Dynamic Trading with Predictable Returns and Transaction Costs

Nicolae Gârleanu
UC Berkeley, CEPR, and NBER

Lasse H. Pedersen
New York University, Copenhagen Business School,
AQR Capital Management, CEPR, and NBER

Motivation: Dynamic Trading

- Active investors – e.g., hedge funds, mutual funds, proprietary traders, individuals, other asset managers – try to
  - predict returns
  - minimize transactions costs
  - minimize risk

- Dynamic problem: investor trades now and in the future

- Key research questions:
  - What is the optimal trading strategy?
  - Does it work empirically?

Motivating Example

- An investor makes the following predictions:
  - Based on strong fundamentals (low M/B, P/E, low accruals, high and stable earnings, etc.) the annualized expected excess return (alpha) on Centurytel Inc. is 10%.
    - this alpha is expected to last for 2 years
  - Based on recent catalysts, improving fundamentals and pricing, the annualized alpha of Treehouse Foods Inc. is also 10%
    - this alpha is expected to last for half a year
  - Based on recent demand pressure from funds with outflow, the annualized alpha of HJ Heinz Co. is -12%
    - this alpha is expected to last for 2 weeks
  - These and other signals are collected for numerous securities
- All these stocks are positively correlated
- The investor has estimated the trading cost (incl. market impact) for these stocks based on past experience
- The investor makes a similar analysis every day

Results: Aim in Front of the Target

- Closed-form optimal dynamic trading strategy

Results: Aim in Front of the Target

- ▷ Closed-form optimal dynamic trading strategy
- ▷ Two portfolio principles:
  1. Aim in front of the target
  2. Trade partially towards the current aim

Results: Aim in Front of the Target

- ▷ Closed-form optimal dynamic trading strategy
- ▷ Two portfolio principles:
  1. Aim in front of the target
  2. Trade partially towards the current aim

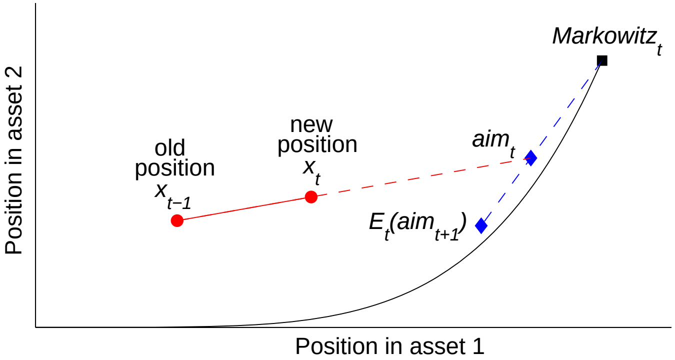

Results: Aim in Front of the Target

- ▷ Closed-form optimal dynamic trading strategy
- ▷ Two portfolio principles:
  1. Aim in front of the target
  2. Trade partially towards the current aim
- ▷ "Aim portfolio":
  – Weighted average of current and future expected Markowitz portfolios
  – Predictors with slower mean reversion: more weight
- ▷ Application to commodity futures: superior net returns

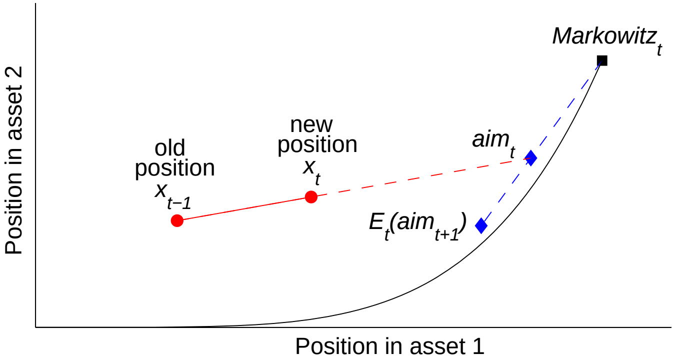

Aim in Front of the Target: Finance and Beyond

I. Dynamic trading: aim in front of the target

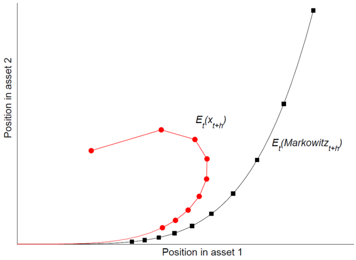

II. "Skate to where the puck is going to be" – Wayne Gretzky

III. Shooting: Lead the duck

IV. Missile systems: Lead homing guidance

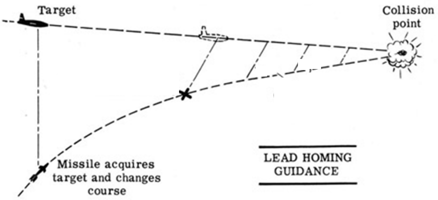

Related Literature

- Optimal trading with transactions costs, no predictability
  - Constantinides (86), Amihud and Mendelson (86), Vayanos (98), Liu (04)
- Predictability, no transactions costs
  - Merton (73), Campbell and Viceira (02)
- Optimal trade execution with exogenous trade:
  - Perold (88), Almgren and Chriss (00)
- Numerical results with time-varying investment opportunity set
  - Jang, Koo, Liu, and Loewenstein (07), Lynch and Tan (08)
- Quadratic programming
  - Used in macroeconomics (Ljungqvist and Sargent (04)) and other fields: solve up to Ricatti equations
  - Grinold (06)

Outline of Talk

- Basic model
- Optimal portfolio strategy: Aim in front of the target
- Persistent price impact
- Application: Commodity futures
- If time permits: Continuous-time model
- If time permits: Equilibrium

Discrete-Time Model

Returns: $$r^s_{t+1} = \underbrace{\sum_k \beta^{sk} f^k_t}_{=E_t(r^s_{t+1})} + u^s_{t+1}$$

Risk: $$\mathrm{var}_t(u_{t+1}) = \Sigma$$

Alpha decay: $$\Delta f^k_{t+1} = -\sum_j \Phi^{kj} f^j_t + \varepsilon_{t+1}$$

Transaction costs: $$TC(\Delta x_t) = \frac{1}{2}\Delta x_t^\top \Lambda \Delta x_t$$

Assumption A: $$\Lambda = \lambda\Sigma$$

Objective: $$\max_{x_t} \mathrm{E} \sum_t (1-\rho)^{t+1} \left(x_t^\top r_{t+1} - \frac{\gamma}{2} x_t^\top \Sigma x_t\right) - \frac{(1-\rho)^t}{2} \Delta x_t^\top \Lambda \Delta x_t$$

Solution Method: Dynamic Programming

Introduce value function V that solves the Bellman equation:

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

Proposition

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

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

The coefficient matrices A_{xx}, A_{xf}, A_{ff} can be solved explicitly and A_{xx} is positive definite.

Trade Partially Towards the Aim

Proposition (Trade Partially Towards the Aim)

i)The optimal dynamic portfolio x_t *is:*

$$x_t = x_{t-1} + \Lambda^{-1} A_{xx} \left(aim_t - x_{t-1} \right)$$

with "trading rate" $\Lambda^{-1} A_{xx}$ *and*

$$aim_t = A_{xx}^{-1} A_{xf} f_t$$

Trade Partially Towards the Aim

Proposition (Trade Partially Towards the Aim)

i)The optimal dynamic portfolio x_t is:

$$x_t \;=\; x_{t-1} + \Lambda^{-1} A_{xx} \left(aim_t - x_{t-1} \right)$$

with "trading rate" $\Lambda^{-1} A_{xx}$ and

$$aim_t \;=\; A_{xx}^{-1} A_{xf} f_t$$

ii) Under Assumption A, the trading rate is the scalar

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

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

What is the Target and What is the Aim?

- What is the moving target, i.e., the optimal position in the absence of transaction costs?

$$Markowitz_t = (\gamma\Sigma)^{-1} Bf_t$$

- What is the aim portfolio?

Aim in Front of the Target

Proposition (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, letting $z = \gamma/(\gamma + a)$:

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

(ii) The aim portfolio is the weighted average of the current and future expected Markowitz portfolios. Under Assumption A,

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

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

Aim in Front of the Target: Illustration

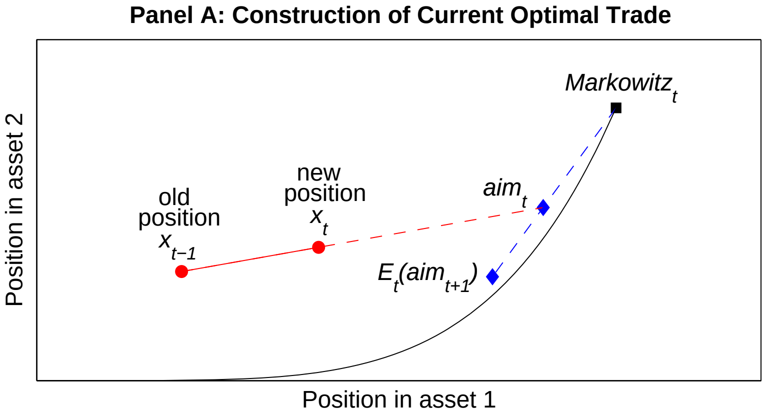

Aim in Front of the Target: Illustration

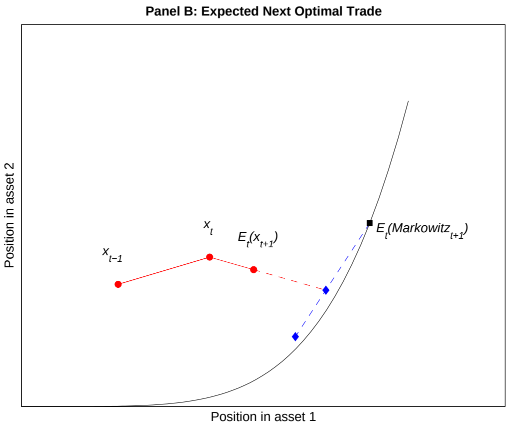

Aim in Front of the Target: Illustration

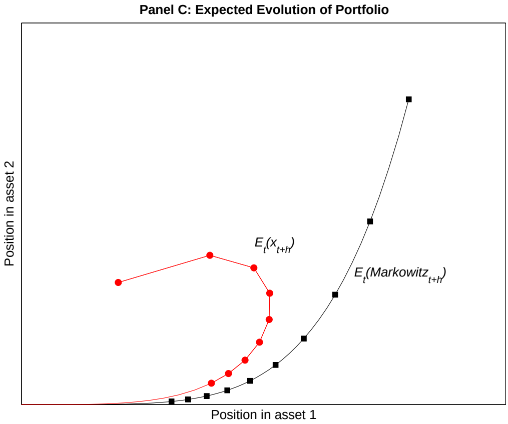

Weight Signals Based on Alpha Decay

Proposition (Weight Signals Based on Alpha Decay)

(i) Under Assumption A, the aim portfolio is:

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

(ii) If the matrix Φ *is diagonal,* $\Phi = \mathrm{diag}(\phi^1, ..., \phi^K)$*, then the aim portfolio is:*

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

I.e., the aim pf. is the Markowitz pf. with factors f_t^k
scaled down based on their own alpha decay given by Φ.

Weight Signals Based on Alpha Decay: Illustration

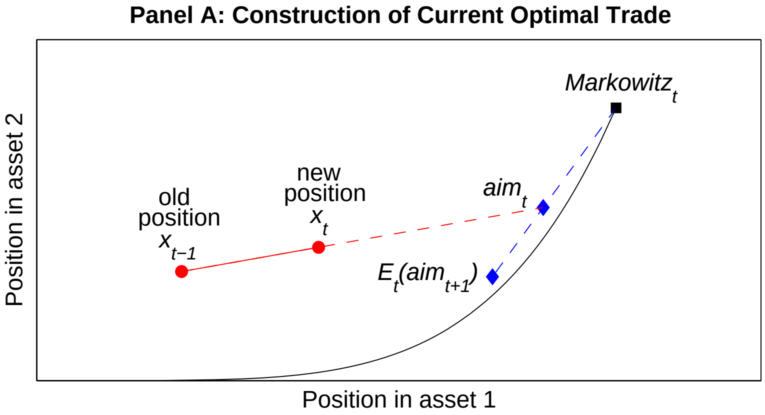

Position Homing In

Proposition (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,

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

Example: Timing a Single Security

A security has risk $\Sigma = \sigma^2$ and return

$$r_{t+1} = \underbrace{\sum_k \beta^k f_t^k}_{=E_t(r_{t+1})} + u_{t+1}$$

The optimal strategy is

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

Example: Relative-Value Trades w/ Security Characteristics

Each security s (e.g., IBM) has its own characteristics $f_t^{i,s}$ (e.g., its value and momentum) and characteristics predict returns for all securities, with the same coefficients:

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

Each characteristic has the same mean-reversion speed for all securities

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

Example: Relative-Value Trades w/ Security Characteristics

Each security s (e.g., IBM) has its own characteristics $f_t^{i,s}$ (e.g., its value and momentum) and characteristics predict returns for all securities, with the same coefficients:

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

Each characteristic has the same mean-reversion speed for all securities

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

The optimal characteristic-based strategy is

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

Example: Static Model

When the future is completely discounted ($\rho = 1$), objective is

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

Example: Static Model

When the future is completely discounted ($\rho = 1$), objective is

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

Solution

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

No choice of γ, λ recovers the dynamic solution.

Example: Signals (Equally) Valuable for K Days

Suppose:

- All factors equally good $B = (\beta, ..., \beta)$
- Today's yesterday is tomorrow's day-before-yesterday:

$$\begin{aligned} f^1_{t+1} &= \varepsilon^1_{t+1} \\ f^k_{t+1} &= f^{k-1}_t \quad \text{for } k > 1 \end{aligned}$$

Example: Signals (Equally) Valuable for K Days

Suppose:

- All factors equally good $B = (\beta, ..., \beta)$
- Today's yesterday is tomorrow's day-before-yesterday:

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

Optimal strategy:

$$x_t = \left(1 - \frac{a}{\lambda}\right) x_{t-1} + \frac{a}{\lambda} \frac{\beta}{\sigma^2 (1 - z)} \sum_k \left(1 - z^{K+1-k}\right) f_t^k,$$

where $z = a/(a + \gamma) < 1$.

More weight to recent signals even if they don't predict better.

Persistent Transaction Costs Model

Proposition

With temporary and persistent transaction costs, the optimal portfolio x_t is

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

which tracks an aim portfolio, $aim_t = M^{aim} y_t$, that depends on the return-predicting factors and the price distortion.

Persistent Transaction Costs Model

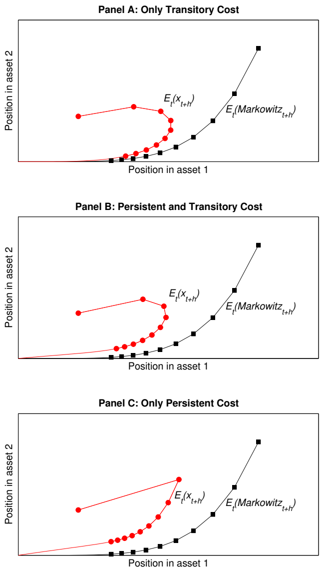

Application: Dynamic Trading of Commodity Futures

Data on liquid futures without tight price limits 01/01/1996 – 01/23/2009:

- Aluminum, Copper, Nickel, Zinc, Lead, Tin from London Metal Exchange (LME)
- Gas Oil from the Intercontinental Exchange (ICE)
- WTI Crude, RBOB Unleaded Gasoline, Natural Gas from New York Mercantile Exchange (NYMEX)
- Gold, Silver is from New York Commodities Exchange (COMEX)
- Coffee, Cocoa, Sugar from New York Board of Trade (NYBOT)

Predicting Returns and Other Parameter Estimates

Pooled panel regression:

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

Alpha decay:

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

Risk: Σ estimated using daily price changes
Absolute risk aversion: $\gamma = 10^{-9}$
Time discount rate: $\rho = 1 - \exp(-0.02/260)$
Transactions costs: $\lambda = 3 \times 10^{-7}$, as well as $\lambda^{high} = 10 \times 10^{-7}$

Performance of Trading Strategies Before and After TCs

	Panel A: Benchmark Transaction Costs		Panel B: High Transaction Costs	
	Gross SR	**Net SR**	**Gross SR**	**Net SR**
Markowitz	0.83	-9.38	0.83	-10.11
Dynamic optimization	**0.63**	**0.60**	**0.58**	**0.53**
Static optimization				
Weight on Markowitz = 10%	0.63	0.00	0.63	-1.45
Weight on Markowitz = 9%	0.62	0.10	0.62	-1.10
Weight on Markowitz = 8%	0.62	0.20	0.62	-0.78
Weight on Markowitz = 7%	0.62	0.29	0.62	-0.49
Weight on Markowitz = 6%	0.62	0.36	0.62	-0.22
Weight on Markowitz = 5%	0.61	0.43	0.61	0.00
Weight on Markowitz = 4%	0.60	0.48	0.60	0.19
Weight on Markowitz = 3%	**0.58**	**0.51**	0.58	0.33
Weight on Markowitz = 2%	0.52	0.49	**0.52**	**0.39**
Weight on Markowitz = 1%	0.36	0.34	0.36	0.31

Positions in Crude and Gold Futures

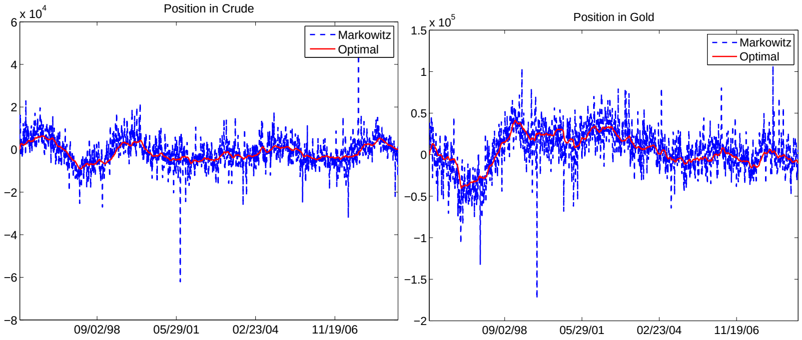

Optimal Trading in Response to Shock to 5-Day Return-Predicting Signal

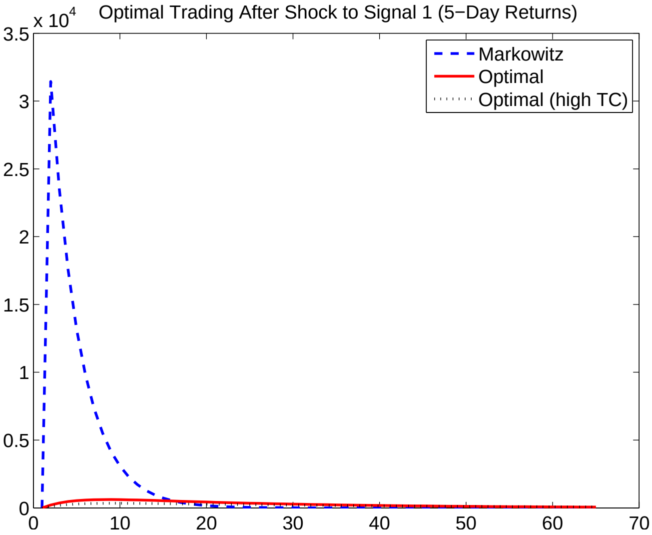

Optimal Trading in Response to Shock to 1-Year Return-Predicting Signal

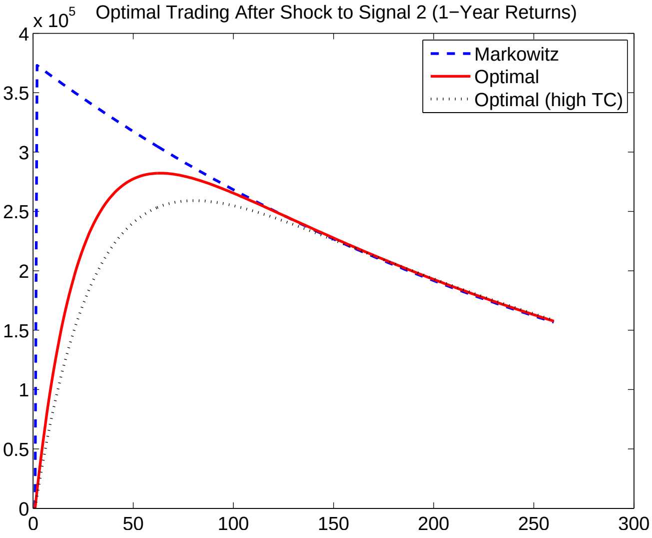

Optimal Trading in Response to Shock to 5-Year Return-Predicting Signal

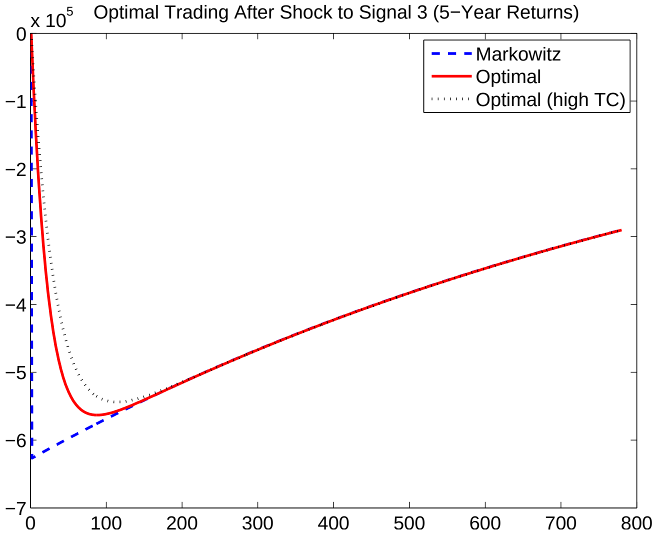

Continuous-Time Model

Returns:	dp_t	$=$	$(r^f p_t + \underbrace{Bf_t}_{=\alpha_t})dt + du_t$
Risk:	$\text{var}_t(du_t)$	$=$	$\Sigma\, dt$
Alpha decay:	df_t	$=$	$-\Phi f_t dt + d\varepsilon_t$
T-cost with $dx_t = \tau_t dt$:	$TC(\tau_t)$	$=$	$\frac{1}{2}\tau_t^\top \Lambda \tau_t$

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

Continuous-Time Solution Method

Value function V solves Hamilton-Jacoby-Bellman equation:

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

Maximizing with respect to the trading intensity results in

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

Conjecture and verify quadratic value function:

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

Optimal Continuous-Time Strategy

Proposition

The optimal trading intensity $\tau = \frac{dx_t}{dt}$ *is*

$$\tau_t = \Lambda^{-1} A_{xx} \left[aim_t - x_t \right],$$

where

$$aim_t = A_{xx}^{-1} A_{xf} f_t.$$

The coefficients are

$$\begin{aligned} A_{xx} &= -\frac{\rho}{2}\Lambda + \Lambda^{\frac{1}{2}} \left(\gamma \Lambda^{-\frac{1}{2}} \Sigma \Lambda^{-\frac{1}{2}} + \frac{\rho}{4} I \right)^{\frac{1}{2}} \Lambda^{\frac{1}{2}} \\ \mathrm{vec}(A_{xf}) &= \left(\rho I + \Phi^{\top} \otimes I_K + I_S \otimes (A_{xx} \Lambda^{-1}) \right)^{-1} \mathrm{vec}(B) \end{aligned}$$

Optimal Continuous-Time Strategy

Proposition

If $\Lambda = \lambda\Sigma$, then the optimal trading intensity $\tau_t = \frac{dx_t}{dt}$ is

$$\tau_t = \frac{a}{\lambda}\left[aim_t - x_t\right]$$

where

$$\begin{aligned} aim_t &= (\gamma\Sigma)^{-1} B \left(I + a\Phi/\gamma\right)^{-1} f_t \\ a &= \frac{-\rho\lambda + \sqrt{\rho^2\lambda^2 + 4\gamma\lambda}}{2}. \end{aligned}$$

If each factor's alpha decay only depends on itself, then:

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

Persistent Price Impact: Model (Continuous Time)

Transaction price:	$\bar{p}_t$	$=$	$p_t + D_t$
Price distortion:	dD_t	$=$	$-RD_t dt + C\tau_t dt$
Returns:	$d\bar{p}_t$	$=$	$\underbrace{Bf_t dt + dD_t}_{=\alpha_t dt} + du_t$
Risk:	$\mathrm{var}_t(du_t)$	$=$	$\Sigma\, dt$
Alpha decay:	df_t	$=$	$-\Phi f_t dt + d\varepsilon_t$
T-cost with $dx_t = \tau_t dt$:	$TC(\tau_t)$	$=$	$\frac{1}{2}\tau_t^\top \Lambda \tau_t$

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

Persistent Price Impact: Solution

Proposition (Temporary and Persistent Transaction Costs)

If temporary costs are positive, the optimal trading strategy is

$$\tau_t = M^{rate} \times (aim_t - x_t)\,,$$

with

$$aim_t = M^{aim} \times (f_t, D_t)\,.$$

Proposition (Purely Persistent Costs)

If temporary costs are zero, the optimal portfolio is

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

The optimal portfolio is a linear combination of the current price distortion, the current Markowitz pf., and an exponentially weighted average of future expected Markowitz pfs.

Conclusion: Aim in Front of the Target

- Derive the closed-form optimal dynamic portfolio strategy
  1. Aim in front of the target
  2. Trade partially towards the current aim at constant rate
- Give more weight to persistent factors
- Superior net returns in application
- Lots of potential applications of this model
  - Equilibrium: large high-frequency alphas

Equilibrium Model with Noise Traders

Noise trader positions:

$$
\begin{aligned}
dz_t^l &= \kappa\left(f_t^l - z_t^l\right)dt \\
df_t^l &= -\psi_l f_t^l dt + dW_t^l.
\end{aligned}
$$

State variables: $f \equiv (f^1, ..., f^L, z)$ where $z_t = \sum_l z_t^l$.

Given this, mean-reversion matrix Φ is

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

Equilibrium:

$$
x_t = -z_t
$$

Equilibrium Alphas

Proposition

The market is in equilibrium iff

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

The coefficients $\lambda\sigma^2\kappa(\psi_k + \rho + \kappa)$ *increase in the mean-reversion parameters* ψ_k *and* κ *and in the trading costs* $\lambda\sigma^2$*.*

Equilibrium Alphas

Proposition

The market is in equilibrium iff

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

The coefficients $\lambda\sigma^2\kappa(\psi_k + \rho + \kappa)$ increase in the mean-reversion parameters ψ_k and κ and in the trading costs $\lambda\sigma^2$.

I.e., noise trader selling ($f_t^k < 0$) increases the alpha, and especially if its mean reversion is faster and if the trading cost is larger.

Appendix: Connection Discrete-Continuous Time

Proposition

Consider discrete-time models with parameters defined to depend on the time interval Δ_t in the following way:

$$
\begin{aligned}
\hat{\Sigma}(\Delta_t) &= \Sigma\Delta_t \\
\hat{\Omega}(\Delta_t) &= \Omega\Delta_t \\
\hat{\Lambda}(\Delta_t) &= {\Delta_t}^{-1}\Lambda \quad \textit{or} \quad \hat{\lambda}(\Delta_t) = {\Delta_t}^{-2}\lambda \\
\hat{B}(\Delta_t) &= B\Delta_t \\
\hat{\Phi}(\Delta_t) &= 1 - e^{-\Phi\Delta_t} \\
\hat{\rho}(\Delta_t) &= 1 - e^{-\rho\Delta_t} \\
\hat{\gamma}(\Delta_t) &= \gamma \\
\hat{R}(\Delta_t) &= I - e^{-R\Delta_t} \\
\hat{C}(\Delta_t) &= C.
\end{aligned}
$$

Then the discrete-time solution converges to the continuous-time solution as Δ_t approaches zero.

Appendix: Micro-Foundation

- Agent trades with specialized, risk-averse intermediaries
- Intermediaries require *fixed* length of time to place inventory with end users
- The implied transaction cost to the trader is linear in Δ_t^{-1} and quadratic in Δx