 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. The optimal new position is achieved by trading partially towards this aim portfolio. Another [Figure 1] 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).

While new to finance, these portfolio principles have close analogues in other fields such as the guidance of missiles towards moving targets, hunting, and sports. The most famous example from the sports world is perhaps the following quote, illustrated in Panel D of Figure 1:

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

Similarly, hunters are reminded to “lead the duck” when aiming their weapon as seen in Panel E.²

Panel B of Figure 1 illustrates the expected trade at the next trading date, and Panel C shows how the optimal position is expected to chase the Markowitz portfolio over time. The expected path of the optimal portfolio resembles that of a guided missile chasing an enemy airplane in so-called “lead homing” systems, as seen in Panel F.

The optimal portfolio is forward-looking and depends critically on each return predictor’s mean-reversion speed (alpha decay). To see this in Figure 1, note the convex J-shape of the expected path of the Markowitz portfolio: The Markowitz position in security 1 decays more slowly than that in security 2, as the predictor that currently “likes” security 1 is more persistent. Therefore, the aim portfolio loads more heavily on 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 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 new implication of our model. It arises because transaction costs imply that the investor cannot easily change his portfolio and therefore must consider his optimal portfolio both now and in the future. In contrast, absent transaction costs, the investor can reoptimize at no cost and needs to consider only current investment opportunities without regard to alpha decay.

Our specification of transaction costs is sufficiently rich to allow for both purely transitory and persistent costs. With persistent transaction costs, the price changes due to the trader’s market impact persist for a while. 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.

Finally, we illustrate our results empirically in the context of commodity futures markets. We use returns over the past five days, 12 months, and five years to predict returns. The five-day signal is quickly mean-reverting (fast alpha decay), the 12-month signal mean-reverts more slowly, whereas the five-year signal is the most persistent. We calculate the optimal dynamic trading strategy taking transaction costs into account and compare its performance to both the optimal portfolio ignoring transaction costs and 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 an 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 I describes how our paper contributes to the portfolio selection literature that starts 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 II lays out the model with temporary transaction costs and the solution method. Section III shows the optimality of aiming in front of the target and trading partially towards the aim. Section IV solves the model with persistent transaction costs. Section V provides a number of theoretical applications, while Section VI applies our framework empirically to trading commodity futures. Section VII concludes. All proofs are in the appendix.

I. Related Literature

A large literature studies portfolio selection with return predictability in the absence of trading costs (see, for example, Campbell and Viceira (2002) and references therein). Alpha decay plays no role in this literature, 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 transaction cost implications for the equity premium.³ Equilibrium models with trading costs include Amihud and Mendelson (1986), Vayanos (1998), Vayanos and Vila (1999), Lo, Mamaysky, and Wang (2004), and Gârleanu (2009), as well as 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-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 the size of transaction costs the trader is willing to accept (that is, 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 sets by Balduzzi and Lynch (1999), Lynch and Balduzzi (2000), Jang et al. (2007), and Lynch and Tan (2011). 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, for example, Perold (1988), Bertsimas and Lo (1998), Almgren and Chriss (2000), Obizhaeva and Wang (2006), and Engle and Ferstenberg (2007)).

Finally, quadratic programming techniques are also used in macroeconomics and other fields, and usually the solution comes down to algebraic matrix Riccati equations (see, for example, Ljungqvist and Sargent (2004) and references therein). We solve our model explicitly, including the Riccati equations.

II. 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,⁴ 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 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.⁵ 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, $\frac{1}{2} \Lambda$, for convenience) is a multidimensional version of Kyle’s lambda, which can also be justified by inventory considerations (for example, Grossman and Miller (1988) or Greenwood (2005) for the multiasset case). While this transaction-cost specification is chosen partly for tractability, the empirical literature generally finds transaction costs to be convex (for example, Engle, Ferstenberg, and Russell (2008), Lillo, Farmer, and Mantegna (2003)), with some researchers actually estimating quadratic trading costs (for example, 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 1. 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 (2009). 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 rate and γ is the risk aversion coefficient.⁶

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} , and A_{ff} are stated explicitly in (A.15), (A.18), and (A.22), and A_{xx} is positive definite.*⁷

III. 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,

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

(ii) *Under Assumption 1, 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 a fixed fraction towards the aim (that is, 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. Put 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. A larger absolute risk aversion can also be viewed as a smaller investor, for whom transaction costs play a smaller role and who therefore trades closer to her aim.

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),

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

As is well known, the Markowitz portfolio is the tangency portfolio appropriately leveraged depending on the risk aversion γ .

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 1, this can be written as follows, letting $z = \gamma/(\gamma + a)$:*

$$aim_t = z Markowitz_t + (1 - z) E_t(aim_{t+1}). \quad (12)$$

(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 1,

$$aim_t = \sum_{\tau=t}^{\infty} z(1 - z)^{\tau-t} E_t(Markowitz_{\tau}). \quad (13)$$

The weight z of the current Markowitz portfolio decreases with the 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 (expected) Markowitz portfolios, where the weight on the current (and near-term) Markowitz portfolio is larger if transaction costs are smaller.

The optimal trading policy is illustrated in detail in Figure 1 (as discussed briefly in the introduction). Since expected returns mean-revert, the expected Markowitz portfolio

converges to its long-term mean, illustrated at the origin of the figure. We see that the aim portfolio is a weighted average of the current and future Markowitz portfolios (that is, the aim portfolio lies in the convex cone of the solid curve). 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 asset 1. The optimal new position is found by moving partially towards the aim portfolio as seen in the figure.

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. Naturally, a more persistent factor has a larger effect on future Markowitz portfolios than a factor 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 1, the aim portfolio is the Markowitz portfolio built as if the signals f were scaled down based on 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 :

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

(iii) A persistent factor i is scaled down less than a fast factor j , and the relative weight of i compared to that of j increases in the transaction cost, that is, $(1 + \phi^j a / \gamma) / (1 + \phi^i a / \gamma)$ increases in λ .

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, $(\gamma\Sigma)^{-1}Bf_t$. Hence, transaction costs imply that one optimally trades only part of the way towards the aim, and that the aim downweights each return-predicting factor more the higher is its alpha decay ϕ^k . Downweighting 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 1. The convexity 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. Put differently, 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 exponen-*

tially weighted average of past aim portfolios. Under Assumption 1,

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

We see that the optimal portfolio is an exponentially weighted average of current and past aim portfolios. Clearly, the history of past expected returns affects the current position, since the investor trades patiently to economize on transaction costs. One reading of the proposition is that the investor computes the exponentially weighted average of past aim portfolios and always trades *all the way* to this portfolio (assuming that his initial portfolio is right, otherwise the first trade is suboptimal).

IV. Persistent Transaction Costs

In some cases the impact of trading on prices may have a nonnegligible *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 large relative to the 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. The following proposition characterizes the optimal portfolio strategy.⁸

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, $aim_t = M^{aim}y_t$. The aim portfolio 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 given 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 2 illustrates graphically the optimal trading strategy with temporary and persistent price impacts. Panel A uses the parameters from Figure 1, Panel B has both temporary and persistent transaction costs, while Panel C has a purely persistent price impact.⁹ Specifically, suppose that Kyle's lambda for the temporary price impact is $\Lambda = w\tilde{\Lambda}$ and Kyle's lambda for the persistent price impact is $C = (1 - w)\tilde{\Lambda}$, where we vary w to determine how much of the total price impact is temporary versus 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).

We see that the optimal portfolio policy with persistent transaction costs also tracks the Markowitz portfolio while aiming in front of the target. It can be shown more generally that the optimal portfolio under a persistent price impact depends on the expected future Markowitz portfolios (that is, aims in front of the target). This is similar to the case of a temporary price impact, but what is different with a 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 is more willing to incur it early to start collecting the benefits of a better portfolio. (The resilience still makes it cheaper to postpone part of the trade, however). Furthermore, the cost of buying a position and selling it shortly thereafter is much smaller with a persistent price impact.

V. Theoretical Applications

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

Example 1: Timing a Single Security

A 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 1 ($\Lambda = \lambda\Sigma$) 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 (in contrast, in the general model above, all of the factors could influence all of the securities). For instance, one can imagine that each security is associated with a value characteristic (for example, 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 (for example, IBM's book-to-market) and β^i is the predictive ability of characteristic i (that is, 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 $f_t^i = (f_t^{i,1}, \dots, f_t^{i,S})^\top$, for example, 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 express B using the Kronecker product:

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

Thus, $E_t(r_{t+1}) = Bf_t$. 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

Consider an investor who performs a static optimization involving current expected returns, risk, and transaction costs. Such an investor simply solves

$$\max_{x_t} x_t^\top E_t(r_{t+1}) - \frac{\gamma}{2} x_t^\top \Sigma x_t - \frac{\lambda}{2} \Delta x_t^\top \Sigma \Delta x_t, \quad (28)$$

with 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. The first shortcoming of the static portfolio (point (i)), namely that 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, the second shortcoming (point (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 VI 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 only γ and λ , but also the expected returns $E_t(r_{t+1}) = Bf_t$ 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:

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

Put 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

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

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.

VI. 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.

A. 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 I, we [Table I]

collect data on Aluminum, Copper, Nickel, Zinc, Lead, and Tin from the London Metal Exchange (LME), Gasoil from the Intercontinental Exchange (ICE), WTI Crude, RBOB Unleaded Gasoline, and Natural Gas from the New York Mercantile Exchange (NYMEX), Gold and Silver from the New York Commodities Exchange (COMEX), and Coffee, Cocoa, and Sugar from the New York Board of Trade (NYBOT). (Note that we exclude futures on various agriculture and livestock that have tight price limits.) We consider the sample period 01/01/1996 to 01/23/2009, for which we have data on all the above commodities.¹⁰

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 I provides summary statistics on each contract’s average price, the standard deviation of price changes, the contract multiplier (for example, 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 contract nears maturity and the next contract becomes more liquid, the trader “rolls” into the next contract, that is, 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 in this fashion. 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 the 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 I. Specifically, we obtain an estimate of the average daily volume of the most liquid contract traded electronically and outright (that is, not including calendar-spread trades) in December 2010 from an asset manager based on underlying data from Reuters.

B. Predicting Returns and Other Parameter Estimates

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

$$r_{t+1}^s = \underset{(0.17)}{0.001} + \underset{(2.22)}{10.32} f_t^{5D,s} + \underset{(2.82)}{122.34} f_t^{1Y,s} - \underset{(-1.79)}{205.59} f_t^{5Y,s} + u_{t+1}^s, \quad (31)$$

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 change 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; to avoid dividing by a number close to zero, the standard deviations are winsorized below the 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 reversal over long horizons.¹¹ 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 five-day signal, a 206-day half-life for the one-year signal, and a 700-day half-life for the five-year signal.¹²

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 one for an agent with \$1 billion under management. We set the time discount rate to $\rho = 1 - \exp(-0.02/260)$, corresponding to a 2% 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) provide a similar estimate.) Further, Greenwood (2005) finds evidence that a 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 is 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 I 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 (for example, the standard deviation is 1.0×10^{-6}), reflecting the fact 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 (for example, by using these calibrated λ 's in the diagonal of the Λ matrix), we also need to estimate the spillover effects (that is, the off-diagonal elements). Since Greenwood (2005) provides the only estimate of these transaction cost spillovers 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 theoretically.

We also consider a more conservative transaction cost estimate of $\lambda = 10 \times 10^{-7}$. This more conservative analysis can be interpreted as providing the trading strategy of a larger investor (that is, we could equivalently reduce the absolute risk aversion γ).

C. 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 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 an aim portfolio that depends on signals’ alpha decays), and we consider 10 different trading speeds as seen in Table II. Hence, under the static optimization, the updated portfolio is a [Table II] weighted average of the Markowitz portfolio (with weight denoted “weight on Markowitz”) and the current portfolio.

Table II 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 Sharpe ratio before transaction costs is achieved by the Markowitz strategy. The optimal and static portfolios have similar drops in gross Sharpe ratio 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 outperformance is that the dynamic strategy gives less weight to the five-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 five-day signal), or it trades so slowly that 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 3 depicts 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 larger when the expected return is large, but moderates the speed and magnitude of trades. [Figure 3]

D. 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 4 shows the responses to shocks to each return-predicting factor, namely, the five-day factor, the one-year factor, and the five-year factor. [Figure 4]

The first panel shows that the Markowitz strategy immediately jumps up after a shock to the five-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 one-year factor. The Markowitz strategy jumps up and decays, whereas the optimal position increases more smoothly and catches up as the Markowitz strategy starts to decay. The third panel shows the same for the five-year signal, except that the effects are slower and with opposite sign, since five-year returns predict future reversals.

VII. 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 trade-off between risk and return, but is 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 excessively 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.

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, that is, 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. Appendix: Proofs

In what follows we make repeated use of the notation

$$\bar{\rho} = 1 - \rho \tag{A.1}$$

$$\bar{\Lambda} = \bar{\rho}^{-1} \Lambda \tag{A.2}$$

$$\bar{\lambda} = \bar{\rho}^{-1} \lambda. \tag{A.3}$$

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

$$\begin{aligned} E_t[V(x_t, f_{t+1})] &= -\frac{1}{2}x_t^\top A_{xx}x_t + x_t^\top A_{xf}(I - \Phi)f_t + \frac{1}{2}f_t^\top (I - \Phi)^\top A_{ff}(I - \Phi)f_t \\ &\quad + \frac{1}{2}E_t(\varepsilon_{t+1}^\top A_{ff}\varepsilon_{t+1}) + A_0. \end{aligned} \tag{A.4}$$

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

$$\begin{aligned} J_t &= \gamma \Sigma + \bar{\Lambda} + A_{xx} \\ j_t &= (B + A_{xf}(I - \Phi))f_t + \bar{\Lambda}x_{t-1} \\ d_t &= -\frac{1}{2}x_{t-1}^\top \bar{\Lambda}x_{t-1} + \frac{1}{2}f_t^\top (I - \Phi)^\top A_{ff}(I - \Phi)f_t + \frac{1}{2}E_t(\varepsilon_{t+1}^\top A_{ff}\varepsilon_{t+1}) + A_0. \end{aligned} \tag{A.5}$$

The maximum value is attained by

$$x_t = J_t^{-1}j_t, \tag{A.6}$$

which 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:¹³

$$-\bar{\rho}^{-1}A_{xx} = \bar{\Lambda}(\gamma\Sigma + \bar{\Lambda} + A_{xx})^{-1}\bar{\Lambda} - \bar{\Lambda} \quad (\text{A.7})$$

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

$$\begin{aligned} \bar{\rho}^{-1}A_{ff} &= (B + A_{xf}(I - \Phi))^\top(\gamma\Sigma + \bar{\Lambda} + A_{xx})^{-1}(B + A_{xf}(I - \Phi)) \\ &\quad + (I - \Phi)^\top A_{ff}(I - \Phi). \end{aligned} \quad (\text{A.9})$$

The existence of a solution to this system of Riccati equations can be established using standard results, for example, 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 (A.7) as

$$\bar{\rho}^{-1}Z = I - (\gamma M + I + Z)^{-1}, \quad (\text{A.10})$$

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

$$Z^2 + Z(I + \gamma M - \bar{\rho}I) = \bar{\rho}\gamma M \quad (\text{A.11})$$

$$\left(Z + \frac{1}{2}(\gamma M + \rho I)\right)^2 = \bar{\rho}\gamma M + \frac{1}{4}(\gamma M + \rho I)^2, \quad (\text{A.12})$$

resulting in

$$Z = \left(\bar{\rho}\gamma M + \frac{1}{4}(\rho I + \gamma M)^2\right)^{\frac{1}{2}} - \frac{1}{2}(\rho I + \gamma M) \quad (\text{A.13})$$

$$A_{xx} = \bar{\Lambda}^{\frac{1}{2}} \left[\left(\bar{\rho}\gamma M + \frac{1}{4}(\rho I + \gamma M)^2\right)^{\frac{1}{2}} - \frac{1}{2}(\rho I + \gamma M) \right] \bar{\Lambda}^{\frac{1}{2}}, \quad (\text{A.14})$$

that is,

$$A_{xx} = \left(\bar{\rho}\gamma\bar{\Lambda}^{\frac{1}{2}}\Sigma\bar{\Lambda}^{\frac{1}{2}} + \frac{1}{4}(\rho^2\bar{\Lambda}^2 + 2\rho\gamma\bar{\Lambda}^{\frac{1}{2}}\Sigma\bar{\Lambda}^{\frac{1}{2}} + \gamma^2\bar{\Lambda}^{\frac{1}{2}}\Sigma\bar{\Lambda}^{-1}\Sigma\bar{\Lambda}^{\frac{1}{2}}) \right)^{\frac{1}{2}} - \frac{1}{2}(\rho\bar{\Lambda} + \gamma\Sigma). \quad (\text{A.15})$$

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 (A.8). To write the solution explicitly, we note first that, from (A.7),

$$\bar{\Lambda}(\gamma\Sigma + \bar{\Lambda} + A_{xx})^{-1} = I - A_{xx}\Lambda^{-1}. \quad (\text{A.16})$$

Using the general rule that $\text{vec}(XYZ) = (Z^\top \otimes X)\text{vec}(Y)$, we rewrite (A.8) in vectorized form:

$$\text{vec}(A_{xf}) = \bar{\rho} \text{vec}((I - A_{xx}\Lambda^{-1})B) + \bar{\rho}((I - \Phi)^\top \otimes (I - A_{xx}\Lambda^{-1})) \text{vec}(A_{xf}), \quad (\text{A.17})$$

so that

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

Finally, A_{ff} is calculated from the linear equation (A.9), which is of the form

$$\bar{\rho}^{-1}A_{ff} = Q + (I - \Phi)^\top A_{ff}(I - \Phi) \quad (\text{A.19})$$

with

$$Q = (B + A_{xf}(I - \Phi))^\top (\gamma\Sigma + \bar{\Lambda} + A_{xx})^{-1} (B + A_{xf}(I - \Phi)) \quad (\text{A.20})$$

a positive-definite matrix.

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

$$A_{ff,ij} = \frac{\bar{\rho}Q_{ij}}{1 - \bar{\rho}(1 - \Phi_{ii})(1 - \Phi_{jj})}. \quad (\text{A.21})$$

In general,

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

One way to see that A_{ff} is positive-definite is to iterate (A.19) 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 $\lambda > 0$, the solution to the value function coefficients is $A_{xx} = a\Sigma$, where a solves a simplified version of (A.7):

$$-\bar{\rho}^{-1}a = \frac{\bar{\lambda}^2}{\gamma + \bar{\lambda} + a} - \bar{\lambda}, \quad (\text{A.23})$$

or

$$a^2 + (\gamma + \bar{\lambda}\rho)a - \lambda\gamma = 0, \quad (\text{A.24})$$

with solution

$$a = \frac{\sqrt{(\gamma + \bar{\lambda}\rho)^2 + 4\gamma\lambda} - (\gamma + \bar{\lambda}\rho)}{2}. \quad (\text{A.25})$$

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

Note that a is symmetric in $(\lambda\rho(1-\rho)^{-1}, \gamma)$. Consequently, a increases in λ if and only if it increases in γ . Differentiating (A.25) with respect to λ , one gets

$$2\frac{da}{d\lambda} = -\bar{\rho}^{-1}\rho + \frac{1}{2} \frac{(2(\gamma + \bar{\lambda}\rho) + 4\gamma)}{\sqrt{(\gamma + \bar{\lambda}\rho)^2 + 4\gamma\lambda}}. \quad (\text{A.26})$$

This expression is positive if and only if

$$\bar{\rho}^{-2}\rho^2((\gamma + \bar{\lambda}\rho)^2 + 4\gamma\lambda) \leq ((\gamma + \bar{\lambda}\rho)\bar{\rho}^{-1}\rho + 2\gamma)^2, \quad (\text{A.27})$$

which is verified to hold with strict inequality as long as $\bar{\rho}\gamma > 0$.

Finally, note that a/λ is increasing in γ and homogeneous of degree zero in (λ, γ) , so that applying Euler's theorem for homogeneous functions gives

$$\frac{d}{d\lambda} \frac{a}{\lambda} = -\frac{d}{d\gamma} \frac{a}{\lambda} < 0. \quad (\text{A.28})$$

Proof of Proposition 3. We show that

$$aim_t = (\gamma\Sigma + A_{xx})^{-1} (\gamma\Sigma \times Markowitz_t + A_{xx} \times E_t(aim_{t+1})) \quad (\text{A.29})$$

by using (8), (A.8), and (A.7) successively to write

$$\begin{aligned} aim_t &= A_{xx}^{-1} A_{xf} f_t \\ &= A_{xx}^{-1} \Lambda (\gamma\Sigma + \bar{\Lambda} + A_{xx})^{-1} (\gamma\Sigma \times Markowitz_t + A_{xx} \times E_t(aim_{t+1})) \\ &= (\gamma\Sigma + A_{xx})^{-1} (\gamma\Sigma \times Markowitz_t + A_{xx} \times E_t(aim_{t+1})). \end{aligned} \quad (\text{A.30})$$

To obtain the last equality, rewrite (A.7) as

$$(\Lambda - A_{xx}) \Lambda^{-1} (\gamma\Sigma + \bar{\Lambda} + A_{xx}) = \bar{\Lambda} \quad (\text{A.31})$$

and then further

$$\gamma\Sigma + A_{xx} = (\gamma\Sigma + \bar{\Lambda} + A_{xx}) \Lambda^{-1} A_{xx}, \quad (\text{A.32})$$

since $A_{xx} \Lambda^{-1} \Sigma = \Sigma \Lambda^{-1} A_{xx}$. equation (12) follows immediately as a special case.

For part (ii), we iterate (A.29) 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). Given that a increases in λ , z decreases in λ . Furthermore, z increases in γ if and only if a/γ decreases, which is equivalent (by symmetry) to a/λ decreasing in λ .

Proof of Proposition 4. In the case $\Lambda = \lambda\Sigma$, equation (A.8) is solved by

$$\begin{aligned} A_{xf} &= \lambda B((\gamma + \bar{\lambda} + a)I - \lambda(I - \Phi))^{-1} \\ &= \lambda B((\gamma + \bar{\lambda}\rho + a)I + \lambda\Phi)^{-1} \\ &= B\left(\frac{\gamma}{a} + \Phi\right)^{-1}, \end{aligned} \tag{A.33}$$

where the last equality uses (A.23). The aim portfolio is

$$aim_t = (a\Sigma)^{-1} B\left(\frac{\gamma}{a} + \Phi\right)^{-1} f_t, \tag{A.34}$$

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

For part (iii), we use the result shown above (proof of Proposition 2) that a increases in λ , which implies that $(1 + \phi^j a/\gamma)/(1 + \phi^i a/\gamma)$ does whenever $\phi^j > \phi^i$.

Proof of Proposition 5. Rewriting (7) as

$$x_t = (I - \Lambda^{-1}A_{xx})x_{t-1} + \Lambda^{-1}A_{xx} \times aim_t \tag{A.35}$$

and iterating this relation backwards gives

$$x_t = \sum_{\tau=-\infty}^t (I - \Lambda^{-1}A_{xx})^{t-\tau} \Lambda^{-1}A_{xx} \times aim_\tau. \tag{A.36}$$

Proof of Proposition 6. We start by defining

$$\begin{aligned}\Pi &= \begin{bmatrix} \Phi & 0 \\ 0 & R \end{bmatrix}, & \tilde{C} &= (1-R) \begin{bmatrix} 0 \\ C \end{bmatrix}, \\ \tilde{B} &= \begin{bmatrix} B & -(R+r^f) \end{bmatrix}, \\ \tilde{\Omega} &= \begin{bmatrix} \Omega & 0 \\ 0 & 0 \end{bmatrix}, & \tilde{\varepsilon}_t &= \begin{bmatrix} \varepsilon_t \\ 0 \end{bmatrix}.\end{aligned}\tag{A.37}$$

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

$$E_t[y_{t+1}] = (I - \Pi)y_t + \tilde{C}(x_t - x_{t-1}).\tag{A.38}$$

The conjectured value function is

$$V(x_{t-1}, y_t) = -\frac{1}{2}x_{t-1}^\top A_{xx}x_{t-1} + x_{t-1}^\top A_{xy}y_t + \frac{1}{2}y_t^\top A_{yy}y_t + A_0,\tag{A.39}$$

so that

$$\begin{aligned}E_t[V(x_t, y_{t+1})] &= -\frac{1}{2}x_t^\top A_{xx}x_t + x_t^\top A_{xy} \left((I - \Pi)y_t + \tilde{C}(x_t - x_{t-1}) \right) + \\ &\quad \frac{1}{2} \left((I - \Pi)y_t + \tilde{C}(x_t - x_{t-1}) \right)^\top A_{yy} \left((I - \Pi)y_t + \tilde{C}(x_t - x_{t-1}) \right) + \\ &\quad \frac{1}{2}E_t[\tilde{\varepsilon}_{t+1}^\top A_{yy}\tilde{\varepsilon}_{t+1}] + A_0.\end{aligned}\tag{A.40}$$

The trader consequently chooses x_t to solve

$$\begin{aligned}\max_x &\left\{ x^\top \tilde{B}y_t - x^\top (R + r^f)C(x - x_{t-1}) - \frac{\gamma}{2}x^\top \Sigma x \right. \\ &\quad + \frac{1}{2}\bar{\rho}^{-1} \left(x^\top Cx - x_{t-1}^\top Cx_{t-1} - (x - x_{t-1})^\top \Lambda(x - x_{t-1}) \right) \\ &\quad - \frac{1}{2}x^\top A_{xx}x + x^\top A_{xy} \left((I - \Pi)y_t + \tilde{C}(x - x_{t-1}) \right) \\ &\quad \left. + \frac{1}{2} \left((I - \Pi)y_t + \tilde{C}(x - x_{t-1}) \right)^\top A_{yy} \left((I - \Pi)y_t + \tilde{C}(x - x_{t-1}) \right) \right\},\end{aligned}\tag{A.41}$$

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

$$J = \frac{1}{2}(J_0 + J_0^\top) \quad (\text{A.42})$$

$$J_0 = \gamma\Sigma + \bar{\Lambda} + (2(R + r^f) - \bar{\rho}^{-1})C + A_{xx} - 2A_{xy}\tilde{C} - \tilde{C}^\top A_{yy}\tilde{C} \quad (\text{A.43})$$

$$j_t = \tilde{B}y_t + (\bar{\Lambda} + (R + r^f)C)x_{t-1} + A_{xy}\left((I - \Pi)y_t - \tilde{C}x_{t-1}\right) + \tilde{C}A_{yy}\left((I - \Pi)y_t - \tilde{C}x_{t-1}\right) \quad (\text{A.44})$$

$$\equiv S_x x_{t-1} + S_y y_t \quad (\text{A.45})$$

$$d_t = -\frac{1}{2}x_{t-1}^\top \bar{\Lambda} x_{t-1} - \frac{1}{2}\bar{\rho}^{-1}x_{t-1}^\top C x_{t-1} + \frac{1}{2}\left((I - \Pi)y_t - \tilde{C}x_{t-1}\right)^\top A_{yy}\left((I - \Pi)y_t - \tilde{C}x_{t-1}\right). \quad (\text{A.46})$$

Here,

$$S_x = \bar{\Lambda} + (R + r^f)C - A_{xy}\tilde{C} - \tilde{C}^\top A_{yy}\tilde{C} \quad (\text{A.47})$$

$$S_y = \tilde{B} + A_{xy}(I - \Pi) + \tilde{C}^\top A_{yy}(I - \Pi). \quad (\text{A.48})$$

The value of x attaining the maximum is given by

$$x_t = J^{-1}j_t, \quad (\text{A.49})$$

and the maximal value is

$$\frac{1}{2}j_t^\top J^{-1}j_t + d_t = V(x_{t-1}, y_t) - A_0 \quad (\text{A.50})$$

$$= -\frac{1}{2}x_{t-1}^\top A_{xx}x_{t-1} + x_{t-1}^\top A_{xy}y_t + \frac{1}{2}y_t^\top A_{yy}y_t. \quad (\text{A.51})$$

The unknown matrices have to satisfy a system of equations encoding the equality of all coefficients in (A.51). Thus,

$$-\bar{\rho}^{-1}A_{xx} = S_x^\top J^{-1}S_x - \bar{\Lambda} - \bar{\rho}^{-1}C + \tilde{C}^\top A_{yy}\tilde{C} \quad (\text{A.52})$$

$$\bar{\rho}^{-1}A_{xy} = S_x^\top J^{-1}S_y - \tilde{C}^\top A_{yy}(I - \Pi) \quad (\text{A.53})$$

$$\bar{\rho}^{-1}A_{yy} = S_y^\top J^{-1}S_y + (I - \Pi)^\top A_{yy}(I - \Pi). \quad (\text{A.54})$$

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

$$x_t = x_{t-1} + \underbrace{(I - J^{-1}S_x)}_{M^{rate}} \left(\underbrace{(I - J^{-1}S_x)^{-1} (J^{-1}S_y) y_t}_{aim_t = M^{aim} y_t} - x_{t-1} \right). \quad (\text{A.55})$$

Notes

¹Panel A-C of the figure are based on simulations of our model. We are grateful to Mikkel Heje Pedersen for Panels D–F. Panel F is based on “Introduction to Rocket and Guided Missile Fire Control,” Historic Naval Ships Association (2007).

²We thank Kerry Back for this analogy.

³Davis and Norman (1990) provide a more formal analysis of Constantinides’ model. Also, Gârleanu (2009) and Lagos and Rocheteau (2009) show how search frictions and payoff mean-reversion impact how close one trades to the static portfolio. Our model also 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).

⁴The unconditional mean excess returns are also captured in the factors f . For example, 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, $\varepsilon_t^1 = 0$.)

⁵The 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 $\bar{\Lambda}$, that is, $\Lambda = (\bar{\Lambda} + \bar{\Lambda}^\top)/2$, generates the same trading costs as $\bar{\Lambda}$.

⁶Put differently, the investor has mean-variance preferences over the change in his wealth W_t each time period, net of the risk-free return: $\Delta W_{t+1} - r^f W_t = x_t^\top r_{t+1} - TC_{t+1}$.

⁷Note that A_{xx} and A_{ff} can always be chosen to be symmetric.

⁸We assume that the objective (19) is concave and a non-explosive solution exists. A sufficient condition is that γ is large enough.

⁹The parameters used in Panel A of Figure 2, and Panels A–C of Figure 1, are $f_0 = (1, 1)^\top$, $B = I_{2 \times 2}$, $\phi_1 = 0.1$, $\phi_2 = 0.4$, $\Sigma = I_{2 \times 2}$, $\gamma = 0.5$, $\rho = 0.05$, and $\Lambda = 2\Sigma$. The additional parameters for Panels B–C of Figure 2 are $D_0 = 0$, $R = 0.1$, and the risk-free rate given by $(1 + r^f)(1 - \rho) = 1$. As further interpretation of Figure 2, 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 low resiliency, and Panel B has two kinds of price impact with, respectively, high and low resiliency.

¹⁰Our 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.

¹¹Erb and Harvey (2006) document 12-month momentum in commodity futures prices. Asness, Moskowitz, and Pedersen (2008) confirm this finding and also document five-year reversals. These results are robust and hold for both price changes and returns. Results for five-day momentum are less robust. For instance, for certain specifications using percent returns, the five-day coefficient switches sign to reversal. This robustness is not important for our study, however, due to our focus on optimal trading rather than out-of-sample return predictability.

¹²The 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 five-day signal.

¹³Remember that A_{xx} and A_{ff} can always be chosen to be symmetric.

References

- Acharya, Viral, and Lasse Heje Pedersen, 2005, Asset pricing with liquidity risk, *Journal of Financial Economics* 77, 375–410.
- Almgren, Robert, and Neil Chriss, 2000, Optimal execution of portfolio transactions, *Journal of Risk* 3, 5–39.
- Amihud, Yakov, and Haim Mendelson, 1986, Asset pricing and the bid-ask spread, *Journal of Financial Economics* 17, 223–249.
- Asness, Cliff, Tobias Moskowitz, and Lasse Heje Pedersen, 2008, Value and momentum everywhere, *Journal of Finance* (forthcoming).
- Balduzzi, Pierluigi, and Anthony W. Lynch, 1999, Transaction costs and predictability: some utility cost calculations, *Journal of Financial Economics* 52, 47–78.
- Bertsimas, Dimitris, and Andrew W. Lo, 1998, Optimal control of execution costs, *Journal of Financial Markets* 1, 1–50.
- Breen, William J., Laurie S. Hodrick, and Robert A. Korajczyk, 2002, Predicting equity liquidity, *Management Science* 48, 470–483.
- Brunnermeier, Markus K., and Lasse H. Pedersen, 2005, Predatory trading, *Journal of Finance* 60, 1825–1863.
- Campbell, John Y., and Luis M. Viceira, 2002, *Strategic Asset Allocation Portfolio Choice for Long-Term Investors*. (Oxford University Press, Oxford, UK).
- Carlin, Bruce I., Miguel Lobo, and S. Viswanathan, 2008, Episodic liquidity crises: Cooperative and predatory trading, *Journal of Finance* 62, 2235–2274.
- Constantinides, George 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, Robert, and Robert Ferstenberg, 2007, Execution risk, *Journal of Portfolio Management* 33, 34–45.
- Engle, Robert, Robert Ferstenberg, and Jeffrey Russell, 2008, Measuring and modeling execution cost and risk, *Working paper, University of Chicago*.
- Erb, Claude, and Campbell R. Harvey, 2006, The strategic and tactical value of commodity futures, *Financial Analysts Journal* 62, 69–97.
- Gârleanu, Nicolae, 2009, Portfolio choice and pricing in imperfect markets, *Journal of Economic Theory* 144, 532–564.
- Gârleanu, Nicolae, Lasse Heje Pedersen, and Allen Poteshman, 2009, Demand-based option pricing, *Review of Financial Studies* 22, 4259–4299.
- Greenwood, Robin, 2005, Short and long term demand curves for stocks: Theory and evidence, *Journal of Financial Economics* 75, 607–650.
- Grinold, Richard, 2006, A dynamic model of portfolio management, *Journal of Investment Management* 4, 5–22.
- Grossman, Sanford, and Merton Miller, 1988, Liquidity and market structure, *Journal of Finance* 43, 617–633.
- Heaton, John, and Deborah Lucas, 1996, Evaluating the effects of incomplete markets on risk sharing and asset pricing, *Journal of Political Economy* 104, 443–487.
- Lagos, Ricardo, and Guillaume Rocheteau, 2009, Liquidity in asset markets with search frictions, *Econometrica* 77, 403–426.
- Lillo, Fabrizio, J. Doyne Farmer, and Rosario N. Mantegna, 2003, Master curve for price-impact function, *Nature* 421, 129–130.

- Liu, Hong, 2004, Optimal consumption and investment with transaction costs and multiple assets, *Journal of Finance* 59, 289–338.
- Ljungqvist, Lars, and Thomas Sargent, 2004, *Recursive Macroeconomic Theory, 2nd edition*. (MIT press, Cambridge, MA).
- Lo, Andrew, Harry Mamaysky, and Jiang Wang, 2004, Asset prices and trading volume under fixed transaction costs, *Journal of Political Economy* 112, 1054–1090.
- Longstaff, Francis A., 2001, Optimal portfolio choice and the valuation of illiquid securities, *The Review of Financial Studies* 14, 407–431.
- Lynch, Anthony, and Sinan Tan, 2011, Explaining the magnitude of liquidity premia: The roles of return predictability, wealth shocks, and state-dependent transaction costs, *Journal of Finance* 66, 1329–1368.
- Lynch, Anthony W., and Pierluigi Balduzzi, 2000, Predictability and transaction costs: the impact on rebalancing rules and behavior, *Journal of Finance* 55, 2285–2309.
- Markowitz, Harry M., 1952, Portfolio selection, *Journal of Finance* 7, 77–91.
- Obizhaeva, Anna, and Jiang Wang, 2006, Optimal trading strategy and supply/demand dynamics, *Working paper, MIT*.
- Oehmke, Martin, 2009, Gradual Arbitrage, *Working paper, Columbia*.
- Perold, Andre, 1988, The implementation shortfall: Paper versus reality, *Journal of Portfolio Management* 14, 4–9.
- Vayanos, Dimitri, 1998, Transaction costs and asset prices: A dynamic equilibrium model, *Review of Financial Studies* 11, 1–58.
- Vayanos, Dimitri, and Jean-Luc Vila, 1999, Equilibrium interest rate and liquidity premium with transaction costs, *Economic Theory* 13, 509–539.

Table I
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 to 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.

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 II

Performance of Trading Strategies Before and After Transaction Costs.

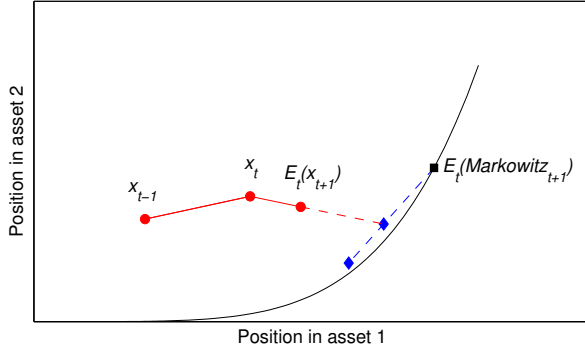
This table shows the annualized Sharpe ratio gross (“Gross SR”) and net (“Net SR”) of trading costs for the optimal trading strategy in the absence of trading costs (“Markowitz”), our optimal dynamic strategy (“Dynamic”), and a strategy that optimizes a static one-period problem with trading costs (“Static”). Panel A illustrates these results for a low transaction cost parameter, while Panel B uses a high one.

	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


Panel A. Constructing the current optimal portfolio




Panel B. Expected optimal portfolio next period




Panel C. Expected future path of optimal portfolio



Panel D. “Skate to where the puck is going to be”



Panel E. Shooting: lead the duck



Panel F. Missile systems: lead homing guidance




Figure 1: Aim in front of the target. Panels A–C show the optimal portfolio choice with two securities. The Markowitz portfolio is the current optimal portfolio in the absence of transaction costs: the target for an investor. It is a moving target, and the solid curve shows how it is expected to mean-revert over time (towards the origin, which could be the market portfolio). Panel A shows how the optimal time- t trade moves the portfolio from the existing value x_{t-1} towards the aim portfolio, but only part of the way. Panel B shows the expected optimal trade at time $t + 1$. Panel C shows the entire future path of the expected optimal portfolio. The optimal portfolio “aims in front of the target” in the sense that, rather than trading towards the current Markowitz portfolio, it trades towards the aim, which incorporates where the Markowitz portfolio is moving. Our portfolio principle has analogues in sports, hunting, and missile guidance as seen in Panels D–F.

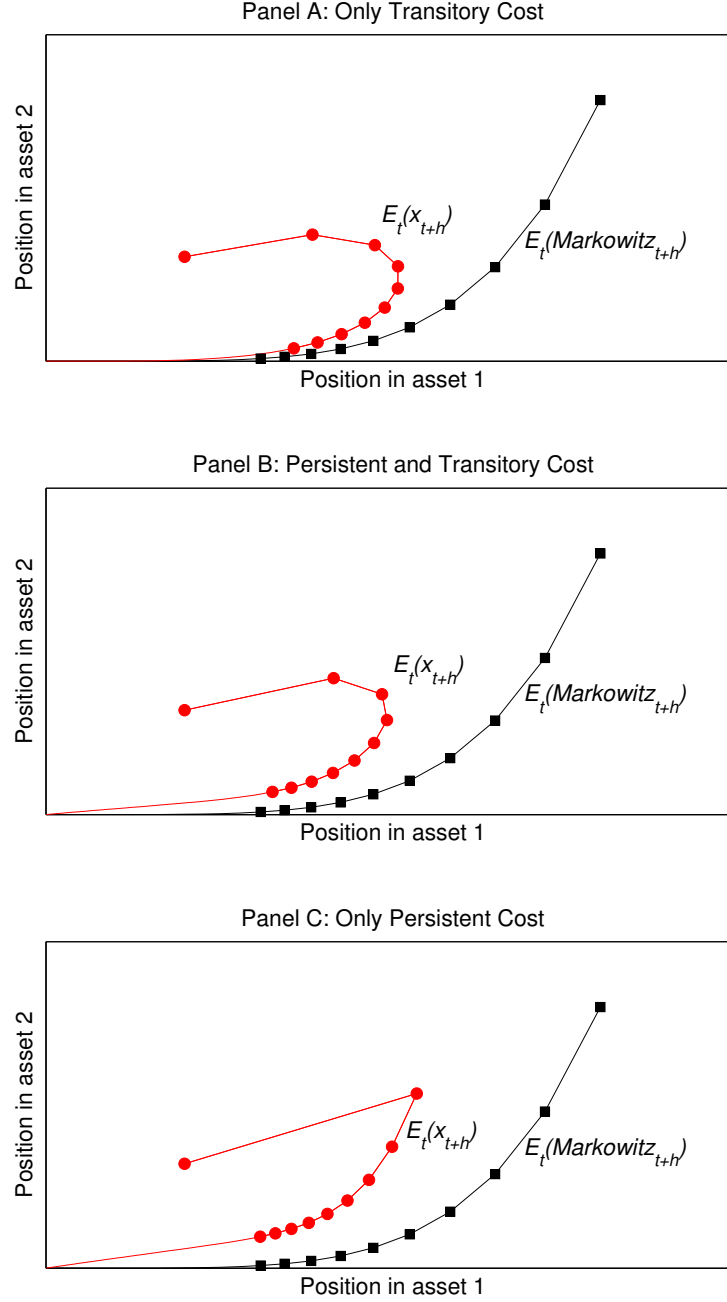


Figure 2: Aim in front of the target with persistent costs. This figure shows the optimal trade when part of the transaction cost is persistent. In Panel A, the entire cost is transitory, as in Figure 1 (A–C). In Panel B, half of the cost is transitory, while the other half is persistent, with a half-life of 6.9 periods. In Panel C, the entire cost is persistent.




Figure 3: 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”).

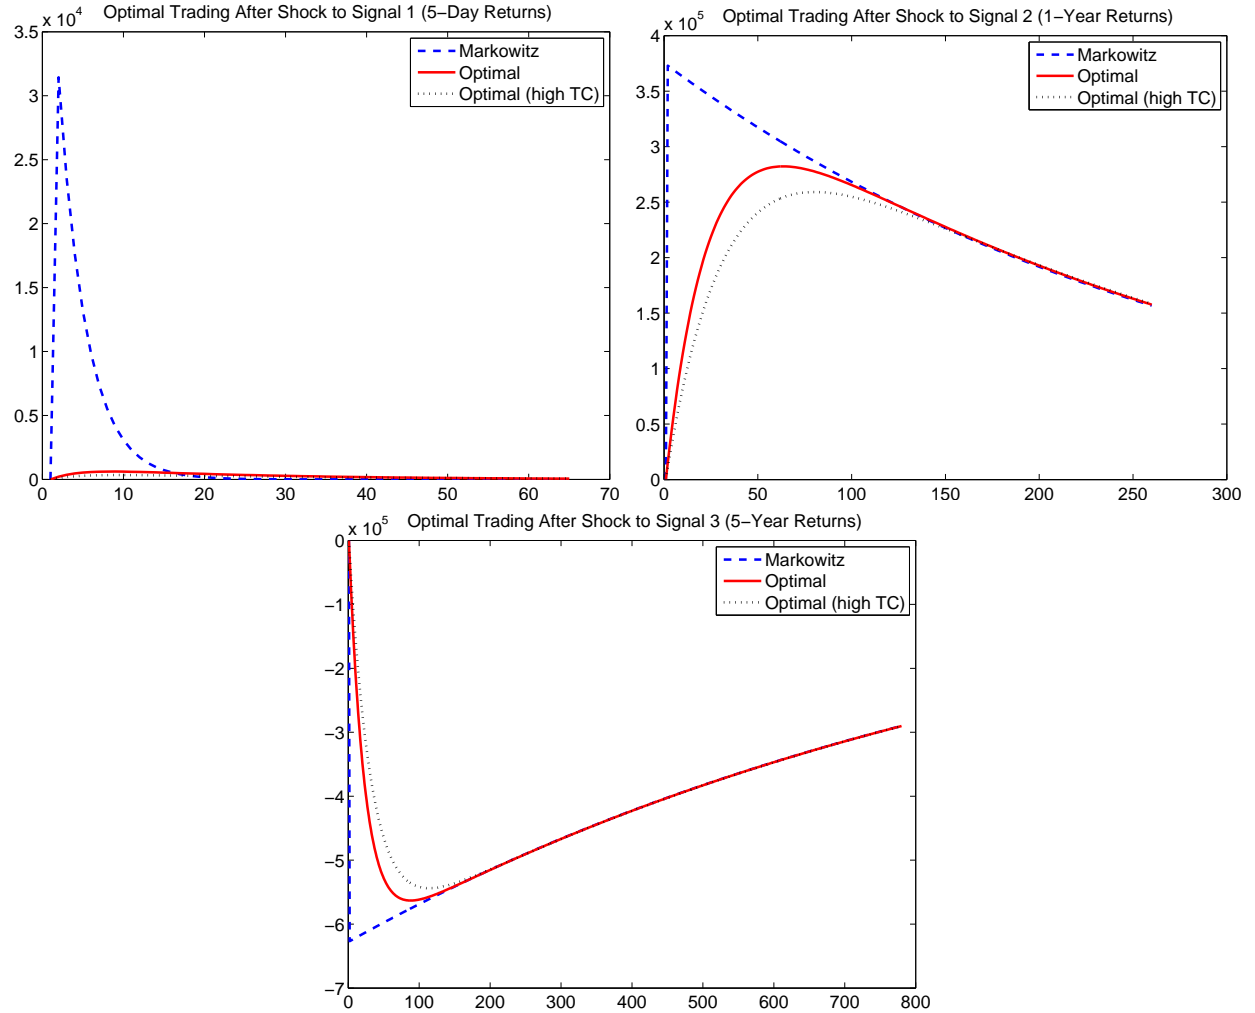


Figure 4: Optimal trading in response to shocks to return-predicting signals. This figure shows the response in the optimal position following a shock to a return predictor as a function of the number of days since the shock. The top left panel considers a shock to the fast five-day return predictor, the top right panel considers a shock to the one-year return predictor, and the bottom panel to the five-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.